

This article was downloaded by: [Levy, Simon D]

On: 27 May 2011

Access details: Access Details: [subscription number 938095617]

Publisher Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Connection Science

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713411269>

Compositional connectionism in cognitive science II: the localist/distributed dimension

Ross W. Gayler^a; Simon D. Levy^b

^a Philosophy Program, La Trobe University Bundoora, VIC, Australia ^b Computer Science Department, Washington and Lee University Lexington, USA, VA

Online publication date: 27 May 2011

To cite this Article Gayler, Ross W. and Levy, Simon D.(2011) 'Compositional connectionism in cognitive science II: the localist/distributed dimension', Connection Science, 23: 2, 85 – 89

To link to this Article: DOI: 10.1080/09540091.2011.587505

URL: <http://dx.doi.org/10.1080/09540091.2011.587505>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

EDITORIAL

Compositional connectionism in cognitive science II: the localist/distributed dimension

1. Background

In October 2004, approximately 30 connectionist and nonconnectionist 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).

2. Topics and goals of the second workshop

The August 2010 workshop focused on what may now be the major issue in connectionism and computational cognitive neuroscience: the debate between proponents of localist representations (e.g. Plate 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 was 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 solicited submissions from both localist and distributed modellers, as well as those whose work bypasses this distinction or challenges its importance. We expected vigorous and exciting debate on this topic, and we were not disappointed.

Specifically, our call for participation encouraged discussion on 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) 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?
- (3) What constraints can neuroscience research bring to the distributed/localist debate? What can this debate contribute to the interpretation of neuroscientific research?
- (4) Are some cognitive functions more plausibly seen as localist, and others more plausibly distributed?
- (5) Do distributed (or localist) models scale more easily than localist (or distributed) models to realistically large problems?

3. Participants

Our call for participation yielded eight submitted papers, including two plenary papers, and around 40 participants.

3.1. Plenary speakers

Our plenary speakers, Chris Eliasmith and John Hummel, are notable for the extent of their disagreement on fundamental point (1) above, viz., the meaning of the distributed/localist distinction.

3.1.1. *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 (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 behaviours (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, co-authored with Charles Anderson Eliasmith and Anderson (2003), describing a general framework for modelling neurobiological systems in a realistic manner.

The paper included in this special issue, co-authored with Terrence Stewart and Trevor Bekolay, re-examines the question of localist versus distributed neural representations using a biologically realistic framework, and suggests that claims of a set of neurons being localist or distributed cannot be made sense of without specifying the particular stimulus set used to examine the neurons.

3.1.2. *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 behaviour? 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 recognise and categorise objects? To address such questions, Hummel and his students and collaborators have developed a symbolic neural network model LISA (Learning and Inference with Schemas and Analogies; Hummel and Holyoak (2003) of analogical mapping, analogy- and rule-based inference, and schema induction. They recently generalised the LISA model to account for aspects of cognitive development, especially the development of relational concepts and relational representations, in the DORA model of Doumas et al. (2008).

In the paper included in this special issue, Hummel argues that role/filler binding is necessary to get symbolic processing from a neural computing architecture. Specifically, he argues that vector addition is both the best way to accomplish dynamic binding and an essential part of the proper treatment of symbols in a neural architecture, thereby casting doubt on the feasibility of completely holistic computation (topic (2) above).

3.2. *Other speakers*

Our other six speakers presented work addressing the topics mentioned in our call for participation.

Jeff Bowers (Department of Psychology, University of Bristol) examines the controversial 'grandmother cell' hypothesis (that a single neuron can represent an entire entity or concept). He argues that grandmother-cell theories are often mischaracterised and that when such mischaracterisations are avoided, grandmother cell theories constitute a viable account of how knowledge is represented in the brain.

Cynthia Henderson and James McClelland (Department of Psychology, Stanford University) present a connectionist model of the simultaneous perception of multiple objects, exploring how multiple objects can be perceived correctly in normal subjects given sufficient time but can give rise to illusory conjunctions under damage or time pressure. The model uses a classic hidden-layer PDP architecture that learns distributed representations of input/output relations; however, the hidden layers for each pathway are trained to produce either localist or distributed outputs. Characteristics of the outputs related to this distinction lead to important differences between the two sections of the network.

Trent Kriete and David Noelle (Cognitive and Information Sciences, University of California, Merced) studied the ability of a leading model of the working memory circuits of the prefrontal cortex (PFC) to learn componential representations from limited experience. The model was trained to represent, actively maintain, and apply a sequence of stimulus-response rule instructions, and the ability of the network to systematically generalise to novel sequences was assessed. Their model suggests that systematic generalisation arises from (1) distributed neural codes within PFC stripes and in more posterior brain areas, (2) slot-filler structures built from isolated pools of neural units that encode different representational components, supported by the isolated neural stripes in PFC anatomy, and (3) the temporally extended and sequential application of selective PFC stripe contents to the rest of the brain, controlled by a learned output gating mechanism.

Mark Reimers (Department of Biostatistics, Virginia Commonwealth University) addresses topic (3) directly, discussing the constraints that neuroscience research bring to the distributed/localist debate. His paper draws on the literature of neuroscience to suggest that a more subtle way of describing how different brain regions contribute to a behaviour, in terms of individual learning and in terms of degrees of importance, may render the current debate

moot: both sides of the ‘localist’ versus ‘distributed’ debate emphasise different aspects of biology.

Pat Simen (Princeton Neuroscience Institute) evaluates the pros and cons of a cognitive modelling approach based strictly on localist representations, by examining a sequential decision making architecture developed to model the neural basis of problem solving. This architecture seems not only amenable to distributed representations in its graded layers, but actively in need of them in order to support levels of model complexity on the scale of symbolic systems. The architecture requires a combinatorially explosive number of localist units as problem complexity increases. Simen discuss the possibility of preventing combinatorial explosion by binding low-level representations into high-level representations through temporal synchrony, using the same dynamics that underlie decision making in the architecture.

Frank van der Velde (Leiden University and Twente University) and Marc de Kamps (Leeds University) discuss the localist/distributed dimension (and related topics) from the perspective of a connectionist architecture that integrates grounding, productivity and dynamics. Concerning the localist/distributed dimension, the issue in terms of this overall architecture is not whether representations are local or distributed, but whether they are grounded. In particular, the issue is whether representations remain grounded when they are a part of a compositional representation.

4. Conclusions

Although we certainly cannot claim to have resolved the localist/distributed debate, we feel that the papers contained in this special issue have moved the debate forward in a useful way. For example, we no longer have reason to believe that these two kinds of representations differ in their basic ability to support compositional structures and systematic operations over their structures; instead, the difference may end up having more to do with the ways that such structures and operations scale up on different architectures. And the view from biology indicates that brain structures and processes do not map straightforwardly onto the localist/distributed distinction, suggesting that the distinction might be better explored at levels involving coalitions of several brain regions or processes. We look forward to the continued discussion and debate afforded by these discoveries.

Ross W. Gayler

Philosophy Program, La Trobe University Bundoora, VIC 3086 Australia

Simon D. Levy

Computer Science Department, Washington and Lee University Lexington, VA 24450 USA

Acknowledgements

The second author is grateful for the support of a Gerry Lenfest Summer Grant at Washington and Lee University, which supported the effort of organising and running this workshop.

References

- Doumas, L., Hummel, J., and Sandhofer, C. (2008), ‘A theory of the discovery and predication of relational concepts’, *Psychological Review*, 115, 1–43.
- Eliasmith, C., and Thagard, P. (2001), ‘Integrating structure and meaning: A distributed model of analogical mapping’, *Cognitive Science*, 25(2), 245–286.
- Eliasmith, C., and Anderson, C. (2003), *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems*, Cambridge, MA: MIT Press.

- Hummel, J.E., and Holyoak, K. (2003), 'A symbolic-connectionist theory of relational inference and generalization', *Psychological Review*, 110, 220–264.
- Levy, S.D., and Gayler, R. (2004), 'Compositional Connectionism in Cognitive Science: Papers from the 2004 AAAI Symposium', in *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', in *Encyclopedia of Cognitive Science*, Macmillan Reference Ltd.