Nirwan Ansari is at the Center for Communications and Signal Processing, Electrical and Computer Engineering Department, New Jersey Institute of Technology, University Heights, Newark, NJ 07102.

Reference

Widrow, B., & Stearns, S. D. (1985). *Adaptive Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall.

Understanding Hidden Connections

Connections between levels of description of perception. James L. McClelland. *Proceedings of the International Joint Conference on Neural Networks, 1,* 743-747. Hillsdale, NJ: Lawrence Erlbaum Associates. Washington, DC, January 15-19, 1990.

Reviewed by Dominic W. Massaro

A tension exists in psychological theory around the issue of level of description. Information processing theory traditionally described performance in terms of several relatively molar algorithms and heuristics that take the processor from input to output. An alternative espoused by McClelland, many scientists, and this journal is to describe behavior as emerging from the processing in a large network of neuron-like computing elements. Resolution of this tension will take some time, but we can expect it will depend on which level of description fits best with the behavior of interest.

Psychology is the science of mind and behavior. Neurophysiology is the science of brain and behavior. Because mind depends on brain, we might expect to be able to reduce the former to the latter. Although mind might emerge from brain, an important question is whether behavior is better understood as caused by mind—hidden psychological operations of the information processing variety—or brain. My hunch is that the neural level is, in many cases, simply too far removed from the behavior of interest. Some support for this hunch comes from developments in physics involving chaos and nonlinear dynamics, in which a more global level of description has advantages over a more reductionistic level. Although networks might be ideally suited for describing neural events, neural theory does not easily generate valuable predictions for behavioral events. By valuable predictions, I mean those that help explain the phenomenon, those that are informative, and those that are falsifiable. Explanation occurs when something unknown is described in terms of something known. Informative predictions are those that make a difference. And finally, falsifiable predictions are those that can be rejected given some set of observations of behavior. (Note that one needs to distinguish between the predictions of "network" theories that implement psychological principles and predictions which are solely a function of the network itself. Thus, the original interactive activation (IAC) model (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1981) implemented the psychological principle of reading a letter string by combining the activations of several words in memory (Glushko, 1979). Many of the predictions of the IAC model tend to emerge because of this psychological principle, rather than because of the actual implementation of the principle in the network.)

Psychologists have discovered empirically that stimulus-response or input-output functions are too coarse to develop informative theory. In psychophysics, for example, both perceptual and decision processes govern observable behavior, and each must be accounted for in both experimentation and theory. In human performance, several relatively sequential stages of processing intervene between some stimulus, such as a spoken word, and behavior, such as pronouncing the spoken word. The operating thesis of connectionists is that the neural level of description can produce a better characterization of the hidden processes between the observable input and observable output than the characterization given by information processing. In IAC, two-way connections between units are assumed. If the units are taken to represent sources of information, then IAC models do not maintain a single feedforward flow through the processes of evaluation, integration, and decision. The reason is that integration feeds back to evaluation and changes the outcome of evaluation. The net result of interactive activation combined with a nonlinear activation function is that discriminability of different levels of one source of information varies depending on the nature of other sources of information. The discriminability between two levels is enhanced when another source of information points in the opposite direction, in comparison to the case in which the two sources of information point in the same direction. This prediction has been falsified in several experiments (Massaro, 1979, 1989). Faced with this falsification, McClelland pursued the obvious strategy of trying to repair the model. His goal, however, became one of predicting a stimulus-response function—additivity of stimulus and context—without questioning the nature of the intervening processes underlying this function.

McClelland's strategy might have been reasonable if our empirical data base were limited to this simple result of additivity in a categorization task. Evidence from other tasks, however, informs us about the hidden processes, and these processes appear to be inconsistent

with the processes assumed in stochastic interactive activation networks. Additivity is not limited to stimulus input and context, but also applies to multiple sources of stimulus input. Two sources of information—such as auditory and visual perception of speech—produce the same type of independence observed between input stimulus and context. In addition, this classical result is also found in rating tasks in which subjects rate the degree to which the target is one alternative or another. If subjects are assumed to use the same decision rule as in categorization tasks, then the stochastic interactive activation model (SIAC) is necessarily a categorical model and, therefore, is contradicted by the distribution of these rating judgments (Massaro, 1987, pp. 142-148). If subjects are assumed to use the quantitative relationship between the activations of the two alternatives in making their rating judgments, then SIAC cannot predict the classical result. This follows from the fact that the activations themselves do not give the classical result, and it is only the assumed variability and the best-one-wins (BOW) decision rule that allow the classical result to be predicted.

McClelland (in press) compromised several attractive features of interactive activation in his revised models. Interactive activation was valued partly because it specified the time course of information processing. The derivation and prediction of the classical result, however, was carried out for the networks after they had reached equilibrium. In the original model, activation functions—the currency of interactive activation—were considered to be psychologically real and meaningful. It was these activations that predicted behavior. The variability and decision rule assumed by stochastic interactive activation models, however, now makes the form of these activation functions irrelevant (Massaro & Cohen, 1990). The only important property of the function is where the activation function for one response alternative crosses over the function for the other response alternative.

Theoretical and empirical analyses of IAC models and the fuzzy logical model of perception (FLMP) have yielded useful information relevant to several important issues in the description of how sensory information and context jointly determine perceptual recognition. The empirical results support independent contributions of these two sources of information. Interactive activation necessarily introduces nonindependence, because of two-way excitatory connections among units in different layers and because of inhibitory connections among units within a layer. This nonindependence at the activation level can be neutralized at the decision stage by adding noise to the input or during processing along with a BOW decision rule. McClelland constructed an IAC network to generate hypothetical results that are consistent with independence. This same network, however, provided a significantly poorer description of actual empirical results than that provided by the FLMP (Massaro & Cohen, 1990). Thus, SIAC models have not been shown to give an adequate account of asymptotic performance.

Notwithstanding their poorer fit to asymptotic performance, there has been a good reason to value interactive-activation accounts. Interactive activation is attractive because it specifies the dynamics of information processing that lead to asymptotic performance. Surprisingly, however, there have been few direct tests of dynamical predictions of the model. In most cases, predictions are generated by letting the model run to asymptote (or to equilibrium in the case of Boltzmann machines). In fact, McClelland's implementation of SIAC models is based on the assumption that processing occurs for a *fixed* number of cycles and that a decision is made based on the activations at this time. Thus, it is essential to test the IAC models against performance measures of the dynamics of perceptual processing.

Backward masking was used to provide performance measures at different points during perceptual processing, making it possible to observe how stimulus information and context interact over time. The empirical results revealed that context can have a substantial influence even at very short processing intervals. In addition, the relative contribution of stimulus information increases over processing time. Neither of these empirical results are easily predicted by IAC models. Several cycles of processing are necessary before top-down activation can influence the representation of the target letter. In order to predict context effects at short processing times, the SIAC model requires a very short time (2 msec) for each processing cycle. Given this short time per cycle, however, the activations become asymptotic too quickly, and changes in performance with additional processing time cannot be predicted.

Given interactive activation, context can eventually overwhelm sensory information about the target. This prediction is contradicted by both phenomenological experience and experimental results. In everyday encounters, additional processing enhances our perception of sensory information in context rather than degrading it: We are more likely to notice a misspelling in a word when we look at it carefully than when we read it rapidly. In experiments varying both target information and context, contextual effects are larger to the extent that the sensory information specifying the target is ambiguous (Massaro, 1979). The SIAC model, on the other hand, cannot easily predict a decrease in context effects with an increase of processing time.

Dominic W. Massaro is at the Program in Experimental Psychology, University of California, Santa Cruz, Santa Cruz, CA 95064. *The writing of this paper was supported, in part, by grants from the Public Health Service (PHS R01 NS 20314), the National Science Foundation (BNS*

8812728), and the graduate division of the University of California, Santa Cruz. Massaro thanks Michael M. Cohen for valuable comments.

References

Glushko, R. J. (1979). The organization and activation of orthographic knowledge in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance, 5,* 674-691.

Massaro, D. W. (1979). Letter information and orthographic context in word perception. *Journal of Experimental Psychology: Human Perception and Performance, 5,* 595-609.

Massaro, D. W. (1987). *Speech perception by ear and eye: A paradigm for psychological inquiry.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Massaro, D. W. (1989). Testing between the TRACE model and the fuzzy logical model of speech perception. *Cognitive Psychology, 21,* 398-421.

Massaro, D. W., & Cohen, M. M. (1990). Interactive activation and the interaction of sensory information and context in perception. Submitted for publication.

McClelland, J. L. (in press). Stochastic interactive activation and the effects of context on perception. *Cognitive Psychology.*

McClelland, J. L., & Rumelhart, D. E., (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review, 88,* 375-407.

Rumelhart, D. E., & McClelland, J. L., (1981). An interactive activation model of context effects in letter perception: Part 2. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review, 89,* 60-94.

The Author Responds

Massaro's comments misrepresent my paper, "Connections between levels of description of perception." The aims of my paper are not well described, and there are several distortions of fact.

In "Connections," I took it for granted that there are two useful levels of description of perception: one in terms of global, whole-system performance, and the other in terms of the underlying processing activity, which I assume is best characterized in terms of parallel-distributed processing (PDP). Indeed, the point of the paper was to show that a particular class of PDP models can provide mechanisms that implement frequently replicated patterns of human performance in perceptual identification tasks. Massaro seems to indicate that I feel that macro-level characterizations should be discarded in favor of micro-level, PDP accounts. I do not. Rather, I feel that PDP may eventually help us understand how macro-level characterizations of performance may arise from the underlying microstructure.

Massaro takes exception to my claim that interactive activation models actually can account for whole-system performance in useful ways. For the sake of brevity I will not enumerate the many ways in which I find Massaro to have distorted this class of models, the facts, or both. I will reply to just one claim. Massaro says: "McClelland constructed an IAC network to generate hypothetical results that are consistent with independence. This same network, however, provided a significantly poorer description of actual empirical results than that provided by the FLMP." The poorer description was one obtained by Massaro and Cohen (1990). It turns out that this poorer description is not a general fault of IAC but of Massaro and Cohen's particular choice of how to parameterize effects of different contexts on performance.

To me it seems like a very good bet that we will someday understand how the macrostructure of perception emerges from a PDP-like microstructure. Certainly, there is more work to do, but at this point in the development of this line of research things look very promising for a PDP account.

James L. McClelland is at the Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213-3890.

Reference

Massaro, D. W., & Cohen, M. M. (1990). Interactive activation and the interaction of sensory information and context in perception. Submitted for publication.