Cognitive, Affective, & Behavioral Neuroscience
2002, 2 (3), 187-213

The impact of synaptic depression following brain damage: A connectionist account of "access/refractory" and "degraded-store" semantic impairments

STEPHEN J. GOTTS and DAVID C. PLAUT
Carnegie Mellon University, Pittsburgh, Pennsylvania
and Center for the Neural Basis of Cognition, Pittsburgh, Pennsylvania

Neuropsychological studies of patients with acquired semantic impairments have yielded two distinct and contrasting patterns of performance in a spoken-word/picture–matching task (Warrington & Cipolotti, 1996). Patients labeled *access/refractory* are strongly influenced by presentation rate, semantic relatedness of distractors, and repetition, yet they seem relatively unaffected by lexical frequency. *Degraded-store* patients, on the other hand, are strongly affected by lexical frequency but are less affected by presentation rate, semantic relatedness, or repetition. Our account of these patterns of performance is based on the distinction between two different types of neurological damage: (1) damage to neuromodulatory systems that function to amplify neural signals while suppressing normal refractory-like effects and (2) damage to connections between groups of neurons that encode semantic information and are sensitive to frequency/familiarity. We present a connectionist model that learns to map spoken-word input to semantic representations and that incorporates a particular form of neural refractoriness referred to as *synaptic depression*, as well as a simple form of neuromodulation. We show that the model is capable of accounting for the contrasting patterns of semantic impairment under these two different forms of damage and, furthermore, demonstrate how it is capable of handling several documented cases that are exceptions to the basic patterns of impairment. Several predictions and limitations of the present model are discussed.

An important goal in the enterprise of explaining human cognition is to characterize the nature and structure of *semantic knowledge*—general conceptual, functional, and factual knowledge about the world that is gradually acquired and abstracted over many individual experiences (Tulving, 1972). Knowing that a *cat* is a member of a larger class of things that we call animals, that a *hammer* is a tool used to pound nails into wood, and that the earth is round and not flat are all examples of semantic knowledge. The semantic system is thought to relate information across different sensory and motor modalities, according it a central role in a wide variety of important cognitive behaviors and tasks, such as language and visual–motor interaction (e.g., Caramazza, Hillis, Rapp, & Romani, 1990; Morton, 1981; Plaut, in press; Riddoch, Humphreys, Coltheart, & Funnell, 1988; Shallice, 1988). In developing theories of normal semantic processing and representation, one potentially useful set of empirical constraints comes from studies of individuals with brain damage that selectively impairs performance on tasks requiring semantic knowledge. Indeed, these studies have produced some of the most counterintuitive findings observed in neuropsychology, and many researchers have taken this to suggest that such findings are particularly useful constraints on theorizing (Shallice, 1988).

One of the more controversial sets of constraints centers around contrasting patterns of semantic impairment observed in two different populations of brain-damaged patients. One group of patients is strongly influenced by word frequency and consistently identifies the same stimuli correctly. The other group of patients is less influenced by frequency and performs inconsistently, in that they are temporarily worse under conditions of short intertrial intervals, close semantic distractors, and repeated stimuli. This contrasting pattern has led some researchers to propose that there are two distinct forms of semantic impairment, one resulting from damage directly to semantic representations (a *degraded-store* impairment) and the other resulting from damage to semantic access processes that makes them abnormally refractory (an *access* or *refractory* impairment; e.g., Warrington & McCarthy, 1983; Warrington & Shallice,

S.J.G. is supported by an NSF Graduate Research Fellowship. D.C.P. is supported by NIH Grant MH55628. We thank Tim Shallice, Lisa Cipolotti, Jay McClelland, Bobby Klatzky, Carson Chow, and members of the Carnegie Mellon PDP Research Group for discussions and/or helpful comments on earlier drafts of the paper. Correspondence concerning this article should be addressed to S. J. Gotts, Department of Psychology, Baker Hall, Carnegie Mellon University, Pittsburgh, PA 15213 (e-mail: gotts@cnbc.cmu.edu).

1979). Although a number of cases have been documented that fit broadly within this framework, it has been criticized on both empirical and theoretical grounds. Rapp and Caramazza (1993) pointed out that the relevant patients have not been assessed on all of the same stimulus factors and that some of them appear to exhibit a mixing of access and degraded-store patterns. They also argued that the theory of semantic access and representation implicit in the proposal was too underspecified to be scientifically useful. Although Warrington and Cipolotti (1996) have responded to the lack of empirical validity, the theoretical concerns of Rapp and Caramazza have yet to be completely allayed.

The present paper provides a computational theory of semantic processing that is capable of addressing the variety of patient effects associated with the access/degraded-store distinction. Our account is based on a distinction between two different types of neurological damage that can selectively affect the semantic system. One type involves damage to neuromodulatory systems that normally function to enhance neural signals that are otherwise attenuated by automatic refractory processes. Such damage can have a selective effect because the fiber pathways to different cortical regions from subcortical neuromodulatory centers are broadly segregated and can be selectively disrupted (e.g., Selden, Gitelman, Salamon-Murayama, Parrish, & Mesulam, 1998). Another type involves damage directly to neurons that encode semantic information. We present a neural network model that learns to map spoken-word input to semantic representations and incorporates a particular form of neural refractoriness, referred to as *synaptic depression*. A simple form of neuromodulation that is consistent with known effects of acetylcholine and norepinephrine serves to amplify activity while reducing refractory effects owing to synaptic depression. Damage to frequency-sensitive connection strengths that spares neuromodulation yields a degraded-store pattern. The model produces a strong frequency effect, is consistent in which words it correctly identifies, and shows little effect of presentation rate or repetition. An access/refractory pattern, on the other hand, is produced by damage that reduces the presence of neuromodulatory factors. Synaptic depression is stronger, resulting in large effects of presentation rate and repetition, as well as inconsistent responding. These effects are most severe when the stimuli being repeated are semantically related, because such stimuli activate many of the same neurons and synaptic depression can build up across stimuli. Synaptic depression is also stronger when activity in the network is initially higher (as is the case for high-frequency words), counteracting normal frequency effects. Under different combinations of neuromodulatory damage and damage to connections, it is possible to account for patients who do not fit cleanly into either patient group.

In the remainder of the paper, we first will review patient data associated with access and degraded-store semantic impairments. This will be followed by a discussion of refractory processes in the normal brain and how they interact with neuromodulatory agents, such as acetylcholine and norepinephrine. We then will present the model and simulation experiments, followed by a discussion of the model's predictions and limitations.

Access and Degraded-Store Semantic Impairments

A range of studies have demonstrated selective impairments on semantic tasks. Warrington (1975) documented the performance of 3 patients suffering from progressive temporal lobe atrophy. All 3 patients were markedly impaired on matching pictures to words, naming pictures, and giving verbal definitions to orally presented words, whereas they performed at near-normal levels on tests of general intelligence, working memory, language functioning, and perception. They were quite unlike amnesic patients, in that they were well oriented in time and place, did not have a tendency to repeat conversational topics, and were able to refer forward and backward to important events in their lives. On further examination, it became clear that the patients were particularly impaired at identifying stimuli with low-frequency names and that they were often able to provide superordinate category information for stimuli that they could not identify. A follow-up investigation of one of the patients (E.M.) showed that she consistently identified the same items correctly or incorrectly over repeated testing (Coughlan & Warrington, 1981). Frequency effects and the relative preservation of general semantic knowledge common to many objects have since been broadly reported in the literature on acquired semantic impairments in patients with a variety of etiologies, such as semantic dementia (probable Pick's disease), herpes simplex encephalitis, and dementia of the Alzheimer's type (e.g., Breedin, Saffran, & Coslett, 1994; Chertkow, Bub, & Seidenberg, 1989; Cipolotti & Warrington, 1995; Done & Gale, 1997; Hodges, Graham, & Patterson, 1995; Hodges, Salmon, & Butters, 1992; Lambon Ralph, Graham, Ellis, & Hodges, 1998; Parkin, 1993; Sartori & Job, 1988; Silveri & Gainotti, 1988; Warrington & Shallice, 1984; although see Funnell, 1995, and Tyler & Moss, 1998, for exceptions). Although response consistency has been investigated less often, in studies in which it has been examined, the co-occurrence of consistency, frequency effects, and the relative preservation of general semantic knowledge has been observed (e.g., Chertkow et al., 1989; Hodges et al., 1992; Silveri & Gainotti, 1988; Warrington & Shallice, 1984). Consistency and frequency effects have also been observed to co-occur in anomic aphasic patients without marked semantic impairments (e.g., Howard, 1995; Lambon Ralph, 1998).

In contrast, other patients with semantic impairments have exhibited highly inconsistent, variable performance across repeated testing, some with weaker or nonexistent effects of word frequency. For example, Warrington and Shallice (1979) characterized the performance of a dyslexic patient (A.R.) who had difficulty naming letters, words, or objects on visual presentation but was able to

name from verbal description. A.R.'s word-reading performance was unaffected by lexical frequency and exhibited a high degree of inconsistency across repeated testing. He was able to provide semantic information about words that he could not read, and his reading performance improved significantly when he was cued with a semantically related auditory probe word. Warrington and McCarthy (1983) studied the semantic impairment of a global aphasic patient (V.E.R.) who had suffered a large infarction of the left middle cerebral artery that resulted in profound language comprehension and production problems. Using a word/picture-matching task to assess comprehension, they found that she was strongly influenced by presentation rate, performing much worse with shorter delays between stimuli (2 vs. 30 sec); she exhibited a weak but significant effect of lexical frequency, performed worse when stimuli within a block of trials were all highly related semantically, and performed inconsistently across stimuli repeated within a block, performing well at the beginning of the block and progressively worse across repetitions (termed a *serial position* effect), recovering over longer delays between blocks. The fact that V.E.R.'s performance was relatively spared on picture–object matching led Warrington and McCarthy (1983) to argue that her deficit was primarily one of auditory verbal comprehension. This basic pattern of performance has been replicated with a number of other global aphasic patients (Y.O.T., Warrington & McCarthy, 1987; J.M., Forde & Humphreys, 1995), along with the lack of a significant frequency effect in some cases (H.E.C., Cipolotti & Warrington, 1995; M.E.D., McNeil, Cipolotti, & Warrington, 1994; A1 and A2, Warrington & Cipolotti, 1996; see also Forde & Humphreys, 1995, Experiment 6).

Warrington, Shallice, and colleagues proposed that these two somewhat different patterns of impairment—consistent responding with marked frequency effects versus inconsistent responding with weak or absent frequency effects—might result from two substantively different types of semantic impairment (Shallice, 1988; Warrington & McCarthy, 1983; Warrington & Shallice, 1979, 1984). In particular, they drew a distinction between damage directly to semantic representations, referred to as a *degraded-store* deficit, and damage to modality-specific access pathways and processes that spared semantic representations themselves, referred to as an *access* deficit. They reasoned that a degraded-store deficit should be accompanied by consistent responding, significant frequency effects, and a hierarchical breakdown of semantic knowledge, if one assumes a permanent loss of knowledge and a more robust representation of familiar stimuli and general, superordinate category information. Damage to access processes, on the other hand, might lead to somewhat stochastic access from trial to trial, giving rise to inconsistent responding across repetitions and reducing frequency effects, since the stochastic influences could be unrelated to word frequency and other stimulus characteristics. They further hypothesized that spared representations under an access deficit should support priming/cuing effects when primed/cued from a different modality, whereas severe damage to semantic representations should not. Warrington and McCarthy (1983, 1987; see also Cipolotti & Warrington, 1995) later refined the notion of a semantic access deficit to involve *refractoriness*, a reduction in the ability to utilize the semantic system efficiently for a period of time following activation. They claimed that refractoriness was sufficient to explain characteristics of the global aphasic performance, such as the effects of presentation rate and serial position, as well as inconsistent responding and reduced or absent frequency effects. It would also be possible to explain the effects of semantic relatedness if one were to assume a gradient of refractoriness within a semantic category.[1] A summary of the patient data relating to access/refractory and degraded-store deficits is provided in Table 1.

Although the theoretical distinction between deficits of access and deficits of storage may seem reasonable and justified, some researchers have taken issue with the distinction and have raised a number of challenges. Notably, Rapp and Caramazza (1993) put forward two strong criticisms: (1) The empirical validity of the distinction was far from established, insofar as several of the performance characteristics/criteria had not been assessed in both proposed patient types and several cases existed showing a mixing of the access and the degraded-store patterns, and (2) in the absence of a more specific theoretical proposal as to the nature of stored representations and access mechanisms, the distinction is of little scientific value. On the first criticism, they reviewed the performance of 2 putative degraded-store patients, P.W. (Howard, 1985) and K.E. (Hillis, Rapp, Romani, & Caramazza, 1990), who responded with a high degree of consistency, yet who showed no effect of frequency (see Table 1). Similarly, six Alzheimer's patients documented by Chertkow et al. (1989) showed several characteristics of the degraded-store pattern yet exhibited strong semantic priming effects. Several putative access patients also inappropriately showed effects of frequency (e.g., Patients C.A.V., V.E.R., and J.C.U.; Table 1). They argued that these instances of criteria *mixing* are problematic for the distinction. If one wants to suggest either (1) that perhaps these cases are *mixed disorders*, having damage to both semantic representations and modality-specific access pathways, or (2) that the list of criteria is, perhaps, incomplete or some of the criteria have been erroneously included, it becomes very difficult to make progress scientifically. Indeed, Shallice (1988) admitted that what is needed in order to address some of these issues is a well-established model of the semantic access process.

Warrington and Cipolotti (1996) responded to the lack of empirical validity of the distinction, at least for a subset of patients and criteria. They demonstrated a contrasting pattern of impairment in the comprehension performance of 6 patients: 2 global aphasic patients who suffered a large left-hemisphere vascular lesion (Patient A1) and an intrinsic cerebral tumor (Patient A2) and 4 semantic dementia patients who suffered focal left tem-

Table 1
Summary of Degraded-Store and Access/Refractory Patient Data

Patient	Etiology	Task(s)	*Rate* (*F* < *S*)	*Semantic Relatedness*	*Frequency* (*HF* > *LF*)	*Consistency*	*Serial Position*	Priming/ Cuing	Hierarchical (Super > Sub)
					Degraded-Store				
A.B.	Cortical atrophy	Defining, AFC			$p < .001$				$p < .001$
E.M.	Cortical atrophy	Defining, AFC			$p < .001$	$p < .001$ ($C = .58$)			$p < .01$
C.R.	Cortical atrophy	Defining, AFC			$p < .01$				$p < .001$
J.B.R.	Herpes simplex encephalitis	Name/define words and pictures			$p < .001$	$p < .001$ ($C = .61$)			$p < .001$
S.B.Y.	Herpes simplex encephalitis	Name/define words and pictures			$p < .02$	$p < .05$ ($C = .40$)			$p < .001$
L.A.	Herpes simplex encephalitis	Name/define words and pictures			$p < .01$	$p < .001$ ($C = .54$)			Super and semantic errors
$n = 22$	DAT (Alzheimer's)	Naming, sorting, WPM, defining			$p < .0001$	$p < .0001$			Impaired at sub (L2, L3) but not super
S1	Atrophy (left hemisphere)	WPM	n.s.	$p < .001$ (between category)	$p < .001$	$p < .001$	n.s.		
S2	Atrophy (left temporal lobe)	WPM	n.s.	$p < .001$ (between category)	$p < .001$	$p < .001$	n.s.		
S3	Atrophy (left temporal lobe)	WPM	n.s.	$p < .01$ (within and between category)	$p < .001$	$p < .01$	n.s.		
S4	Atrophy (left temporal lobe)	WPM	n.s.	$p < .001$ (between category)	$p < .001$	$p < .001$	n.s.		
P.W.	**Left MCA stroke**	**Naming pictures, reading**			**n.s.**	**Highly consistent**			
***n* = 6**	**DAT (Alzheimer's)**	**Picture naming WPM, AFC, lexical decision**			***p* < .01**	***p* < .005**		***p* < .001**	**Impaired at sub (<.0005) but not super**
K.E.	**Left MCA stroke (frontoparietal)**	**Oral/written naming, reading**			**n.s.**	***ON:* < .01** ***WN:* < .0001** ***Read:* < .0005**			

Table 1 (Continued)

Patient	Etiology	Task(s)	*Rate* (*F* < *S*)	*Semantic Relatedness*	*Frequency* (*HF* > *LF*)	*Consistency*	*Serial Position*	Priming/ Cuing	Hierarchical (Super > Sub)
				Access/Refractory					
A.R.	Intracerebral abscess (left parietal)	Reading			n.s.	n.s. (C = .31)		p < .01	n.s.
Y.O.T.	Left MCA stroke (temporoparietal)	WPM	p < .001	p < .02 (within category)		n.s.	p < .01		
M.E.D.	Lesions (left posterior frontoparietal)	Spoken–written word matching	p < .001	p < .01 (within category)	n.s.	n.s.	p < .02		
J.M.	Stroke (left temporoparietal)	WPM	p < .02	p < .05 (within category)	n.s.	n.s.	p < .05		
H.E.C.	Stroke (left hemisphere)	WPM	p < .01	p < .02 (within category)	n.s.	n.s.	p < .02		
A2	Multifocal cerebral tumor	WPM	p < .01	p < .01 (within category)	n.s.	n.s.	p < .0001		
A1	**Left MCA stroke**	**WPM**	***p* < .001**	***p* < .001 (within and between category)**	**n.s.**	**n.s.**	**n.s.**		
C.A.V.	**Cerebral tumor (left posterior)**	**Reading**			***p* < .01**	**n.s. (C = .23)**		**Strong priming**	
V.E.R.	**Left MCA stroke (frontoparietal)**	**Word–object matching**	***p* < .001**	***p* < .01 (within category)**	***p* < .01**	**n.s.**	***p* < .05**		
J.C.U.	**Haematoma (left frontotemporal)**	**Picture naming**			***p* < .025**	**n.s.**		***p* < .001**	

Note—The five criteria validated by Warrington and Cipolotti (1996) that are most relevant to subsequent model performance (rate, semantic relatedness, frequency, consistency, and serial position) are shown in italics. Patients who have exhibited "mixed" patterns are shown in boldface (see the text for explanation). Significance values shown either were those originally reported or, in some cases, were calculated on the basis of the data provided. Empty spaces denote the criteria that were not evaluated for a given patient (or where information was unavailable). The list of behavioral tasks used with each patient represents the subset used to acquire one or more of the included results. WPM, word–picture matching; AFC, alternative forced choice; *F*, fast; *S*, slow; *HF*, high frequency; *LF*, low frequency; Super, superordinate; Sub, subordinate. Patients: A.B., C.R., Warrington (1975); E.M., Warrington (1975); Coughlan and Warrington (1981); J.B.R., S.B.Y., Warrington and Shallice (1984); L.A., Silveri and Gainotti (1988); n = 22, Hodges, Salmon, and Butters (1992); A1, A2, S1–S4, Warrington and Cipolotti (1996); P.W., Howard (1985); Patterson and Marcel (1977); n = 6, Chertkow, Bub, and Seidenberg (1989); K.E., Hillis, Rapp, Romani, and Caramazza (1990); A.R., Warrington and Shallice (1979); Y.O.T., Warrington and McCarthy (1987); M.E.D., McNeil, Cipolotti, and Warrington (1994); J.M., Ferrand and Humphreys (1996); Forde and Humphreys (1995, 1997); H.E.C., Cipolotti and Warrington (1995); Warrington and Cipolotti (1996); C.A.V., Warrington (1981); V.E.R., Warrington and McCarthy (1983); J.C.U., Howard and Orchard-Lisle (1984).

poral lobe atrophy with a diagnosis of probable Pick's disease (Patients S1–S4). On a spoken-word/picture–matching task, the performance of the 2 global aphasic patients was worse with a faster rate of presentation (1 vs. 15 sec between trials) yet was unaffected by lexical frequency (see Table 1). When trials presented under a fast rate were considered, performance was inconsistent across within-block repetition, and one patient (A2) exhibited a significant serial position effect, with poorer performance following repetition.[2] In contrast, the semantic dementia patients performed much worse with low-frequency stimuli and were consistent in which stimuli they correctly identified. They were unaffected by presentation rate or the within-block repetition of stimuli. Both groups of patients performed worse when stimuli within a block were semantically related, although the semantic distance that elicited effects was different for the two groups (within- and between-category manipulations yielded significant effects in the global aphasics, but only between-category manipulations produced effects in the semantic dementia patients). Critically, the two groups of patients did not differ in overall level of correct performance on the task, undermining any unifying explanation that appeals to severity of impairment.

The fact that Warrington and Cipolotti (1996) did observe a contrasting pattern of impairment when the two sets of patients were tested on the same task and on the same set of items provides a response to Rapp and Caramazza's (1993) criticism that the distinction lacked empirical validity, at least for the criteria of *rate*, *semantic relatedness*, *frequency*, *serial position*, and *consistency*. With regard to the criticism that theories of access and representation were too underspecified to be scientifically useful, Warrington and Cipolotti argued that the concepts of refractoriness and permanent loss of knowledge are specific enough to account for the pattern of effects generated by both sets of patients. They also proposed a neurophysiological basis for the two different types of impairment: (1) They suggested that a refractory impairment may result from vascular lesions and tumors because such damage could conceivably give rise to anoxia, producing increased neural refractory periods and greater probability of conduction failure, and (2) they suggested that a degraded-store impairment may result from actual structural damage to neurons and cell death, rather than from conduction failure.

The promise of this proposal is that a single process—namely, neural refractoriness—may ultimately account for the entire pattern of effects associated with the global aphasic patients. Permanently damaged representations would also appear to account for the lack of many of these effects in semantic dementia patients. However, the nature of process and representation is still left largely unspecified. As a result, it is not entirely clear whether the proposal can capture the full pattern of effects. For example, whether or not a refractory impairment predicts the lack of a frequency effect will depend on the specifics of how refractoriness is instantiated and how words are processed and represented. Without a more explicit characterization, it is also unclear to what extent a mixing of access and degraded-store patterns raises problems for the account. There are examples of mixed patterns even when one considers only global aphasic performance (see Patients A1 and V.E.R.; Table 1). Do such mixed patterns undermine the proposal, or do they lie as points on a continuum within its scope? What is needed to address these shortcomings is an explicit model of semantic processing and a more formal characterization of refractory processes in that model.

In our view, Warrington and Cipolotti (1996) took an important first step in suggesting a neurophysiological basis for the differences observed across patients. Our account of these phenomena builds on this first step and appeals to a range of neurophysiological findings in humans and animals, placing them within a computationally explicit neural network model of semantic processing. Indeed, we will argue that attempting to understand the neurophysiological basis of these semantic effects provides critical insight into the functional differences between patients and, hence, into the functional principles of the normal semantic system.

We propose that an access/refractory pattern results from damage to neuromodulatory systems that normally reduce a common neural refractory process known as synaptic depression. Synaptic depression is mechanistically distinct from anoxia and is expected to influence information processing in the neurologically intact brain, as well as under conditions of neurological damage. By formulating a computational instantiation of synaptic depression and neuromodulation, we are able to evaluate quantitatively the proposal's ability to address Warrington and Cipolotti's (1996) empirical findings for the validated criteria, as well as to address departures of individual patients from the basic contrasting pattern. In order to better motivate the inclusion of these processes in our model, we will now discuss refractory processes in the normal brain and how they are affected by neuromodulation.

Refractory Processes in the Normal Brain

In a broad sense, a neural refractory process is any process that is dependent on neural activity and leads to the temporary reduction of subsequent neural responses. Warrington and Cipolotti (1996) have pointed to anoxia as a refractory process that might underlie the impaired performance of access/refractory patients. Although this may be a plausible suggestion in the present context, there are other refractory processes in the neurologically intact brain that have been studied in some detail, using functional neuroimaging and extracellular and intracellular neural recording techniques. Interestingly, many of these studies report effects that are consistent with a refractory process's occurring at the same time scale as effects in access/refractory patients and in a variety of relevant neocortical brain regions. For example, in a recent event-related fMRI study with human subjects, Jiang, Haxby, Martin, Ungerleider, and Parasuraman (2000) demonstrated that, during performance of a delayed matching-to-sample task, blood flow (and presumably neural ac-

tivity) in the ventral temporal, occipital, and parietal cortices was reduced within a block of stimuli for repeated items. This blood flow decrease was stimulus specific, built up over four to five stimulus repetitions, occurred regardless of whether the repeated stimuli were targets or distractors, and largely recovered between blocks of trials. Similar blood flow decreases have been observed following stimulus repetition in a number of other fMRI and PET studies with humans (e.g., Buchel, Coull, & Friston, 1999; Buckner et al., 1995; Schacter, Alpert, Savage, Rauch, & Albert, 1996; Wagner, Maril, & Schacter, 2000), and neural recording experiments with awake behaving and anesthetized monkeys have documented comparable decreases in firing rate following repetition (e.g., Baylis & Rolls, 1987; Haenny & Schiller, 1988; Li, Miller, & Desimone, 1993; Miller, Gochin, & Gross, 1991; Miller, Li, & Desimone, 1991, 1993; Muller, Metha, Krauskopf, & Lennie, 1999). In addition, neural recording experiments have revealed that firing rate decreases are greater for stimuli that initially induce higher firing rates (Li et al., 1993; Miller et al., 1993). The close association of stimulus repetition with reduced neural activity has led some researchers to refer to this phenomenon as *repetition suppression* (Desimone, 1996; Wiggs & Martin, 1998).

An effect that is present at individual synapses in the neocortex appears to mirror most, if not all, of the empirical properties of repetition suppression. When an individual presynaptic neuron fires repetitively, its effect on a postsynaptic neuron decreases, an effect known as *synaptic depression* (e.g., Abbott, Varela, Sen, & Nelson, 1997; Tsodyks & Markram, 1997; Varela, Song, Turrigiano, & Nelson, 1999). Synaptic depression is present at both excitatory and inhibitory synapses throughout the neocortex in a wide variety of animal species (e.g., the primary visual cortex, Abbott et al., 1997; Finlayson & Cynader, 1995; Varela et al., 1997; Varela et al., 1999; the sensorimotor cortex, Thomson, Deuchars, & West, 1993; the somatosensory cortex, Galarreta & Hestrin, 1998; Tsodyks & Markram, 1997). Because synaptic depression occurs at individual synapses, it is likely to yield a stimulus-specific neural response decrement: Only those synapses that have been recently activated in the processing of a previous stimulus will be depressed following repetition. Synapses involved in the processing of other stimuli will be unaffected. Studies of the time course of synaptic depression suggest that recovery largely occurs within 3–4 sec of stimulation, although complete recovery can require as much as 1–2 min (e.g., Finlayson & Cynader, 1995; Galarreta & Hestrin, 1998; Tsodyks & Markram, 1997; Varela et al., 1997; Varela et al., 1999). If stimuli are repeated within the time window of recovery, synaptic depression can build up with repetition. Like repetition suppression effects, synaptic depression appears to be very automatic, in that it requires only presynaptic activity to occur and is commonly observed in experiments with cortical slices that have been removed from the brain. The degree of synaptic depression that is observed is roughly proportional to the firing rate of the presynaptic cell (i.e., the greater the presynaptic firing rate, the stronger the depression; see Abbott et al., 1997, Tsodyks & Markram, 1997, Varela et al., 1997, and Varela et al., 1999, for detailed model fits). Indeed, synaptic depression appears to depend on the probability of presynaptic transmitter release. Synapses for which transmitter release is less likely (given presynaptic spiking) undergo less depression, and pharmacological manipulations that reduce transmitter release also reduce synaptic depression (Tsodyks & Markram, 1997). Although the biological mechanisms that underlie synaptic depression are not completely clear, potential candidates include presynaptic transmitter depletion, down-regulation of transmitter release by presynaptic autoreceptors, and other cumulative inactivation of presynaptic release processes (e.g., McLean & Palmer, 1996; Senn, Markram, & Tsodyks, 2001; Varela et al., 1999).

Neuromodulation and Refractory Processes

Neuromodulatory systems in the brain differ from standard chemical neurotransmission in that they exert slower, longer-lasting, and often spatially more diffuse effects on cell function. Neuromodulators can up- or down-regulate transmitter release, resting membrane potential, adaptation effects that reduce spiking, the rate of synaptic modification, and a whole host of other cellular processes on time scales ranging from tens of milliseconds to hours (see Hasselmo, 1995, for a review). Of the various neuromodulators, acetylcholine and norepinephrine are particularly relevant for the present discussion, because they are known to reduce excitatory and inhibitory neurotransmitter release in a variety of cortical and neocortical regions (e.g., Brocher, Artola, & Singer, 1992; Dodt, Pawelzik, & Zieglgansberger, 1991; Hasselmo & Bower, 1992; Jahr & Nicoll, 1982; Pitler & Alger, 1992; Tsodyks & Markram, 1997; Vidal & Changeux, 1993). Because synaptic depression depends on the probability of transmitter release, acetylcholine and norepinephrine would be expected to reduce synaptic depression in the neocortex, an effect that has been confirmed experimentally for acetylcholine (e.g., Gil, Connors, & Amitai, 1997; Tsodyks & Markram, 1997).[3] At the same time, both substances block firing-rate adaptation effects that normally attenuate spiking rates, leading to higher and more sustained spiking (Barkai & Hasselmo, 1994; Krnjevic, Pumain, & Renaud, 1971; Madison & Nicoll, 1984, 1986; Schwindt, Spain, & Crill, 1992; Woody & Gruen, 1987). These two basic actions, combined with the findings that acetylcholine and norepinephrine enhance the magnitude of synaptic plasticity effects (e.g., Brocher et al., 1992; Hopkins & Johnston, 1988; Huerta & Lisman, 1993) and reduce transmitter release more at intracortical than at feedforward synapses (e.g., Gil et al., 1997; Hasselmo & Bower, 1992; Hasselmo, Linster, Patil, Ma, & Cekic, 1997), suggest that both neuromodulators enhance the processing and learning of bottom-up sensory inputs. Given these cellular actions, it is perhaps not surprising that acetylcholine and norepinephrine are implicated in attentional and memory processing

(e.g., Coull, 1998; Hasselmo, 1995; Marrocco & Davidson, 1998; Robbins, 1997). Norepinephrine also appears to be integral to arousal, detection of changes in the external environment, aspects of emotional processing, and executive function (Aston-Jones & Bloom, 1981; Coull, 1994, 1998; Ressler & Nemeroff, 1999; Servan-Schreiber, Printz, & Cohen, 1990). Increased levels of acetylcholine appear to be associated with the presentation of novel or behaviorally relevant sensory stimuli (e.g., Acquas, Wilson, & Fibiger, 1996; Butt, Testylier, & Dykes, 1997; Miranda, Ramirez-Lugo, & Bermudez-Rattoni, 2000).

The brain's supply of acetylcholine and norepinephrine is provided by subcortical nuclei of the basal forebrain and brain stem (a series of basal forebrain nuclei, Ch1–Ch4, along with the reticular formation in the brain stem, supplies acetylcholine; the locus coeruleus in the brainstem supplies norepinephrine). These nuclei send axonal projections up through the white matter tracts in fiber bundles that broadly innervate regions of the cortex. This fact, combined with the observation that these diffuse projections release their chemicals at axonal varicosities that affect many synapses simultaneously in a brain region, has led to the conclusion that acetylcholine and norepinephrine are likely to be involved in more global shifts of the information-processing state (e.g., Hasselmo, 1995). However, at least in the case of acetylcholine, separate projections are sent to different neocortical areas and are spatially segregated to the extent that damage to white matter owing to disease processes or stroke might create a much more selective and localized neuromodulatory deficit (e.g., Selden et al., 1998).

SIMULATION

Our account of the contrasting patterns of behavioral effects associated with the access/degraded-store distinction hinges on the notion that there can be two substantively different types of neurological damage that can selectively impair semantic processing. One type involves damage to neuromodulatory systems that normally function to enhance neural signals that are otherwise attenuated by synaptic depression. Such damage can selectively affect performance in semantic tasks because the ascending neuromodulatory fiber pathways to cortical regions subserving semantic processing (e.g., the temporal cortex) are broadly segregated in the white matter and can be selectively disrupted (e.g., Selden et al., 1998). Another type involves damage to the neurons in semantic brain regions and to the cortical fibers that connect them.

We present a connectionist model that learns to map the phonology of artificially generated spoken words to their meanings. The model is composed of relatively simple, neuron-like processing units that engage in parallel interactions by way of weighted connections. Units in the model are organized into groups or layers that represent the different types of information to be associated. These groups are intended to be roughly construed as the different neocortical brain regions subserving phonological and semantic processing. The model incorporates synaptic depression and some basic neuromodulatory actions of acetylcholine and norepinephrine that influence depression and overall neural excitability. According to our account, following training, damage to connection strengths that spares neuromodulation should yield a degraded-store pattern. The model should produce a strong frequency effect, because the damaged connections are sensitive to the training frequency of particular words. Spared neuromodulation should reduce the impact of synaptic depression, yielding consistent performance across repetitions and little effect of presentation rate. An access/refractory pattern, on the other hand, should be produced under damage that reduces the presence of neuromodulatory factors. The impact of synaptic depression will be stronger, resulting in large effects of presentation rate and repetition, as well as inconsistent responding. These effects should be most severe when the stimuli being repeated are semantic associates, because semantically similar stimuli activate many of the same neurons and synaptic depression can build up across stimuli. Synaptic depression should also be stronger when activity in the network is initially higher (as is the case for high-frequency words), counteracting normal frequency effects.

Rather than modeling these behavioral effects in the abstract, a specific data set was chosen: Experiment 2 of Warrington and Cipolotti (1996). This experiment (along with Experiment 3) directly contrasts the effects of rate, semantic relatedness, frequency, consistency, and serial position for both groups of patients on the same task and on the same stimuli. For the sake of simplicity, we chose to model only the auditory word comprehension aspects of the spoken-word/picture–matching task. This was largely motivated by observations that access/refractory patients tend to perform at or near ceiling on picture/picture and picture/object matching, yet poorly with spoken-word/picture, written-word/picture, and spoken-word/written-word matching (Cipolotti & Warrington, 1995; Forde & Humphreys, 1995; McNeil et al., 1994; Warrington & McCarthy, 1983; although see Forde & Humphreys, 1997). The implications of this simplification will be considered more fully in the General Discussion section.

Network Architecture

The architecture for the present model is graphically depicted in Figure 1. Thirty phonological units representing spoken-word input are connected in a feedforward manner to 200 hidden units, which are in turn connected in a feedforward manner to 150 semantic units. The phonological units are roughly intended to correspond to brain regions subserving the sensory processing of speech. This may include parts of the left posterior temporal cortex classically known as Wernicke's area (Brodmann's Areas [BAs] 22/21), as well as parts of the left inferior frontal and posterior basal temporal cortices (BA 44–46, 37/19; e.g., Price & Friston, 1997). Similarly, the

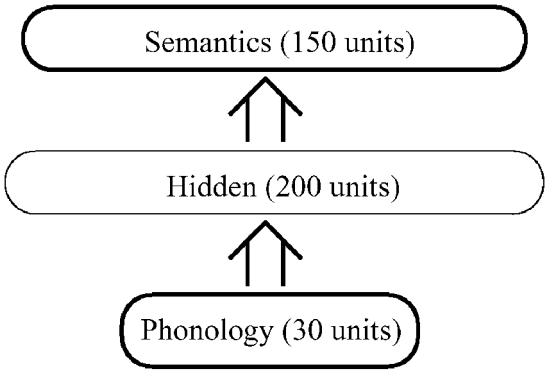

Figure 1. Network architecture.

semantic units are intended to correspond to semantic brain regions, such as the inferior temporal and the posterior parietal cortices (BA 20, 28/38, 39, and 47; e.g., Demonet et al., 1992; Pugh et al., 1996; see Price, 1998, for a review). Hidden units are intended to correspond to brain regions that are intermediate between phonological and semantic regions, although it is less clear at this point which set of regions might play this role. The model is restricted to feedforward connectivity for the sake of simplicity and to reduce training time. A fuller instantiation of the theory would involve feedback connections to allow semantic processing to influence phonological processing (e.g., Plaut, McClelland, Seidenberg, & Patterson, 1996). We believe that these interactions are important in other contexts, although results in this context do not hinge on their absence and preliminary results in smaller models that include feedback connections are comparable to those presented here.

Synaptic Depression

We adapted the spiking-based implementation of synaptic depression, developed by Varela et al. (1999), for use in a connectionist model. This approach had the virtue that the parameter values controlling the magnitude and time course of depression were derived by fitting actual neurophysiological data, rather than by fiat for our purposes. It also forced us to relate aspects of connectionist units more explicitly to properties of real neurons.

Synaptic depression was implemented as a dynamic scaling factor, ranging between 0 and 1, on the output of each unit in the network. Dynamics were dictated by the levels of presynaptic unit activity and neuromodulation (discussed below). With high presynaptic activity, the synaptic depression scaling factor moves toward 0 from an initial value of 1, reducing the impact of the unit on the activities of other units. The depression scaling factor then recovers exponentially back to 1 in the absence of presynaptic activity (see the Appendix for mathematical details). Unit activity here ranges between 0 and 1 and is intended to represent the proportion of maximal neural firing rate in a large population of similarly tuned cells,[4] where a value of 1 was chosen to correspond to 30 spikes per second (Hz), a value comparable to average peak rates for in vivo neural recording studies of neocortical cells (e.g., Miller, Li, & Desimone, 1991, 1993; Rainer & Miller, 2000). To promote consistency with this conceptualization, a negative bias was included on the input to sigmoid unit activity to yield low baseline activity in the absence of input (see Servan-Schreiber et al., 1990, for a similar approach). Although there are a number of differences between more standard connectionist models and biophysically based spiking neural networks (see Barkai & Hasselmo, 1994, and Crick & Asanuma, 1986, for discussions), the dynamics of synaptic depression hinge primarily on presynaptic firing rate, and our explicit inclusion of this value allowed us to produce depression dynamics in a connectionist model that were virtually identical to those observed in more detailed biophysical models (results not shown).

Neuromodulation

As was discussed above, two of the main actions of acetylcholine and norepinephrine in the neocortex are to suppress transmitter release presynaptically and to increase the sensitivity of cells postsynaptically to excitatory inputs. We implemented both of these actions, since they were expected to influence the extent to which network dynamics are dominated by synaptic depression. When transmitter release is suppressed, the amount of synaptic depression is reduced. Because the sensitivity of postsynaptic cells to excitatory inputs is enhanced at the same time, neural firing will be sustained with less attenuation. If the concentration of neuromodulators is reduced owing to damage, the network should be dominated more by synaptic depression and should exhibit refractory-like effects.

The presynaptic effect of neuromodulation is to scale down the value of presynaptic activity, simulating presynaptic suppression of transmitter release and reducing the buildup of synaptic depression. The release scaling factor, ρ, follows a decreasing sigmoid function of the level of neuromodulator, M, from a maximum value of 1 for large negative values of M down to a minimum value of .2 for large positive values of M (see the Appendix for details). The minimum value of ρ here was chosen to be consistent with the experimental results of Hasselmo and Bower (1992), which showed that transmitter release under high concentrations of acetylcholine saturated at 20%–30% of the levels measured in the absence of acetylcholine. At large positive values of M, ρ scales down the presynaptic unit activity by .2, leading to a weaker buildup of synaptic depression. However, this also has the effect of scaling down the impact of the presynaptic activity on other units in the network. Therefore, at high levels of neuromodulation, synaptic depression is less extreme, although input activity to subsequent units is reduced overall. This is depicted graphically in Figure 2 for a single synapse with time-varying (sinusoidal) presynaptic activity.

In contrast, for large negative values of M (low levels of neuromodulation), synaptic depression is more extreme, and the resulting input activity to other units is

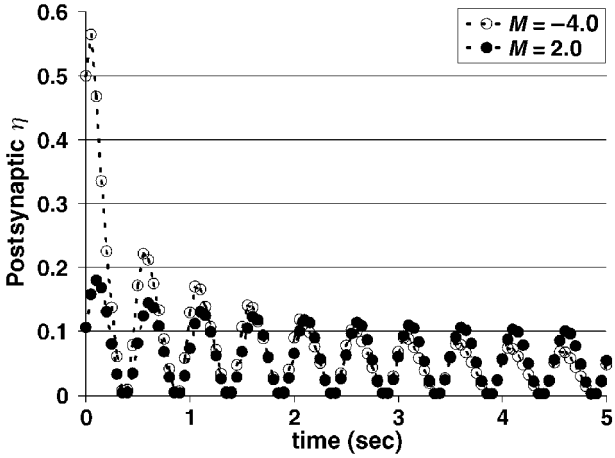

Figure 2. Presynaptic effect of neuromodulation on transmitter release and the buildup of synaptic depression. Depicted on the *y*-axis is the postsynaptic net input (η) for a single synapse (equal to presynaptic activity, sinusoidally modulated, multiplied by a weight of 1.0, and scaled by the release factor, ρ (*M*), and the depression factor, *D*; see the Appendix for details). Two different levels of neuromodulation are shown for comparison: very low (*M* = −4.0) and moderately high (*M* = 2.0).

initially much higher, because transmitter release is not suppressed. However, after repeated stimulation, the input to other units is comparable to that for high levels of neuromodulation and can even cross over to lower input values.

The postsynaptic effect of neuromodulation is to block firing-rate adaptation effects that otherwise reduce sensitivity to excitatory inputs and attenuate spiking. This means that neuromodulation leads to higher and more sustained spiking than when identical input in the absence of neuromodulation is given. We simulate these changes in sensitivity by including a multiplicative scaling factor, *g* (for *gain*), on the input to each unit. The value of *g* is an increasing sigmoid function of *M*, the level of neuromodulation; *g* ranges between a minimum, g_{min}, for large negative values of *M* and g_{max} for large positive values of *M* (see the Appendix for details). Our choices of g_{min} and g_{max} were relatively unconstrained, although we chose a ratio of g_{max}/g_{min} that was plausible given experimental results (2.0 for our simulations). For comparison, Barkai and Hasselmo (1994) observed two to three times the control spiking response following the introduction of acetylcholine. [5] To remain consistent with observations that acetylcholine and norepinephrine enhance spiking responses mainly to depolarizing/excitatory input, *g* was applied only to input that exceeded the baseline input value set by the negative bias for each unit. The enhancing effect of neuromodulation on postsynaptic sensitivity is shown in Figure 3 for different values of *M*.

Although our implementation of neuromodulation is highly simplified, it is broadly consistent with a range of empirical findings and previous approaches to simulating neuromodulatory mechanisms (e.g., Cohen & Servan-Schreiber, 1992; Gil et al., 1997; Hasselmo & Bower, 1992; Tsodyks & Markram, 1997; although see Barkai & Hasselmo, 1994).

Training Procedure

Input patterns. One hundred twenty-eight artificial "words," each composed of a pattern of on and off unit activities, were constructed to be presented to the network. For each word, 4 of the 30 phonological units were randomly selected to be active. The only constraint on the randomized pattern generation was that each pattern be unique.

Target patterns. One hundred twenty-eight target artificial semantic patterns, each corresponding to a single input word, were designed by picking 10 of the 150 semantic units to be on. These units were chosen so that patterns clustered into eight different semantic categories, 16 patterns per category. Within a category, 8 of the 16 patterns were randomly generated to be *closely* related to one another (average pairwise normalized dot product of 0.493), and 8 were generated to be *distantly* related (average pairwise normalized dot product of 0.266). The relatedness between different categories was much lower (average pairwise normalized dot product of 0.044). Each target semantic pattern was then paired at random with an input word in order to instantiate the assumption that the phonology of a word is more or less arbitrarily related to its meaning (see Plaut & Shallice, 1993b, for a discussion).

The present model was trained using an iterative version of the back-propagation learning procedure known as *back-propagation through time* (Rumelhart, Hinton, & Williams, 1986; Williams & Peng, 1990). Half of the 128 training patterns were assigned to be high frequency (presented 20 times as often during training) and half low frequency, crossing training frequency with semantic relatedness (close vs. distant). The level of neuro-

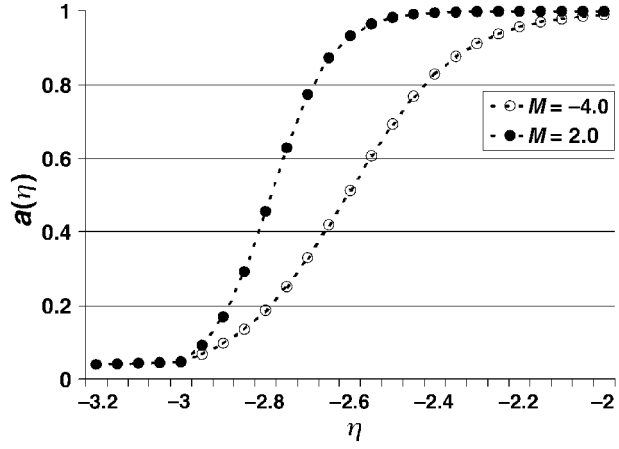

Figure 3. Effect of neuromodulation on postsynaptic activity. Shown on the *y*-axis is unit activity, $a(\eta)$ [function of net input, η, with the gain, $g(\ )$, dependent on neuromodulation, *M*; see the Appendix for details]. The increase in gain/sensitivity is apparent for moderately high levels of neuromodulation (*M* = 2.0), as compared with very low levels (*M* = −4.0).

modulation (M) during training was set to a moderate positive value (+2.0 in our simulations) so that there would be a moderate degree of synaptic depression being applied throughout learning.

Time course of a single training pattern. The time course of presentation for a single word was composed of two distinct periods. The first period served as the proxy for response–stimulus interval (RSI) in the spoken-word/picture–matching task or, more generally, as the latency between words in the speech stream. In this period (4–60 units of network time sampled uniformly, where 4 units of network time is assumed to equal one actual second), zeros were presented across the phonological input units, and no target values were presented to the semantic units. During the second period of time (2 units of network time, or 500 msec), the input word was presented with its corresponding target pattern presented at semantics. Presentation of the next training trial followed immediately, with no unit reinitialization.

Lesioning Procedure

By hypothesis, damage to connections between the hidden and the semantic layers with normal levels of neuromodulation (M) should produce a degraded-store pattern of performance. Low levels of neuromodulation should, instead, lead to an access/refractory pattern. "Damage" was carried out by randomly zeroing a proportion of synaptic strengths between hidden and semantic units (ranging from 0% to 95% in steps of 5%). A neuromodulatory deficit was simulated simply by reducing the value of M (ranging from +2.0, used during training in the intact or "normal" system, down to −4.0 in steps of 1.0; the sigmoidal nature of the pre- and postsynaptic neuromodulatory functions guaranteed asymptotic values near ±4.0). Each combination of damage-proportion/neuromodulation was repeated 80 times during testing (10 times for the testing of each of the eight semantic categories) in order to ensure a stable estimate of performance.

Testing Procedure

Following each instance of damage, the trained network was presented with four types of arrays of four words each, as in Warrington and Cipolotti (1996): close/high-frequency, close/low-frequency, distant/high-frequency, and distant/low-frequency. In each testing block, all four words in one of the arrays were probed three times in a pseudorandom order and at a fixed presentation rate (RSI of either 4 or 60 time units, representing 1 vs. 15 sec). Each array was presented at both a fast rate and a slow rate.

The pattern of semantic activity generated by an input word was compared with the target semantic patterns of all four words in the array. The best match was taken to be the network's response, unless the sum squared error (*SSE*) between actual and best-match target was larger than an arbitrary criterion value (7.0 for our simulations). In these cases, responses were selected at random, with a likelihood that was directly proportional to the amount that the criterion was exceeded (up to a likelihood of 1.0 at $SSE = 10.0$). This appropriately penalized cases in which unit activities (i.e., firing rates) were so low or divergent from known patterns that they would not be reasonably expected to support processing at points in the cognitive system subsequent to semantics. We would suggest that patients, in the absence of reliable information available from the stimulus, may similarly "guess" in alternative forced-choice paradigms.[6]

Results

The network was trained until the activity for each semantic unit at the end of stimulus presentation was on the correct side of 0.5 (either greater than 0.5 for an "on" unit or less than 0.5 for an "off" unit). The network reached this level of accuracy after approximately 500 passes through the entire training set. Following training, we sampled the full range of neuromodulatory deficits and damage to connections between hidden and semantic units. This allowed us to examine our hypotheses that an access/refractory pattern would be associated with a neuromodulatory deficit and that a degraded-store pattern would be associated with damage to connection strengths, sparing neuromodulation. It also allowed us to assess the model's ability to produce a mixing of access and degraded-store patterns. We will first demonstrate the model's ability to account for the contrasting pattern of semantic impairment observed by Warrington and Cipolotti (1996). This will be followed by a demonstration of the model's ability to account for patients who are exceptions to this basic contrasting pattern.

Evaluation of the Model's Ability to Account for Patient Performance

Consistent with our hypotheses, effects associated with access/refractory patient performance were produced under a severe neuromodulatory deficit with little or no damage to connections, whereas degraded-store patient effects were produced under normal levels of neuromodulation with severe damage to connections. Progressively surveying damage combinations from one extreme to the other and keeping overall correct performance in a comparable range (around 50% correct), we observed a gradual transition between the two types of performance patterns. Figures 4–7 establish these points for the stimulus factors of rate, semantic relatedness, frequency, and serial position; performance of individual patients from Warrington and Cipolotti (1996) are also shown for comparison. For example, Figure 4 shows the proportion correct under fast (RSI = 1 sec) and slow (RSI = 15 sec) presentation rates for each damage combination.

The model produces rate effects under severe neuromodulatory damage (e.g., $M = -4$ to -2) that are comparable in magnitude to those exhibited by the access/refractory patients A1 and A2. As neuromodulatory damage becomes less severe and damage to connections becomes more severe (left to right in the figure), the rate effects diminish. Under severe damage to connections with normal levels of neuromodulation ($M = 2$, % le-

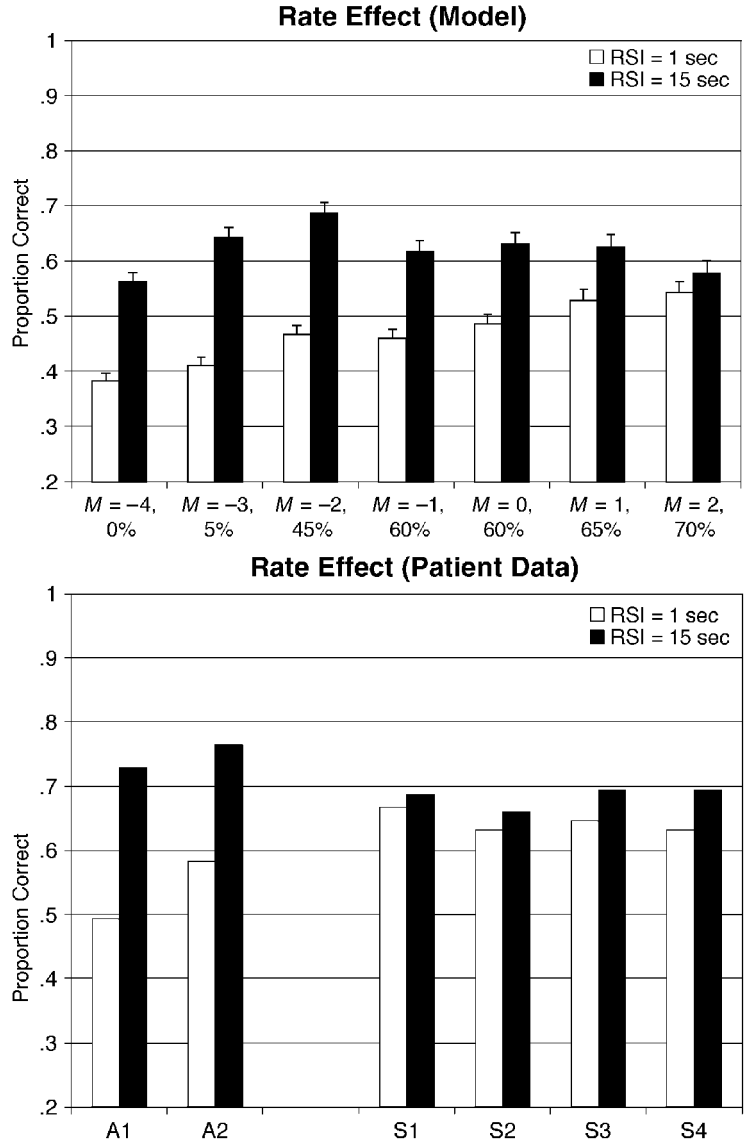

Figure 4. Effect of rate in the model under different damage combinations and the same effect in the patient data. The values of *M* and percentages of lesioned connections are listed for each damage combination.

sion = 70%), the lack of rate effect is comparable to that exhibited by the degraded-store patients, S1–S4.

The effect of semantic relatedness for different damage combinations is shown in Figure 5, along with the patient performance. Although the effects in the model are of somewhat smaller magnitudes than those exhibited by the patients, they show the appropriate pattern. Performance is better for semantically distant arrays, relative to close arrays, under severe neuromodulatory damage and little or no damage to connections. As neuromodulatory damage becomes less severe and damage to connections becomes more severe, the semantic relatedness effects gradually diminish and/or reverse.

Figure 6 shows effects of frequency for the different damage combinations, along with the patient data. Effects of frequency are small for severe neuromodulatory damage, and they grow to large magnitudes that are comparable to those of Patients S1–S4 as neuromodulatory damage becomes less severe and damage to connections becomes more severe. Interestingly, moderate damage to both connections and neuromodulation (e.g., $M = -2$ to -1) produces frequency effects that are quite substantial, while producing rate and semantic relatedness effects at the same time (discussed in more detail in the next section).

Effects of serial position under a fast rate of presentation (RSI = 1 sec) are shown in Figure 7 for each damage combination. It should be noted that the serial position effects for the individual patients in Warrington and Cipolotti (1996) are not reported in terms of proportion correct for each within-block repetition but, instead, as

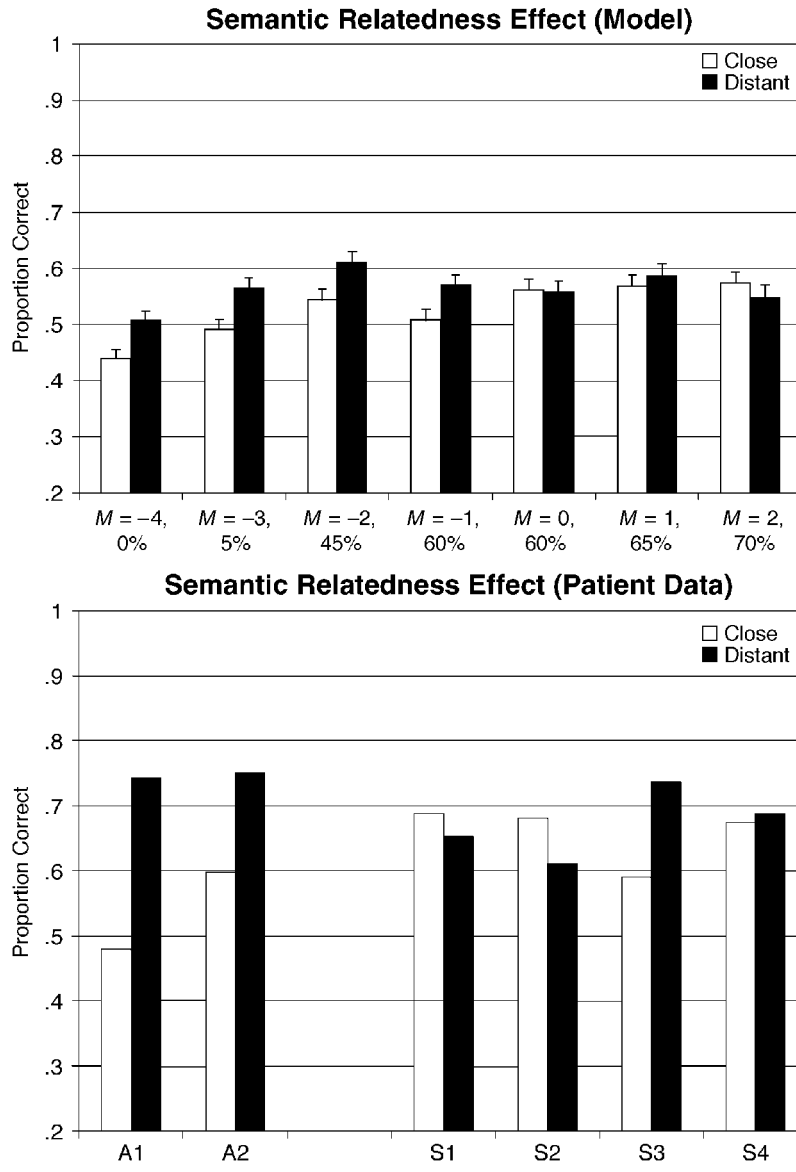

Figure 5. Effect of semantic relatedness in the model under different damage combinations and the same effect in the patient data.

the number of times a response changed from correct to incorrect (or vice versa) across the first two repetitions. More direct comparisons with the patient data will be made below in the statistical analyses of serial position. For now, it is sufficient to observe that performance decreases across within-block repetitions under severe neuromodulatory damage. As neuromodulatory damage becomes less severe and damage to connections becomes more severe, serial position effects gradually diminish.

Out of the seven different combinations of damage that are depicted in Figures 4–7, we focused our statistical analyses of model performance on two particular combinations that typify the two patient performance patterns: $M = -3$, % lesion = 5% for the access/refractory pattern and $M = 2$, % lesion = 70% for the degraded-store pattern. Repeated measures analyses of variance (ANOVAs) were run first with both damage combinations included together to evaluate the interaction of stimulus factors with damage combination. Separate ANOVAs were then conducted for each damage combination individually. Analyses were run using both damage repetition and training items as the random factor, although for simplicity we report the results of the analyses over item data, since the two methods produced comparable results and the item analyses tended to be slightly more conservative (the standard error bars shown in Figures 4–7 are based on item data). For more direct comparisons with patient results, we also calculated χ^2 statistics on the basis of the mean estimates of proportion correct over the same number of experimen-

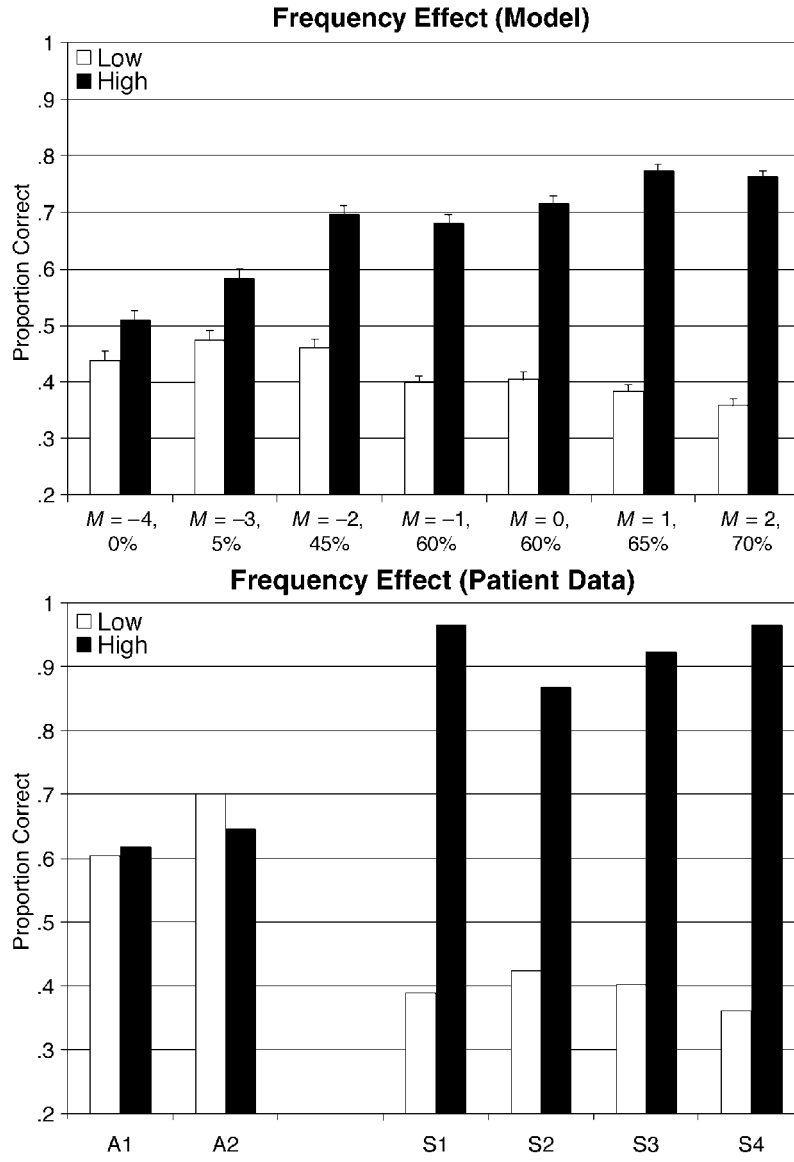

Figure 6. Effect of frequency in the model under different damage combinations and the same effect in the patient data.

tal trials as that administered to the patients in Warrington and Cipolotti's (1996) Experiment 2.[7]

A 2 (damage) × 2 (rate) × 2 (semantic relatedness) × 2 (frequency) repeated measures ANOVA was conducted on item data, with semantic relatedness and frequency as between factors and damage and rate as within factors. This was then followed by similar 2 (rate) × 2 (semantic relatedness) × 2 (frequency) ANOVAs calculated for each damage combination separately. There was a significant main effect of damage [$F(1,124) = 4.67$, $p < .04$], indicating that average performance was better for $M = 2$, % lesion = 70% than for $M = -3$, % lesion = 5% (.56 vs. .53 correct). However, this difference was not large enough to reach significance in a χ^2 analysis over the same number of trials administered to patients [$\chi^2(1, N = 288) = 0.52$, $p > .3$], suggesting that the two damage combinations are reasonably matched for overall performance. These levels of performance are slightly lower than the average levels of Warrington and Cipolotti's (1996) patients (S1–S4, .66; A1–A2, .64), although they are within the range of the most impaired patient [A1, .61 correct; $M = 2$, 70% vs. A1, $p > .2$; $M = -3$, 5% vs. A1, $p > .05$].

Rate effects. The damage × rate interaction was highly significant [$F(1,124) = 131.39$, $p < .0005$]. This indicated that the rate effect was indeed larger for severe neuromodulatory damage [$F(1,124) = 232.08$, $p < .0005$] than for severe damage to connections [$F(1,124) = 12.23$, $p < .002$]. In χ^2 analyses equivalent to those conducted on the patient data by Warrington and Cipolotti (1996),

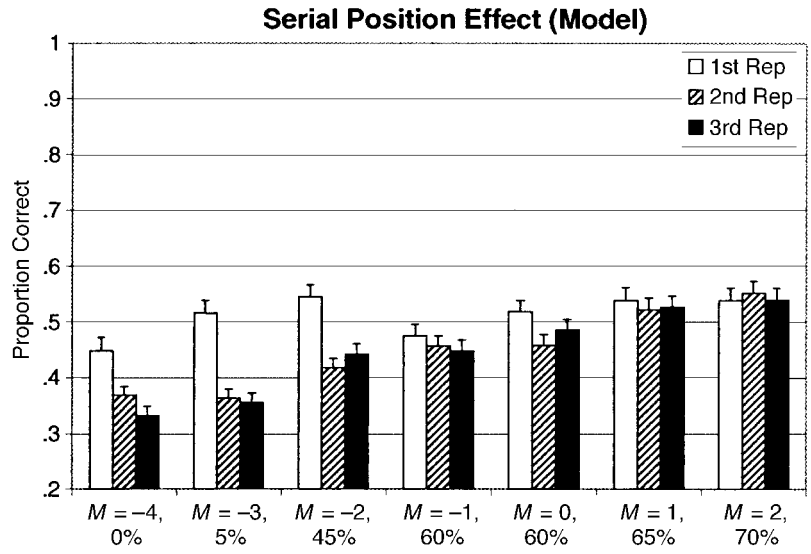

Figure 7. Effect of serial position in the model under different damage combinations (fast presentation rate conditions, response–stimulus interval = 1 sec).

the magnitude of the rate effect under severe neuromodulatory damage was large enough to reach significance [$\chi^2(1) = 14.67$, $p < .001$], whereas it failed to reach significance under severe damage to connections. These results are quite comparable to those reported for the patients.

Semantic relatedness effects. As with rate, there was a significant damage × semantic relatedness interaction [$F(1,124) = 10.40$, $p < .003$]. This showed that the semantic relatedness effect was larger under severe neuromodulatory damage [$F(1,124) = 9.321$, $p < .004$] than under severe damage to connections ($p > .2$). The magnitudes of the effects were not large enough to yield significant results in the χ^2 analyses either under severe neuromodulatory damage or under severe damage to connections. It appears, then, that although the model produces the appropriate pattern of performance, the magnitudes of the effects are somewhat smaller than those exhibited by the patients. If one considers a stronger manipulation of semantic relatedness, it is possible to observe larger effects. For example, performance on semantically unrelated or *very distant* arrays (each stimulus taken from a different semantic category) was significantly better than that on close arrays under a severe neuromodulatory deficit [$\chi^2(1) = 7.28$, $p < .01$]. Although the semantic relatedness effects under severe damage to connections also appeared to be larger, they were still too small to reach significance in the χ^2 analyses.

Frequency effects. The damage × frequency interaction was highly significant [$F(1,124) = 91.03$, $p < .0005$]. Frequency effects were larger for severe damage to connections [$F(1,124) = 412.27$, $p < .0005$] than for severe neuromodulatory damage [$F(1,124) = 20.39$, $p < .0005$]. χ^2 analyses equivalent to those conducted on the patient data showed that the frequency effect under severe damage to connections was highly significant [$\chi^2(1) = 45.99$, $p < .001$], whereas the effect under severe neuromodulatory damage failed to reach significance. These results appear to be comparable to those reported for the patients.

In addition to the basic interactions of damage combination and stimulus factors, we also observed several other, more detailed effects. There was a significant three-way damage × rate × frequency interaction [$F(1,124) = 10.67$, $p < .002$]. This reflected the lack of a rate × frequency interaction under severe neuromodulatory damage but a significant interaction under severe damage to connections (greater frequency effect under a slow rate than under a fast rate; $p < .0005$). There was also a significant semantic relatedness × frequency interaction under severe neuromodulatory damage [$F(1,124) = 4.172$, $p < .05$], indicating that the frequency effect was smaller in the close condition than in the distant condition.

Serial position effects. We conducted a 2 (damage) × 3 (within-block repetition) repeated measures ANOVA over item data in the fast presentation rate conditions, as well as separate one-way ANOVAs on within-block repetition for each damage condition. There was a highly significant damage × within-block repetition interaction [$F(2,254) = 21.95$, $p < .0005$]. The effect of within-block repetition was significant for severe neuromodulatory damage [$F(2,254) = 31.85$, $p < .0005$], but not for severe damage to connection strengths. In order to compare the model's performance more directly with that of the patients, we also conducted McNemar change tests on the first two within-block repetitions in fast rate conditions over the same number of trials as were administered to patients (see Siegel, 1956, and Warrington & Cipolotti, 1996, for further discussions). The results are presented in Table 2 for both damage combinations, along with the patient results. Shown are the number of times that a re-

Table 2
Effect of Serial Position in the Model and in Patient Data from Warrington and Cipolotti (1996)

Damage/Patient	✓ x	x ✓	$\chi^2(1)$	p
$M = -3$, 5%	25.4	13.3	3.84	.05
A1	16.0	18.0	0.03	<.90
A2	24.0	6.0	12.03	<.0001
$M = 2$, 70%	11.9	13.0	0.05	<.90
S1	9.0	7.0	0.06	<.90
S2	7.0	14.0	1.71	<.20
S3	15.0	11.0	0.34	<.70
S4	7.0	7.0	0.00	<1.00

Note— ✓ x, number of times a response was first correct and then incorrect; x ✓, number of times a response was first incorrect and then correct.

sponse was first correct and then incorrect (✓ x) versus incorrect and then correct (x ✓). There was a significant serial position effect under severe neuromodulatory damage [$\chi^2(1) = 3.84$, $p = .05$], but not under severe damage to connections. With the exception of Patient A1, who failed to show a significant serial position effect (discussed in more detail below), our results are comparable to those reported for the patients.

Consistency effects. We conducted a consistency analysis over trials in the fast rate conditions that was identical to that conducted by Warrington and Cipolotti (1996) for their patients. First, we calculated the distribution of correct/incorrect responses across all three within-block repetitions (shown in Table 3 for the two damage combinations).

We used the average level of performance for fast rate trials to generate distributions expected by chance, using the binomial expansion (assuming independence of individual trials). The observed and chance distributions were then compared with χ^2 analyses, where large χ^2 values indicate large deviations from the distributions expected by chance and more consistent responding (see Faglioni & Botti, 1993, for a critical discussion of standard response consistency measures as applied to the study of access and storage impairments). There was a significant effect of response consistency under severe damage to connections [$\chi^2(3) = 13.46$, $p < .01$], but not under severe neuromodulatory damage. As a further measure of consistency, we also calculated Pearson contingency coefficients (C), using the first two within-block stimulus repetitions over the same number of trials as in the consistency/serial position analyses reported by Warrington and Cipolotti (1996; $N = 240$; see Siegel, 1956, for a discussion of the contingency coefficient). The values of C range between 0 (low consistency) and .71 (high consistency) for 2 × 2 contingency tables (number correct/incorrect by first/second presentation). This afforded comparisons with consistency calculations that had been reported for some of the other *access* and *degraded-store* cases (e.g., Warrington & Shallice, 1979, 1984). The results from these analyses were similar to those based on the binomial expansion. Under severe damage to connections, there was significant response consistency at both fast ($C = .33$, $p < .01$) and slow ($C = .41$, $p < .001$; average $C = .37$) presentation rates. In contrast, response consistency under severe neuromodulatory damage failed to reach significance at either presentation rate (fast, $C = .02$, $p > .8$; slow, $C = .18$, $p > .1$; average $C = .10$). These values of C are somewhat lower than those observed in patient studies, although the average values of C for each damage combination do fall on different sides of the criterion value suggested by Warrington and Shallice to distinguish between *access* and *degraded-store* impairments ($C = .35$; discussed in Rapp & Caramazza, 1993).

The analyses above provide clear evidence that our model is capable of generating, at least qualitatively, the same contrasting patterns of semantic impairment as those observed by Warrington and Cipolotti (1996) over each of the five patient effects (rate, semantic relatedness, frequency, serial position, and consistency). With the exception of the semantic relatedness effect for Patients A1 and A2, the model also appeared to produce a good quantitative fit to patient performance. In order to evaluate the quantitative fit more explicitly, we correlated the model's values of proportion correct in each of the experimental conditions with those reported for the patients. Average proportion correct values for each patient *type* were calculated by combining data across individual patients (A1–A2 and S1–S4). For effects of serial position and consistency, we used the reported χ^2 values (also averaged for each patient type). Since the proportion correct and the χ^2 values were on radically different scales, we transformed all of the values to z scores to prevent any artificial inflation of the correlation measure. The correlation between model and patient data was found to be highly significant [$r(14) = .94$, $p < .001$; $R^2 = .89$], indicating that the model was capable of accounting for approximately 89% of the variance in the patient data.

Model's Ability to Account for Exceptions to the Refractory/Degraded-Store Patterns

Warrington and Cipolotti (1996) demonstrated that it is possible to observe contrasting patterns of semantic impairment in some patients for certain criteria. However, other patients, tested on these same criteria, have been shown to exhibit a mixing of access/refractory and degraded-store patterns. For example, Patient V.E.R. showed a significant frequency effect (degraded-store characteristic), along with significant rate, semantic relatedness, and serial position effects and inconsistent responding (access/refractory characteristics; Warrington & McCarthy, 1983; see Table 1). Patients C.A.V. (Warrington, 1981) and J.C.U. (Howard & Orchard-Lisle, 1984) similarly exhibited frequency effects, along with inconsistent responding. In contrast, Patients P.W. (Howard, 1985) and K.E. (Hillis et al., 1990) both exhibited the lack of a significant frequency effect (access/refractory characteristic), along with significant effects of response consistency (degraded-store characteristic). Patient A1 (Warrington & Cipolotti, 1996) showed no effect of serial position (degraded-store characteristic) but

Table 3
Effect of Response Consistency in the Model and in the Patient Data from Warrington and Cipolotti (1996)

Damage/Patient	Response	✓✓✓	x x x	✓✓x	✓x x	$\chi^2(3)$	$p<$
$M = -3$, 5%	Observed	7.8	18.7	21.9	31.7	0.70	.90
	Chance	5.6	16.3	23.9	34.2		
A1	Observed	20.0	14.0	19.0	27.0	6.60	.10
	Chance	11.0	9.0	31.0	29.0		
A2	Observed	26.0	8.0	24.0	22.0	3.82	.30
	Chance	19.0	5.0	35.0	22.0		
$M = 2$, 70%	Observed	25.4	16.4	15.9	22.3	13.46	.01
	Chance	12.8	7.6	32.3	27.2		
S1	Observed	39.0	23.0	9.0	9.0	40.8	.001
	Chance	17.0	5.0	34.0	23.0		
S2	Observed	36.0	15.0	14.0	15.0	18.9	.001
	Chance	20.0	4.0	32.0	21.0		
S3	Observed	34.0	12.0	20.0	14.0	12.8	.01
	Chance	22.0	3.0	35.0	19.0		
S4	Observed	41.0	17.0	8.0	14.0	32.5	.001
	Chance	21.0	4.0	35.0	20.0		

did show all other aspects of the access/refractory performance pattern. Do these exceptional cases constitute new forms of cognitive deficit, or can they be accounted for with the same proposal used to address the more purely contrasting patterns? We demonstrate in this section that, although our model is not capable of producing every arbitrary combination of patient effects, it is capable of producing each of these observed exception patterns.

Frequency effects without consistency. Figures 4–7 reveal that there are a number of points in the space of damage possibilities at which frequency effects overlap with other aspects of the access/refractory pattern. For example, with a moderate to severe neuromodulatory deficit and moderate damage to connections, such as $M = -2$, % lesion = 45%, there are significant effects of rate [$F(1,124) = 235.18$, $p < .0005$; $\chi^2(1) = 13.27$, $p < .001$], semantic relatedness [$F(1,124) = 11.25$, $p < .002$; although $\chi^2(1) = 1.06$, $p > .3$], frequency [$F(1,124) = 138.93$, $p < .0005$; $\chi^2(1) = 15.37$, $p < .001$], and serial position [$F(2,254) = 18.58$, $p < .0005$; McNemar, $\chi^2(1) = 2.86$, $p < .10$] and no significant effect of response consistency [$\chi^2(3) = 1.43$, $p > .5$; $C = .08$, $p > .3$]. In general, significant frequency effects can be observed, along with inconsistent responding in the model, as long as neuromodulatory damage is severe enough to lead to marked refractory behavior and damage to connections is severe enough to lead to larger frequency effects. Indeed, moderate to strong frequency effects are observed throughout much of the space of damage combinations, showing dramatic reductions only for severe neuromodulatory damage with little or no damage to connections or where performance is near floor/ceiling.

Consistency without frequency effects. Patients K.E. (Hillis et al., 1990) and P.W. (Howard, 1985) have both been documented as exhibiting significant response consistency across tasks without significant frequency effects. This pattern, when taken with those of V.E.R., C.A.V., J.C.U., and the patients studied by Warrington and Cipolotti (1996), would appear particularly problematic for the access/degraded-store distinction, because it suggests that frequency and consistency may be observed in all combinations (see Table 1). Although there are potentially important differences between K.E.'s and P.W.'s deficits and those of the patients examined by Warrington and Cipolotti (not to mention differences in methods of testing), the present model is also capable of yielding above-chance response consistency without significant frequency effects.

One of the more notable aspects of the model's performance is that, under neuromodulatory deficits with little or no damage to connections, frequency effects can still be absent at slower presentation rates that do not allow synaptic depression to accumulate much across stimuli. For example, $M = -3$, % lesion = 5% yielded a frequency effect under a slow presentation rate (RSI = 15 sec) of 69% (high frequency) versus 59% (low frequency) correct, which was not large enough to reach significance in the χ^2 analyses. This is possible because synaptic depression accumulates and recovers at two very different time scales, one on the order of a few hundred milliseconds and the other on the order of seconds to tens of seconds (see the Appendix for details). When fast presentation rates are used, both the fast and the slow components of synaptic depression have an impact on processing. When slower presentation rates are used, the slow component of synaptic depression never has a chance to build up, and only the fast component has an impact on processing. For a certain fraction of high- and low-frequency stimuli, the fast component of synaptic depression reduces activities to such low values that they no longer support reliable performance. Since synaptic depression builds up more quickly for larger presynaptic activities, high-frequency words with larger values of corresponding net input tend to be affected more than low-frequency words, counteracting the benefits of greater learning to a certain extent and leading to smaller frequency effects.

The fact that the fast component of synaptic depression recovers relatively quickly suggests that there are experimental conditions in which it may reliably affect

the same stimuli, simultaneously producing above-chance response consistency and diminished frequency effects. If presentation rates are slower or if words repeated in close succession are semantically unrelated, the slow component of synaptic depression will have less of an opportunity to build up. This will leave only the impact of the fast component of synaptic depression, yielding the same effect on a stimulus each time it is presented. Indeed, these intuitions appear to be borne out in the model under a severe neuromodulatory deficit with little or no damage to connections ($M = -3$, 5%). When arrays composed of semantically unrelated words were used, responses exhibited above-chance consistency for RSIs longer than 1 sec [e.g., RSI = 5 sec, $\chi^2(3) = 9.74$, $p < .05$; $C = 0.17$, $p < .01$], yet at the same time, there was no significant effect of frequency [RSI = 5 sec, $\chi^2(1) = 2.33$, $p > .1$]. The choice of unrelated arrays is actually more comparable to the experimental conditions in which most neurological patients are tested, since a range of stimuli with different meanings is often used (particularly in assessments of reading and writing abilities). We would suggest that although somewhat different testing methods were employed in the studies of the "exception" patients, K.E. and P.W., the present simulation may provide some insight into how their pattern of performance is possible. It also raises the question as to whether K.E. and P.W. might have exhibited more inconsistent responding under slightly different experimental conditions (e.g., blocking semantically related stimuli at fast rates).

Access/refractory pattern without serial position effect. Patient A1 (Warrington & Cipolotti, 1996) showed no effect of serial position (✓ x, 16; x ✓, 18) but showed significant effects of rate and semantic relatedness without effects of frequency or consistency. Is the lack of a serial position effect inconsistent with an account that relies on refractory processes? A closer examination of the serial position effect in our model would suggest that the answer is no. Under a severe neuromodulatory deficit with no damage to connections, $M = -4$, % lesion = 0% (shown in Figures 4–7), the serial position effect across the first two within-block repetitions was relatively weak and failed to reach significance when analyzed with the McNemar change test (✓ x = 22.1, x ✓ = 15.7: $p > .2$). However, there was also a strong effect of rate [$\chi^2(1) = 8.67$, $p < .01$] and a lack of consistent performance ($p > .95$), along with weak effects of frequency ($p > .2$) and semantic relatedness ($p > .2$). A more extreme example can be seen for $M = -1$, % lesion = 60% (also shown in Figures 4–7), where performance was inconsistent ($p > .3$) and strongly affected by presentation rate [$\chi^2(1) = 6.59$, $p < .02$], but for which there was only a slight hint of a serial position effect [✓ x = 17.1, x ✓ = 15.6: $p > .8$]. These reduced serial position effects are possible mainly because refractory effects owing to synaptic depression build up gradually across individual semantically related words, rather than at the courser scale of within-block repetition. As unit activities decrease substantially under severe damage to connections or neuromodulation, synaptic depression builds up less, and refractory effects attenuate to a certain extent (as in Figure 2). If near-asymptotic poor performance is reached in the fast rate conditions after the first 1–2 stimuli within a block, the average for the first within-block repetition (stimuli 1–4) is very close to that of the second within-block repetition, producing reduced serial position effects. However, if activity values are larger (as under more moderate damage), the refractory effects owing to synaptic depression are also larger. Under these circumstances, it is possible to observe larger serial position effects, asymptoting by the third within-block repetition or even showing a slight recovery from the second to the third repetition. Examples of both of these patterns can be observed in Figure 7 (e.g., $M = -3$, 5%; $M = -2$, 45%). In general, recovery from the second to the third repetition in the model appears to be associated with particularly strong rate and serial position effects (results not shown). On this point, it is interesting to note that a similar rebound effect has been apparent in the performance of certain access/refractory patients, such as M.E.D., studied by McNeil et al. (1994, Experiment 3), and J.M., studied by Forde and Humphreys (1995, Experiment 13; 1997, Experiment 2). The results of our simulations suggest that this may be a real effect in the patients, although it would likely require a large number of experimental trials to yield significant results.

Although we have shown above that our model is capable of accounting for the basic contrasting patterns of semantic impairment observed by Warrington and Cipolotti (1996), as well as the major patterns of exception, it does not have so many degrees of freedom that it can produce any arbitrary pattern of performance by simply choosing the appropriate combination of the two damage types. Figure 8 shows concisely the directions and magnitudes of the rate, semantic relatedness, frequency, and serial position effects throughout the entire space of damage combinations (neuromodulation levels of $M = +2$ down through $M = -4$; damage to connections ranging from 0% to 95%).

Figure 8 reveals that the magnitudes of the rate, semantic relatedness, and serial position effects are maximal under moderate to severe neuromodulatory deficits with little or no damage to connections. These effects contrast with the frequency effect, which shows maximal values under normal levels of neuromodulation with severe damage to connections. Although there are clearly damage combinations that lead to a mixture of the different effects, note that it is not possible to observe large reversals of the rate, frequency, or serial position effects under any damage combination. Indeed, the most striking aspect of these plots is that effects that increase with one damage type tend to decrease with the other, limiting the extent to which mixed effects of arbitrary magnitudes can be observed.

GENERAL DISCUSSION

We have proposed that two different types of neurological damage underlie the contrasting patterns of se-

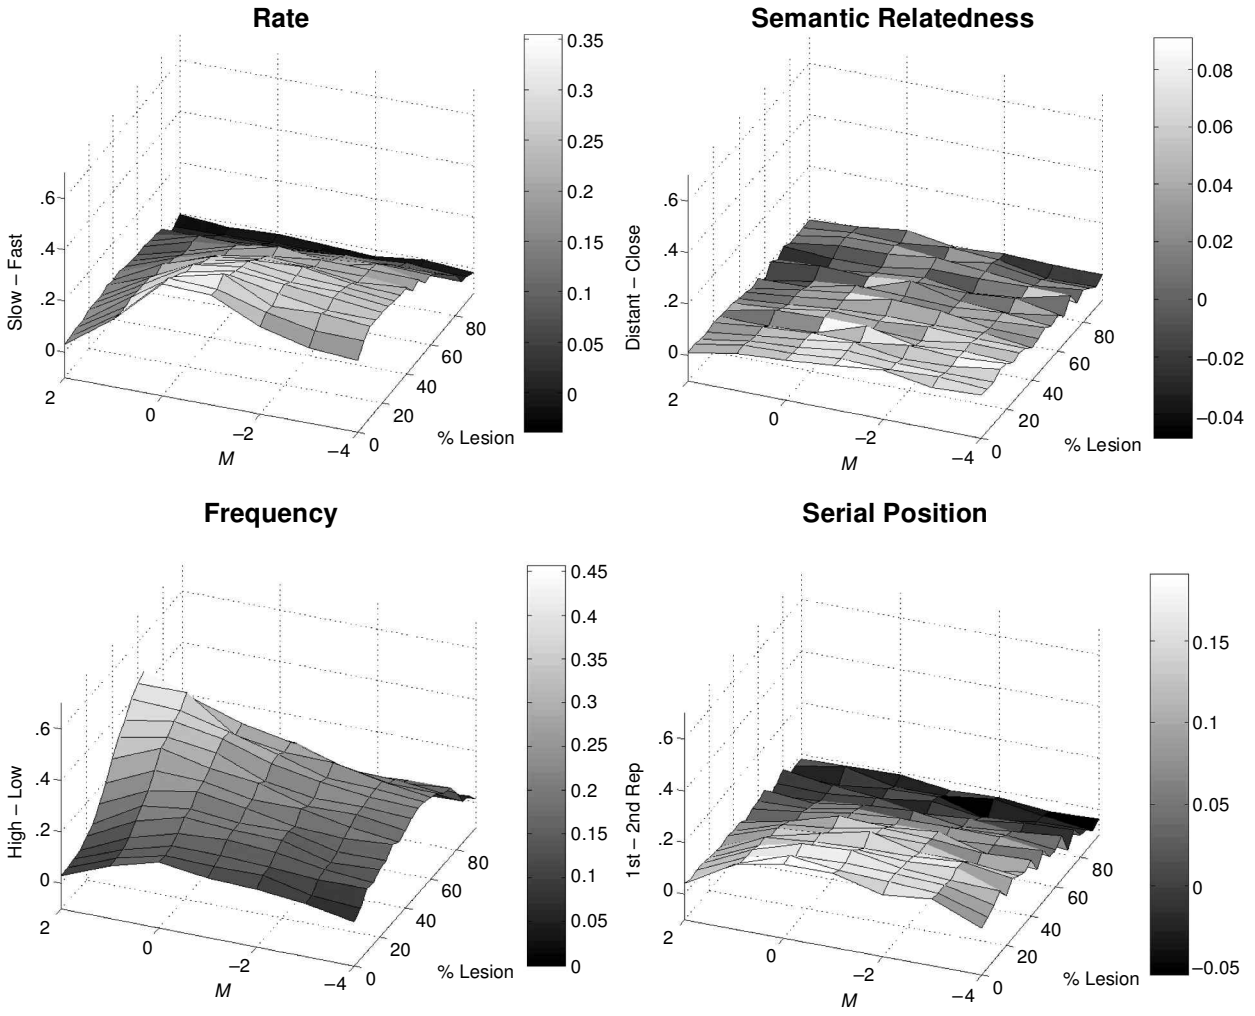

Figure 8. Surface plots showing the directions and magnitudes of the rate, semantic relatedness, frequency, and serial position effects for the entire space of damage combinations. For each effect, the *x*- and *y*-axes represent the two damage types (neuromodulation, *M*, and damage to connections, *% lesion*), and the *z*-axis represents the magnitude of the effect in terms of proportion correct (e.g., the slow minus the fast condition for the rate effect, distant − close for the semantic relatedness effect, etc.). The magnitudes of the effects are also represented in gray scale, where white corresponds to the maximum effect value and black corresponds to the minimum value; color bars to the right of each plot indicate scale.

mantic impairment exhibited by access/refractory and degraded-store patients. One type involves damage to neuromodulatory systems that normally enhance and sustain neural signals by simultaneously increasing postsynaptic sensitivity to excitatory inputs and reducing presynaptic transmitter release, which then indirectly diminishes the automatic refractory processes of synaptic depression that depend on transmitter release. The other type involves damage to the neural cells and connections that participate in coding semantic information. In a connectionist model trained to map spoken-word input to semantic representations, we have shown that damage to neuromodulation that spares connections is capable of reproducing the access/refractory pattern of impairment, whereas damage to connections that spares neuromodulation is capable of reproducing the degraded-store pattern. Under neuromodulatory damage sparing connections, there was a strong effect of presentation rate (fast < slow), an effect of semantic relatedness (close < distant), a much diminished effect of frequency (high > low), a significant serial position effect (first > second presentation), and a lack of significant response consistency. In contrast, damage to connections sparing neuromodulation yielded little or no effect of presentation rate, semantic relatedness, or serial position but yielded strong effects of frequency (high > low) and response consistency.

In the model, it is possible to understand the mechanistic basis of these contrasting patterns of semantic impairment in some detail. Perhaps the more straightforward of the two patterns to understand is the degraded-store pattern. The two most striking aspects of this pattern are the markedly impaired identification of low-frequency relative to high-frequency words and the high degree of re-

sponse consistency across multiple repetitions of the same words. The frequency effects in our model occur for the same reasons that they occur in other connectionist models of lexical processing (e.g., Plaut et al., 1996; Seidenberg & McClelland, 1989). At the beginning of training, the weight values are initialized to small random values. As training progresses, high-frequency words are learned more quickly, because the network has more chances to reduce the error present across the semantic units, and consequently, the corresponding weight values and net inputs grow to larger values more rapidly. When connections are lesioned (set to zero) following training, net input values decrease back toward zero, but low-frequency words are disproportionately affected because the weights and net inputs start out at smaller values. Performance is consistent across multiple repetitions of the same words, because refractory effects owing to synaptic depression are not large enough under normal levels of neuromodulation to lead to a large number of errors. Instead, performance is relatively deterministic and is influenced mainly by which connections happen to be spared following lesioning. The reduced impact of synaptic depression also explains the lack of rate and serial position effects and, to a certain extent, the reduced semantic relatedness effects. The semantic relatedness of words within a stimulus array was not unimportant, as was evidenced by elevated performance on arrays with semantically unrelated words (this effect was also shown for Patients S1–S4 in Warrington & Cipolotti, 1996). It is possible that the semantic relatedness manipulation (close vs. distant) was not strong enough during training to lead to large effects, a possibility that was supported by the lack of large semantic relatedness effects under any combination of damage types.

The mechanistic basis of the access/refractory pattern of performance is quite different from that of the degraded-store pattern. Refractory effects owing to synaptic depression are much larger under severe neuromodulatory deficits. Transmitter release is no longer strongly suppressed, allowing synaptic depression to build up across semantically related stimuli that share some of the same active units. The sensitivity or gain of postsynaptic units to their excitatory inputs is also simultaneously decreased. The combined effects of heightened presynaptic depression and reduced postsynaptic sensitivity lead to lower values of postsynaptic activity and a larger number of errors. At a fast presentation rate, the effects of synaptic depression are particularly marked, because both the fast and the slow components of synaptic depression can build up across semantically related stimuli (discussed in the Results section above and in the Appendix). Significant effects of presentation rate are observed because the slow component of synaptic depression builds up much less at a slow rate (RSI = 15 sec), giving rise to larger values of net input and fewer errors. Significant serial position effects are observed at a fast presentation rate, because the slow component of synaptic depression tends to build up gradually across semantically related stimuli. This means that the performance on the first within-block repetition is better, on average, than that on the second repetition. Large serial position effects contribute to the lack of significant response consistency, but random factors in testing also contribute. For example, on the first within-block repetition, an error may occur because a stimulus followed another, highly related stimulus and synaptic depression had a large impact on processing. On the second within-block repetition, the same stimulus might be identified correctly because it followed a less related stimulus, allowing synaptic depression to recover in between. The stochastic nature of responding when error on a stimulus was too large contributed to the lack of response consistency, to a certain degree, as well (see the Testing Procedure section, above).

Synaptic depression also underlies the larger semantic relatedness effects and reduced frequency effects that are part of the access/refractory pattern. During training, semantically similar stimuli develop similar internal representations that share many of the same units because of the similarities present in the semantic error signals (see Hinton & Shallice, 1991, Plaut, 1991, and Plaut & Shallice, 1993a, 1993b, for further discussions). Synaptic depression can build up to a greater extent and give rise to more errors when more units are shared, as tends to be the case for words that are closely related, as compared with those that are distantly related or unrelated. Although there was little evidence of an interaction of semantic relatedness with presentation rate under severe neuromodulatory damage, this appeared to be due to floor effects in the fast presentation rate conditions. Greater semantic relatedness effects were observed under a fast presentation rate with weaker neuromodulatory damage (e.g., $M = -1$, 15%). Synaptic depression is also responsible for the lack of strong frequency effects. Synaptic depression builds up more quickly for larger presynaptic activities, causing high-frequency words with larger values of corresponding net input to suffer more than low-frequency words. This property of synaptic depression counteracts the benefits of greater practice and learning to a certain degree, and diminished frequency effects result. As was discussed above, this same property explains the lack of frequency effects even at slow presentation rates, because the fast component of synaptic depression builds up rapidly during the processing of individual stimuli.

In addition to accounting for the basic access/refractory and degraded-store patterns of semantic impairment observed by Warrington and Cipolotti (1996), the model is capable of explaining several documented exceptions to the basic patterns. Patients V.E.R. (Warrington & McCarthy, 1983), C.A.V. (Warrington, 1981), and J.C.U. (Howard & Orchard-Lisle, 1984) all exhibited significant lexical frequency effects along with inconsistent responding. This pattern arises in the model under a neuromodulatory deficit with some degree of damage to connections. Patients P.W. (Howard, 1985) and K.E. (Hillis et al., 1990) both exhibited consistency in their responses,

without significant frequency effects. The model yields this pattern of behavior under severe neuromodulatory damage that spares connections, as long as some time is allowed to pass between highly similar stimuli. When stimuli presented in close succession are unrelated and activate few of the same units, the slow component of synaptic depression does not build up across stimuli, and only the effect of the fast component on each stimulus individually remains. Although Patients P.W. and K.E. were tested under somewhat different circumstances than were the patients in Warrington and Cipolotti (1996; e.g., consistency was assessed across tasks separated by long delays), our present demonstration provides a plausible hypothesis as to how this pattern occurred in the patients. Finally, Patient A1 (Warrington & Cipolotti, 1996) showed little effect of serial position, while exhibiting all of the other aspects of an access/refractory impairment. The model produces this pattern of impairment for a severe neuromodulatory deficit with no damage to connections. The slow component of synaptic depression in the model builds up not only across actual repetitions, but also across semantically related stimuli that activate many of the same units. Severe damage reduces activity, leading to weaker synaptic depression and weaker differences across repetitions. In contrast, strong serial position effects are observed under more moderate damage, because activities are higher and synaptic depression builds up much more across related stimuli. The model in these circumstances is capable of exhibiting a slight rebound in performance from the second to the third within-block repetitions, as has been observed in some access/refractory patients (e.g., M.E.D., McNeil et al., 1994; J.M., Forde & Humphreys, 1995, 1997).

Additional Aspects of Access/Refractory and Degraded-Store Patient Performance

Although the present model is quite successful at reproducing the empirical effects observed by Warrington and Cipolotti (1996), there are other findings that have been integral to the access/degraded-store distinction in the past that we have not addressed here. Notably, we have not attempted to simulate effects of priming/cuing or the preservation of superordinate over subordinate semantic knowledge exhibited by some patients. These effects were originally proposed as criteria by Warrington and Shallice (1979) to distinguish between deficits of access and storage. Part of the reason that we have not considered such effects is that there is much less direct evidence that they contrast for different patients tested on the same task and with the same items. This is compounded by the fact that there are good reasons to expect significant priming/cuing effects and the preservation of superordinate semantic knowledge in both access/refractory and degraded-store patients. As long as semantic knowledge is only partially degraded under damage, priming and cuing effects are expected to exist to some extent. Indeed, there are several demonstrations of preserved priming or cuing in patients that otherwise exhibit characteristics of a degraded-store impairment (e.g., Chertkow et al., 1989; Humphreys & Rumiati, 1998; Moss, Tyler, Hodges, & Patterson, 1995; Tyler & Moss, 1998). It is difficult to know whether effects of priming or cuing are weaker in patients exhibiting other aspects of a degraded-store pattern without a reasonably direct comparison to effects shown by patients exhibiting the access/refractory pattern on the same stimulus set. Since all the patients tested by Warrington and Cipolotti exhibited semantic relatedness effects to some extent (in which finer semantic distinctions were impaired), it also seems unclear whether one group of patients should show or not show the relative preservation of superordinate semantic information over subordinate. At least with patients suffering from progressive dementias, such as S1–S4, one might expect different semantic relatedness effects at different points in the progression of the disease, calling into question any strong claims about the presence or absence of effects (e.g., Tyler & Moss, 1998). It is interesting to note, in this regard, that Patient S3, like Patients A1 and A2, performed significantly better on semantically distant than on semantically close stimulus arrays and better on very distant than on distant arrays (Warrington & Cipolotti, 1996, Experiments 2 and 5). With respect to the model, we would expect some degree of priming/cuing effects and relative preservation of superordinate knowledge under either type of damage explored here. These issues could be examined more explicitly in the future by taking response time measures from the model (e.g., Plaut & Booth, 2000), training it to handle probe questions about particular stimulus properties (e.g., Rogers & McClelland, in press), and simulating additional modalities of sensory input (e.g., Plaut, in press).

Several of the access/refractory and degraded-store patients on record have exhibited a relatively selective impairment in identifying members of particular semantic categories (e.g., Warrington & McCarthy, 1983, 1987; Warrington & Shallice, 1984). Indeed, one of the access/refractory patients (A2) and two of the degraded-store patients (S1 and S3) examined by Warrington and Cipolotti (1996) showed significant effects of semantic category (animals, objects, and foods) in their spoken-word/picture–matching performance. Although the stimuli in each category were not closely matched on all of the various stimulus factors that have been shown to confound the evaluation of category-specific semantic impairments (e.g., Funnell & De Mornay Davies, 1996; Funnell & Sheridan, 1992; see Caramazza & Shelton, 1998, for a discussion), the fact that some patients exhibited significant effects of category and others did not when presented with exactly the same stimulus set suggests that both access/refractory and degraded-store patterns of impairment may be relatively confined to certain categories. We have not simulated the category-specific aspects of some of these patients' impairments, because we view these aspects as somewhat orthogonal to the main issues explored here. A number of the documented cases have not exhibited significant effects of category yet have shown one of the two basic patterns of impairment. For example, Patients H.E.C. (Cipolotti & Warrington,

1995) and A1 (Warrington & Cipolotti, 1996) showed no significant effects of category while exhibiting the access/refractory performance pattern. Similarly, Patients S2 and S4 (Warrington & Cipolotti, 1996) showed no category effects along with the degraded-store pattern. Others have used connectionist models to simulate category-specific disorders with some degree of success, either by assuming anatomical segregation of feature types that are strongly associated with members of particular semantic categories (e.g., Farah & McClelland, 1991) or by further incorporating differential inter-correlations of features for different categories (e.g., Devlin, Gonnerman, Andersen, & Seidenberg, 1998). We would suggest that as long as different semantic categories (or feature types that are associated with particular categories) are relatively segregated anatomically, it would be possible to observe the access/refractory or degraded-store patterns along with some degree of category specificity. The lack of category specificity would then be associated with broader, more diffuse damage. Further work is needed, however, to establish such points more clearly.

We chose to simulate only the auditory/verbal components of the spoken-word/picture–matching task on the basis of observations that most access/refractory patients are globally aphasic and have been relatively spared on picture/picture matching (Cipolotti & Warrington, 1995; McNeil et al., 1994; Warrington & Cipolotti, 1996; Warrington & McCarthy, 1983, 1987). The degraded-store patients S1–S4 tested by Warrington and Cipolotti were also relatively spared at picture/picture matching. Since the picture/picture–matching task requires patients to match visually dissimilar examples of objects with the same name, it has been argued that the task requires intact high-level visual structural descriptions of objects and, perhaps, intact semantic representations to a certain degree. However, Forde and Humphreys (1995, 1997) have presented evidence that for at least 1 patient (J.M.), refractory effects can occur using only picture stimuli and can transfer bidirectionally across pictures and written words. Another study with Patient J.M. has shown that refractory phenomena can also transfer across languages, from French to English and from English to French (Ferrand & Humphreys, 1996). These findings indicate that refractory effects are not limited to verbal material and may exist at a more central semantic locus. Although we have not simulated them here, such findings are quite compatible with our theoretical account of the access/refractory pattern. Cholinergic and noradrenergic neuromodulatory systems project broadly throughout the cortex, and neuromodulatory deficits are expected to impact a variety of processing loci within the brain. Neuromodulatory deficits can be relatively restricted to particular cortical regions to the extent that the projections are segregated spatially and can be damaged selectively by stroke or other disease processes. Selden et al. (1998) have shown recently that cholinergic pathways from the basal forebrain are indeed segregated into separate medial and lateral projections and that the lateral projections are further segregated into two separate divisions. Each pathway of projections innervates a different set of cortical regions and courses through regions of the white matter that are frequently involved in cerebrovascular and demyelinating diseases. It is interesting to note that both Patients A1 and A2 had extensive white matter damage (owing to stroke in A1 and a multifocal tumor in A2), and previous access/refractory patients suffered strokes of the left middle cerebral artery that affected large regions of white, as well as gray, matter. However, the inferotemporal cortex is one cortical region that also receives blood from the posterior cerebral artery, allowing a middle cerebral artery stroke to spare the cells most directly involved in semantic processing while simultaneously deinnervating the neuromodulatory inputs to those cells that ascend through the white matter. This may help to explain the replicable nature of the access/refractory patterns across a number of different patients. One might also expect such cases to be rare, with damage frequently affecting both neuromodulatory projections and cortical cells together. Nevertheless, it is plausible that neuromodulatory deficits contribute to many different cognitive impairments, some with a semantic locus and some with other loci. Consistent with this possibility, McCarthy and Kartsounis (2000) have recently demonstrated marked refractory behavior in a patient (F.A.S.) with naming difficulties that appeared to result from a postsemantic deficit. Crutch and Warrington (2001), in their study of Patient V.Y.G., have also observed refractory behavior in reading words that was not apparent in object naming, simple oral repetition, or spontaneous conversation.

Theoretical Considerations

We endorse a general view of semantic processing in which the semantic system mediates between different sensory and motor modalities in the service of performing a wide variety of cognitive behaviors and tasks. Our abilities to understand spoken and written language, to converse about ongoing events in the visual environment, and to utilize and interact with objects appropriately all depend to a large degree on our long-term semantic knowledge about the structure of the world around us. We suggest that gradual neural learning mechanisms in these varied behavioral contexts ultimately yield distributed semantic representations that reflect the shared similarity structure of the individual modalities, since these similarities aid learning and naturally promote generalization of semantic knowledge to new objects and events (Plaut, in press). For convenience, we have not simulated all of the various behaviors that a mature semantic system must surely support but have, instead, trained our model with semantic representations that contain the crucial similarities that we believe emerge in learning and that drive semantic effects in normal subjects and brain-damaged patients (see also Hinton & Shallice, 1991; Plaut & Booth, 2000; Plaut & Shallice, 1993a, 1993b). Recent attempts to simulate aspects of concept

acquisition have provided convergent support for the basic soundness of these simplifications (Plaut, in press; Rogers & McClelland, in press).

The present model also supports a view of perceptual and semantic processing in which the dynamics of processing can be altered in different behavioral circumstances. Although we have dealt, in the present simulations, only with changes in tonic levels of neuromodulation owing to damage, under more normal circumstances slight decreases in neuromodulation may be beneficial. A moderate degree of synaptic depression can aid in the detection of novel stimuli by automatically downplaying stimulus representations that have been activated recently. Higher levels of neuromodulation and weaker synaptic depression, on the other hand, can yield heightened, sustained levels of activity that permit extended processing of behaviorally relevant stimuli. Dynamic changes in neuromodulation that are contingent on task context and behavioral performance could promote efficient perceptual and semantic processing, allowing lower, transient activity to suffice when possible. Such ideas are consistent with evidence that the release of acetylcholine and norepinephrine is enhanced during the presentation of novel or behaviorally relevant stimuli, as well as following poor behavioral performance (Acquas et al., 1996; Aston-Jones & Bloom, 1981; Butt et al., 1997; Coull, 1994, 1998; Miranda et al., 2000; Ressler & Nemeroff, 1999). In addition to dynamic changes in neuromodulatory levels, top-down neural activity from the prefrontal cortex may help to sustain perceptual and semantic processing in more posterior cortical regions that would otherwise attenuate owing to synaptic depression (see Grossberg, 1976, 1980, for similar ideas).

CONCLUSIONS

We have presented a computational theory of semantic processing that incorporates the biological principles of synaptic depression and neuromodulation, and we have shown that a distinction between damage to neuromodulatory systems and damage to connections that encode semantic knowledge provides a good qualitative and quantitative explanation of the contrasting patterns of semantic impairment exhibited by access/refractory and degraded-store patients. We were also able to demonstrate that several of the main exceptions to the two basic performance patterns lie within the scope of behaviors that the model produces under damage. Importantly, the performance of the model is constrained, in the sense that effects that increase with one damage type tend to decrease with the other. Moreover, the fact that the magnitude and time course of synaptic depression and neuromodulation have been directly constrained by neuroscience observations allows the model to make empirical predictions about the dynamics of neural activity, as well as about behavioral performance under repeated stimulus processing. The results of this work are quite promising and demonstrate well the potential utility of taking constraints from biology in order to understand cognitive-level processes.

However, it is important to point out that much work is left to be done. The back-propagation learning procedure used in the present model is not very biologically plausible, and it does not incorporate some of the known influences of neuromodulation on learning, influences that play important roles in other models of cortical function (e.g., Hasselmo, 1994). We view this as a temporary condition until more biologically based rules can be developed that are capable of learning the necessary classes of mappings (e.g., O'Reilly, 1996). There is also more work to be done in understanding the relationship between connectionist units and neural activity. The equations that we have used for unit activation are quite similar to equations that have been derived for population firing rate by averaging the activity over large numbers of spiking neurons with similar tuning properties and random firing (e.g., Gerstner, 1995, 1998; Wilson & Cowan, 1972). The critical assumptions that afford the simplified firing-rate description may not always hold, such as when neurons fire synchronously (i.e., nonrandomly; see Amit & Tsodyks, 1991, Gerstner, 1995, and Laing & Chow, 2001, for discussions). Finally, further work is needed to understand the roles of synaptic depression and neuromodulation in cognitive processing more generally. Some recent work has demonstrated that synaptic depression can facilitate, as well as impair, stimulus processing under certain conditions, and it may play an important role in short-term behavioral priming effects (e.g., Gotts & Chow, 2001; see also Desimone, 1996; Wiggs & Martin, 1998). How such properties might interact with neuromodulatory mechanisms or how neuromodulation might be dynamically adjusted to promote better behavioral performance is unknown. Nevertheless, the success of the present model suggests that it can serve as a useful starting point for addressing these issues.

REFERENCES

Abbott, L. F., Varela, K., Sen, K., & Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science*, **275**, 220-223.

Acquas, E., Wilson, C., & Fibiger, H. C. (1996). Conditioned and unconditioned stimuli increase frontal cortical and hippocampal acetylcholine release: Effects of novelty, habituation, and fear. *Journal of Neuroscience*, **16**, 3089-3096.

Amit, D. J., & Tsodyks, M. V. (1991). Quantitative study of attractor neural network retrieving at low spike rates: I. Substrate—spikes, rates and neuronal gain. *Network*, **2**, 259-273.

Aston-Jones, G., & Bloom, F. E. (1981). Norepinephrine-containin g locus coeruleus neurons in behaving rats exhibit pronounced responses to non-noxious environmental stimuli. *Journal of Neuroscience*, **1**, 887-900.

Barkai, E., & Hasselmo, M. E. (1994). Modulation of the input/output function of rat piriform cortex pyramidal cells. *Journal of Neurophysiology*, **72**, 644-658.

Baylis, G. C., & Rolls, E. T. (1987). Responses of neurons in the inferior temporal cortex in short-term and serial recognition memory tasks. *Experimental Brain Research*, **65**, 614-622.

Breedin, S. D., Saffran, E. M., & Coslett, H. B. (1994). Reversal of the concreteness effect in a patient with semantic dementia. *Cognitive Neuropsycholog y*, **11**, 617-660.

Brocher, S., Artola, A., & Singer, W. (1992). Agonists of cholinergic and noradrenergic receptors facilitate synergistically the induction of long-term potentiation in slices of rat visual cortex. *Brain Research*, **573**, 27-36.

Buchel, C., Coull, J. T., & Friston, K. J. (1999). The predictive value of changes in effective connectivity for human learning. *Science*, **283**, 1538-1541.

Buckner, R. L., Petersen, S. E., Ojemann, J. G., Miezen, F. M., Squire, L. S., & Raichle, M. E. (1995). Functional anatomical studies of explicit and implicit memory retrieval tasks. *Journal of Neuroscience*, **15**, 12-29.

Butt, A. E., Testylier, G., & Dykes, R. W. (1997). Acetylcholine release in rat frontal and somatosensory cortex is enhanced during tactile discrimination learning. *Psychobiology*, **25**, 18-33.

Caramazza, A., Hillis, A. E., Rapp, B. C., & Romani, C. (1990). The multiple semantics hypothesis: Multiple confusions? *Cognitive Neuropsychology*, **7**, 161-189.

Caramazza, A., & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain: The animate–inanimate distinction. *Journal of Cognitive Neuroscience*, **10**, 1-34.

Chertkow, H., Bub, D., & Seidenberg, M. (1989). Priming and semantic memory loss in Alzheimer's disease. *Brain & Language*, **36**, 420-446.

Cipolotti, L., & Warrington, E. K. (1995). Toward a unitary account of access dysphasia: A single case study. *Memory*, **3**, 309-332.

Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, **99**, 45-77.

Coughlan, A. K., & Warrington, E. K. (1981). The impairment of verbal semantic memory: A single case study. *Journal of Neurology, Neurosurgery, & Psychiatry*, **44**, 1079-1083.

Coull, J. T. (1994). Pharmacological manipulations of the alpha(2)-noradrenergic system: Effects on cognition. *Drugs & Aging*, **5**, 116-126.

Coull, J. T. (1998). Neural correlates of attention and arousal: Insights from electrophysiology, functional neuroimaging and psychopharmacology. *Progress in Neurobiology*, **55**, 343-361.

Crick, F. H. C., & Asanuma, C. (1986). Certain aspects of the anatomy and physiology of the cerebral cortex. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models* (pp. 333-371). Cambridge, MA: MIT Press.

Crutch, S. J., & Warrington, E. K. (2001). Refractory dyslexia: Evidence of multiple task-specific phonological output stores. *Brain*, **124**, 1533-1543.

Demonet, J. F., Chollet, F., Ramsay, S., Cardebat, D., Nespoulous, J. L., Wise, R., Rascol, A., & Frackowiak, R. (1992). The anatomy of phonological and semantic processing in normal subjects. *Brain*, **115**, 1753-1768.

Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences*, **93**, 13494-13499.

Devlin, J. T., Gonnerman, L. M., Andersen, E. S., & Seidenberg, M. S. (1998). Category-specific semantic deficits in focal and widespread brain damage: A computational account. *Journal of Cognitive Neuroscience*, **10**, 77-94.

Dodt, H. U., Pawelzik, H., & Zieglgansberger, W. (1991). Actions of noradrenaline on neocortical neurons in vitro. *Brain Research*, **545**, 307-311.

Done, D. J., & Gale, T. M. (1997). Attribute verification in dementia of the Alzheimer type: Evidence for the preservation of distributed concept knowledge. *Cognitive Neuropsychology*, **14**, 547-571.

Faglioni, P., & Botti, C. (1993). How to differentiate retrieval from storage deficit: A stochastic approach to semantic memory modeling. *Cortex*, **29**, 501-518.

Farah, M. J., & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Journal of Experimental Psychology: General*, **120**, 339-357.

Ferrand, L., & Humphreys, G. W. (1996). Transfer of refractory states across languages in a global aphasic patient. *Cognitive Neuropsychology*, **13**, 1163-1191.

Finlayson, P. G., & Cynader, M. S. (1995). Synaptic depression in visual cortex tissue slices: An in vitro model for cortical neuron adaptation. *Experimental Brain Research*, **106**, 145-155.

Forde, E. M. E., & Humphreys, G. W. (1995). Refractory semantics in global aphasia: On semantic organisation and the access-storage distinction in neuropsychology. *Memory*, **3**, 265-307.

Forde, E. M. E., & Humphreys, G. W. (1997). A semantic locus for refractory behavior: Implications for access–storage distinctions and the nature of semantic memory. *Cognitive Neuropsychology*, **14**, 367-402.

Funnell, E. (1995). Objects and properties: A study of the breakdown of semantic memory. *Memory*, **3**, 497-518.

Funnell, E., & De Mornay Davies, P. (1996). JBR: A reassessment of concept familiarity and a category-specific disorder for living things. *Neurocase*, **2**, 461-474.

Funnell, E., & Sheridan, J. (1992). Categories of knowledge: Unfamiliar aspects of living and nonliving things. *Cognitive Neuropsychology*, **9**, 135-153.

Galarreta, M., & Hestrin, S. (1998). Frequency-dependent synaptic depression and the balance of excitation and inhibition in the neocortex. *Nature Neuroscience*, **1**, 587-594.

Gerstner, W. (1995). Time structure of the activity in neural network models. *Physical Review E*, **51**, 738-758.

Gerstner, W. (1998). Populations of spiking neurons. In W. Maass & C. M. Bishop (Eds.), *Pulsed neural networks* (pp. 261-295). Cambridge, MA: MIT Press.

Gil, Z., Connors, B. W., & Amitai, Y. (1997). Differential regulation of neocortical synapses by neuromodulators and activity. *Neuron*, **19**, 679-686.

Gotts, S. J., & Chow, C. C. (2001, November). *Mechanisms underlying neural synchronization and improved efficiency in networks of spiking neurons: An account of the relationship between repetition suppression and behavioral priming*. Poster presented to the Meeting of the Society for Neuroscience, San Diego.

Grossberg, S. (1976). Adaptive pattern classification and universal recoding: II. Feedback, expectation, olfaction, and illusions. *Biological Cybernetics*, **23**, 187-202.

Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review*, **87**, 1-51.

Haenny, P. E., & Schiller, P. H. (1988). State dependent activity in monkey visual cortex: I. Single cell activity in V1 and V4 on visual tasks. *Experimental Brain Research*, **69**, 225-244.

Hasselmo, M. E. (1994). Runaway synaptic modification in models of cortex: Implications for Alzheimer's disease. *Neural Networks*, **7**, 13-40.

Hasselmo, M. E. (1995). Neuromodulation and cortical function: Modeling the physiological basis of behavior. *Behavioural Brain Research*, **67**, 1-27.

Hasselmo, M. E., & Bower, J. M. (1992). Cholinergic suppression specific to intrinsic not afferent fiber synapses in rat piriform (olfactory) cortex. *Journal of Neurophysiology*, **67**, 1222-1229.

Hasselmo, M. E., Linster, C., Patil, M., Ma, D., & Cekic, M. (1997). Noradrenergic suppression of synaptic transmission may influence cortical signal-to-noise ratio. *Journal of Neurophysiology*, **77**, 3326-3339.

Hillis, A. E., Rapp, B. C., Romani, C., & Caramazza, A. (1990). Selective impairment of semantics in lexical processing. *Cognitive Neuropsychology*, **7**, 191-243.

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

Hodges, J. R., Graham, N., & Patterson, K. (1995). Charting the progression of semantic dementia: Implications for the organization of semantic memory. *Memory*, **3**, 363-395.

Hodges, J. R., Salmon, D. P., & Butters, N. (1992). Semantic memory impairment in Alzheimer's disease: Failure of access or degraded knowledge? *Neuropsychologia*, **30**, 301-314.

Hopkins, W. F., & Johnston, D. (1988). Noradrenergic enhancement of long-term potentiation at mossy fiber synapses in the hippocampus. *Journal of Neurophysiology*, **59**, 667-687.

Howard, D. (1985). *The semantic organization of the lexicon: Evidence from aphasia*. Unpublished doctoral dissertation, University of London.

HOWARD, D. (1995). Lexical anomia: Or the case of the missing lexical entries. *Quarterly Journal of Experimental Psychology*, **48A**, 999-1023.

HOWARD, D., & ORCHARD-LISLE, V. M. (1984). On the origin of semantic errors in naming: Evidence from the case of a global aphasic. *Cognitive Neuropsychology*, **1**, 163-190.

HUERTA, P. T., & LISMAN, J. E. (1993). Heightened synaptic plasticity of hippocampal CA1 neurons during a cholinergically induced rhythmic state. *Nature*, **364**, 723-725.

HUMPHREYS, G. W., & RUMIATI, R. I. (1998). When joys come not in single spies but in battalions: Within-category and within-modality identification increases the accessibility of degraded stored knowledge. *Neurocase*, **4**, 111-126.

JAHR, C. E., & NICOLL, R. A. (1982). Noradrenergic modulation of dendrodendritic inhibition in the olfactory bulb. *Nature*, **297**, 227-229.

JIANG, Y., HAXBY, J. V., MARTIN, A., UNGERLEIDER, L. G., & PARASURAMAN, R. (2000). Complementary neural mechanisms for tracking items in human working memory. *Science*, **287**, 643-646.

KRNJEVIC, K., PUMAIN, R., & RENAUD, L. (1971). Mechanism of excitation by acetylcholine in cerebral cortex. *Journal of Physiology*, **215**, 247-268.

LAING, C. R., & CHOW, C. C. (2001). Stationary bumps in networks of spiking neurons. *Neural Computation*, **13**, 1473-1494.

LAMBON RALPH, M. A. (1998). Distributed versus localist representations: Evidence from a study of item consistency in a case of classical anomia. *Brain & Language*, **64**, 339-360.

LAMBON RALPH, M. A., GRAHAM, K. S., ELLIS, A. W., & HODGES, J. R. (1998). Naming in semantic dementia: What matters? *Neuropsychologia*, **36**, 775-784.

LI, L., MILLER, E. K., & DESIMONE, R. (1993). The representation of stimulus familiarity in anterior inferior temporal cortex. *Journal of Neurophysiology*, **69**, 1918-1929.

MADISON, D. V., & NICOLL, R. A. (1984). Control of the repetitive discharge of rat CA1 pyramidal neurons in vitro. *Journal of Physiology*, **354**, 319-331.

MADISON, D. V., & NICOLL, R. A. (1986). Actions of noradrenaline recorded intracellularly in rat hippocampal CA1 pyramidal neurons, in vitro. *Journal of Physiology*, **372**, 221-244.

MARROCCO, R. T., & DAVIDSON, M. C. (1998). Neurochemistry of attention. In R. Parasuraman (Ed.), *The attentive brain* (pp. 35-50). Cambridge, MA: MIT Press.

MCCARTHY, R. A., & KARTSOUNIS, L. D. (2000). Wobbly words: Refractory anomia with preserved semantics. *Neurocase*, **6**, 487-497.

MCLEAN, J., & PALMER, L. A. (1996). Contrast adaptation and excitatory amino acid receptors in cat striate cortex. *Visual Neuroscience*, **13**, 1069-1087.

MCNEIL, J. E., CIPOLOTTI, L., & WARRINGTON, E. K. (1994). The accessibility of proper names. *Neuropsychologia*, **32**, 193-208.

MILLER, E. K., GOCHIN, P. M., & GROSS, C. G. (1991). Habituation-like decrease in the responses of neurons in inferior temporal cortex of the macaque. *Visual Neuroscience*, **7**, 357-362.

MILLER, E. K., LI, L., & DESIMONE, R. (1991). A neural mechanism for working memory and recognition memory in inferior temporal cortex. *Science*, **254**, 1377-1379.

MILLER, E. K., LI, L., & DESIMONE, R. (1993). Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *Journal of Neuroscience*, **13**, 1460-1478.

MIRANDA, M. I., RAMIREZ-LUGO, L., & BERMUDEZ-RATTONI, F. (2000). Cortical cholinergic activity is related to the novelty of the stimulus. *Brain Research*, **882**, 230-235.

MORTON, J. (1981). The status of information processing models of language. *Philosophical Transactions of the Royal Society of London: Series B*, **295**, 387-396.

MOSS, H. E., TYLER, L. K., HODGES, J. R., & PATTERSON, K. (1995). Exploring the loss of semantic memory in semantic dementia: Evidence from a primed monitoring study. *Neuropsychology*, **9**, 16-26.

MULLER, J. R., METHA, A. B., KRAUSKOPF, J., & LENNIE, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science*, **285**, 1405-1408.

O'REILLY, R. C. (1996). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, **8**, 895-938.

PARKIN, A. J. (1993). Progressive aphasia without dementia: A clinical and cognitive neuropsychological analysis. *Brain & Language*, **44**, 201-220.

PATTERSON, K. E., & MARCEL, A. J. (1977). Aphasia, dyslexia and the phonological coding of written words. *Quarterly Journal of Experimental Psychology*, **29**, 307-318.

PITLER, T. A., & ALGER, B. E. (1992). Cholinergic excitation of GABAergic interneurons in the rat hippocampal slice. *Journal of Physiology*, **450**, 127-142.

PLAUT, D. C. (1991). *Connectionist neuropsychology: The breakdown and recovery of behavior in lesioned attractor networks*. Unpublished doctoral dissertation, Carnegie Mellon University, School of Computer Science. (Also available as Tech. Rep. No. CMU-CS-91-185)

PLAUT, D. C. (in press). Graded modality-specific specialization in semantics: A computational account of optic aphasia. *Cognitive Neuropsychology*.

PLAUT, D. C., & BOOTH, J. R. (2000). Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. *Psychological Review*, **107**, 786-823.

PLAUT, D. C., MCCLELLAND, J. L., SEIDENBERG, M. S., & PATTERSON, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, **103**, 56-115.

PLAUT, D. C., & SHALLICE, T. (1993a). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, **10**, 377-500.

PLAUT, D. C., & SHALLICE, T. (1993b). Perseverative and semantic influences on visual object naming errors in optic aphasia: A connectionist account. *Journal of Cognitive Neuroscience*, **5**, 89-117.

PRICE, C. J. (1998). The functional anatomy of word comprehension and production. *Trends in Cognitive Sciences*, **2**, 281-288.

PRICE, C. J., & FRISTON, K. J. (1997). Cognitive conjunction: A new approach to brain activation experiments. *NeuroImage*, **5**, 261-270.

PUGH, K. R., SHAYWITZ, B. A., SHAYWITZ, S. E., CONSTABLE, R. T., SKUDLARSKI, P., FULBRIGHT, R. K., BRONEN, R. A., SHANKWEILER, D. P., KATZ, L., FLETCHER, J. M., & GORE, J. C. (1996). Cerebral organization of component processes in reading. *Brain*, **119**, 1221-1238.

RAINER, G., & MILLER, E. K. (2000). Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron*, **27**, 179-189.

RAPP, B. C., & CARAMAZZA, A. (1993). On the distinction between deficits of access and deficits of storage: A question of theory. *Cognitive Neuropsychology*, **10**, 113-141.

RESSLER, K. J., & NEMEROFF, C. B. (1999). Role of norepinephrine in the pathophysiology and treatment of mood disorders. *Biological Psychiatry*, **46**, 1219-1233.

RIDDOCH, M. J., HUMPHREYS, G. W., COLTHEART, M., & FUNNELL, E. (1988). Semantic systems or system? Neuropsychological evidence re-examined. *Cognitive Neuropsychology*, **5**, 3-25.

ROBBINS, T. W. (1997). Arousal systems and attentional processes. *Biological Psychology*, **45**, 57-71.

ROGERS, T. T., & MCCLELLAND, J. L. (in press). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.

RUMELHART, D. E., HINTON, G. E., & WILLIAMS, R. J. (1986). Learning representations by back-propagating errors. *Nature*, **323**, 533-536.

SAHIN, M., BOWEN, W. D., & DONOGHUE, J. P. (1992). Location of nicotinic and muscarinic cholinergic and mu-opiate receptors in rat cerebral neocortex: Evidence from thalamic and cortical lesions. *Brain Research*, **579**, 135-147.

SANCHEZ-VIVES, M. V., NOWAK, L. G., & MCCORMICK, D. A. (2000). Cellular mechanisms of long-lasting adaptation in visual cortical neurons in vitro. *Journal of Neuroscience*, **20**, 4286-4299.

SARTORI, G., & JOB, R. (1988). The oyster with 4 legs: A neuropsychological study on the interaction of visual and semantic information. *Cognitive Neuropsychology*, **5**, 105-132.

SCHACTER, D. L., ALPERT, N. M., SAVAGE, C. R., RAUCH, S. L., & AL-

BERT, M. S. (1996). Conscious recollection and the human hippocampal formation: Evidence from positron emission tomography. *Proceedings of the National Academy of Sciences*, **89**, 1837-1841.

SCHWINDT, P. C., SPAIN, W. J., & CRILL, W. E. (1992). Calcium-dependent potassium currents in neurons from cat sensorimotor cortex. *Journal of Neurophysiology*, **67**, 216-226.

SEIDENBERG, M. S., & MCCLELLAND, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, **96**, 523-568.

SELDEN, N. R., GITELMAN, D. R., SALAMON-MURAYAMA, N., PARRISH, T. B., & MESULAM, M.-M. (1998). Trajectories of cholinergic pathways within the cerebral hemispheres of the human brain. *Brain*, **121**, 2249-2257.

SENN, W., MARKRAM, H., & TSODYKS, M. (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Computation*, **13**, 35-67.

SERVAN-SCHREIBER, D., PRINTZ, H., & COHEN, J. D. (1990). A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science*, **249**, 892-895.

SHALLICE, T. (1988). *From neuropsychology to mental structure*. Cambridge: Cambridge University Press.

SIEGEL, S. (1956). *Nonparametric statistics*. New York: McGraw-Hill.

SILVERI, M. C., & GAINOTTI, G. (1988). Interaction between vision and language in category-specific semantic impairment. *Cognitive Neuropsychology*, **5**, 677-709.

THOMSON, A. M., DEUCHARS, J., & WEST, D. C. (1993). Large, deep layer pyramid–pyramid single axon EPSPS in slices of rat motor cortex display paired pulse and frequency-dependent depression, mediated presynaptically and self-facilitation, mediated postsynaptically. *Journal of Neurophysiology*, **70**, 2354-2369.

TSODYKS, M. V., & MARKRAM, H. (1997). The neural code between cortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*, **94**, 719-723.

TULVING, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory* (pp. 381-403). New York: Academic Press.

TYLER, L. K., & MOSS, H. E. (1998). Going, going, gone . . . ? Implicit and explicit tests of conceptual knowledge in a longitudinal study of semantic dementia. *Neuropsychologia*, **36**, 1313-1323.

VARELA, J. A., SEN, K., GIBSON, J., FOST, J., ABBOTT, L. F., & NELSON, S. B. (1997). A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *Journal of Neuroscience*, **17**, 7926-7940.

VARELA, J. A., SONG, S., TURRIGIANO, G. G., & NELSON, S. B. (1999). Differential depression at excitatory and inhibitory synapses in visual cortex. *Journal of Neuroscience*, **19**, 4293-4304.

VIDAL, C., & CHANGEUX, J. P. (1993). Nicotinic and muscarinic modulations of excitatory synaptic transmission in the rat prefrontal cortex in vitro. *Neuroscience*, **56**, 23-32.

WAGNER, A. D., MARIL, A., & SCHACTER, D. L. (2000). Interactions between forms of memory: When priming hinders new episodic learning. *Journal of Cognitive Neuroscience*, **12**(Suppl. 2), 52-60.

WARRINGTON, E. K. (1975). The selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology*, **27**, 635-657.

WARRINGTON, E. K. (1981). Concrete word dyslexia. *British Journal of Psychology*, **72**, 175-196.

WARRINGTON, E. K., & CIPOLOTTI, L. (1996). Word comprehension: The distinction between refractory and storage impairments. *Brain*, **119**, 611-625.

WARRINGTON, E. K., & MCCARTHY, R. A. (1983). Category specific access dysphasia. *Brain*, **106**, 859-878.

WARRINGTON, E. K., & MCCARTHY, R. A. (1987). Categories of knowledge: Further fractionations and an attempted integration. *Brain*, **110**, 1273-1296.

WARRINGTON, E. K., & SHALLICE, T. (1979). Semantic access dyslexia. *Brain*, **102**, 43-63.

WARRINGTON, E. K., & SHALLICE, T. (1984). Category specific semantic impairments. *Brain*, **107**, 829-854.

WIGGS, C. L., & MARTIN, A. (1998). Properties and mechanisms of perceptual priming. *Current Opinion in Neurobiology*, **8**, 227-233.

WILLIAMS, R. J., & PENG, J. (1990). An efficient gradient-based algorithm for on-line training of recurrent network trajectories. *Neural Computation*, **2**, 490-501.

WILSON, H. R., & COWAN, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, **12**, 1-24.

WOODY, C. D., & GRUEN, E. (1987). Acetylcholine reduces net outward currents measured in vivo with single electrode voltage clamp techniques in neurons of the motor cortex of cats. *Brain Research*, **424**, 193-198.

NOTES

1. Warrington and Shallice (1979) originally proposed that patients with an access impairment should not exhibit evidence of a *hierarchical* breakdown of semantic information (i.e., no better at verifying general superordinate features shared across category members than more subordinate, member-specific features), because the stochastic influences that dominate their poor performance would be uncorrelated with feature type. It is unclear how this proposal squares with the common observation that these patients perform worse when distractors in the matching-to-sample tasks are semantically related.

2. Another patient (H.E.C.; Cipolotti & Warrington, 1995) has shown a significant serial position effect when tested with the same items and procedures. This patient, like A2, also exhibited significant rate and semantic relatedness effects, along with the lack of a frequency effect (see Table 1).

3. Although Gil et al. (1997) observed effects comparable to those of Tsodyks and Markram (1997) at intracortical synapses, they found reduced or opposite effects at thalamocortical synapses. Thalamocortical relay cells express a kind of cholinergic receptor, nicotinic, that is largely absent at intracortical synapses that express mainly muscarinic receptors (e.g., Sahin, Bowen, & Donoghue, 1992). Since synaptic depression effects are large at cortical synapses that involve mainly muscarinic receptors, acetylcholine would be expected to reduce synaptic depression throughout the cortex.

4. The connectionist equations used here are similar in form to *firing-rate* equations that have been derived as spatially or temporally averaged descriptions of large networks of spiking neurons, where the firing times of individual neurons are noisy or random (e.g., Amit & Tsodyks, 1991; Gerstner, 1995, 1998; Laing & Chow, 2001; Wilson & Cowan, 1972).

5. Barkai and Hasselmo (1994) emphasized the importance of dynamic changes in postsynaptic excitability. For simplicity, we have ignored these dynamics in the present context, because the larger changes in excitability that they consider build up/recover nearly completely within 100 msec (somewhat slower and considerably weaker changes recover within 500–600 msec). These changes are nearly instantaneous, relative to the time scale of updates to unit activities in our simulations (every 50 msec), and the refractory phenomenon that we consider occurs on the longer time scale of seconds. There is a recent demonstration that firing-rate adaptation effects can be much slower (Sanchez-Vives, Nowak, & McCormick, 2000), although we anticipate that adding such effects will not substantially alter the behavior of the model, because they are expected to be stronger under low levels of neuromodulation and weaker under normal levels.

6. A similar approach has been taken by Plaut and Shallice (1993b) in simulating omission errors in visual object naming. The basic idea is that if a pattern of activity is too far away from any known pattern (past some criterion value), it cannot support correct identification performance.

7. As was discussed by Warrington and McCarthy (1983) and Warrington and Cipolotti (1996), the independence assumption of the χ^2 test is unlikely to be satisfied when refractory phenomena are present at time scales longer than the duration of individual trials. We present these results mainly for purposes of comparison with the values reported for patients.

APPENDIX

Synaptic Depression

As was discussed above, synaptic depression is implemented as a dynamic scaling factor on the output activity of each unit in the network. The equation used to calculate unit activity is as follows:

$$a_i = \frac{1}{1 + e^{-g(\eta_i, M)}},$$

where a_i is the activation of the ith postsynaptic unit, $g(\,)$ is a scaling function that implements the postsynaptic effects of neuromodulation (M; discussed below), and η_i is the net input to the ith unit. The model incorporates the effects of synaptic depression in the equation used to calculate the net input of units:

$$\frac{d\eta_i}{dt} = -\eta_i + \sum_j a_j \rho(M) w_{ij} D_j,$$

where a_j is the activity of the jth presynaptic unit, $\rho(M)$ is a scaling function that implements the presynaptic effect of neuromodulation (M) on transmitter release (see below), w_{ij} is the *weight* or synaptic strength from the jth to the ith unit, and D_j is the synaptic depression scaling factor ranging between 0 and 1. In the Varela et al. (1999) model, D_j has both a faster and a slower component at excitatory synapses in the primary visual cortex (referred to here as $D_{j[f]}$ and $D_{j[s]}$, respectively; $D_j = D_{j[f]} D_{j[s]}$). The equations given by Varela et al. (1999) for each of these components are as follows:

Each time a presynaptic spike occurs,

$$D_{j[k]} \rightarrow d_{[k]} D_{j[k]} \qquad (0 < d_{[k]} < 1);$$

otherwise,

$$\frac{dD_{j[k]}}{dt} = \frac{1 - D_{j[k]}}{\tau_{[k]}},$$

where $d_{[k]}$ determines the amount of depression for each pre-synaptic spike (.78 for the fast component, .97 for the slow component), and $\tau_{[k]}$ determines the rate of recovery back to a value of 1 (0.634 sec for the fast component, 9.3 sec for the slow component). We approximated these two equations with a single equation, weighting the depression and recovery terms by an estimate of the likelihood of a spike versus no spike for a given value of a_j:

$$\frac{dD_{j[k]}}{dt} = a_{\max} \cdot a_j \cdot \rho(M)(d_{[k]} - 1) D_{j[k]} + \frac{[1 - a_{\max} \cdot a_j \cdot \rho(M)](1 - D_{j[k]})}{\tau_{[k]}},$$

where $a_{\max}$ is the maximum firing rate (spikes/msec) to which a unit activity value of 1 is supposed to correspond (0.03 in our simulations, equivalent to 30 Hz). Our values for both $d_{[k]}$ and $\tau_{[k]}$ were taken directly from the Varela et al. (1999) model and, as such, are not free parameters. Note that D_j actually follows the product $a_j \cdot \rho(M)$, rather than a_j by itself, instantiating the notion that synaptic depression depends on transmitter release.

Neuromodulation

We capture the presynaptic suppression of transmitter release by neuromodulation (M) with the decreasing sigmoid function $\rho(M)$:

$$\rho(M) = \rho_{\min} + \frac{(1 - \rho_{\min}) e^{-M}}{1 - e^{-M}},$$

where $\rho_{\min}$ is a constant representing the minimum value of $\rho(M)$ (.2 for our simulations); $\rho(M)$ is 1 for large negative M and $\rho_{\min}$ for large positive M. We capture the increases in postsynaptic sensitivity to excitatory input with the function $g(\eta_i, M)$, referred to above in the equation for postsynaptic unit activity a_i:

$$g(\eta_i, M) = \eta_i \left[g_{\min} + \frac{(g_{\max} - g_{\min})}{1 + e^{-M}} \right] + \beta_i,$$

where $g_{\min}$ is the minimum value by which the net input (η_i) is scaled, $g_{\max}$ is the maximum value by which η_i is scaled, and β_i is a bias term for the ith unit (< -3.0 in our simulations to set low baseline unit activity). To remain consistent with observations that acetylcholine and norepinephrine enhance spiking responses mainly to depolarizing/excitatory input, $g(\,)$ is only applied if $\eta_i > 0$; $g(\,)$ simply returns $\eta_i + \beta_i$ for $\eta_i < 0$.

(Manuscript received May 23, 2001;
revision accepted for publication May 14, 2002.)