An Extended Local Connectionist Manifesto: Embracing Relational and Procedural Knowledge

Lokendra Shastri

International Computer Science Institute

1947 Center Street

Berkeley, CA 94704 U.S.A

shastri@icsi.berkeley.edu

http://www.icsi.berkeley.edu/~shastri

Abstract

The author has performed an important service by dispelling several myths and misconceptions concerning the localist approach. The localist position and computational model presented in the target article, however, are overly restrictive and do not address the representation of complex conceptual items such as events, situations, actions, and plans. Working toward the representation of such items leads to a more sophisticated and articulated view of the localist approach.

The author has performed an important service to the connectionist and cognitive science community by dispelling several myths and misconceptions concerning the localist approach. Over the years, these myths and misconceptions have acquired the status of self-evident truths and it has become almost obligatory for connectionist models to display the "distributed representations inside" logo in order to lay claim to being the genuine article.

I enthusiastically concur with the bulk of what the author has to say, but I would like to elaborate on the localist approach outlined in the target article based on my own involvement with the approach. In my opinion, the localist position and the localist computational model presented in the target article are overly restrictive. The author has focused primarily on the representation of entities that can be expressed as soft (weighted) conjunctions of features. The localist model described in Section 4 deals almost exclusively with the acquisition and retrieval of higher-level concepts (nodes) that are soft conjunctions of lower-level features. Even the more advanced example discussed in Section 4.5 focuses on learning associations between such entities. This limited representational focus is also reflected in the examples of entities enumerated by the author, namely, "words, names, persons, etc." (target article, page 9). What the author leaves out are more complex conceptual items such as events, situations, actions, and plans which form the grist of human cognition.

Events, situations, actions, and plans involve *relational* and *procedural* knowledge, and hence, cannot be encoded as mere soft conjunctions of features; their encoding requires more *structured* representations. Working toward a representation of such complex and structured items leads to a more articulated view of the localist approach than the one presented in the target article. I will briefly comment on this view. For more details, the reader is referred to specific models that instantiate this view (see, Shastri, 1991; 1997; 1999a; 1999b; 1999c; Shastri and Ajjanagadde, 1993; Shastri, Grannes, Narayanan, and Feldman, in press).

In the enriched representational context of events, situations, actions, and plans the operative representational unit is often a *circuit* of nodes rather than a node. Moreover, only some of the nodes in such a circuit correspond to cognitively meaningful entities (as the latter are characterized in Section 2.5, page 9). Most of the other nodes in the circuit serve a *processing* function or perform an ancillary representational role. For example, such nodes glue together simpler items in systematic ways to form composite relational items, they provide a handle for systematically accessing specific components of a composite item, and they allow actions and plans to be expressed as partially-ordered structures of sub-actions and sub-plans. Thus the encoding of an event (E1) "John gave a book to Mary" in long-term memory would involve not only nodes corresponding to cognitively meaningful entities such as John, Mary, book, giver, recipient, and object, but also *functionally* meaningful nodes such as: a node for asserting belief in E1, a node for querying E1, *binder* nodes for encoding role-entity bindings in E1 (for example, a node for encoding the binding ⟨giver=John⟩), *binding-extractor* nodes for selectively retrieving role-fillers in E1 (for example, a node for activating "John" in response to the activation of the role "giver" in the context of E1), and nodes for *linking* the encoding of E1 to a generic perceptual-motor schema for the *give* action. Furthermore, the localist encoding of the *give* schema would involve specific nodes and circuits for encoding a partially ordered sequence of perceptual-motor sub-actions comprising the *give* action.

In the expanded localist framework individual nodes continue to have well-defined localist interpretations. However, these interpretations are best couched in terms of a node's *functional* significance rather than its semantic significance (cf. Section 2.5, page 9).

The learning framework presented by the author has a strong overlap with work on recruitment learning (Wickelgren, 1979; Feldman, 1982; Shastri, 1988; 1999b; 1999c; Diederich, 1989; Valiant, 1994). The architecture described in Figure 9 of the target article is in many ways analogous to that sketched out in (Shastri, 1988, pp 182-192). This overlap merits further exploration. In the recruitment learning framework, learning occurs within a network of quasi-randomly connected nodes. Recruited nodes are those nodes that have acquired a distinct meaning or functionality by virtue of their *strong* interconnections to other recruited nodes and/or other sensorimotor nodes. Nodes that are not yet recruited are nodes "in waiting" or "free" nodes. Free nodes are connected via weak links to a large number of free, recruited, and/or sensorimotor nodes. These free nodes form a primordial network from which suitably connected nodes may be recruited for representing new items. For example, a novel concept y which is a conjunct of existing concepts x_1 and x_2 can be encoded in long-term memory by "recruiting" free nodes that receive links from both x_1 and x_2 nodes. Here recruiting a node simply means strengthening the weights of links incident on the node from x_1 and x_2 nodes. In general, several nodes are recruited for each item.

The recruitment process can transform quasi-random networks into structures consisting of nodes tuned to specific functionalities. Typically, a node receives a large number of links, and hence, can potentially participate in a large number of functional circuits. If however, the weights of selected links increase, and optionally, the weights of other links decrease, the node can become more selective and participate in a limited number of functional circuits.

In (Shastri, 1999b; 1999c) it is shown that recruitment learning can be firmly grounded in the biological phenomena of *long-term potentiation* (LTP) and *long-term depression* (LTD) that involve rapid, long-lasting, and specific changes in synaptic strength (Bliss and Collingridge, 1996; Linden, 1994). Moreover, as explained in (Shastri, 1999c) the specification of a learning algorithm amounts to choosing a suitable network architecture and a set of appropriate parameter values for the induction of LTP and LTD.

The recruitment learning framework also offers an alternate explanation for the age-of-acquisition effect discussed in Section 4.4. It suggests that (on an average) more cells are recruited for items acquired earlier in a learning cycle than for items acquired later in the cycle. Thus items acquired earlier in the learning cycle have greater neuronal mass and it is this greater mass that gives these items their competitive edge.

To conclude, the author must be commended for systematically and comprehensively presenting a strong case for the localist models. The localist position and computational model presented in the target article, however, can be enriched by considering the representation of complex items involving relational and procedural knowledge. Work on representing such items leads to a more articulated view of the localist approach than that presented in the target article.

References

Bliss, T.V.P. and G.L. Collingridge (1993). A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361, 31–39.

Diederich, J. (1989). Instruction and high-level learning in connectionist networks. *Connection Science* 1, 161–180.

Feldman, J. A. (1982). Dynamic connections in neural networks. *Bio-Cybernetics* 46, 27–39.

Feldman, J. A. (1989). Neural Representation of Conceptual Knowledge. In *Neural Connections, Mental Computation* ed. L. Nadel, L.A. Cooper, P. Culicover, and R.M. Harnish. MIT Press.

Linden, D.J. (1994). Long-term synaptic depression in the mamallian brain. *Neuron* 12, 457–472.

Shastri, L. (1988). *Semantic Networks: An evidential formalization and its connectionist realization.* Morgan Kaufmann, Los Altos/Pitman Publishing Company, London, 1988.

Shastri, L. (1991). Relevance of Connectionism to AI: A representation and reasoning perspective. In *Advances in Connectionist and Neural Computation Theory*, vol. 1, J. Barnden and J. Pollack (Ed.), Ablex, Norwood, N.J. 1991. pp 259–283.

Shastri, L. (1996). Temporal Synchrony, Dynamic Bindings, and SHRUTI: a representational but non-classical model of reflexive reasoning, *Behavioral and Brain Sciences* 19:2, 331-337, 1996.

Shastri, L. (1997). A Model of Rapid Memory Formation in the Hippocampal System, *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, pp. 680–685, Stanford University, CA. August 1997.

Shastri, L. (1999a). Advances in SHRUTI — A neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony., *Applied Intelligence*, 11, 79–108, 1999.

Shastri, L. (1999b). Recruitment of binding and binding-error detector circuits via long-term potentiation, *Neurocomputing*, 26-27, 865–874, 1999.

Shastri, L. (1999c). A Biological Grounding of Recruitment Learning and Vicinal Algorithms. Technical Report TR-99-009, International Computer Science Institute, Berkeley, CA, 94704. April, 1999.

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

Shastri, L., D.J. Grannes, S. Narayanan, and J.A. Feldman (in press). A Connectionist Encoding of Parameterized Schemas and Reactive Plans, in *Hybrid Information Processing in Adaptive Autonomous vehicles,* G.K. Kraetzschmar and G. Palm (Eds.), Lecture Notes in Computer Science, Springer-Verlag, Berlin.

Valiant, L. (1994). *Circuits of the mind.* New York : Oxford University Press.

Wickelgren, W. A. (1979). Chunking and Consolidation: A Theoretical Synthesis of Semantic Networks, Configuring in Conditioning, S-R Versus Cognitive Learning, Normal Forgetting, the Amnesic Syndrome, and the Hippocampal Arousal System. *Psychological Review*, 86 (1): 44-60.

Endnotes

1. This work was supported in part by NSF grants SBR-9720398 and ECS-9970890.