Psychological Review
1992, Vol. 99, No. 1, 45-77

Copyright 1992 by the American Psychological Association, Inc.
0033-295X/92/$3.00

Context, Cortex, and Dopamine: A Connectionist Approach to Behavior and Biology in Schizophrenia

Jonathan D. Cohen
Carnegie Mellon University, University of Pittsburgh, and Stanford University

David Servan-Schreiber
Carnegie Mellon University and University of Pittsburgh

Connectionist models are used to explore the relationship between cognitive deficits and biological abnormalities in schizophrenia. Schizophrenic deficits in tasks that tap attention and language processing are reviewed, as are biological disturbances involving prefrontal cortex and the mesocortical dopamine system. Three computer models are then presented that simulate normal and schizophrenic performance in the Stroop task, the continuous performance test, and a lexical disambiguation task. They demonstrate that a disturbance in the internal representation of contextual information can provide a common explanation for schizophrenic deficits in several attention- and language-related tasks. The models also show that these behavioral deficits may arise from a disturbance in a model parameter (gain) corresponding to the neuromodulatory effects of dopamine, in a model component corresponding to the function of prefrontal cortex.

Schizophrenia is marked by a wide variety of behavioral deficits, including disturbances of attention, language processing, and problem solving. At the same time, findings of biological abnormalities in schizophrenia continue to accumulate, including disturbances in specific neurotransmitter systems (e.g., dopamine and norepinephrine) and anatomic structures (e.g., prefrontal cortex and hippocampus). For the most part, however, the behavior and biology of schizophrenia have remained separate fields of inquiry. Despite a modern consensus that information-processing deficits in schizophrenia are the result of underlying biological abnormalities, few efforts have been made to specify exactly how these phenomena relate to one another.

In this article we address this issue by drawing on a recent development within cognitive science: the parallel distributed processing, or *connectionist,* framework. This framework provides a means for building computer simulation models of performance in specific behavioral tasks. However, connectionist models differ from other computer simulation models of behavior in their use of information-processing mechanisms that incorporate important features of biological computation. Using this framework, it is possible to develop models that explore the effects of biologically relevant variables on behavior. In this article we explore the ability of such models to explain aspects of schizophrenic behavior in terms of specific underlying biological disturbances.

At the behavioral level, we focus on schizophrenic disturbances of selective attention and language. We describe a set of connectionist models that simulate both normal and schizophrenic patterns of performance in three experimental tasks: two that tap attentional performance (the Stroop task and the continuous performance test) and one that measures language-processing abilities (a lexical disambiguation task). The models make use of a common set of information-processing mechanisms and show how a number of seemingly disparate observations about schizophrenic behavior can all be related to a single functional deficit: a disturbance in the internal representation of context.

Furthermore, the models suggest that this functional deficit may be explained by a specific biological disturbance: a reduction in the effects of dopamine in prefrontal cortex. First, we show how a particular parameter of the models can be used to simulate the neuromodulatory effects of dopamine at the cellular level. We then describe the results of disturbing this parameter within a module corresponding to the function of prefrontal cortex in each of the three behavioral simulations. In each case,

---

The order of authorship is alphabetical; both authors made equal contributions to the research and preparation of this article. Preliminary versions of this work were presented at the 142nd Annual Meeting of the American Psychiatric Association in May 1989 and at the Annual Meeting of the Cognitive Science Society in July 1990.

This work was supported by a National Institute of Mental Health (NIMH) Physician Scientist Award (MH00673) to Jonathan D. Cohen and an NIMH Individual Fellow Award (MH09696) to David Servan-Schreiber. Part of this work was also supported by a research grant from the Scottish Rite Schizophrenia Research Program, N.M.J., U.S.A., to Jonathan D. Cohen and a research fellowship from Western Psychiatric Institute and Clinic to David Servan-Schreiber.

We gratefully acknowledge the helpful suggestions made by David Galin, Steve Matthysse, James McClelland, Benoit Mulsant, and David Plaut during the development of our ideas and the preparation of this article. In addition, John Csernansky, Len Horowitz, Roy King, Tracy Kristoffersen, David Spiegel, Elisabeth Targ, and Sue Theimann participated in the design and implementation of the experimental study of schizophrenic language performance that we reported. We also thank our three reviewers for their careful consideration of our manuscript and their many useful suggestions.

Correspondence concerning this article should be addressed to Jonathan D. Cohen, Department of Psychology, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, or to David Servan-Schreiber, Western Psychiatric Institute and Clinic, 3811 O'Hara Street, Pittsburgh, Pennsylvania 15213. Electronic mail may be sent to jc5e@andrew.cmu.edu or DSS@cs.cmu.edu.

this disturbance leads to changes in performance that quantitatively match those observed for schizophrenics in the corresponding tasks. Taken together, these findings suggest that a number of the disturbances of selective attention and language found in schizophrenia may result from a common information-processing deficit (a disturbance in the internal representation of context) which, in turn, may be explained by a single biological abnormality (a reduction of dopaminergic activity in prefrontal cortex).

Before proceeding, it is important to acknowledge that the models we present in this article focus on a specific set of behavioral and biological phenomena associated with schizophrenia. Although we recognize that these represent only a subset of the many disturbances observed in schizophrenia, we suggest that they form a coherent subcomponent of this illness. Our hope is that a more precise account of this subcomponent will help delimit its role in schizophrenia and provide a framework for tackling other phenomena such as hallucinations and delusions.

We begin by reviewing data concerning cognitive and biological deficits in schizophrenia that are relevant to our argument. There are four components to this argument: (a) Schizophrenics' performance in a variety of cognitive tasks indicates a decreased ability to use context for choosing appropriate behavior; (b) frontal cortex is directly involved in, and necessary for the internal representation and maintenance of, context information; (c) schizophrenia is associated with abnormalities of frontal cortex; (d) the normal function of frontal cortex relies on the activity of the mesocortical dopamine system, which also appears to be disturbed in schizophrenia. Following our review of the empirical literature, we present a set of connectionist models that show how the behavioral phenomena can be causally related to specific biological mechanisms.

Background

Cognitive Deficits in Schizophrenia

Disturbances in the Internal Representation of Context

A large number of experiments have revealed schizophrenic deficits in information-processing tasks (for example, see Chapman, 1980). Although these encompass a variety of different processing domains—including selective attention, signal detection, memory, language processing, and problem solving—we believe that many of these may reflect a common underlying deficit: a degradation in the ability to construct and maintain an internal representation of context. By an *internal representation of context* we mean information held in mind in such a form that it can be used to mediate an appropriate behavioral response. This can be a set of task instructions, a specific prior stimulus, or the result of processing a sequence of prior stimuli (e.g., the interpretation resulting from processing a sequence of words in a sentence). By this definition, context information is relevant to but does not form part of the content of the actual response. This distinguishes context information from the kind of information traditionally thought to be stored in short-term memory. We usually think of short-term memory as storing recently presented information, the identity of which must later be retrieved—"declarative" representations in the sense used by Anderson (1983). In contrast, we think of internal representations of context as information stored in a form that allows it to mediate a response to the stimulus other than the simple reporting of its identity. Although it is possible, in principle, that the same representations (and mechanisms) could be involved in both cases, there is evidence that, in fact, the internal representation of context can be dissociated from short-term memory. For example, human infants younger than 6 months show evidence of knowing that an object is hidden behind a cover before they can use that information to reach for the object (Baillargeon, 1990). Similarly, schizophrenics show normal performance on a number of short-term memory tasks (e.g., Larsen & Fromholt, 1976; Oltmanns & Neale, 1975); however, in tasks that rely on the internal representation of context, they consistently show deficits. To illustrate this, we focus on three different tasks: a selective attention task, a signal detection task, and a lexical disambiguation task. In each case, we consider the role that the internal representation of context plays in the task and how schizophrenic deficits in both attention and language performance may be related to a degradation in this internal representation.

Attentional Tasks

Since its definition as an illness, schizophrenia has been associated with deficits of selective attention (Bleuler, 1911; Kraeplin, 1950). Investigators who have focused on the phenomenology of schizophrenia have often reported that patients appear to be highly distractible and unable to screen out irrelevant stimuli from the environment (e.g., Garmezy, 1977; Lang & Buss, 1965; McGhie, 1970; McGhie & Chapman, 1961). A laboratory task that has been used extensively to study selective attention is the Stroop task (Stroop, 1935; for reviews, see Dyer, 1973, and MacLeod, 1991), and several experiments have been conducted using this task with schizophrenics.

The Stroop task. This task taps a fundamental aspect of selective attention: the ability to respond to one set of stimuli even when other, more compelling stimuli are available. The standard paradigm consists of two subtasks. In one, subjects name the color of the ink in which a word is printed. In the other, subjects read the word aloud while ignoring ink color. Three types of stimuli are used: conflict stimuli, in which the word and the ink color are different (e.g., the word RED in green ink); congruent stimuli, in which they are the same (e.g., the word RED in red ink); and control stimuli. The control stimuli for word reading are typically color words printed in black ink; for color naming they are usually a row of XXXXs printed in a particular color. The subjective experience of performing this task is that word reading is much easier, and there is no difficulty in ignoring the color of the ink. In contrast, it is much harder to ignore the word when the task is to name ink color.

These experiences are reflected in the time it takes for subjects to respond to stimuli of each type (see Figure 1). Three basic effects are observed: (a) Word reading is faster than color naming, (b) ink color has no effect on the speed of word reading, and (c) words have a large effect on color naming. For example, subjects are slower to respond to the color red when the word GREEN is written in red ink than when the word RED or

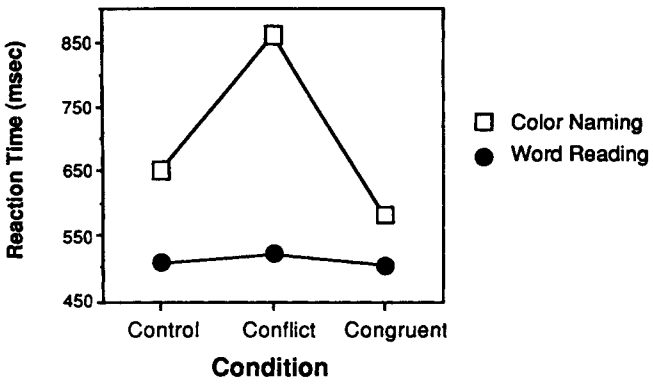

Figure 1. Performance in the standard Stroop task (after Dunbar & MacLeod, 1984). (Data are average reaction times to stimuli in each of the three conditions of the two tasks.)

a series of Xs appear in red ink. Thus, normal subjects have a harder time selectively attending to colors—and ignoring words—than the reverse. This is commonly referred to as the Stroop effect. If schizophrenics suffer from a deficit in selective attention, then they should show a larger Stroop effect; that is, they should be even worse than normal subjects in ignoring word information and should therefore show a greater interference effect.

Table 1 reports data from three empirical studies comparing normal and schizophrenic performance in the Stroop task.[1] Performance of control subjects conformed with the standard findings in this task: Subjects were faster at reading words than at naming colors, and words interfered with color naming. Schizophrenics also showed this pattern of results. However, in all three studies schizophrenics differed significantly from controls in two important ways: (a) Schizophrenics showed an overall slowing of responses, and (b) they showed a statistically disproportionate slowing of responses in the interference condition of the color-naming task. As we noted earlier, a deficit in selective attention would predict this increase in interference. However, because an overall slowing of reaction time is also observed, the significance of this increase in interference must be questioned: This may simply reflect an unanticipated effect of general slowing of performance, rather than the effects of a specific attentional deficit (see Chapman & Chapman, 1978, for a discussion of differential vs. generalized deficits). This issue has not been resolved in the literature. Later we will show how a simulation model of this task can shed new light on this issue by helping to distinguish the effects of a general slowing from those of a specific attentional deficit.

Considerations of the Stroop effect typically focus on the role of selective attention. However, a reliable internal representation of context is also crucial to this task. In order to respond to the appropriate dimension of the stimulus, the subject must hold in mind the task instructions for that trial. These provide the necessary context for interpreting the stimulus and generating the correct response. In Stroop experiments trials are typically blocked by task (e.g., all color naming, or all word reading) so that the proper context is consistent and regularly reinforced. This places only moderate demands on the maintenance of an internal representation of context. However, in other attentional tasks—such as certain variants of the continuous performance test—this is not the case.

The continuous performance test (CPT). This task (Rosvold, Mirsky, Sarason, Bransome, & Beck, 1956) has been used extensively to study attentional deficits in schizophrenics. In the CPT, subjects are asked to detect a target event among a sequence of briefly presented stimuli and to avoid responding to distractor stimuli. The target event may be the appearance of a single stimulus (e.g., "Detect the letter X appearing in a sequence of letters"), or a stimulus appearing in a particular context (e.g., "Respond to X only when it follows A"). The percentages of correctly reported targets (hits) and erroneous responses to distractors (false alarms) are used to compute a measure of the subject's ability to discriminate target from nontarget events (d') independent of response criterion (cf. Green & Swets, 1966). Schizophrenics (and often their biological relatives) show lower hit rates and similar or higher false alarm rates compared to normal subjects and patient controls, indicating poorer signal detection ability (e.g., Cornblatt, Lenzenweger, & Erlenmeyer-Kimling, 1989; Erlenmeyer-Kimling & Cornblatt, 1978; Kornetsky, 1972; Nuechterlein, 1983, 1984; Rutschmann, Cornblatt, & Erlenmeyer-Kimling, 1977; Spohn, Lacoursiere, Thomson, & Coyne, 1977). The fact that schizophrenics show impaired signal detection performance, independent of response criterion, indicates that their poorer performance is not due simply to a lack of motivation (e.g., ignoring the task altogether) or to arbitrary responding (Swets & Sewall, 1963). It is interesting that their deficit is most apparent in versions of the task that make high processing demands: when stimuli are degraded or when information about the previous stimulus is necessary. For example, in one version—the "CPT-double"—a target event consists of the consecutive reoccurrence of any stimulus. As such, the previous stimulus provides the context necessary to decide whether or not the current one is a target; inability to use this context will impair performance. Schizophrenics perform especially poorly in this and similar versions of the task (Cornblatt et al., 1989; Nuechterlein, 1984). Note that here, in contrast to the Stroop task, elements of context change from trial to trial, so that there is additional demand placed on maintaining an internal representation of context.

Other measures of attention. In addition to the Stroop task and the CPT, there are a number of other information-processing paradigms in which schizophrenics exhibit performance deficits that have been related to selective attention, including the span of apprehension task (Neale, 1971), dichotic listening tasks (Spring, 1985; Wielgus & Harvey, 1988), and a variety of

[1] To our knowledge, there are only five studies reported in the literature in which schizophrenics were tested using the standard Stroop task (Abramczyk, Jordan, & Hegel, 1983; Grand, Steingart, Freedman, & Buchwald, 1975; Mirsky et al., 1984; Wapner & Krus, 1960; Wysocki & Sweet, 1985). Only four of these report reaction times, and one involved only 4 subjects (Mirsky et al., 1984). The data for these 4 subjects, although statistically unreliable, conformed to the overall pattern of our predictions: Subjects showed disproportionate amounts of interference. It is interesting that this worsened when they were taken off of medication. The data for the three remaining studies reporting reaction times appear in Table 1.

Table 1
Performance of Normal and Schizophrenic Subjects in Three Studies Using the Stroop Task

Condition	Wapner & Krus (1960)	Abramczyk, Jordan, & Hegel (1983)	Wysocki & Sweet (1985)
	Normal controls		
Word reading	0.39	0.43	0.44
Color naming	0.57	0.60	0.64
Color naming interference	0.98	1.00	1.13
	Schizophrenics		
Word reading	0.57	0.50	0.52
Color naming	0.78	0.77	0.84
Color naming interference	1.51	1.40	1.49

Note. All of these studies used the original form of the Stroop task, in which subjects are given three cards, one with color words written in black ink (word reading), one with color patches or XXXs printed in different colors (color naming), and one with color words each written in a conflicting ink color (color naming interference). Data are presented as the average response time (in seconds) for stimuli of each type.

reaction time tasks (see Nuechterlein, 1977, for a review of the early literature, and Borst & Cohen, 1989, and R. Cohen, Borst, & Rist, 1989, for more recent work). Interpretations of schizophrenic performance in these tasks still frequently refer to Shakow's (1962) original formulation in terms of major and minor sets: Normal subjects are able to adopt a "major set" that takes account of all of the various factors involved in performing a task; schizophrenics are unable to do so, relying instead on a "minor set" that takes account of only a limited set of factors (e.g., the most recent events). Shakow (1962) argued that this is indicative of "the various difficulties created by context [*sic*] . . . It is as if, in the scanning process which takes place before the response to a stimulus is made, the schizophrenic is unable to select out the material relevant for optimal response" (p. 25). As yet, however, there is no generally accepted understanding of the specific information-processing mechanisms that are involved in maintaining an attentional set and that explain their relationship to schizophrenic disturbances in the processing of context.

Schizophrenic Language Deficits

Perhaps the clearest demonstration of deficits in the processing of context can be found in studies of schizophrenic language performance. The classic example of this comes from Chapman, Chapman, and Miller's (1964) study of schizophrenics' interpretation of lexical ambiguities. They found that schizophrenics tended to interpret the strong (dominant) meaning of a homonym used in a sentence even when context provided by the sentence mediated the weaker (subordinate) meaning. For example, given the sentence "The farmer needed a new *pen* for his cattle," schizophrenics interpreted the word *pen* to mean "writing implement" more frequently than did control subjects. They did not differ from control subjects in the number of unrelated meaning responses that were made (e.g., interpreting "pen" to mean "fire truck"), nor did they differ in the number of errors made when the strong meaning of the homonym was correct. These findings have been replicated in a number of studies (e.g., Benjamin & Watt, 1969; Blanley, 1974; J. D. Cohen, Targ, Kristoffersen, & Spiegel, 1988; Strauss, 1975).

Other studies of language performance also indicate that schizophrenics make poor use of context, including studies using cloze analysis (guessing the words deleted from a transcript of speech—e.g., Salzinger, Portnoy, & Feldman, 1964; Salzinger, Portnoy, Pisoni, & Feldman, 1970), speech reconstruction (ordering sentences that have been randomly rearranged—Rutter, 1979), and cohesion analysis (examining the types of references used in speech—e.g., Harvey, 1983; Rochester & Martin, 1979). (For reviews of this literature see Cozolino, 1983; Maher, 1972; Schwartz, 1982.) Whereas a disturbance in the internal representation of context may not account for all of the language and thought disturbances that have been observed in schizophrenics (e.g., idiosyncratic verbalizations, clang associations, or neologisms), it may be directly related to at least one aspect of the clinical presentation of this illness: the concreteness that is characteristic of schizophrenic thought processes. For example, the inability to evoke subtler, but contextually appropriate, meanings of common words may explain overly literal interpretation of proverbs and metaphors.

In language processing, as in attentional tasks, schizophrenics seem to suffer particularly from a restriction in the temporal range over which they are able to use context. Thus, for example, Salzinger et al. (1964, 1970) found that schizophrenics and normal subjects performed comparably well in "clozing" speech (i.e., guessing words deleted from a sample of normal speech) when contextual cues were local (e.g., when the missing word was surrounded by only two or three words). However, when remote cues were possible (e.g., when the missing word was surrounded by larger numbers of words), normal subjects improved in their ability to predict the word, whereas schizophrenics did not. This suggests that normal subjects were able to make use of the additional context provided by more remote cues but that schizophrenics could not. Conversely, Salzinger also showed that it is easier for normal subjects to cloze small segments of schizophrenic speech than larger ones. This implies that broader segments of schizophrenic discourse do not add contextual constraint, presumably because schizophrenics produce contextual references that span more limited segments of speech. On the basis of these data, Salzinger (1971) proposed an immediacy hypothesis which stated that "the behavior of schizophrenic patients is more often controlled by stimuli which are immediate in their spatial and temporal environment than is that of normals" (p. 608).

Recently, we tested the idea that schizophrenics are restricted in the temporal range over which they can process linguistic context (J. D. Cohen et al., 1988). We designed a lexical ambiguity task, similar to the one used by Chapman et al. (1964), in which we could manipulate the temporal parameters involved.

Subjects were presented with sentences made up of two clauses; each clause appeared one at a time on a computer screen. One clause contained an ambiguous word in neutral

context (e.g., "you need a PEN"), and the other clause provided disambiguating context (e.g., "in order to keep chickens" or "in order to sign a check"). Clauses were designed so that they could be presented in either order: context first or context last. The ambiguity in each sentence always appeared in capital letters so that it could be identified by the subject. Ambiguities were used that had previously been shown to have a strong (dominant) and a weak (subordinate) meaning, and a context clause was designed for each of the two meanings. Subjects were presented with the sentences and, following each, were asked to interpret the meaning of the ambiguity as it was used in the sentence.

Sentences were distributed across three conditions: (a) weak meaning correct, context last; (b) weak meaning correct, context first; and (c) strong meaning correct, context first. For example, a given subject would have seen the ambiguity PEN in one of the three following conditions and then chosen his or her response from the list of possible meanings:

(a) without a PEN
[*clear screen, pause*]
you can't keep chickens

—*or*—

(b) you can't keep chickens
[*clear screen, pause*]
without a PEN

—*or*—

(c) you can't sign a check
[*clear screen, pause*]
without a PEN

[*clear screen, pause*]

The meaning of the word in capital letters is:

a writing implement	[*dominant meaning*]
a fenced enclosure	[*subordinate meaning*]
a kind of truck	[*unrelated meaning*]

The results of this study (shown in Figure 2) corroborated both Chapman et al.'s (1964) original findings and the explanation of their findings in terms of a restriction in the temporal range over which schizophrenics are able to use context. As Chapman et al. found, schizophrenics made significantly more dominant meaning errors than did controls when the weak meaning was correct. However, this occurred only when the context came first, as in Condition (b). When context came last, schizophrenics correctly chose the weak meaning. This was the only type of error that reliably distinguished schizophrenics from controls. Thus, schizophrenics appear to have had difficulty using context but only when it was temporally remote (i.e., came first), and not when it was more recently available (i.e., came last). This effect is consistent with Salzinger's (1971) immediacy hypothesis. Moreover, it suggests that an impairment observed in language tasks may be similar in nature to the impairments observed in attentional tasks: a difficulty in maintaining and using the internal representation of context to control action.

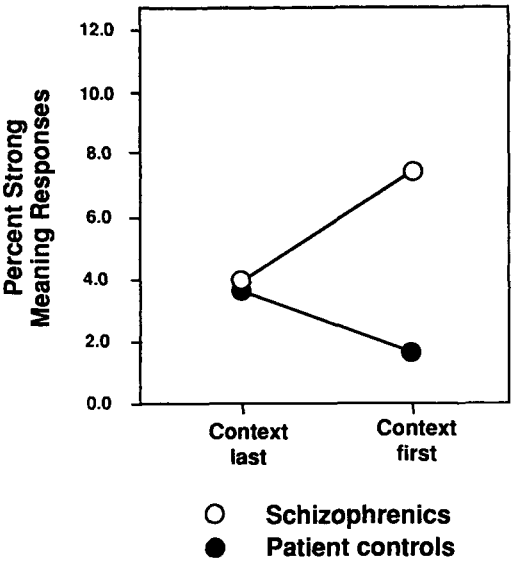

Figure 2. Medians for the rates of strong-meaning responses for schizophrenics and patient controls when the weak meaning was correct. (Because of the low overall rate of weak-meaning responses when the strong meaning was correct, and of unrelated responses in all conditions, as well as the lack of any significant differences between groups in these types of errors, these data are not shown.)

Biological Disturbances in Schizophrenia

In parallel to research on schizophrenic information-processing deficits, there has been intensive research on biological abnormalities in this illness. Some of these involve systems that are believed to play a central role in the construction and maintenance of internal representations of context: prefrontal cortex and the mesocortical dopamine system.

Prefrontal Cortex and the Internal Representation of Context

Recent studies have begun to supply direct evidence that frontal areas are involved in maintaining internal representations of context for the control of action. For example, in neurophysiological studies, Fuster (1980, 1985a, 1985b), Goldman-Rakic (1987), and others (e.g., Barone & Joseph, 1989a, 1989b) have observed cells in prefrontal cortex that are specific to a particular stimulus and response and that remain active during a delay between these. These investigators have argued that neural patterns of activity are maintained in prefrontal cortex that encode the temporary information needed to guide a response. At the behavioral level, these authors and others (e.g., Damasio, 1979; Mishkin & Pribram, 1955; Passingham, 1985; Rosenkilde, 1979; Rosvold, Szwarcbart, Mirsky, & Mishkin, 1961; Stuss & Benson, 1984) have reported data suggesting that prefrontal cortex is needed to perform tasks involving delayed responses to ambiguous stimuli. Diamond and Goldman-Rakic (1989) have emphasized that prefrontal representations are required, in particular, to overcome reflexive or previously reinforced response tendencies in order to mediate a contextually relevant—but otherwise weaker—response (see also Dia-

mond, 1985, 1990a, 1990c; Diamond & Doar, 1989). Diamond cites extensive data from lesion studies in adult monkeys and from developmental studies in human and monkey infants that use a variety of behavioral tasks (including object retrieval, visual paired comparisons, delayed response, and the A $\bar{B}$ task). Results from these and many previous studies suggest that prefrontal cortex is directly involved in maintaining representations that inhibit reflexive or habitually reinforced behaviors to attain a goal.

For example, in the A $\bar{B}$ task (pronounced "A not B"; Piaget, 1937/1954) subjects observe a desired object being hidden at one of two locations that are identical in appearance. Following a delay—during which fixation is drawn away from the hiding place—subjects are allowed to retrieve the object. The object is hidden at the same location until the subject has successfully retrieved it some number of times, after which the hiding place is switched. Normal adult monkeys and 5-year-old human children can successfully retrieve the object with delays between hiding and retrieval of 2 min or more. Monkeys with lesions of prefrontal cortex, as well as human infants younger than 6 months (in whom the frontal lobes are still in a rudimentary stage of development), can perform the task successfully only if there is no delay between the cue and test phases. With delays as short as 2 s, they show perseveration of previously reinforced responses: a return to the location at which the object was last retrieved (Diamond & Goldman-Rakic, 1989). This pattern of errors is specific to subjects lacking prefrontal cortex and is not found with lesions of the hippocampus or parietal lobes, where performance is either normal or at chance (Diamond, 1990b, 1990c).

Thus, it appears that prefrontal cortex is responsible for maintaining a representation (the location of the hidden object) required to inhibit a dominant response (return to the most recently rewarded location). Note, furthermore, that these findings help distinguish the internal representation of context from memory for specific associations between stimuli and responses. They indicate that these two functions are supported by different neural structures, with prefrontal cortex involved only in the former. It is precisely because lesions of prefrontal cortex affect internal representations of context and not associative memory that perseverations based on learned associations can occur. In contrast, lesions that involve other areas subserving such associations (e.g., hippocampus or parietal lobes) result in random, rather than perseverative, behavior (e.g., Diamond, 1990b).

The performance deficits observed for infants and frontally lesioned monkeys on delay tasks are similar to those observed for adult frontal lobe patients on the Wisconsin Card Sort Task (WCST; Grant & Berger, 1948). In this task, subjects are presented with a series of cards containing figures that vary in shape, color, and number. They are asked to sort the cards into piles according to a rule that the experimenter has in mind (e.g., "separate the cards by color"). However, subjects are not explicitly told the rule for sorting; rather, they are given feedback for each card as to whether or not they have sorted it properly. Normal subjects discover the rule quickly. Once they have demonstrated that they know it (i.e., by correctly sorting a certain number of cards in a row) the experimenter switches the rule, and the subject is required to discover the new rule. Patients with damage to the frontal lobes do poorly on this task (e.g., Milner, 1963; Nelson, 1976; Robinson, Heaton, Lehman, & Stilson, 1980). Although they are able to discover the first rule without too much difficulty, they are unable to switch to a new one: They continue to sort according to the old rule.

As in delay tasks, this behavior can be viewed as a failure to use contextual information—in this case, feedback from the experimenter—to overcome a response pattern that was correct on previous trials. Furthermore, there are additional indications from these tasks that a specific failure to use internal representations of context is involved, as distinct from a disturbance in declarative, or short-term, memory. In both the WCST and in delayed response tasks, subjects have been observed who show perseveratory behavior despite indications that they remember the relevant prior information. Thus, subjects in the WCST will sometimes comment that they know their perseveratory response is incorrect even as they carry it out (Goldberg, Weinberger, Berman, Pliskin, & Podd, 1987). Similarly, in the A $\bar{B}$ task, subjects have been observed looking at the cued (new) location while reaching for the old (incorrect) one (Diamond & Goldman-Rakic, 1989). These kinds of observations support a dissociation between declarative, or short-term, memory on the one hand, and the internal representation of context needed to actually control the response, on the other. We assume that it is the latter—representation of response-specific contextual information—that is mediated by prefrontal cortex.

Note that in both the WCST and the A $\bar{B}$ task, subjects with poor prefrontal function are not impaired in their ability to learn the basic elements of the task. Rather, they are impaired in their ability to use an internal representation of context to override the effects of prior experience in the task. This characterization of frontal lobe function fits well with clinical descriptions of the "disinhibition syndrome" that often accompanies frontal lobe pathology (e.g., Stuss & Benson, 1984). It is also consistent with difficulties that have been observed for frontal lobe patients in performing the Stroop task (Perret, 1974) and similar tasks in clinical use (e.g., the "go–no-go" paradigm) that require the subject to use task instructions to inhibit a dominant response tendency.

Finally, physiological measures have begun to provide converging evidence for the role of prefrontal cortex in supporting internal representations of context. Using measures of regional cerebral blood flow (rCBF), Weinberger and his collaborators (Berman, Illowsky, & Weinberger, 1988; Weinberger, Berman, & Chase, 1988; Weinberger, Berman, & Zec, 1986) have demonstrated that, in normal subjects, prefrontal metabolism correlates with WCST performance. Furthermore, this correlation is specific to prefrontal cortex (vs. other cortical areas). This finding corroborates the results of neuropsychological studies that link WCST performance with frontal lobe function (e.g., Nelson, 1976; Robinson et al., 1980). Weinberger's group also showed that not all cognitive tasks requiring effort and concentration are accompanied by such an increase in prefrontal activity. For example, during the Raven Progressive Matrices test—in which the task-relevant information is visually available at all times—metabolism increased in parietal and occipital areas but not in frontal areas.

Other tasks that rely on internal representations of context also appear to activate prefrontal cortex. R. M. Cohen and his

colleagues (R. M. Cohen et al., 1987; R. M. Cohen, Semple, Gross, Holcomb, et al., 1988) used positron emission tomography (PET) to measure regional brain metabolism during performance of an auditory-discrimination version of the CPT. They found an increase in prefrontal metabolism in normal subjects, which correlated with performance on the task: Subjects who made more commission errors (false alarms) showed less of an increase in metabolism in prefrontal areas. Not all studies examining frontal lobe function during CPT performance have yielded positive results (e.g., Berman, Zec, & Weinberger, 1986). However, differing results may be attributable to differences in the actual tasks and conditions that were run. We will return to this issue in the General Discussion section.

In summary, evidence from attentional tasks (e.g., the CPT and the Stroop task), problem-solving tasks (e.g., the WCST and the A $\bar{B}$ task), and from physiological imaging studies suggests that areas of the frontal cortex support the representation of information needed for response selection. A disturbance in this representation is most apparent when experimental tasks involve competing, prepotent responses. These dominant responses may have developed during the task itself (as in the WCST and the A $\bar{B}$ task), or they may have existed prior to the experiment (e.g., the Stroop task). A disturbance in the prefrontal representation manifests as a bias toward prepotent, but task-inappropriate, response tendencies (e.g., interference in the Stroop task; perseveratory patterns in the WCST and the A $\bar{B}$ task). The data reviewed earlier concerning schizophrenic performance deficits fit with this profile: an insensitivity to context and a dominant response tendency. It is not surprising to find, therefore, that frontal lobe deficits have been implicated in schizophrenia.

Frontal Deficits in Schizophrenia

The idea that the frontal lobes may be involved in schizophrenia is not new. Kraeplin (1950), who first defined this illness (as *dementia praecox*), wrote

> On various grounds it is easy to believe that the frontal cortex, which is especially well developed in man, stands in closer relation to his higher intellectual abilities, and that these are the faculties which in our patients invariably show profound loss. (p. 219)

Schizophrenics show typical frontal lobe deficits on neuropsychological tests (see Kolb & Whishaw, 1983, for a review), including the WCST (e.g., Malmo, 1974) and the Stroop task (e.g., Abramczyk et al., 1983; Wapner & Krus, 1960; Wysocki & Sweet, 1985). Several studies using imaging and electrophysiological techniques have also provided evidence suggesting frontal involvement in schizophrenia. Using rCBF, Ingvar and Franzen (1974; Franzen & Ingvar, 1975) reported abnormal perfusion of frontal areas in schizophrenics, and Buchsbaum et al. (1982) found abnormalities of glucose utilization localized to similar areas. Andreasen et al. (1986) reported evidence of frontal lobe atrophy in computerized tomographic (CT) images, and other data indicate that ventricular enlargement in schizophrenics (e.g., Andreasen et al., 1986; Weinberger et al., 1980) is associated with frontal lobe atrophy (Morihisa & McAnulty, 1985). Farkas et al. (1984) demonstrated a correlation between abnormal structure (CT) and perfusion (PET) of the frontal lobes, and Morihisa and McAnulty (1985) showed a correlation between structural (CT) and electrophysiological abnormalities.

Not all studies using physiological imaging techniques have found metabolic abnormalities of the frontal lobes in schizophrenics (e.g., Gur et al., 1987). However, in recent studies investigators have begun to use these techniques to examine frontal activity under specific behavioral conditions. Weinberger et al. (1986) demonstrated abnormal perfusion of prefrontal cortex during performance of the WCST. Similarly, R. M. Cohen et al. (1987; R. M. Cohen, Semple, Gross, Nordahl, et al., 1988) showed that schizophrenics fail to show the normal pattern of increased perfusion of prefrontal cortex during performance of a version of the CPT. These studies suggest that anatomic and physiological deficits of frontal cortex may indeed be associated with some of the behavioral deficits that have been observed in schizophrenics.

Dopamine, Prefrontal Cortex, and Schizophrenia

The hypothesis that frontal lobe dysfunction is involved in schizophrenia fits well with the prevailing neurochemical and psychopharmacological data concerning this illness. Prefrontal cortex is a primary projection area for the mesocortical dopamine system, a disturbance of which has consistently been implicated in schizophrenia (e.g., Losonczy, Davidson, & Davis, 1987; Meltzer & Stahl, 1976; Nauta & Domesick, 1981). The dopamine hypothesis is one of the most enduring biological hypotheses concerning schizophrenia. Evidence for this hypothesis comes from a variety of sources. Perhaps the strongest argument is the chemical specificity of the neuroleptics, which are used to treat the symptoms of schizophrenia. In vitro studies have demonstrated that neuroleptics have a specific affinity for dopamine binding sites and that this affinity is correlated with their clinical potency (B. M. Cohen, 1981; Creese, Burt, & Snyder, 1976; Snyder, 1976). Furthermore, drugs that influence dopamine activity in the central nervous system—such as amphetamines and L-dopa—exacerbate symptoms in psychotic patients (Angrist, Peselow, Rubinstein, Wolkin, & Rotrosen, 1985; Davidson et al., 1987; Janowsky, Huey, Storms, & Judd, 1977; Lieberman et al., 1984) and may induce psychosis in nonpsychotic individuals (e.g., Janowsky & Rich, 1979; Snyder, 1972). Studies of the plasma (Bowers, Heninger, & Sternberg, 1980; Pickar et al., 1984) and cerebrospinal fluid (Sedvall, Fyro, Nyback, Wiesel, & Wode-Helgodt, 1974) of schizophrenics have revealed abnormal levels of dopamine metabolites. Finally, several postmortem studies have found evidence for an elevation in the number of dopamine receptors in schizophrenics compared to controls (e.g., Cross, Crow, & Owen, 1981; Seeman et al., 1984; Tyrone & Seeman, 1980), and this elevation has been found to correlate with the previous experience of hallucinations and delusions (Crow et al., 1984).

Whereas different investigators have argued that central dopamine activity is either reduced (e.g., Early, Posner, Reiman, & Raichle, 1989; Karoum, Karson, Bigelow, Lawson, & Wyatt, 1987) or increased (e.g., Creese et al., 1976; Snyder, 1972, 1976) in schizophrenia, one hypothesis is that both conditions may occur (either within or across individuals) and that each is associated with a different psychopathological profile. For exam-

ple, both Crow (1980) and Mackay (1980) suggested that the symptoms of schizophrenia can be divided into two subtypes, one that reflects dopamine overactivity (positive symptoms—e.g., hallucinations and delusions) and another that reflects dopamine underactivity (negative symptoms—e.g., motivational difficulties, cognitive impairment, paucity in content of speech, flattening of affect, and deficits in social functioning). Several authors have argued that it is the negative symptoms of schizophrenia that are most often associated with frontal lobe deficits (e.g., Andreasen, Flaum, Swayze, Tyrrell, & Arndt, 1990; Andreasen et al., 1986, 1991; Goldman-Rakic, 1991; Levin, 1984). This is consistent with mounting evidence that mesocortical dopamine activity in prefrontal cortex is directly related to cognitive function and that a reduction of this activity can produce many of the cognitive deficits observed in schizophrenics. Thus, McCulloch, Savaki, McCulloch, Jehle, and Sokoloff (1982) showed that activation of mesocortical dopaminergic neurons increases metabolic activity in the prefrontal cortex of animals. Conversely, lesions of the same dopamine projections reduce metabolism in prefrontal cortex and impair cognitive functions usually associated with this brain region, such as the execution of search strategies or delayed-alternation tasks (Oades, 1981; Simon, Scatton, & Le Moal, 1980). For example, rhesus monkeys could not perform a delayed-alternation task following selective destruction of dopamine terminals in prefrontal cortex (Brozoski, Brown, Rosvold, & Goldman, 1979). This deficit was as severe as that following full surgical ablation of the same area of cortex. Moreover, performance recovered almost entirely with dopamine agonists such as L-dopa and apomorphine. Similar findings have been reported with respect to attentional impairments (e.g., Corwin, Kanter, Watson, Heilman, Valenstein, & Hashimoto, 1986) and, more recently, working memory deficits (Sawaguchi & Goldman-Rakic, 1991). Finally, studies of human patients suffering from Parkinson's disease—in which dopamine function is markedly impaired—provide similar evidence: Even when these patients do not display clinically significant cognitive deficits, they display impairments on the WCST similar to those observed in frontal lobe subjects (Bowen, Kamienny, Burns, & Yahr, 1975). This deficit is less pronounced in patients receiving the dopamine precursor L-dopa, which often has therapeutic efficacy in reestablishing dopaminergic tone.

In view of these findings, several authors have proposed that reduced dopaminergic tone in prefrontal cortex may be associated with frontal lobe abnormalities in schizophrenia and may be responsible for several of the cognitive deficits that have been observed. Levin (1984) reviewed a wide variety of behavioral data in support of this hypothesis (also see Goldman-Rakic, 1991; Levin, Yurgelun-Todd, & Craft, 1989). Recently, physiological data have also begun to accumulate. Weinberger, Berman, and Illowsky (1988) reported that levels of the dopamine metabolite homovanillic acid in the cerebrospinal fluid of schizophrenics show a strong correlation with prefrontal activity during WCST performance. In another study, Geraud, Arne-Bes, Guell, and Bes (1987) were able to reverse the metabolic hypofrontality observed in schizophrenics on PET by administration of a dopamine agonist. Thus, there is growing evidence that dopamine is closely related to the activity of prefrontal cortex and that a disturbance in this system may be involved in schizophrenic cognitive deficits. What is lacking, however, is a coherent account of these findings in terms of mechanisms that link biological processes with performance in behavioral tasks.

Summary

We have reviewed evidence suggesting (a) that schizophrenics suffer from an inability to construct and maintain internal representations of context for the control of action, (b) that prefrontal cortex plays a role in maintaining such representations, (c) that an intact mesocortical dopamine system is necessary for the normal function of prefrontal cortex, and (d) that disturbances of both prefrontal cortex and the mesocortical dopamine system appear to be involved in schizophrenia. However, despite a growing recognition that these findings are related, no theory has yet been proposed that answers the following question: How does a disturbance of dopamine activity in prefrontal cortex lead to the pattern of cognitive deficits observed in schizophrenia? In the section that follows, we describe simulation models developed within the connectionist framework that attempt to answer this question.

Connectionist Simulations

We have constructed a set of simulation models that show how the connectionist framework can be used to link the biological and behavioral processes just discussed. In this section, we present three such models, which simulate performance in three of the experimental tasks discussed previously: the Stroop task, the CPT, and the lexical disambiguation task. In each case, we show that a simulation of reduced dopamine effect in a component of the model identified with the function of prefrontal cortex results in performance deficits that are quantitatively similar to those observed for schizophrenics in the corresponding tasks. As background for understanding these models, we first provide a brief overview of the connectionist framework.

The Connectionist Framework

The principles of connectionism, or parallel distributed processing (McClelland & Rumelhart, 1986; Rumelhart & McClelland, 1986b), provide a framework for building computer models that can simulate cognitive phenomena. At the same time, these principles are meant to capture the salient details of the mechanisms underlying information processing in the brain. They can be roughly divided into those having to do with processing and those having to do with training.

Processing

Each unit in a connectionist network (see Figure 3) is a simple summing device: It accumulates inputs from other units and adjusts its output in response to these inputs. Typically, units are grouped into modules, and modules are connected into pathways. Information is represented as the pattern of activation over the units in a module. The activation of each unit is a real-valued number varying continuously between a minimum and maximum value, which can be thought of as the unit's

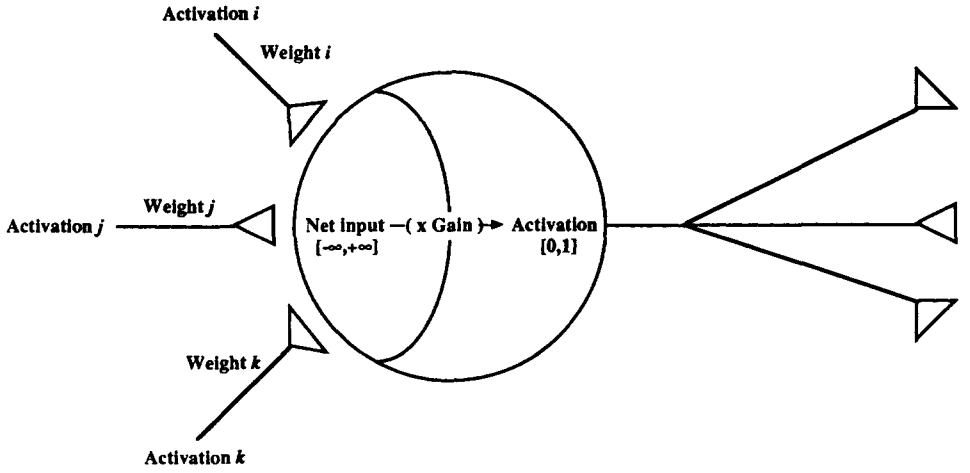

Figure 3. Schematic representation of a typical unit in a connectionist system.

probability of firing. The responsivity of each unit is scaled by its *gain* parameter, which serves as a multiplier for the effects of excitatory and inhibitory inputs to the unit. Processing occurs by the propagation of signals (spread of activation) among units within and between modules. This occurs via the connections that exist between units. The connections between the units of different modules constitute processing pathways.

Training

The ability of this type of system to perform a given task depends on its having an appropriate set of connection weights in the pathway that runs from the input modules to the output modules relevant to the task. The connections in a pathway are set by learning. Although a number of different connectionist learning techniques have been described, the generalized delta rule, or back propagation algorithm (Rumelhart, Hinton, & Williams, 1986), is in widest use. In brief, this involves the following series of operations: (a) present an input pattern to the network; (b) allow activation to spread to the output level; (c) compute the difference (error) for each output unit between its current activation and the one desired (i.e., the one specified by the target, or teaching pattern); (d) "back propagate" these error signals toward the input units. The back propagation algorithm provides a way for each unit in a pathway to compute the adjustment it must make to its connection weights so as to best reduce the error at the output level. A common criticism of this algorithm is that it is not biologically plausible. That is, it is difficult to imagine that real neural systems rely on the back propagation of error signals for learning. However, back propagation implements the general phenomenon of *gradient descent*—the gradual reduction of error by incremental adjustments in connection weights. Gradient descent has proven to be a powerful concept for describing many of the details concerning human learning behavior (e.g., J. D. Cohen, Dunbar, & McClelland, 1990). Thus, it may be that back propagation offers a reasonable approximation of the type of learning that occurs in neural systems, even if the actual algorithm is different.

It is important to recognize also that most connectionist models are not intended to be detailed circuit diagrams of actual neural networks. Rather, like statistical mechanical models in physics and chemistry, connectionist models are designed to capture those features of a lower level system (information-processing mechanisms in the brain) that are most relevant at a higher level of analysis (cognition and behavior). Thus, an important goal in constructing such models is to make it possible to examine the effects of biological variables on behavior without having to reproduce the entire brain.

We believe that developing biologically plausible models of information processing will (a) lead to more realistic models of cognitive phenomena and (b) make it possible to relate behavior directly to biological processes (for an in-depth discussion, see Rumelhart & McClelland, 1986c). Connectionist models have begun to show promise along these lines in their ability to explain a variety of phenomena at both the biological and behavioral levels. These include the computation of spatial orientation from retinal and eye-position information (Zipser & Andersen, 1988), the computation of object shape from shading information (Lehky & Sejnowski, 1988), the acquisition of regular and irregular verb forms in English (Rumelhart & McClelland, 1986a), text-to-speech translation and disturbances of this phenomenon in surface dyslexia (Patterson, Seidenberg, & McClelland, 1989; Seidenberg & McClelland, 1989), and access to word meaning from word form in deep dyslexia (Hinton & Shallice, 1991).

Using the connectionist framework, we developed simulation models of three tasks relevant to research on schizophrenia: the Stroop task, the continuous performance test, and the lexical disambiguation task described earlier. Each model was designed to simulate normal performance in one of these tasks. Although the models differ in the details necessary to capture differences between these tasks, all three rely on a common set of information-processing principles (as just described) and, in particular, share a common mechanism for representing and maintaining context. In each case, this mechanism relies on a specific module that we identify with the function of the prefrontal cortex. After establishing the ability of each model to capture normal performance of its particular task, we exam-

ined the effects that reducing the gain of units in the context module had on performance in order to explore the hypothesis that such a disturbance can account for schizophrenic deficits in that task. We begin our description of the models by showing how the physiological influence of dopamine can be simulated by changes in the gain parameter of individual units. We then describe simulations which show that a change in gain in the module used to represent context can account for differences between normal and schizophrenic performance in the Stroop, CPT, and lexical disambiguation tasks.

Simulation of the Physiological Effects of Dopamine

In contrast to other neurotransmitter systems, such as amino acids or peptides, the anatomy and physiology of dopamine systems are not suited to the transmission of discrete sensory or motor messages. Rather, like other catecholamine systems, dopamine systems are in a position to modulate the state of information processing in entire brain areas over prolonged periods of time. Several anatomical and physiological observations support this contention. Dopamine neurons originate in discrete nuclei localized in the brain stem, and their fibers project radially to several functionally different areas of the central nervous system. The baseline firing rate of these neurons is low and stable, and the conduction velocity along their fibers is slow. These characteristics result in a steady state of transmitter release and relatively long-lasting postsynaptic effects that are conducive to modulatory influences. Most important, recent evidence suggests that, at least under certain circumstances, the effect of dopamine release is not to directly increase or reduce the firing frequency of target cells (e.g., Chiodo & Berger, 1986; Schneider, Levine, Hull, & Buchwald, 1984); rather, like norepinephrine, dopamine can modulate the response properties of postsynaptic cells such that both inhibitory and excitatory responses to other afferent inputs are potentiated. Some investigators have described this effect as an increase in the "signal-to-noise ratio" of the cells' behavior (Foote, Freedman, & Oliver, 1975) or an "enabling" of the cells' response (Bloom, Schulman, & Koob, 1989).

The modulatory effects of dopamine have been investigated mostly in the striatum, where they are similar to those observed for norepinephrine elsewhere. The results of investigations conducted directly in the prefrontal cortex are less clear. Some studies (Bunney & Aghajanian, 1976; Ferron, Thierry, Le Douarin, & Glowinski, 1984; Reader, Ferron, Descarries, & Jasper, 1979; Bunney & Sesack, 1987) report a potentiation of inhibitory responses but a reduction of excitatory responses. However, in these studies the amount of dopamine released also produced a direct decrease in the baseline firing rate of the cells. This direct decrease in baseline firing rate has also been observed in striatal cells, but only when large amounts of dopamine were released and not for smaller amounts (Chiodo & Berger, 1986). Thus, the reduction of excitatory responses in prefrontal cortex that has been reported may be related to the use of high concentrations of dopamine.[2] The effects of smaller concentrations—which do not affect baseline firing rate—have not been tested. More consistent with the idea of dopamine-induced *potentiation* are the results of two other studies, conducted in primate prefrontal cortex (Aou, Nishino, Inokuchi, & Mizuno, 1983; Sawaguchi & Matsumura, 1985), in which inhibitory as well as excitatory effects of dopamine iontophoresis were found.[3]

In our models, we assume that the effects of dopamine on cells in prefrontal cortex—at concentrations relevant to the behavioral tasks we are interested in—are similar to the effects that have been observed in striatal cells: a potentiation of cell responses. Two caveats are warranted by this assumption. First, although potentiation is compatible with existing data, as just noted, it has not been substantiated directly in physiological studies of prefrontal cortex. Second, the mechanism by which dopamine potentiates both excitatory and inhibitory inputs has not yet been elucidated. In particular, it is not clear whether this is a direct effect of dopamine on the target cell or whether potentiation arises from local interactions among inhibitory and excitatory cells. We will consider both of these issues further in the General Discussion section. In the simulations described in this section we have not concerned ourselves with the detailed mechanisms of how potentiating effects arise at the cellular level; we have focused instead on their functional significance by capturing the effects of potentiation with a single parameter.

In the models we simulate the action of dopamine by changing gain, which is a parameter of the function that relates a unit's input to its activation value. Thus, first we assume that the relationship between the input to a neuron and its firing rate can be simulated as a nonlinear function relating the net input of a unit in the model to its activation value. Physiological experiments suggest that in biological systems the shape of this function is sigmoid, with its steepest slope around the baseline firing rate (e.g., Burnod & Korn, 1989; Freeman, 1979). The same experiments also indicate that small increments in excitatory drive produce changes in firing frequency that are greater than those produced by equivalent increments in inhibitory input. These properties can be captured by the logistic function with a constant negative bias (see Figure 4, gain = 1.0):

$$\text{activation} = \frac{1}{1 + e^{-(\text{gain}^*\text{net}) + \text{bias}}}.$$

The potentiating effects of dopamine can then be simulated by increasing the gain parameter of the logistic function. As Figure 4 illustrates, with a higher gain (gain = 2.0) the unit is more sensitive to afferent signals, whereas its baseline firing rate (net input = 0) remains the same. Elsewhere we have shown that such a change in gain can simulate a number of different catecholaminergic effects at both the biological and behavioral

[2] For example, in Reader, Ferron, Descarries, and Jasper (1979) both the concentration of dopamine in the micropipette and the intensity of the iontophoretic current were almost one order of magnitude greater than the corresponding concentrations and current intensity used in the Chiodo and Berger (1986) study. Moreover, the effect of dopamine at these concentrations was to completely inhibit spontaneous firing in the target cells.

[3] It is worth noting that different dopamine receptor subtypes predominate in the striatum and the prefrontal cortex. Whereas the potentiating effects of dopamine have been studied primarily within the striatum (where D2 receptors predominate), D1 receptors in this region have been shown to mediate potentiating effects (Hu & Wang, 1988).

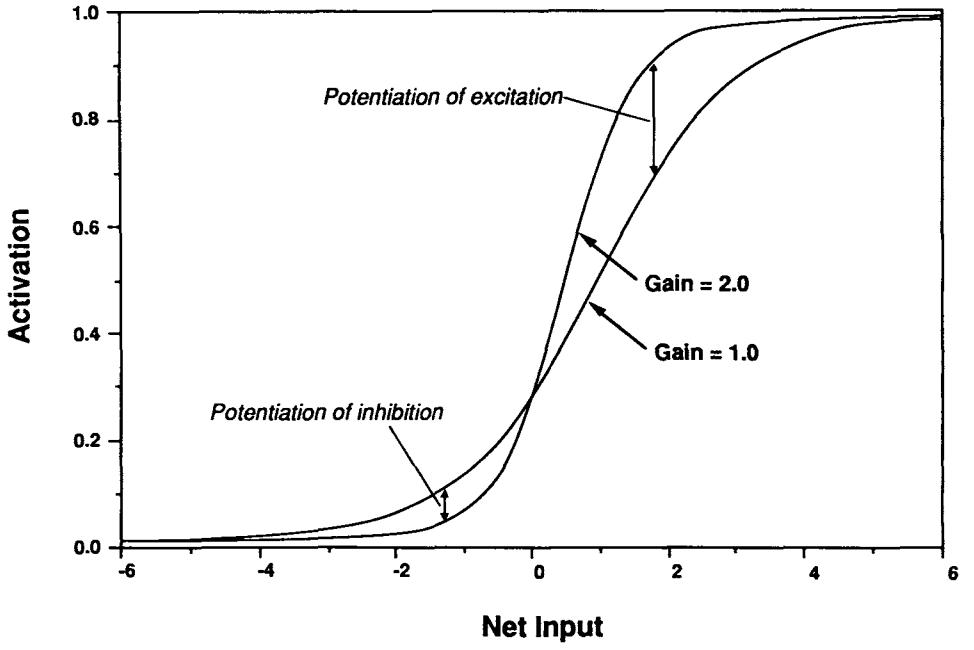

Figure 4. The influence of the gain parameter on the logistic activation function of an individual unit. (Note that with an increase in gain, the effect of the net input on the unit's activation is increased, whereas the reverse is true with a decrease in the gain. These effects simulate the consequences of dopamine release on target neurons in the central nervous system.)

levels (e.g., the influence of catecholamines on the receptive field of individual units, the influence of amphetamines on stimulus detection in humans, and stimulus response generalization in both humans and rats; see Servan-Schreiber, 1990; Servan-Schreiber, Printz, & Cohen, 1990).

To simulate the effect of a neuromodulator, such as dopamine, we change gain equally for all units in the model that are assumed to be influenced by that neuromodulator. For example, the mesocortical dopamine system has extensive projections to prefrontal cortex. To model the action of dopamine in this brain area, we change the gain of all units in the module corresponding to this area. In the following models, decreased dopaminergic activity in prefrontal cortex was simulated by reducing the gain of units in the module used to represent and maintain context.

A Connectionist Model of Selective Attention (the Stroop Effect)

Architecture and Processing

Elsewhere we have described a connectionist model of selective attention that simulates human performance in the Stroop task (J. D. Cohen et al., 1990). In brief, this model consists of two processing pathways, one for color naming and one for word reading (see Figure 5). Simulations are conducted by activating input units corresponding to stimuli used in an actual experiment (e.g., the input unit representing the color red in the color naming pathway) and allowing activation to spread through the network. This leads to activation of the output unit corresponding to the appropriate response (e.g., "red"). Reaction time is considered to be linearly related to the number of cycles it takes for an output unit to accumulate a specified amount of activation.

Training

The model is trained to produce the appropriate behavior by presenting it with the input patterns for each of the responses it is expected to make and using the back propagation learning algorithm to adjust the connection weights accordingly. During training, the model is given more experience with (i.e., a greater number of training trials on) the word-reading task than the color-naming task. This corresponds to the common assumption that human adults have had more experience generating a verbal response to written words than to colors they see. Because of this, the connection weights in the word-reading pathway become greater than those in the color-naming pathway. As a result, when the network is presented with conflicting inputs in the two pathways (e.g., the word RED and the color green), it responds preferentially to the word input. Of course, human subjects are able to overcome this tendency and respond to the color instead of the word when requested to do so. To capture this effect in the model, a set of units are included that represent the intended behavior (i.e., color naming vs. word reading). Thus, the specification of a particular task is represented by the appropriate pattern of activation over a set of "task demand" units. These are connected to the intermediate units in each of the two pathways and modulate their responsivity. For example, when the pattern corresponding to "color naming" is activated

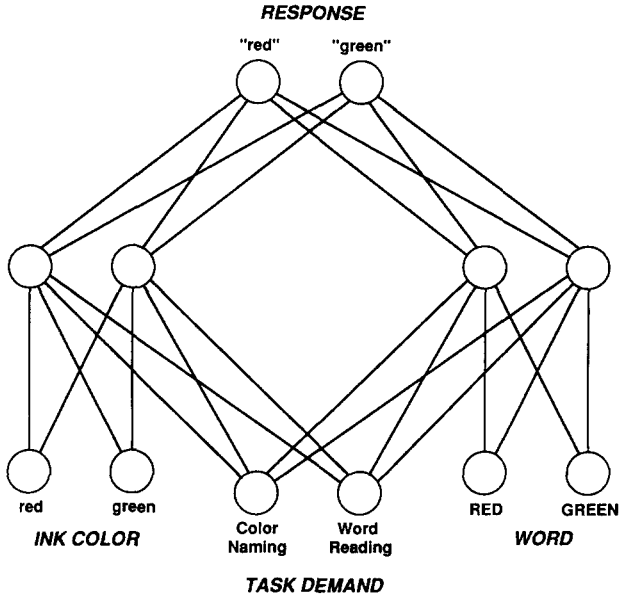

Figure 5. Network architecture. (Units at the bottom are input units, and units at the top are the output [response] units. From "On the Control of Automatic Processes: A Parallel Distributed Processing Account of the Stroop Effect" by J. D. Cohen, K. Dunbar, and J. L. McClelland, 1990, *Psychological Review, 97,* p. 336. Copyright 1990 by the American Psychological Association.)

over the task demand units, activation spreading from these units has a sensitizing effect on processing units in the color pathway while it "desensitizes" units in the word pathway. This modulates the flow of information in the two pathways, favoring the color pathway. The result is that although the connection strengths in the color pathway are weaker, a signal presented to this pathway is able to overcome the otherwise dominant response mediated by the word pathway. In other words, the model is able to selectively attend to information in the task-relevant pathway. Note that spreading activation and attentional modulation are not different processes. Attentional modulation of both pathways is a consequence of activation spreading from the task demand units to the intermediate units in each pathway. Thus, both the "direct processing" of information as well as its attentional modulation rely on the same mechanisms of processing (see J. D. Cohen et al., 1990, for a more detailed discussion).

Simulation

This simple model is able to simulate a wide variety of empirical phenomena associated with the Stroop task. It captures all of the phenomena depicted in Figure 1 (asymmetry in speed of processing between word reading and color naming, the immunity of word reading to the effects of color, the susceptibility of color naming to interference and facilitation from words, and greater interference than facilitation effects). It also captures the influence of practice on interference and facilitation effects, the relative nature of these effects (i.e., their dependence on the nature of a competing process), stimulus onset asynchrony effects, and response set effects (see J. D. Cohen et al., 1990).

The model also clarifies the relationship between attention and the internal representation of context. Stimuli that vary in more than one dimension are inherently ambiguous (e.g., "Should I respond to the word or the color?"). Task instructions provide the context necessary to disambiguate the stimulus and choose the appropriate response. Furthermore, task instructions must be represented internally because, as we have said, the stimuli themselves do not indicate which task to perform. In the model, this internal representation was captured as a pattern of activation in the task demand module. This had a direct attentional effect: It was responsible for the model's selecting one pathway for the processing of information and not the other. Thus, the model suggests that attentional selection can be thought of as the mediating effects that the internal representation of context has on processing.

These ideas are directly relevant to schizophrenic deficits. If prefrontal cortex is responsible for maintaining the internal representation of context, and if schizophrenia involves a disturbance of frontal lobe function, then we should be able to simulate schizophrenic deficits in the Stroop task by disturbing processing in the task demand module. More specifically, if frontal lobe dysfunction in schizophrenia is due to a reduction in the activity of its dopaminergic supply, then we should be able to simulate this by reducing the gain of units in the task demand module.

Panel B of Figure 6 shows the results of this simulation, in which the gain of only the task demand units was reduced; all other units were unperturbed. This change in the context (i.e., task demand) module produces effects similar to those observed for schizophrenics: an increase in overall response time, with a disproportionate increase on color-naming conflict trials. It is important to emphasize here that this simulation was conducted without making any changes to the original Stroop model (J. D. Cohen et al., 1990) other than manipulating the gain of units in the task demand module, as motivated by our hypothesis.

It is interesting that the simulation shows that a lesion restricted to the mechanism for representing context can produce an *overall* degradation of performance in addition to the expected specific attentional deficit (i.e., increased interference). The overall slowing occurs because, according to this model, all processes rely on attention to some degree (see J. D. Cohen et al., 1990, for a detailed discussion). Disproportionate slowing occurs in the conflict condition because weaker, less automatic processes (e.g., color naming) rely more on selective attention (i.e., mediating context) than do stronger, more automatic ones (e.g., word reading), particularly when these are in competition.

The model also allows us to address a problem that frequently besets the interpretation of schizophrenic deficits. Recall the argument that, given an overall degradation of performance, it is difficult to know whether poor performance in a particular experimental condition is due to a specific deficit or to a more generalized one responsible for the overall degradation. This reflects a limitation in our ability to attribute cause when the information we have about a system is restricted to its behavior and we lack any knowledge about underlying mechanisms. However, the model provides us with a tool for specify-

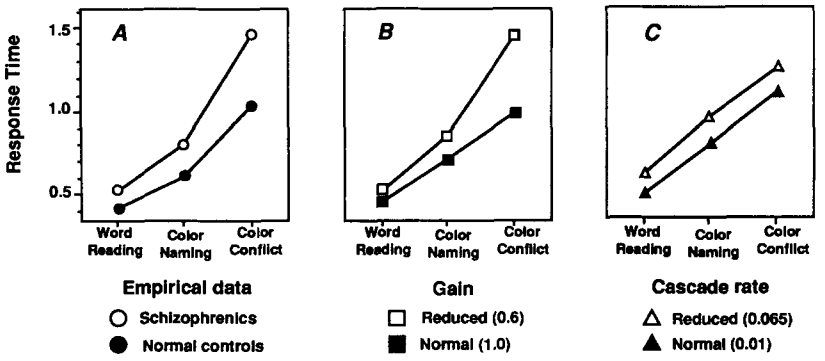

Figure 6. Stroop task performance for normal and schizophrenic subjects, and results from simulations manipulating the gain parameter (task demand units only) and cascade rate (all units) in the network. (Empirical data are the response times [in seconds] for stimuli in each condition of the Stroop task, averaged over the three empirical studies reported in Table 1. Simulation data are the number of cycles required for processing stimuli of each type, averaged over 1,000 trials of each type and scaled by a constant [0.02] to facilitate comparison between these and the empirical data. The 5% confidence intervals for all simulation data are all less than 0.02.)

ing possible mechanisms and studying the behavior they produce. Earlier, we described a mechanism for a specific attentional deficit: a disturbance in the context module. We can compare this to a more generalized deficit—one that produces an overall slowing of response—by decreasing the rate at which information accumulates for units in the network. This is determined by a parameter called the *cascade rate.* We tested for the effects of a generalized deficit of this sort by reducing the cascade rate for all units in network. The reduction in cascade rate was selected by matching the degree of slowing in the word-reading condition of the simulation to the amount of slowing observed for schizophrenics relative to control subjects in the empirical data. The results of this manipulation are shown in Panel C of Figure 6. There is an overall slowing of response, but no disproportionate slowing in the interference condition. Thus, slowing the processing rate of units throughout the network is unable to account for schizophrenic performance in the interference condition of the task. We explored other deficits that produce an overall slowing of response (e.g., an increase in the response threshold) with similar results. In contrast, as we noted above, impairment of the context module produces both effects: Slowing occurs in all conditions but is most pronounced in the interference condition.

Chapman and Chapman (1978) pointed out the danger in assuming that degraded performance in a particular task condition necessarily reflects a deficit in processes related to that condition. If the condition is also the hardest one for normal subjects (as is the case for the interference condition of the color-naming task), then even a disproportionate degradation of performance in that condition could be caused by a generalized deficit (i.e., one that is not specific to any particular processing component). We have tried to show how a simulation model can help us deal with this problem. Our model demonstrates that a specific attentional deficit provides a better account for the data than a number of possible generalized deficits. Furthermore, the model provides a new interpretation of the data, reversing the typical view: It shows how a general degradation in performance can arise from a specific deficit, rather than the other way around. To our knowledge, this possibility has not been considered in the literature. Of course, our results do not preclude the possibility that some other model could account for the findings in terms of a different deficit—specific or generalized. However, by providing an explanation of the findings in terms of an explicit set of information-processing mechanisms, we have set a threshold for explanation that must be met by competing alternatives. Furthermore, we have shown how simulation models can be used to deal with the problem of differential deficits described by Chapman and Chapman. When tasks (or conditions) differ in difficulty, it is still possible to compare competing hypotheses by specifying the mechanisms believed to underlie the deficit and comparing their ability to account for the empirical data.

Finally, as discussed earlier, the model relates a disturbance in selective attention directly to the processing of context. Selective attention is viewed as the effects that context has on processing, and a failure to maintain an appropriate contextual representation (e.g., the task demand specification) leads directly to a failure in selective attention. In the Stroop task, this manifests as an increased susceptibility to interference in the color-naming task. This interference, in turn, reflects the increased influence of dominant response processes (e.g., word reading) that occurs with the weakening of attention. Schizophrenic performance has often been characterized as reflecting a dominant response tendency (e.g., Chapman et al., 1964; Maher, 1972), although no specific mechanism has previously been proposed for this. We will return to this issue later in our discussion of schizophrenic language performance.

Simulation of the Continuous Performance Test (CPT)

The Stroop model shows how contextual information and its attentional effects can be represented in a connectionist model and how a specific disturbance in this mechanism can explain important aspects of schizophrenic performance. One question

we might ask is: How general are these findings? Here, we extend the principles described in the Stroop model to account for performance in the CPT.

As we discussed earlier, schizophrenics show consistent deficits in the CPT. This is particularly true for variants in which demand is placed on the active maintenance of internal representations of context. For example, in the CPT-double, a target consists of any consecutive reoccurrence of a stimulus (e.g., a B immediately following a B). Thus, subjects must be able to use a representation of the previous stimulus as context for responding to the subsequent one. We have argued that such response-specific representations are maintained in prefrontal cortex and that this function of prefrontal cortex is impaired in schizophrenics. Indeed, schizophrenics perform poorly in the CPT-double and similar versions (Cornblatt et al., 1989; Nuechterlein, 1984). We suggest that, like deficits in the Stroop task, this impairment can be explained by a reduction of dopaminergic tone in prefrontal cortex resulting in an impairment of the internal representation of context. If this is so, then we should be able to simulate schizophrenic deficits in the CPT-double using the same manipulation used to produce deficits in the Stroop task: a reduction of gain in the module responsible for representing and, in this case, maintaining context. To test this, we constructed a network to perform the CPT-double.

The network consisted of four modules: an input module, an intermediate (associative) module, a module for representing the prior stimulus, and a response module (see Figure 7; details of the simulation are described in the Appendix). The input module was used to represent the visual features of individual letters. Stimulus presentation was simulated by activating the input units corresponding to the features of the stimulus letter. This produced a unique pattern of activation for each letter. The network was trained to record the presentation of a given input pattern by activating the appropriate unit in the prior stimulus module. In addition, the network was trained to activate the unit in the response module whenever a stimulus letter appeared twice or more in a row. To do this, however, the network must be able to use the information stored in the prior stimulus module. To make this possible, a set of connections was added from the prior stimulus module back to the intermediate module. Intermediate units could thus receive "bottom up" information from the feature units (representing the current input) and "top down" information from the prior stimulus units. This allowed the network to compare the current and previous inputs and thereby learn to activate the response unit whenever these were the same—that is, whenever two consecutive letters were identical.

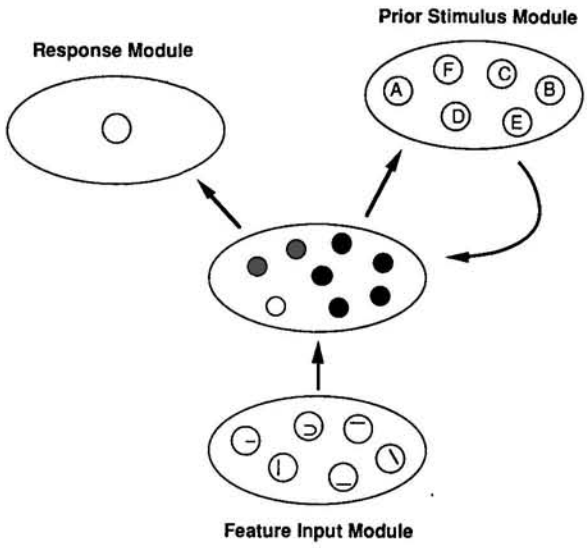

Figure 7. Network used to simulate the continuous performance test-double. (Note the bidirectional connections between units in the intermediate and prior stimulus modules.)

Note that the prior stimulus units in this model played the same role as the task demand units did in the Stroop model. The representation over the prior stimulus units in the CPT model provided the context for disambiguating the response to a particular pattern of input, just as the task demand units did in the Stroop model. The only important difference is that the context in the CPT model was determined by the previous input and therefore changed from trial to trial. We should emphasize, however, that the function of the prior stimulus module should not be thought of as a form of declarative, short-term memory. For the sake of clarity, we have labeled information in this module as individual letters. However, we imagine that in actuality this information is stored in a form that can be used to govern response selection but that may not be suitable for identification or reporting of the actual letter.

Following training, the network was able to perform the CPT-double task perfectly for a set of 10 different stimuli. To simulate the performance of normal subjects—who typically miss on 17% of trials and produce false alarms on 5% of trials (see Figure 8)—we added noise to processing. Noise in neural systems is usually attributed to afferent signals that are independent of the relevant stimulus. To simulate this distortion of input, we added a small amount of random, normally distributed noise to the net input of each unit on every processing cycle. The amount of noise was adjusted to match the performance of the network with that of human subjects. The results of this simulation also appear in Figure 8 (gain = 1.0). Then, to simulate schizophrenic performance, we disturbed processing in the prior stimulus module—which was responsible for representing and maintaining context—by decreasing the gain of these units to the same level used in the Stroop simulation (0.6). No other changes were made to the model. The percentage of misses increased to 44.9%, and false alarms increased slightly to 8.9%. These results closely match those from empirical studies of schizophrenic subjects (see Figure 8).

Although some authors have interpreted schizophrenic performance in the CPT in terms of a deficit in sensory processing, our model suggests an alternative hypothesis: Performance deficits are due to a degradation in the internal representation required—as context—for processing the current stimulus. We hypothesize that this representation is maintained in prefrontal cortex and is directly influenced by changes in the dopaminergic supply to this area. This hypothesis is consistent with our account of Stroop performance and with disturbances of language processing, which we turn to next.

Simulation of Context-Dependent Lexical Disambiguation

In the two previous simulations, we focused on different aspects of schizophrenic performance deficits and the mecha-

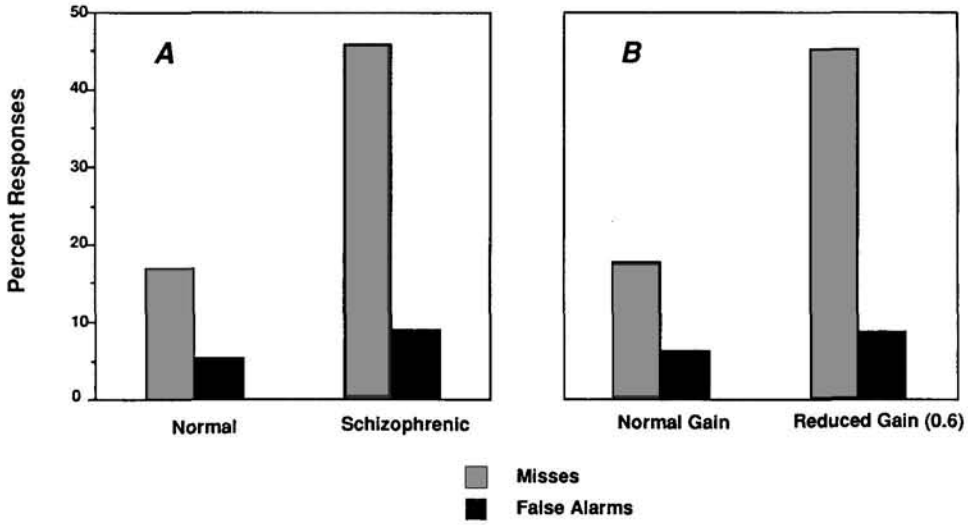

Figure 8. Percentage of misses and false alarms for (A) normal and schizophrenic subjects in the continuous performance test, and (B) the simulation run with normal and reduced gain on units in the prior stimulus module. (Empirical data are taken from Cornblatt, Lenzenweger, and Erlenmeyer-Kimling [1989]. That article reported *d'* and ln beta values; we obtained the values for misses and false alarms from B. Cornblatt directly [personal communication, March 1990]. In addition, Cornblatt et al. [1989] distinguished between "false alarms" [responses to stimuli similar to the target] and "random errors"; because both types of errors consist of responses to nontarget stimuli, we have combined these and considered them together as false alarms. Simulation data are based on 1,000 trials run in each condition. The 5% confidence intervals for the normal gain condition were ±2.3% for misses and ±1.7% for false alarms; for the reduced gain condition they were ±3.1% for misses and ±1.8% for false alarms.)

nisms involved. The Stroop simulation showed how a disturbance in the internal representation of context can lead to dominant response tendencies. In the CPT simulation, we introduced a mechanism for generating and maintaining previously generated representations of context and showed how a disturbance in this mechanism could account for experimental results. The lexical disambiguation task (described earlier) allows us to explore both of these factors (dominant response biases and maintenance of the internal representation of context) within a single paradigm. The results of our study using this task replicated the finding of others that schizophrenics show a tendency to respond to the dominant meaning of lexical ambiguities even when context confers the weaker, less frequent meaning. However, our results suggested that this tendency is significant only when context is temporally remote, implicating a deficit in the maintenance of context. Here, we show how this language-processing deficit can be simulated using the same principles that we used to account for schizophrenic performance in the Stroop task and the CPT.

Architecture and Processing

To simulate performance in the lexical disambiguation task, we constructed a network (see Figure 9; details of the simulation are described in the Appendix) with the same basic architecture as the CPT model. In this case, the input module was used to represent lexical stimuli (e.g., the word PEN). The network was trained to associate patterns of activation in this module with patterns in two other modules: a response module and a discourse module. Patterns in the response module specified the meaning of the input word (e.g., "writing implement"), whereas the discourse module was used to represent the topic of the current sequence of inputs (e.g., the meaning of a sentence or phrase, as opposed to the meaning of individual words). The intermediate module functioned as a semantic module, encoding an internal representation for the meaning of the input that could be used to generate an appropriate response in the output module and a relevant discourse representation in the discourse module. Analogous to the CPT model, there were two-way connections between the semantic module and the discourse module. Thus, once a discourse representation had been activated (e.g., by a prior input pattern), it could be used to influence the processing of subsequent stimuli in the semantic module. This provided the mechanism by which context could be used to resolve lexical ambiguity.

Training

The model was trained to produce an output and discourse representation for 30 different input words, some of which were ambiguous. In the case of ambiguous words, the model was trained to produce the response and discourse patterns related to one meaning (e.g., PEN → "writing implement" and *writing;* we use uppercase letters to denote input words, quotation marks to denote output responses, and italics to denote discourse representations) more than the other (e.g., PEN → "fenced enclosure" and *farming*). This asymmetry of training was analogous to training in the Stroop model (words more than colors), with a comparable result: When presented with an ambiguous input word, the network preferentially activated the

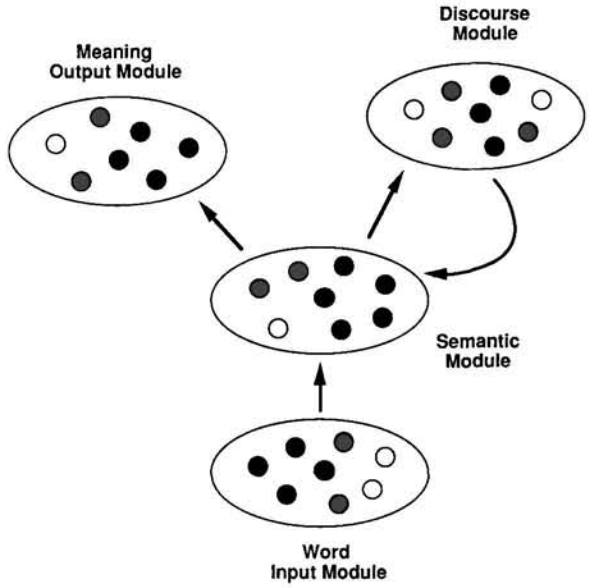

Figure 9. Schematic diagram of the language-processing model. (Patterns of activation over the units in the input module are assumed to represent the current sensory stimulus [e.g., the orthographic code for a written word], whereas the output module is assumed to represent the information necessary to generate an overt response [e.g., the phonological code needed to pronounce the meaning of the word]. Note that the connections between the semantic and discourse modules are bidirectional.)

strong (more frequently trained) response and discourse representations. To permit access to the weaker meaning, the network was sometimes presented with an ambiguous word along with one of its associated discourse representations (e.g., PEN and *farming*)[4] and trained to generate the appropriate response (i.e., "fenced enclosure"). Finally, the network was trained on a set of context words, each of which was related to one meaning of an ambiguity; these words (e.g., CHICKEN) were trained to produce their own meaning as the response ("fowl") as well as a discourse representation that was identical to the corresponding meaning of the related ambiguity (*farming*).

The combined effect of these training procedures was that when an ambiguous word was presented and no representation was active over the discourse units, the output was a blend of the two meanings of the word, with elements of the more frequently trained (dominant) meaning being more active than the other (subordinate) meaning. However, when a discourse representation was active, the model successfully disambiguated the input and activated only the relevant meaning response.

Simulations

First, we tested the model's ability to simulate—in very simple form—the use of context in natural language processing. Most words in English have more than one meaning; therefore, language processing relies on context provided by prior stimuli to disambiguate current ones. In the model, this occurred by constructing a discourse representation in response to each lexical input, which could then be used as context for processing subsequent stimuli. We tested the model for this ability by first presenting it with a context word (e.g., CHICKEN), allowing activation to spread through the network, and then presenting the ambiguity (e.g., PEN) and observing the output. Note that, in this case, the model was not directly provided with a discourse representation. Rather, it had to construct this from the first input and then use it to disambiguate the second. Tested in this way with all context-word/ambiguous-word pairs (e.g., either CHICKEN or PAPER followed by PEN), the model was able to consistently generate the meaning response appropriate for the context.

To simulate performance in our lexical disambiguation experiment, the model was presented with pairs of context and ambiguous words (representing the clauses used in the experiment) in either order (context first or last). Following each pair, the network was probed with the ambiguous word, simulating the subjects' process of reminding themselves of the ambiguity and choosing its meaning. To simulate the variability observed for human performance in this task, a small amount of noise was added to the activation of every unit in the model at each time step of processing. The amount of noise was adjusted so that the simulation produced an overall error rate comparable to that observed for control subjects in the experiment. The model's response in each trial was considered to be the meaning representation that was most active over the output units after the probe was presented. To simulate schizophrenic performance, we introduced a disturbance identical to the one used in the Stroop and CPT models: a reduction in gain of units in the context (discourse) module to the same level used in the other models (0.6). No other changes were made to the model. The results of this simulation show a strong resemblance to the empirical data (see Figure 10). They demonstrate both significant effects: (a) In the reduced gain mode, the simulation made about as many more dominant response errors as did schizophrenic subjects; however, (b) as with human subjects, this occurred only when context came first—gain did not have a significant effect when context came last. Gain also had little effect on other aspects of performance (e.g., number of unrelated errors) and in other conditions of the task (e.g., when the dominant meaning was correct), which corresponds well to the empirical findings. Thus, the model appears to be specific in its behavior; that is, it demonstrated performance deficits only in those conditions in which schizophrenics showed similar deficits, and not in conditions in which schizophrenic performance was similar to that of normal subjects.

The model provides a clear view of the relationship between dominant response bias, internal representation of context, and a reduction of gain. When gain is reduced in the context module, the representation of context is degraded; as a consequence, it is more susceptible to the cumulative effects of noise. If a contextual representation is used quickly, this effect is less sig-

[4] Recall that the discourse module is connected to the semantic module with two-way connections, so that the discourse module can be used as either an input module or an output module, depending on whether the representation in this module is explicitly specified by the experimenter or is allowed to develop in response to activation it receives from the semantic module.

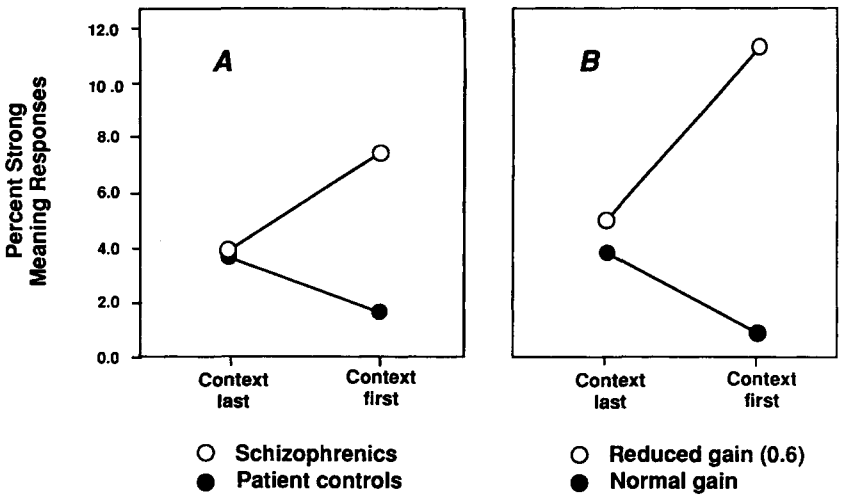

Figure 10. Percentage of strong-meaning responses when the weak meaning was correct for subjects in the empirical study and for the simulation. (The rates of unrelated errors [not shown] and of weak-meaning responses when the strong meaning was correct were the same in the normal and reduced gain conditions of the simulation and were of the same magnitude as those observed for human subjects [about 1%–2%]. The 5% confidence intervals for the simulation data are ±1.1% for context last and ±0.6% for context first in the normal gain condition, and ±1.3% for context last and ±2.0% for context first in the reduced gain condition.)

nificant, and the representation is sufficient to overcome a dominant response bias. However, if time passes (as when context is presented first), the effects of noise accumulate, and the representation is no longer strong enough to reliably mediate the weaker of two competing responses. It is interesting that the cumulative effects of noise are offset by a priming effect when gain is normal in the discourse module. That is, when the internal representation of context is sufficiently robust, then its occurrence before the ambiguity allows it to prime the correct meaning, leading to better performance than when context follows the ambiguity. Note that a trend toward this effect is also observed in control subjects.

General Discussion

We began by reviewing behavioral data concerning deficits of selective attention and language processing in schizophrenia. We also reviewed data which indicate that prefrontal cortex and its dopaminergic supply are important for the construction and maintenance of internal representations of context and that disturbances in these systems are involved in schizophrenia. We then showed how the connectionist framework can be used to relate these findings to one another. We presented three models that (a) simulated quantitative aspects of performance in the Stroop task, the CPT, and a lexical disambiguation task; (b) elucidated the role of the internal representation of context—and its relationship to attention—in these various tasks; (c) related behavior in these tasks to biological processes; and (d) identified a specific disturbance in these processes that could account for schizophrenic patterns of performance. The models touch on a number of important issues concerning cognition in both normal subjects and schizophrenics and the biological processes involved. We discuss these in this section and consider some of the limitations of our models. We then compare our models with others addressing similar issues. We conclude with a discussion of some general issues concerning the modeling endeavor at large.

Attention and the Internal Representation of Context

The Stroop task and the CPT are commonly thought of as measures of selective attention, whereas the lexical disambiguation task is most naturally thought of as a measure of context effects in language processing. Our models suggest, however, that there is a close relationship between selective attention and the internal representation of context. The attentional effects observed in our simulations of the Stroop task and the CPT resulted directly from the influence of context. In both cases, the use of context led to the selection of the appropriate response to an otherwise ambiguous stimulus. Similar processes were at play in the lexical disambiguation task: Here as well context was necessary for the selection of an appropriate response. This similarity between the attentional and language tasks was demonstrated by our ability to simulate performance in these different tasks using the same basic mechanisms for representing and processing context in each case. Thus, the models contribute to our understanding of the cognitive processes involved in these tasks in two important ways: (a) The models suggest that selective attention can be thought of as the influence that context has on the selection of task-appropriate information for processing, and they are explicit about the mechanisms by which this occurs. (b) The similarity of these mechanisms across tasks suggests that, although at the surface these tasks may appear to be very different from one another, they are governed by a common set of information-processing principles. This should not be taken to suggest, however, that

the actual processing pathways are the same for all three tasks. They involve different levels of information processing (from letter recognition to the access of semantic and discourse-level knowledge) and even different domains (color vs. word form recognition in the Stroop task). We have no doubt that the processing pathways involved in each task differ in ways not captured by our current models. However, these differences do not appear to be relevant to the dimensions that distinguish between normal and schizophrenic performance in these tasks. Indeed, it is precisely the simplifications introduced by the models that helped bring the commonalities among the tasks and the parameters relevant to schizophrenia into focus.

Disturbance in the Internal Representation of Context

Viewing selective attention as the effects of context also helps organize several findings in the schizophrenia literature. We were able to show that a single disturbance in the mechanisms underlying the processing of context can account for a number of attention and language deficits in schizophrenics—phenomena that have often been treated separately in the literature. From this, we would predict that performance should correlate across tasks that rely on the internal representation of context.

Previous attempts to examine cross-task correlations of schizophrenic cognitive deficits have produced conflicting results. On the one hand, Kopfstein and Neale (1972) reported only small correlations between five different tasks that were presumed to tap attentional mechanisms (a reaction time task, a size estimation task, the Benjamin proverbs test, the Goldstein-Scheerer object sorting task, and an auditory discrimination task), and Asarnow and MacCrimmon (1978) found that performance on the simple CPT-X did not correlate with performance on the span of apprehension test. On the other hand, Kornetsky and Orzack (1978) did find that the poorer performers on the CPT-X were also more affected by the previous (nonpredictive) preparatory interval on the Shakow reaction time task. The disparate nature of these findings has led investigators to assume that "attention" may not be a general mechanism. That is, the measures used in these studies might have tapped different components of attention, or other information-processing mechanisms altogether. However, little work has ensued to explain, in more specific terms, the nature of and differences between these mechanisms. This may reflect one of the major difficulties that has confronted research in this area: the lack of a theoretical framework within which to compare and select tasks.

Our models offer an approach to this problem, by specifying the mechanisms underlying at least one component of attention (the effects of context) and relating this directly to task performance. In particular, they identify two task dimensions that are relevant to attentional effects and schizophrenic deficits: (a) the relative strength of competing responses and (b) the demands placed on the internal representation of context. The models show that (a) control of dominant response tendencies is sensitive to degradation in the internal representation of context and (b) maintenance of the internal representation of context is also directly affected by its initial degradation. Reducing the gain of units in the context module attenuates context representations; this, in turn, renders them more susceptible to the cumulative effects of noise. Thus, in a noisy system, an impairment in the ability to represent information leads to an impairment in the ability to maintain it. Table 2 categorizes the tasks we have considered along these two dimensions (strength of competing responses and maintenance of the internal representation of context).

Tasks in which subjects need to keep only a set of instructions or a single stimulus in mind place the least demand on maintaining context. That is, when task instructions or a target stimulus remain constant throughout the task, they are reinforced by performance on each trial and therefore are less susceptible to degradation with time. For example, in the CPT-X (detect any occurrence of an X) the subject needs to remember only the target stimulus. Similarly, in the standard Stroop paradigm—in which trials are blocked by task—the instructions remain constant (respond to color or respond to word). These tasks are shown on the left side of Table 2. In contrast, in the CPT-double and CPT-AX subjects must remember the previous stimulus—which changes from trial to trial—in addition to the task instructions, increasing the demand placed on maintenance of the internal representation of context. This is also true of the lexical disambiguation task when context comes first. These are shown on the right side of Table 2. Attentional effects related to the processing of context should be most evident in these tasks.

The second dimension of Table 2 concerns the influence of competing response tendencies. In some tasks, all potential responses are of equivalent strength. For example, subjects are equally familiar with the letters used in standard variants of the CPT; these tasks are shown at the top of Table 2. In other tasks, however, the strength of one response is much greater than the strength of the other. This is due to different amounts of experience either with different features of the stimulus (as in the Stroop task: colors vs. words) or with different responses to the same stimulus (as in the lexical disambiguation task). In our simulations, these differences were captured by differential amounts of training on competing stimulus–response associations. Tasks with response strength asymmetries are shown at the bottom of Table 2. Whereas contextual effects can be observed whenever a stimulus is associated with more than one response, tasks in which competing responses are of unequal strength will be most sensitive to these effects. For example, in the CPT-double, a complete failure to use context would result in performance at chance (because, in the absence of context,

Table 2
Role of Context and Response Strength in Cognitive Tasks

Response strength	Reliance on context	
	Less	More
Equivalent	CPT-X	CPT-AX CPT-double
Asymmetric	Stroop task (interference condition)	Lexical disambiguation task (context first)

Note. CPT = continuous performance test; X = "detect any occurrence of an X"; AX = "detect any occurrence of an X following an A"; double = "detect consecutive reoccurrence of any stimulus."

the competing responses have equal strength). A much stronger effect would be observed in the Stroop and lexical disambiguation tasks: a consistent elicitation of the stronger response even when it is inappropriate. Thus, these latter tasks should provide the most sensitive measure of attentional effects related to the internal representation of context.

Table 2 can help explain the pattern of schizophrenic deficits that have been reported in these various tasks. To the extent that a disturbance in the internal representation of context is involved, we would expect performance to be least affected in tasks at the top and to the left of Table 2. Existing data support at least one implication of this analysis: Schizophrenics show fewer and less reliable deficits in the CPT-X than in the CPT-double or the CPT-AX (Cornblatt et al., 1989; Nuechterlein, 1984). We may also be able to explain one of the failures to correlate across measures of attention. Asarnow and MacCrimmon (1978) found no relationship between performance in the CPT-X and the span of apprehension test. As in the CPT-X, target stimuli in the span of apprehension test remain constant throughout the task; this task belongs with CPT-X in the upper left of Table 2. Because these tasks are the least sensitive to deficits in the internal representation of context, we would expect them to be the least likely to reveal a correlation.

Most important, Table 2 provides a rational approach for the design of new studies to evaluate cross-task correlations. Schizophrenics should show the greatest deficits, and therefore the greatest correlations, when the tasks involve both internal representation of context and a response strength asymmetry (i.e., dominant response tendencies that lead to contextually inappropriate performance). We have begun to provide support for this prediction with the results of our lexical disambiguation task. However, it should be possible to demonstrate increased sensitivity to schizophrenic deficits in each of the other tasks—and correlations among them—by varying them along the appropriate dimensions of Table 2. For example, response strength asymmetry could be introduced into the CPT by varying the frequency with which targets appear. Conversely, demand placed on the internal representation of context could be increased in the Stroop task by varying the instructions from trial to trial and presenting the stimuli at various delays following the instructions. These task manipulations should increase both their sensitivity to schizophrenic deficits and the likelihood of detecting cross-task correlations.

We should be clear, of course, that a disturbance in the internal representation of context may be only one of several disturbances underlying schizophrenic cognition. Indeed, we have focused on a circumscribed set of experimental findings in this article. However, as we have suggested, these may represent the cognitive correlates of the negative symptoms of schizophrenia —an important clinical component of this illness. It is worth noting that negative symptoms involve a number of affective disturbances (e.g., flattening of affect, emotional withdrawal, and motivational difficulties). In their present form, the models do not provide an explanation of the relationship between cognitive and affective disturbances. However, in specifying a set of relevant cognitive variables they provide an anchor point for the design of studies in which such relationships can be explored in the future.

A related issue concerns the apparent discrepancy between the severity of schizophrenic symptoms—even when manifestations of the illness are limited to negative symptoms—and the comparatively mild impairment that is observed in cognitive tasks. Because of this discrepancy, it is sometimes suggested that the cognitive deficits observed in laboratory tasks may be epiphenomenal and not central to the disorder. However, we believe that a disturbance in the internal representation of context may be directly responsible for at least some of the clinical features of schizophrenia. For example, normal functioning in social situations requires the constant monitoring and use of environmental cues as determinants for appropriate behavior. Impairment in social functioning is one of the diagnostic criteria for schizophrenia (American Psychiatric Association, 1987), which is often characterized by the expression of impulses in socially inappropriate contexts. The reason for this may be the failure to use social and environmental cues as context for (a) the selection of socially appropriate behavior and (b) the suppression of more compelling but socially inappropriate behavior. Given the multiplicity, complexity, and often subtlety of contextual cues that must be processed in social circumstances, and the variety of competing responses from which to select, social functioning may place significantly greater demands on the ability to use context to control behavior than do cognitive tasks studied in the laboratory. A test of these ideas provides another important challenge for future research.

At present, our models fail to address the positive symptoms of schizophrenia, such as hallucinations and delusions. It is possible, however, that the mechanisms we have discussed may be relevant to the generation of these symptoms when applied to other brain areas (e.g., temporal cortex or limbic structures). For example, an *increase* in the gain parameter (corresponding to an increase in dopaminergic activity) results in active and contrasted patterns of activation on the output layer of a network regardless of the strength (i.e., degree of activation) of the input. When such output patterns are produced in the absence of any meaningful input (e.g., in response to noise) the network might be considered to display misperceptions or misinterpretations that resemble the phenomena of hallucinations and delusions. A related argument has been offered by Hoffman (1987), which we will discuss later. The possibility that schizophrenia may be associated with an increase in gain in some brain areas at the same time as a decrease in gain in frontal cortex is also considered further on.

As we have indicated at several points, it is important to distinguish the hypothesis that schizophrenic performance deficits arise from a disturbance in internal representation of context from the possibility that these simply reflect an impairment in short-term memory. Earlier, we referred to findings that schizophrenics perform normally on measures of short-term memory (Larsen & Fromholt, 1976; Oltmanns & Neale, 1975) and that in some circumstances they show awareness of information that they are unable to use to respond appropriately (Goldberg et al., 1987).[5] We have proposed a distinction

[5] We should note that schizophrenic deficits have been observed in tasks purported to measure short-term memory, but only when the number of stimuli to be recalled exceeds the usual 5–7-item span of short-term memory (Calev, 1983; Gjerde, 1983; Koh, 1978). These so-

between memory for the identity of a prior stimulus (short-term memory) and the active maintenance of information in a form that allows the stimulus to control other kinds of responses (internal representation of context). Schizophrenics demonstrate a dissociation in these two abilities, which suggests that distinct mechanisms may underlie these two functions. Our models are an attempt to specify the mechanism underlying the processing of context and to show how a specific disturbance in this mechanism can account for the kinds of performance deficits demonstrated by schizophrenics.

Finally, it is important to consider the rationale behind the actual manipulation used to induce a disturbance in the processing of context within the models. This disturbance was introduced by reducing the gain of units in the context module. In principle, there are other parameters that could be used to generate similar results (e.g., a reduction in the number of units in the context module, accelerated decay of activation, or an increase in the level of noise in this module). The particular manipulation we used—reducing gain—was motivated by the biological factors thought to underlie schizophrenic deficits and their interpretation in computational terms (Servan-Schreiber et al., 1990). We discuss this more extensively later, but we wish to emphasize here, however, that it is precisely this parameter (gain) that provides a point of convergence between behavioral and biological data. That is, the role of the models is not only to account for behavioral and biological phenomena in their own right but also to explore how these two different sets of data can be accounted for using the same set of information-processing mechanisms. We have identified a parameter—gain—that satisfies constraints provided by empirical data at both levels concerning schizophrenic deficits. Whether there are other parameters that can provide a similar integration of behavioral and biological data is a question for future studies.

Generalized Versus Specific Deficits

At several points we have referred to a frequently raised issue in schizophrenia research: To what extent does a particular set of findings reflect a generalized deficit as opposed to a deficit in a specific component of processing? For example, the widespread finding of an increase in reaction time is typically considered to reflect a generalized deficit. However, it is difficult to know the meaning of this hypothesis without defining it in specific information-processing terms. Widlöcher and Hardy-Bayle (1989) discussed the importance and difficulty of conceptualizing the generalized deficits observed in schizophrenia in terms of specific mechanisms. As they emphasized, a generalized deficit must still reflect a disturbance of some kind, somewhere in the system. Our model of the Stroop effect allowed us to specify one type of generalized deficit (a slowing of the processing rate throughout the system) and to compare this with a hypothesis concerning a more specific deficit (a reduction of gain in the module responsible for representing context). The model provided a framework not only in which to make these hypotheses explicit but also in which to compare their ability to quantitatively fit the data. In the case of these two hypotheses, our findings favored the more specific deficit. Furthermore, they showed how this specific deficit could have generalized consequences. This is a reconceptualization of the data: What appeared to be a general effect (overall slowing of response) could be attributed to a circumscribed disturbance in the attentional mechanism. Although an overall increase in reaction time for schizophrenics in other tasks may well be due to more general deficits, the Stroop model showed that this need not always be the case.

Conversely, the simulations showed that an impairment of representations in the context module does not necessarily affect all aspects of performance. Most important, a decrease in gain did not affect performance in conditions where schizophrenics perform as well as normals (e.g., in the strong-meaning condition of the lexical disambiguation task). By extension, we would not expect a decrease in gain to have an effect in simulations of tasks that are performed equally well by schizophrenics and normal controls. In that sense, an impairment of the internal representation of context induces specific, predictable deficits. It does not amount to a degradation of all aspects of processing irrespective of task.

Prefrontal Cortex, Dopamine, and the Internal Representation of Context

Perhaps the most common characterization of frontal lobe function concerns its involvement in goal-oriented activity: the planning and sequencing of complex actions (e.g., Bianchi, 1922; Duncan, 1986; Luria, 1966; Shallice, 1982). The construction and maintenance of internal representations of context can be seen as an important component of these facilities. The actions associated with a particular goal may, in other contexts, be relatively infrequent or "weak" behaviors. Such actions require the maintenance of an internal representation of the goal, or of goal-related knowledge, to favor their execution and to suppress competing, possibly more compelling, behaviors. For example, we have all struggled with the urge to scratch a mosquito bite. Resisting this urge relies on actively accessing the knowledge that if the bite is left alone it will heal more quickly. This knowledge can be thought of as the context needed to control behavior, and it must be maintained or the prepotent response tendency (scratching the bite) will prevail (e.g., as it might during sleep or while absorbed in another activity).

Fuster (1980, 1985a, 1985b) and Goldman-Rakic (1987) provided direct evidence concerning the role of prefrontal cortex in relating information over space and time. Diamond (1990a) emphasized the importance of the role that prefrontal cortex plays in inhibiting "prepotent" (dominant) response tendencies. Our models show how these information-processing functions may be implemented in biologically plausible mechanisms and how they may be modulated by dopaminergic activity. For example, in our simulation of the lexical disambiguation task, the discourse module supported a representation that was built up in the course of processing. This allowed the model to process later elements of the sentence in the context of ones it had seen earlier. The models also showed how such contextual information permitted the expression of a weaker response in the pres-

called "supra-span" tasks typically involve specific encoding strategies such as chunking or elaborative rehearsal that may well rely on internal representations of context and not just short-term memory.

ence of a stronger (more dominant) one. Thus, we were able to account for two important functions that have been attributed to prefrontal cortex in terms of a specific component in our models. Moreover, they suggested an explicit mechanism for dopaminergic effects in prefrontal cortex. By maintaining or increasing the gain of neurons in this area, dopamine may help augment contextual representations against a background of noise. This, in turn, would lead to better preservation of contextual information over time and more effective control over dominant response tendencies.

Because the models make the relation between dopamine, prefrontal cortex, and the internal representation of context explicit, predictions can be made about the interplay between these factors that can be tested in empirical studies. Here we provide three examples of such predictions: how dopamine agonists may affect prefrontal activity, the relation between prefrontal activity and the processing of context in the CPT, and the relation between prefrontal cortex function and language processing.

Dopamine Agonists and Prefrontal Metabolic Activity

Mesocortical projections form a major component of the dopamine system. From this, it might be expected that dopamine agonists would have the general effect, in normal subjects, of increasing metabolic activity in prefrontal cortex. Our models make a somewhat different prediction: To the extent that a task does not rely heavily on the maintenance of an internal representation of context and involves a set of routine responses, we would predict that the administration of dopamine agonists would *not* have any effect on the metabolism of prefrontal cortex during the performance of such a task. This is because our models specify that the effect of dopamine release is to potentiate the response of target cells to afferent signals. In tasks that do not rely heavily on the internal representation of context, we assume that there are fewer signals arriving in prefrontal areas and therefore that the activity of units in these areas should be relatively unchanged. However, during performance of a task that does rely on the internal representation of context, the effect of dopamine agonists should be to substantially increase metabolism in prefrontal cortex. Thus, we predict an interaction between task type and drug condition. Our predictions have received preliminary support from data reported by Geraud et al. (1987), who found that the prefrontal activity of normal subjects at rest was not increased by the dopamine agonist Pirebdil. However, the second part of our prediction has not yet been tested: that an agent such as Pirebdil would increase prefrontal activity in the same subjects during a task requiring internal representation of context.

Prefrontal Activity During the CPT

We have argued that schizophrenic deficits on the CPT can be attributed to frontal lobe dysfunction. In some reports, however, indices of metabolic activity in the prefrontal cortex during performance on the CPT have failed to differentiate between schizophrenic subjects and controls. Berman et al. (1986) reported the absence of any correlation between CPT performance and prefrontal activation in either schizophrenics or normal controls; neither group showed significant prefrontal enhancement during the task. These results are in conflict with the findings of R. M. Cohen et al. (1987; R. M. Cohen, Semple, Gross, Holcomb, et al., 1988), who reported a significant correlation between prefrontal metabolism and CPT performance. The analysis of task dimensions relevant to CPT performance that we presented earlier may provide a reconciliation of these findings. Berman et al. used two variants of the CPT: the simple CPT-X, which makes fewer demands on the internal representation of context, and a version of the CPT-AX with interstimulus intervals (ISIs) of 0.8 s or less. At such short ISIs, the association between A and X can be encoded through direct reinforcement; we assume that reinforcement learning does not rely on prefrontal areas (see earlier section, *Prefrontal Cortex and the Internal Representation of Context*). In fact, as subjects' performance improved, Berman et al. attempted to increase the difficulty of the CPT-AX by reducing the ISI even further. Such an increase in the event rate of the task has been shown to impair performance in normal subjects (Parasuraman, 1979). However, according to our analysis, this increase in difficulty is unrelated to the specific difficulty that the CPT presents for schizophrenics. Rather, it is when the duration between the contextual cue (here the letter A) and the potential target (X) is increased that we would expect schizophrenics to show difficulty with the task, that is, as their internal representation of context fails to reliably bridge the gap between the two stimuli. These observations may explain why no specific increase in prefrontal metabolic activity was observed during performance of the two variants of the CPT used by Berman et al.

In contrast, R. M. Cohen et al. used an auditory CPT in which subjects were asked to detect the softest of three tones of equal frequency. Tones were presented at 2-s intervals. In this task, the target can be identified only in reference to the other stimuli, which provide the context for deciding whether or not to respond. Because the subjects need to integrate over several previous trials (at least two) in order to make the relevant comparison, and because of the longer ISI (2 s), this task places greater demands on the internal representation of context than does the version used by Berman et al. (1986). It is not surprising, therefore, that in this study a correlation was found between prefrontal activity and CPT performance. On the basis of these arguments, we can make the following prediction: The ability of standard versions of the CPT-AX and the CPT-double to differentiate schizophrenic from normal subjects should depend on ISI. When the ISI is 1 s or less, schizophrenic performance should not be dramatically impaired compared to that of normal subjects. However, at longer ISIs (e.g., 5 s) normal subjects should do better (because the event rate goes down), whereas schizophrenics' performance should be degraded (because maintenance of the internal representation of context is now required). Moreover, in normal subjects, CPT performance may not correlate with prefrontal metabolic activity at short ISIs (as Berman et al. found), but it should correlate during blocks of trials at longer ISIs.

Prefrontal Cortex and Language Performance

Finally, our models suggest that prefrontal cortex plays a specific and important role in language processing. This has sev-

eral implications. First, it suggests that other disorders that involve prefrontal cortex (e.g., neurologic patients with lesions of this area) may show language deficits of the sort we have described (we temper our claim here because neurological patients may show modularized deficits; see discussion further on under *Specificity of Frontal Deficits to Schizophrenia*). It also suggests that prefrontal metabolic activity should correlate with performance on language tasks that rely heavily on the internal representation of context. This represents an exciting area for future research.

Before concluding this section, we should point out that our models have not yet been directly applied to the A $\bar{B}$ task or the Wisconsin Card Sort Test, both of which have been traditionally associated with frontal lobe function and, in the case of the latter, schizophrenic deficits. These tasks (especially the WCST) involve processes of problem solving and hypothesis testing that are not captured by our models in their present form. Nevertheless, our models suggest an interpretation of frontal deficits on these tasks that could, in principle, be captured in a simulation model (see for example Dehaene & Changeux, 1991). As we noted earlier, efficient performance in both of these tasks requires that subjects overcome the tendency to repeat response patterns that were correct on previous trials. Thus, both demand that context (e.g., placement of the object on the current trial) be used to control a response tendency (return to prior location) that has gained strength over the course of previous trials. Failure to do so would result in the patterns of perseveration observed (A $\bar{B}$ errors, or failure to switch sorting principle in the WCST). The difference between these tasks and the tasks with response strength asymmetries that we have simulated (Stroop and lexical disambiguation tasks) is that, in the A $\bar{B}$ and the WCST, response strength asymmetries develop within the task, rather than existing a priori. If, however, training of the response pathways was allowed to occur during task performance, then experience on previous trials could lead to the development of response strength asymmetries that could then compete with recent contextual information to determine the response. In fact, Dehaene and Changeux (1989) have proposed a network model of behavioral and electrophysiological data in delayed response tasks, including the A $\bar{B}$ task. This model exhibits principles that are similar to what we have just described. In their model, a low-level associational module is responsible for mapping stimuli onto responses and is subject to training on each experimental trial. A higher level module—which can memorize task conditions or perform rule induction—selects or modulates actions performed by the lower level. In this model, the higher level module is assumed to perform the function of the prefrontal cortex. A $\bar{B}$-type errors arise when this module is impaired, and responses are governed to a greater degree by the training experience of the low-level association module. The similarities between our models and the ones these authors have described—developed independently and with regard to different empirical phenomena—lend strong support to the generality of the principles involved.

Gain as a Model of Dopamine Effects in the Prefrontal Cortex

As we discussed earlier, most evidence for a potentiating, or gainlike, effect of dopamine comes from studies of striatal cells. Studies of dopamine in prefrontal cortex have largely indicated inhibitory effects. However, concentrations were used that also produced inhibitory effects in striatal cells. Studies using lower concentrations—at which potentiating effects are observed in the striatum—have not yet been performed. In this regard, our models offer a prediction. By assuming—in our simulations—that dopamine has the same effects in prefrontal cortex as it does in the striatum, we are able to account for a variety of behavioral findings associated with schizophrenia. The success of our models at the behavioral level may be taken as a prediction concerning the validity of the biological assumptions upon which they are based. That is, our models predict that dopamine should have the same gainlike effect in prefrontal cortex that has been observed for cells in the striatum. In this respect, the models provide theoretical guidance for further studies at the physiological level.

In addition, they offer a new interpretation of some existing data and suggest directions for future simulation work. Recently, structural characteristics of dopamine synapses on pyramidal cells in prefrontal cortex and the striatum have been described which suggest that they are predominantly inhibitory (e.g., Freund, Powell, & Smith, 1984; Goldman-Rakic, Leranth, Wallns, Mons, & Geffard, 1989; Pickel, Beckley, Joh, & Reis, 1981). This seems to contradict the potentiating effects of dopamine that are known to occur (at least within the striatum). However, Goldman-Rakic et al. also observed inhibitory dopaminergic synapses on gamma-aminobutyric acid inhibitory interneurons in the same region. This suggests the possibility that potentiating effects might arise from an ensemble of local computations. For example, if the interneurons that receive inhibitory dopaminergic synapses turn out to be the same ones that themselves inhibit target pyramidal cells, then dopamine release would not only directly inhibit target cells but would also inhibit their inhibitory interneurons. That is, dopamine release might result in both inhibition (directly) as well as disinhibition (indirectly) of pyramidal cells. The dual influence of these effects could result in a disregulation, or potentiation, of both inhibitory and excitatory inputs.[6] In our current models, potentiating effects were simulated as a direct change in a processing parameter (gain) of individual units. This is because, so far, we have concerned ourselves primarily with the functional consequences of potentiating effects and less with the specific details of the neurophysiological mechanisms from which they arise. However, in future work it may be worthwhile to implement potentiation as the interaction of local circuit effects and to explore what implications this has at the behavioral level.

Finally, we should clarify our rationale for selecting the particular value of gain that we used in the simulations. We do not know of any quantitative estimates of dopamine deficits in the prefrontal cortex of schizophrenics. Therefore, we simply picked a value for the gain parameter that produced performance deficits comparable to those of schizophrenics in the three different tasks. The actual value that we picked (0.6) provided the best simultaneous fit (i.e., using a single value across simulations) to the data from the three different tasks.

---

[6] We would like to thank David Lewis for suggesting this possibility.

Biological Disturbances in Schizophrenia

Is Dopamine Increased or Decreased in Schizophrenia?

We have argued that certain cognitive deficits in schizophrenics can be explained by a reduction of dopamine activity in prefrontal cortex. This may seem to be at odds with what is known about the effects of antipsychotic (neuroleptic) medications. As we discussed earlier, neuroleptics that tend to improve thought disorder also improve performance on cognitive tasks. For example, performance on the CPT improves with long-term neuroleptic therapy (Spohn et al., 1977), and R. M. Cohen, Semple, Gross, Nordahl, et al. (1988) showed that the correlation between prefrontal activity and CPT performance was restored in schizophrenic subjects treated with antipsychotic medications. Yet neuroleptic medications are commonly thought to reduce dopamine activity by blocking its postsynaptic effects (e.g., Snyder, Banerjee, Yamamura, & Greenberg, 1974). This seems to challenge our hypothesis: How could cognitive deficits attributed to a reduction of dopamine activity be ameliorated by antidopaminergic agents? Evidence gathered over the past decade suggests some possible answers to this question.

First, evidence from animal studies suggests the existence of complex feedback relations between cortical and subcortical dopamine activity. For example, Pycock, Kerwin, and Carter (1980) showed that a selective destruction of dopamine afferents in rat frontal cortex results in a state of chronic dopamine hyperactivity in subcortical areas. This indicates that states of dopamine hypo- and hyperactivity can coexist in different brain regions and may even be causally related (see also, Deutch, Clark, & Roth, 1990; Grace, 1991; Tassin, Simon, Glowinski, & Bockaert, 1982). It is tempting to use this finding to explain the coexistence of positive (e.g., hallucinations, delusions) and negative symptoms of schizophrenia and to reconcile theories postulating increased or decreased dopamine activity in schizophrenia (e.g., see Weinberger, 1987, and Grace, 1991).

Second, studies of the effect of neuroleptics on dopamine synthesis have suggested that the mesolimbic and mesocortical dopamine systems respond differently to chronic administration of these medications (for reviews see Bannon, Freeman, Chiodo, Bunney, & Roth, 1987, and Grace, 1991). These have shown—in rodents, primates, and humans—that tolerance to activation of synthesis develops rapidly in the striatal and limbic areas, whereas it develops slowly and remains limited in prefrontal cortex (Bacopoulos, Spokes, Bird, & Roth, 1979; Roth, Bacopoulos, Bustos, & Redmond, 1980; Scatton, 1977; Scatton, Boireau, Garret, Glowinski, & Julou, 1977). Moreover, during chronic administration of neuroleptics, most dopamine cells enter a state of depolarization inactivation. However, a small number of cells remain active, and the majority of these have been identified as mesocortical cells projecting to prefrontal cortex (Chiodo & Bunney, 1983). Overall, these data suggest that dopamine tone in prefrontal areas is less affected by neuroleptics than are limbic and striatal dopamine. The net result of neuroleptic administration might actually be to enhance dopamine activity in the prefrontal cortex, at least relative to its activity in other brain regions. This would lead us to expect that neuroleptics would, at worst, have no influence on the cognitive deficits we have addressed and, at best, lead to improvements by strengthening context effects relative to more automatic (e.g., subcortically mediated) responses.

Specificity of Frontal Deficits to Schizophrenia

We have argued that a reduction of dopamine in prefrontal cortex reduces the dynamic range of units in this area, and we have begun to explore the behavioral effects of this disturbance. However, we have not yet compared these effects to those that might result from other possible disturbances, such as the actual loss of units that might result from neurologic damage. One factor that distinguishes schizophrenia from neurological damage to the frontal lobes, and that is brought into focus by our models, is the diffuse neuromodulatory nature of the lesion in schizophrenia. Assuming that representations within the prefrontal cortex are somehow modularized (e.g., by modality or by level of analysis), then we would expect to see different behavioral profiles for patients with focal lesions in different regions of prefrontal cortex. In contrast, the diffuse disturbances resulting from decreased dopaminergic tone in schizophrenia should result in more consistent patterns of deficit across subjects. This idea is supported by the results of Stroop task performance in these two groups. Studies of schizophrenics using this task have produced reasonably consistent results (see Table 1). In contrast, studies of frontal lobe patients have yielded conflicting results (e.g., Perrett, 1974, vs. Shallice, 1982), which may be due to heterogeneity in the specific sites lesioned. In contrast, other disturbances of dopamine neuromodulation—to the extent that they affect prefrontal cortex—should be associated with the kinds of cognitive deficits we have discussed. As we noted earlier, performance of Parkinson's patients on the WCST provides preliminary support for this claim. Additional studies along these lines are clearly warranted.

Other Biological Systems

There is little doubt that disturbances of systems other than prefrontal cortex and the mesocortical dopamine system are involved in schizophrenia. Other brain regions have been implicated, such as the hippocampus (e.g., Conrad, Abebe, & Austin, 1991; Kovelman & Scheibel, 1984) and various subcortical structures including the thalamus (e.g., Crosson & Hughes, 1987), the globus pallidus (e.g., Early, Reiman, Raichle, & Spitznagel, 1987), and the basal ganglia (Stevens, 1973), as have neurotransmitters other than dopamine, such as norepinephrine (e.g., Lake et al., 1980; van Kammen et al., 1989) and serotonin (e.g., Geyer & Braff, 1987). Indeed, the involvement of other systems may account for the significant differences between schizophrenia and other illnesses that affect dopamine, such as Parkinson's disease. In their present form, our models are limited in the scope of biological systems that they address. Nevertheless, by delineating aspects of the physiology and behavioral consequences of reduced dopaminergic tone in prefrontal cortex, they help refine our understanding of this component of schizophrenia. For example, applications of the models to other tasks may be useful in identifying aspects of behavior that will and will not be affected by variations in dopaminergic activity in prefrontal cortex. Furthermore, our model of the neuromodulatory effects of dopamine provides a starting point for

exploring the involvement of dopaminergic disturbances in other areas, such as the effects of an increase in dopamine activity that has been hypothesized for limbic and subcortical areas. By developing models that correspond to the functions of these brain areas (e.g., declarative memory, perception, and motor control), it may be possible to account for symptoms of schizophrenia not addressed in this article. Finally, by providing an example of how important features of biological processes can be captured within the connectionist framework, and how these can be related to specific behavioral phenomena, the models provide a framework for exploring the role of neuromodulatory systems in other forms of illness and their relation to information-processing behavior.

Comparison With Other Models of Schizophrenia

An overwhelming number of theories have been proposed to account for the various cognitive and biological abnormalities observed in schizophrenia. Here, we focus on those that are most directly related to our own—either by methodology or claims—and that help delineate the specific contributions of our approach.

Broadbent's Attentional Filter and Its Breakdown in Schizophrenia

Perhaps the most common theory of cognitive dysfunction in schizophrenia draws on the filter model of selective attention first proposed by Broadbent (1958, 1971). According to this model, multiple stimuli are registered by the sensory organs and enter a short-term store. At this point, stimuli are passed through a filter that provides access to a limited-capacity channel in which further processing takes place. The filter is set by past experience (e.g., conditional probabilities based on past events) and feedback provided by processing in the limited-capacity channel. Investigators who have focused on the phenomenology of schizophrenia (e.g., Garmezy, 1977; Lang & Buss, 1965; McGhie, 1970; McGhie & Chapman, 1961) have suggested that patients experience a difficulty in screening out irrelevant stimuli and that this may be due to a breakdown in the filtering mechanism. Schizophrenics would thus experience one of two states: either a state of stimulus overload in which all stimuli gain equal access to the limited-capacity channel, or a shutdown of information intake in which all stimuli are equally blocked from accessing that channel.

Our models relate to this conception in several ways. First, the models provide an explicit set of mechanisms for stimulus selection and access to response systems. However, there is no dedicated filter in these models. Rather, a filterlike effect emerges from the interaction of stimulus processing with processing of context when both are channelled through a common layer of intermediate (associative) units. The models suggest how this filtering of incoming information may be implemented in neural structures. Furthermore, they identify ways in which catecholaminergic systems may influence the selection of information and the consequences of their disruption. Specifically, the models demonstrate that weakening the top-down source of stimulus selection (i.e., the internal representation of context) by reducing dopaminergic tone to the prefrontal cortex would not result in a complete disorganization of stimulus processing. Rather, degradation would follow a distinctive pattern, in which stronger responses begin to dominate weaker ones, with decreasing sensitivity to context provided by the task.

Joseph, Frith, and Waddington (1979)

These authors described a mathematical network model that, like our own, relates neural function to higher cognition. They focused on neurotransmitter interactions presumed to support attentional functions. Their model assumes that the dopamine system itself acts as a filter for external inputs and shows how excessive dopamine activity results in exhaustion of inhibitory mechanisms and ultimately a breakdown of filtering functions. This model demonstrates how a neural network can be constructed which performs a filtering function on the basis of simple excitatory and inhibitory interactions. However, Joseph et al. did not relate disturbances of this filtering mechanism to schizophrenic performance in specific behavioral tasks. Because of this, it is difficult to evaluate their model's ability to explain quantitative aspects of cognitive performance.

Hoffman (1987)

Hoffman described a set of computer simulations that display behaviors considered to be analogous to several of the positive symptoms of schizophrenia (loosening of associations, blocking, and hallucinations). The simulations used fully interconnected Hopfield-type networks as a model of human associative memory processes. During the training phase, the network was taught a set of associations. In the test phase, an input state was specified by activating a subset of the processing units. The network was then allowed to cycle until it settled into a stable configuration of activations. This end-state represented the memory that was accessed from the input specification; this was based on the pattern of connections between units that was learned during training. Hoffman showed that when such a network was forced to encode an excessive number of associations ("memory overload"), specific disturbances of processing occurred: The system often settled into memory states that were inappropriate given the input (hallucinations) or into states that did not correspond to any of the previously encoded associations (loosening of associations). Thus, these simulations related the positive symptoms in schizophrenia to a specific disturbance in the computational mechanisms of the model. Hoffman suggested that this disturbance—memory overload—may arise in schizophrenics as a consequence of a reduced neuronal mass in the prefrontal cortex. In this respect, Hoffman's model can be considered complementary to those we have presented, addressing a different set of symptoms and pathophysiological processes relevant to schizophrenia. However, as with the model suggested by Joseph et al. (1979), these models have not yet been applied to the simulation of quantitative aspects of behavior. This may be due, in part, to the complex and often inaccessible nature of positive symptoms, which pose serious difficulties for quantification. Indeed, this remains a challenge for all approaches to research on the full range of symptoms characteristic of schizophrenia.

Goldman-Rakic (1987), Levin (1984), Weinberger (1987), and Weinberger and Berman (1988)

These authors have emphasized the possible role of frontal lobe dysfunction in schizophrenia. Weinberger's laboratory in particular has provided important empirical support for this hypothesis. Furthermore, they have specifically suggested that a deficit in the dopaminergic innervation of the prefrontal cortex is responsible for performance impairments in tasks such as the WCST and, from a clinical perspective, for the negative symptoms of schizophrenia. Our models concur with these hypotheses and extend them by proposing a specific set of mechanisms that explain the relationship between a disturbance in dopamine activity, the function of prefrontal cortex, and task performance. This has allowed us to address quantitative aspects of performance in a number of behavioral tasks and to provide a unified account of schizophrenic performance in different tasks in terms of a common underlying deficit. Although we have not yet applied our models to performance on the WCST, we discussed how schizophrenic deficits on this task could be related to a disturbance in the processing of context, and we suggested how the models could be extended to test this idea. We would add, however, that the WCST is a complex task (involving learning, the interpretation of social cues, and the choice of strategies for hypothesis testing and problem solving). Although data from this task have helped break ground concerning the relationship between prefrontal cortex and schizophrenia, we feel that future progress in this area will require complementing this approach with new strategies. In particular, it is crucial that we begin to use simpler tasks that lend themselves more readily to cognitive analysis in terms of specific processing mechanisms so that we can begin to test these mechanisms in adequately controlled ways. We have made an effort to show how simulation models can help guide us in this direction.

The Role of Computational Modeling

A common objection to computer models of behavioral data is that such models have so many parameters that they can always be optimized to fit the data. For example, the number of units, the amount of training on different pathways, or the gain parameter could presumably be adjusted in each model to produce better fits. Such fits would not serve a useful purpose because they would be a reflection not of the correspondence between observed data and a mechanism of interest but only of the skill of the modeler.

In the simulations discussed here, the different parameters were indeed adjusted separately for each model. However, we did this only when attempting to fit the behavior of normal (or control) subjects. Once these parameters were determined, they were fixed; only the gain parameter was changed to simulate the behavior of schizophrenics. Furthermore, the gain parameter was reduced by exactly the same amount and in the same functional location in all three models. Thus, although it is true that the models were optimized to fit the normal data for each task, the fit to schizophrenic data was obtained by affecting only the variable (gain) motivated by our theory of schizophrenic deficits.

A related question is often asked: How do models contribute to an understanding of the data they simulate? After all, the data already exist, and the principles or ideas captured by a model can often be expressed more simply without the use of a computer program (indeed, one might contend that this must be so if the ideas are of any general value). McClelland (1988) provided an articulate reply to this question in describing the relevance of models to empirical investigations in psychology. He pointed out that models can (a) bring seemingly disparate empirical phenomena together under a single explanation, (b) provide new interpretations of existing findings, (c) reconcile contradictory evidence, and (d) lead to new predictions. Throughout the present discussion, we have tried to show how our models realize these different goals. For example, by showing that a disturbance in the internal representation of context can explain impairments of selective attention, language processing, and overall reaction time in schizophrenia, our models bring these seemingly disparate phenomena together under a single, unifying explanation. By revealing that an overall increase in reaction time could arise from a specific rather than a generalized information-processing deficit, they provide a reinterpretation of the data. They suggest a reconciliation of contradictory findings with respect to the CPT and prefrontal activation. And they lead to new predictions concerning normal and schizophrenic performance on behavioral tasks, as well as to predictions about dopamine effects on prefrontal metabolism. McClelland also emphasized the role that models play in formalizing theoretical concepts. By committing a set of ideas to a computer program and examining their ability to account for quantitative data, the ideas are put to a rigorous test of both their internal coherence and the resolution of their explanatory power.

Most important, however, is the role that modeling plays in the discovery process. At times the insights provided by a model may seem, in hindsight, to be obvious or not to have required the effort involved in constructing a computer simulation. On other occasions, one may be concerned with the possible circularity of a theory based on a model that has presumably been designed with the theory in mind. Usually, however, such perceptions fail to recognize that the insight and the emerging theory came from the process of developing the model itself. The three models described in this article were actually developed independently and for different purposes. The Stroop model was developed to account for normal performance in this task (J. D. Cohen et al., 1990); the CPT simulation was developed to explore gain as a model of catecholaminergic effects on behavior (Servan-Schreiber et al., 1990); and the language model was inspired by our work on the processing of ambiguous stimuli in recurrent networks (Cleeremans, Servan-Schreiber, & McClelland, 1989; Servan-Schreiber, Cleeremans, & McClelland, 1991). It was only when we compared the mechanisms at work in these different models that we realized how all relied on common principles of processing. This suggested a hypothesis about the relationship between biological and behavioral factors in schizophrenia. In this way, the models provided an important vehicle for the discovery—and not just the testing—of new ideas.

Conclusion

We have tried to show how the connectionist framework can be brought to bear on the relationship between some of the

biological and cognitive disturbances characteristic of schizophrenia. The models we have presented suggest that a common information-processing deficit underlies impaired performance in selective attention and language-processing tasks and relate this deficit to decreased dopaminergic activity in prefrontal cortex. The models, and the simulations based on them, rely on many simplifying assumptions and provide, at best, a coarse approximation of the mechanisms underlying both normal and schizophrenic behavior. Although accounting for empirical data is a primary goal in the development of computer simulation models, McClelland (1988) argued that this may not be the only basis for their evaluation. Models are useful if they offer new interpretations of empirical phenomena, unify previously unrelated observations, reconcile conflicting findings, and predict new empirical facts. We have indicated how our models—simple as they are—may fulfill these different functions. In so doing, we hope that these models will help provide a more refined and integrated approach to the riddle of behavioral and biological disturbances in schizophrenia.

References

Abramczyk, R. R., Jordan, D. E., & Hegel, M. (1983). "Reverse" Stroop effect in the performance of schizophrenics. *Perceptual and Motor Skills, 56,* 99–106.

American Psychiatric Association. (1987). *Diagnostic and statistical manual of mental disorders* (3rd ed., rev). Washington, DC: Author.

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

Andreasen, N. C., Flaum, M., Swayze, V. W., Tyrrell, G., & Arndt, S. (1990). Positive and negative symptoms in schizophrenia. *Archives of General Psychiatry, 47,* 615–621.

Andreasen, N. C., Nasrallah, H. A., Dunn, V., Olson, S. C., Grove, W. M., Ehrhardt, J. C., Coffman, J. A., & Crossett, J. H. W. (1986). Structural abnormalities in the frontal system of schizophrenia. *Archives of General Psychiatry, 43,* 136–144.

Angrist, B., Peselow, E., Rubinstein, M., Wolkin, A., & Rotrosen, J. (1985). Amphetamine response and relapse risk after depot neuroleptic discontinuation. *Psychopharmacology, 85,* 277–283.

Aou, S., Nishino, H., Inokuchi, A., & Mizuno, Y. (1983). Influence of catecholamines on reward-related neuronal activity in monkey orbitofrontal cortex. *Brain Research, 267,* 165–170.

Asarnow, R. F., & MacCrimmon, D. J. (1978). Residual performance deficit in clinically remitted schizophrenics: A marker of schizophrenia? *Journal of Abnormal Psychology, 87,* 597–608.

Bacopoulos, N. G., Spokes, E. G., Bird, E. D., & Roth, R. H. (1979). Antipsychotic drug action in schizophrenic patients: Effects on cortical dopamine metabolism after long-term treatment. *Science, 205,* 1405–1407.

Baillargeon, R. (1990, August). *Young infants' physical knowledge.* Paper presented at the 98th Annual Convention of the American Psychological Association, Boston, MA.

Bannon, M. H., Freeman, A. S., Chiodo, L. A., Bunney, B. S., & Roth, R. H. (1987). The electrophysiological and biochemical pharmacology of the mesolimbic and mesocortical dopamine neurons. In L. I. Iversen, S. D. Iversen, & S. H. Snyder (Eds.), *Handbook of psychopharmacology* (Vol. 19, pp. 329–374). New York: Plenum Press.

Barone, P., & Joseph, J. P. (1989a). Prefrontal cortex and spatial sequencing in macaque monkey. *Experimental Brain Research, 78,* 447–464.

Barone, P., & Joseph, J. P. (1989b). Role of the dorsolateral prefrontal cortex in organizing visually guided behavior. *Brain Behavior and Evolution, 33,* 132–135.

Benjamin, T. B., & Watt, N. F. (1969). Psychopathology and semantic interpretation of ambiguous words. *Journal of Abnormal Psychology, 74,* 706–714.

Berman, K. F., Illowsky, B. P., & Weinberger, D. R. (1988). Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia: Further evidence for regional and behavioral specificity. *Archives of General Psychiatry, 45,* 616–622.

Berman, K. F., Zec, R. F., & Weinberger, D. R. (1986). Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia: II. Role of neuroleptic treatment, attention and mental effort. *Archives of General Psychiatry, 43,* 126–135.

Bianchi, L. (1922). *The mechanism of the brain and the function of the frontal lobes.* Edinburgh, Scotland: Livingstone.

Blanley, P. H. (1974). Two studies on the language behavior of schizophrenics. *Journal of Abnormal Psychology, 83,* 23–31.

Bleuler, E. (1911). *Dementia praecox, or the group of schizophrenias.* New York: International Universities Press.

Bloom, F. E., Schulman, J. A., & Koob, G. F. (1989). Catecholamines and behavior. In U. Trendelenburg & N. Wiener (Eds.), *Handbook of experimental pharmacology* (Vol. 90, chap. 11, pp. 27–88). Berlin: Springer-Verlag.

Borst, U., & Cohen, R. (1989). Filling the preparatory interval with temporal information or visual noise: The crossover effect in schizophrenics and controls. *Psychological Medicine, 19,* 865–874.

Bowen, F. P., Kamienny, R. S., Burns, M. M., & Yahr, M. D. (1975). Parkinsonism: Effects of L-dopa treatment on concept formation. *Neurology, 25,* 701–704.

Bowers, M. B., Jr., Heninger, G. R., & Sternberg, D. (1980). Clinical processes and central dopaminergic activity in psychotic disorders. *Communications in Psychopharmacology, 4,* 177–188.

Broadbent, D. E. (1958). *Perception and communication.* Elmsford, NY: Pergamon Press.

Broadbent, D. E. (1971). *Decision and stress.* San Diego, CA: Academic Press.

Brozoski, T. J., Brown, R. M., Rosvold, H. E., & Goldman, P. S. (1979). Cognitive deficit caused by regional depletion of dopamine in prefrontal cortex of Rhesus monkey. *Science, 205,* 929–931.

Buchsbaum, M. S., Ingvar, D. H., Kessler, R., Waters, R. N., Cappelletti, J., van Kammen, D. P., King, C. A., Johnson, J. L., Manning, R. G., Flynn, R. W., Mann, L. S., Bunny, W. E., & Sokoloff, L. (1982). Cerebral glucography with positron tomography. *Archives of General Psychiatry, 39,* 251–259.

Bunney, B. S., & Aghajanian, G. K. (1976). Dopamine and norepinephrine innervated in the rat prefrontal cortex: Pharmacological differentiation using microiontophoretic techniques. *Life Sciences, 19,* 1783–1792.

Bunney, B. S., & Sesack, S. R. (1987). Electrophysiological identification and pharmacological characterization of dopamine sensitive neurons in the rat prefrontal cortex. In L. A. Chiodo & A. S. Freeman (Eds.), *Neurophysiology of dopaminergic systems—Current status and clinical perspectives* (pp. 129–140). Grosse Pointe, MI: Lakeshore.

Burnod, Y., & Korn, H. (1989). Consequences of stochastic release of neurotransmitters for network computation in the central nervous system. *Proceedings of the National Academy of Science (U.S.A.), 86,* 352–356.

Calev, A. (1983). Anti-cholinergic drugs and memory. *British Journal of Psychiatry, 143,* 422–423.

Chapman, L. J. (1980). Recent advances in the study of schizophrenic cognition. *Schizophrenia Bulletin, 5*(4), 69–81.

Chapman, L. J., & Chapman, J. P. (1978). The measurement of differential deficit. *Journal of Psychiatric Research, 14,* 303–311.

Chapman, L. J., Chapman, J. P., & Miller, G. A. (1964). A theory of verbal behavior in schizophrenia. In B. A. Maher (Ed.), *Progress in*

experimental personality research (Vol. 1, pp. 135–167). San Diego, CA: Academic Press.

Chiodo, L. A., & Berger, T. W. (1986). Interactions between dopamine and amino acid-induced excitation and inhibition in the striatum. *Brain Research, 375,* 198–203.

Chiodo, L. A., & Bunney, B. S. (1983). Typical and atypical neuroleptics: Differential effects of chronic administration on the activity of A9 and A10 midbrain dopaminergic neurons. *Journal of Neuroscience, 3,* 1607–1619.

Cleeremans, A., Servan-Schreiber, D., & McClelland, J. L. (1989). Finite state automata and simple recurrent networks. *Neural Computation, 1,* 372–381.

Cohen, B. M. (1981). Dopamine receptors and antipsychotic drugs. *McLean Hospital Journal, 1*(2), 95–115.

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.

Cohen, J. D., Targ, E., Kristoffersen, T., & Spiegel, D. (1988). *The fabric of thought disorder: Disturbances in the processing of context.* Unpublished manuscript.

Cohen, R., Borst, U., & Rist, F. (1989). *Cross-over and cross-modality effects in schizophrenia: Some old and some new data.* Unpublished manuscript.

Cohen, R. M., Semple, W. E., Gross, M., Holcomb, H. H., Dowling, M. S., & Nordahl, T. E. (1988). Functional localization of sustained attention: Comparison to sensory stimulation in the absence of instruction. *Neuropsychiatry, Neuropsychology, and Behavioral Neurology, 1,* 2–20.

Cohen, R. M., Semple, W. E., Gross, M., Nordahl, T. E., DeLisi, L. E., Holcomb, H. H., King, A. C., Morihisa, J. M., & Pickar, D. (1987). Dysfunction in a prefrontal substrate of sustained attention in schizophrenia. *Life Sciences, 40,* 2031–2039.

Cohen, R. M., Semple, W. E., Gross, M., Nordahl, T. E., Holcomb, H. H., Dowling, M. S., & Pickar, D. (1988). The effect of neuroleptics on dysfunction in a prefrontal substrate of sustained attention in schizophrenia. *Life Sciences, 43,* 1141–1150.

Conrad, A. J., Abebe, R., & Austin, R. (1991). Hippocampal pyramidal cell disarray in schizophrenia as a bilateral phenomenon. *Archives of General Psychiatry, 48,* 413–417.

Cornblatt, B., Lenzenweger, M. F., & Erlenmeyer-Kimling, L. (1989). A continuous performance test, identical pairs version: II. Contrasting attentional profiles in schizophrenic and depressed patients. *Psychiatry Research, 29,* 65–85.

Corwin, J. V., Kanter, S., Watson, R. T., Heilman, K. M., Valenstein, E., & Hashimoto, A. (1986). Apomorphine has a therapeutic effect on neglect produced by unilateral dorsomedial prefrontal cortex lesions in rats. *Experimental Neurology, 94,* 683–698.

Cozolino, L. J. (1983). The oral and written productions of schizophrenic patients. In B. A. Maher (Ed.), *Progress in experimental personality research* (Vol. 12, pp. 101–152). San Diego, CA: Academic Press.

Creese, I., Burt, D. R., & Snyder, S. H. (1976). Dopamine receptor binding predicts clinical and pharmacological potencies of antischizophrenic drugs. *Science, 192,* 481–483.

Cross, A. J., Crow, T. J., & Owen, F. (1981). 3-H-flupenthixol binding in post-mortem brains of schizophrenics: Evidence for a selective increase in dopamine D2 receptors. *Psychopharmacology, 74,* 122–124.

Crosson, B., & Hughes, C. W. (1987). Role of the thalamus in language: Is it related to schizophrenic thought disorder? *Schizophrenia Bulletin, 13*(4), 605–621.

Crow, T. J. (1980). Molecular pathology of schizophrenia: More than one disease process? *British Medical Journal, 137,* 383–386.

Crow, T. J., Cross, A. J., Johnson, J. A., Johnstone, E. C., Joseph, M. H., Owen, F., Owens, D. G. C., & Poulter, M. (1984). Catecholamines and schizophrenia: An assessment of the evidence. In A. R. Liss (Ed.), *Catecholamines: Neuropharmacology and central nervous system-therapeutic aspects* (pp. 11–20). New York: A. R. Liss.

Damasio, A. R. (1979). The frontal lobes. In K. M. Heilman & E. Valenstein (Eds.), *Clinical neuropsychology* (pp. 339–375). New York: Oxford University Press.

Davidson, M., Keefe, R. S. E., Mohs, R. C., Siever, L. J., Losonczy, M. F., Horvath, T. B., & Davis, K. L. (1987). L-dopa challenge and relapse in schizophrenia. *American Journal of Psychiatry, 144,* 934–938.

Dehaene, S., & Changeux, J. P. (1989). A simple model of prefrontal cortex function in delayed-response tasks. *Journal of Cognitive Neuroscience, 1,* 244–261.

Dehaene, S., & Changeux, J. P. (1991). The Wisconsin Card Sorting Test: Theoretical analysis and modeling in a neuronal network. *Cerebral Cortex, 1,* 62–79.

Deutch, A. Y., Clark, W. A., & Roth, R. H. (1990). Prefrontal cortical dopamine depletion enhances the responsiveness of mesolimbic dopamine neurons to stress. *Brain Research, 521,* 311–315.

Diamond, A. (1985). Developmental of the ability to use recall to guide action, as indicated by infants' performance on A $\bar{B}$. *Child Development, 56,* 868–883.

Diamond, A. (1990a). The development and neural bases of memory functions as indexed by the A $\bar{B}$ and delayed response tasks in human infants and infant monkeys. In A. Diamond (Ed.), *The development and neural bases of higher cognitive functions* (pp. 267–317). New York: NY Academy of Science Press.

Diamond, A. (1990b). Rate of maturation of the hippocampus and the developmental progression of children's performance on the delayed non-matching to sample and visual paired comparison tasks. In A. Diamond (Ed.), *The development and neural bases of higher cognitive functions* (pp. 394–426). New York: NY Academy of Science Press.

Diamond, A. (1990c). Developmental time course in human infants and infant monkeys and the neural bases of inhibitory control in reaching. In A. Diamond (Ed.), *The development and neural bases of higher cognitive functions* (pp. 637–676). New York: NY Academy of Science Press.

Diamond, A., & Doar, B. (1989). The performance of human infants on a measure of frontal cortex function, the delayed response task. *Developmental Psychobiology, 22*(3), 271–294.

Diamond, A., & Goldman-Rakic, P. S. (1989). Comparison of human infants and rhesus monkeys on Piaget's A $\bar{B}$ task: Evidence for dependence on dorsolateral prefrontal cortex. *Experimental Brain Research, 74,* 24–40.

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.

Duncan, J. (1986). Disorganisation of behaviour after frontal lobe damage. *Cognitive Neuropsychology, 3*(3), 271–290.

Dyer, F. N. (1973). The Stroop phenomenon and its use in the study of perceptual, cognitive, and response processes. *Memory and Cognition, 1,* 106–120.

Early, T. S., Posner, M. I., Reiman, E. M., & Raichle, M. E. (1989). Hyperactivity of the left striato-pallidal projection: Part I. Lower level theory. *Psychiatric Development, 2,* 85–108.

Early, T. S., Reiman, E. M., Raichle, M. E., & Spitznagel, E. L. (1987). Left globus pallidus abnormality in never-medicated patients with schizophrenia. *Proceedings of the National Academy of Science (U.S.A.), 84,* 561–563.

Erlenmeyer-Kimling, L., & Cornblatt, B. (1978). Attentional measures in a study of children at high-risk for schizophrenia. *Journal of Psychiatric Research, 14,* 93–98.

Farkas, T., Wolf, A. P., Jaeger, J., Brodie, J. D., Christman, D. R., &

Fowler, J. S. (1984). Regional brain glucose metabolism in chronic schizophrenia. *Archives of General Psychiatry, 41,* 293–300.

Ferron, A., Thierry, A. M., Le Douarin, C., & Glowinski, J. (1984). Inhibitory influence of the mesocortical dopaminergic system on spontaneous activity or excitatory response induced from the thalamic mediodorsal nucleus in the rat medial prefrontal cortex. *Brain Research, 302,* 257–265.

Foote, S. L., Freedman, R., & Oliver, A. P. (1975). Effects of putative neurotransmitters on neuronal activity in monkey auditory cortex. *Brain Research, 86,* 229–242.

Franzen, G., & Ingvar, D. H. (1975). Abnormal distribution of cerebral activity in chronic schizophrenia. *Journal of Psychiatric Research, 12,* 199–214.

Freeman, W. J. (1979). Nonlinear gain mediating cortical stimulus–response relations. *Biological Cybernetics, 33,* 243–247.

Freund, T. F., Powell, J. F., & Smith, A. D. (1984). Ultrastructural immunocytochemical localization of tyrosine hydroxylase in the neostriatum. *Neuroscience, 13,* 1189–1215.

Fuster, J. M. (1980). *The prefrontal cortex.* New York: Raven Press.

Fuster, J. M. (1985a). The prefrontal cortex and temporal integration. In A. Peters & E. G. Jones (Eds.), *Cerebral cortex* (pp. 151–177). New York: Plenum Press.

Fuster, J. M. (1985b). The prefrontal cortex, mediator of cross-temporal contingencies. *Human Neurobiology, 4,* 169–179.

Garmezy, N. (1977). The psychology and psychopathology of attention. *Schizophrenia Bulletin, 3,* 360–369.

Geraud, G., Arne-Bes, M. C., Guell, A., & Bes, A. (1987). Reversibility of hemodynamic hypofrontality in schizophrenia. *Journal of Cerebral Blood Flow and Metabolism, 7,* 9–12.

Geyer, M. A., & Braff, D. L. (1987). Startle habituation and sensorimotor gating in schizophrenia and related animal models. *Schizophrenia Bulletin, 13*(4), 643–688.

Gjerde, P. F. (1983). Attentional capacity dysfunction and arousal in schizophrenia. *Psychological Bulletin, 93,* 57–72.

Goldberg, T. E., Weinberger, D. R., Berman, K. F., Pliskin, N. H., & Podd, M. H. (1987). Further evidence of dementia of the prefrontal type in schizophrenia? *Archives of General Psychiatry, 44,* 1008–1014.

Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In F. Plum (Ed.), *Handbook of physiology—The nervous system, V* (pp. 373–417). Bethesda, MD: American Physiological Society.

Goldman-Rakic, P. S. (1991). Prefrontal cortical dysfunction in schizophrenia: The relevance of working memory. In B. Caroll (Ed.), *Psychopathology and the brain* (pp. 1–23). New York: Raven Press.

Goldman-Rakic, P. S., Leranth, C., Wallns, S. M., Mons, N., & Geffard, M. (1989). Dopamine synaptic complex with pyramidal neurons in primate cerebral cortex. *Proceedings of the National Academy of Science, 86,* 915.

Grace, T. (1991). Phasic versus tonic dopamine release and the modulation of dopamine system responsivity: A hypothesis for the etiology of schizophrenia. *Neuroscience, 41,* 1–24.

Grand, S., Steingart, I., Freedman, N., & Buchwald, C. (1975). Organization of language behavior and cognitive performance in chronic schizophrenia. *Journal of Abnormal Psychology, 84,* 621–628.

Grant, D. A., & Berger, E. A. (1948). A behavioral analysis of degree of reinforcement and ease of shifting to new responses in a Weigl type card sorting problem. *Journal of Experimental Psychology, 38,* 404–411.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* Huntington, NY: Robert E. Krieger.

Gur, R. E., Resnick, S. M., Gur, R. C., Alavi, A., Caroff, S., Kushner, M., & Reivich, M. (1987). Regional brain function in schizophrenia. II. Repeated evaluation with positron emission tomography. *Archives of General Psychiatry, 44,* 125–129.

Harvey, P. D. (1983). Speech competence in manic and schizophrenic psychoses: The association between clinically rated thought disorder and cohesion and reference performance. *Journal of Abnormal Psychology, 92,* 368–377.

Hinton, G. E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review, 98,* 74–95.

Hoffman, R. E. (1987). Computer simulations of neural information processing and the schizophrenia–mania dichotomy. *Archives of General Psychiatry, 44,* 178–188.

Hu, X. T., & Wang, R. Y. (1988). Comparison of effects of D_1 and D_2 dopamine receptor agonists on neurons in the rat caudate putamen: An electrophysiological study. *The Journal of Neuroscience, 8,* 4340–4348.

Ingvar, D. H., & Franzen, G. (1974). Abnormalities of cerebral flow distribution in patients with chronic schizophrenia. *Acta Psychiatrica Scandinavica, 50,* 425–462.

Janowsky, D. S., Huey, L., Storms, L., & Judd, L. L. (1977). Methylphenidate hydrochloride effects on psychological tests in acute schizophrenic and nonpsychotic patients. *Archives of General Psychiatry, 34,* 189–194.

Janowsky, D. S., & Rich, C. (1979). Amphetamine psychosis and psychotic symptoms. *Psychopharmacology, 65,* 73–77.

Joseph, M. H., Frith, C. D., & Waddington, J. L. (1979). Dopaminergic mechanisms and cognitive deficits in schizophrenia. *Psychopharmacology, 63,* 273–280.

Karoum, F., Karson, C. N., Bigelow, L. B., Lawson, W. B., & Wyatt, R. J. (1987). Preliminary evidence of reduced combined output of dopamine and its metabolites in chronic schizophrenia. *Archives of General Psychiatry, 44,* 604–607.

Koh, S. D. (1978). Remembering of verbal material by schizophrenic young adults. In S. Schwartz (Ed.), *Language and cognition in schizophrenia* (pp. 55–99). Hillsdale, NJ: Erlbaum.

Kolb, B., & Whishaw, I. Q. (1983). Performance of schizophrenic patients on tests sensitive to left or right frontal, temporal, or parietal function in neurologic patients. *Journal of Mental and Nervous Disease, 171,* 435–443.

Kopfstein, J. H., & Neale, J. M. (1972). A multivariate study of attention dysfunction in schizophrenia. *Journal of Abnormal Psychology, 3,* 294–298.

Kornetsky, C. (1972). The use of a simple test of attention as a measure of drug effects in schizophrenic patients. *Psychopharmacologia (Berlin), 8,* 99–106.

Kornetsky, C., & Orzack, M. H. (1978). Physiological and behavioral correlates of attention dysfunction in schizophrenic patients. *Journal of Psychiatric Research, 14,* 69–79.

Kovelman, J. A., & Scheibel, A. B. (1984). A neurohistological correlate of schizophrenia. *Biological Psychiatry, 19,* 1601–1621.

Kraeplin, E. (1950). *Dementia praecox and paraphrenia* (J. Zinkin, Trans.). New York: International Universities Press.

Lake, C. R., Sternberg, D. E., van Kammen, D. P., Ballenger, J. C., Ziegler, M. G., Post, R. M., Kopin, I. J., & Bunney, W. E. (1980). Schizophrenia: Elevated cerebrospinal fluid norepinephrine. *Science, 207,* 331–333.

Lang, P. J., & Buss, A. H. (1965). Psychological deficit in schizophrenia: II. Interference and activation. *Journal of Abnormal Psychology, 70,* 77–106.

Larsen, S. F., & Fromholt, P. (1976). Mnemonic organization and free recall in schizophrenia. *Journal of Abnormal Psychology, 85,* 61–65.

Lehky, S. R., & Sejnowski, T. J. (1988). Network model of shape-from-shading: Neural function arises from both receptive and projective fields. *Nature, 333,* 452–454.

Levin, S. (1984). Frontal lobe dysfunctions in schizophrenia—II. Impairments of psychological and brain functions. *Journal of Psychiatry Research, 18*(1), 57–72.

Levin, S., Yurgelun-Todd, D., & Craft, S. (1989). Contributions of clinical neuropsychology to the study of schizophrenia. *Journal of Abnormal Psychology, 98,* 341–356.

Lieberman, J. A., Kane, J. M., Gadaleta, D., Brenner, R., Lesser, M. S., & Kinon, B. (1984). Methylphenidate challenge as a predictor of relapse in schizophrenia. *American Journal of Psychiatry, 141,* 633–638.

Losonczy, M. F., Davidson, M., & Davis, K. L. (1987). The dopamine hypothesis of schizophrenia. In H. Y. Meltzer (Ed.), *Psychopharmacology: The third generation of progress* (pp. 715–726). New York: Raven Press.

Luria, A. R. (1966). *Higher cortical functions in man.* London: Tavistock.

Mackay, A. V. P. (1980). Positive and negative schizophrenic symptoms and the role of dopamine. *British Journal of Psychiatry, 137,* 379–386.

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin, 109,* 163–203.

Maher, B. A. (1972). The language of schizophrenia: A review and interpretation. *British Journal of Psychiatry, 120,* 3–17.

Malmo, H. P. (1974). On frontal lobe functions: Psychiatric patient controls. *Cortex, 10,* 231–237.

McClelland, J. L. (1988). Connectionist models and psychological evidence. *Journal of Memory and Language, 27,* 107–123.

McClelland, J. L., & Rumelhart, D. E. (Eds.). (1986). *Parallel distributed processing explorations in the microstructure of cognition. Vol. 2: Psychological and biological models.* Cambridge, MA: MIT Press.

McClelland, J. L., & Rumelhart, D. E. (1988). *Explorations in parallel distributed processing. A handbook of models, programs, and exercises.* Cambridge, MA: MIT Press.

McCulloch, J., Savaki, H. E., McCulloch, M. C., Jehle, J., & Sokoloff, L. (1982). The distribution of alterations in energy metabolism in the rat brain produced by apomorphine. *Brain Research, 243,* 67–80.

McGhie, A. (1970). Attention and perception in schizophrenia. In B. A. Maher (Ed.), *Progress in experimental personality research* (Vol. 5, pp. 1–36). San Diego, CA: Academic Press.

McGhie, A., & Chapman, J. (1961). Disorders of attention and perception in early schizophrenia. *British Journal of Medical Psychology, 34,* 103–116.

Meltzer, H. Y., & Stahl, S. M. (1976). The dopamine hypothesis of schizophrenia: A review. *Schizophrenia Bulletin, 2,* 19–76.

Milner, B. (1963). Effects of different brain lesions on card sorting. *Archives of Neurology, 9,* 90–100.

Mirsky, A. F., DeLisi, L. E., Buchsbaum, M. S., Quinn, O. W., Schwerdt, P., Siever, L. J., Mann, L., Weingartner, H., Zec, R., Sostek, A., Alterman, I., Revere, V., Dawson, S. D., & Zahn, T. P. (1984). The Genain quadruplets: Psychological studies. *Psychiatry Research, 13,* 77–93.

Mishkin, M., & Pribram, K. H. (1955). Analysis of the effects of frontal lesions in monkey. I. Variations on delayed alternation. *Journal of Comparative and Physiological Psychology, 338,* 492–495.

Morihisa, J. M., & McAnulty, G. B. (1985). Structure and function: Brain electrical activity mapping and computed tomography in schizophrenia. *Biological Psychiatry, 20,* 3–19.

Nauta, J. H. W., & Domesick, V. B. (1981). Ramifications of the limbic system. In S. Matthysse (Ed.), *Psychiatry and the biology of the human brain: A symposium dedicated to Seymour S. Kety* (pp. 165–188). New York: Elsevier/North-Holland.

Neale, J. M. (1971). Perceptual span in schizophrenia. *Journal of Abnormal Psychology, 77,* 196–204.

Nelson, H. E. (1976). A modified card sorting test sensitive to frontal lobe defects. *Cortex, 12,* 313–324.

Nuechterlein, K. H. (1977). Reaction time and attention in schizophrenia: A critical evaluation of the data and theories. *Schizophrenia Bulletin, 3,* 373–428.

Nuechterlein, K. H. (1983). Signal detection in vigilance tasks and behavioral attributes among offspring of schizophrenic mothers and among hyperactive children. *Journal of Abnormal Psychology, 92,* 4–28.

Nuechterlein, K. H. (1984). Information processing and attentional functioning in the developmental course of schizophrenic disorders. *Schizophrenia Bulletin, 10,* 160–203.

Oades, R. D. (1981). Dopaminergic agonist and antagonistic drugs in the ventral tegmentum of rats inhibit and facilitate changes of food search behaviour. *Neuroscience Letters, 27,* 75–80.

Oltmanns, T. F., & Neale, J. M. (1975). Schizophrenic performance when distractors are present: Attentional deficit or differential task difficulty? *Journal of Abnormal Psychology, 84,* 205–209.

Parasuraman, R. (1979). Memory load and event rate control sensitivity decrements in sustained attention. *Science, 205,* 924–927.

Passingham, R. E. (1985). Memory of monkeys (*Macaca mulatta*) with lesions in prefrontal cortex. *Behavioral Neuroscience, 99,* 3–21.

Patterson, K., Seidenberg, M. S., & McClelland, J. L. (1989). Connections and disconnections: Acquired dyslexia in a computational model of reading processes. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology:* Oxford: Clarendon Press.

Perret, E. (1974). The left frontal lobe of man and the suppression of habitual responses in verbal categorical behaviour. *Neuropsychologia, 12,* 323–330.

Piaget, J. (1954). *The origins of intelligence in children.* New York: Basic Books. (Original work published 1937)

Pickar, D., Labarca, R., Linnoila, M., Roy, A., Hommer, D., Everett, D., & Paul, S. M. (1984). Neuroleptic-induced decrease in plasma homovanillic acid and antipsychotic activity in schizophrenic patients. *Science, 225,* 954–957.

Pickel, V. M., Beckley, S., Joh, T. H., & Reis, D. J. (1981). Tyrosine hydroxylase-immunoreactive boutons in synaptic contact with identified striatonigral neurons, with particular reference to dendritic spines. *Brain Research, 225,* 373–385.

Printz, H., & Servan-Schreiber, D. (1990). *Foundations of a computational theory of catecholamine effects* (Tech. Rep. CMU-CS-90-105). Pittsburgh, PA: Carnegie Mellon University, School of Computer Science.

Pycock, C. J., Kerwin, R. W., & Carter, C. J. (1980). Effect of lesion of cortical dopamine terminals on subcortical dopamine receptors in rats. *Nature, 286,* 74–77.

Reader, T. A., Ferron, A., Descarries, L., & Jasper, H. H. (1979). Modulatory role for biogenic amines in the cerebral cortex: Microiontophoretic studies. *Brain Research, 160,* 217–229.

Robinson, A. L., Heaton, R. K., Lehman, R. A. W., & Stilson, D. W. (1980). The utility of the Wisconsin Card Sorting Test in detecting and localizing frontal lobe lesions. *Journal of Consulting and Clinical Psychology, 48,* 605–614.

Rochester, S. R., & Martin, J. R. (1979). *Crazy talk: A study in the discourse of schizophrenic speakers.* New York: Plenum Press.

Rosenkilde, K. E. (1979). Functional heterogeneity of the prefrontal cortex in the monkey: A review. *Behavioral Neural Biology, 25,* 301–345.

Rosvold, K. E., Mirsky, A. F., Sarason, I., Bransome, E. D., & Beck, L. H. (1956). A continuous performance test of brain damage. *Journal of Consulting Psychology, 20,* 343–350.

Rosvold, K. E., Szwarcbart, M. K., Mirsky, A. F., & Mishkin, M.

(1961). The effect of frontal-lobe damage on delayed response performance in chimpanzees. *Journal of Comparative and Physiological Psychology, 54,* 368–374.

Roth, R. H., Bacopoulos, N. G., Bustos, G., & Redmond, D. E. (1980). Antipsychotic drugs: Differential effects on dopamine neurons in basal ganglia and mesocortex following chronic administration in human and non-human primates. In F. Cattabeni, G. Racagni, P. F. Spano, & E. Costa (Eds.), *Advances in biochemical psychopharmacology* (Vol. 24, pp. 513–520). New York: Raven Press.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by backpropagating errors. *Nature, 323,* 533–536.

Rumelhart, D. E., & McClelland, J. L. (1986a). On learning the past tense of English verbs. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models* (pp. 216–271). Cambridge, MA: MIT Press.

Rumelhart, D. E., & McClelland, J. L. (Eds.). (1986b). *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1: Foundations.* Cambridge, MA: MIT Press.

Rumelhart, D. E., & McClelland, J. L. (1986c). PDP models and general issues in cognitive science. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1: Foundations* (pp. 110–146). Cambridge, MA: MIT Press.

Rund, B. R., & Landro, R. (1990). Information processing: A new model for understanding cognitive disturbances in psychiatric patients. *Acta Psychiatrica Scandinavica, 81,* 305–316.

Rutschmann, J., Cornblatt, B., & Erlenmeyer-Kimling, L. (1977). Sustained attention in children at risk for schizophrenia. *Archives of General Psychiatry, 34,* 571–575.

Rutter, D. R. (1979). The reconstruction of schizophrenic speech. *British Journal of Psychiatry, 134,* 356–359.

Salzinger, K. (1971). An hypothesis about schizophrenic behavior. *American Journal of Psychotherapy, 25,* 601–614.

Salzinger, K., Portnoy, S., & Feldman, R. S. (1964). Verbal behavior of schizophrenic and normal subjects. *Annals of the New York Academy of Sciences, 105,* 845–860.

Salzinger, K., Portnoy, S., Pisoni, D. B., & Feldman, R. S. (1970). The immediacy hypothesis and response-produced stimuli in schizophrenic speech. *Journal of Abnormal Psychology, 76,* 258–264.

Sawaguchi, T., & Goldman-Rakic, P. S. (1991). D1 dopamine receptors in prefrontal cortex: Involvement in working memory. *Science, 251,* 947–950.

Sawaguchi, T., & Matsumura, M. (1985). Laminar distributions of neurons sensitive to acetylcholine, noradrenaline and dopamine in the dorsolateral prefrontal cortex of the monkey. *Neuroscience Research, 2,* 255–273.

Scatton, B. (1977). Differential regional development of tolerance to increase in dopamine turnover upon repeated neuroleptic administration. *European Journal of Pharmacology, 46,* 363–369.

Scatton, B., Boireau, A., Garret, C., Glowinski, J., & Julou, L. (1977). Action of the palmitic ester of pipotiazine on dopamine metabolism in the nigro-striatal, meso-limbic and meso-cortical systems. *Naunyn-Schmiedeberg's Archives of Pharmacology, 296,* 169–175.

Schneider, J. S., Levine, M. S., Hull, C. D., & Buchwald, N. A. (1984). Effects of amphetamine on intracellular responses of caudate neurons in the cat. *Journal of Neuroscience, 4,* 930–938.

Schwartz, S. (1982). Is there a schizophrenic language? *Behavioral and Brain Sciences, 5,* 579–626.

Sedvall, G., Fyro, B., Nyback, H., Wiesel, F. A., & Wode-Helgodt, B. (1974). Mass fragmentometric determination of HVA in lumbar cerebrospinal fluid of schizophrenic patients during treatment with antipsychotic drugs. *Journal of Psychiatric Research, 11,* 75–80.

Seeman, P., Ulpian, C., Bergeron, C., Riederer, P., Jellinger, K., Gabriel, E., Reynolds, G. P., & Tourtelotte, W. W. (1984). Bimodal distribution of dopamine receptor densities in brain of schizophrenics. *Science, 225,* 728–731.

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed development model of word recognition and naming. *Psychological Review, 96,* 523–568.

Servan-Schreiber, D. (1990). *From physiology to behavior: Computational models of catecholamine modulation of information processing* (PhD thesis; Tech. Rep. CMU-CS-90-167). Pittsburgh, PA: Carnegie Mellon University, School of Computer Science.

Servan-Schreiber, D., Cleeremans, A., & McClelland, J. L. (in press). Graded state machines: The representation of temporal contingencies in simple recurrent networks. *Machine Learning.*

Servan-Schreiber, D., Printz, H. W., & Cohen, J. D. (1990). A network model of catecholamine effects: Gain, signal-to-noise ratio and behavior. *Science, 249,* 892–895.

Shakow, D. (1962). Segmental set: A theory of the formal psychological deficit in schizophrenia. *Archives of General Psychiatry, 6,* 1–17.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London, 298,* 199–209.

Simon, H., Scatton, B., & Le Moal, M. (1980). Dopaminergic A10 neurones are involved in cognitive functions. *Nature, 286,* 150–151.

Snyder, S. H. (1972). Catecholamines in the brain as mediators of amphetamine psychosis. *Archives of General Psychiatry, 27,* 169–179.

Snyder, S. H. (1976). The dopamine hypothesis of schizophrenia: Focus on the dopamine receptor. *American Journal of Psychiatry, 133,* 197–202.

Snyder, S. H., Banerjee, S. P., Yamamura, H. I., & Greenberg, D. (1974). Drugs, neurotransmitters and schizophrenia. *Science, 184,* 1243–1253.

Spohn, H. E., Lacoursiere, R. B., Thomson, K., & Coyne, L. (1977). Phenothiazine effects on psychological and psychophysiological dysfunction in chronic schizophrenics. *Archives of General Psychiatry, 34,* 633–644.

Spring, B. (1985). Distractibility as a marker of vulnerability to schizophrenia. *Psychopharmacology Bulletin, 21,* 509–512.

Stevens, J. R. (1973). An anatomy of schizophrenia? *Archives of General Psychiatry, 29,* 177–189.

Strauss, M. E. (1975). Strong meaning-response bias in schizophrenia. *Journal of Abnormal Psychology, 84,* 293–298.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18,* 643–662.

Stuss, D. T., & Benson, D. F. (1984). Neuropsychological studies of the frontal lobes. *Psychological Bulletin, 95,* 3–28.

Swets, J. A., & Sewall, S. T. (1963). Invariance of signal detectability over stages of practice and levels of motivation. *Journal of Experimental Psychology, 66,* 120–126.

Tassin, J. P., Simon, H., Glowinski, J., & Bockaert, J. (1982). Modulations of the sensitivity of DA receptors in the prefrontal cortex and the nucleus accumbens: Relationship with locomotor activity. In R. Collu, J. R. Ducharme, A. Barbeau, & G. Tobis (Eds.), *Brain peptides and hormones* (pp. 17–30). New York: Raven Press.

Tyrone, L., & Seeman, P. (1980). Elevation of brain neuroleptic/dopamine receptors in schizophrenia. *American Journal of Psychiatry, 137,* 191–197.

van Kammen, D. P., Peters, J., van Kammen, W. B., Nugent, A., Goetz, K. L., Yao, J., & Linnoila, M. (1989). CSF norepinephrine in schizophrenia is elevated prior to relapse after haloperidol withdrawal. *Biological Psychiatry, 26,* 176–188.

Wapner, S., & Krus, D. M. (1960). Effects of lysergic acid diethylamide,

and differences between normals and schizophrenics on the Stroop color-word test. *Journal of Neuropsychiatry, 2,* 76–81.

Weinberger, D. R. (1987). Implications of normal brain development for the pathogenesis of schizophrenia. *Archives of General Psychiatry, 44,* 660–669.

Weinberger, D. R., & Berman, K. F. (1988). Speculation on the meaning of metabolic hypofrontality in schizophrenia. *Schizophrenia Bulletin, 14,* 157–168.

Weinberger, D. R., Berman, K. F., & Chase, T. N. (1988). Mesocortical dopaminergic function and human cognition. *Annals of the New York Academy of Science, 537,* 330–338.

Weinberger, D. R., Berman, K. F., & Illowsky, B. P. (1988). Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia: A new cohort of evidence for a monoaminergic mechanism. *Archives of General Psychiatry, 45,* 606–615.

Weinberger, D. R., Berman, K. F., & Zec, R. F. (1986). Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia: I. Regional cerebral blood flow evidence. *Archives of General Psychiatry, 43,* 114–125.

Weinberger, D. R., Bigelow, L. B., Kleinman, J. E., Klein, S. T., Rosenblatt, J. E., & Wyatt, R. J. (1980). Cerebral ventricular enlargement in chronic schizophrenia. *Archives of General Psychiatry, 37,* 11–13.

Widlöcher, D., & Hardy-Bayle, M. C. (1989). Cognition and control of action in psychopathology. *European Bulletin of Cognitive Psychology, 9,* 583–615.

Wielgus, M. S., & Harvey, P. D. (1988). Dichotic listening and recall in schizophrenia and mania. *Schizophrenia Bulletin, 14*(4), 689–700.

Wysocki, J. J., & Sweet, J. J. (1985). Identification of brain-damaged, schizophrenic, and normal medical patients using a brief neuropsychological screening battery. *International Journal of Clinical Neuropsychology, 7*(1), 40–49.

Zipser, D., & Andersen, R. A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature, 331,* 679–684.

Appendix

Details of the Continuous Performance Test and Lexical Disambiguation Simulations

Here we describe details specific to the continuous performance test (CPT) and lexical disambiguation simulations. Details of the Stroop simulation are reported in J. D. Cohen, Dunbar, and McClelland (1990). All simulations shared the following features. For each condition of each task, 1,000 trials were run to obtain reliable estimates of performance in the presence of processing noise. Noise was introduced to simulate the variability of human subjects in these tasks; in the absence of such noise, processing is deterministic and therefore performance is perfect on all tasks. To provide a measure of the reliability of the simulation results, the 5% confidence intervals were computed for a run of 1,000 trials in each condition of the task. For each model, parameters were adjusted to fit the model as closely as possible to normal subjects' performance in the corresponding task. The parameters for the model were then fixed (with the exception of the response threshold in the CPT—see explanation following), and only the value of the gain parameter was adjusted to simulate the performance of schizophrenics in that task. The value of gain chosen to simulate schizophrenic performance was constrained to be the same across all conditions of all tasks. A value was picked (0.6) that provided the best overall fit to the data from the three tasks.

Details of the CPT Simulation

Architecture

The network contained three layers: an input layer of 22 units, a hidden layer of 30 units, and an output layer of 11 units. Each unit was connected through feed-forward connections to all the units in the subsequent layer. Output units had no outgoing connections.

The input layer was divided into two pools. The first pool constituted the input layer proper. Each of the 12 units in this pool represented a particular letter feature (i.e., one of the 12 segments of the grid presented in Figure A1). A pattern of activation over this pool therefore represented a particular letter. The other pool of 10 units—called the "context" pool—was used to present a copy of the pattern of activation over the output layer at the previous time step. The output layer was also divided into two pools: 10 "prior stimulus" units and 1 response unit. Each of the prior stimulus units represented one particular letter from A to J (the first unit coded for A, the last one for J). The last unit of the output layer was the "response unit." Its activation was not copied onto the context pool at each time step.

The activation of the units in the input pool proper was always determined by the identity of the letter being presented (e.g., all of the units corresponding to the features of A were clamped on with an activation of 1.0 when the letter A was presented). Activation of the units in the context pool—copies of the activations of the units in the output layer at the previous step—were also clamped on with the presentation of each letter. The activation of any other unit was determined by multiplying the activation of each unit in the previous layer by the weight of the corresponding connection and summing these products over all such connections. This quantity, called the "net input," was then passed through the logistic activation function to determine the activation value. In the control condition of the simulation, the logistic function had a gain of 1.0 and a fixed negative bias of −4.0.

Training

Before training, all of the weights in the network were randomly selected from the interval [−1.0, +1.0]. The network was trained with the back-propagation learning algorithm as implemented in the PDP software provided by McClelland and Rumelhart (1988). One third of all trials during training consisted of target trials (consecutive identical letters). Apart from this constraint, letters were randomly selected on each trial.

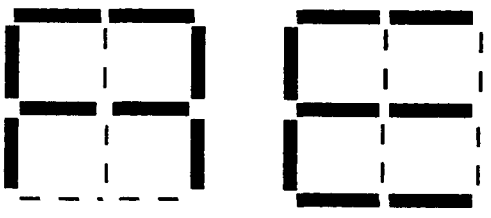

Figure A1. The patterns corresponding to the letters A and E. (Each unit in the input layer proper corresponded to 1 of the 12 segments of a feature grid. The pattern of activation over the input layer proper depended on the overlap between a given letter and the feature grid.)

(*Appendix continues on next page*)

Following the presentation of a letter, the output of the network—that is, the activation of the prior stimulus units and of the response unit—was compared to the desired target and the weights were adjusted using the back-propagation algorithm. The learning rate was 0.01, and the momentum was 0.9. No noise was added to the net input during training. The network was trained until the response unit was activated above 0.6 by all of the 10 possible pairs of consecutive identical letters, and activated below 0.1 in all other cases. This was achieved after 7,000 training trials.

Testing

During testing, a noise term was added to the computation of the net input into each unit. This noise term was sampled independently for each unit and on each trial from a normal Gaussian distribution. It was then multiplied by a scaling factor that was the same for all units in the network. After each letter presentation, the activation of the response unit was compared to a threshold value. When this activation was above threshold in the absence of a target event, the trial was classified as a false alarm. The percentage of false alarms reported in the text corresponds to the total number of false alarms divided by the total number of nontarget events. Conversely, when the activation of the response unit was below threshold in the presence of a target event, the trial was classified as a miss. The percentage of misses corresponds to the total number of misses divided by the total number of target events.

The value of the response threshold and of the noise scaling factor were adjusted until the number of misses and false alarms matched the performance of normal controls reported in Cornblatt, Lenzenweger, and Erlenmeyer-Kimling (1989). Note that there is only one value of the pair [response threshold, noise scaling factor] for which the performance of the model can match the performance of Cornblatt et al.'s normal subjects. This is because the noise scaling factor is solely responsible for the degree of confusability between target and nontarget events, whereas the value of the response threshold affects only the relative balance between misses and false alarms for a given level of discriminability. The final value of the threshold was 0.46, and the noise scaling factor was 0.185. These values were then used to simulate normal performance.

After the noise scaling factor was determined, we reduced the gain of the letter identification units. We then adjusted the response threshold until the number of misses and false alarms matched the data of schizophrenic subjects reported by Cornblatt et al. Note again that changes in the gain of the letter identification units affect the confusability between target and nontarget events, whereas the value of the threshold affects only the relative balance between misses and false alarms. With a gain of 0.6, the performance of the model provided a good fit to that of schizophrenic subjects when the response threshold was 0.425. We felt justified in adjusting the threshold in addition to reducing the gain parameter because it can be easily shown that, for a given task, the optimal response threshold in a network of neural-like elements changes as the gain changes (see Printz & Servan-Schreiber, 1990; Servan-Schreiber, Printz, & Cohen, 1990). Similarly, it is commonly assumed that subjects who vary in their ability to discriminate between target and nontarget events adjust their threshold as a function of this ability to maximize their payoffs given the structure of the signal detection task.

Details of the Lexical Disambiguation Simulation

Architecture

The architecture of the language model was similar to that of the CPT model. Processing involved the same sequential presentation of stimuli and copy of the previous output. The input layer was divided into a "word input module" containing 30 units and a context pool of 20 units (each of which corresponded to a unit in the discourse module, to be mentioned shortly). The hidden layer contained 25 units. The output layer was divided into a "discourse module" of 20 units and a "meaning output module" of 40 units.

The network was used to represent 10 ambiguities with two meanings each, and a set of related context words related to each of the two meanings of the ambiguities. Thus, each of the 30 units in the word input module corresponded to one particular word: 10 corresponded to ambiguous words (e.g., PEN); and 20 corresponded to each of the two nonambiguous (context) words related to one meaning of an ambiguity (e.g., PAPER and CHICKEN). Each of the 20 units in the discourse module corresponded to a specific topic of discourse (e.g., FARMING or WRITING) related to one meaning or the other of an ambiguity and its related context word. Finally, the units in the meaning output module, like the response unit in the CPT model, could be monitored to determine the response of the network to input stimuli. Each of these units corresponded to a particular meaning of the word presented on the word input module (e.g., writing implement, fenced enclosure, a kind of fowl). Because there were 10 ambiguities, each with two different meanings, as well as 20 context words, each with their own meanings, there were 40 units in the meaning output module.

Activation of a unit in the word input module represented the word being presented (i.e., the unit corresponding to that word was clamped on with an activation of 1.0, whereas all other units were clamped off with an activation of 0.0). Activation of the units in the context pool—copies of the activations of the units in the discourse module at the previous step—were also clamped with the presentation of each word. When the first word of a sequence was presented, activations on the context pool were reset to 0.5. The activation of all other units in the model was determined by computing their net input and passing it through the logistic function.

Training

Before training, all of the weights in the network were randomly selected from the interval [−1.0, +1.0]. The network was trained with the back-propagation learning algorithm with a learning rate of 0.01 and a momentum of 0.9. No noise was added to the net input during training.

The training set consisted of 10 ambiguous words for which one meaning is more common than another (e.g., PEN, BANK). In addition, there were 20 context words—one for each meaning of each ambiguous word (e.g., PAPER and CHICKEN for the word PEN). Two types of training trials were used. In one type, a unit was activated in the word input module but no units were activated in the context module. The network was allowed to settle and was then trained to activate the units in the meaning output and discourse modules corresponding to the meaning of the input word. In the case of ambiguous input words, the network was sometimes trained to activate the meaning output and discourse units corresponding to one meaning of the input and sometimes trained to activate the units corresponding to the other. The ratio of training trials for common to uncommon meanings of ambiguous words was adjusted to provide the best fit to the empirical data for control subjects when the network was tested. This was achieved when the network was trained to produce the common meaning of an ambiguity 10 times more often than the uncommon meaning when no context was available. In the other type of training trial, the network was trained to produce the relevant meaning of an ambiguous input word when context was available. In this case, both an input word unit and a context unit were activated, and the network was trained to activate the units in the meaning output and discourse modules corresponding to the context-relevant meaning of the input. Thus, when context was present in the input, the network was trained to produce only the correct (i.e., context-relevant) meaning of the ambigu-

ity. Training the network five times more often on context trials related to the common than the uncommon meaning of an ambiguity provided an acceptable fit to the empirical data for control subjects. The network was trained equally often on all other stimuli (i.e., input words related to each of the meanings of each ambiguity—such as PAPER and CHICKEN—with and without context).

During training, stimuli were presented in random order and in the relative proportions just described. Training proceeded until the total sum square error on the entire set of training stimuli was less than 0.5. This was achieved after 5,000 epochs (i.e., passes through the entire training set).

Testing

During testing, a noise term was added to the computation of the net input into each unit in the network (except the input units). This noise term was sampled independently for each unit and on each trial from a normal Gaussian distribution. It was then multiplied by a scaling factor. The value of this factor was adjusted so that the error rates in the model matched the performance of control subjects in the empirical study. Such a match was obtained with a value of 0.85 for the noise scaling factor.

In each test trial, three stimuli were presented to the network in sequence. The first two stimuli corresponded to each of the two clauses (context and ambiguity) in the empirical study. On some trials the ambiguity was presented first, followed by the context; on others, the opposite order was used (trials were distributed across conditions in the same proportions as in the empirical study—see text). The third stimulus—always the ambiguity being tested—was used to probe the network's interpretation of the ambiguity. After presenting this probe, the unit most active in the meaning output module was recorded. If this corresponded to the context-relevant meaning of the ambiguity, the response was classified as "correct." If the response corresponded to the other meaning of the ambiguity, it was classified as "incorrect-related," and if it did not correspond to either meaning of the ambiguity, it was classified as "incorrect-unrelated." The model was tested 100 times with each of the 10 ambiguities in each of the conditions of the empirical study (i.e., context first, related to weak meaning; context last, related to weak meaning; and context first, related to strong meaning) for a total of 1,000 trials. After parameters of the model (e.g., the training frequency ratios and noise scaling factor) had been adjusted to best match the empirical data for control subjects, we then reduced the gain of the units only in the discourse module, from 1.0 to 0.6, and ran the same test. This produced the pattern of results reported in Figure 10 of the text.

Received April 20, 1990
Revision received April 15, 1991
Accepted April 18, 1991 ■

Zahn-Waxler Appointed New Editor, 1993–1998

The Publications and Communications Board of the American Psychological Association announces the appointment of Carolyn Zahn-Waxler as editor of *Developmental Psychology.* Zahn-Waxler is associated with the National Institute of Mental Health. As of January 1, 1992, manuscripts should be directed to

Carolyn Zahn-Waxler
4305 Dresden Street
Kensington, Maryland 20895

Manuscript submission patterns make the precise date of completion of the 1992 volume uncertain. The current editor will receive and consider manuscripts through December 1991. Should the 1992 volume be completed before that date, manuscripts will be redirected to the incoming editor for consideration in the 1993 volume.