Time phases, pointers, rules and embedding

John A. Barnden

Extracted with minor edits from *Behavioral and Brain Sciences, 16*(3), pp.451–452 (1993).

(It's a commentary on the target article by Lokendra Shastri & Venkat Ajjanagadde: "From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony" in same issue of the journal, pp.417–451.)

Binding by time phases is an interesting special case of the following very general (temporary) binding method: To bind two things, mark them in roughly the same way. Let's call this the similar-mark approach. Note that it could apply to non-connectionist as well as connectionist systems. In Shastri & Ajjanagadde's (S&A's) case, we may take the marks to be the oscillatory patterns of excitation acquired by argument nodes and so on. Two marks are similar enough to constitute a binding if they have sufficiently similar phases (and frequencies). So, S&A's method is a special case of the approach of temporarily binding connectionist nodes or subnetworks together by dynamically making them hold activation patterns that are similar enough in some specific sense. That is, the time-phase method is a special case of "pattern-similarity association" or PSA (Barnden & Srinivas 1991). This is in turn a connectionist special case of the similar-mark approach.

A benefit of considering the time-phase method in the context of PSA and similar-mark binding in general is that we see the close relationship to the technique of "associative addressing' widely used in specialized computer hardware as an alternative to pointers. With this technique two memory areas can be temporarily "linked" together by placing identical or sufficiently similar bit-strings somewhere within them. Such a bit-string extracted from one place can be used to find the other place or places that contain that bit-string or suitably similar ones. In sum, S&A's system, which is one of the few connectionist systems that can actually perform inferencing of any respectable complexity, turns out to rest on a binding scheme quite strongly related to a conventional computer technique, but without using any close analogue of pointers.

We are being led here to the question of what happens to the notion of a pointer when we move away from computers. This question is examined to some extent by Barnden and Srinivas (1991). One can define a pointer in a connectionist system to be a temporary system substate (e.g., activation pattern) that identifies a permanently existing place in the system. However, without specific architectural assumptions we cannot say what a "place" is, other than by resting on the excessively restrictive option of a place just being a single node or on the excessively loose option of a place being any subset of the system's nodes. Going back to S&A, if the phase assigned to the John node or assembly, say, were fixed for all time, then phases could be regarded as pointers, because they would permanently identify such nodes or assemblies. However, S&A allow phases to be dynamically assigned, so they are probably radically different from "pointers" under any usefully narrow construal of that word. The signature scheme in ROBIN (Lange & Dyer 1989) is more pointerlike, because signatures are statically assigned.

One major benefit of similar-mark techniques is that they allow bidirectional binding in two senses. (1) A binding could be conceptually bidirectional; one might, for instance, say that if several S&A argument nodes have the same phase then they are bidirectionally bound to each other. (2) A binding, though perhaps conceptually unidirectional, could be used bidirectionally. For instance, a node oscillating at a certain phase might broadcast its oscillation to other nodes, thereby causing

similar-phased nodes to light up in some special way, but the same thing could be done starting at any of those nodes (see also Touretzky 1990). By contrast, computer pointers can only efficiently be used in one direction.

A disadvantage of many similar-mark schemes, however, is that if a binding is conceptually unidirectional, one needs something extra to specify direction. That something could be highly implicit in the overall system architecture; thus, the binding between an S&A argument node and the John node is, arguably, conceptually unidirectional, and that fact is implicitly respected in the whole way that the system operates. However, if one needed unidirectional bindings between argument nodes for some reason, one would need to do something more than simply give them the same phase.

A concern I have about many connectionist systems, including S&A's, is that they may face difficulty in encompassing certain important types of reasoning, including some reflexive types [reflexive in the sense of a reflex]. S&A claim it is unlikely that the input propositions to reflexive reasoning episodes can be dynamically arising rules. I take the point of their syllogism example, but there are more mundane examples that are not so easily disposed of. For instance, suppose someone says, "All the people at the party were toothbrush salespersons. Some of them even had their sample cases with them." The obvious inference that those cases contained toothbrushes seems no less a candidate for being dubbed "reflexive" than do the inferences in S&A's Little Red Riding Hood and Colombian drug enforcement agency example. Yet one of the input propositions is the universally quantified one about all the people in the room, and this can be viewed as a dynamically arising rule.

Now suppose someone says, "Tom thought that the milk was sour. He went out to buy some more." One needs to be able to apply the general knowledge that sour milk tends to be unusable for certain purposes, together with Tom's reported thought, in order to understand Tom's motive in going out for more milk. Is such reasoning not reflexive? That is, I am suggesting that input propositions and reasoning episodes involving them can be embedded in propositional attitude contexts (among other sorts of context, such as counterfactual ones) without making the reasoning nonreflexive. This makes the task of connectionistically implementing reflexive reasoning yet more complex.

Embedding and dynamically arising rules are discussed further in Barnden (1992, pp. 149-78). Because very few workers in connectionism, or indeed critics of connectionism, have even paid lip service to the issues, my comments are hardly a strike against S&A specifically.

References:

Barnden, J. A. (1992) Connectionism, generalization and propositional attitudes: A catalogue of challenging issues. In: *The symbolic and connectionist paradigms: Closing the gap*, ed. J. Dinsmore. Erlbaum.

Barnden, J. A. & Srinivas, K. (1991) Encoding techniques for complex information structures in connectionist systems. *Connection Science 3*(3): 263–309.

Lange, T. E. & Dyer, M. G. (1989) High-level inferencing in a conneetionist network. *Connection Science 1*(2): 181–217.

Touretzky, D. (1990) BoltzCONS: Dynamic symbol structures in a connectionist network. *Artificial Intelligence 46* (1-2), 5–46.