

Multiple Simultaneous Interpretations of Ambiguous Sentences

Peter Norvig
University of California, Berkeley

INTRODUCTION

This paper is concerned with the problem of semantic and pragmatic interpretation of ambiguous sentences. We start by offering a simple yet commonly adopted interpretation strategy:

Strategy 1: Apply syntactic rules to the sentence to derive a set of parse trees. Next apply semantic rules to the trees to get a set of logical formulae, and discard any inconsistent formulae. Do a pragmatic interpretation of each formula, and give a score to each possibility based on consistency or likelihood in the given context. Finally, choose the interpretation with the highest score.

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 1 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 the following 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 the third 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.

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) behaves like 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 2:

Strategy 2: Do syntactic, semantic, and pragmatic interpretation as in Strategy 1. Retain only the highest-ranking interpretation(s), according to some threshold function. If there is more than one interpretation remaining, alternate between them.

A problem with Strategy 2 is that it assumes all possible interpretations will be considered and ranked. However, many sentences have a prohibitively large or infinite number of interpretations. Consider the following sentence:

- (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

NORVIG

enumerating all possibilities and then choosing among them. Hobbs (1983) has argued that enumeration should be minimized even in syntactic and semantic analysis. Accepting that intuition, we get Strategy 3, which is similar to the approaches used by several recent authors, including Hobbs (1987), Stallard (1987), and Charniak (unpublished).

Strategy 3: 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. Solutions are generated in a roughly best-first manner, and when there is a large drop-off in the ranking of solutions, we stop and the final interpretation alternates between the high ranking one(s).

This might represent an efficient interpretation mechanism, but it doesn't mirror the human interpretation mechanism particularly well. Sentences like (5-9) each have only one good pragmatic interpretation, which would be found easily by Strategy 3. But (5-9) are notoriously hard for humans to get right without at least a conscious sense of having to back up and re-parse the sentence.

- (5) The horse raced past the barn fell.
- (6) The astronomer married a star.
- (7) The rabbi was hit on the temple.
- (8) The landlord painted all the walls with cracks.
- (9) Ross was told what to do by the river.

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, which appeal to performance issues such as limits on available memory space. Schubert (1984, 1986) and Kurtzman (1984) argue convincingly that no simple syntactic preference will do. Rather, many factors must be considered, as in Strategy 4:

Strategy 4: 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 simple 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. Strategies 2-4 assume that the reader eventually arrives at a single interpretation, or a Necker-cube-like alternation between interpretations. But Parsons is saying that his slogan *Fur is for Life* is different. The slogan seems to have two primary interpretations, (10) and (11)

NORVIG

below. But it also has important *connotations*, listed as (12-14), as well as another interpretation, (15), that Parsons presumably wants the public to ignore.

- (10) Fur lasts a lifetime.
- (11) The fur industry is pro-conservation.
- (12) Fur wearers are lively.
- (13) The recipient of a fur may become indebted to the giver for life.
- (14) Life is a good thing; hence fur is a good thing.
- (15) Fur, while on an animal, protects its life.

Although we would not be likely to say that any of (12-14) 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 theory of 'one interpretation at a time,' it appears to be quite common in poetry, politics, and advertising (see Burli-Storz, 1980).

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:

- (16) 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*:

- (17) 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 (18-23) are all used in the interpretation of the six words *go gentle into that good night*: (18) for *go*,

NORVIG

(19) for *gentle*, (20) for *into*, (21) for *good night*, and (22) 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 ambiguity.

- (18) Life is a journey.
- (19) Life is a struggle; death is an adversary.
- (20) Life is 'here'; death is a another world.
- (21) Death is sleep.
- (22) A lifetime is a day; death is night.
- (23) Life is a fire that blazes and burns out.

At this point let us try to modify Strategy 4 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 5a: 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 5b: 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 5b, as it is applied to sentence (24), and its interpretation, the disjunction (24'):

- (24) The chicken is ready to eat.
- (24') $\text{chicken}(x) \ \& \ \text{ready}(x,e) \ \& \ \text{eating}(e) \ \& \ (\text{agent}(e,x) \ | \ \text{patient}(e,x))$

Using Strategy 5b, 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 (24), so we have an argument against 5b. However, that argument only goes through if the proposed logical form (24') is accurate. Suppose we use the following logical form instead:

- $\text{chicken}(x) \ \& \ \text{ready}(x,e) \ \& \ \text{eating}(e) \ \& \ ((\text{agent}(e,x) \ \& \ \text{alive}(x) \ \& \ \text{location}(e,\text{barnyard}) \ \& \ \text{patient}(e,\text{seed}) \ \& \ \dots) \ | \ (\text{patient}(e,x) \ \& \ \text{not alive}(x) \ \& \ \text{location}(e,\text{table}) \ \& \ \text{agent}(e,\text{human}) \ \& \ \dots))$

Then we have two interpretations that cannot be combined coherently, neither under Strategy 5b nor 5a. Thus, we see that for 5b to be feasible, we need to insist on full 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." Now let's try an example that does not bring as much background knowledge into play:

- (25) 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 5b from accepting both interpretations for the phrase, whereas we know that if this were the intended meaning, one would have to use something like

NORVIG

the following:

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

Thus, Strategy 5b as it stands is rejected. To evaluate Strategy 5a, we need to develop a better notion of *sanctioning* a combined interpretation, which we will address in the next section.

JOKES AND PUNS

Consider the following advertisement for Flintstones brand Vitamins:

(27) We are Flintstones kids, ten million strong and growing.

The coordinate *and growing* can attach to either *are* or *ten million strong*, with the respective interpretations that the individual children are growing, or that the number of children is increasing. Most informants recognize 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 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 4, and in the process of trying to choose between them, realizes that both were intended interpretations, and successfully superimposes the two images.

In short, (27) 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 good pun, both of the meanings taken separately). But in (27) 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:

(28) 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 your tires' and 'much rides in the car which is on the tires,' 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):

(29) 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, *familiar* + *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 elicited by the release of suppressed violent or sexual thoughts. That explains, perhaps, why the following is a fairly good

NORVIG

joke, while other lexical and structural ambiguities in this paper are not:

(30) 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 also cases of combined simultaneous interpretation which don't involve poetic license or puns. Consider the use of *book* in (31). *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.

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

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

It can't be that *book* is a single sense implying all these aspects, because in (32), *book* must refer only to the 'plot or sequence of ideas' sense. One could not felicitously use (32) 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.

Len Talmy (1977) provides a good example of image combination in non-ambiguous language. In (33), the single interpretation is 'she traveled lightly and easily through the room and the guests at the party, her path displaying a topology similar to a leaf wafting through air.' (33) 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 the result. Talmy explains just what properties of the verb are maintained, and which are taken from the complements.

(33) 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 (34), which is a consistent use of metaphor, with (35). Sentence (35) provides a topological clash that cannot easily be resolved into a single interpretation, even though the meaning of the two cliches is consistent.

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

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

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

Sentence (36), taken from a newspaper article on real estate prices, is an example of a mixed metaphor with varying effectiveness; some find it to be fine, while most report that the topology is all wrong: one should be striving to get out of a rat-hole, not out from under it.

NORVIG

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 experimentally derived human preference results, and (3) uses a combination mechanism that is needed for non-ambiguous language as well.

ACKNOWLEDGMENT

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, *Deliberate Ambiguity in Advertising*, Bern, 1980.
- Frazier, Lyn and Janet Dean Fodor, "The sausage machine: A new two-stage parsing model," *Cognition*, 6(4), December 1978, 291-325.
- Freud, Sigmund, *Jokes and Their Relation to the Unconscious*, Translated by James Strachey, 1960, Norton, NY, originally published 1916.
- Hobbs, Jerry, "Representing ambiguity," *Proceedings of the Second West Coast Conference on Formal Linguistics*, Stanford University, 1983.
- Hobbs, Jerry, and Paul Martin, "Local pragmatics," *Proceedings of the International Joint Conference on Artificial Intelligence*, 1987.
- Hockett, Charles F., "Two models of grammatical description," *Word*, 1954, 386-399.
- Jackendoff, Ray, *Consciousness and the Computational Mind*, MIT Press, 1987.
- Kimball, John, "Seven principles of surface parsing in natural language," *Cognition*, 2(1), 1973, 15-47.
- Kurtzman, Howard S., "Ambiguity resolution in the human syntactic parser: An experimental study," *Coling*, 1984.
- Lakoff, George, and Mark Turner, *More Than Cool Reason*, to appear, 1988.
- Minsky, Marvin, "Jokes and the logic of the cognitive unconscious," *AI memo No. 603*, MIT AI Lab, 1980.
- Schubert, Lenhart K., "On parsing preferences," *Proceedings of Coling84*, Stanford, CA, July 1984.
- Schubert, Lenhart K., "Are there preference trade-offs in Attachment Decisions?," *Proceedings of AAAI*, 1986, 601-605.
- Stallard, David, "The logical analysis of lexical ambiguity," 25th Annual Meeting of the ACL, 1987, 179-185.
- Talmy, Leonard, "Rubber-sheet cognition in language," *Chicago Linguistics Society*, 13, 1977, 612-628.