UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Where Do Relations Come From?

Permalink
https://escholarship.org/uc/item/4qh579wz

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 19(0)

Authors
Colunga-Leal, Eliana
Gasser, Michael

Publication Date
1997

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Where Do Relations Come From?

Eliana Colunga-Leal (ECOLUNGA@CS.INDIANA.EDU)
Michael Gasser (GASSER@CS.INDIANA.EDU)
Computer Science Department, Cognitive Science Program
Indiana University, Lindley Hall 215, Bloomington, IN 47405 USA

Natural language, and much of human knowledge, seems to revolve around relations. But while relations are part and parcel of symbolic accounts of language and thought, it is not so clear where they stand when we move down into perception. How would a visual system, along with a cognitive/linguistic system, learn to map input in which there are no explicit relations onto expressions such as *the book is on the bed*?

In connectionist networks, the problem of representing relations is an instance of the *binding problem*: how is the system to keep track of which features go with which? One popular solution to this problem is *synchronization* (Shastri and Ajjanagadde, 1993), that is, the alignment of units along a dimension separate from their activation, a "relative phase angle" (RPA). This alignment effectively represents the notion of "same thing."

There are two ways in which it is possible to know a relation. First, one can know that objects of a certain sort tend to be associated with objects of another sort; that is, one can make inferences about kinds of objects on the basis of other objects. This sort of knowledge could be implemented in the form of connections between units that are outfitted with RPAs as well as activation. But more complex inferences are also possible, inferences about what kinds of objects play particular role in relations or about how particular relations tend to be associated with other relations. An example of the former is the inference that if an object A is on an object B, then B tends to have a flat upper surface. An example of the latter is the inference that if an object A is on an object B, then A tends to be smaller than B. However, there is no way to incorporate the ON relation into a single conventional network unit, and thus no way for a connection to represent the sort of inferences we are interested in.

What is needed for this second sort of relational knowledge, we argue, is explicit *micro-relations*, which are activated to the extent that two different objects with particular sorts of features are activated. We present a simple connectionist model of the learning of spatial concepts and spatial language in which micro-relations are pre-wired clusters of units. The network, a generalized Hopfield network augmented with RPAs, is trained to auto-associate input patterns, either visual patterns or words or both.

In the visual system we assume the conventional division between the "what" and "where" components but extend this division into the lexicon as well. The what side is responsible for identifying object categories, ultimately nouns, in visual input, while the where side is responsible for identifying spatial relations, ultimately spatial words, in visual input. In an initial phase of training, the network is exposed to visual inputs only, and it learns the relational regularities in the inputs in its layer of spatial micro-relations. During a second phase, spatial phrases are presented along with the visual inputs, and the network learns to associate prepositions with spatial micro-relations. Note that the relationship between spatial expressions such as *a book on the bed* and their meanings is an example of a relation-to-relation inference and seems to require explicit micro-relations to be represented and learned. The figure shows the representation of the phrase *book on bed* across a subset of the units in the network (only a few of the relevant connections appear).

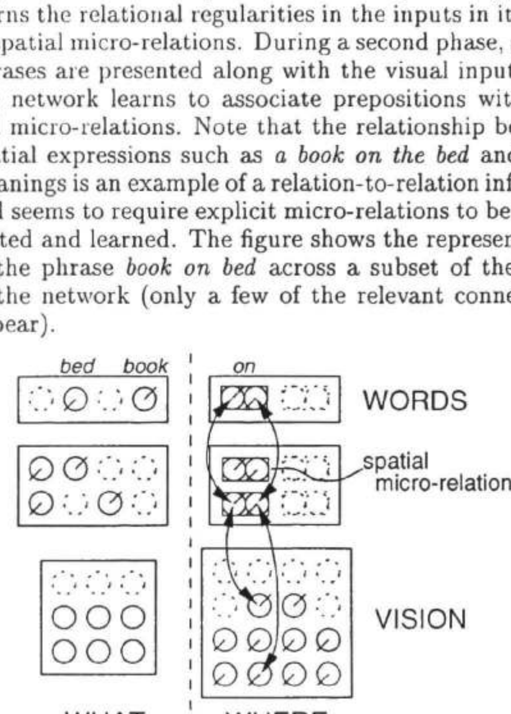

In the model, spatial relational words can only be learned if the appropriate relational regularities are learned in the spatial micro-relation layer. Several predictions follow from this property of the model. The model also has implications for the mutual influence of language and visual-spatial cognition.

References

Shastri, L. & Ajjanagadde, V. (1993). From simple associations so systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences, 16*, 417–494.