

MODELS OF NEUROMODULATION AND INFORMATION PROCESSING DEFICITS IN SCHIZOPHRENIA

David Servan-Schreiber and Jonathan D. Cohen

School of Computer Science and Department of Psychology
Carnegie Mellon University

and

Western Psychiatric Institute and Clinic
University of Pittsburgh

Abstract

This paper illustrates the use of connectionist models to explore the relationship between biological variables and cognitive deficits. The models show how empirical observations about biological and psychological deficits can be captured within the same framework to account for specific aspects of behavior. We present simulation models of three attentional and linguistic tasks in which schizophrenics show performance deficits. At the cognitive level, the models suggest that a disturbance in the processing of context can account for schizophrenic patterns of performance in both attention and language-related tasks. At the same time, the models incorporate a mechanism for processing context that can be identified with the function of the prefrontal cortex, and a parameter that corresponds to the effects of dopamine in the prefrontal cortex. A disturbance in this parameter is sufficient to account for schizophrenic patterns of performance in the three cognitive tasks simulated. Thus, the models offer an explanatory mechanism linking performance deficits to a disturbance in the processing of context which, in turn, is attributed to a reduction of dopaminergic activity in prefrontal cortex.

Schizophrenia is marked by a wide variety of behavioral deficits, including disturbances of attention, language processing and problem solving. Schizophrenia is also characterized by biological abnormalities, 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 paper we address this issue by drawing upon the framework of connectionist models. This framework provides theoretical concepts that are intermediate between the details of neuroscientific observations and the box-and-arrow diagrams of traditional information processing or neuropsychological theories. We explore the ability of connectionist models to explain aspects of schizophrenic behavior in terms of specific underlying biological disturbances. At the behavioral level, the models address both normal and schizophrenic 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 use 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 processing of context.

The models also suggest that this functional deficit may be explained by a specific biological disturbance: a reduction in the effects of the neurotransmitter dopamine in prefrontal cortex (PFC). First, we show how a particular parameter of the models can be used to simulate the neuromodulatory effects of dopamine at the neuronal level. We then present the results of simulations in which this parameter is disturbed within a module corresponding to the function of PFC. In each of the three behavioral simulations, this disturbance leads to changes in performance that quantitatively match those observed for schizophrenics in the corresponding tasks. These findings suggest that a number of the disturbances of attention and language found in schizophrenia may all result from a disturbance in the processing of context which, in turn, may be explained by a single biological abnormality, a reduction of dopaminergic activity in PFC.

The background of the models presented in this paper spans a large and diverse literature concerning cognitive deficits in schizophrenia, the anatomy and physiology of dopamine systems, the neurophysiology and

The order of authorship was determined by the toss of a coin. This work was supported by a NIMH Physician Scientist Award (MH00673) to JDC and a NIMH Individual Fellow Award (MH09696) to DSS. Correspondence concerning this paper should be addressed to either author at Carnegie Mellon University, Pittsburgh, PA 15213.

neuropsychology of frontal cortex, and the role of these biological systems in schizophrenia. A full review of these data is beyond the scope of this communication (see Cohen and Servan-Schreiber, 1989, for a more comprehensive review). Here, we highlight five empirical observations concerning information processing and biological deficits in schizophrenia. The purpose of our simulations is to show how these observations — which range from the biological to the cognitive level — can be articulated within the same model and account for specific aspects of behavior. These observations are: 1) Schizophrenics' performance in a variety of cognitive tasks indicates a decreased ability to use context information for selecting appropriate behavior. By *context*, we mean information that is relevant to, but is not part of the content of a behavioral response. This can be task instructions or specific previous stimuli that determine correct behavior. 2) Prefrontal cortex (PFC) is directly involved in, and necessary for the representation and maintenance of context information. 3) The normal function of PFC relies on the activity of the mesocortical dopamine system. 4) Dopamine has a modulatory effect on the activity of PFC, by influencing the responsiveness, or *gain*, of cells in this brain region; and 5) Schizophrenia is associated with abnormalities of both frontal cortex and dopamine activity.

Disturbances in the Processing of Context in Schizophrenia

The Stroop task. This task taps a fundamental attentional phenomenon: the ability to respond to one set of stimuli, even when other, more compelling stimuli are available. The 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. Normal subjects are less able to attend selectively to colors (i.e., ignore words) than the reverse. If schizophrenics suffer from an attentional deficit then this effect should be exacerbated; that is, they should be less able to ignore word information, and should show a greater interference effect. This prediction is supported by studies of schizophrenic performance in the Stroop task (Abramczyk et al., 1983; Wapner & Krus, 1960). However, because an overall slowing of reaction time is also observed, the significance of an increase in interference has been called into question: This may simply reflect an unanticipated effect of general slowing of performance, rather than of a specific attentional deficit. This issue has not been resolved in the literature. Below, we will show how a simulation model of this task can help 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, the processing of context is also central 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. However, in other attentional tasks — such as the continuous performance test — this is not the case.

The Continuous Performance Test. In this task (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 stream of other letters), or a stimulus appearing in a particular context (e.g. respond to 'X' preceded by 'A'). Schizophrenics (and often their biological relatives) are typically impaired in their ability to discriminate between target and distractors on this task, compared to normal and patient controls (e.g., Kornetsky, 1972; Nuechterlein, 1984). This deficit is most apparent in versions of the task that make high processing demands. For example, in the 'CPT-double' a target event consists of two consecutive identical letters. Memory for the previous letter provides the context necessary to evaluate the significance of the current letter; inability to use this context would impair performance. Schizophrenics perform especially poorly in this and similar versions of the task.

Schizophrenic Language Deficits. Schizophrenics also show poor use of context in language processing. Chapman et al. (1964) first demonstrated this in a 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 suggested 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 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 was correct. Recently, we tested the idea that schizophrenics are restricted in the *temporal range* over which they can process linguistic context (Cohen et al., 1989). We designed a lexical ambiguity task, similar to the one used by Chapman and his colleagues, 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"), while 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. 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*; 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 their response from the list of possible meanings:

- (a) you can't keep chickens [clear screen / pause] without a PEN (weak meaning, context first)
— or —
(b) without a PEN [clear screen / pause] you can't keep chickens (weak meaning, context last)
— or —
(c) you can't sign a check [clear screen / pause] without a PEN (strong meaning, context first)
[clear screen / pause]

The meaning of the word in capital letters is:

a writing implement (strong meaning)	a fenced enclosure (weak meaning)	a kind of truck (unrelated meaning)
---	--------------------------------------	--

The results of this study (shown in Figure 7) corroborated both the Chapmans' original findings, and the explanation of their findings in terms of an inability to maintain context. Thus, as the Chapmans found, schizophrenics made significantly more dominant meaning errors than did controls when the weak meaning was correct. However, this only occurred when the context came *first* (i.e., less recently — condition B above). When context came last (i.e., more recently), schizophrenics correctly chose the weak meaning. This was the only type of error that reliably distinguished schizophrenics from controls. These findings suggest that the impairment observed in language tasks may be similar in nature to the impairments observed in attentional tasks: a difficulty in representing and maintaining context.

PFC, Context, and Dopamine

Several studies suggest that frontal areas are specifically involved in maintaining context information for the control of action. For example, at the neurophysiological level, Fuster (1980) and Goldman-Rakic (1985) have observed cells in PFC that are specific to a particular stimulus and response, and that remain active during a delay between these. They have argued that neural patterns of activity are maintained in PFC which encode the temporary information needed to guide a response. At the behavioral level, these authors and Diamond (e.g., 1989) have also reported data showing that PFC is needed to perform tasks involving delayed responses to ambiguous stimuli. Diamond has emphasized that prefrontal memory is required, in particular, to overcome reflexive or previously reinforced response tendencies in order to mediate a contextually relevant — but otherwise weaker — response.

Furthermore, it has been shown that dopaminergic innervation of PFC is necessary for this brain region to maintain contextual information. Experimental lesions in animals, or clinical lesions in humans, to this dopaminergic supply can mimic the effect of lesions to the PFC itself on behavioral tasks requiring memory for context (e.g., Brozoski et al., 1979).

Neuromodulatory Effects of Dopamine

Several anatomical and physiological observations support the idea that catecholamines such as dopamine and norepinephrine modulate information processing in the brain. Dopamine and norepinephrine neurons originate in discrete nuclei localized in the brain stem and their fibers project radially to several functionally different areas of the CNS. 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 post-synaptic effects that are consistent with a modulatory role. Most importantly, recent evidence suggests that the effect of dopamine release is not to directly increase or reduce the firing frequency of target cells (e.g., Chiodo and Berger, 1986). Rather, like norepinephrine, dopamine seems to modulate the response properties of post-synaptic cells such that both inhibitory and excitatory responses to other afferent inputs are *potentiated*. This effect has been described as an increase in the ‘signal-to-noise ratio’ of the cells’ behavior or an ‘enabling’ of its response (e.g., Foote et al., 1975).

PFC and Dopamine in Schizophrenia

The behavioral data reviewed earlier concerning schizophrenic performance deficits indicates an insensitivity to context, and a dominant response tendency. This is consistent with evidence that schizophrenia is associated with frontal lobe impairment. Schizophrenics show typical frontal lobe deficits on standard neuropsychological tests, including the Wisconsin Card Sort Test (WCST) (e.g., Malmö, 1974) and the Stroop task (as described above). In addition, imaging and electrophysiological studies suggest an atrophy and abnormal metabolism in the frontal lobes

of schizophrenics (e.g., Ingvar and Franzen, 1974). Recent studies have even demonstrated abnormal metabolism in the PFC of schizophrenics specifically during performance on tasks requiring memory for context such as the Wisconsin Card Sort Task and a variant of the CPT (Weinberger et al., 1986; R.M. Cohen et al., 1987). This work confirms that anatomic and physiological deficits of frontal cortex may indeed be associated with some behavioral deficits observed in schizophrenics.

Frontal lobe dysfunction in schizophrenia fits well with the prevailing neurochemical and psychopharmacological data concerning this illness. The PFC is a primary projection area for the mesocortical dopamine system, a disturbance of which has consistently been implicated in schizophrenia (e.g., Meltzer & Stahl, 1976). In view of these findings, several authors have proposed that reduced dopaminergic tone in PFC may be associated with hypofrontality in schizophrenia, and may be responsible for several of the cognitive deficits that have been observed (e.g., Weinberger & Berman, 1988).

Summary

We referred to evidence that schizophrenics inadequately maintain context for the control of action; that the PFC plays a role in maintaining context; that an intact mesocortical dopamine system is necessary for normal PFC function; and finally, that the mesocortical dopamine system is affected in schizophrenia. Despite a growing recognition that these observations are related, no theory has yet been proposed which explains — in terms of causal mechanisms — the relationship between disturbances in PFC and dopamine on the one hand, and behavioral deficits on the other. In the remainder of this paper, we present a set of connectionist models that simulate schizophrenic performance in the tasks described above.

Simulation of the Physiological Effects of Dopamine

In the models, the action of dopamine is simulated as a change in a parameter of the function relating a unit's input to its activation value. To do so, we first assume that the relationship between the input to a neuron and the neuron's frequency of firing can be simulated as a non-linear function relating the net input of a model unit 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., Freeman, 1979). The same experiments also indicate that small increments in excitatory drive result in greater changes in firing frequency than equivalent increments in inhibitory input. These properties can be captured by the logistic function with a constant negative bias:

$$\text{activation} = \frac{1}{1 + e^{-(\text{gain} * \text{net}) + \text{bias}}} \quad (\text{see Figure 1, Gain} = 1.0)$$

The potentiating effects of dopamine can be simulated by increasing the gain parameter of the logistic function. As Figure 1 (Gain = 2.0) illustrates, with a higher gain the unit is more sensitive to afferent signals, while 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 levels (e.g., Servan-Schreiber, Prinz & Cohen, 1990).

In order to simulate the effect of a neuromodulator, 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 models described below, decreased dopamine activity in prefrontal cortex was simulated by reducing the gain of units in the module used to represent and maintain context. In all three models, simulation of schizophrenic performance was conducted by reducing gain from a normal value of 1.0 to a lower value in the range 0.6-0.7.

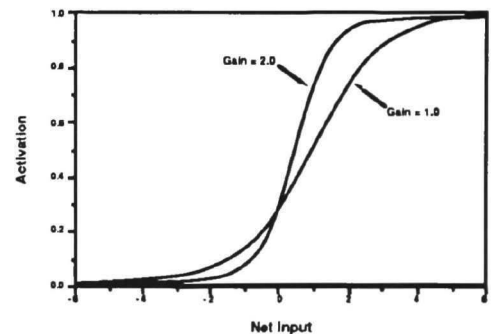


Figure 1

Simulation of the Stroop effect

Elsewhere, we have described a connectionist model of selective attention that simulates human performance in the Stroop task (Cohen, Dunbar & McClelland, in press). In brief, this model consists of two processing pathways, one for color naming and one for word reading and a task demand module that can selectively facilitate processing in either pathway (see Figure 2). Simulations are conducted by activating input units corresponding to stimuli used in an actual experiment (e.g., the input unit in the color naming pathway representing the color red) and the appropriate task demand unit. Activation is then allowed to spread through the network. This leads to activation of the output unit corresponding to the appropriate response (e.g., "red").

This simple model is able to simulate an impressive number of empirical phenomena associated with the Stroop task. It captures the four basic effects (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), as well as the influence of practice on interference and facilitation effects, the relative nature of these effects, response set effects and stimulus onset asynchrony effects (see Cohen et al., in press).

This model also exhibits behaviors that make it relevant to understanding schizophrenic disturbances of attention, and their relationship to the processing of context. The model shows how attention can be viewed as the effect that context has on selecting the appropriate pathway for responding. Here, context is provided by the task instructions. Thus, when subjects are presented with conflicting input in two dimensions (e.g., the word GREEN in red ink), they respond to one dimension and not the other, depending upon the *context* in which it appears (i.e., the task: color naming or word reading). If frontal cortex is responsible for maintaining this context, and if schizophrenia involves a disturbance of frontal lobe function, then we should be able to simulate schizophrenic performance 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 input, then we should be able to simulate this by reducing the gain of units in the task demand module.




Figure 2

Figure 3 shows the results of such a simulation, in which the gain of only the task units was reduced; all other units were unperturbed. This change in the context (task demand) module produced effects similar to those observed for schizophrenics: an increase in overall response time, with a disproportionate increase for color naming interference trials. Thus, the model shows that a lesion restricted to the mechanism for processing context can produce both an overall degradation in performance as well as the expected attentional deficit.

The model also allows us to compare the effects of this specific disturbance to those of a more general disturbance, addressing a common difficulty in schizophrenia research. It is often argued that, in the presence of a general degradation of performance in schizophrenics (e.g., overall slowing of response), it is difficult to know whether degradation in a particular experimental condition is due to a specific deficit or a more generalized one. However, this difficulty arises primarily when the mechanisms for the deficits involved have not been specified. The model provides us with a tool for doing this. Above, we described the mechanism for a specific attentional deficit related to disturbances of dopamine activity in PFC. To compare this to a more generalized deficit, we induced overall slowing in the model by decreasing the rate at which information accumulates for each unit (cascade rate); this was done for *all* units in the model (third panel in Figure 3). A lower cascade rate induced an overall slowing of response, but no disproportionate slowing in the interference condition. In contrast, the specific disturbance in context representation produced both effects. Thus, the context hypothesis provides a better account for the data than at least one type of generalized deficit. We have explored others (e.g., an increase in the response threshold), with similar results.




Figure 3

Simulation of the Continuous Performance Test

The Stroop model shows how contextual information and its attentional effects can be represented in a connectionist model, and how a specific biologically relevant disturbance in this mechanism can explain aspects of schizophrenic performance. One question we might ask is: How general are these findings? Here, we extend the principles applied 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 a demand is placed on memory for context. For example, in the CPT-double, a target consists of any consecutive re-occurrence of a letter (e.g., a 'B' immediately following a 'B'). Thus, subjects must remember the

previous letter, which provides the necessary context for responding to the subsequent one. Schizophrenics perform poorly in this task. This may be due to an impairment in the processing of context that, like deficits in the Stroop task, might be explained by a reduction of dopaminergic tone in prefrontal cortex. 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 letter identification module and a response module (see Figure 4). 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. The network was trained to associate these input patterns with the corresponding letter units in the letter identification 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. This was made possible by introducing a set of connections from the letter units back to the intermediate units. This allowed the network to store and use information about the previous as well as the current stimulus (see Cohen & Servan-Schreiber, 1989 for a more complete description of training and processing in this model). Note that there is a direct analogy between the role played by the letter units in this model, and the role played by the task demand units in the Stroop model. The representation over the letter 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. In the CPT model, however, context was determined by the previous input, and therefore changed from trial to trial.

Following training, the network was able to perform the CPT-double task perfectly for a set of 26 different stimuli. To simulate the performance of normal subjects — who typically miss on 13% of trials and false alarm on 1% of trials (see Figure 5a) — noise was added to the net input, with the amount adjusted to match the performance of the network with that of human subjects. The results of this simulation appear in Figure 5b (gain = 1.0). Then, to simulate schizophrenic performance, we disturbed processing in the letter module — which was responsible for representing and maintaining context — by decreasing the gain of these units by an amount comparable to the amount used in the Stroop simulation (0.66). The percentage of misses increased to 20%, while false alarms increased slightly to 1.1%. These numbers closely match the results of empirical observations of schizophrenic subjects.

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 memory trace required — as context — for processing the current stimulus. We assume that this memory trace 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 that we turn to next.




Figure 4

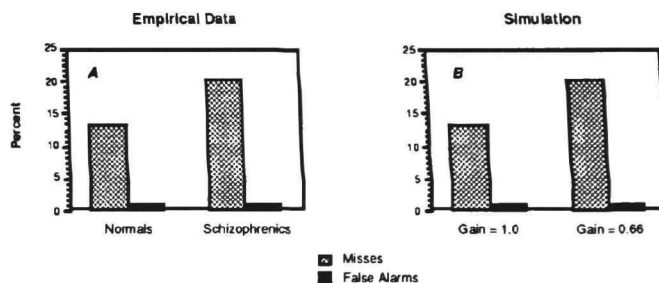


Figure 5

Simulation of Context-Dependent Lexical Disambiguation

The language model (Figure 6) incorporates elements of the two previous simulations. The network was similar to the CPT model. It was trained to associate input patterns representing lexical stimuli (e.g., the word PEN) to patterns in two output modules: a response module and a discourse module. Patterns in the response module specified the meaning of the input words (e.g., “writing implement”), while the discourse module represented the topic of the current sequence of inputs (i.e., the meaning of the sentence, rather than the meaning of individual words). As in the CPT model, there were two-way connections between the intermediate module and the context (discourse) module. Thus, once a discourse representation had been activated by an 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.

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*) more than the other (e.g., PEN → “fenced enclosure” and *FARMING*). This asymmetry of training was similar to that of the Stroop model (trained on words more than colors), with a comparable result: when presented with an ambiguous input word, the network preferentially activated the strong (more frequently trained) response and discourse representations. To permit access to the weaker meaning, the network was sometimes presented with an ambiguous word as input along with one of its associated discourse representations (e.g., PEN and *FARMING*), 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 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 effects of these training procedures was that when an ambiguous word was presented and there was no representation active over the discourse units, the output was a blend of the two meaning 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.




Figure 6

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 was achieved 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 word related to one of the meanings of an ambiguity (e.g., CHICKEN), allowing activation to spread through the network, 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 experiment, the model was presented with pairs of context and ambiguous words (representing the clauses used in the experiment) in either order. 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. At each time step of processing, a small amount of noise was added to the activation of every unit. The amount of noise was adjusted so that the simulation produced an overall error rate comparable to that of control subjects. The model's response on each trial was considered to be the meaning that was most active over the output units after the probe was presented. To simulate schizophrenic performance, we introduced a disturbance analogous to the one in the CPT model: a reduction in gain of units in the context module. The results of this simulation (shown in Figure 7, along with the results from the empirical study) show a strong resemblance to the empirical data. They demonstrate both significant effects: a) in the low gain mode, the simulation makes about as many more dominant response errors as do the schizophrenic subjects; however, b) as with the human subjects, this only occurred when context came first. (The number of unrelated errors — not shown in Figure 7 — was approximately the same in both the low gain and normal gain mode.) The model provides a clear view of this relationship between dominant response bias and memory. 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, these effects are less, and the representation is sufficient to overcome a dominant response bias. However, if time

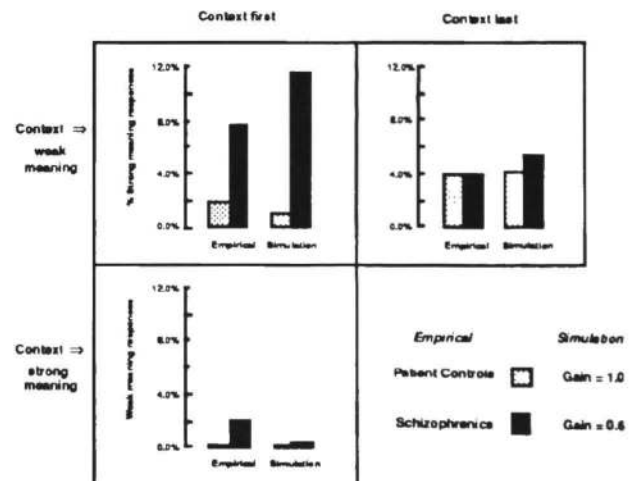


Figure 7

passes (as when context is presented first), the effects of noise accumulate, and the representation is no longer strong enough to mediate the weaker of two competing responses.

Conclusion

The three models we have presented showed how the connectionist framework can be used to link previously unrelated biological and behavioral observations. This was achieved by bringing a common set of mechanisms to bear simultaneously on physiological and psychological phenomena. Specifically, the models served several purposes: a) they simulated quantitative aspects of performance in three previously unrelated behavioral tasks; b) they elucidated the role of processing of context in both the attentional and linguistic tasks; c) they related processing of context to biological processes; and d) they showed how a specific disturbance at the biological level could account for schizophrenic patterns of performance.

References

- Abramczyk RR, Jordan DE & Hegel M (1983). "Reverse" Stroop effect in the performance of schizophrenics *Perceptual and Motor Skills* 56:99-106.
- Brozoski TJ, Brown RM, Rosvold HE, & Goldman PS (1979). Cognitive deficit caused by regional depletion of dopamine in prefrontal cortex of Rhesus monkey. *Science* 205:929-931.
- Chapman LJ, Chapman JP & Miller GA (1964). A theory of verbal behavior in schizophrenia. In Maher BA (Ed.), *Progress in Experimental Personality Research, Vol 1*. New York: Academic Press.
- Chiodo LA & Berger TW (1986). Interactions between dopamine and amino acid-induced excitation and inhibition in the striatum. *Brain Research* 375:198-203.
- Cohen JD, Dunbar K & McClelland JL (1990). On the control of automatic processes: A parallel distributed processing model of the Stroop effect. *Psychological Review*
- Cohen JD & Servan-Schreiber D (1989). A parallel distributed processing approach to behavior and biology in schizophrenia. Technical Report AIP-100. Department of Psychology, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213.
- Cohen JD, Targ E, Kristoffersen T & Spiegel D (1990). The fabric of thought disorder: Disturbances in the processing of context. In submission; available upon request.
- Cohen RM, Semple WE, Gross M, Nordahl TE, DeLisi LE, Holcomb HH, King AC, Morihisa JM & Pickar D (1987). Dysfunction in a prefrontal substrate of sustained attention in schizophrenia. *Life Sciences* 40:2031-2039.
- Diamond A (1989), *The development and neural bases of higher cognitive functions*. New York: NY Academy of Science Press
- Foote SL, Freedman R & Oliver AP (1975). Effects of putative neurotransmitters on neuronal activity in monkey auditory cortex. *Brain Research* 86:229-242.
- Freeman WJ (1979). Nonlinear gain mediating cortical stimulus-response relations. *Biological Cybernetics* 33:243-247.
- Fuster JM (1980). *The prefrontal cortex*. Raven Press.
- Goldman-Rakic PS (1985). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. *Handbook of Physiology - The Nervous System* 5:373-417.
- Ingvar DH & Franzen G (1974). Abnormalities of cerebral flow distribution in patients with chronic schizophrenia. *Acta Psychiatrica Scandinavia* 50:425-462.
- Kornetsky C (1972). The use of a simple test of attention as a measure of drug effects in schizophrenic patients. *Psychopharmacologia (Berl.)* 8:99-106.
- Malmo HP (1974). On frontal lobe functions: Psychiatric patient controls. *Cortex* 10:231-237.
- Meltzer HY & Stahl SM (1976). The dopamine hypothesis of schizophrenia: a review. *Schizophrenia Bulletin* 2:19-76.
- Nuechterlein KH (1984). Information processing and attentional functioning in the developmental course of schizophrenic disorders. *Schizophrenia Bulletin* 10:160-203.
- Servan-Schreiber D, Printz H, and Cohen JD (1990) A neural network model of catecholamine effects: Enhancement of signal detection performance is an emergent property of changes in individual unit behavior. in DS Touretzky (ed.) *Advances in Neural Information Processing Systems 2*. San Mateo, CA: Morgan Kaufman.
- Wapner S & Krus DM (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 DR & Berman KF (1988). Speculation on the meaning of metabolic hypofrontality in schizophrenia. *Schizophrenia Bulletin* 14:157-168.
- Weinberger DR, Berman KF & Zec RF (1986). Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia: I. Regional cerebral blood flow evidence. *Archives of General Psychiatry* 43:114-125.