A parallel distributed processing approach to automaticity

JONATHAN D. COHEN, DAVID SERVAN-SCHREIBER
Carnegie Mellon University and University of Pittsburgh

JAMES L. McCLELLAND
Carnegie Mellon University

We consider how a particular set of information processing principles, developed within the parallel distributed processing (PDP) framework, can address issues concerning automaticity. These principles include graded, activation-based processing that is subject to attentional modulation; incremental, connection-based learning; and interactivity and competition in processing. We show how simulation models, based on these principles, can account for the major phenomena associated with automaticity, as well as many of those that have been troublesome for more traditional theories. In particular, we show how the PDP framework provides an alternative to the usual dichotomy between automatic and controlled processing and can explain the relative nature of automaticity as well as the fact that seemingly automatic processes can be influenced by attention. We also discuss how this framework can provide insight into the role that bidirectional influences play in processing: that is, how attention can influence processing at the same time that processing influences attention. Simulation models of the Stroop color-word task and the Eriksen response-competition task are described that help illustrate the application of the principles to performance in specific behavioral tasks.

This special issue surveys current thinking about the concept of automaticity. In this article, we consider this issue within the context of a set of principles of information processing formulated in the broad framework of parallel distributed processing (PDP). We will show how these principles make it possible to construct models that capture the major phenomena of automaticity, as well as many findings that have been seen as problematic for the usual dichotomy between automatic and controlled processes. In particular, we will show how our framework allows us to capture the inescapable conclusions that (a) automaticity is a relative matter, and (b) processes that are automatic by some criteria are nevertheless susceptible to interference and influences of attention. We will also show that the principles provide ways of understanding bidirectional influences between processing and at-

AMERICAN JOURNAL OF PSYCHOLOGY
Summer 1992, Vol. 105, No. 2, pp. 239–269

tention: that is, that attention influences processing while at the same time processing influences attention.

The article is structured as follows. First, we present the processing principles. Then, we consider the basic phenomena of automaticity and illustrate how these phenomena can be captured in a model of performance in the Stroop interference task. This model incorporates most but not all of the principles, and as we shall explain, it now appears that incorporation of the rest of the principles would allow the model to account for the mutual dependency of processing and attention, and to overcome several specific empirical shortcomings. We then illustrate the usefulness of the full set of principles by applying them to an interesting pattern of data from the Eriksen response-competition paradigm that could not be accounted for by the Stroop model. The discussion considers several general issues related to attention in light of models based on these principles, including the concept of "resources" and the distinction between automatic and controlled processes.

Principles of information processing

McClelland (1992) has articulated a small set of basic principles that appear to provide a promising framework for modeling a broad range of information processing phenomena. These principles presuppose that information processing takes place in a PDP system (Rumelhart, Hinton, & McClelland, 1986). A PDP system is simply a system in which processing occurs through the interactions of a large number of simple, interconnected processing elements called units. These elements may be organized into modules, each containing a number of units; sets of modules may be organized into pathways, each containing a set of interconnected modules. Pathways may overlap, in that they may contain modules in common. Processing in a PDP system occurs by the propagation of activation among the units, via weighted connections. The knowledge that governs processing is stored in the weights of the connections, and the effects of experience on information processing are captured by changes to the connection weights.

The PDP framework is extremely broad, and can be used to address a very wide range of different modeling goals, from efforts to capture the detailed properties of specific neural circuits to efforts to solve problems in artificial intelligence that have not yielded to more traditional symbolic approaches.

The PDP framework has also been applied to psychological modeling, and it has been extremely useful in this regard; but it is sufficiently broad that it does not provide adequate guidance or constraint without further assumptions. To constrain the further development of our own theoretical efforts, we have constructed the following provisional list of principles:

1. The activation of each unit is a graded, sigmoid function of its input.
2. Activation propagates gradually in time.
3. The activation process is intrinsically variable.
4. Learning (by connection adjustment) is also gradual, and is driven by differences between the obtained activation value and the one representing the correct response.
5. Attentional influences occur through the modulation of processing in one or more pathways as a result of the pattern of activation in another.
6. Between-module connections are bidirectional and excitatory, so that processing is interactive.
7. Within-module connections are bidirectional and inhibitory, so that processing is competitive.

We do not go into the full motivation for each principle in this article, because this would take us too far afield; this is spelled out in McClelland (1992). We focus instead on the relevance of the principles to issues of automaticity and attention.

We also want to stress that we do not take this set of principles as the final word. Rather, we take it as a provisional starting-place and guide for research. No doubt there are other principles in addition to these, and some or all of the principles will require further refinement.

The principles are stated in qualitative terms, without specific detailed quantitative assumptions. Although particular models must be formulated in terms of specific quantitative assumptions, we have found repeatedly that these details are relatively unimportant. It does not matter, for example, what the exact form of the graded sigmoid function is, or whether the intrinsic noise is Gaussian or uniformly distributed in a bounded interval.

Basic aspects of automaticity

As other authors in this volume point out, the term automaticity encompasses a number of different phenomena that often vary from one definition to another. Nevertheless, there are a core set of phenomena that seem to recur in most discussions of automaticity:

1. an increase in speed of performance with practice following a power law
2. diminishing requirements for attention with practice, with
3. a concomitant release from attentional control—or involuntariness (i.e., the involuntariness of automatic processes)

4. immunity from interference with competing processes, and
5. the requirement that practice be "consistently mapped" for these phenomena to develop.

Many discussions have treated automaticity as an all-or-none phenomenon. That is, a process is either automatic or "controlled." A classic example of this is the widely accepted account of the Stroop effect (e.g., Posner & Snyder, 1975): Word reading is considered to be automatic because it is fast, it produces interference even when subjects attempt to ignore the word, and it is not subject to interference by ink color. In contrast, color naming is considered to be controlled because it is slower, it can be voluntarily inhibited (thereby failing to interfere with word reading), and it is subject to interference.

Recent evidence suggests, however, that the attributes of automaticity can develop gradually with practice and, furthermore, that they may depend on the context in which they are evaluated. For example, MacLeod and Dunbar (1988) demonstrated that color naming shows all of the attributes of automaticity when it is placed in competition with a novel task, such as producing color words as names that have been arbitrarily assigned to shapes (see Figure 1). However, extensive practice with shape naming led to a gradual reversal of interference effects, with the color-naming task eventually reassuming its traditional role as the slower task, subject to but not able to produce interference. These findings suggest that there is a continuum of automaticity, and that speed of processing and interference effects may indicate the relative position of two tasks along this continuum, rather than necessitating that one be automatic and the other controlled.

The Stroop Model

In this section, we describe a PDP model that captures all of these aspects of automaticity, as they arise in the Stroop color-interference task, in terms of the first five principles of information processing listed above. As we shall see, the model accounts for a large number of basic and sometimes puzzling findings in ways that directly reflect the operation of the principles enumerated. After presenting the model and these successes, we will turn to a number of further considerations that implicate the remaining principles of interactivity and competition.

The model is shown in Figure 2. In brief, it consists of two processing pathways, one for color naming and one for word reading, and a task demand module that can selectively facilitate processing in either pathway. Simulations are conducted by activating input units corre-

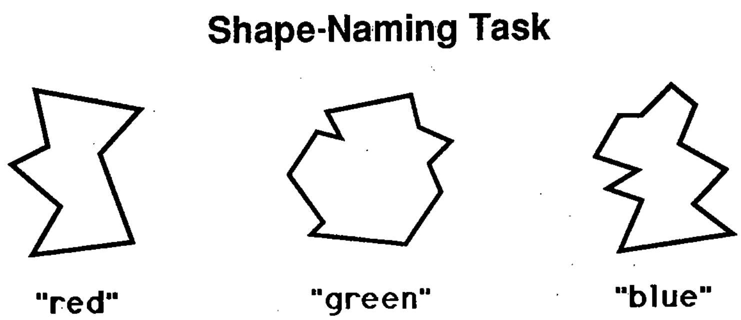

Figure 1. Training stimuli of the type used by MacLeod and Dunbar (1988) for the shape-naming task. Each of four shapes was assigned an arbitrary color name, which the subjects had to learn. *Note.* Figure after Cohen et al., 1990. Copyright 1990 by the American Psychological Association. Reprinted by permission.

sponding to stimuli used in an actual experiment (e.g., the input unit in the color-naming pathway representing the color red) and the appropriate task demand unit. Activation is then allowed to spread through the network. This leads to activation of the output unit corresponding to the appropriate response (e.g., *red*). Reaction time is linearly related to the number of cycles it takes for an output unit to accumulate a specified amount of activation. Training was more extensive on the word-reading than on the color-naming task, capturing the assumption that subjects have more extensive experience with the former than with the latter. Similar results would obtain if the network were given more consistent training on one task than the other, in agreement with the observation that consistency as well as amount of practice is important for the development of automaticity (e.g., Schneider & Shiffrin, 1977). This simple model is able to capture a number of empirical findings associated with the Stroop task (see Figure 3) and the development of automaticity in general.

Empirical and simulation results

Speed improvements and the power law. The model provides a straightforward account of the relationship between practice and speed. Additional training on the word-reading task resulted in the development of larger connection weights in that pathway, and therefore more rapid spread of activation along that pathway, with a corresponding decrease in reaction time. In addition, the model demonstrates the universal finding that, with practice, speed increases (and standard deviation decreases, Logan, 1988) according to a power law. This stems from two of our principles: incremental, difference-based learning; and a graded, sigmoidal activation function. The model was

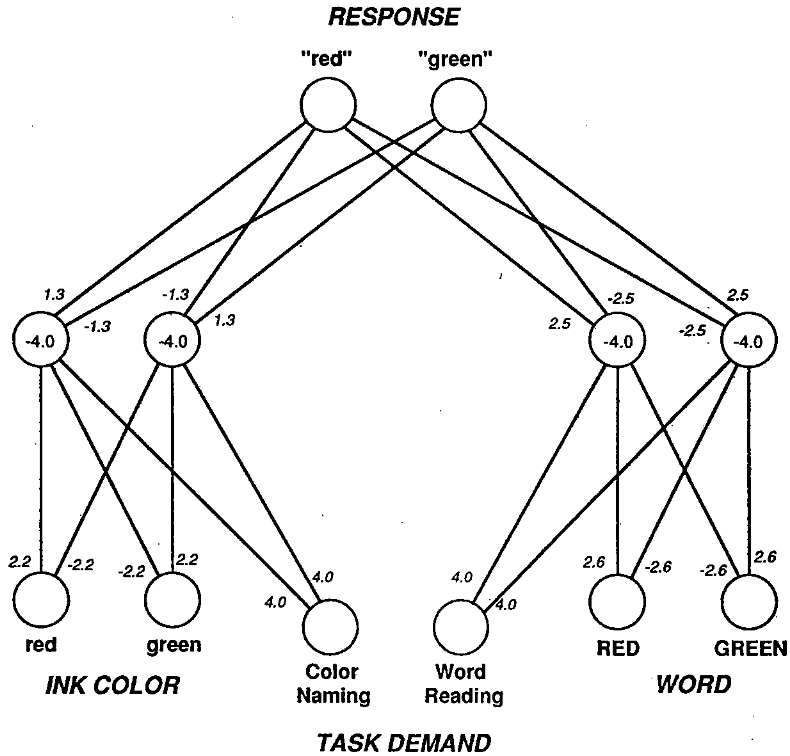

Figure 2. Diagram of the network used for the Stroop model, showing the connection strengths after training on the word-reading and color-naming tasks. Strengths are shown next to connections; biases on the intermediate units are shown inside the units. Attention strengths (i.e., from task demand units to intermediate units) were fixed, as were biases for the intermediate units. The values were chosen so that when the task demand unit was on, the base input for units in the corresponding pathway was 0.0, while the base input to units in the other pathway was in range of −4.0 to −4.9, depending upon the experiment (see text). *Note.* Figure after Cohen et al., 1990. Copyright 1990 by the American Psychological Association. Reprinted by permission.

trained using the back propagation learning algorithm of Rumelhart, Hinton, and Williams (1986). The details of the algorithm are not relevant here, but the fact that the algorithm is incremental, and that the sizes of the changes that are made are proportional to the magnitude of the difference between actual and desired output, is relevant. The amount that each connection weight is changed in each training trial is proportional to how much the asymptotic activations of the response units in the network differ from the desired output, which in this case is taken to be maximal activation of 1.0 for the correct response unit, and minimal activation of 0.0 for all other responses.

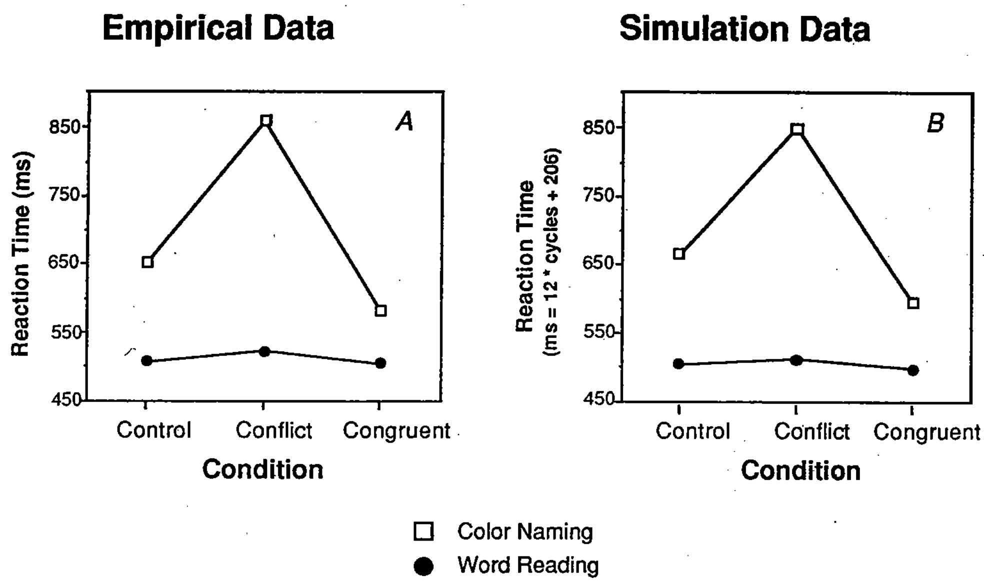

Figure 3. Performance data for the standard Stroop task. Panel A shows data from an empirical study (after Dunbar & MacLeod, 1984). Panel B shows the results of the model's simulation of this data. *Note.* Figure after Cohen et al., 1990. Copyright 1990 by the American Psychological Association. Reprinted by permission.

Early in training, this difference is likely to be large, so sizable changes will be made to the connection strengths. As the appropriate set of strengths develops, the error gets smaller and therefore so do the changes made to the connections.

A deceleration of speedup with practice also results from the fact that as connections get stronger, subsequent increases in strength have less of an influence on activation (and therefore reaction time). This is because of the sigmoidal shape of the activation function (see Figure 4): Once a connection (or set of connections) is strong enough to produce an activation close to 0.0 or 1.0, further changes will have little effect on that unit. Thus, smaller changes in strength, as well as the smaller effects that such changes have, progressively reduce the speedup of reaction time that occurs with practice. In our simulations, this pattern of diminishing returns adheres to the form of a log-log relationship; however, a formal analysis of these factors, as well as their relationship to the power law, remains to be done.

Interference effects. As seen in Figure 3, the model also reproduces the relative amounts of interference and facilitation observed in the word-reading and color-naming tasks. These effects are attributable to a pair of interacting factors: the relative strengths of the connections in the two competing pathways, and the modulatory effects that at-

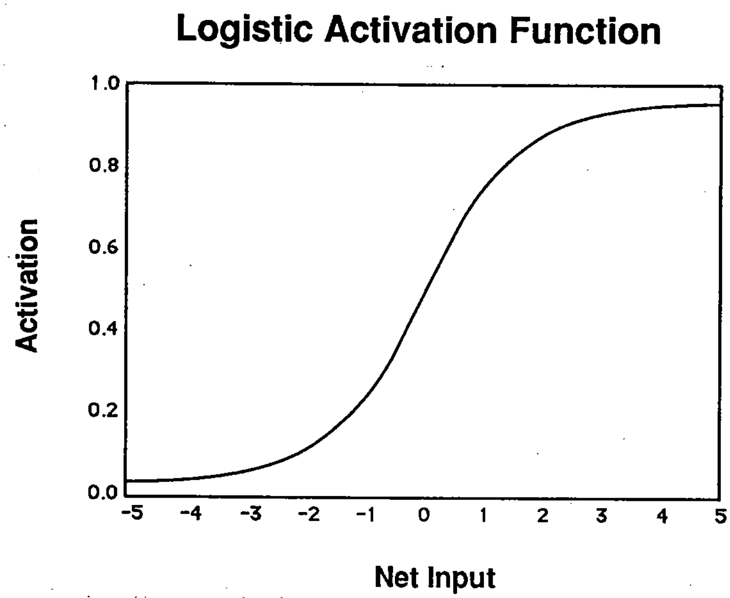

Figure 4. The logistic function, an example of a graded, sigmoid function. Note that the slope of this function is greatest when the net input is 0.0, and decreases when the net input is large in either the positive or negative directions.

tention has on processing in these pathways. Attention is implemented in the model as a pattern of activation over units in the task demand module (see Figure 2). The units in this module have connections to the intermediate units in each processing pathway such that activation of the unit for a given task sends input to the intermediate units in the corresponding pathway. This input increases the activation level of the corresponding intermediate units from a very low value, where the activation function is relatively flat, to a higher value, where the slope of the activation function is steeper and units are more sensitive to their input. Thus, the attentional mechanism takes advantage of the sigmoid shape of the activation function, to produce a modulatory influence on processing: Failure to allocate attention to a particular pathway reduces, but does not completely eliminate, stimulus-driven activation in that pathway. The amount of activation in an unattended pathway depends upon the strength of connections in that pathway. This is seen in the results of the Stroop simulation.

When the task is to name the color, connection weights in the word pathway are sufficient to allow some activation to flow along this pathway, enough so that when the word agrees with the color there is facilitation, and when it conflicts there is interference. This flow of activation along the word pathway, in the absence of attention, captures the involuntary or "automatic" nature of this process. In contrast, when the task is to name the word, there is an almost undetectable amount of interference. The reason for this is that in this case activation builds up very fast through the word pathway, because of the combined effects of the strong connections and the increased sensitivity due to attention. This rapid increase in activation through the word pathway has the effect of minimizing the effect that other factors can have on the time to reach the response threshold.

The phenomena that have been discussed so far, speedup with practice, and an asymmetry of interference effects between processes with different amounts of practice, are readily accounted for by a number of other theories (cf. Anderson, 1983; Logan, 1980, 1988). However, the simple model we have presented accounts for a number of other phenomena that have not, to date, been explained by other means.

Asymmetry of facilitation versus interference. First, note in the empirical data that the size of the interference effect is significantly greater than the facilitation effect. This is a general finding in the Stroop task and its equivalents (Dunbar & MacLeod, 1984). The model faithfully reproduces this effect. Although the details of the interactions that produce this effect are beyond this discussion (see Cohen, Dunbar, & McClelland, 1990), an important factor is the nonlinearity of the activation function (see Figure 4). This imposes a ceiling on the activation of the correct response unit, which leaves less room for an excitatory response to congruent information than for an inhibitory response to conflicting information coming from the competing pathway.

This unanticipated consequence of the use of a saturating activation function is noteworthy in that it shows that the asymmetry may be accounted for without assuming that facilitation and interference arise from distinct processing mechanisms, as proposed by some authors (e.g., Glaser & Glaser, 1982; MacLeod & Dunbar, 1988). Although it remains possible that separate mechanisms are involved, the model we present demonstrates that this is not necessarily the case. The failure of previous theories to account for this asymmetry in terms of a single mechanism may well be due to their reliance, either explicitly or implicitly, on strictly linear processing mechanisms.

Stimulus onset asynchrony effects. Another anomaly that has confronted Stroop theorists concerns the finding that stimulus onset asynchrony (SOA) has little impact on the Stroop effect. Thus, even when the color is presented well before the word (400 ms), it still fails to produce interference with word reading (Glaser & Glaser, 1982). Automaticity theory can explain this finding (color naming is controlled, therefore it can be inhibited), but no process model has succeeded in reproducing this effect. Furthermore, as we have seen (and will return to shortly), there are problems in assuming that color naming is truly a controlled process. Our model addresses this phenomenon by demonstrating that interference effects depend directly on the strength of processing, and not on the relative finishing times of the two tasks. When attention is withdrawn from the weaker pathway, it produces less activation at the output level than does the stronger pathway under the same conditions. As a result, weaker pathways produce less interference, independent of finishing time.

Relative nature of automaticity. As mentioned above, MacLeod and Dunbar (1988) showed that the pattern of interference effects associated with a particular task can depend heavily on the context in which it is performed. Thus, when compared with a novel task, such as shape naming, color naming may actually appear automatic. The model can account for this finding in terms of the relative strengths of competing pathways. When a new pathway is added—to represent the shape-naming process—and given connection strengths weaker than those in the color pathway, processing in the new pathway is subject to interference (and facilitation) from color information. Color naming, on the other hand, is not influenced by information in the new pathway. Thus, as observed in the MacLeod and Dunbar experiment, color naming reverses roles.

After demonstrating these initial effects, MacLeod and Dunbar went on to train their subjects on the shape-naming task over a period of 20 days. As expected, reaction time decreased according to a power law. At the end of training, when shape naming had become faster than color naming, interference effects had also reversed. We were able to accurately simulate their reaction time findings on a trial-for-trial basis, as well as the different patterns of interference effects on the first and last days of training. Thus, the model not only supports the notion of a continuum of automaticity, but provides an explicit set of information processing mechanisms underlying this continuum. These mechanisms accurately simulate the concurrent changes in reaction time and interference effects that occur with practice and link these changes to the gradual changes in connection strengths that occur with difference-based learning.

Requirements for attention. Attention is implemented as a graded, modulatory influence on processing. This means that information can flow along pathways, even when there is no allocation of attention. This was the case for word information, which was able to influence the color-naming process, even when no attention was allocated to the word pathway. This is consistent with the automatic nature of word reading. However, contrary to traditional views of automaticity, this does not mean that attention has no effect on automatic processes. To the contrary, in the model, attention has a strong influence on "automatic" processes such as word reading: Although some information "leaks" through the unattended channel—influencing the speed of the response—it is not enough to determine the actual content of the response. Only with the allocation of attention can a process, even if it is a strong one, be carried to completion. The degree to which a process relies on attention is determined by the strength of the underlying pathway (i.e., the connections in that pathway). This is shown in Figure 5A, which compares the word-reading and color-naming processes under varying degrees of attentional allocation. The model shows that requirements for attention are also influenced by the strength of a competing process. This is shown in Figure 5B, which compares the attentional requirements of the color-naming process when it faces competition from processes of varying strength. We should be clear that the ideas that attention is a modulatory process and that even automatic processes may rely on attention are not new ones. For example, Treisman (1960) proposed an attenuation theory which claimed that messages outside of the focus

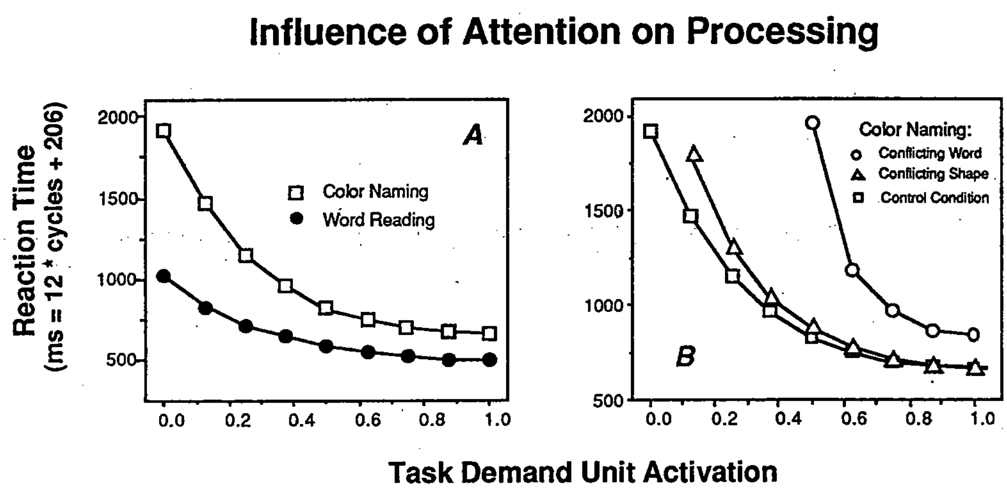

Figure 5. Influence of attention on processing. Panel A shows differences in the requirements for attention between color naming and word reading, and the effect on these two processes of reducing activation of the task demand unit. Panel B shows the different requirements for attention of the color-naming process when it must compete with a stronger process (word reading) and a weaker one (shape naming, early in training). *Note.* Figure after Cohen et al., 1990. Copyright 1990 by the American Psychological Association. Reprinted by permission.

of attention were not completely shut out; rather, the flow of information was simply "attenuated" on the unattended channel. Furthermore, Kahneman and Treisman (1984) argued strongly against what they termed the "strong automaticity claim": that automatic processes have no requirements for attention. The model helps support these claims by committing them to a specific set of information-processing mechanisms that can account for the empirical data and that help extend these ideas to encompass related, but previously unintegrated phenomena (e.g., SOA and practice effects).

Interactivity and Competition in Attention

Thus far we have seen how the Stroop model was able to capture several phenomena associated with automaticity and attention as they emerge in the Stroop task, in terms of several of the principles enumerated at the outset. However, the Stroop model did not incorporate the principles of interactivity and competition; processing was strictly feed-forward, whereas interactivity and competition are inherently bidirectional processes. In previous work these principles have been exploited in models of context effects in perception (McClelland, 1992; McClelland & Rumelhart, 1981). We argue here that they should also be incorporated into thinking about automaticity and attention.

A primary reason for incorporating interactivity is that there appear to be bidirectional influences between stimulus processing and attention. One particularly interesting example of this comes from an intriguing experiment by Brunn and Farah (1991). They examined the effects that familiar stimuli (words) can have on the allocation of attention in patients with hemilateral neglect. Such patients have right-sided parietal lesions, and tend to neglect stimuli appearing in the left half of space. To quantify this effect, Brunn and Farah asked such a patient to indicate the midpoint of a horizontal line. She marked the line well to the right of center, indicating neglect of its left end. Brunn and Farah next showed the patient a horizontal line beneath a string of letters, as shown in Figure 6. When the string of letters formed a random sequence, the patient bisected the line as before. But when the string of letters formed a word, the patient bisected the line much closer to its true midpoint. The study strongly suggests that the word stimulus elicits attention to the entire spatial region occupied by the word, thereby inducing the subject to notice the part of the line that might otherwise have been neglected.

On any account in which attention was a top-down process (as it is

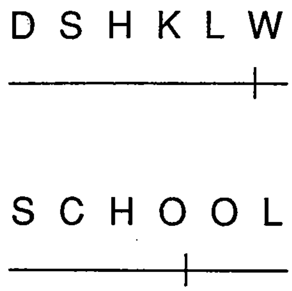

Figure 6. Stimuli used in the Brunn and Farah (in press) experiment, with the kind of responses generated by their subject.

in our Stroop model), this finding must seem perplexing. If attention is a top-down process, then why does the nature of the stimulus influence it? At the same time, the finding seems perplexing on any model in which perception is strictly bottom up (as in our Stroop model again). For if perception is bottom up, then surely the left-hand letters of the word suffer as much from neglect as the left-hand letters of a random string. Why then can their presence lead to a reorientation of perception?

A straightforward account for these findings can be offered in terms of a model that incorporates interactivity—both in perception, as in the interactive activation model, and in attention, as recently proposed by Phaf, Van der Heijden, and Hudson (1990). The idea is sketched in Figure 7. Three modules are shown, one representing position-specific feature patterns (the letters in the string), one representing familiar objects (words), and one representing the focus of spatial attention (locations). We assume that in patients with neglect, spatial attention is ordinarily biased to the right; therefore, there is more activation for right-sided locations in the attention module. When a random letter string is presented, this bias leads to stronger activation of the letters on the right, and because no familiar object is activated, that is the end of the matter. But when a word is shown, the letters on the right, plus weak activations from the letters on the left, lead to the activation of a representation for the whole word in the familiar objects module. This in turn feeds activation back to the position-specific feature level, strengthening the activations of the letters in the ordinarily neglected field. These strengthened feature-level activations then lead to a strengthening in the activation of the location representations associated with the ordinarily neglected field. As a result attention itself is allocated more evenly across the field.

Thus far we have considered evidence for the interactivity as-

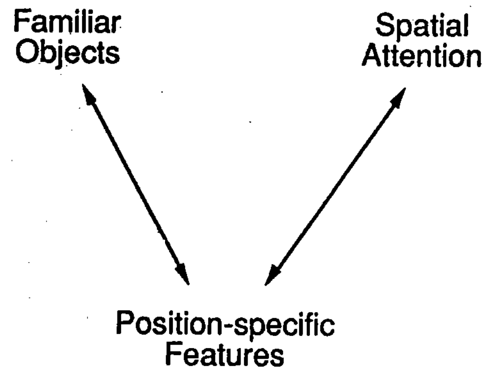

Figure 7. Interactivity in perception, localization, identification, and attention.

sumption. What about the assumption that there is competition among alternatives at each level? This assumption was originally introduced into the framework to capture the winner-take-all character of perceptual processes, in which the alternative that best satisfies the combined constraints imposed by bottom-up and top-down influences becomes most active and suppresses all competitors. However, there are several reasons to feel that competition may play a strong role not only in perception but in attention. For one thing, the presentation of one stimulus tends to divert attention away from other stimuli. This kind of attention diversion is particularly apparent in visual search after extended practice with a constant target set (e.g., Schneider & Shiffrin, 1977). The practiced targets come to command attention, as is easily shown after changing to a new target set. Now the practiced targets are distractor items. When such items appear in the display, they appear to prevent the subject from noticing members of the new target set. Thus, when attention is attracted to one item it appears simultaneously to be withdrawn from others.

That an attention-demanding stimulus diverts attention from other targets is naturally captured in terms of competitive or mutually inhibitory interactions between units representing alternative loci of attention. In the model shown above, the presentation of the distracting stimulus tends to activate the attention units for the location containing the distracting stimuli. These in turn inhibit attention to other loci. Competition, if it is present at the position-specific feature level as well, would tend to have a direct suppressive effect at that level too. Competition at either the perceptual level or attentional level, or both, could be the reason that target detection is generally faster and more accurate when the target is presented alone rather than in the presence of other stimuli.

We have only begun to explore the roles of interactivity and competition in our simulations of attentional phenomena. Below we report one simulation that we have conducted recently (Servan-Schreiber, 1990). The simulation does not consider learning, but otherwise incorporates all the principles enumerated at the beginning of this article. These are used to address, in detail, an interesting pattern in the time course of processing seen in some recent experiments using a response-competition task originally described by Eriksen & Eriksen (1974). These data cannot be accounted for by the feed-forward Stroop model described above, but can easily be captured by a model that adds the principles of interactivity and competition.

The Eriksen task

The Eriksen task has been studied extensively in behavioral as well as psychophysiological experiments and is particularly well suited to the detailed study of attentional effects in choice reaction-time situations (Coles & Gratton, 1986; Coles, Gratton, Bashore, Eriksen, & Donchin, 1985; Gratton, Coles, Sirevaag, Eriksen, & Donchin, 1988). In this task, subjects are asked to respond with a different hand to two different target letters (*S* or *H*) that appear in the middle of a three- or five-letter stimulus array. In the compatible condition, all letters are identical (i.e., *HHHHH* or *SSSSS*), whereas in the incompatible condition the central letter is different from the surrounding letters (i.e., *HHSHH* or *SSHSS*). All stimuli have an equal probability of being presented. As in the Stroop task, subjects are slower and make more errors in the incompatible condition.

In psychophysiological studies, responses are recorded when subjects squeeze a dynamometer to 25% of maximal force. Because electromyographic activity in both arms is also recorded, information about activity in either response channel is available even when it is not associated with an overt response. In recent studies, Gratton et al. (1988) have also used recordings of event-related potentials over the motor cortex to provide information about covert response preparation in the absence of overt muscular activity.

The overt performance of subjects on this task, together with electroencephalogram (EEG) and electromyogram (EMG) recordings, sheds light on the coupling between sensory processing of the stimuli and response selection over time. We will start by reviewing the empirical observations that have helped constrain the development of our model.

Graded and continuous evaluation processes. EEG recordings have been used to argue that responses can occur before stimulus evaluation is complete. This conclusion is based on P300 recordings showing that reaction times can be shorter than P300 latency. Also, EEG and

EMG recordings have shown that both the correct and incorrect responses can be activated on the same trial. This suggests a continuous, parallel flow of information from stimulus analysis to response selection, rather than a single stage for stimulus evaluation followed by a response selection process (Coles & Gratton, 1986; Gratton et al., 1988).

Competition between response channels. The delay between EMG activity and squeeze response in the correct channel can be plotted as a function of EMG activity in the incorrect channel. There is a systematic positive relationship between the two: When EMG activity in the incorrect channel is greater, correct responses are delayed. This observation provides clear evidence for a competitive interaction between the two response processes.

Delayed effect of attention. When accuracy of responses is plotted against reaction time, the shape of this time-accuracy curve is not the same for the compatible and incompatible conditions. In the compatible condition, accuracy starts at 50% (random response) for very short reaction times and rises monotonically to an asymptote close to 100% correct. However, in the incompatible condition—which requires selective attention to the central letter of the stimulus array—performance is at chance initially but then drops significantly below chance level before it rises to asymptote (see Figure 9A). This "dip" in the time-accuracy curve suggests that, at very short latencies, unattended but salient stimuli (i.e., the flankers) tend to influence response processes. It is as if the mediation of task-appropriate responses through spatial attention required additional processing time.

Fixed response criterion. The covert activity in the motor cortex area that engenders overt muscular responses can be evaluated using the contingent negative variation (CNV) wave of an EEG recording. This CNV activity is lateralized to the cortical area contralateral to the overt response. The magnitude of the CNV is related to motor preparation for the overt response. It is possible to measure the difference between the two CNV waves on each side and to follow this difference over time from the warning stimulus to after the response execution. This measure provides an indication of relative response activation. The data show that regardless of condition or speed of response, there appears to be a fixed degree of asymmetry which, when exceeded, leads to an overt response. This result suggests that subjects use a fixed response criterion at all reaction times and in all conditions. In turn, this suggests that the variability in reaction times and the shape of the speed-accuracy curve is not due to a variable threshold but rather to the interplay between random activity in the system (noise) and a process of progressive accumulation of evidence about the target.

A model of the Eriksen task

Architecture. The network is composed of three modules, with inhibitory connections among the units within each module, and excitatory connections between units of different modules (see Figure 8). The input module contains six units, one for each letter (*H* or *S*) in each of three positions (*left, center, and right*). The output module contains only two units, one for each response (*H* or *S*). Input units have excitatory connections to the corresponding output units (e.g., all *H* input units are connected to the *H* response unit). Finally, a third module—the attention module—contains three units that each represent one of the three input positions. Each of these units has bidirectional excitatory connections to the two input units coding for *H* or *S* in the corresponding position (e.g., the *center* attention unit is connected to both *H* and *S* in the central position). When one of these attention units is activated, the network can selectively enhance the activation of input letters in the corresponding location. The positive connection weights between the units in the different modules

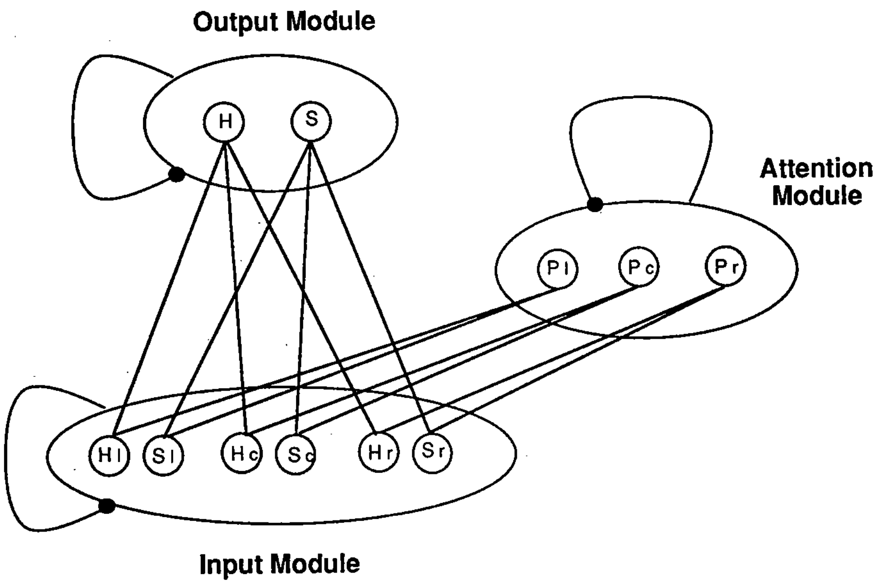

Figure 8. Schematic representation of the network used to simulate the data of Gratton et al. (1988). The subscripts *l*, *c*, and *r* refer to left, center, and right, respectively. Between-module connections are excitatory only. In addition, connections between the input and attention modules are bidirectional. Within each module, each unit inhibits every other unit.

and the negative weights between the units within each module were all set by hand such that, in the absence of noise, the system would reach a stable state in which the correct output unit is active and the other output units are inhibited (both in the compatible and incompatible conditions).

Intrinsic noise. As in the Stroop model, variability in the system's performance relies on an independently sampled Gaussian noise term added to the net input of each unit at each cycle of processing.

Simulation of a trial. Simulations begin with several preparatory processing cycles, before the presentation of the stimulus input. Because the task requires the subjects to identify only the central letter, excitatory input is provided to the center attention unit at the beginning of the preparatory period and is left on throughout the trial. This in turn primes the position-specific letter units for the central position. Because of the noise in the system, the activations of all of the units tend to vary randomly during the preparatory interval. On occasion, the response threshold can actually be reached during this preparatory interval. If so, the response is classified as premature, and the trial is aborted (human subjects also make such responses). The stimulus array is presented as a fixed input into the input units. This remains constant until a response is recorded. A failure to respond after 100 cycles is recorded as an omission.

Response mechanism. A response is recorded when the activation of one of the two output units reaches a fixed threshold.

Parameter selection. In addition to the basic architectural assumptions, the model has a number of free parameters. These include the values of excitatory and inhibitory weights for each module, the amount of net input provided to input units and to the attention unit, the number of cycles preceding the beginning of a trial, and the value of the response threshold. However, the richness of the data greatly constrains the selection of parameters in the model. We attempted to fit simultaneously the mean reaction time for each condition (compatible and incompatible), the average accuracy of each condition, the number of premature responses (less than 1%), the number of omissions (less than 1%), the proportion of responses in each of seven reaction time bins for each condition, and the accuracy for each of these seven reaction time bins in each condition.

The results of the simulation are summarized in Figure 9. Following the method of Gratton et al. (1988), we divided the trials of the simulation into seven reaction time bins (on the basis of the number of cycles). A simple linear regression was used to establish the correspondence between number of cycles in the simulation and the reaction time in the empirical data.

Figure 9 shows that the simulation captured all of the important aspects of the data: (a) the monotonic approach to asymptote of the accuracy curve in the compatible condition; (b) the dip in the accuracy curve in the incompatible condition; (c) the overall shape of the reaction time distribution; and (d) the greater number of responses in the later bins in the incompatible condition than in the compatible condition.

How do each of these four effects arise in the model? First, consider the compatible condition. Initially, the only source of activity in the network is from the random noise associated with the input to each unit. The early part of the reaction time distribution reflects this random activity. However, as time passes, activation provided by the stimulus spreads from the input units to the corresponding response units, causing response accuracy to rise progressively toward asymptote.

In the incompatible condition, external input is provided to two incorrect letter units and only one correct letter unit in the input module. Because of this, the incorrect response unit tends to receive more activation early in processing than the correct response unit. However, the attention unit for the center position is also receiving some initial input. This tends to activate both center letter units in the input module. However, only one of these is receiving external input from the stimulus. The other, though it receives excitatory input from the center attention unit, is inhibited by all of the other active input units and is therefore rapidly inactivated. Ultimately, the mutual excitation between the center attention unit and the center letter unit allows this unit to dominate the other two and, in turn, to activate the correct response unit. It is the delay required for this interaction to take place that accounts for the dip observed in the time-accuracy curve.

Note that the account suggested by the model contrasts with other attempts to explain the dip discovered by Coles et al. (1985). For example, faced with the limitations of box-and-arrow models of information processing, these investigators have had to rely on two subprocesses to explain this phenomenon: an early direct process responsible for providing information on the identity of display elements independent of location, and a second, slower, process that provides identity information tied to particular locations. In our account, location-independent information and location-specific information arise successively from a single system. The interaction between stimulus feature information and the allocation of attention to a particular location is such that processing is dominated early on by the totality of information arising from all locations; only gradually

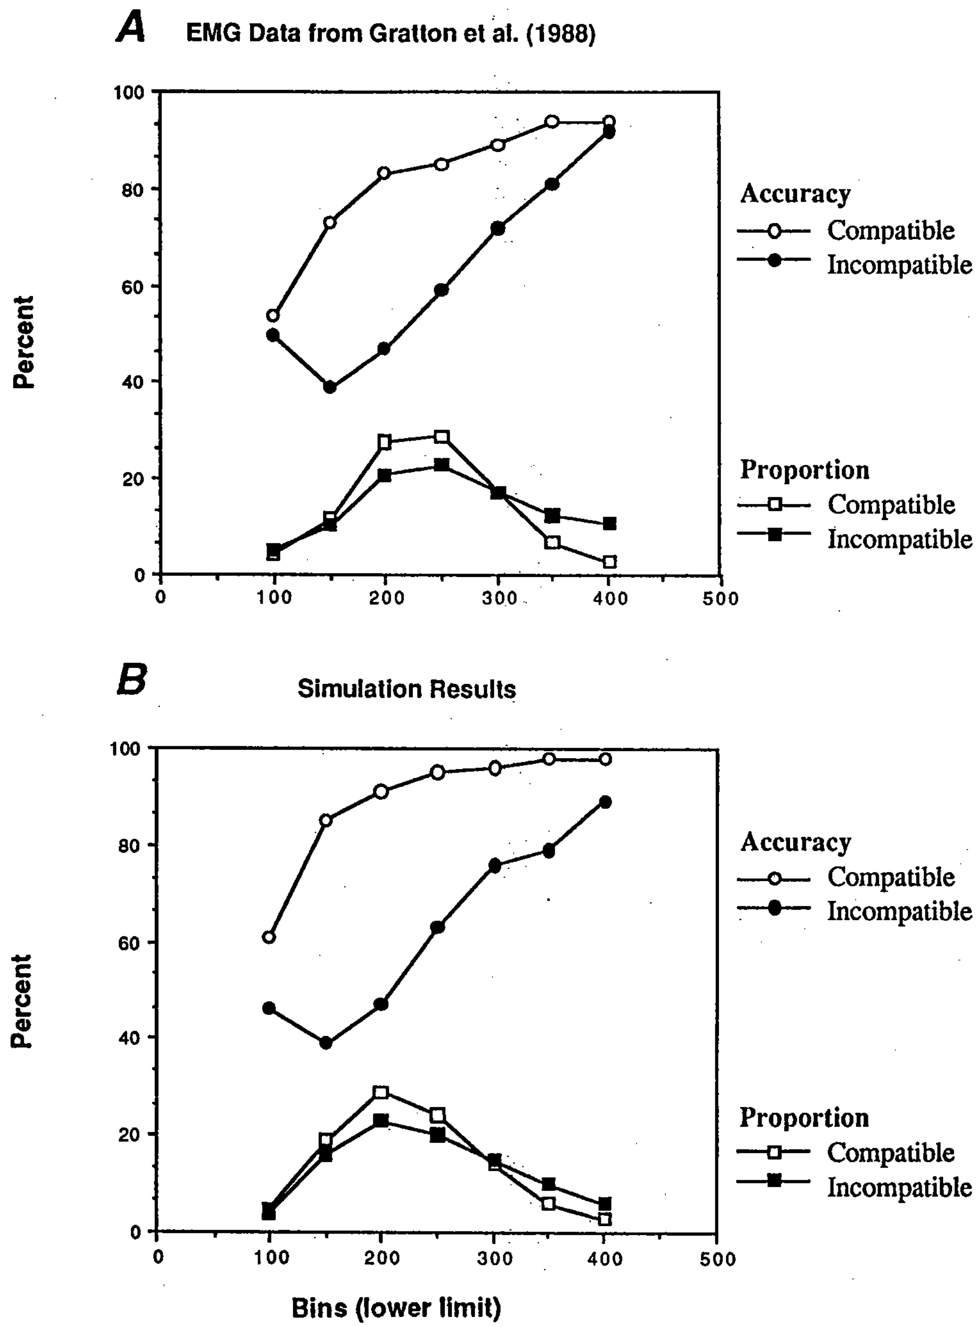

Figure 9. Comparison of the empirical results of Gratton et al. (1988) with the performance of the model. In each case, responses were divided by response time into seven 50-ms bins. Data points in the lower part of each graph are the proportion of responses occurring in each bin. Data points in the upper part of each graph are the proportion of correct responses in each bin. The original empirical data were graciously provided by Gabriele Gratton and Michael Coles.

does the competition mechanism allow the active units in the attended location to suppress the units in unattended locations, thereby permitting a correct response.

The explanation we have given would seem completely ad hoc if it were formulated in terms of unanalyzed boxes and arrows, because it would simply amount to stipulating that there is a single box that produces both position-independent and position-specific identity information, and that the former is produced more slowly than the latter. Without a description of the mechanisms inside the box, the explanation becomes a mere restatement of the data. Yet with a set of principles guiding our conception of the mechanisms assumed to underlie processing within each module, it is possible to see in fact that a single module can produce just such a pattern at its outputs.

An additional comment is required here concerning the response mechanism used in the model. Gratton et al. (1988) suggested that subjects emitted an overt response when the difference between the CNV waves over each motor area reached a fixed threshold. In the model, we did not compare the difference between the activation of the two response units to a threshold. We simply compared the activation of the most active unit to a fixed threshold. Yet, post hoc analyses showed that when a response unit reached threshold in the model, the difference between the activation of the two response units was consistently the same whether the stimulus array was compatible or incompatible ($M = 0.22$, $SD = 0.055$). This suggests that the response selection mechanism used in the model also results in a strong correlation between response emission and the difference between activation levels of the two response channels. This is because the two response units in the model have reciprocal inhibitory connections. Hence, their activation levels are not independent; the more active one unit becomes, the more inhibited the other becomes.

The overall shape of the reaction time distribution in the model, with the largest number of responses at intermediate bins, arises from the interaction of information about the stimulus and random noise. In the first bins, responses occur only when random noise spuriously accumulates in favor of one of the two responses. In the last bins, responses are delayed because noise spuriously strengthens the incorrect response unit—which inhibits the correct response unit—or directly reduces the net input into the correct unit. Both of these events are comparatively rare because they rely on the noise terms of many different units in the network having the same valence (with respect to the response units) simultaneously (or a single noise term being extremely large, or a large noise term having the same valence for many consecutive cycles, etc.).

Finally, the larger number of responses in the later bins seen in the incompatible condition compared with the compatible condition is due to the influence of the two input units that provide activation to the incorrect response unit. A greater activation of the incorrect response unit results in a direct inhibition of the correct response unit, which delays the latter's approach to threshold.

DISCUSSION

The role of the seven principles

We have presented two models that exhibit many basic aspects of automatic processes and the control of such processes via attention. However, there are differences between the two models: The first is strictly a feed-forward model, and highlights the role of incremental, difference-reducing connection adjustment processes; the second is fully interactive and competitive, exploiting bidirectional excitatory connections between levels and bidirectional inhibitory connections within levels, though it looks at performance without regard to the learning process.

The next step for this research is to unify the two models, capturing all of the aspects of attention and automaticity discussed here within a single model that encompasses all of the principles. One reason we have not yet taken this step is that effective, computationally plausible learning rules for networks with bidirectional connections have only recently become available (Hinton, 1989; Movellan, 1990; Peterson & Anderson, 1987), and we are just now beginning to incorporate them into our work. These algorithms retain the incremental, difference-reducing character of back propagation (without requiring the propagation of error information backward through time to calculate weight changes for networks with recurrent connections, as has been the case for back propagation networks).

All seven of the principles enumerated in the introduction have played a role in our simulations. The first three principles—sigmoidal activation function, gradual propagation of activation, and intrinsic noise—seem to be basic prerequisites to the modeling of performance in information-processing tasks. The next two—incremental, difference-driven connection adjustment and control by modulation—combined with the first three principles give rise to the gradual emergence of automaticity together with the strong but far from absolute control over processing that is exerted by attentional influences. These five principles play a central role in explaining the core phenomena of automaticity that have concerned us here. But the last two principles—competition and interactivity—are also relevant to issues of attention. The principle of competition is in fact partially incorporated in the Stroop model, because in that model it is the difference in activation between the most active response unit and its competitors that is used to trigger a response. This principle is more thoroughly integrated into the model of the conditional accuracy functions in the Eriksen response-competition task, and plays a key role in allowing the correct response to eventually dominate even when, initially, the incorrect response is more strongly activated.

The principle of interactivity, which is incorporated in the simulation of Gratton et al. (1988), may not be crucial in this particular case. The key aspect of this model is the competitive inhibition between alternatives, rather than the presence of bidirectional excitatory connections. The role of interactivity in processing has been argued elsewhere (Dell, 1985; McClelland & Elman, 1986; McClelland & Rumelhart, 1981).

We argued above that interactivity plays a role in attention, though we have not yet had the opportunity to develop simulations of tasks in which this plays a crucial role. Principally, the role of interactivity in attention is to provide a means where attention itself, albeit largely a matter of top-down control, may be partially under the control of stimuli themselves.

The idea that information processing is interactive can lead to a blurring of the traditional distinction between attentional and perceptual mechanisms. The distinction actually disappears in the recent model of Phaf et al. (1990). These authors use a modeling framework very close to the one we describe here to argue that the mechanisms of attention and perception are in fact one and the same. They note that when multiple stimuli are shown, subjects can be instructed to select one to respond to on the basis of any property of the object. In this view, location is just a property like any other (color, shape, etc.). Thus, when shown a blue triangle to the left of a red square, subjects can select the blue object, the triangle, the left object, etc. Phaf et al. use the same mechanism we use here to select by location, but also add mechanisms to select by color, shape, etc., in exactly the same way. Furthermore, the analyzers that are used to select for color, shape, or location top down are the same ones that are used to represent the perception of these items when they are activated bottom up. Thus the interactive attentional model of Phaf et al. actually obliterates the classical distinction between perception and attention and views them as simply different aspects of the function of a single interactive processing system. This fits squarely with the view that emerges from both the Stroop model and our model of the Eriksen paradigm: Attentional information can be treated like information of any other kind, and attentional effects can be attributed to the modulatory influence that one source of information has on any other.

Based on the foregoing, it appears that all of the principles enumerated at the beginning of this article play a role in automaticity and the attentional control of processing. In the remainder of this section, we consider the approach we have taken here in relation to general issues in attention research and in relation to other models of automatic processes and their control by attention. In particular, we focus on two issues that often seem to be at the core of theoretical discussions about automaticity and attention: the distinction between controlled and automatic processes, and the notion of capacity.

Controlled versus automatic processing

The Stroop model strongly suggests that color naming and word reading can be seen as relying on qualitatively similar processes, and that differences in speed and interference effects can be attributed to the relative strengths of the connections underlying each pathway. This view differs from the traditional notion that the Stroop effect demonstrates the properties of two qualitatively different types of processing: controlled (color naming) and automatic (word reading). This does not mean, however, that we reject the idea that qualitatively different kinds of processing exist. Indeed, we assume that very early in training on novel tasks—before connection weights of an appreciable degree have had a chance to develop in the relevant pathway—subjects rely on a different set of mechanisms than they eventually come to rely on with practice. In Cohen et al. (1990), we called this type of processing "indirect," to capture the fact that it may be mediated by explicit consideration of verbal instructions or verbally mediated associations, and to distinguish it from "direct" processing, in which no such mediation is involved (i.e., there is a "direct" pathway from stimulus to response, such as those in the Stroop model for color naming and word reading). Furthermore, we assume that the type of processing underlying early performance shares many of the attributes traditionally associated with controlled, or strategic processing: It is slow, highly susceptible to interference from distracting tasks, and relies heavily on attention. The important point, however, is that these attributes can continue to be exhibited by tasks even after they have received extensive practice, when they are placed in competition with other, even more highly practiced tasks. The terms direct and indirect map only partially, then, onto the classical distinction between controlled and automatic; they are meant rather to convey the kind of processing that we believe underlies each type. The correspondence between these terms and the traditional ones is shown in Figure 10.

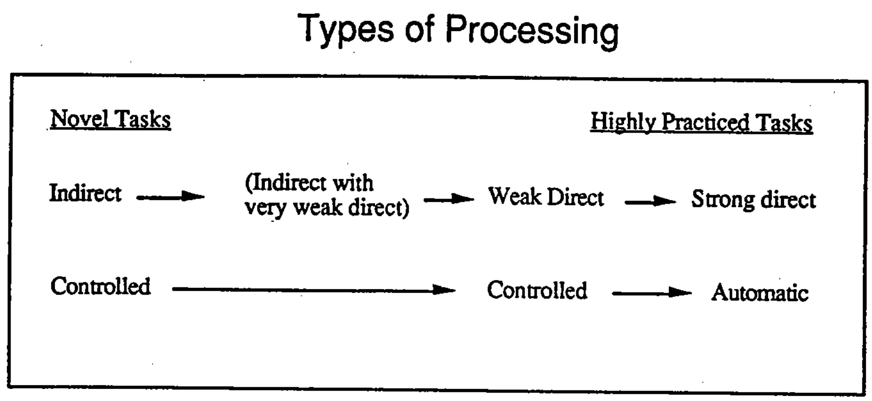

Figure 10. Relationship of the proposed distinction between direct versus indirect processing and the traditional distinction between controlled versus automatic processing. *Note.* Figure after Cohen et al., 1990. Copyright 1990 by the American Psychological Association. Reprinted by permission.

This distinction between direct versus indirect processing provides an appropriate context in which to consider the relationship between our approach and an approach based on production systems, such as the ACT* model described in this issue by Anderson (1992). In our view, each approach has its own natural domain of application. As evidenced by the success of the Stroop model, the PDP approach seems naturally suited to capturing the progressive changes that occur with extensive practice and that lead to increases in automaticity. The strengthening process used to account for these changes emerges naturally from a system in which processing is connection based, and learning involves the gradual adjustment of these connections. Productions, on the other hand, are inherently discrete in nature—one either has a production for something or one does not. Although strengthening mechanisms can be tacked onto such models, they are not really an intrinsic feature of the approach.

Indirect processing, however, presents a somewhat different perspective. This may well be describable in terms of procedural rules that can be flexibly sequenced to accomplish novel tasks. The ability to model performance in higher level cognitive tasks in terms of a composition of separate, rulelike parts is a primary motivation for the use of production systems in psychological models (Anderson, personal communication). The fundamental insight underlying the production system approach is that many skills, especially ones that are unfamiliar or that are complex and high level (e.g., multidigit arithmetic), can be decomposed into a set of simple, discrete strategies or rules, and that these can be conveniently and effectively represented as productions.

In our view, it makes good sense to characterize direct processes in terms of connectionist mechanisms and to characterize indirect processes in terms of productions. We do believe that, ultimately, *all* processing relies on connection-based pathways, and that high-level skills should be representable within the PDP framework. The point we make here is only that for some purposes, a higher level characterization, somewhat removed from the underlying processing mechanisms, may capture the essential features of some processes in a succinct way. In other cases, a finer grain of analysis may provide the more natural and perspicuous account.

Attention and capacity

The approach we have taken to automaticity also sheds light on the notion of attentional capacity. There are two prevailing views concerning this issue that have been described in the literature and that, on the surface, would seem to be in conflict. The traditional view holds that "controlled" processing relies on a central, limited-capacity attentional mechanism, whereas automatic processes are independent of this mechanism and compete only when they lead to conflicting responses (e.g., Posner & Snyder, 1975; Shiffrin & Schneider, 1977). The problem with this view is that *all* processes (including putatively automatic ones) can be shown to rely to some extent on the allocation of attention (e.g., Kahneman & Chajczyk, 1983; Kahneman & Henik, 1981). In contrast, other authors have proposed a "multiple resources" view (e.g., Allport, 1982; Hirst & Kalmar, 1987; Logan, 1985; Navon & Gopher, 1979; Wickens, 1984), which postulates that all processes require resources of some kind, but that these are "local" and that there may be many different types. According to this view, competition (and interference) arises when two tasks place simultaneous demands on the same set of resources. The problem with this view is that neither the nature of attention nor the nature of the resources postulated are specified. Our approach offers a reconciliation of these two perspectives, and can address the problems that confront each.

The models we have presented show how attention can be seen to modulate processes that by traditional criteria would be considered to be automatic (e.g., word reading in the Stroop task). At the same time, they show how the requirements for attention can vary both *among* processes (color naming vs. word reading; see Figure 5A) and, for a given process, depending upon the context in which it occurs (color naming with a conflicting word vs. a conflicting shape; see Figure 5B). However, attention is not given a unique status within our framework. Rather attentional information is represented and processed like information of any other type: as a pattern of activation over a set of units in a module. In this respect, the processing associated with attention is governed by the same principles and constraints that govern all other types of processing. One of these constraints is the competition that can arise when two different sources of information compete for representation within a module. If information arriving from different pathways generates disparate patterns of activation within a module, then the two processes will compete for representation within that module. Thus, the processing capacity of that module can be thought of as being limited: It cannot support the full processing of both signals at once. This property of the system can account for the limited capacity of attention in the traditional view, and for the notion of competition for resources in the multiple resources view.

When a given module plays an attentional (i.e., modulatory) role for some other set of processes (such as the task demand module does for color naming and word reading in the Stroop model), we are led to a perspective that is very similar to the traditional one. That is, competing representations within an attentional module will manifest as a limitation in attentional capacity: Both representations will be degraded. Of course of the processes that rely on these attentional representations, stronger ones will be less influenced by this degradation than weaker ones (see Figure 5), consistent with the view that the more automatic a process is, the less it will rely on attention. However, our approach differs from the traditional approach in that it allows there to be more than one attentional mechanism (module) within the system, and that different processes may rely on different such modules. The extent to which limitations in attentional capacity will affect performance will depend on the particular processes involved in the task (or set of tasks), the extent to which these processes rely on attentional resources, and whether the attentional resources are the same or different for the various processes involved.

The perspective shifts when we focus on modules that are directly involved in a processing pathway; that is, modules which lie in the pathway along with information flows from input to output. Such modules may be involved in one or more pathways (e.g., the response module in the Stroop model), and there may be many such points of intersection between pathways. When disparate information arrives from different sources within such modules, interference occurs. This seems to capture the main thrust of the multiple resources view: Tasks will interfere to the extent that they compete for local resources. However, the principles underlying our approach allow us to go beyond the multiple resources view, by specifying the exact nature of these resources and their limitations: Resources are sets of units whose activations are used to represent information; their capacity is limited by the competition for activation that is assumed to exist between units within a module. These principles allow us to capture the type of interference phenomena that arise when information from two sources converge on a common module.

CONCLUSION

At the outset of this article we enumerated seven principles of information processing that constrain the more general PDP framework. We then showed how these principles can be used to account for a number of the major phenomena associated with automaticity: gradual development with practice; concomitant improvements in speed (and a reduction of variance) that follow a power function; reduced reliance on, but not complete autonomy from, the effects of attention; the relative nature of interference effects; and the interacting influences of stimulus information and attentional allocation on responding. We presented two computational models to demonstrate the ability of the principles to account for empirical data concerning these phenomena. It is important to emphasize, however, that it is not the details of these implementations that we consider to be important (e.g., the specific activation function used, or the shape of the noise distribution), but rather the principles upon which they are based (a sigmoid activation function, and variability of processing). Indeed, we believe that these principles can be used to account for a wide variety of findings in the psychological literature (see McClelland, 1992) that go beyond the phenomena of automaticity discussed in this article.

Notes

Correspondence concerning this article should be addressed to Jonathan D. Cohen, Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213. Received for publication May 14, 1991; revision received August 2, 1991.

References

Allport, D. A. (1980). Attention and performance. In G. I. Claxton (Ed.), *Cognitive psychology: New directions* (pp. 112–153). London: Routledge & Kegan Paul.

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R. (1992). Automaticity and the ACT* theory. *American Journal of Psychology, 105,* 165–180.

Brunn, J. L., & Farah, M. J. (1991). The relation between spatial attention and reading: Evidence from the neglect syndrome. *Cognitive Neuropsychology, 8,* 59–75.

Cohen, J. D., Dunbar K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing model of the Stroop effect. *Psychological Review, 97,* 332–361.

Coles, M. G. H., & Gratton, G. (1986). Cognitive psychophysiology and the study of states and processes. In G. R. J. Hockey, A. W. K. Gaillard, & M. G. H. Coles (Eds.), *Energetics and human information processing* (pp. 409–424). Boston, MA: Kluwer Academic.

Coles, M. G. H., Gratton, G., Bashore, T. R., Eriksen, C. W., & Donchin, E. (1985). A psychophysiological investigation of the continuous flow model of human information processing. *Journal of Experimental Psychology: Human Perception and Performance, 11,* 529–553.

Dell, G. S. (1985). Positive feedback in hierarchical connectionist models: Applications to language production. *Cognitive Science, 9,* 3–23.

Dunbar, K., & MacLeod, C. M. (1984). A horse race of a different color: Stroop interference patterns with transformed words. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 622–639.

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of target letter in a non-search task. *Perception & Psychophysics, 16,* 143–149.

Glaser, M. O., & Glaser, W. R. (1982). Time course analysis of the Stroop phenomenon. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 875–894.

Gratton, G., Coles, M. H., Sirevaag, E. J., Eriksen, C. W., & Donchin, E. (1988). Pre- and poststimulus activation of response channels: A psychophysiological analysis. *Journal of Experimental Psychology: Human Perception and Performance, 14,* 331–344.

Hinton, G. E. (1989). Deterministic Boltzmann learning performs steepest descent in weight-space. *Neural Computation, 1,* 143–150.

Hirst, W., & Kalmar, D. (1987). Characterizing attentional resources. *Journal of Experimental Psychology: General, 116,* 68–81.

Kahneman, D., & Chajczyk, D. (1983). Tests of the automaticity of reading: Dilution of Stroop effects by color-irrelevant stimuli. *Journal of Experimental Psychology: Human Perception and Performance, 9,* 497–509.

Kahneman, D., & Henik, A. (1981). Perceptual organization and attention. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization* (pp. 181–211). Hillsdale, NJ: Erlbaum.

Kahneman, D., & Treisman, A. (1984). Changing views of attention and automaticity. In R. Parasuraman, R. Davies, & J. Beatty (Eds.), *Varieties of attention* (pp. 29–61). New York: Academic Press.

Logan, G. D. (1980). Attention and automaticity in Stroop and priming tasks: Theory and data. *Cognitive Psychology, 12,* 523–553.

Logan, G. D. (1985). Skill and automaticity: Relations, implications, and future directions. *Canadian Journal of Psychology, 39,* 367–386.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review, 95,* 492–527.

MacLeod, C. M., & Dunbar, K. (1988). Training and Stroop-like interference: Evidence for a continuum of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 126–135.

McClelland, J. L. (1992). Toward a theory of information processing in graded, random, interactive networks. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence and cognitive neuroscience—A Silver Jubilee Volume.* Cambridge, MA: MIT Press.

McClelland, J. L., & Elman, J. L. (1986). Interactive processes in speech perception: The TRACE model. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 58–121). Cambridge, MA: MIT Press.

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.

Movellan, J. R. (1990). Contrastive Hebbian learning in the continuous Hopfield model. In D. Touretzky, J. Elman, T. Sejnowski, & G. Hinton (Eds.), *Connectionist models: Proceedings of the 1990 summer school* (pp. 10–18). San Mateo, CA: Morgan Kaufmann.

Navon, D., & Gopher, D. (1979). On the economy of the human processing system. *Psychological Review, 86,* 214–255.

Peterson, C., & Anderson, J. A. (1987). A mean field theory learning algorithm for neural networks. *Complex Systems, 1,* 995–1019.

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

Posner, M. I., & Snyder, C. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition* (pp. 55–85). Hillsdale, NJ: Erlbaum.

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In D. E. Rumelhart, J. L. McClelland, and the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 45–76). Cambridge, MA: MIT Press.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 318–362). Cambridge, MA: MIT Press.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review, 84,* 1–66.

Servan-Schreiber, D. (1990). *From physiology to behavior: Computational models* *of catecholamine modulation of information processing* (Doctoral dissertation, Tech. Rep. cmu-cs-90-167). Pittsburgh, PA: Carnegie Mellon University, School of Computer Science.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84,* 127–190.

Treisman, A. M. (1960). Contextual cues in selective listening. *Quarterly Journal of Experimental Psychology, 12,* 242–248.

Wickens, D. D. (1984). Processing resources in attention. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of attention* (pp. 63–102). Orlando, FL: Academic Press.