Think Globally, Connect Locally

Localist Connectionist Approaches to Human Cognition
by Jonathan Grainger and Arthur M. Jacobs (Eds.)
Mahwah, NJ: Erlbaum, 1998. 357 pp. ISBN 0-8058-2555-8. $45.00

Review by Dominic W. Massaro and Michael J. Wenger

Early in the 20th century, Lashley (1929) offered the notions of equipotentiality and mass action to describe the general principles by which the nervous system codes for psychological experience (see also, Hilgard, 1951). The first of these constructs referred to the idea that learning and performance could not be localized to particular structural states or changes in the cerebral cortex. The second referred to the idea that a wide range of cortical areas that are pertinent to particular behaviors contribute to performance. To their time, these ideas were somewhat controversial, existing in the context of a concern for the localization of brain function, and offering the alternative view that psychological states were instead the result of patterned, dynamic, cooperative activities across a variety of cortical locations. In the 1980s, the "rediscovery" of distributed neural computation (cf., Anderson, 1995; Anderson & Rosenfeld, 1998) offered an alternative to flowchart models that localized perceptual and cognitive structures and processes to boxes and arrows.

Since that time, connectionist or neural modeling has made significant advances, in both theory and application (e.g., Bishop, 1995; Golden, 1996; Grossberg, 1978, 1991). However, the complexity and opacity of many of the resulting models (proving other problems) have caused some to wonder whether they have truly allowed for any significant increase in the level of understanding of the phenomena being modeled (referred to as the "Dennis paradox," Lewandowsky, 1993; Massaro, 1988; McCloskey, 1992; also being pertinent). Where the flowchart models allowed for apparently intuitive "big picture" views of the psychological problem, the neural models' wide range of mathematical and functional possibilities (i.e., their multiple levels of specification) seemed to obscure the system's basics, though it is possible that this was due as much to a lack of use of the tools available to understand neural models as a general level (e.g., Bishop, 1995; Golden, 1996) as to anything inherent to the models themselves.

As an alternative, rather than attempt to construct models that can address the entire psychological phenomenon, perhaps understanding can be achieved by starting with assumptions about structures and processes (such as those that were essential in the flowchart models), then, using the techniques of neural computation to precisely model that apply those assumptions in a formal and possibly more constrained way. These founding assumptions are the "localist" representations referred to in the title of this volume, with an influential precedent by Rumelhart and McClelland (1981) providing a grounding exemplar for much of the work that has come since. The contributions to this volume (which span the range of contemporary cognitive psychology) illustrate how this general approach can be applied.

The editors have done a commendable job of laying out a general approach to psychological theory building and testing. Importantly, they have also done much to show how connectionist modeling can be linked to information processing approaches in psychological science. They describe general principles of theory construction (referred to as system and design principles) that allow for some of the type of cross-model comparisons that have generally only been allowed for in psychological metatheories, such as signal detection theory (e.g., Egan, 1975; Green & Swets, 1966; Macmillan & Creelman, 1991; Swets, 1964), its multidimensional

DOMINIC W. MASSARO, Laboratory of Cognitive Psychology, University of Arizona, Aix-en-Provence, France

MICHAEL J. WENGER, Department of Psychology, University of California, Santa Cruz, California 95064. Electronic mail may be sent to massaro@fuzzy.ucsc.edu.

generalization (e.g., Ashby & Townsend, 1986; Kadlec & Hicks, 1998; Kadlec & Townsend, 1992; Thomas, 1996; Wickens & Olzak, 1989) into stochastic systems theory (e.g., Schweickert, Fisher, & Gold stein, 1992; Schweickert & Townsend, 1989; Schweickert & Wang, 1993; Townsend, 1972; Townsend & Nozawa, 1995; Townsend & Schweickert, 1989; formal models of utility and decision making (e.g., see the contributions in Marley, 1997) and refinements and extensions of functional measurement (e.g., Anderson, 1981; Massaro, 1998). The volume as a whole also suggests how the assumption of meaningful, interpretable structures, when given computational instantiation, may retain the type of biological plausibility that has long been argued as a strength of neural computation (e.g., O'Reilly, 1998).

The editors have done a commendable job of laying out a general approach to psychological theory building and testing.

Readers new to the notion of "localist connectionism" may be somewhat puzzled by the distinction between this variant of connectionist modeling and the more widely known distributed representations (e.g., McClelland & Rumelhart, 1986). All of the models described in this volume have the standard across treatments of connectionist models, particularly that basic unit of theoretical currency, the node. Within the distributed framework, the node has no straightforward a priori interpretation, accepting functional roles that might suggest an interpretation in terms of higher level constructs only after the network has been trained. In contrast, within the localist framework, nodes are given interpretable status at the outset. Higher level localist constructs exist at the outset as "hard-wired" components of the models. The editors, to their credit, acknowledge that the distinction between distributed versus localist—as at best imprecise. Still, as the editors argue, and as the various chapters illustrate, there is much that can be learned by relying on this form of formal theorizing to guide psychological research, and there are many crucial questions (predicted and unpredicted) that are brought to focus.

Broad Application, Warts and All

One of the strengths of this approach is its utility in addressing problems across the range of cognitive that characterize contemporary cognitive psychology and cognitive science. The initial application is to aspects of selective attention, and the remaining chapters include discussions of basic aspects of pattern recognition, word processing, and stimulus similarity (as pertinent to categorization and decision making). Each of these applications, to varying degrees, illustrates how the general approach to model construction, validation, and testing can be instantiated in particular regions of experimental psychology. For the discerning reader, each chapter provides examples of the advantages and potential pitfalls of the general approach.

Houghton and Tipper's chapter on selective attention provides an initial and intuitive illustration of how it is possible to derive strong, falsifiable predictions from a localist connectionist model. Exploration of their model leads to the strong prediction of an absence of negative priming (the finding of increases in response latency to targets previously presented as distractors at a specific point in time under specific experimental conditions. This, of course, is the ideal for any modeling approach. Theoretical explorations leading to distinct, empirically testable predictions. Yet, this initial application also illustrates how the localist approach may not avoid some of the problems that plagued the verbal "boxology" models. For example, the assumption of grouping of visual information into meaningful object representations as a point of departure for the modeling (47ff.) harkens to similar types of assumptions as starting points in "boxology" models, such as Deutsch and Young's (1966) logic model for facial processing.

In fact, this same model (Bruce & Young, 1986) provides the starting point for the model of facial recognition described in the chapter by Burton. For all of its acknowledged creativity in addressing the shortcomings of earlier approaches, Burton's model highlights the potential lack of constraints in positing localist representations. In particular, the redundancy of representation in this model (some thing that is not particular to Burton's approach, to be fair) does not appear to be assessed for its necessity. Instead, it appears to be justified on the basis of sufficiency for allowing the model to account for particular patterns. The lack of a more thorough analyses of this model would seem to underscore the strength of some of the predictions claimed by Burton (e.g., see pp. 84–85).

The chapter by Praamstra and Peters takes the reader from applications of the localist strategy in the processing of complex visual forms and faces to an application to the processing of complex auditory events (speech). This chapter raises the intriguing and complex issue of parameter estimation and interpretation. Earlier advocates of the localist approach, in modeling speech had great faith in specifying the parameter values for a model a priori, without reliance on new behavioral data to guide parameter estimation. The strategy adopted by Praamstra and Peters is to use a parameter estimation procedure guided by a measure of goodness-of-fit to empirical data. The issues raised here are deep (as discussed in the final chapter by Myung and Pitt, in Myung & Pitt, 1997), and this chapter reveals that a concern for these issues within the specific domain of word processing is related to the more general concern with issues of model estimation and validation.

The chapter that follows (authored by Jacobs, Rey, Ziegler, and Grainger) continues with the concern for language processing, but now at the level of visual word recognition. An important theoretical contribution of this chapter is to examine the nature of the "readout" of information at different levels of representation in such tasks, and the influence of this information in tasks such as those illustrating the word superiority effect. Previous models for this made it (in particular the McClelland and Rumelhart interactive activation model), under this reading, occur at the letter level. This assumption necessitated inferences that were shown to be invalid (Massaro, 1975), and the model presented in this chapter avoids this problem by adopting the strategy of allowing readout to occur at multiple levels of representation. In addition, the model is extended to make predictions at finer levels of analysis (from letter, e.g., see Spieler & Balota, 1997) than those addressed by earlier models.

Dijkstra and van Heuven develop a model for bilingual language processing that addresses, using the localist strategy, the perennial question of whether bilinguals possess two independent lexical stores that can be accessed, or whether there is a unitary, integrated lexicon in which representations from both languages are simultaneously activated. This chapter nicely illustrates how the localist strategy can allow for the detailed consideration of convergences with other areas of cognitive psychology (e.g., information processing approaches). For example, an important conclusion from model tests is that knowledge from both languages is simultaneously activated and used in processing words from one of the languages. The point of contact here is with the more general observation (Massaro, 1995) that language processing (and multiple other aspects of cognition) reveal important influences of multiple sources of information.

The connection to other approaches to modeling is also well illustrated in the work by Page and Norris on immediate serial recall. This chapter applies the localist strategy to the question of immediate serial recall. This is a question that has been addressed in some detail and with rather pronounced success in the information processing literature, specifically in work by Baddeley (1981; 1986). Page and Norris show how the localist approach can be used to provide connectionist instantiations of the same types of assumptions about memory that ground the information processing models. This type of work then shows how it might be possible to connect more traditional work in information processing psychology to computational models (see an extended discussion of such possibilities in O'Reilly, Wiegel, & Townsend, in press).

A similar link to information processing approaches can be found in the chapter by Schade and Eikmeyer, who discuss the manner in which object descriptions can be modeled, their assumption of sequentialization and the specification of "control nodes" important for the modeling of ordering in linguistic structures, harkens to the need for specification of both control structures and logical or information structures discussed thoroughly in Atkinson and Shiffrin (1968). In addition, this chapter provides another example of how system principles can be used to examine metatheoretical questions—questions that extend across models in a general way (e.g., 239ff.).

Nickerson's chapter on similarity and alignment extends the localist approach, and illustrates how questions of mixed self-organization (learning) can be accommodated. In one sense, then, this chapter illustrates the bridge to the more traditional distributed neural computation approach as well as to the information processing approach. Another strength of this chapter is its illustration of how strong predictions can be obtained using a localist modeling approach. In this case with the effect of object alignment contrasting in the general impact of object similarity (e.g., 298ff.).

Comparisons, Tests, and Standards for Theorizing

At the outset, the editors express a concern for principled model comparisons and evaluation. This concern is given concrete form with the inclusion of a chapter that can be considered somewhat unique in this context. The chapter is contributed by Myung and Pitt and reviews an approach to model evaluation first reported by Myung and Pitt (1997). One of the most consistent criticisms of the connectionist models has been the ability to create models that cannot be falsified. Myung and Pitt substantiate this, presenting a set of demonstrations of this possibility that mirror a similar demonstration reported by Massaro (1988). The authors propose that it might be possible to falsify localist models by assessing predictions of a model across the complete parameter space of the model. This chapter, like the paper it extends, may represent a foundational advance for this strategy of modeling, in that it presents a theoretical basis for establishing methods of model evaluation and comparison.

Thus, this final chapter reinforces the rather bold step taken by the editors. Not only have they illustrated a general approach to modelling cognitive phenomena, but they have suggested principles for the very construction and evaluation. In that sense, the volume can be seen as a collective illustration of the potential and power that can be obtained with a general metatheoretical approach. Certainly it is possible to show how the editors' goals and standards are not met in many of the chapters. The same could be said for applications of any of the other metatheoretical approaches in psychology. As a cumulative record, our science advances on the shoulders of both our successes and shortcomings. The type of general theoretical approach espoused by the authors and editors allows for discussion of both of these. This is the fundamental strength and contribution of this volume. □

References

Anderson, J. A. (1998). Learning arithmetic with a neural network: Seven times seven is about fifty. In D. Scarborough & S. Sternberg (Eds.), *An invitation to cognitive science: Methods, models and conceptual issues* (Vol. 4, pp. 201–250). Cambridge, MA: MIT Press.

Anderson, J. A., & Rosenfeld, E. (1998). *Talking nets: An oral history of neural networks*. Cambridge, MA: MIT Press.

Anderson, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press.

Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological Review, 93*, 154–179.

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2, pp. 89–195). New York: Academic Press.

Bishop, C. M. (1995). *Neural networks for pattern recognition*. Oxford, England: Clarendon.

Boring, E. G. (1950). *History of experimental psychology*. New York: Appleton-Century-Crofts.

Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology, 77*, 305–327.

Egan, J. P. (1975). *Signal detection theory and ROC analysis*. New York: Academic Press.

Golden, R. M. (1996). *Mathematical methods for neural network analysis and design*. Cambridge, MA: Bradford.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

Grossberg, S. (1976). Competition, decision, and consensus. *Journal of Mathematical Analysis and Applications, 66*, 470–493.

Grossberg, S. (1991). Nonlinear neural networks: Principles, mechanisms, and architectures. In G. A. Carpenter & S. Grossberg (Eds.), *Pattern recognition by self-organizing neural networks* (pp. 35–109). Cambridge, MA: MIT Press.

Kadlec, H., & Hicks, C. L. (1998). Invariance of perceptual spaces and perceptual separability of stimulus dimensions. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 80–101.

Kadlec, H., & Townsend, J. T. (1992). Implications of marginal and conditional detection parameters for the separabilities and dependencies of perceptual dimensions. *Journal of Mathematical Psychology, 36*, 325–374.

Lashley, K. S. (1929). *Brain mechanisms*

and intelligence. Chicago: University of Chicago Press.

Lewandowsky, S. (1993). The rewards and hazards of computer simulation. *Psychological Science, 4*, 236–243.

Macmillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide*. Cambridge, England: Cambridge University Press.

Marley, A. A. J. (Ed.). (1997). *Choice, decision, and measurement: Essays in honor of R. Duncan Luce*. Mahwah, NJ: Erlbaum.

Massaro, D. W. (1975). Letter information and orthographic context in word perception. *Journal of Experimental Psychology: Human Perception and Performance, 5*, 595–609.

Massaro, D. W. (1988). Some criticisms of connectionist models of human performance. *Journal of Memory and Language, 27*, 213–234.

Massaro, D. W. (1996). *Perceiving talking faces: From speech perception to a behavioral principle*. Cambridge, MA: Bradford.

McClelland, J. L., & Rumelhart, D. E. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: Bradford.

McCloskey, M. (1991). Networks and theories: The place of connectionism in cognitive science. *Psychological Science, 2*, 387–395.

Myung, I. J., & Pitt, M. A. (1997). Applying Occam's razor in modeling cognition: A Bayesian approach. *Psychonomic Bulletin & Review, 4*, 79–95.

Naime, J. S. (1991). Positional uncertainty in long-term memory. *Memory & Cognition, 19*, 332–340.

Naime, J. S. (1992). The loss of positional certainty in long-term memory. *Psychological Science, 3*, 199–203.

O'Reilly, R. C. (1998). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation, 8*, 895–938.

O'Toole, A. J., Wenger, M. J., & Townsend, J. T. (in press). Quantitative models of perceiving and remembering faces: Precedents and possibilities. In M. J. Wenger & J. T. Townsend (Eds.), *Computational, geometric, and process perspectives on facial cognition: Contexts and challenges*. Mahwah, NJ: Erlbaum.

Rumelhart, D. E., & McClelland, J. L. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review, 88*, 375–407.

Schweickert, R., Fisher, D. L., & Goldstein, W. M. (1992). *General latent network theory: Structural and quantitative analysis of networks of cognitive processes* (Tech. Rep. No. 92-1). Lafayette, IN: Purdue University Mathematical Psychology Program.

Schweickert, R., & Townsend, J. T. (1989). A trichotomy method: Interactions of factors prolonging sequential and concurrent mental processes in stochastic PERT networks. *Journal of Mathematical Psychology, 33*, 328–347.

Schweickert, R., & Wang, Z. (1993). Effects of response bias on factors selectively influencing processes in additive task networks with OR gates. *British Journal of Mathematical and Statistical Psychology, 46*, 1–30.

Spieler, D. H., & Balota, D. A. (1997). Bringing computational models of word naming down to the item level. *Psychological Science, 8*, 411–416.

Swets, J. A. (1964). *Signal detection and recognition by human observers*. New York: Wiley.

Thomas, R. D. (1996). Gaussian general recognition theory and perceptual independence. *Psychological Review, 103*, 192–200.

Townsend, J. T. (1972). Some results concerning the identifiability of parallel and serial processes. *British Journal of Mathematical and Statistical Psychology, 25*, 168–199.

Townsend, J. T., & Nozawa, G. (1995). On the spatio-temporal properties of elementary perception: An investigation of parallel, serial, and coactive theories. *Journal of Mathematical Psychology, 39*, 321–359.

Townsend, J. T., & Schweickert, R. (1989). Toward the trichotomy method: Laying the foundation of stochastic mental networks. *Journal of Mathematical Psychology, 33*, 309–327.

Wickens, T. D., & Olzak, L. A. (1992). Three views of association in concurrent detection. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 29–54). Hillsdale, NJ: Erlbaum.