

Position Paper on Common-sense and Formal Semantics

Geoffrey Nunberg
Xerox PARC and CSLI, Stanford

1. A philological excursus

I'm not sure what I'm doing on this panel, but I thought it would be helpful if we could start at the beginning. It's interesting to note that both the dictionary and common sense were eighteenth-century inventions. This is no coincidence; in fact, it's entirely appropriate that the most celebrated blow that was ever struck on behalf of common sense was delivered by Dr. Johnson himself, when he kicked a stone in refutation of Berkeley's idealism. If you'll indulge me, I'd like briefly to explore this connection, with the promise that I'll wind up by drawing a moral about natural-language understanding.

First for "common sense." To be sure, the phrase had been in use since the fourteenth century, but it was not until the eighteenth century that it acquired three important philosophical uses, which we confuse at our peril. First, there was the understanding of common sense as the faculty of judgement possessed by the average person as part of his birthright. This is the sense of the phrase that Priestly had in mind when he wrote (1775): "Common sense...in common acceptation...has long been appropriated...to that capacity for judging of common things that persons of middling capacities are capable of." In this sense, common sense is variously a general or universal faculty, "the same in every time and clime," as Hume (somewhere) put it; to avoid confusion here, I'll sometimes refer to this as the faculty of common sense. Now for many philosophers of the period, it was an article of doctrine that natural and moral law were accessible to common sense, and hence that the body of common man could achieve both the degree of understanding of the world and the refinement of moral judgment necessary to self-government. This was the view of Hume, of Johnson, and of course of Thomas Paine, who entitled his famous treatise on natural rights *Common Sense*.

The second use of the phrase was more narrowly associated with the Scottish School of philosophy (also called the "Philosophers of Common Sense"); on this understanding it is a universal endowment, not just of critical faculties, but of beliefs. As the philosopher Hamilton put it, common sense in this interpretation is "the complement of those cognitions or convictions which we receive from nature; which all men possess in common...." This group made common sense the test of all philosophical doctrines, and in particular referred to the common-sense belief in the reality of the physical world in the course of their Berkeley-bashing (a

position to which Hume, unlike Johnson, was not sympathetic). More recently, this is the sense of the phrase that a lot of AI researchers appear to have in mind when they talk about "common-sense knowledge" as a body of generally accepted beliefs about the world.

The third eighteenth-century use of "common sense" is somewhat archaic nowadays, and usually appeared with the definite article, as "the common sense." Here, the common sense is the body of beliefs and values possessed by a community as a whole--what we might now call "received wisdom," or "custom" or perhaps what Durkheim referred to as the "collective representation." For now, I'll call this the "collective sense," to avoid confusion. Roughly, you could think of this as a kind of systematic representation that exists only collectively, and which is realized only in interaction.

Now for the dictionary. Before the eighteenth century, such dictionaries as existed were mostly lists of "hard words," and no one seemed to mind the absence of a dictionary of the modern sort, which would set out the rules for using all the words in general currency. It's a complicated matter to explain why this changed in Johnson's time. In part, he was animated by a widespread contemporary sentiment that knowledge was expanding and becoming specialized beyond the ability of any one person to master it (recall that this was also the time at which the "division of labor" entered the general discussion). But he was also sensitive to the doctrines about common sense and politics that I mentioned a moment ago. What he hoped to do, finding the language "copious without order," was nothing less than to systematize the collective sense of the English-speaking community, or at least such part of it as was embodied in the uses of words, so as to make it generally accessible to the ordinary reader. This step was crucial to the democratic program, which relied on the ability of the common man to coordinate an informed discourse about the world as it was and as it ought to be; to use words, if not in exactly the same ways, then at least in line with general practice.

Note that the possibility of succeeding in this enterprise rested on two assumptions about the role that the individual faculty of common sense would have to play in the process. First, as I noted earlier, Johnson and his contemporaries assumed that the ordinary individual was capable of grasping that part of the collective sense embodied in the definition (or as he called it, the "explanation") of a word; as he put it in an essay in the *Idler*, the most obscure doctrines of philosophy and science would be "found to contain nothing more than very plain truths," which could be "delivered in plain language." The second assumption is somewhat subtler. Johnson had to assume as well that readers would come generally to acknowledge the authority of his dictionary, or in more humility, the authority of the authors he cited in explication of the uses of words. So far as I know, Johnson himself never addressed this problem head-on, but the importance of common sense in this process was acknowledged by his contemporary George Campbell, a rhetorician associated with the common-sense school, in his *Philosophy of Rhetoric* (1776): "The source...of that preference which distinguisheth good use from bad in language, is a natural propensity of the human mind to believe that those are the best judges of the proper signs, and of the proper applications of them, who understand best the things which they represent." It is fair to say, then, that eighteenth-century

beliefs about the faculty of common sense both made the *Dictionary* necessary, and made it possible.

2. Common sense and the social

Now let me wrench this out of context. What makes the eighteenth century interesting from a linguistic point of view is not the peculiarities of its situation--the fact, for example, that an actual written dictionary was first produced then, or that the enterprise was connected to explicit political notions. The latter is important only in that it forced people to become self-consciously aware of certain problems of language, and set up a scheme for taking them on. (Analogously, it was in the eighteenth century that philosophers first became aware of the market as a social type, though certainly markets themselves had been around long before.) In particular, they made an explicit effort to talk about problems of linguistic coordination. And stripped of local details, their formulation of the situation is applicable to all discourse in all languages. You could say that all linguistic communities are democracies, in the sense that deference to linguistic authority is always consensual.

If these questions haven't loomed very large in recent research in linguistics, AI and related fields, it's only because those fields have been wed to a doctrine of what Putnam calls "methodological solipsism." I'll come back to this point presently, but right now I want to try to show how these considerations bear on problems of natural-language interpretation.

Suppose we put the problem in a schematic way. On the one hand, you have this extensive body of knowledge and assumptions--the collective sense-- which underlies the use of natural-language expressions. A part of this knowledge is actually possessed by all discourse participants when they interpret utterances--this is what constitutes their "common-sense beliefs" in the accepted use of the term. But the rest of this material is not represented by most users when they use most expressions, though parts of it are available to some speakers in some contexts. And yet this information is at least tacitly accepted as licencing all uses of all expressions. Most talk, that is, is carried out on credit, not cash on the barrel head, but the success of particular exchanges, as well as the overall coordination of the larger discourse, depends crucially on a general belief that the credit is sound; that the linguistic chits can be cashed in if circumstances required. And as with any sound currency--and linguistic currency is more stable even than the Swiss Franc--the interesting and somewhat paradoxical consequence is that people almost never do call in the notes, because it's taken for granted that they're good.

Some examples may help. Now the only area in which the social determination of meaning has been considered at all is in the use of natural-kind terms (though no one seems to know what to make of these, either). But I want to stay away from these cases, because in fact the situation here is rather complicated, and is obscured in Putnam's account. What is remarkable is that no one has given much thought to the fairly obvious point that the linguistic division of labor is pervasive in any direction you look. Take words from any domain: *grade-A prime*, *felony*, *squeeze play*, *sentence*, *jazz*, *braise*, or *hacker*; you'll find that most people

use these words in imperfect knowledge of the things they are held to denote, and in implicit deference to the collective sense that regulates their use.

This in turn raises several questions for theories of natural-language understanding. First, how do people get away with talking about the world in the absence of what economists would call "perfect information" about the meanings of the words they use? The answer is that the task of regulating words and world is taken care of socially. I don't know how to make a Peking Duck (or do we call it a Beijing Duck nowadays?), not even declaratively. But man and boy, I can't count the number of times I've succeeded in ordering Peking Ducks, or recommending Peking Ducks, or warning people off Peking Ducks. And I can get away with this because there's someone in the back room whom I trust to keep the phrase "Peking Duck" in line with its denotation.

How do you go about assigning an interpretation to my utterances about Peking Ducks, then? There are several aspects to the problem. In the first place, how do you decide what information I associate with the phrase when I use it, and hence what inferences you're entitled to draw about my internal state? Note that this is not the same as asking how much I know about Peking Ducks: Nor is it the same as asking how much my interlocutor and I can presuppose as common knowledge about Peking Ducks. To take another example (since as I said, I don't know a lot about Peking Ducks), suppose I'm a gallery owner and I'm trying to sell a particular acrylic painting to a client. I may know all sorts of things about acrylics: that they're quick-drying, easy to wash off, less noxious and less smelly than oil paints, and lacking in a certain tactility. And maybe the client knows that too. But in this context, none of that information is relevant to our interests, and you'd be making a big mistake if you tried to grind out all of the inferences that that kind of knowledge gives rise to. Whereas if I'm recommending a shift from oils to acrylics for a painter who has just come down with a rash, you would have to include some of that information to get the interpretation right.

In another way, however, the collective sense does play a role in the interpretation of all utterances, even when I am ignorant of it. Whatever my internal state vis-a-vis the world, I make certain social commitments about the world when I use an expression, and these are determined by the collective sense. Say you go into a restaurant and order a bottle of Sauternes, knowing only that it's a white wine; then when it comes, you discover that it's very sweet. Now you can ask the waiter to take it back on those grounds, I suppose, and he will oblige you if he's a good guy, but you'll be a little embarrassed about the exchange, because you realize he is not *obliged* to make the world conform to your idiosyncratic representation of it. He could say to you, "Listen, you ordered Sauternes, you got Sauternes. If you don't like it, get up, I need the table."

Now suppose we ask what kinds of linguistic authority we recognize as legitimate, and on what grounds? Or to put it another way, what kinds of information do we put in the collective sense? Here is where Johnson's assumptions about common sense come into play, and where Putnam comes a cropper. Because you can't just stick in any information that experts happen to have on hand; in fact, you can't even say by fiat who the experts are going to be. Suppose there's some weird kind of grape that's a member of the same biological kind as the grape that

Sauternes wine is made from, and suppose that that grape produces a wine that tastes and looks exactly like Gatorade. If the waiter brings me a bottle of that stuff, I send it back with full peremptory indignation, because I am certainly not committed to defer to biological authority in this matter if it is indifferent to my gastronomic interests. I defer to authority only where common sense tells me it is in my interest to do so. And I require of such authority that it define categories in ways that are consistent with my common-sense appraisal of what those interests are. Until you have fleshed out this story about common-sense deference, you can't determine just what commitments people are undertaking when they talk about the world.

3. A Note on Common-sense and formal semantics

As I noted, nobody has spent much time addressing the sorts of problems I've touched on here. The reason for this, I would suggest, is that linguistics and AI have been working under a certain methodological assumption that makes these things hard to get at. As Chomsky put it in his well known formulation of the goals of linguistic theory, linguistics is concerned with the competence of "an ideal speaker-listener in a homogeneous speech-community, *who knows its language perfectly....*" And while virtually every other clause of this formulation has been raked over by one or another critic, this one has escaped almost unnoticed, for good reason. The fact is that when you look at the sorts of linguistic phenomena most formal linguists have been interested in, and especially when you look at them with an eye towards describing their systematic formal properties, it isn't going to be particularly interesting to consider the problems raised by variation and imperfect knowledge. For one thing, variation in these domains doesn't often coincide with any interesting differences in social interests or social knowledge. Then too, from a purely structural point of view, an incomplete system is not interestingly different from a complete one. Questions about the social determination of meaning only arise when you look at domains in which linguistic representations co-vary in an interesting way with the social differentiation of knowledge.

This suggests an interesting way of understanding a part of the distinction between the methodologies of formal and common-sense semantics. I suspect that formal semantics is capable of making substantial progress in just those domains in which the social differentiation of knowledge is not relevant; in those domains, that is, in which common-sense beliefs and the collective sense amount to pretty much the same thing. But once you leave those domains for the wilds of the lexicon, you are going to have to do a lot of common-sense inferencing before you can even say what the meanings of expressions are, and how those meanings are relevant in their various ways to utterance interpretation.