Grounded Concepts Without Symbols

George Lakoff

Linguistics Department
University of California at Berkeley
Berkeley, California 94720
lakoff@cogsci.berkeley.edu

Since the publication of Leonard Talmy's classic paper, *The Relation of Grammar to Cognition* in 1978, the study of spatial relations concepts has become a major industry in cognitive science. Talmy has worked out the basic outlines (Talmy, 1983, 1988), and enormous additional insights have been provided by Ron Langacker (1987), Susan Lindner (1981), Eugene Casad (1982), Annette Herskovits (1986), Claudia Brugman (1988) and Laura Janda (1986). An entire research group has been set up at Nijmegan under Steve Levinson to extend these results. Spatial relations concepts stand as one of the most richly studied areas of conceptual structure.

Among the basic results obtained so far are these: Spatial relations concepts in the languages of the world appear to be decomposable into a relatively small number of primitive "image schemas" -- containers (as with IN and OUT), paths (as with FROM, ALONG, TO), contact (as with ON), and other schemas such as relative distance, front-back, up-down, center-periphery, etc. These image-schemas have been shown by Talmy to have topological, orientational, and force dynamic properties, which, as I have observed, give rise to their inferential properties and characterize the logic of space (Lakoff, 1987, Case study 2).

In the best philosophical work on grounding that I know, Mark Johnson's *The Body in the Mind*, Johnson has argued convincingly that image-schemas are grounded in the body and arise from preconceptual structures of experience (Johnson, 1987). He argues further that such grounding is sufficient in the case of image-schemas to define their inferential structure and constrain the possibilities for conceptual combination.

Conceptual metaphor is another area where there has been a boom in empirical research over the past decade. The basic results are these: (1) Abstract concepts are characterized via metaphorical mappings from more concrete concepts. (2) Metaphorical mappings preserve image-schemas, and (3) Abstract inferences are metaphorical mappings of spatial inferences. (4) Conceptual metaphorical mappings are grounded in everyday experiences -- bodily and interpersonal experiences. (Lakoff and Johnson, 1980; Lakoff, 1987, Case Study 1; Turner, 1987; Lakoff and Turner, 1989; Sweetser, 1990; Turner, 1991; Lakoff, in press.)

Put together, the study of spatial relations concepts and conceptual metaphor provide the following picture of conceptual structure: Spatial relations concepts are directly grounded in the body and the logic of spatial concepts comes out of the topological, orientational, and force dynamic nature of those concepts. These are extended to abstract concepts via metaphorical mappings, which are themselves grounded in everyday experience.

These results raise a question: Exactly what does it mean for a concept to be grounded in the body? What, precisely, is grounding? And, most important, how do the inferential properties of concepts arise from their grounding? Recent research by Terry Regier at Berkeley has begun to answer those questions.

Regier has built a structured connectionist model that learns and represents a large range of spatial relations concepts in the world's languages. The model learns to pair spatial relations terms with visual images (represented in an N-by-N array of neural units -- a much oversimplified model of a retinotopic map). The crucial aspect of the model that accomplishes the pairing of terms with images is a structured connectionist architecture that models aspects of the brain's visual system: topographic maps, orientation-sensitive cells, and center-surround architectures. In a sense, the model builds in the relevant aspects of the body to provide bodily grounding for spatial relations concepts. Thus, the topological properties of spatial relations concepts arise from the use of topographic maps, the orientational properties from the use of orientation-sensitive cells, and so on. In this way the properties of the spatial relations concepts are "embodied."

For example, Regier's model uses a structured connectionist model of a topographic map of the visual field to compute interiors of objects. The relation IN is characterized in terms of overlap with and INTERIOR, and OUT is characterized in terms of nonoverlap. A center-surround architecture is used to characterize CONTACT, as in the concept ON. Orientation-sensitive cells are used to characterize vertical orientation, which plays a major part in the characterization of ABOVE.

Regier's model represents concepts without "symbols" in the ordinary sense of the term, as used in physical symbol systems or formal logics. The model contains concepts which are grounded directly in the structure of the model. The constraints on the concepts are not characterized using symbolic representations, but rather are built into the grounding of the model. The spatial concepts structure visual scenes in the form of retinotopic maps directly -- something that symbols cannot do.

All this suggests that "symbols," in the sense of physical symbol systems, are not needed at all to represent concepts. Spatial concepts, and other concepts with a bodily grounding, should be characterizable nonsymbolically with models like Regier's, and abstract concepts can be characterized by metaphorical mappings from such inherently grounded concepts. The key to making all this work for both learning and representation is the use of structured connectionist models of the sort long proposed by Jerry Feldman, where the structure characterizes the relevant aspects of bodily grounding. (Feldman, to appear)

What are "Symbols"?

The term "symbol" as used in cognitive science is taken from certain technical disciplines: symbolic logic, the theory of formal grammars, and physical symbol systems. In all of these, a symbol is an arbitrary mathematical element, which can be represented by an arbitrary "sign" -- say a letter of the alphabet or a binary number.

In symbolic logic, symbols are concatenated into strings called well-formed formulas, and deductions -- or sequences of such strings -- are characterized by logical axioms and meaning postulates (which are also strings of symbols) and by rules of inference (which are sequences of strings). The well-formedness principles, logical axioms, and meaning postulates state constraints on how the symbols can be used. These are stated in purely syntactic terms, without referring to the meaning of the symbols at all.

All of these symbols and symbol strings are meaningless. We are taught that the symbols are to stand for concepts -- but they are not concepts in themselves. Symbols in symbolic logic are purely syntactic in nature, and have no inherent meaning. The meaning of the symbols never enters into the formal deductions that characterize "inferences."

The same is true of the symbols in physical symbol systems of the sort used in classical AI. To characterize an inference in classical AI, one begins with a string of symbols and applies a computer program (a sequence of symbol strings) to it to produce another string of symbols. All of this is purely syntactic. The meaning of the symbols never enters the picture. What are called "inferences" in classical AI are thus meaning-free: they are simply the results of string manipulations. Moreover, the symbols in themselves do not in any way constrain how they can be manipulated by computer programs in the production of "inferences."

In classical AI, as in symbolic logic, symbols are taken as standing for concepts. But it is vital to remember that they are not concepts in themselves. Concepts are inherently meaningful, while symbols are not. In designing an AI system, syntactic constraints on the combinations of symbols are placed on the system, so that symbols will combine in ways that the designer of the system intends. In a system like Regier's, however, such syntactic constraints are unnecessary because the very way the concepts are grounded constrains how the concepts can combine. In short, the grounding of the concepts characterizes what the concepts mean, thus constraining how the concepts can combine with other concepts. No symbols are needed to do this. Human concepts also are constrained in how they can combine with other concepts by what they inherently mean. Human concepts therefore have very different properties than do the arbitrary symbols of physical symbol systems.

Physical symbol systems are quite useful, and certainly not harmful, in computer science. But when one moves to cognitive science, where the issue is how the mind really works, the picture changes. In AI-style models of the mind, symbols are commonly taken as *being* concepts. That is a mistake, since concepts are inherently meaningful, while symbols are not.

A major challenge for cognitive science is to give a theory of what concepts are -- how they are grounded, how their inherent meaning is a consequence of their grounding, and how constraints on their conceptual combination follows from their grounding. Regier's model does this for spatial relations concepts. It does it by embodying the concepts in the model -- through the use of biologically-based structured connectionism. The spatial relations concepts are those aspects of the model that play a crucial role in computing the spatial relations.

Incidentally, Regier's model contains some things that may look like symbols, but technically are not symbols, in the sense of physical symbol systems. Regier's output nodes characterize spatial terms that name the spatial relations concepts computed by the model. Those nodes represent names for grounded concepts. Those concepts exist independently of their names and their inferential properties are built into the way they are characterized in the model. The topological properties come out the use of topographic maps; the orientational properties come out of the use of orientation-sensitive cells.

One could not just take the names for those concepts and use them the way symbols are used in physical symbol systems. That is, one could not just apply arbitrary computer programs to them to yield "inferences." Because they are names for grounded, and hence, inherently meaningful concepts, they have inherent constraints that follow from their grounding on what inferences they can enter into. "Symbols," in the sense of physical symbol syste' .s, are not like that.

Thus, what Regier's model provides are GROUNDED CONCEPTS WITHOUT SYMBOLS! This is exactly what the theory of concepts requires. And it gives a precise idea of what it means for concepts to be embodied.

Overview

This paper is far more than mere philosophical speculation about the nature of conceptual grounding. It is based, first, on the extensive empirical findings on the nature of spatial relations concepts and on conceptual metaphor. Those empirical findings are at odds with classical symbolic approaches to what concepts are. Second, it is based on an actual computational model in which spatial relations concepts are grounded in a biologically-motivated way. That model is not a mere kluge, but incorporates the important insight that the topological and orientational properties of spatial concepts arise from topographic maps and orientation-sensitive cells in the human visual system.

At present, Regier's model is the only one that can learn and represent spatial relations concepts in such a way as to pair spatial terms with visual scenes. I find it unlikely that any other model will be successful without incorporating Regier's fundamental insights about the biological basis of the grounding of spatial relations concepts.

Any appropriate reply to this paper should take into account both the empirically discovered properties of spatial relations concepts and the fact that Regier's model, which incorporates biologically-based insights about the physical origin of those properties, actually works, where no other model does.

Finally, this paper should be distinguished sharply from approaches to what has been called "symbol-grounding." The symbol-grounding approach appears to accept the idea that concepts require a symbolic representation in which principles of conceptual combination are stated. It is that idea about the nature of concepts that this paper denies. Regier's approach suggests that the whole level of symbolic manipulation is unnecessary, since the meaning of concepts is a consequence of conceptual grounding, and it is that grounding that places constraints on conceptual combination.

A Final Note

I began this paper with a discussion of empirical research concerning spatial relations concepts and conceptual metaphor. Research in these areas is very much at odds with the symbolic view of concepts, and points toward a theory of concepts grounded in the body and in bodily interactions in the world. Anyone at all interested in grounding should be familiar with this literature, since it places a considerable range of constraints on what grounding can be.

References

Brugman, Claudia. 1988. *The Story of* Over: *Polysemy, Semantics, and the Structure of the Lexicon. New York: Garland.*

Casad, Eugene. 1982. *Cora Locationals and Structured Imagery*. Ph.D. Disstertation. University of California at San Diego.

Feldman, Jerome. To appear. Structured Connectionist Models and Language Learning. International Computer Science Institute, University of California at Berkeley.

Hershkovits, Annette. 1986. *Language and Spatial Cognition: An interdisciplinary study of prpositions in English*. Cambridge: Cambridge University Press.

Janda, Laura. 1986. *A Semantic Analysis of the Russian Verbal Prefixes ZA-, PERE-, DO-, and*

OT-. Band 192. Slavistische Beitrage. Munchen: Verlag Otto.

Johnson, Mark. 1987. *The Body in the Mind.* Chicago: University of Chicago Press.

Lakoff, George and Mark Johnson. 1980. *Metaphors We Live By.* Chicago: University of Chicago Press.

Lakoff, George. 1987. *Women, Fire, and Dangerous Things.* Chicago: University of Chicago Press.

Lakoff, George and Mark Turner. 1989. *More Than Cool Reason.* Chicago: University of Chicago Press.

Lakoff, George. In Press. *The Contemporary Theory of Metaphor.* In A. Ortony (ed.), *Metaphor and Thought.* Cambridge: Cambridge University Press.

Langacker, Ronald. 1987. *Foundations of Cognitive Grammar: Vol I.* Stanford: Stanford University Press.

Lindner, Susan. 1981. *A Lexico-Semantic Analysis of Verb-Particle Constructions with UP and OUT.* Ph. D. Dissertation. University of California, San Diego.

Regier, Terry. 1992. The Acquisition of Lexical Semantics for Spatial Terms: A Connectionist Model of Perceptual Categorization. Ph.D. Thesis, Department of Computer Science, University of California at Berkeley. Technical Report TR-92-062, International Computer Science Institute, 1947 Center Street, Berkeley, CA 94704-1105. E-mail: regier@icsi.berkeley.edu

Sweetser, Eve. 1990. *From Etymology to Pragmatics.* Cambridge: Cambridge University Press.

Talmy, Leonard. 1978. The Relation of Grammar to Cognition. In D. Waltz (ed.), *Proceedings of TINLAP-2 (Theoretical Issues in Natural Language Processing). Champaign, Ill.: Coordinated Science Laboratory, University of Illinois.*

Talmy, Leonard. 1983. How Language Structures Space. in H. Pick and L. Acredolo, eds., Spatial Orientation: Theory, Research, and Application. New York: Plenum Press.

Talmy, Leonard. 1988. Force Dynamics in Language and Cognition. In *Cognitive Science* 12, 49-100.

Turner, Mark. 1987. *Death is the Mother of Beauty.* Chicago: University of Chicago Press.

Turner, Mark. 1991. *Reading Minds.* Princeton: Princeton University Press.