In T. Inui, & J. L. McClelland (Eds.) (1996).
Attention and Performance XVI: Information Integration in Perception and Communication. Cambridge, MA: MIT Press.

24 Integration of Information: Reflections on the Theme of Attention and Performance XVI

James L. McClelland

24.1 ANALYSIS AND SYNTHESIS

In the early 1980s the idea that key aspects of cognition arise from the independent activity of autonomous modules was popular in several areas of the cognitive and neural sciences. This divide-and-conquer approach was represented by Chomsky's (1980) position on the autonomy of language in general and syntax in particular, by Marr's (1982) emphasis on independent computation of surface properties from each of several distinct visual cues, and by Fodor's (1983) advocacy of modularity as a general principle of brain organization.

These ideas were advances over earlier approaches that attempted to encompass all of cognition or behavior in terms of a few very general principles. Paradigms like Gestalt psychology, genetic epistemology, and behaviorism were all radically different, but they each offered a broad framework intended to encompass a wide range of diverse phenomena. All of these approaches ultimately proved unsatisfying, and more modular approaches may have arisen in part because of the perception that one of the problems with these approaches lay in the very attempt to generalize so broadly. Chomsky certainly taught us a great deal more about language than Skinner, and Marr certainly took us beyond the vague holism of the Gestaltists toward a much more explicit computational understanding of the extraction of shape from visual cues. Focus on a specific domain and attention to its details, as if it were autonomous, has led to insights that have sharpened our understanding of language and perception. But now that these advances have been achieved, it may be time to go beyond modularity and consider integration again.

The sciences of mind and brain have seen an alternation between global approaches and more modular approaches before. In the nineteenth century, and into the twentieth, ideas about the organization of function in the brain alternated between varieties of equipotentialism and much more localist treatments. In my own scientific education, the history of this alternation was used to make two distinct and equally important points. Teitelbaum (1967), an important physiological psychologist of the 1960s, emphasized the importance of analysis *and* synthesis. First, we must carve the bird at its joints, he argued, and analyze the structure and functional contributions of each of its parts. But if we are to understand how it can fly we must reconnect these parts, and see how they work together. And Luria (1966, 1973), who is of course still widely known for his many contributions to cognitive neurology, stressed the idea of dynamic functional systems. He used the findings of the localizationists to show that each part has its own special role. But he noted the poverty of considering these parts in isolation and insisted that they must be seen as working in concert to achieve system-level functions such as perception, communication, and action.

In view of these points, it is encouraging to see just how much contemporary research builds on the contributions of the modularists in an effort to understand how the parts of the cognitive system work together. The present volume brings together a collection of such researchers, who consider the problem of integration from multiple perspectives and in multiple contexts. What they share—nearly all of them—is a commitment to the effort to understand how complex, intelligent functions are synthesized from simpler parts. Their efforts build on the analytic insights that arise from the modular approaches mentioned above, as well as on continuing work that combines analysis with an effort to understand the behavior of the parts in context.

24.2 INTEGRATION IN PERCEPTION, COMMUNICATION, ATTENTION, AND ACTION

When Toshio Inui proposed the theme of Attention and Performance XVI, he certainly touched a responsive chord in me, and in many other members of the executive committee. The theme of information integration turned out to provide a touchstone for a number of excellent contributions that fit together in many different ways. The goal of these remarks is to bring out some of the threads that link the chapters—and, in a few places, to touch on relevant work not presented in the symposium. While the exact shape of a full account of information integration in perception, communication, and action is not yet clearly in view, I hope these comments will help the reader weave some of the threads together and envision something of the fabric of this ultimate synthesis. At the least, I hope they will provide pointers to some of the main findings and ideas presented by the participants in this meeting.

The five broad questions Inui raised in his introduction form an excellent framework for my discussion, since they cut across the various specific topics covered in the conference and provide a structure in which the crucial points of nearly all the papers fit very neatly. However, I will order the questions differently for expository convenience, and I will add one more question at the end of the list:

• What are the limitations on the integration of information? Are there impenetrable modules in the cognitive system and, if so, what are their boundaries?

• Are there supramodal representations for integration across modalities, or is integration merely mutual constraint between modality or dimension specific representations?

• Are there dimensions that provide a strong organizing frame for integration? Are time and space examples of such dimensions?

• What role does attention play in the integration of information?

Limits on Integration

What are the limitations on the integration of information? Are there impenetrable modules in the cognitive system and, if so, what are their boundaries?

In the literature spawned by the emphasis on modularity in the late 1970s and 1980s, many findings were taken as evidence for the impenetrability of modules of various kinds. A great deal of this work focused on the study of language comprehension as it unfolds on-line during reading and listening. Leading this wave, Swinney (1979) explored the effect of context on the resolution of lexical ambiguity:

Rumor had it that, for years, the government building had been plagued with problems. The man was not surprised when he found several spiders, roaches, and other bugs in the corner of the room.

He probed for activations of appropriate and inappropriate meanings of the ambiguous word ("bugs") using a lexical decision probe related to the appropriate meaning (e.g., "ant") the inappropriate meaning (e.g., "spy") or unrelated to either meaning (e.g., "table"). When the probe occurred immediately after the ambiguous word, lexical decision times were faster to both spy and ant compared to the unrelated condition, and the difference between the priming effects in the two cases was not significant. This null finding was then taken to support the claim that immediate lexical access during language comprehension was unaffected by context. Lexical access then became one of the impenetrable submodules of the language processing system.

Swinney's (1979) experiment exemplifies a common research approach taken during this period: (1) manipulate a possible source of constraining influence on some process; (2) when a null effect is found, declare the process impenetrable to the source of constraining influence. Similar studies were carried out by other investigators, examining possible effects of varibables such as animacy and semantic plausibility on the assignment of syntactic constituent structure during sentence processing (Clifton and Ferreira 1989). As in the Swinney experiment, when null effects were found, the process was declared impenetrable to the manipulated variable. The net effect of these studies and others with similar logic was to suggest that lexical access, syntactic parsing, and several other processes were essentially autonomous modules impervious to contextual influences.

Yet even at the peak of this modularist wave, them were some indications that the processes in question might not be so impenetrable after all. In particular, Simpson (1981) found that context did influence initial activation of meanings of ambiguous words, if other factors (particularly, the relative frequency of the two alternative readings of the ambiguous word) were relatively balanced. Similarly, Taraban and McClelland (1988) found that plausibility affected initial syntactic parsing decisions in cases where other factors (syntactic cues favoring one or the other alternative parse) were relatively balanced. Reviewing this evidence, I suggested in a paper presented at Attention and Performance XII (McClelland 1987) that both lexical ambiguity resolution and syntactic ambiguity resolution might in fact be sensitive to contextual as well as intrinsic factors.

Several other groups of investigators have explored the possibility that lexical and syntactic aspects of language processing might exploit contextual as well as intrinsic factors. This tradition goes back at least to Miller (1962), and can he traced through the work of Bever (1970), Rumelhart (1977) and Marslen-Wilson (1987). Recent proponents include many of the participants in this symposium. MacDonald (chap. 17) provides a definitive statement of the current state of development of this approach. She characterizes lexical access and syntactic structure assignment as constraint satisfaction processes that exploit multiple sources of graded information. On this view, influences from various sources may differ in strength, but no one source is necessarily decisive—all sources contribute to the accumulation of evidence for each of the available alternatives (see below). She indicates how syntactic structure assignment may hinge on lexical ambiguity resolution and explains why contextual effects only emerge when other sources of constraint are in relative balance (see also Massaro, chap. 16).

In summary, a considerable body of recent literature on lexical and syntactic ambiguity resolution is consistent with the view that two of the processes prominently discussed as examples of autonomous processes impervious to contextual influences may not in fact be autonomous or impenetrable after all. While there still may be room for skepticism about the exact nature and timing of the influences of context on these processes, the current evidence certainly appears consistent with a constraint satisfaction view.

Of course lexical access and syntactic structure assignment are only two among many candidate processes, and there have certainly been claims of impenetrability in other domains besides these. Since the contributors to the present volume have nearly all been concerned with exploring the integration of information, it will probably come as no surprise that the emphasis is on instances in which a process does appear to be susceptible to multiple sources of information. Several of these cases may be surprising to those who adopt modular views of the architecture of cognition. Here are several of the most striking instances of information integration reported in the chapters of this volume:

1. Acoustic and visual cues are combined to determine the outcome of the process of phoneme identification (Massaro, chap. 16).

2. Visual, auditory, and tactile cues all contribute to the specification of the target of attention and the contents of the spotlight of attention (Driver and Grossenbacher, chap. 9; Hikosaka et al., chap. 10; Shimojo et al., chap. 23).

3. Proprioceptive information from multiple relevant joints ranging from the ankle to the neck is integrated with visual information in the specification of the location of objects in external space (Roll et al., chap. 12) and in the specification of the location of objects with respect to the body (Colby, chap. 7; Graziano and Gross, chap. 8).

4. Visual and spoken language inputs are integrated in determining the intended referent in a task involving instructed action in context (Tanenhaus et al., chap. 18).

In the face of the first three points a modularist might suggest that there are modules but these are defined not so much in terms of traditional sense modalities like proprioception, audition, vision, and so on, as in terms of broad functional systems, dedicated to language comprehension or action coordinated with objects in external space. But the work of Tanenhaus et al. (chap. 18) makes it clear that there is integration between the spatial and the linguistic modalities: language directs attention in space in real time, and spoken input and space work together to constrain linguistic interpretation. Altmann (chap. 19) extends this point further, drawing of the work of Sopena (1991) and his own recent simulations of artificial language learning to argue that a constraint called the "late closure constraint," which has often been taken to be a purely syntactic constraint, might well arise because language use and language acquisition occur in situations where interpretation of language and interpretation of the environment are mutually constraining. In particular, Altmann suggests that late closure may arise from an effect of language on directing attention to objects, together with a tendency to remain focused on one object until directed elsewhere. The general point is that language guides interpretation of the environment as much as the environment guides interpretation of language.

Along with all these indications that the functional systems underlying verbal and spatial cognition integrate many different sources of information, some of the papers in this volume do suggest some limits on the integration of information. Treisman and DeSchepper (chap. 2) show that there is enough perceptual integration even of unattended objects to produce long-lasting, object-specific aftereffects, but there is nevertheless a limit to the sophistication of the object-specific representations that are formed outside of attention; these representations seem to be simple representations of an enclosing contour rather than full structural descriptions. Hummel and Stankiewicz (chap. 5) suggest that the formation of full structural descriptions, providing an integrated representation of the object that could be recognizable across a transformation like mirror reversal, requires focal attention to the object for an extended period of time.

Another example of limits on integration is provided by studies reviewed by MacDonald (chap. 17), showing that exploitation of context may vary dramatically with individual differences in comprehension skill. While skilled readers often show robust context effects in both lexical and syntactic ambiguity resolution, poor readers often show little or no effect of context (Just and Carpenter 1992). The exact interpretation of the source of these individual differences is not entirely clear; Just and Carpenter suggest that maintaining context and processing current input impose competing demands, and that better comprehenders are better able to cope with these competing demands, and so are better able to maintain a representation of the context. Indeed, one possible reason for the strong contextual influences reported by Tanenhaus et al. (chap. 18) is that the context is available visually in their experiments, making reliance on maintaining it in memory unnecessary.

A final example of limits on integration is provided by Marks and Armstrong (chap. 11). They consider the visual and haptic perception of length and report two findings relevant to information integration. First, they find that although both visual and haptic perception of length exhibit distortions, these distortions are not strictly homologous; and second, they find that adaptation within one modality, which modulates these distortions, does not transfer to the other modality. These findings appear to suggest that visual and haptic cues are carried by separate input systems, each subject to independent adaptation, even if the results are ultimately combined to constrain representations of the shape and layout of objects in space.

Principles of Integration

Are there common principles of integration that span domains in which integration is necessary?

To address this question, it is useful to have some sort of common framework for thinking about different examples of information integration. One way to frame the issue of integration that encompasses integration in perception, comprehension, attention, and planning of action is to think of each of these processes as leading to the construction or selection of something we will call a "specification." We can view perception as the construction of a specification of a perceptual interpretation of a distal stimulus; we can view comprehension as the construction of a specification of a conceptual interpretation of a linguistic expression; we can view attention as the construction of a specification of an object of attention; and we can view planning as the construction of a specification of one action (or action sequence) out of the many possible actions that might be taken.

In many models specification is treated as a matter of selection of one of N enumerable alternatives. I will begin by considering this simpler case and then comment briefly on a generalization of the concept. For the case of integration for selection, there appears to be broad consensus among the participants in Attention and Performance XVI on three basic points:

1. Integration for selection depends on accumulation of support for each alternative from a range of different sources, including prior knowledge, current expectations, and one or more sources of input information.

2. The support contributed to each alternative by each source is a matter of degree, and the total support for each alternative that results from the combination of all of the sources is itself a continuous function of the support provided by all of the sources.

3. Selection occurs through a competition among the alternatives, where the outcome depends upon the relative amount of total support for each alternative, compared to all of the other alternatives.

These points are explicit (to varying degrees) in Massaro's fuzzy logical model of perception (chap. 16), in Bülthoff and Yuille's concept of competitive priors (chap. 3), in Rosenbaum et al.'s model of posture selection (chap. 13), in MacDonald's constraint satisfaction model of syntactic ambiguity resolution (chap. 17), and in Duncan's integrated competition model of selective attention (chap. 21), and they are implicit in several other places. Indeed, versions of these ideas can be found in many earlier works where integration of different sorts of information is considered, including Treisman (1964), Morton (1969), Rumelhart (1977), and even Marcus (1980). Of these, Marcus's 1980 book is especially interesting historically, in that it is generally regarded as a classic example of an effort to develop a highly autonomous syntactic parser in the context of the Chomskian view of syntax. Yet Marcus argues forcefully for the need to consider both semantic and syntactic constraints in parsing decisions and provides a very convincing argument in favor of a set of principles that are basically equivalent to the three points just listed. He notes that both syntactic constraints and semantic constraints must be matters of degree because neither is dominant in every case; and he suggests that selection in cases of conflict depends on the total amount of support for each alternative, whether it comes from syntax, semantics, or a combination of the two.

Of course, points 1–3 provide a very general framework that leaves room for important differences in details. Massaro (chap. 16) emphasizes that the contributions of various sources of support are generally treated as independent, while Bülthoff and Yuille (chap. 3) explicitly question this assumption. This contrast will be considered in more detail below.

It should also be noted that the exact formulation of the process of selection varies as well. The candidate selection schemes are (1) choose the best alternative (i.e., the one with the greatest total support) or (2) choose an alternative probabilistically, setting the probability of choosing a particular alternative proportional to its total support. These different proposals are discussed by Bülthoff and Yuille (chap. 3) and Massaro (chap. 16), respectively. These ideas are not very easy to distinguish empirically. The second scheme is naturally probabilistic, while the first gives rise to probabilistic performance if there is any noise in the inputs or the process that accumulates the total support. Rosenbaum et al. (chap. 13) introduce another possibility: construct a blended alternative by taking a weighted average of the alternatives, where the weights are proportional to the total support for each. This proposal goes beyond the mere selection of a single alternative, and provides a useful way of allowing some degree of generalization to novel cases. It is particularly useful when the alternatives are particular values of continuous parameters (such as the muscle length targets as in Rosenbaum et al. (chap. 13)), and the candidates with appreciable amounts of support are sufficiently close together in the space encompassing all possible alternatives.

As already suggested, the idea of selection of an existing alternative, or even of averaging similar alternatives, is unlikely to prove fully adequate to capture the creativity and flexibility of human perception, comprehension, and planning processes. What is needed is a more general framework. Classically (Fodor and Pylyshyn 1988), cognitive scientists have supposed that what is specified is a structural description—a structure consisting of a hierarchical arrangement of items assigned to roles. Each item in the structural description is either an embedded structural description or an atom. While this framework is surely vastly more powerful than the selection of one alternative out of *N*, it has not proven to be very tractable for capturing the integration of graded constraints. An alternative approach is to think of the specification as a pattern of activation over an ensemble of processing units. The settling of a network into an attractor state, consisting of a pattern of selected (active) elements, can then be seen as an implementation of such a selection process. This settling process can be regarded as a process of hill climbing in some ensemble measure of the goodness of the entire state (Hopfield 1982). While each element may participate in many patterns, the overall pattern can represent a novel configuration; the selection of the particular pattern that ends up being specified clearly depends on a graded synthesis and competition process that accords with the three points enumerated above (Rumelhart et al. 1986). Several methods for capturing complex hierarchical structure in such patterns of activation have been proposed (see the papers in Hinton 1991). A third approach is to view the object of specification as a sequence of states. While the item selected at a given point in a sequence may not be novel at all, the concatenation of several individually selected items can give rise to novel sequences. Given that sequences may be unbounded in length, productivity is thereby assured. Although many traditional sequential models do not provide for the graded synthesis of multiple constraints, there are now many sequential architectures that do (Jordan 1986; St. John and McClelland 1990; Jordan and Rumelhart 1992). In this volume, Kawato's model of the specification of a sequence of motor commands (chap. 14) is a good example of a specification process in which the selection of the motor command at each time point depends on a graded synthesis of multiple constraints, including the motor commands programmed for adjacent time points. It seems likely that the future study of information integration will come to focus more and more sharply on specifications that involve patterns and/or sequences rather than simple selections of one out of *N* alternatives.

Constraint Propagation for Integration

How do constraints propagate efficiently for integration of information?

This is one of the central questions addressed by Attention and Performance XVI and gives rise to two more specific questions:

1. What is the time grain of constraint propagation?
2. How is the integration of constraints organized computationally?

There seems to be a degree of consensus on the first question, and a good deal of debate about the second.

Regarding the time grain, it appears that the propagation of constraints takes place *continuously*, so that the total support for each alternative in a selection task is continuously updated. Propagation is also apparently very rapid, so that selection and action become possible very shortly after sufficient support from a combination of sources is available. These points seem to be reasonably well established for auditory and visual on-line comprehension tasks, where they are supported through a large body of research by Marslen-Wilson and his colleagues (e.g., Marslen-Wilson 1987), and many models of language perception and comprehension incorporate these assumptions. Tanenhaus et al. (chap. 18) provide a very nice illustration of these points. When a subject must select a target for action from a combination of visual and auditory inputs, selection can occur very shortly after the auditory input is sufficient to rule out all but one of the alternatives consistent with the visual information. This shows that information from the auditory and the visual modalities is integrated in real time in the on-line comprehension process. The on-line character of the integration process is also reflected in the N400 component of the evoked potential, which reflects the combination of syntactic, semantic, lexical, and even episodic factors that influence the identification of words in context (see Kutas and King, chap. 20).

A further point relevant to the time course of integration is presented by MacDonald (chap. 17). She notes that when an ambiguity is encountered in language, the contextual information necessary to resolve the ambiguity often comes afterward, not before. In keeping with her claim that a particular source of information only has an effect when the alternatives for selection are at relatively balanced and intermediate levels of activation, she finds that subsequent context can influence ambiguity resolution, but only if it arrives close enough in time so that the alternatives that it affects are still in balance. As time passes, whichever alternative is stronger tends to dominate, and subsequent sources of constraint become ineffective.

Regarding the organization of the computations that lead to integration, several chapters in *Attention and Performance XVI* take explicit stands on this question. There appear to be essentially three kinds of views:

1. The feedforward view. According to the feedforward view, processing propagates in one direction—either from input to some central decision-making system for perceptual and comprehension tasks or from some internal specification of the desired action to a specification of the overt response, for motor control tasks. In this approach integration occurs bottom-up, and constraints from multiple sources are integrated only at central levels.

2. The interactive view. According to the interactive approach, constraints are propagated bidirectionally through a multilayered processing system. In this view integration occurs at every level of processing, based on constraints propagated top-down and bottom-up.

3. The intrinsic integration view. According to this approach, many if not all of the levels of processing envisioned in the first two approaches are seen as descriptive conveniences rather than actual separate parts of the mechanism. Instead, a pattern of activation over an ensemble of units is taken to encompass several levels of representation simultaneously.

Visual word recognition models provide examples of each of these three views. Feedforward models (e.g., Paap et al. 1982) propagate information from features to letters to words; and interactive models (e.g., McClelland and Rumelhart 1981) propagate information in both directions. Golden (1986) provides an example of an intrinsic integration model, in that his model has a single layer of units. Each unit corresponds to a visual feature of the letter in a particular position within the word; thus, to process four-letter words, there must be four sets of visual feature units. A letter, in this model, is a particular pattern over the feature units in one of these sets, and a word is a pattern over all of these sets. The network is trained with a connectionist learning rule to settle into states corresponding to particular words. In this model, integration across levels is intrinsic, in that the same pattern specifies at once the features, the letters, and the whole word.[1]

In my own scientific career I have explored all three of these alternatives (McClelland 1976, 1979; McClelland and Rumelhart 1981, 1985). This is not the place to attempt an extensive discussion of the merits of these alternatives. However, a few remarks may be in order to clarify the current state of play in the debate.

The *feedforward view* can be seen as consistent with many modular approaches to cognitive functions: Stimulus-driven processes provide separate sources of support for various alternatives in a selection task, and these sources are then combined with contextual influences and/or constraints imposed by task instructions to determine the outcome. A version of the feedforward view, put forward by Massaro (chap. 16), has been justified by the adequacy of his fuzzy logical model of perception (FLMP) and similar models to account for a large body of data from a wide range of experiments on integration of information. These studies show that subjects often choose among responses based on the combination of two or more separately manipulated sources of information in a way that suggests that the subjects treat each source of information as though it provides independent evidence for each of the possible alternatives. Massaro takes this independence as support for the view that each source of information is analyzed separately from all the others, and the results of these analyses are then simply multiplied together to yield the total support for each alternative. Responses are then chosen probabilistically, with the probability of choosing a particular response equal to the total support for that response, divided by the sum over all alternatives of the total support for each.

Two issues related to this proposal can be noted. First, as Bülthoff and Yuille argue in chapter 3, and as Frisby, Buckley and Freeman illustrate with experimental data (chap. 4), it is often the case that the interpretation of one source of information depends on the status of other sources. For example, Frisby et al. find that in estimating surface orientation, subjects may rely heavily on global linear perspective cues provided by the surface as a whole, when other cues indicate the surface is planar; when other cues indicate that the surface has local curvature, the global linear perspective cue is not used. (In fact this is quite sensible because the global surface orientation is not predictive of local surface orientation if the surface has local curvature.) Bülthoff and Yuille suggest that cases of violation of independence might reflect "incorrect modularization." They point to the existence of numerous cases where the interpretation of one set of cues depends critically on the state of other cues; they argue that these constraints can only be properly captured by considering them jointly rather than independently.

Second, it should be noted that what counts as an independent source of constraint is quite flexible in Massaro's model (chap. 16). Consider, for example, the sources of constraint on the identification of a target letter when it occurs in the context of other letters. Specifically, consider the identification of a target display element as a *c* or an *e* when it is followed by the letters *_oin*. In the FLMP the visual features of the target display element itself constitute one source of information, and the context, taken as a whole, constitutes another. In this case, if the target display element were a *c* it would form a word with the context, so the context is taken as providing more support for the *c* alternative than for the *e* alternative. While this is clearly a correct statement of the facts of the matter, something is lacking: namely, a specification of the computations performed on the visual inputs arising in the other letter positions that result in this support. This aspect of the FLMP has always made me feel that we should view it as characterizing the influence of different sources of information on the *outcome* of processing—when these are in fact independent—but we should continue to ask mechanistic questions about how the computation is actually performed.

The *interactive view* of the integration of information is proposed by Inui in his introductory chapter. It is explicitly adopted by Kawato (chap. 14) for integrating constraints in motor control, and by Duncan (chap. 21) for integrating multiple influences on the allocation of attention. The interactive activation approach to perception and language processing (Rumelhart 1977; McClelland and Rumelhart 1981; McClelland 1987, 1991) is part of the background of this Attention and Performance meeting; Inui, Kawato, and Duncan provide complementary motivations for selecting an interactive approach.

Inui (chap. 1) notes the arrangement of the brain into separate areas, each apparently specialized for the representation of a particular type of information in a particular format, and points to the bidirectional connections between these areas as support for the idea that integration occurs through the bidirectional propagation of influences.

Kawato (chap. 14) is concerned with integration of constraints arising from the specification of a reaching task (e.g., move the tip of your finger from point A to point B in external space) and constraints arising from the motor system (e.g., make a movement that minimizes the total amount of change of the commands to the muscles). He argues that the computation of a movement sequence that jointly satisfies both sets of constraints is only possible if the different constraints are imposed on representations of the movement in different frames of references. The constraint to move from A to B must be imposed and evaluated in a frame of reference where the dimensions are the axes of external three-dimensional space, while the constraint to minimize the total amount of change in the commands to the muscles must be imposed on a representation of the movement in a frame of reference where the dimensions correspond to the motor commands to each of the muscles. Kawato argues that a computation that propagates constraints bidirectionally between these two frames of reference is the most efficient way to settle on a sequence of motor commands conforming to both sets of constraints. A similar argument for aspects of vision is alluded to by Inui (chap. 1).

Duncan's suggestion in chapter 21 that attention is an emergent property of an activation and competition process relies heavily on the idea that the propagation of influences on attention is bidirectional. One of the great attractions of this approach (and the precursor offered by Phaf, van der Heijden, and Hudson 1990) is that the architecture used for perceptual processing also provides the mechanism for the propagation of attentional influences. This proposal can be seen as a contemporary—and much more physiologically based—version of Neisser's (1967) suggestion that what is attended is essentially what is incorporated into the percept by the constructive process of perceptual interpretation.

Duncan's proposal seems to presuppose acceptance of an interactive account of constraint propagation. After reading chapter 16, however, we certainly cannot say that an interactive account is universally accepted; Massaro provides an extensive critique of the particular version of the interactive approach represented by the stochastic interactive activation model (McClelland 1991). Some of his particular points are addressed elsewhere (Movellan and McClelland 1995). Here I will just make one key observation.

It was shown in McClelland (1991) that the stochastic interactive activation model implements exactly the computations specified by Massaro's FLMP. McClelland's 1991 model differs from the original version (McClelland and Rumelhart 1981) in two ways. First, it assumes that the outcome of the interactive activation process is variable, either due to variability in the inputs to the interactive activation process or to inherent randomness in the processing itself, or both. Second, it assumes that response selection occurs by simply selecting the processing unit corresponding to the most active alternative—the one receiving the most support from the overall interactive activation process. The fact that the stochastic interactive activation model implements the computation prescribed by FLMP may seem surprising; it certainly contradicts the claim of Massaro (1989) that bidirectional propagation of activation cannot capture the independent influences of different sources of information. However, this result follows from well-known facts about the probability distributions over the possible states of stochastic, bidirectionally connected networks (e.g., Boltzmann machines), and from the limits imposed on these probability distributions by the architectural constraints built into the interactive activation model.

The architectural constraint that leads the interactive activation model to conform to the FLMP is worth understanding because it relates to questions about what sorts of brain architecture would be expected to lead to independence as captured by FLMP and what sorts of brain architecture would lead to strong interactions of different sources of constraint. The crucial variable is not a matter of whether there are bidirectional influences; it is instead a matter of whether the influences that act at a particular level are *structurally independent*. Two sources of influence are said to be structurally independent if there are no direct connections between the channels that carry these influences to the processing units representing the alternatives. Bidirectional connections into and out of the units representing the alternatives may allow the sources to influence each other, but only indirectly, via the units representing the alternatives. Structural independence is violated only if there are other ways in which the two sources of influence on the alternatives can influence each other. We can illustrate this notion of structural independence using the interactive activation model (fig. 24.1 depicts the model more completely and accurately than the more familiar figure from McClelland and Rumelhart 1981). Consider again the situation in which an observer has to make a choice between *e* and *c* as the interpretation of a particular target display character presented as the first element of a four-character display, where the remaining elements are *_oin*. In the interactive activation model, the target display

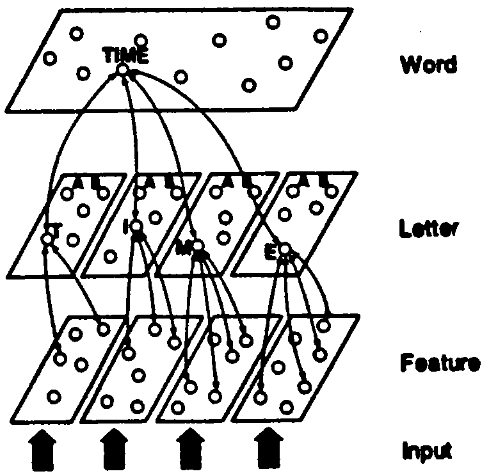

Figure 24.1 Interactive activation model of McClelland and Rumelhart (1981), showing four letter-processing channels each consisting of feature- and letter-level units, together with overarching word level. Units within rectangle are mutually inhibitory. Connections shown illustrate presence of bidirectional excitatory connections between mutually consistent features and letters, and between mutually consistent letters and words; only connections relevant to processing of single word "TIME" are shown. There is one unit for each letter-feature and each letter in each letter processing channel, and there is one unit for each word at word level. From McClelland (1985).

character and the context exert structurally independent influences on the units representing the letter alternatives in the first letter position: the units that carry information about the contents of the first display position to the alternatives have no direct connections with the units that carry information about the context. Interestingly, this independence is simultaneously true for all four letter positions. For each, the direct featural information and the contextual influences are structurally independent. It is also simultaneously true that at the word level, the contents of each letter position exert independent influences on each word alternative. Contextual influences arising from other sources that have structurally independent projections to the word level would likewise exert effects on word identification independent of those arising from the letter level.

What this argument shows, then, is that stochastic interactive activation implements an integration process simultaneously at different levels, and at different positions within a given level, and that the different sources of information exert their influences independently at all these different points in the processing system at the same time so long as these different sources are carried by structurally independent processing channels. Thus, it would appear that the stochastic interactive activation process, operating through bidirectional connections among units in networks that conform to the structural independence assumption, could represent the desired implementation of Massaro's FLMP.[2]

With this result in hand, we can now ask, why should the perceptual processing system ever adhere to the principle of structural independence? If processing is interactive, why should the interactions ever be constrained in this way? Is this, in fact, a residual form of imposed architectural modularity? Even though activation can propagate bidirectionally, so that higher-level influences can affect lower levels, the computation is nevertheless still tightly disciplined. Perhaps this limited form of interactive propagation of activation is structurally imposed by the brain to control the complexity of neural computation?

Structurally imposed constraints on interactive activation are certainly apparent in the brain; indeed, the topographic organization of many areas of the brain, and the fact that the propagation of information between maps maintains independence between moderately separated locations within each map, seem to suggest that there is some imposed modularity. However, it is worth noting that the topographic organization is refined and maintained by adaptive synaptic modification processes. Even if one begins with much fuller connectivity, adaptive synaptic modification processes will lead to structural independence where this reflects an actual statistical independence in the environment (Linsker 1986; Miller, Keller, and Stryker 1989). Structural independence may thus arise through adaptive connection modification processes in cases where the sources of constraint actually are conditionally independent. In this view, the ultimate basis for independence lies in the structure of the environment.

We finally turn to *the intrinsic integration view* with these considerations as context. This view was not discussed by any of the participants at Attention and Performance XVI, but it has become a force in connectionist circles since the advent of powerful learning algorithms such as backpropagation. As previously noted, the key insight to emerge from this connectionist work is the realization that standard cognitive units (letters, words, phrases, schemata, etc.) may be only descriptive conveniences, applied to characterize emergent properties of processing systems rather than reified directly within them. This idea is represented in connectionist models that span a range of specific applications (Elman 1990; Elman and Zipser 1987; Golden 1986; Rumelhart et al. 1986; St. John and McCelland 1990; Van Orden, Pennington, and Stone 1990). The burden of the simulations is to show how behavior thought to indicate the existence of various sorts of cognitive units can arise from a network with minimal prior structural commitments through an adaptive learning process. In this view, letters, words, and phrases do not exist on separate levels but are simply the coherent and relatively independent subpatterns of larger, more coherent whole patterns. Likewise, conceptually distinct types of constraints, such as syntactic versus semantic constraints on sentence structure or lexical versus phonotactic constraints on phoneme identity, are not reified in architectural distinctions in such models. Several models of this type (Golden 1986; Shillcock et al. 1992; St. John and Gernsbacher 1995) have been able to account for a number of findings previously accounted for by more structured interactive networks that honor the units and levels of the classicists.

Evidence from physiology might help determine to what extent different levels of description of structure offered by psychologists and linguists are represented separately in the brain, and to what extent they communicate bidirectionally. This evidence may be available for constraints on perception and action, which can be studied in primates but is more limited where the function in question only exists in humans, as in the case of language. There is a vast literature on the neuropsychology of language, together with a growing body of studies of language in the brain based on the use of non-invasive techniques, but the limitations of these methods leave considerable room for uncertainty about the representations used. One of the findings reported by Kutas and King (chap. 20) using event-related potentials (ERPs) illustrates some of the difficulties. Kutas and King found that variations in syntactic complexity lead to differences in brain activity recorded at frontal cites. But these findings do not tell us whether the actual computation of complex constructions occurs in frontal areas, or whether they indicate a generalized frontal response to variations in task difficulty. The authors are therefore cautious in their interpretation of the meaning of their findings for the functional organization of language processing.

In any case, as already suggested, the organization of the visual and somatosensory areas of the brain appears to suggest a great deal of discipline to the interactions among processing units. It is clearly not the case that the brain is, as Touretzky (personal communication) once put it, an "undifferentiated mass of connectoplasm," nor did it start that way at birth, to be shaped only by the structure present in experience. At the same time, however, evidence of sensitivity to structure in the behavior of a system should not necessarily be taken to mean that descriptive constructs such as phonemes and words will necessarily be reflected in hard-wired processing structures of the sort suggested by the interactive activation models.

Representations for Integration

Are there supramodal representations for integration across modalities, or is integration merely mutual constraint between modality- or dimension-specific representations?

Some of the most exciting ideas and findings reported at Attention and Performance XVI relate to this question. Introspection suggests to many people that we maintain a single, stable representation of the external world across fixations. Auditory, tactile, and visual inputs all contribute to the contents of this representation and so it seems to be supramodal in some sense, and it seems plausible that many of the examples of information integration reviewed in the chapters in this volume could arise from convergence of influences on this stable representation from different sources. Yet several of the papers presented at Attention and Performance XVI suggest a rather different conception of what may be happening in the brain. Instead of maintaining a supramodal representation for the integration of multiple sources of information, the brain may instead maintain many separate representations, each with its own frame of references that interactively constrain each other.

Both Colby (chap. 7) and Graziano and Gross (chap. 8) propose versions of this multiple coordinated reference frames view. A key aspect of this view is the notion that as each body part moves relative to other parts, and as the body moves relative to external space, the brain continually adjusts the mapping between these reference frames, and continually updates each representation to track the locations of objects in each. While the remapping process may seem complex, it has one beautiful result—it allows us to know where external objects are with respect to different parts of our body. If this information is given, the computation of the movement that must be made to bring an object into contact with a particular body part is greatly simplified.

Striking evidence of the remapping process is reported by Graziano and Gross (chap. 8), who recorded from neurons that represent the presence of an object in a receptive field defined with respect to one of the monkey's limbs. As the limb moves, stationary objects in the environment may come in and out of this receptive field. Likewise, as the eye moves, the region of retinocentric space that corresponds to the receptive field of the neuron continually changes. Thus locations in both allocentric and retinocentric space are continually remapped to locations in the limb-centered frame of reference. Graziano and Gross and Colby (chap. 7) both suggest that there may be neurons that code positions of objects with respect to many different body parts—amounting perhaps to scores of parallel and coordinated frames of reference!

A further, crucial finding described by Colby in chapter 7 is the observation that as the animal gets ready to initiate a movement of the body part in question, remapping can occur in anticipation of the effect of the movement—and neurons begin to fire in anticipation of targets that will be in their receptive field after the movement, Thus far, this demonstration has only been made with neurons that encode the locations of target stimuli with respect to the direction of gaze of the eyes. It will be interesting to see if such anticipatory remapping also occurs in other coordinate systems.

What signals enable the remapping process? The finding of Roll et al. (chap. 12) that movement illusions (apparent motion of a body part or of stimuli in extrapersonal space) can be induced by proprioceptive stimulation suggests that proprioception contributes to the remapping process; another source or remapping is "efference copy," recurrent feedback of the motor command. However, anticipatory remappings cannot reflect proprioception. Even efference copy is suspect an a source, though it could provide the relevant information if "intention" is coded neurally in terms of low-rate, anticipatory firing of the same neurons that will ultimately trigger the movement.

Exactly how coordination is maintained between effector-specific references frames remains somewhat obscure. Graziano and Gross (chap. 8) do not deny there may be brain areas that maintain a representation of the observer's place within an allocentric reference frame, but it plays no privileged or crucial role for them in mediating between other frames of reference; rather, it is just one of the many representations that is continually remapped like all of the others. The authors do note that there are areas in the parietal lobes where vision, touch, and proprioception come together, but they do not identify them with specific frames of reference. One possibility is that these areas capture conjunctions of inputs that are useful for mapping between frames of reference, much as hidden units would in a connectionist network.

Dimensions for Integration

Are there dimensions that provide a strong organizing frame for integration? Are time and space examples of such dimensions?

Findings relevant to these questions come from a number of the papers presented at Attention and Performance XVI. I will consider the use of time together with the role of attention in the next section. The use of space as an organizing frame for integration is partially addressed by the answer to the previous question. It appears that there may be many spatial frames of reference all in use at one time, and constraints may be integrated simultaneously within each (cf. the discussion of Kawato's ideas on the integration of constraints in extrinsic and muscle command coordinate systems in section 24.2.3).

A perspective that at first may seem to contradict the idea that constraints are integrated in spatial frames of reference emerges in chapter 6, by Irwin and Andrews. Summarizing twenty years of research on integration of information over successive fixations in vision, Irwin and Andrews note that there is considerable evidence of integration of information over successive fixations, but that this integration appears not to occur within a common spatial frame of reference that survives across eye movements. As evidence of this, they note that integration proceeds just as well, whether or not the stimuli giving rise to the information in the first fixation maintain their spatial position after the eye movement. Integration is even unaffected by alteration of the physical form of the stimulus carrying the constraining information. The authors cite evidence that partial information obtained from a parafoveal presentation of a word on one fixation can facilitate the identification of the same word when it is fixated after an eye movement, whether or not the parafoveal and target presentations maintain a fixed position in external space, and whether or not the word switches from uppercase to lowercase type between fixations. The implication is that integration takes place via the accumulation over fixations of constraints on the abstract identity of items, independent of details of the exact form and spatial location of the stimulus that is the source of these constraints.

Within the context of the proposals of Graziano and Gross (chap. 8), these findings suggest that integration for the specification of the identity of objects occurs within a representational system quite different from the representational systems used to maintain information about the locations of objects in space relative to the body. This suggestion is consistent with the notion that there are separate processing streams for "what" and "where" (as Ungerleider and Mishkin, 1982, put it) or for "what" and "how" (as Goodale, chap. 15, and others in this volume put it), and that integration of "what" information takes place in a representation specifically structured to combine constraints on object identity, just as other representations are specifically structured to combine different sources of constraint on the spatial locations of objects relative to a particular body part. It must be added, however, that considerable work remains to be done to understand how the brain maintains coordination of information in these different representations.

Attention and Integration

What role does attention play in the integration of information?

It may be fitting to end our consideration of a symposium of the Society for Attention and Performance that focused on the integration of information with a consideration of the role of attention in information integration. As Duncan (chap. 21) stresses, every act of attention is an act of selection of some information and suppression of other sources. What purpose does this selection serve? Several of the contributors to this volume have advanced our understanding of this issue.

In Treisman's (1988) feature integration theory (FIT), attention is necessary to select the attributes of a single object so that they may be joined together into a single percept, represented separately from the features of other objects. A key aspect of Treisman's idea is that feature integration requires the sequential allocation of attention to just one object at a time. Each sequential act of attention selects the features of a single object and binds them together, opening an "object file" for the object. The file becomes a basis for immediate report of the features of the object as well as a record of the object in long term memory.

In response to Inui's question on the role of time in information integration (chap. 1), we can observe that in FIT, the activation of two features at the same time causes both to be entered into the object file; the only thing that prevents false conjunctions is the allocation of attention to a single object at a time. Thus time becomes the basis for conjoining features together into object representations.

The idea that time and attention play a special role in forming object representations has now been extended considerably by Hummel and

Stankiewicz (chap. 5). According to their model, there are two forms of perceptual processing of an object: a fast, initial, parallel form of processing that is capable of serving as the basis of recognition if the object has been seen before from the very same viewpoint; and a much slower, inherently time-dependent process, in which a structural description of the object is built up over time by considering each part, and its relation to the object as a whole, in succession. Segmentation itself occurs through a constraint-satisfaction process, in which elements that go together to form a coherent part mutually reinforce each other's activation and suppress simultaneous activation of elements belonging to other parts. In this way, the time over which attention is directed toward an object as a whole comes to be divided into shorter intervals, with attention cycling over the objects' parts across these shorter intervals.

The two modes of processing in the Hummel and Stankiewicz theory have strong parallels to the two modes of processing attested in FIT, where it is assumed that parallel processing can be used to detect the presence of a target, as long as that target is defined as a single feature, rather than a specific conjunction of features that must co-occur together in the same location (see Cohen and Ivry, chap. 22, for an exploration of evidence relevant to these ideas from studies of visual search among orientations and directions of movement). Sequential attention is only necessary to establish the correct bindings of features together in the same object. Much the same applies to the theory of Hummel and Stankiewicz, but with one interesting and important difference. In their model, the parallel processing of an entire object gives rise to a representation that is scale- and translation-invariant, though not invariant with respect to rotation. The representation captures the global shape of the object as a whole and of the shapes and relative locations of its parts, as seen from the observer's viewpoint. This extends the power of the parallel processing mechanism considerably beyond its original conception in FIT, allowing it to provide a basis for accurate recognition of familiar objects seen from familiar viewpoints.

The notion that the initial parallel processing of an item may establish at least a primitive representation of the object is consistent with the findings from the new work reported in Treisman's Association Lecture (Treisman and DeSchepper, chap. 2). They find that just one presentation of an object, even when it is not the focus of attention, gives rise to object specific aftereffects that can persist for many weeks. Because these aftereffects dissociate completely from explicit recognition memory, we would not want to identify them with the formation of an explicit episodic memory for the unattended object. Instead, in keeping with other interpretations of implicit perceptual learning phenomena, we might view them as aftereffects within the perceptual processing system itself.

The suggestion that parallel processing of a display may be more powerful and structured than had previously been thought brings us back again to the question of the need for focal attention. Both Treisman and DeSchepper (chap. 2) and Hummel and Stankiewicz (chap. 5) discuss the possibility that the perceptual processing of unattended objects may be incomplete. The results of Treisman and DeSchepper indicate that focal attention may also be necessary to form an explicit episodic memory. Given that such memories involve the integration of the representation of an object with the situation in which it occurs as well as its relation to the subject as experiencer (Tulving 1983), it may be best to consider these representations constructed within the focus of attention, not as disembodied "object files," but as conjunctive representations of the object in association with other aspects of the external and internal context represented in the distributed pattern of activation that corresponds to whatever is within the span of the subject's attention at the time the object is experienced. It appears that such episodic memories are initially formed in the hippocampus, where the state of activation over all of the higher-level representation systems of the brain may be brought together and interassociated via rapid synaptic modification (see McClelland, McNaughton, and O'Reilly 1995 for a recent discussion of this possibility).

24.3 CONCLUSION

The theme of integration of information in perception, language, attention, and action is very broad, and any attempt to summarize the state of our knowledge about it is surely doomed to oversimplification. Suffice it to say in concluding these reflections that the study of information integration within and across domains suggests many common principles and many issues that deserve much fuller consideration. Attention and Performance XVI certainly did not settle all the issues. I think Toshio Inui would join me in the hope that the meeting and this volume have brought out what the most significant issues are, together with some of the most promising directions that are being pursued to address them.

NOTES

Preparation of this chapter was supported by grants MH00385 and MH47566 from the National Institute of Mental Health. I thank Toshio Inui and David Plaut for comments on the first draft. Correspondence concerning this article may be sent to the author at the Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213, or via E-mail to mcclelland+@cmu.edu.

1. In general, hidden units—units that do not correspond directly to features—may be required if such a model is to be able to learn each word as a separate attractor in a way that successfully mimics the interactive models. However, it is not at all clear that these hidden units will correspond to letters or words per se, or that it would necessarily make sense to view them as constituting a separate level of representation, "above" the feature level.

2. In fact, the architecture of the interactive activation model specifically predicts that some sources of influence on the selection among certain alternatives will not be independent. Consider the influence of separate manipulation of the context letters in positions 2, 3, and 4 on forced-choice identification of the letter in position 1. Because the context letters all

interact with each other at the word level, their influences are not generally expected to be independent. Movellan and McClelland (1995) have confirmed that independence is violated in just these conditions.

REFERENCES

Bever, T. G. (1970). The cognitive basis for linguistic structure. In J. R. Hayes (Ed.), *Cognition and the development of language*, 279–362. New York: Wiley.

Chomsky, N. (1980). Rules and representations. *Behavioral and Brain Sciences, 3*, 1–61.

Clifton, C., and Ferreira, F. (1989). Ambiguity in context. *Language and Cognitive Processes, 4*(SI), 77–104.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science, 14*, 179–211.

Elman, J. L., and Zipser, D. (1987). *Discovering the hidden structure of speech*. Technical report 8701. San Diego: University of California, Institute for Cognitive Science.

Fodor, J. A. (1983). *Modularity of mind: An essay on faculty psychology*. Cambridge, MA: MIT Press.

Fodor, J. A., and Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28*, 3–71.

Golden, R. M. (1986). A developmental neural model of visual word recognition. *Cognitive Science, 10*, 241–276.

Hinton, G. E. (Ed.). (1991). *Connectionist symbol processing*. Cambridge, MA: MIT Press.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, 79*, 2554–2558.

Jordan, M. I. (1986). Attractor dynamics and parallelism in a connectionist sequential machine. In *Proceedings of the 8th Annual Conference of the Cognitive Science Society*, 531–546. Hillsdale, NJ: Erlbaum.

Jordan, M. I., and Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science, 16*(3), 307–354.

Just, M. A., and Carpenter, R. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review, 99*, 122–149.

Linsker, R. (1986). From basic network principles to neural architecture: Emergence of spatial-opponent cells. *Proceedings of the National Academy of Sciences, 83*, 7508–7512.

Luria, A. R. (1966). *Higher cortical functions in man*. New York: Basic Books.

Luria, A. R. (1973). *The working brain*. New York: Penguin.

Marcus, M. P. (1980). *A theory of syntactic recognition for natural language*. Cambridge, MA: MIT Press.

Marr, D. (1982). *Vision*. San Francisco: W. H. Freeman.

Marslen-Wilson, W. D. (1987). Functional parallelism in spoken word recognition. *Cognition, 25*, 71–102.

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

McClelland, J. L. (1976). Preliminary letter identification in the perception of words and nonwords. *Journal of Experimental Psychology: Human Perception and Performance, 2*, 80–91.

McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review, 86,* 287–330.

McClelland, J. L. (1985). Putting knowledge in its place: A scheme for programming parallel processing structures on the fly. *Cognitive Science, 9,* 113–146.

McClelland, J. L. (1987). The case for interactionism in language processing. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading,* 1–36. London: Erlbaum.

McClelland, J. L. (1991). Stochastic interactive activation and the effect of context on perception. *Cognitive Psychology, 23,* 1–44.

McClelland, J. L., McNaughton, B. L., and O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102,* 419–457.

McClelland, J. L., and Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception. I. An account of basic findings. *Psychological Review, 88*(5), 375–407.

McClelland, J. L., and Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General, 114,* 159–188.

Miller, G. (1962). Some psychological studies of grammar. *American Psychologist, 17,* 748–762.

Miller, K. D., Keller, J. B., and Stryker, M. P. (1989). Ocular dominance column development: Analysis and simulation. *Science, 245,* 605–615.

Morton, J. (1969). The interaction of information in word recognition. *Psychological Review, 76,* 165–178.

Movellan, J. R., and McClelland, J. L. (1995). Stochastic interactive activation, Morton's Law, and optimal pattern recognition. Technical Report PDP.CNS.95.4. Pittsburgh: Department of Psychology, Carnegie Mellon University.

Neisser, U. (1967). *Cognitive psychology.* New York: Appleton, Century Croft.

Paap, K. R., Newsome, S. L., McDonald, J. E., and Schvaneveldt, R. W. (1982). An activation-verification model of letter and word recognition: The word superiority effect. *Psychological Review, 89,* 573–594.

Phaf, R. H., van der Heijden, A. H. C., and Hudson, P. T. W. (1990). SLAM: A connectionist model for attention in visual selection tasks. *Cognitive Psychology, 22,* 273–341.

Rumelhart, D. E. (1977). Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and performance VI,* 573–603. Hillsdale, NJ: Erlbaum.

Rumelhart, D. E., Smolensky, P., McClelland, J. L., and Hinton, G. E. (1986). Schemata and sequential thought processes in PDP models. In J. L. McClelland, D. E. Rumelhart, and the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition.* Vol. 2, *Psychological and biological models,* 7–57. Cambridge, MA: MIT Press.

Shillcock, R., Lindsey, G., Levy, J., and Chater, N. (1992). A phonologically motivated input representation for modeling of auditory word perception in continuous speech. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society,* 408–413. Hillsdale, NJ: Erlbaum.

Simpson, G. B. (1981). Meaning dominance and semantic context in the processing of lexical ambiguity. *Journal of Verbal Learning and Verbal Behavior, 20,* 120–136.

Sopena, J. M. (1991). *ERSP: A distributed connectionist parser that uses embedded sequences to represent structure.* Technical report UB-PB-1-91. Barcelona: Department de Psicologia Basica, Universitat de Barcelona.

St. John, M. F., and Gernsbacher, M. A. (1995). Syntactic comprehension: Practice makes perfect and frequency makes fleet. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

St. John, M. F., and McClelland, J. L. (1990). Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence, 46, 46*, 217–257.

Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 18*, 645–659.

Taraban, R., and McClelland, J. L. (1988). Constituent attachment and thematic role assignment in sentence processing: Influences of content-based expectations. *Journal of Memory and Language, 27*, 597–632.

Teitelbaum, P. (1967). *Physiological psychology; Fundamental principles*. Englewood Cliffs, NJ: Prentice-Hall.

Treisman, A. M. (1964). Verbal cues, language, and meaning in selective attention. *American Journal of Psychology, 77*, 206–219.

Treisman, A. (1988). Features and Objects: The Fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology, 40A*, 201–237.

Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford University Press.

Ungerleider, L. G., and Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingli,. M. A. Goodale, and R. J. W. Mansfield (Eds.), *Analysis of sisual behavior*. Cambridge, MA: MIT Press.

Van Orden, G. C., Pennington, B. F., and Stone, G. O. (1990). Word identification in reading and the promise of subsymbolic psycholinguistics. *Psychological Review, 97*(4), 488–522.