Words, meanings, computations

Unedited posts from archives of CSG-L (see INTROCSG.NET):

Date: Mon Nov 02, 1992 3:54 pm PST
Subject: Words, meanings, computations

[From Bill Powers (921102.1230)] Ray Allis (921102.0830) --

I moseyed downstairs just now with a vague notion of putting a question to speakers and listeners on the net: when I talk about meanings being perceptions, am I just talking to myself? And there was your post waiting, showing that at least one person gets the point perfectly, and in fact already understood it. Thank you, Ray. [Reproduced at the end].

I've been reading, or trying to read between bouts of frustration, Gerald M. Edelman's _Bright Air, Brilliant Fire_. It isn't Edelman himself who is the object of my frustration, but the pretentious hogwash (I sound like Tom Bourbon) that he picks apart so carefully. He refers to "The blend of psychology, computer science, linguistics, and philosophy known as cognitive science..." Everything he finds wrong with this occupation can be traced back to not knowing the difference between a word and the perception to which it refers. There is also a tremendous muddle about what's outside the brain and what's inside it. On top of that there is the whole "intentionality" fiasco, which wouldn't even exist if philosophers of science hadn't decided once and for all that purposive behavior couldn't possibly exist. Mary started reading this book first, and she kept exclaiming that Edelman really needs control theory. Now that I've got into it I agree (Mary often leads the way into works like these -- I haven't the patience to get through them unaided).

I'm also trying to write a piece on PCT, and having trouble getting it organized. One thought that came up while reading Edelman, and that your post reinforced, is that what we mean in PCT by "computation" is completely different from what mainstream thinkers mean by it. The mainstream concept of computation is strictly symbol- oriented. You convert the problem into symbols, then manipulate the symbols to arrive at a conclusion, a recommendation, a description, and so on. This is the sort of thing we learn to do in algebra classes, and also when we learn to program digital computers. This is not at all where PCT started.

PCT grew up in the years from 1953 to 1960, while the digital revolution was still gathering speed. Bob Clark and I were running computer models back then, but they weren't digital computer models (that would have been like trying to run a real-time simulation on one of today's $5 pocket calculators). We used a Philbrick analogue computer. This may have more to do with the difficulties of getting PCT across than anything else.

In an analogue computer you don't assign symbols to variables and then do computations with the symbols. You make one variable physical quantity directly represent another. Just by BEING of a certain magnitude, a voltage can represent the magnitude of something else like a force, a speed, a position. An analogue computer setup is an array of wiring connecting physical components. The components make one signal depend on other according to continuous rules of computation such as integration, differentiation, multiplication, addition- subtraction, thresholds, and limits. Those computations are not done with numbers, but with the properties of physical things. A capacitor becomes an integrator when put in a feedback path; it becomes a differentiator when put in series with an input path. A logarithmic converter is made of a diode. A multiplier is made of a diode bridge. An adder or subtractor is an operational amplifier with multiple proportional inputs and proportional feedback. If you want to express the fact that velocity is the integral of acceleration, you provide a voltage that stands for acceleration and make it the input to an integrator. The output voltage of the integrator, properly scaled, varies as the velocity would vary. Run that into another integrator and you get position, again represented as a continuously-variable voltage.

When you put signals into an analog computing network, you get signal variations everywhere in the network, not just at the nominal outputs. All of these variations have meaning, because they stand in continuous relationship to the inputs, according to rules that are the properties of the intervening devices. There isn't a symbol anywhere in the network: just voltages and currents behaving through time. Of course the user of the computer assigns symbolic meanings to the various signals: this is a muscle tension, that is an acceleration, the other is a sum of forces. But the computer itself operates without any such interpretations.

With this orientation, it's not surprising that I should have started out in the beginning thinking of the nervous system as operating directly with signals, not with the names of signals or symbols for their states. It's not surprising that I recognized the difference between a perception and a word that indicates a perception. The perception is just a signal that varies as the inputs to some analogue computing device vary. It doesn't have to have a name. It isn't a member of a category. All signals are alike. Their meaning is in the way they are derived from other signals, or from sensors.

What we were doing in those early years is now called "connectionism." Connectionism is just analogue computing making a comeback.

Ultimately, of course, we had to recognize the existence of words and symbols, and categories, and other things that seem to be perceived as discrete packages. But the same orientation still seems appropriate. A word is an activity in the brain; therefore it must be represented as a signal. A word indicates a meaning when it gives rise, through some continuously-present device, to another signal in a pathway that normally handles direct perceptions. Obviously we can't express the nature of the devices that connect the signals in terms of the signals; that would be like trying to draw a diagram of a radio that's playing music by using musical notation. Words and symbols are the signals handled by these devices; they aren't the devices. In the products of the brain we can recognize relation- ships that suggest processes in the brain, and the processes are created by the computing components, not the signals. The signals and their relationships, which are more signals, are only indications of the processes. I've been trying to express the understanding behind these words for 40 years, and still can't do it right. It's always been obvious to me that the conventional sciences are trying to explain how behavior works by looking at the CONTENT of a brain, the signals that happen to be present in it -- and particularly those that happen to be present in the theoretician or experimenter. The experimenter tries to figure out how one signal behaves when another changes, but fails to realize that it is the functional connections, not the signals, that matter. When we observe a relationship between two perceptions, the important thing is not what that relationship is, or what the perceptions are, but that we ARE PERCEIVING A RELATIONSHIP. That, not the particular signals, is evidence about how the brain works.

This is particularly true at the "cognitive" levels. All of our digital-computer models and the complex mathematics that go with them are evidence about how the brain works, in terms of what it is doing. In themselves they are of little importance. An incorrect deduction is just as good evidence as a correct one. What we should be asking is not what cognitive calculations we can carry out or what new computations they lead to, but WHAT WE ARE DOING WHEN WE CARRY THEM OUT. We can't understand the computer just by looking at more computations. When we let the particular computations occupy our attention, we're missing seeing the principles of computation themselves. It's like starting out to understand the computer and getting hung up on a story that it's printing out. We have plenty of printouts already. We have to start looking for something else, something that explains the very existence of the printouts, not their content.

Well, you see what a bit of encouragement does to me. I don't know that I've managed to say these things any better this time, but it feels good to talk about them.

Is there anybody else out there besides Ray who gets the point?

Best, Bill P.

Date: Mon Nov 02, 1992 10:06 am PST
Subject: Re: Meanings and sentences

[From Ray Allis (921102.0830)] [Bill Powers (921027.0700)]

>> Variable aspects of sentences, for example (like meaning, structure, inflection, etc) are controlled INPUTS, not generated outputs.

> I want to make a modest push to make sure our linguists really truly get this point. They may actually get it, but I don't think we're hearing the result. Sentence construction is not construction of some object out there, or in some vague conceptual space; it's construction of an input, a perception. The mere fact that we know of a sentence shows that it's a perception. The same goes for grammar, for any regularity we PERCEIVE in language. The relationships between different levels of analysis of language are relationships among levels of perception, not levels of output production. We have to guess what the production processes are, because they're outputs and we don't perceive outputs. We perceive only their perceptual consequences: language is perception.

>[From Bill Powers (921025.1900)]

> P.S. Given three nouns A, B, and C, you can create the sentences

> The A B'd the C, and
The B A'd the C.

> The captain hogged the ice-cream, and
The hog captained the ice-cream.

> The only way to judge that one is allowable in ordinary discourse and the other is not is through the meanings of the words. That can't have anything to do with language itself, because the fact that hogs don't captain ice-cream is an accident of experience; given other experiences, we might see immediate sense in the second sentence but not the first, depending on what perceptions "hog" and "captain" are attached to, and what we can imagine hogs and captains doing. There's something about the way words are presented that connects to the way meanings are evoked in juxtaposition and sequence -- that's the input function above. But once the meanings are evoked, the rules that apply aren't those of language, but of perceived reality in general. We don't judge that hogs can't captain things because of language, but because of our experiences with hogs and captains.

> That must sound like utter nonsense.

> Best, Uncle Bill P.

"But once the meanings are evoked, the rules that apply aren't those of language, but of perceived reality in general."

Gee, "Uncle Bill", I used to think I was the only one who thought in this utterly nonsensical way. The implications of this and some other related trains of thought are, in my experience, as invisible to "artificial intelligence" people as PCT is to psychologists. Maybe you'll be interested in this piece of something I haven't even *tried* to publish.

9.2 An example of analogy

An example of analogy (or more properly simile) is the phrase "Life is like a river". This sort of example has been criticized as being unrealistically difficult for computers. This phrase and many others are difficult even for humans. But that is the goal of AI, isn't it? The duplication and extension of human intellectual abilities?"

"Life is like a river." The "meaning" of this string of symbols has to do with time and change and irreversibility, and with some things sensed and not so easily put into words, and this is the problem. Some or all of the similarity we perceive is non-linguistic, so that we cannot express these perceptions very well. (I am standing in the shallows, cold water pressing against my shins through the rubber of my hip boots, with a fly rod in my right hand.) I cannot express at all well in words my experience of "fitness" when I consider the similarity of the river and life. Much of my feeling is just that, feeling; non-verbal motor and sensory perception. Like the ability to ride a bicycle, experiences are not directly transferrable from mind to mind but must be acquired by each individual. We can help each other, but we cannot give such experiences to others in the same way we can tell them how to multiply two numbers (i.e. manipulate deductive systems).

My concept and interpretation of this simile will certainly be different from yours, but sufficiently similar, we both believe, to enable communication. Our belief that we have communicated is an act of induction, based on our perceptions of behavior in others as being similar to our own behavior, and the hypothesis that similar actions are produced by similar motivations.

"Correctness" and "fit" have very different meanings here than in symbolic logic, no matter how "fuzzy" such logics are made. After all, what sense does it make to attempt to position symbols on scales of characteristics the symbols themselves do not have?

Similarity and analogy are relationships among representations (Section 7). The use of deduction (the treatment of the relationships among symbols) sacrifices the ability to perceive "likeness". You lose all the meaning!

Best to you, Ray