Full Day Workshop
Compositional Connectionism in Cognitive Science II: The Localist/Distributed Dimension

Ross W. Gayler (r.gayler@gmail.com)
Philosophy Program, La Trobe University
Bundoora, VIC 3086 Australia

Simon D. Levy (levys@wlu.edu)
Computer Science Department, Washington & Lee University
Lexington, VA 24450 USA

Keywords: Distributed representations, localism, structured representations, connectionism, cognitive modeling

Background

In October 2004, approximately 30 connectionist and non-connectionist researchers gathered at a AAAI symposium to discuss and debate a topic of central concern in artificial intelligence and cognitive science: the nature of compositionality. The symposium offered participants an opportunity to confront the persistent belief among traditional cognitive scientists that connectionist models are, in principle, incapable of systematically composing and manipulating the elements of mental structure (words, concept, semantic roles, etc.). Participants met this challenge, with several connectionist models serving as proofs of concept that connectionism is indeed capable of building and manipulating compositional cognitive representations (Levy and Gayler, 2004).

Topic and Goals

For this workshop we will focus on what may now be the major issue in connectionism and computational cognitive neuroscience: the debate between proponents of localist representations (e.g. Page, 2000), in which a single unit or population of units encodes one (and only one) item, and proponents of distributed representations, in which all units participate in the encoding of all items (see Plate, 2002, for an overview).

The aim of this workshop is to bring together researchers working with a wide range of compositional connectionist models, independent of application domain (e.g. language, logic, analogy, web search), with a focus on what commitments (if any) each model makes to localist or distributed representation. We have solicited submissions from both localist and distributed modelers, as well as those whose work bypasses this distinction or challenges its importance. We expect vigorous and exciting debate on this issue.

Specifically, the workshop will seek to address the following topics:

(1) What do we mean by "localist" / "distributed" in terms of the relationship between connectionist units and the items they represent?
(2) To what extent is the term "connectionist" still valid? Has the distributed / localist dimension supplanted the symbolic / connectionist dimension as the major axis of difference among cognitive models?
(3) How plausible and feasible is "holistic" computation, in which an entire structure is manipulated with sensitivity to its constituent parts without being decomposed into those parts? Does this feasibility depend on whether the representation is localist / distributed?
(4) Are temporal-synchrony-firing models necessarily localist?
(5) What constraints can neuroscience research bring to the distributed / localist debate? What can this debate contribute to the interpretation of neuroscientific research?
(6) Are some cognitive functions more plausibly seen as localist, and others more plausibly distributed?
(7) Do distributed (or localist) models scale more easily than localist (or distributed) models to realistically large problems?
(8) If two connectionist models, one distributed and the other localist, both account reasonably well for the same phenomenon, how can we judge between them?
(9) What mathematical principles (fractals, holography, chaos, etc.) can be borrowed from physics and other sciences to shed light on the nature of connectionist mental representations?

Importance of Workshop

This workshop addresses fundamental representational issues that cut across disciplines. Participants who have expressed an intention to submit a paper or attend come from a variety of backgrounds, including cognitive science, neuroscience, philosophy, psychology, engineering, and computer science. The success of our 2004 workshop and the continued interest of its participants make us feel that this workshop will contribute to the advancement of the field.

Plenary Speakers

Chris Eliasmith (Department of Philosophy, University of Waterloo): *How to build a brain: From single neurons to cognition*.

Chris Eliasmith's main interests lie at the intersection of theoretical/computational neuroscience and philosophy of mind. In his early work with Paul Thagard (Eliasmith and Thagard, 2001), Chris developed a psychological-level computer model called Drama which used distributed representations in modeling human analogical abilities. His more recent work in computational neuroscience includes large-scale simulation of specific animal behaviors (rodent path integration, zebrafish/lamprey locomotion, basal ganglia) using attractor networks and statistical inference. His computational cognitive science research includes working memory, context-sensitive linguistic inference, and emotion and decision making. On the philosophy side, Chris has worked on representational content/meaning, concepts, the philosophy of cognitive science, and the philosophy of neuroscience. In 2003, MIT Press published his book, coauthored with Charles Anderson (Eliasmith & Anderson, 2003), describing a general framework for modeling neurobiological systems in a realistic manner.

John Hummel (Department of Psychology, University of Illinois): *The proper treatment of symbols in a neural architecture*.

John Hummel and his students are interested primarily in the representation and processing of relations. How can neural architectures (brains or artificial neural networks) generate, represent and manipulate relational structures? How does the human mind's solution to this problem manifest itself in observable behavior? One line of research concerns the representation of relational structure in visual perception: How do we represent the relations among an object's parts or features? Under what circumstances, and in what form, will we explicitly encode these relations, and how does our encoding affect the manner in which we recognize and categorize objects? To address such questions, John and his students and collaborators have developed a symbolic neural network model – LISA (Learning and Inference with Schemas and Analogies; Hummel & Holyoak, 2003) – of analogical mapping, analogy- and rule-based inference, and schema induction. They recently generalized the LISA model to account for aspects of cognitive development, especially the development of relational concepts and relational representations (the DORA model; Doumas, Hummel & Sandhofer, 2008).

Program and Publication Plans

Abstracts have been solicited from participants of the 2004 symposium, researchers known to be interested in this area, and readers of relevant research mailing lists. We expect to select approximately 10 abstracts to be presented as papers at the workshop. The program will be available at http://www.cs.wlu.edu/~levy/cogsci2010/

We propose to follow up the workshop with a special issue in Cognitive Science or other relevant journal.

References

Doumas. L. A. A., Hummel, J. E., & Sandhofer, C. M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review,* 115, 1-43.

Eliasmith, C., & Anderson, C. H. (2003). *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems*. MIT Press.

Eliasmith, C., & Thagard, P. (2001). Integrating structure and meaning: A distributed model of analogical mapping. *Cognitive Science,* 25 (2), 245-286.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review,* 110, 220-264.

Levy, S. D., &. Gayler, R. (eds.) (2004). *Compositional Connectionism in Cognitive Science: Papers from the 2004 AAAI Symposium*. Technical Report FS-04-03, AAAI Press.

Page, M. (2000). Connectionist modelling in psychology: a localist manifesto. *Behavioral and Brain Sciences,* 23 (4), 443-467.

Plate. T. (2002). Distributed representations. *Encyclopedia of Cognitive Science*, Macmillan Reference Ltd.