On the Interaction of Selective Attention and Lexical Knowledge: A Connectionist Account of Neglect Dyslexia
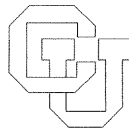
Michael Mozer Marlene Behrmann*

CU-CS-441-89 August 1989

University of Colorado at Boulder

DEPARTMENT OF COMPUTER SCIENCE

On the Interaction of
Selective Attention and Lexical Knowledge:
A Connectionist Account of Neglect Dyslexia

Michael C. Mozer
Marlene Behrmann

CU-CS-441-89   August 1989

Michael C. Mozer
*Department of Computer Science
and Institute of Cognitive Science
University of Colorado
Boulder, CO 80309-0430
USA*

e-mail: mozer@boulder.colorado.edu

Marlene Behrmann
*Department of Psychology
and Sunnybrook Medical Centre
University of Toronto
Toronto, Ontario M5S 1A1
CANADA*

e-mail: marlene@psych.toronto.edu

Abstract

Neglect dyslexia, a reading impairment acquired as a consequence of brain injury, is traditionally interpreted as a disturbance of selective attention. Patients with neglect dyslexia may ignore the left side of an open book, the beginning words of a line of text, or the beginning letters of a single word. These patients provide a rich but sometimes contradictory source of data regarding the locus of attentional selectivity. We have reconsidered the patient data within the framework of an existing connectionist model of word recognition and spatial attention. We show that the effects of damage to the model resemble the reading impairments observed in neglect dyslexia. In simulation experiments, we account for a broad spectrum of behaviours including the following: (1) when two noncontiguous stimuli are presented simultaneously, the contralesional stimulus is neglected (extinction); (2) explicit instructions to the patient can reduce the severity of neglect; (3) stimulus position in the visual field affects reading performance; (4) words are read much better than pronounceable nonwords; (5) the nature of error responses depends on the morphemic composition of the stimulus; and (6) extinction interacts with lexical knowledge (if two words are presented that form a compound, e.g., COW and BOY, the patient is more likely to report both than in a control condition, e.g., SUN and FLY). The convergence of findings from the neuropsychological research and the computational modeling sheds light on the role of attention in normal visuospatial processing, supporting a hybrid view of attentional selection that has properties of both early and late selection.

On the Interaction of Selective Attention and Lexical Knowledge:
A Connectionist Account of Neglect Dyslexia

Neglect dyslexia, a reading impairment acquired as a consequence of brain injury, is traditionally interpreted as a disturbance of selective attention. Neglect dyslexia patients may ignore the left side of an open book, the beginning words of a line of text, or the beginning letters of a single word, even when all the visual information appears in an intact region of their visual field (Bisiach, 1988; Caplan, 1987; Ellis, Flude, & Young, 1987).[1] Several explanations have been advanced in the neuropsychological literature to account for the underlying deficit. One explanation argues that neglect arises from an impairment in the level of arousal of the brain. Following a lesion, the damaged hemisphere is hypoaroused and fails to process the incoming information adequately (Kinsbourne, 1987); alternatively, the equilibrium between the hemispheres is disrupted and attention generally shifts to the ipsilesional side. A second explanation postulates a more specific deficit, defined in terms of the elementary operations of spatial attention. On this account, the mechanism responsible for selecting between competing inputs is damaged, producing systematic deficits in visuospatial functioning (Posner, 1988; Posner & Petersen, 1989).

To understand the impact of the attentional deficit on the processing of printed material, it is important to appreciate the various roles that have been postulated for attention in normal processing. The most conspicuous role of selective attention is to control the order of processing so that words on a page are read from left to right, top to bottom. Attention may also be critical in focusing on one or a small number of words at a time to allocate limited processing resources and prevent crosstalk in the system (Mozer, 1988a; Shallice, 1988). A quite different function suggested for attention is that of integrating the results of processing performed by independent subsystems (Treisman & Gelade, 1980). In reading, this function of attention may serve to tie letter or word identities to locations (Ellis et al., 1987; Mozer, 1989).

Although it is generally accepted that attention plays a central role in reading, the locus of attentional selectivity remains controversial. For nearly three decades, cognitive psychologists have debated whether attentional selection occurs *early* or *late* in the processing of visual information. Early-selection views (Broadbent, 1958; Treisman, 1969) derive their name from the assertion that selection occurs early in the sequence of processing stages, prior to stimulus identification and semantic processing. In contrast, late-selection views (e.g., Deutsch & Deutsch, 1963; Norman, 1968; Schiffrin & Schneider, 1977) posit that all stimuli are processed at least to the stage of identification before selection occurs. Early selection generally assumes that selection is based on low-level features such as location or color, whereas late-selection is based on high-level features such as stimulus identity or meaning (Pashler & Badgio, 1985). According to the early selection view, disturbance of normal attentional functioning should result in the degradation or attenuation of low-level representations. According to the late selection view, the disturbance should come into play only after the stimulus has been processed to a high level.

The findings in the literature on neglect dyslexia provide a rich but seemingly contradictory source of data regarding the locus of the attentional deficit. On the one hand, neglect dyslexia has been shown occur with respect to a retinal coordinate frame, as opposed to an intrinsic object-centered frame.[2] For example, vertically presented words are not subject to neglect; 180 degree rotation of words leads to neglect with respect to the left of the retinal frame, not the object-centered frame;[3] retinal location of a word affects performance, even in the right visual field — the further to the right a word is presented relative to fixation, the better it is reported (Behrmann, Black, Moscovitch, & Mozer, 1989; Ellis et al., 1987). These findings suggest an attentional disruption occurring at an early stage of analysis for the following reason. The initial encoding of the visual world is certainly retinotopic, and one can argue on computational grounds that object recognition requires as a precondition a recoding of the perceptual data

---

[1] Neglect dyslexia, as hemispatial neglect, occurs more frequently following lesions to the right hemisphere than to the left hemisphere; consequently, the descriptions in this paper refer to left-sided neglect.

[2] We use the term *retinal coordinate frame* loosely to describe a reference frame that depends on head and/or body position of the observer, not just eye position. See Ladavas (1987) and Farah, Brunn, Wong, Wallace, and Carpenter (1989) for further discussion of possible reference frames used in spatial attention.

[3] However, see Barbut and Gazzaniga (1987) and Hillis and Caramazza (1989) for an alternative conceptualization.

into an object-centered representation (Hinton, 1981; Marr, 1982). Thus, if attentional selection operates on a retinotopic encoding, it must operate prior to object recognition.

On the other hand, there is contradictory evidence indicating that the attentional disruption occurs at later stages of analysis. For example, neglect is less severe for words than nonwords (Brunn & Farah, 1989; Sieroff, Pollatsek, & Posner, 1988); the nature of error responses depends on the the morphemic composition of the stimulus (Behrmann et al., 1989); and extinction interacts with higher-order stimulus properties — if two words are presented that form a compound, e.g., COW and BOY, the patient is more likely to neglect the left word than in a control condition, e.g., SUN and FLY (Behrmann et al., 1989).

These paradoxical results rule out simple early and late selection views of attention. The early-selection view cannot explain why selection may depend on higher-order properties of the stimuli. The late-selection view is contrary to the finding that neglect depends on the position and orientation of the word in the visual field.

A hybrid view of attentional selection: MORSEL

We have reconsidered the phenomena of neglect dyslexia in light of a connectionist model of two-dimensional object recognition and spatial attention called *MORSEL* (Mozer, 1987, 1988a, 1988b). MORSEL was originally designed to account for a broad spectrum of psychological data, including perceptual errors that arise when several objects appear simultaneously in the visual field, facilitatory effects of context and redundant information, and attentional phenomena. MORSEL appears to be a particularly appropriate framework in which to consider neglect dyslexia because the theory of attention it embodies is a hybrid of early and late selection views. We will show that lesions of MORSEL's attentional system produce symptoms of neglect dyslexia. Before describing this work, we must summarize the components of MORSEL that are essential to the task of word identification.

The essential components of MORSEL (Figure 1) are: (1) a connectionist network that constructs location-invariant representations of visually-presented letters and words, called *BLIRNET*; (2) a network that constructs a consistent interpretation of the perceptual data provided by BLIRNET, called the *pull-out net*; and (3) an *attentional mechanism* (*AM* for short) that guides the efforts of BLIRNET. To illustrate the typical operation of the system, consider a simple example in which MORSEL is shown a display containing two words, PEA and BOY. These words cause a pattern of activity on MORSEL's "retina," which serves as input to BLIRNET as well as to the AM. The AM then focuses on one retinal region, say the location of PEA. Information from that region is processed by BLIRNET, which activates an orthographic representation suggesting that the item is PEA or possibly TEA, PFA, or RER. The pull-out net then selects the most plausible interpretation of BLIRNET's output, based partly on lexical and semantic knowledge, in this case hopefully PEA. The representation at this level of the system encodes the identity of the word but not its retinal location. Location information is recovered from the AM, which indicates the current location of focus. Shape and location information can then be bound together and stored in a visual short-term memory or used however desired by higher-level systems. Next, attention shifts to BOY and this process is repeated.

Input to MORSEL

Presentation of a visual display causes a pattern of activity on MORSEL's "retina." In the current implementation, the retina is a feature map arranged in a 36×6 spatial array, with detectors for five feature types at each point in the array (line segments at four orientations and line-segment terminator detectors). Letters of the alphabet are encoded as an activity pattern over a 3×3 retinal region. For instance, Figure 2 depicts the retinal representation of PEA BOY.

The letter and word recognition system (BLIRNET)

BLIRNET consists of a hierarchy of processing units, starting at the lowest level with location-specific detectors for primitive visual features — the retinal representation — and progressing to a level composed of location-independent detectors for abstract letter identities. Units at intervening levels register successively higher order features over increasingly larger regions of retinotopic space. The effect of this architecture is that both location invariance and featural complexity increase at higher levels of the system.
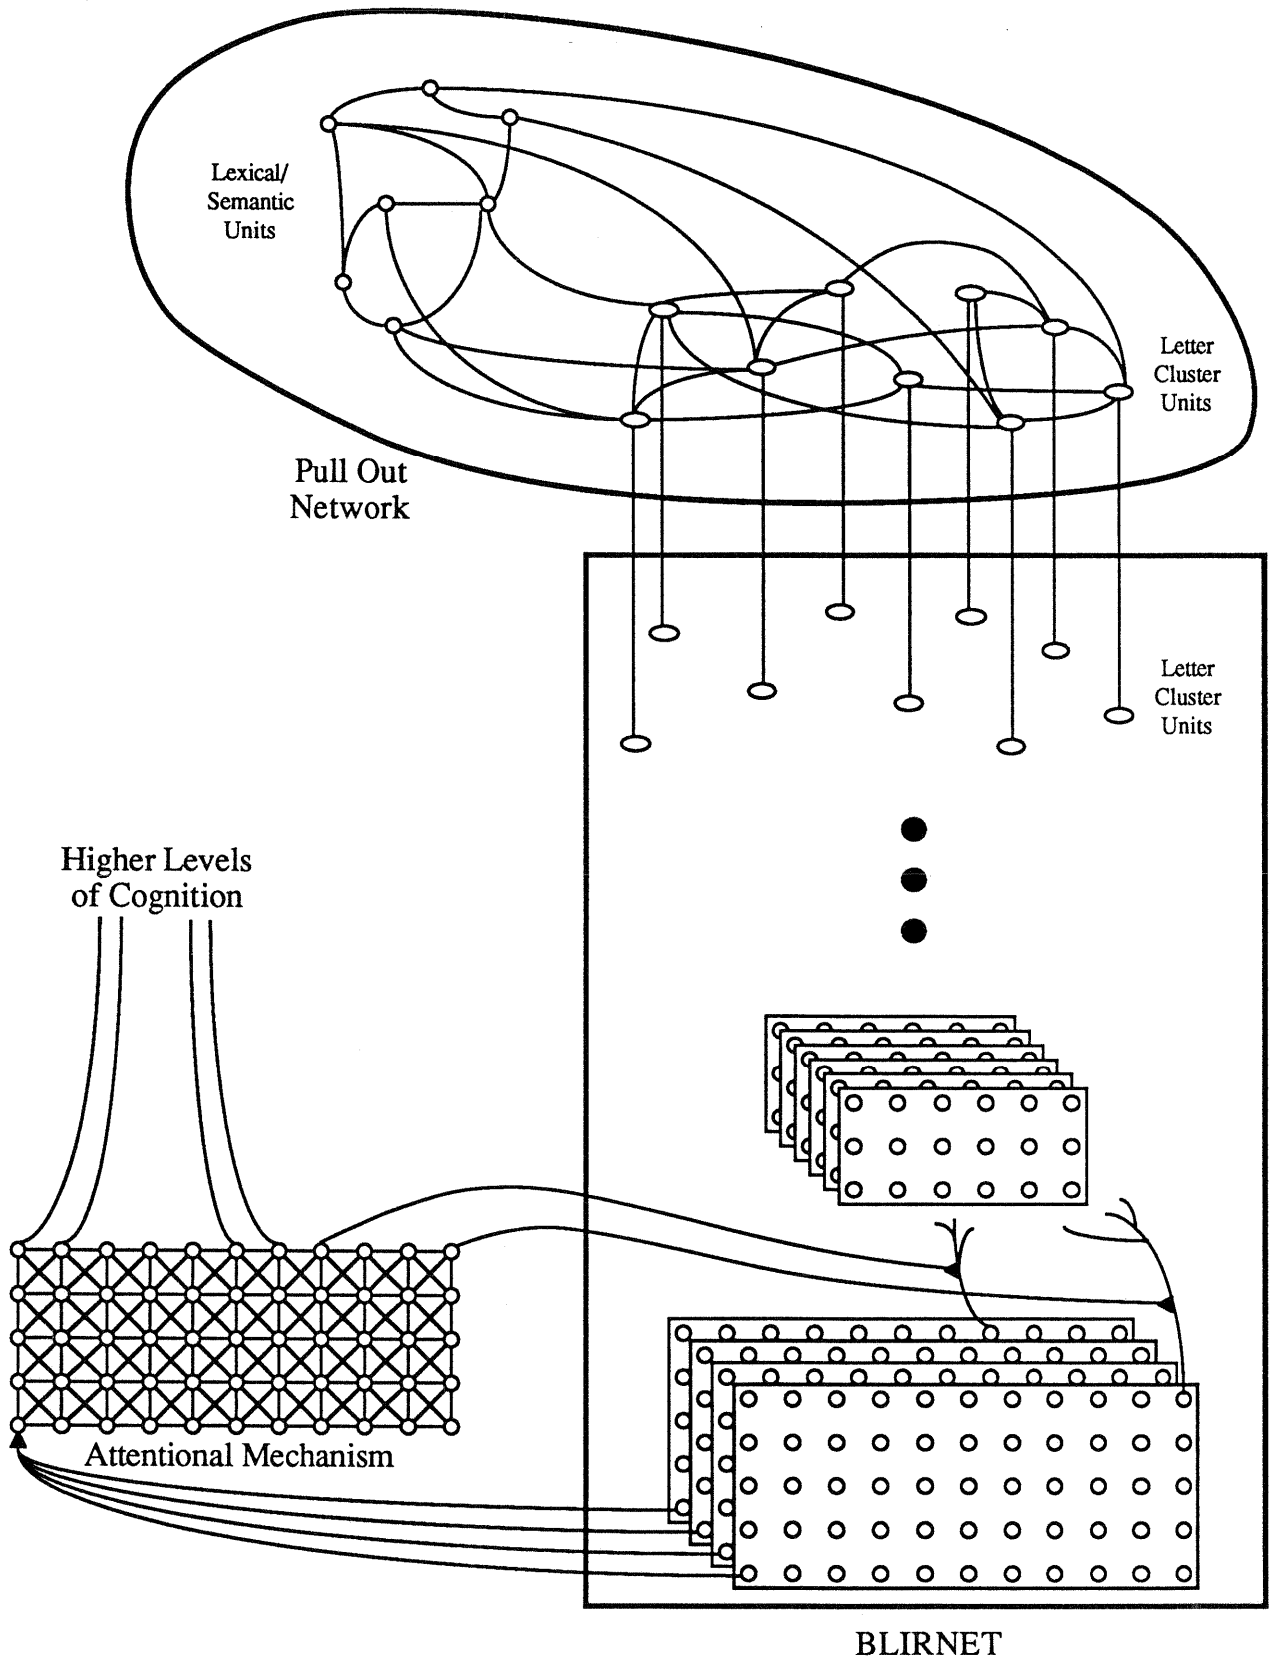
Figure 1. A sketch of the essential components of MORSEL.

Figure 2. The top array shows the superimposed feature activations for a sample input, PEA BOY, positioned on MORSEL's retina. The remaining arrays represent the individual feature maps. Each character in an array represents the activity of a single unit. A "." indicates that the unit is off. A "-", "/", "|", or "\" indicates activity of the corresponding unit in the 0°, 45°, 90°, 135° line segment map, respectively, and "o" indicates activity in the line segment terminator map.

Units in the output layer of BLIRNET have been trained to detect the presence of particular sequences of letters. These *letter-cluster units* respond to local arrangements of letters but are not sensitive to the larger context or the absolute retinal location of the letters. For example, there might be a unit that detects the sequence MON; it would become activated by words like MONEY or DIAMOND.

The letter-cluster units respond to triples of letters in four consecutive slots, either a sequence of three adjacent letters, such as MON, or two adjacent letters and one nearby letter, such as MO_E or M_NE, where the underbar indicates that any single letter may appear in the corresponding position. An asterisk is used to signify a blank space; for example, **M is an M with several spaces to its left. Presentation of MONEY should result in the activation of the following letter-cluster units: **M, **_O, *MO, *_ON, *M_N, MON, M_NE, MO_E, ONE, O_EY, ON_Y, NEY, NE_*, N_Y*, EY*, E_**, and Y**. The representation of words in the output layer of BLIRNET is thus distributed: a word corresponds to a pattern of activity across the letter-cluster units.

In most cases, the letter-cluster coding scheme is *faithful* (Smolensky, 1987), meaning that the set of units associated with a word is unique to that word (but see Prince & Pinker, 1988, for limitations to this type of scheme). The letter-cluster coding scheme also allows for the faithful representation of multiple words in parallel, provided that the words are not too similar (Mozer, 1988a).

The details of BLIRNET's architecture are not particularly important; there are many possible implementations of the same basic idea (e.g., Uhr, 1987; Zemel, Mozer, & Hinton, 1988). The key aspect of BLIRNET is that while it can process multiple letters and words simultaneously in principle, interactions within the network cause a degradation in the quality of analysis as the amount of information to be processed increases. Consequently, when one or more words are presented to BLIRNET, appropriate clusters are not always fully activated and some "spurious" clusters achieve partial activation. These spurious clusters are related to the presented stimuli; they tend to be clusters that would be appropriate if a letter of the stimulus were substituted for a visually-similar letter (e.g., MOV instead of MON), if a letter or two were deleted from or inserted into the stimulus (e.g., ONY or MO_N), or if adjacent letters in the stimulus were transposed (e.g., ENY).
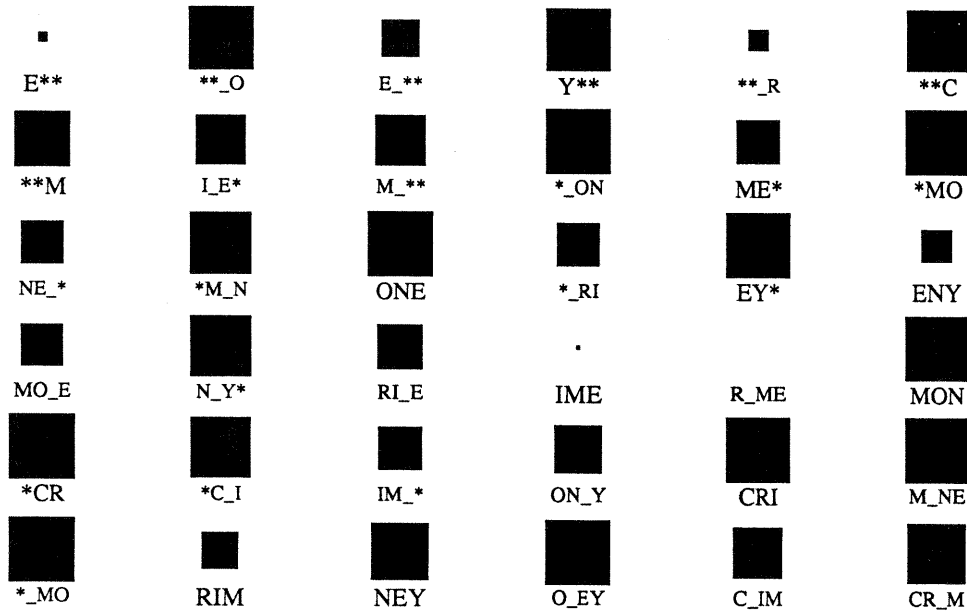
The Pull-Out Network

The noisy pattern of letter-cluster activity produced by BLIRNET is not always easy to interpret. Interpretation is further complicated when several words are processed simultaneously because clusters of one word are entangled with clusters of another. The pull-out network (henceforth, *PO net*) has the task of selecting a set of clusters that represent a single item; it must "clean up" the noise and "disentangle" the hodgepodge of activations from multiple words (Figure 3).
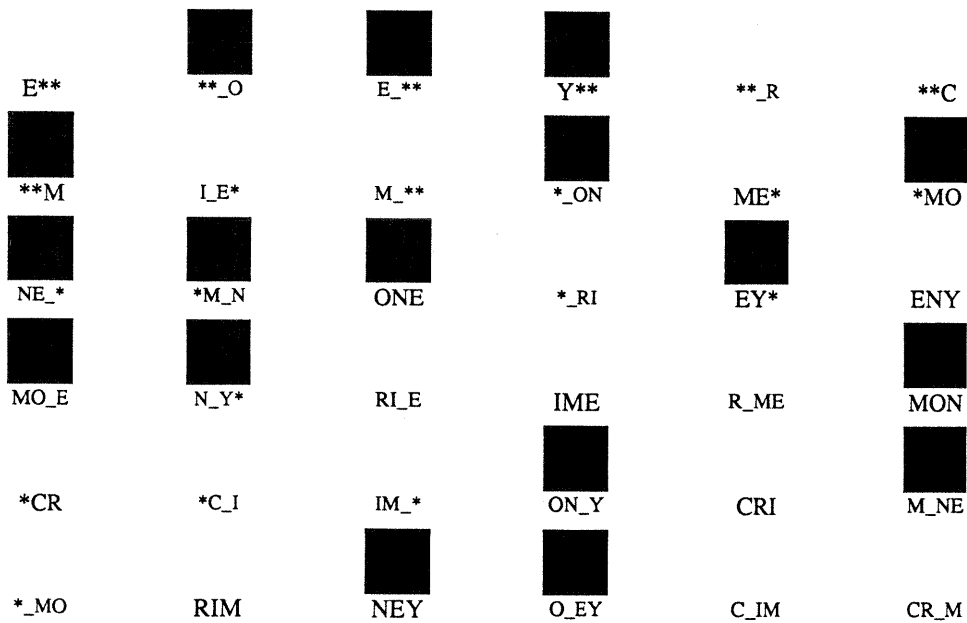
The PO net contains a set of units in one-to-one correspondence with the letter-cluster units of BLIRNET. Each letter-cluster unit excites its corresponding unit in the PO net; thus, the pattern of letter-cluster activity is copied to the PO net. Cooperative and competitive interactions then take place within the PO net to activate a set of letter clusters that exactly correspond to a single letter string. The resulting activity pattern is taken as MORSEL's response.

The basic idea behind the PO net interactions is that *compatible* clusters — ones likely to appear together in a letter string, e.g., MON and ONE — should excite one another and *incompatible* clusters — ones unlikely to appear together, e.g., MON and MOV — should inhibit one another. Thus, the connection strengths are related to how strongly one can predict the presence or absence of one cluster given another cluster. These predictions serve as weak constraints on how the letter clusters might be assembled to form valid strings. The PO net attempts to satisfy as many of these weak constraints as possible while maintaining consistency with the perceptual data. Details of the dynamics are described in Appendix 1. Similar clean up mechanisms have proven useful for recovering information from noisy signals in other connectionist models (Hinton & Shallice, 1989; McClelland & Rumelhart, 1981; Touretzky & Hinton, 1988).

The connections among letter cluster units embody syntactic knowledge about which pairs of clusters can appear together in a letter string. An additional source of information can assist the PO net selection process: higher-order knowledge about valid English words. Some form of lexical or semantic knowledge certainly plays a

E**	**_O	E_**	Y**	**_R	**C
**M	I_E*	M_**	*_ON	ME*	*MO
NE_*	*M_N	ONE	*_RI	EY*	ENY
MO_E	N_Y*	RI_E	IME	R_ME	MON
*CR	*C_I	IM_*	ON_Y	CRI	M_NE
*_MO	RIM	NEY	O_EY	C_IM	CR_M

BLIRNET Activations

E**	**_O	E_**	Y**	**_R	**C
**M	I_E*	M_**	*_ON	ME*	*MO
NE_*	*M_N	ONE	*_RI	EY*	ENY
MO_E	N_Y*	RI_E	IME	R_ME	MON
*CR	*C_I	IM_*	ON_Y	CRI	M_NE
*_MO	RIM	NEY	O_EY	C_IM	CR_M

PO Net Activations

Figure 3. Desired behaviour of the PO net. Suppose that two words, CRIME and MONEY, are processed by BLIRNET, and the resulting activations include those shown in the upper half of the Figure. In this contrived example, clusters of MONEY are, on average, more active than clusters of CRIME. (The area of the black square above a cluster name is proportional to the cluster's activity level.) The lower half of the figure depicts the final activations of letter cluster units in the PO net: clusters of MONEY have been enhanced (e.g., MO_E) and clusters of CRIME suppressed.

role in reading, as abundant evidence suggests that lexical status has a significant effect on performance (e.g., Carr, Davidson, & Hawkins, 1978; McClelland & Johnston, 1977).

The utility of units representing semantic features (*semantic units*) is easiest to envision if word meanings are represented locally, that is, if there is one semantic unit per word meaning. For instance, suppose there was a semantic unit representing the "wealth" sense of MONEY. It would be connected to all clusters of MONEY. Activation of some clusters of MONEY would result in activation of the "wealth" semantic unit, which in turn would reinforce these clusters and help activate the remaining ones. Inhibitory interactions among the semantic units are also necessary to prevent multiple meanings from being activated simultaneously. The end result of the pull out process is then selection of one internally-consistent spelling pattern in the letter-cluster units and one word meaning in the semantic units.

In the current implementation, the semantic unit representation is semi-distributed: there are many semantic units corresponding to each word meaning, but each semantic unit is associated with only one word meaning. Thus, the "semantic" units are actually a lexical representation, albeit a distributed representation, so to be honest we call them *semlex units*. However, the only reason for not constructing a fully-distributed semantic representation is the difficulty of devising a complete set of semantic features.[4]

On grounds of parsimony, we would like to believe that an explicit lexical representation is not necessary; the semantic representation can serve the same function in the pull out process and is necessary in any case to represent word meanings. Further, the architecture we propose — direct association between orthographic and semantic knowledge without mediation by a lexicon — is entirely consistent with Hinton and Shallice's (1989) model of acquired dyslexia.

The Attentional Mechanism

MORSEL has an attentional mechanism, the AM, that controls the amount and temporal order of information flowing through BLIRNET. The AM receives input about where to focus from various sources, resolves conflicting suggestions, and then constructs a "spotlight" centered on the selected region of the retina. The attentional spotlight serves to enhance the activation of input features (such as those depicted in Figure 2) within its bounds relative to those outside. As activity is propagated through BLIRNET, the highlighted region maintains its enhanced status, so that at the output of BLIRNET, letter-cluster units appropriate for the attended item tend to become most active as well. Consequently, the PO net will choose the attended item. Note that attention causes the preferential processing of certain items, but it does not act as an all-or-none filter. Information from the unattended regions of the retina undergoes some degree of analysis by BLIRNET.

The attentional system receives input about where to focus from two sources. First, attention can be guided in a *bottom-up* manner by stimulus information so as to bias selection towards locations where stimuli are actually present. Second, higher-levels of cognition can supply top-down control on the basis of task demands. For instance, if the task instructions are to report the left item in a multi-item display first, selection can be biased towards the left portion of the display initially; if the instructions are to read a page of text, a scanning mechanism can bias selection towards the top-left corner initially, and then advance left to right, top to bottom. (Butter, 1987, argues for a similar distinction between "reflex" and "voluntary" control of attention in humans.)

As shown in Figure 1, the AM is a set of units in one-to-one correspondence with the retinotopic feature maps serving as input to BLIRNET. Activity in an AM unit indicates that attention is focused on the corresponding retinal location and serves to *gate the flow of activity* from the input layer to the next layer of BLIRNET. Specifically, the activity level of an input unit in a given location is transmitted to the next layer with a probability that is monotonically related to the activity of the AM unit the the corresponding location. However, the AM serves only to bias processing: it does not absolutely inhibit activations from unattended regions, but these activations are transmitted with a lower probability.

---

[4] It might seem implausible that a distributed orthographic representation could have any systematic relationship to a distributed semantic representation, but Hinton and Shallice (1989) have demonstrated otherwise in a model that learns orthographic-semantic associations.

Each unit in the AM gets bottom-up input from the corresponding location in all of the retinotopic feature maps, as well as an unspecified top-down input. The dynamics of the AM generate a single, contiguous region of activity over the retinotopic space, with a bias towards locations indicated by bottom-up and top-down inputs. Details of the AM selection process are provided in Appendix 2.

Key Properties of MORSEL

Four properties of MORSEL are essential in accounting for the behaviour of neglect dyslexia patients.

1) Attentional selection by location occurs early in the course of processing. With all other things being equal, there is a preference for locations where stimuli appear.
2) Attention attempts to select a single item. In this regard, an item is defined as a relatively dense bundle of features separated from other bundles by a relatively sparse region. This crude definition does not always suffice, but it allows for early segmentation of the image without higher-order knowledge.
3) Attention gates the flow of activity through the object recognition system. The activities of features outside the attended region are relatively attenuated but not completely suppressed. Consequently, unattended information receives some degree of analysis.
4) After the recognition system has processed the perceptual data in a bottom-up fashion, a clean-up mechanism acts on the resulting representation to recover information that is orthographically and semantically meaningful. This clean-up mechanism can compensate for noise and inaccuracy in the recognition system and in the perceptual data itself.

Any model with these four properties should suffice for the present purpose. There is surely a large class of models with these properties; MORSEL is not unique. The interesting thing about MORSEL is that it was developed to account for a variety of perceptual and attentional data in normal subjects, but, as we will show, it is entirely consistent with the neglect dyslexia data as well. It would be difficult to justify the development of a simulation model as large as MORSEL for the present purpose alone. However, the details of MORSEL had been worked out previously, except for a few tweaks and extensions to specify aspects of the model that were not previously required (e.g., the semlex unit representation).

Many details of MORSEL (e.g., the letter-cluster representation, the operation of BLIRNET) are not critical in the present work. Consequently, we have no strong commitment to the nuts and bolts of MORSEL, only to the framework that it provides. In fact, if we have any commitment at all, it is to the belief that the nuts and bolts are *wrong*. The input representation is not rich enough; the AM dynamics are too brittle; the PO net is not based on a rigorous computational foundation (cf., Hopfield, 1982). Nonetheless, we experimented with a wide variety of alternatives to the mechanisms and parameters reported in this paper, and were pleased to discover that the qualitative behaviour of the model was remarkably insensitive to these details.

Damaging MORSEL to Produce Neglect Dyslexia

We propose that neglect dyslexia results when the bottom-up connections to the AM from the input feature maps are damaged. The damage is graded monotonically, most severe at the left extreme of the retina and least severe at the right (assuming a right hemisphere lesion, as we have throughout the paper).[5] This account may be contrasted with one claiming that the damage to connections in the left field is absolute and connections in the right field are entirely intact.

The consequence of the damage is to affect the probability that features present on the retinotopic input maps are detected by the AM. To the extent that features in a given location are not detected, the AM will fail to focus attention at that location. Note that this is not a "perceptual" deficit, in the sense that if somehow attention can be

---

[5] We deliberately avoid the issue of where eye fixation rests with respect to the retinotopic map, and hence, which input information is processed by which cerebral hemisphere. The only strong claim we wish to make is that, regardless of hemifield, the left-right gradient of damage is present. However, the absolute severity of damage may show a sharp discontinuity when crossing from one hemifield to the other (Mesulam, 1985), and the quantitative nature of the gradient and discontinuity may differ from one patient to another.

mustered, features will be analyzed normally by BLIRNET.

To give the gist of our account, MORSEL and the hypothesized deficit are compatible with the early, peripheral effects observed in neglect dyslexia because the disruption directly affects a low-level representation. MORSEL is also compatible with the late, higher-order effects in neglect dyslexia: The PO net is able to reconstruct the elements of a string that are attenuated by the attentional system via lexical and semantic knowledge.

Simulations of neglect dyslexia

We now turn to a detailed description of the performance of patients with neglect dyslexia and demonstrate through simulation experiments how the lesioned version of MORSEL can account for these behaviours. The patient descriptions and simulation results are grouped according to six basic phenomena. The first three — extinction, modulation of attention by task demands, and the effect of retinal presentation position on accuracy — appear to arise at an early stage of processing, while the last three — relative sparing of words versus nonwords, distinctions in performance within the class of words, and the influence of lexical status on extinction — appear more compatible with a deficit localized at later stages of processing. MORSEL provides a unifying framework to account for these disparate behaviours.

An important finding in neglect dyslexia (and in neuropsychology in general) is that there is great variability in performance across patients. We have chosen a set of phenomena that seem relatively common, and for which some agreement is found in the literature. While it is sensible to begin with phenomena that have been reliably observed, we fully believe that understanding individual differences is likely to be of as much interest as similarities in behaviour.

The extinction effect

A well-documented finding in the literature on neglect is that a patient who can detect a single contralesional stimulus may fail to report that stimulus when a second stimulus appears simultaneously in the ipsilesional space. This phenomenon, termed *extinction*, has been reported to occur with visual, tactile, and auditory stimuli and has a direct analog in reading. When two words are presented simultaneously in the two visual fields, patients tend to neglect the contralesional stimulus. Sieroff and Michel (1987) demonstrated further that with a single word centered across the fovea and subtending the same visual angle as the two non-contiguous words, extinction of information in the contralesional hemifield is less severe. In a similar experiment, Behrmann et al. (1989) showed that a compound word (such as PEANUT) is read better when the two component morphemes (PEA and NUT) are physically contiguous than when they are separated by a single blank space. Further, when the two words are separated by a pound sign (PEA#NUT), performance is still better than in the spaced condition, despite possible perceptual complications introduced by the pound sign, lending additional support to the conclusion that extinction is strongly dependent on the physical separation between items in the display.

The phenomenon of extinction is consistent with the view that the visual attentional system attempts to select one of multiple items in the visual field; in neglect patients, the selection is heavily biased towards the rightmost item. An "item" here can simply be defined by the physical adjacency of its components and physical distinctiveness from its neighbors. (We conjecture that the distinctiveness need not be one of physical separation; any simple property such as color or texture boundaries could suffice.)

MORSEL's AM operates exactly in this manner. In the unlesioned model, when two three-letter words are presented to the AM, attention selects the left word on 41.3% of trials and the right on 40.8%; some combination of the two words is selected on the remaining 17.9% of trials. (See Appendix 3 for details of this and other simulations involving the AM.) In the lesioned model, the right word is nearly always selected because the bottom-up input to the AM from the retinotopic feature maps is degraded for the left word, thereby weakening its support. Figure 4 illustrates the bottom-up input detected by the lesioned AM upon presentation of two three-letter words. Two blobs of activity are apparent, corresponding to the two words, but the left blob is weaker. The consequence of this left-sided degradation can be seen in Figure 5 which shows activities of the AM units over time arising from this input. The AM settles on the right word.

Figure 4. The bottom-up input detected by the lesioned AM upon presentation of two three-letter words. The area of a white square at a given location indicates the relative strength of the input at that location in the retinotopic map. The black dots indicate the locations in the map for which there is no input.

Table 1 shows the distribution of attention in the lesioned model for displays containing two three-letter words. Each row indicates the percent of presentations in which a given combination of letters is selected; "1", "2", and "3" are letters of the left word, "4", "5", and "6" letters of the right word. The right word is selected over 75% of the time, with the remainder of the presentations involving selection of the right word along with the rightmost portions of the left, or selections of only the rightmost portions of the right word. The AM clearly demonstrates extinction of the left item when two words are presented. However, when a single item is presented, either to the normal or the lesioned model, at least some portion of the item will always be attended (as we discuss in more detail below).

In the normal model, when two items are presented, one will be selected arbitrarily. If the AM is allowed to refocus on the same stimulus display, it will select the other item about half the time. Thus, simply by resetting the AM and allowing it to settle again, possibly with a slight inhibitory bias on the location just selected, both display items can be sampled. In the lesioned model, however, refocusing attention is unlikely to alter the selection. As long as the right item is present the left item is prevented from attracting attention; this masking does not occur in the normal model.

Because the AM serves only to bias processing in BLIRNET toward the attended region, as opposed to completely filtering out the unattended information, MORSEL will not necessarily fail to detect the unattended information. This depends on the operation of the PO net, which attempts to combine the outputs of BLIRNET into a meaningful whole. Thus, one cannot directly translate the distribution of attention into a distribution of responses. Nonetheless, the strong right-sided bias will surely affect responses, particularly for simple stimuli which cannot benefit from the PO net's application of higher-order knowledge. For instance, in the task of detecting a single or a pair of simultaneously-presented flashes of light, commonly used to test extinction, responses can only be based on the stimulus strength following attenuation by the AM.

Table 1

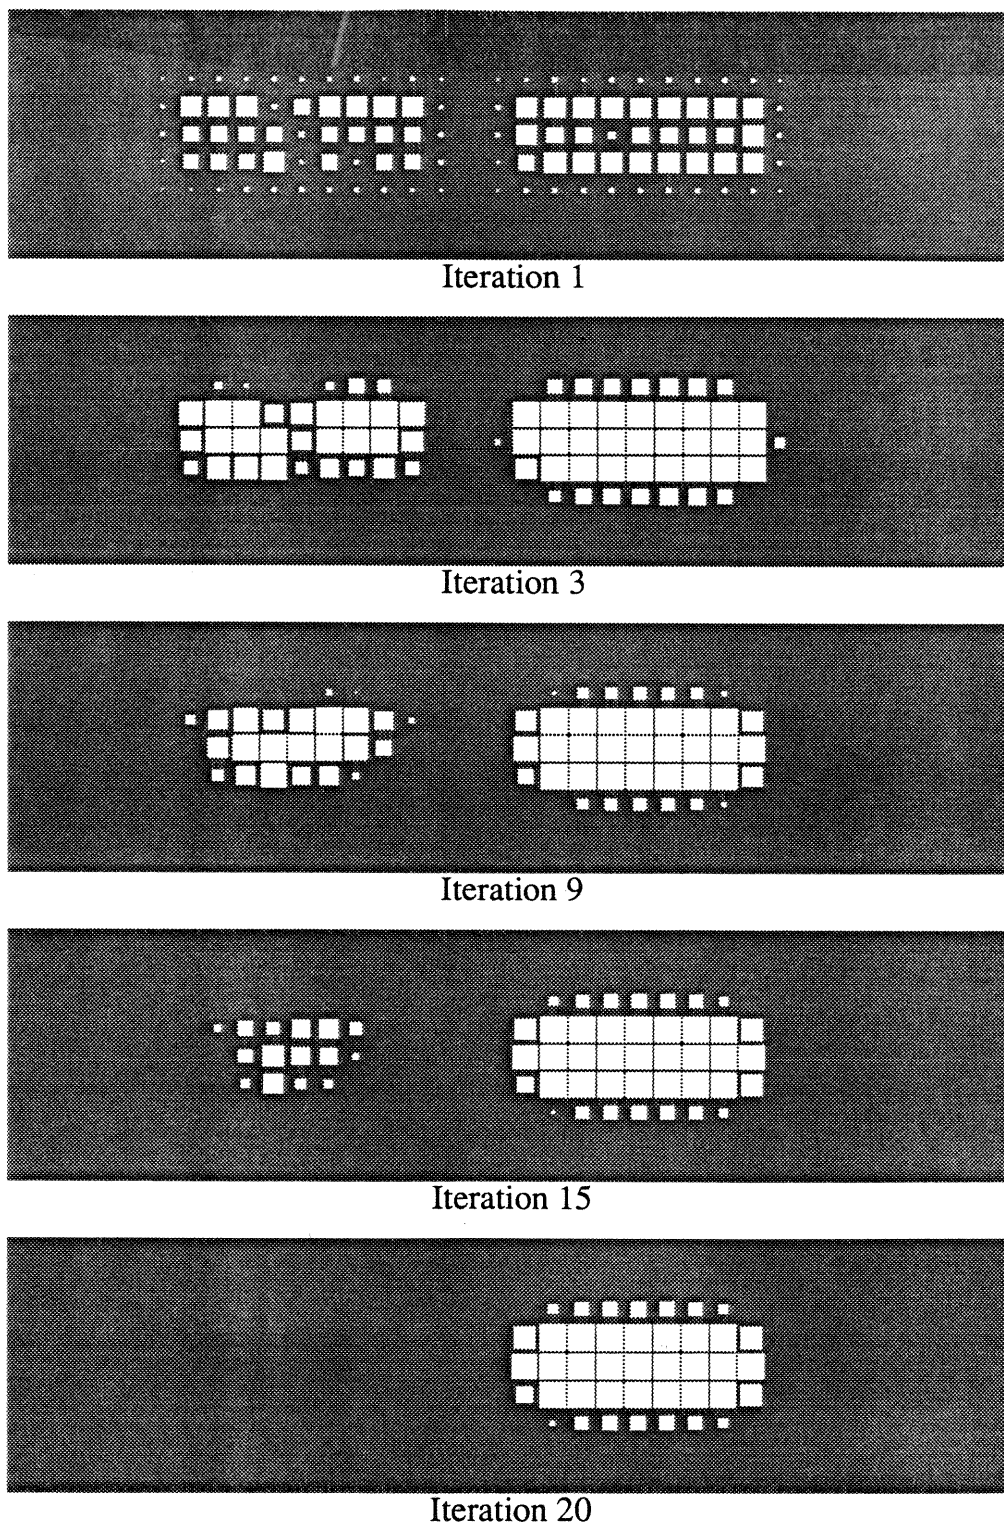| Distribution of Attention in the Lesioned AM for Displays Containing Two Three-Letter Words	
Letters Attended	Relative Likelihood of Attentional State
1 2 3 4 5 6	6.6%
2 3 4 5 6	9.7
3 4 5 6	0.1
4 5 6	76.2
5 6	7.2
6	0.2

Figure 5. Activities of the AM units at several points in time as the right word is selected. By iteration 20, activities within the AM have reached equilibrium.

Modulation of attention by task demands

The strong predominance of right-biased responses in neglect patients can be modulated under certain conditions. Butter (1987) has suggested that the rightward orientation of these patients is a reflexive or involuntary response but that attention can be willfully deployed to the left. Karnath (1988) showed that patients always reported the right-sided stimulus first when given the free choice of order of naming two bilaterally presented stimuli. The left-sided stimulus was often neglected in these cases. When patients were instructed to report the left-sided stimulus first, they were able to report both stimuli. A similar result in the domain of reading was found by Behrmann et al. (1989). One of their patients with neglect dyslexia (AH) reported the left-sided word on only 4% of trials when two words were presented simultaneously. When instructed to report the left-hand word first, AH reported both words correctly on 56% of trials.

An overt attentional shift provided by cuing patients to a stimulus on the left has been shown to overcome the neglect deficit in other tasks too. For example, Riddoch and Humphreys (1983) placed a single letter at each end of a line and instructed their patients to report the identity of the letters prior to bisecting the line. The degree of neglect on the line bisection task was significantly reduced with the additional letter reporting task. The above findings suggest that the distribution of attention can be influenced by task instructions.

In MORSEL, two sources of information can guide attention: bottom up and top down. These two sources simply add together to bias the selection of a location. In the lesioned model, the bottom-up inputs for the left portion of the retina are weakened, but the top-down inputs are undamaged; hence, sufficiently strong top-down "task driven" guidance can compensate for the deficit in bottom-up control of attention. Figure 6 illustrates the effect on the AM when a top-down input to the left field is superimposed on the degraded bottom-up input shown in Figure 4. Without the top-down input, the right word would have been selected (Figure 5). This bias, however, compensates for the bottom-up degradation and the left word is selected.


Figure 6. The upper diagram depicts the top-down input to the AM superimposed on the degraded bottom-up input of Figure 4. The lower diagram depicts the resulting equilibrium state of the AM. The left word is selected here, but without the top-down input, the right word would have been selected (Figure 5).

This example makes the point that the deficit in MORSEL is attentional and not perceptual. A true perceptual deficit would occur if, say, the connections within BLIRNET were lesioned. Our account of neglect dyslexia places the locus of damage outside of the recognition system; further, the effect of the damage on perception can be overcome via alternative routes — the top-down inputs. That neglect is primarily an attentional deficit is widely held in the neuropsychological literature (Heilman, Watson, & Valenstein, 1985; Kinsbourne, 1987; Mesulam, 1981; Posner & Petersen, 1989).

The effect of retinal presentation position on accuracy

One finding in the literature compatible with a deficit at an early stage of processing is that performance changes as a function of stimulus location. Behrmann et al. (1989) presented words to a neglect dyslexia patient with their left edge immediately next to a central fixation point (the *near position*) or in the fourth character position to the right of fixation (the *far position*). Words appearing in the far position were still in the region of high acuity in the patient's intact visual field. The words were three to five letters in length. The patient reported only 28% of the words correctly in the near position, but 44% in the far position. This finding was confirmed with a second set of six and seven letter words in which 39% and 77% of the words were reported correctly from the near and far positions, respectively. Thus, performance improved as the stimuli were displaced farther into ipsilesional space. This result is also obtained using a line bisection task in which the severity of neglect decreased for lines appearing further to the right (Butter, Mark, & Heilman, 1989).

The effect of presentation position argues that attention must be operating at least partially in a retinotopic reference frame, as opposed to an object-centered frame. If neglect occurred with respect to an object-centered frame, the left side of an item might be neglected relative to the right, but the stimulus position in the visual field would not matter.[6] The evidence for attention operating on a retinotopic frame supports an early selection view, i.e., the attentional system chooses among stimuli based on a low-level representation.

While this conclusion is clearly consistent with the architecture of MORSEL, it requires a bit of explanation to see how MORSEL accounts for the effect of presentation position on accuracy. We begin with an overview of the account. Consider first the normal model being shown a single word. Independent of word length, if the letters are arranged sufficiently close to each other, the AM will always select the region of retinotopic space corresponding to the entire word. In the lesioned model, however, the input strength of the left side of the word is less than the right side, often causing the left side to be suppressed in the AM selection process. Consequently, BLIRNET analyzes the word with a relative degradation of the left side. This degradation propagates through BLIRNET, and to the extent that it prevents the PO net from reconstructing the word's identity, accuracy will be higher in the normal model than in the lesioned model. The same reasoning applies with the lesioned model alone when considering presentation of a word on the relative right versus the left. The farther to the right the word appears, the stronger and more homogeneous its bottom-up input to the AM, and the less likely the AM will be to neglect the leftmost letters. Consequently, accuracy will be higher.

Figure 7 illustrates three examples of the AM suppressing the left side of a six-letter word: in the top row, the rightmost five letter positions are attended; in the middle row, four letters are attended; and in the bottom row, three letters are attended. Table 2 summarizes the distribution of attention for a six-letter word presented to the AM in each of three retinal positions. The "standard" position refers to the presentation position used in Figure 7; the shifted positions refer to moving the word one or two letter positions (3 or 6 pixels) to the right of the standard position. As expected, when the word is moved farther to the right, the AM is more likely to focus on its initial letters.

---

[6] Although eye movements have not been carefully controlled for in these studies, the possibility of eye movements cannot fundamentally alter our conclusion concerning the role of the retinotopic frame, for the following reason. If patients were able to foveate on the stimuli then, independent of presentation position, all items would be analyzed in approximately the same retinal position. To explain the effect of presentation position on performance, one would need to postulate that it is easier to move the eyes to a given location than to a location on its relative left. Because eye movements and attention shifts are intertwined, this is tantamount to claiming that attention operates in a retinotopic frame as well.
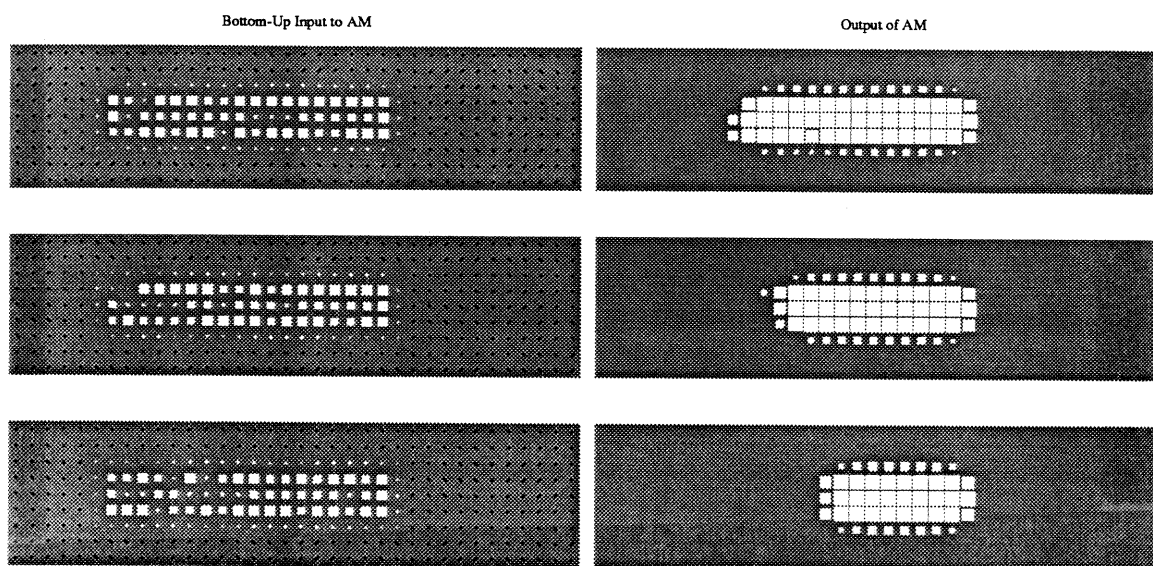
Figure 7. The bottom-up input to the lesioned AM and the resulting AM equilibrium state for three different presentations of a six-letter word. In the top row, the rightmost five letter positions are attended; in the middle row, four letters are attended; and in the bottom row, three letters are attended.

Table 2

Distribution of Attention in the Lesioned AM for Displays Containing One Six-Letter Word			
	Relative Likelihood of Attentional State		
Letters Attended	Standard Position	Shifted Right One Position	Shifted Right Two Positions
1 2 3 4 5 6	8.1%	18.2%	37.2%
2 3 4 5 6	14.6	24.5	31.9
3 4 5 6	30.1	33.7	25.8
4 5 6	33.0	20.0	5.0
5 6	13.9	3.6	0.1
6	0.3	0.0	0.0

The attentional focus produced by the AM affects BLIRNET's processing of a word and ultimately, the accuracy of report. Although we are interested in the accuracy of report, we have chosen not to simulate the detailed operation of BLIRNET for two reasons. First, the version of BLIRNET implemented by Mozer (1988a) was trained to recognize a relatively small set of letter clusters — about 600 of the approximately 6500 needed to represent most English words. The present simulations require a much larger set of letter clusters, and the training procedure is quite computation intensive. Second, the exact activity levels produced by BLIRNET are not critical for the present modeling effort, and in fact, simulation of a large network like BLIRNET obscures the essential properties that are responsible for interesting behaviours. Consequently, rather than simulating BLIRNET, we have incorporated its essential properties into a simple algorithm that determines letter cluster activations for a particular input stimulus and attentional state (see Appendix 4 for further details).

In Figure 8, one can see the simulated activations of various letter cluster units in response to the stimulus PARISH on a trial where the AM has successfully focused on all six letters of the word. Activity levels range from zero to one. The activity of a cluster is indicated by the area of the black square above it. The letter clusters of PAR-ISH (first row of Figure) are highly active. In addition, clusters with letters visually similar to the stimulus word are

Figure 8. Activations of various letter cluster units in response to the stimulus PARISH on a trial where the AM has selected all six letters of the word.

partially activated, for example, RTS, PA_T, RA_I, and DIS, as are clusters that would be appropriate were letters of the stimulus slightly rearranged, for example, AR_I and I_S*. Finally, a bit of noise is thrown into the activation process which creates random fluctuations in the activity pattern.

If only the last three letters of PARISH are attended, the resulting pattern of letter cluster activity looks quite different (Figure 9). Clusters representing the initial segment of the word are less active than in Figure 8. Further, because the initial segment is suppressed, clusters such as **I and **_S will become more active, as if ISH was presented instead of PARISH.

The next stage in processing the stimulus is to feed the output of BLIRNET to the PO net, allow the PO net to settle, and then determine which of a set of alternative responses best matches the final PO net activity pattern. (The procedure for selecting alternative responses — and which letter clusters to include in the PO net simulation — is explained in Appendix 4.) In the case of the fully attended PARISH (Figure 8), the PO net almost always reads out the correct response. In the case of the partially attended PARISH (Figure 9), the PO net often is able to reconstruct the original word; other times it fabricates a left side, reading out instead RADISH or POLISH or RELISH; and occasionally it just reads out the attended portion, ISH, although the influence of the semlex units acts against the read out of nonwords.

To test the effect of stimulus presentation position in MORSEL, we conducted a simulation using six six-letter words: PARISH, BEGGAR, FOSTER, SILVER, MORSEL, and SHADOW. Although the obvious way to test MORSEL is to present a stimulus on the retina, allow the AM to settle, determine the resulting BLIRNET activations, feed these to the PO net and read out a response, we have decoupled the AM and PO net simulations to reduce the computational burden. Running a simulation of the AM alone on a six-letter stimulus, we can determine the probability of the AM selecting a particular combination of letters (the *attentional state*, see Table 2). Independently, the PO net simulation can be run in its entirety for each possible attentional state. The probability of being in attentional state i, $p(state\ i)$, can then be combined with the probability of the PO net responding correctly given a particular attentional state, $p(correct\ |\ state\ i)$, to yield an overall probability of correct response:

$$p(correct) = \sum_i p(state\ i)\,p(correct\ |\ state\ i) .$$

Table 3 presents the results of the PO net simulation on our collection of six words, each line showing the accuracy in a particular attention state. These figures are averaged across the six words and 100 replications of each word. The replications are necessary to obtain a reliable measure of accuracy because noise introduced by BLIRNET can cause different responses on each trial. (See Appendix 4 for further details of the simulation methodology.) The Table indicates that performance drops as fewer letters of the word are attended. Even with only one letter attended, the residual accuracy is quite high, no doubt due to the partial activation of unattended information. Combining the conditional probabilities of Table 3 with the marginal probabilities of being in a given attentional state of Table 2, one obtains an overall probability of correct response: 49% for words presented in the standard position, 63% for words one position to the right, and 79% for words two positions to the right.

Table 3

| Performance of Lesioned MORSEL on Displays Containing One Six-Letter Word	
Letters Attended	Correct Responses Given Attentional State
1 2 3 4 5 6	100%
2 3 4 5 6	85%
3 4 5 6	51%
4 5 6	33%
5 6	17%
6	18%

Figure 9. Activations of various letter cluster units in response to the stimulus PARISH on a trial where the AM has selected the last three letters of the word.

Thus, the peripheral lesion in MORSEL does result in a retinotopic deficit as measured by reading performance. Performance is better than would be expected by examining the distribution of attention alone, thanks to the reconstruction ability of the PO net: Although the entire word is attended on only 8% of trials (for the standard position), the word is correctly reported far more frequently — 49% of trials. Nonetheless, the retinal position of the stimulus does come into play; the PO net is not so effective that accuracy is absolute.

Not surprisingly, when MORSEL does produce an error, the error generally occurs on the left side of a word. For example, with PARISH, the alternative responses include left-sided completions such as POLISH or IRISH and right-sided completions such as PARKER or PARTS, yet the PO net always prefers the left-sided completions. Figure 10 shows a graph of activity over time for the stimulus PARISH on a trial where the AM has selected just the right side — ISH. On this trial, the PO net eventually reads out POLISH.

The account provided by MORSEL suggests that neglect — the difficulty in reading single words — goes hand in hand with extinction — the difficulty in selecting one of two items. Both behaviours are caused by the same underlying deficit. This does not imply, however, that the two behaviours must necessarily co-occur. With a milder gradient of damage than the one we have simulated, MORSEL shows minimal neglect in reading words due to the compensating action of the PO net, yet even a slight right-sided bias leads to extinction. This is consistent with reports in the literature: Neglect and extinction generally co-occur, and on the path to recovery, neglect diminishes in severity, leaving extinction as the only manifestation of the brain damage (Kolb & Whishaw, 1985). At present, there are no data in the domain of reading that challenge MORSEL's claim that both neglect and extinction of words are caused by the same deficit; however, several studies from the general hemispatial neglect literature find a double
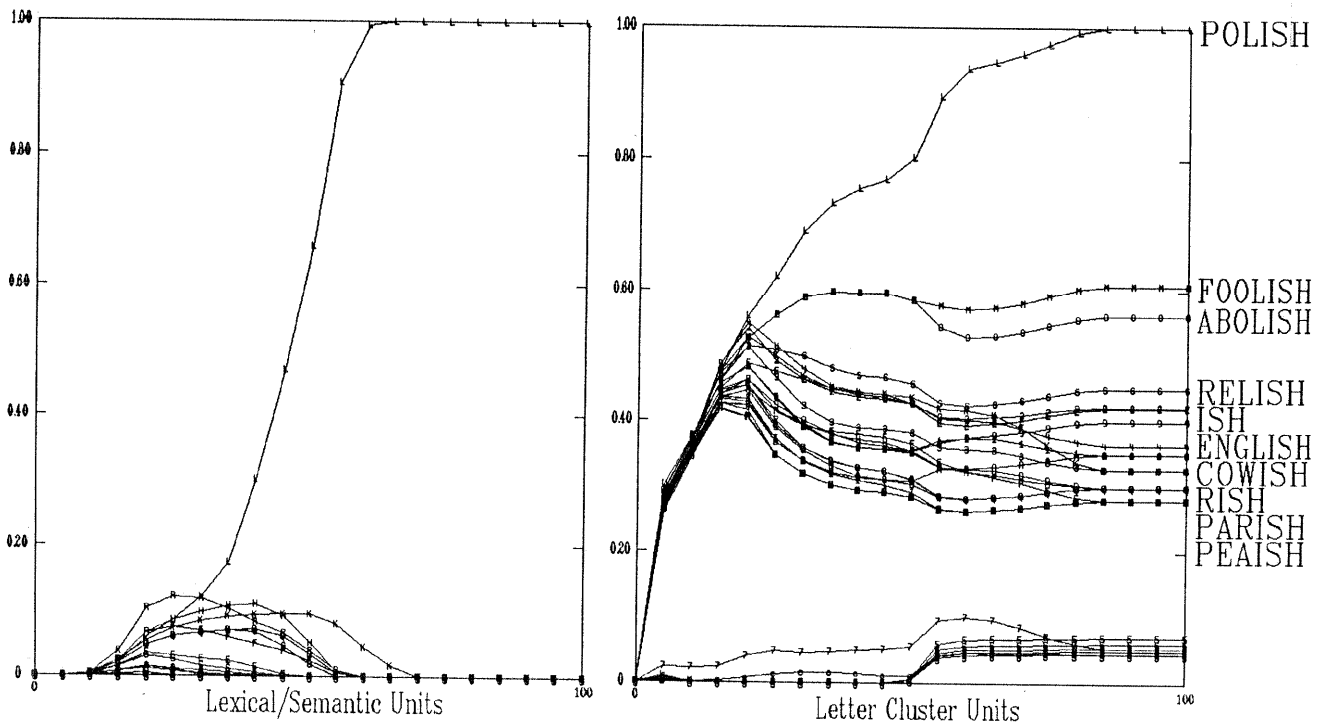
Figure 10. Activity in the PO net as a function of time for the stimulus PARISH on a trial where the AM has selected the last three letters of the word. The PO net eventually reads out POLISH. The left graph summarizes activations of the semlex units over time, and the right graph activations of the letter cluster units. Each trace represents aggregate activity of a particular response and is labeled with a single digit or letter symbol. Traces corresponding to the same word in the two graphs use the same symbol. The ten most active responses in the letter cluster graph are shown on the far right, roughly next to the corresponding trace. The aggregate semlex activity of a response is simply the average activity of its semlex units. The aggregate letter cluster activity of a response is a measure of the activity of the target clusters relative to nontarget clusters (see Appendix 4). Note that due to the distributed word representations, there will always be partial activity of response alternatives similar to the chosen one.

dissociation (Bisiach, Perani, Vallar, & Berti, 1986; Ogden, 1985).

Although the deficit in MORSEL occurs with respect to a retinotopic reference frame, one could easily interpret it as occurring with respect to an object-based frame. This is because, irrespective of the retinal presentation position, the left part of a word tends to be reported more poorly than the right. Indeed, Baxter and Warrington (1983), in finding neglect errors for short as well as long words, suggest that the phenomenon is due to the "faulty distribution of attention to the central representation of a word." MORSEL allows the data to be interpreted from a different perspective, one in which a deficit that produces a retinotopic gradient can lead to a relative difference between the left and right components of an object, independent of the object's size and position.

Relative sparing of words versus nonwords

A general finding in the neglect dyslexia literature is that the reading of words is less affected by neglect than the reading of nonwords. For example, Sieroff, Pollatsek, and Posner (1988) demonstrated that their patients with right parietal lesions showed superior overall performance on words compared to nonpronounceable nonwords. In addition, these patients also exhibited a more marked asymmetry (right side better than the left) for nonwords than for words. The relative superiority of words over nonwords was observed both under brief tachistoscopic presentation of the stimuli and under unlimited exposure duration. This difference in performance as a function of lexical status has also been replicated in several other studies, using pronounceable nonwords (also called *pseudowords*) as well as nonpronounceable nonwords (Behrmann et al., 1989; Brunn & Farah, 1989; Sieroff, 1989).

Sieroff and Posner (1988) reproduced this effect in normal subjects by modulating attention to foveally presented words. They instructed their subjects to report the identity of a cue prior to reading the target; the cue was a single digit appearing to the immediate left or right of the target. As in the case of neglect dyslexia, performance on words is significantly better than on nonwords. The locus of this word superiority effect is controversial. One popular explanation is that word reading is attention free because the orthographic string makes direct contact with its existing lexical entry (LaBerge & Samuels, 1974; Sieroff, Pollatsek, & Posner, 1988). Such a view affords privileged processing status to words. Nonwords, on the other hand, do not benefit from this mode of lexical access and are subject to attentional control. The implication of such a view is that two distinct modes of processing exist. It is not clear, however, where the two paths diverge — at an early level prior to the encoding of the integrated "word form" (Warrington & Shallice, 1980) or as a means for sequential readout of information into phonological or semantic codes (Mewhort, Marchetti, Gurnsey, & Campbell, 1984).

An alternative interpretation, which has been used to account for the perceptual advantage of letters in words over letters in nonwords, is that letters in words are enhanced by feedback from an existing lexical representation (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982). Such feedback does not benefit nonpronounceable nonwords. On this account, letter strings are processed through the same channel independent of lexical status. This account can explain the word advantage in neglect dyslexia: The superiority of words is obtained from the fact that partially-encoded contralesional information may be enhanced by lexical support in the case of words but not in the case of nonwords (Brunn & Farah, 1989; Sieroff, Pollatsek, & Posner, 1988).

This latter account is embodied in MORSEL. Specifically, the PO net acts to recover the portion of a letter string suppressed by the AM using both orthographic knowledge (the connections among letter cluster units) and semantic/lexical knowledge (the connections between letter cluster and semlex units). This gives words an advantage over pseudowords, which lack the support of semantic/lexical knowledge, and a double advantage over nonpronounceable nonwords, which lack the support of orthographic, lexical, and semantic knowledge.

We conducted a simulation study using the lesioned version of MORSEL to compare performance on five-letter words and pseudowords (Table 4). The two conditions differ in that words have an associated representation in the semlex units whereas pseudowords do not. In the first stage of the simulation, we measured the likelihood of the AM attending to a given portion of a five-letter stimulus string (second column of Table 5). Then, the PO net simulation was run for 100 replications of each stimulus in each attentional state to obtain the probability of a correct response for a given stimulus type in a given attentional state (third and fourth columns of Table 5). Combining the AM and PO net simulation results as described for the previous simulation, we obtain an overall probability of correct response: the lesioned MORSEL correctly reported 39% of words but only 7% of pseudowords. In

Table 4

Stimuli for Word/Pseudoword Simulation					
Words			Pseudowords		
CATCH	PRESS	WATER	FATCH	FRESS	SATER
TRUCK	CRIME	MONEY	DRUCK	TRIME	SONEY
FLESH	STICK	FRONT	BLESH	PLICK	DRONT
FRAME	SOUTH	GROUP	TRAME	POUTH	BROUP

Table 5

Performance of Lesioned MORSEL Word/Pseudoword Experiment			
Letters Attended	Relative Likelihood of Attentional State	Correct Responses Given Attentional State	
		Words	Pseudowords
1 2 3 4 5	8%	100%	81%
2 3 4 5	21	79	0
3 4 5	35	19	0
4 5	32	19	0
5	3	21	0

comparison, the neglect dyslexia patient HR studied by Behrmann et al. (1989) correctly reported 66% of words and 5% of pseudowords for stimuli of four to six letters.[7]

To summarize the implications of the current simulation, MORSEL provides a mechanism by which lexical or semantic knowledge can help compensate for noisy sensory data. This results in differential performance for words versus pseudowords because pseudowords do not benefit from such knowledge. MORSEL's account does not require the assumption that words and nonwords are processed along separate channels, or that the processing of words somehow bypasses the attentional system. In MORSEL, the attentional system and the recognition system operate identically for words and nonwords. Ultimately, however, words are less affected by the distribution of attention because of the compensating action of the PO net.

Distinctions in performance within the class of words

Studies examining the lexical status of a letter string have shown a difference in accuracy between words and nonwords, but recent work has found a more subtle influence of psycholinguistic variables on performance. Behrmann et al. (1989) compared performance on words that have a morpheme embedded on the right side — for example, PEANUT, which contains the morpheme NUT, and TRIANGLE, which contains ANGLE — and words having no right-embedded morphemes — for example, PARISH and TRIBUNAL. Although the patient studied by Behrmann et al. showed no difference in accuracy for the two stimulus types, a distinction was found in the nature of the errors produced. The upper portion of Table 6 summarizes the responses of the patient for words that contain right-embedded morphemes (hereafter, *REM words*) and words that do not (*control words*). Words were presented in two positions, either immediately to the right of fixation (the *near condition*) or several letter spaces further to the right (the *far condition*). Responses were classified into three categories: *correct responses, neglect errors* in which the right morpheme or its syllable control is reported — NUT for PEANUT or ISH for PARISH, and all *other errors*. The other errors consist mainly of responses in which the rightmost letters have been reported correctly but alternative

---

[7] HR's data is used for comparison to MORSEL in all simulations. We took little effort to obtain quantitative fits to HR's data for three reasons. First, the data we report is self-contradictory: HR performs quite well in one experiment but then poorly with similar stimulus materials in another. This is because the experiments were conducted sometimes weeks apart, and therefore reflect different stages of recovery of the patient and different overall levels of arousal and motivation. Second, the parameter values used to fit the data of one patient at a particular stage of recovery can hardly be expected to apply to other patients with somewhat different brain lesions. Third, given the number of free parameters of the model — that is, parameters not required in earlier work on MORSEL (e.g., connections involving semlex units, the nature of the attentional deficit) — relative to the small number of data points in this and subsequent simulations, a precise fit should not be considered terribly impressive. The important fact about parameter settings is that the qualitative behaviour of the model is remarkably insensitive to the specific parameter values.

Table 6

Distribution of Responses on Word Reading Task

Neglect Dyslexia Patient (from Behrmann et al., 1989)

Response Type	Near Condition		Far Condition	
	REM Words (e.g., PEANUT)	Control Words (e.g.,PARISH)	REM Words (e.g.,PEANUT)	Control Words (e.g.,PARISH)
Correct Response	43%	40%	79%	76%
Neglect Error	39	4	13	4
Other Error	18	56	9	20

Simulation of Lesioned MORSEL

Response Type	Near Condition		Far Condition	
	REM Words			
	(e.g., PEANUT)	(e.g.,PARISH)	(e.g.,PEANUT)	(e.g.,PARISH)
Correct Response	39%	44%	75%	76%
Neglect Error	32	0	9	0
Other Error	29	56	16	24

letters have been substituted on the left to form an English word — for example, IRISH or POLISH for PARISH (these errors have been termed *backward completions*). In both near and far conditions, overall accuracy is comparable for REM and control words, but neglect errors are the predominant error response for REM words and backward completions for control words. Sieroff, Pollatsek, and Posner (1988) have also studied compound words and found no significant difference in overall accuracy between compound and noncompound words. However, they provide no information about the distribution of error responses.

Our simulation study used twelve compound words and twelve control words — half six letters and half seven — from the stimulus set of Behrmann et al. (Table 7). As in our earlier simulations, the PO net simulation was conducted for each attentional state to obtain the probability of correct and neglect responses for REM and control words conditional upon the attentional state. These conditional probabilities were then combined with the probability of being in each attentional state (measured separately for six and seven letter words) to generate the distribution of responses shown in the lower portion of Table 6. Comparing the upper and lower portions of the Table, it is evident that the model produces the same pattern of results as the patient. The difference in accuracy between near and far conditions confirms the the previous finding concerning of the effect of retinal presentation position. Overall accuracy is about the same for REM and control words. Neglect errors are frequent for REM words, whereas backward completion errors (the primary error type in the "other error" category for the simulation as well as the patient) are most common for control words.

The difference in performance for the two word classes is explained by the action of the semlex units. These units support neglect responses for REM words but not control words. The same effect was responsible for the basic word advantage in the word/pseudoword simulations. However, in the present simulation, the influence of semlex units acts not to increase the accuracy of report for one stimulus type but to bias the network towards one type of error response over another when the perceptual data is not strong enough to allow the PO net to reconstruct

Table 7

Stimuli for Embedded-Morpheme Experiment			
REM Words		Control Words	
Six Letter	Seven Letter	Six Letter	Seven Letter
PEANUT	INKWELL	PARISH	DARLING
SUNSET	SKILIFT	BEGGAR	PROMISE
COWBOY	NEWBORN	FOSTER	CUSTARD
SUNTAN	EARRING	SILVER	TORMENT
OFFSET	NETWORK	MORSEL	GESTURE
SUNDAY	BEDROOM	SHADOW	COMPLEX

the target.

The only discrepancies between the patient and simulation data in Table 6 is that the model produces about a 5% lower neglect error rate uniformly across all conditions and a slight accuracy advantage for control words. The accuracy advantage for control words can be eliminated by adjusting parameters of model, but in truth, the advantage is present for most parameter settings. It is not difficult to see why this is so. Consider the behaviour of the PO net when the AM has selected the last three letters of either PEANUT or PARISH. With PEANUT, the predominant response of the PO net is NUT because the clusters of NUT receive strong support from the semlex units. With PARISH, however, the semlex units do not support ISH but instead favor PARISH or one of the alternative backward completions. If the number of backward completions is relatively small, PARISH is more likely to be read correctly than PEANUT. Such behaviour is at variance with the patient data.

We have two escapes from this dilemma. First, our implementation of the PO net utilizes only a limited number of alternative responses for a given stimulus. This was necessary to make simulations computationally feasible, yet by cutting down on the number of alternative responses, it raises the likelihood of the PO net producing the correct response simply by guessing. Such guessing behaviour occurs when the combination of perceptual data and semlex biases do not strongly agree on a candidate response — the case of PARISH when only ISH is attended. Second, if the semlex units were true semantic units, that is, if they represented semantic features, then for most compound words, the units representing the whole word — PEANUT — would overlap with those representing the semantically-related right-sided morpheme — NUT. No such overlap exists with the current representation. The advantage of overlap is that support for NUT will also support PEANUT, thereby increasing the likelihood of a correct response for compound words.[8]

The influence of lexical status on extinction

The last two sections presented experimental results that were explained by MORSEL in terms of an interaction between attentional selection and higher-order stimulus properties. However, the tie to attentional selection is somewhat indirect because the stimuli were single words or pseudowords and attention is generally thought of as selecting between two competing items, not selecting between portions of a single item.

Using the extinction paradigm, Behrmann et al. (1989) have been able to show that the ability of a neglect dyslexia patient to select the leftmost of two words is indeed influenced by the relation between the words. When the patient was shown pairs of semantically unrelated three-letter words separated by a space, e.g., SUN and FLY, and was asked to read both words, the left word was reported on only 12% of trials; when the two words could be joined to form a compound word, e.g., COW and BOY, the left word was read on 28% of trials. (On all trials where the left word was reported, the right word was also reported.) Thus, it would seem that the operation of attention to select among stimuli interacts with higher-order stimulus properties.

One natural interpretation of this interaction is that the attentional system is directly influenced by semantic or lexical knowledge, as proposed by late-selection theories of attention. MORSEL provides an alternative account in which attention operates at an early stage, but because unattended information is partially processed, later stages can alter the material selected. Thus, one need not posit a direct influence of higher-order knowledge on attentional selection to obtain behaviour in which the two interact.

To describe how MORSEL can account for the interaction, we begin with a description of the lesioned model's behaviour and then turn to simulation results. When two items are presented to the lesioned AM usually the right word is selected (Table 1). Consequently, BLIRNET strongly activates the clusters of BOY when COW BOY is presented, partially activates the clusters of COW and, because BLIRNET has some difficulty keeping track of the precise ordering of letters, weakly activates clusters representing a slight rearrangement of the stimulus letters, OWB and WB_Y. These latter clusters support the word COWBOY. The overall pattern of letter cluster activity is thus consistent with COWBOY as well as BOY. Because both words receive support from the semlex units, the PO net can

---

[8] This will not necessarily lead to a corresponding decrease in neglect responses. Interactions in the PO net are simply too difficult to predict without simulations.

potentially read out either; thus, in the case of COWBOY, the left morpheme is read out along with the right. When the two morphemes cannot be combined to form a word, however, the semlex units do not support the joined-morpheme response, and the PO net is unlikely to read the two morphemes out together.

There is another avenue by which the left morpheme may be read out: the patient may be able to shift attention to the left and reprocess the display. In the experiment of Behrmann et al., this seems a likely possibility because all trials contained two words and the patient's task was to report the entire display contents. Although the patient was not explicitly told that two words were present, the observation of both words on even a few trials may have provided sufficient incentive to try reporting more than one word per trial. The patient may therefore have had a top-down control strategy to shift attention leftward. MORSEL is likewise able to refocus attention to the left on some trials using top-down control of the sort illustrated in Figure 6. This will cause an increase in reports of the left morpheme both for related and unrelated stimulus pairs.

Table 8 lists the 12 word pairs used in the stimulation. *Related morphemes* can be joined to form a compound word; *unrelated morphemes* do not combine in this manner. As with previous simulations, each stimulus was presented 100 times in each attentional state to obtain a probability of reporting the joined morpheme (e.g., COWBOY or SUNFLY) conditional upon the attentional state. These conditional probabilities were then combined with the probability of being in a given attentional state (Table 1) to obtain overall response rates. The left morpheme was reported on 14.1% of trials for related morphemes but only 2.8% for unrelated morphemes. Thus, the strength of lexical/semantic knowledge is sufficient to recover the extinguished information on the left for two morphemes that can be combined to form a word. Figure 11 shows a graph of activity over time for COW BOY on a trial where the AM has focused attention only on BOY. Nonetheless, the semlex units of COWBOY and the partial activations from the left morpheme converge to eventually cause the PO net to read out COWBOY.

Assuming that top-down control of the AM allows MORSEL to shift attention to the left and reprocess the display on, say, 10% of trials, the total percent of trials in which the left morpheme is reported rises to 24.1% for related morphemes and 12.8% for unrelated morphemes. These results are in line with the patient data obtained by Behrmann et al. — 28% and 12%.

Interestingly, on trials in which just the right morpheme is reported, MORSEL occasionally produces left neglect errors, for example, reporting ROY for BOY. Behrmann et al.'s patient produced similar errors. Thus, both left-item extinction and left-sided neglect can be observed on a single trial.

Examining the behaviour of MORSEL, it is apparent that two properties of the stimuli are necessary to obtain a reduction in extinction with related morphemes. These two properties have not yet been tested on neurological patients, but are subtle predictions of MORSEL. First, the physical separation between the morphemes is important: the further apart the two morphemes are, the less activation BLIRNET will produce for the internal clusters of the joined morpheme — e.g., OWB and W_OY of COWBOY. This will reduce the likelihood of the PO net reading out COWBOY. Second, the particular effect we have simulated depends not on a semantic relatedness of the two morphemes, but on the fact that the two morphemes can be joined to form a lexical item. Semantic relatedness alone may allow for a reduction in extinction, but it would not be by exactly the same mechanism.[9] Experimental work is currently underway using the Sieroff and Posner (1988) cuing paradigm to simulate neglect in normals and examine

Table 8

Stimuli for Extinction Simulation			
Unrelated Morphemes		Related Morphemes	
COW NUT	PEA SET	COW BOY	PEA NUT
OFF TAN	SUN FLY	OFF SET	SUN TAN
BAR DAY	SUN BOY	BAR FLY	SUN DAY

---

[9] If the two morphemes are semantically related but do not combine to form a compound word, e.g., BOY and MAN, one mechanism whereby one morpheme could affect the read out of the other morpheme involves priming of the semantic units. That is, activation of the semantic units of MAN will support the related word BOY to some extent. This account requires an elaboration of temporal processing in MORSEL which has not been necessary in the present work.
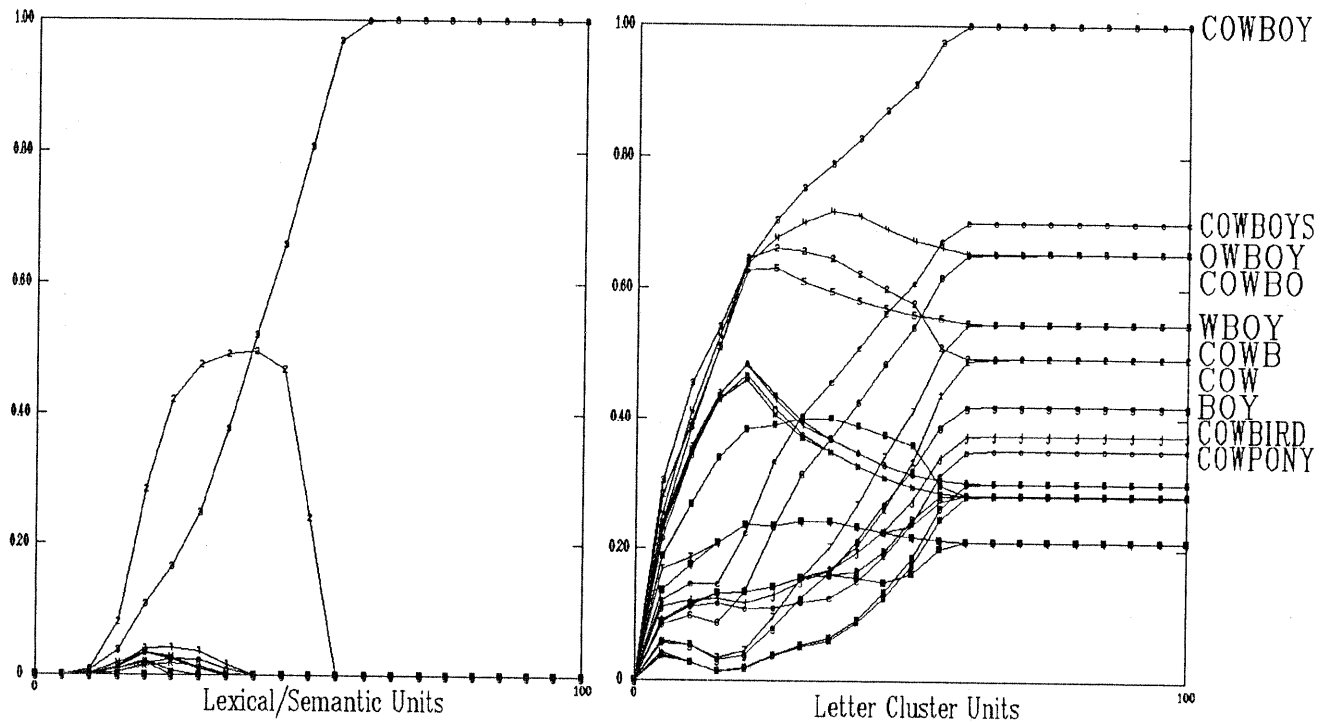
Figure 11. Activity in the PO net as a function of time for the stimulus COW BOY on a trial where the AM has focused attention only on BOY. The PO net eventually reads out COWBOY.

these two predictions of MORSEL.

Discussion

MORSEL was originally developed to explain a variety of psychological phenomena observed in normal subjects. We have demonstrated that damaging the model leads to behaviours observed in neurological patients with neglect dyslexia. A single lesion — to the connections that help draw attention to objects in the visual field — is sufficient to account for a remarkable range of behaviours, some of which are compatible with a deficit at an early stage of processing and others of which might naturally be interpreted as arising at later stages of the system. The fact that MORSEL, in its lesioned state, performs in a manner comparable to neglect dyslexia patients is a further, compelling validation of the model.

Simulations of neuropsychological phenomena, especially acquired dyslexia, have become increasingly popular of late and have generally been successful in modeling pathological performance. Although computational modeling could be feasibly applied to many domains of human performance, it is particularly suitable to the study of acquired dyslexia since research in cognitive neuropsychology has been prolific in recent years and has provided a firm empirical and theoretical database.

The predominant theoretical paradigm in cognitive neuropsychology has been to utilize models of normal cognitive processing for analyzing the locus of the functional lesion in subjects with impaired performance. These models typically consist of box-and-arrow flow diagrams with the underlying assumption that discrete and selective damage may affect a single subsystem without influencing the functioning of other components. Information derived from experiments with brain-damaged subjects is then used to guide and constrain the development of models of normal cognition. According to Seidenberg (1988), models of this sort are limited because they do not incorporate specific proposals about knowledge representation or processing mechanisms. These types of models represent a descriptive, first-order decomposition of tasks such as reading and spelling and thus tend to serve as

recharacterizations of empirical data. Computational models, in which explicit assumptions about processing are made, provide an alternative, more constructive paradigm for examining normal cognition and its breakdown. Examples of this type of research are starting to infiltrate the literature. Simulations of deep dyslexia, with emphasis on mixed semantic and visual errors and category specific preservation (Hinton & Shallice, 1989), and surface dyslexia (Patterson, Seidenberg, & McClelland, 1989) have yielded interesting, counterintuitive results that challenge the more traditional box-and-arrow models.

The common finding associated with these studies, as well as our own, is that complex interactions between the processing subcomponents are observed. These nontransparent interactions are often difficult to account for in the context of box-and-arrow flow diagrams but are more easily explicable in the dynamic framework provided by a working computational model. While MORSEL is made up of a set of discrete and relatively simple components (not too dissimilar from the box-and-arrow, Fodorian modules), damage at one point may have ramifications for the rest of the system. Thus, analyzing each component in isolation provides a restricted view of the overall system. Analyzing the operation of the system in its entirety is far more informative since the net effect of a lesion on behaviour is complicated by interactions among the components.

This wholistic perspective provides us with a solid framework within which to examine the mechanisms underlying neglect dyslexia. Previous neuropsychological studies of the phenomenon have advanced disparate explanations of the deficit. For example, the fact that stimulus position, orientation, and physical features are important determinants of performance has been taken as support for the fact that the attentional deficit arises at peripheral stages of processing (Behrmann et al., 1989; Ellis et al., 1987). This interpretation will not suffice, however, since it cannot explain why lexical and morphemic effects also play an important role in performance. A second group of explanations have been proposed to account for the superiority of words over nonwords and for the role of morphemic composition in reading. One theory falling into this latter group postulates that reading of words is automatic and attention-free and is thus immune from attentional deficits (LaBerge & Samuels, 1974; Sieroff, Pollatsek, & Posner, 1988). Nonwords, however, are subject to such deficits since they necessarily require attention. Clearly, these two types of explanations draw on entirely different theoretical views with little in common.

In the framework of MORSEL, interactions between early and later stages account for these seemingly discrepant views. Although the primary deficit indeed arises at a low-level of processing, higher-order knowledge may compensate for the the peripheral dysfunction. Interactions between attention and higher-order information may then be observed and interpretations that previously appeared contradictory may be brought into alignment. In this way, we have advanced a potential explanation for what appeared to be mutually exclusive phenomena.

MORSEL's framework not only helps to explain the mechanisms of neglect dyslexia, but it also suggests a novel conceptualization of selective attention. The theory of attention embodied in MORSEL cannot be characterized as either early or late selection; MORSEL shows properties of both. The AM performs a type of early selection by location or simple physical attributes (which we did not describe because it was not relevant to the present work). Despite its name, the AM does not have exclusive title over attentional selection; the PO net also serves as a distinct selectional mechanism — one that, for example, prefers words over nonwords. In fact, PO net units can be primed by external inputs in exactly the same manner as the AM to bias PO net selection towards items with particular orthographic, lexical, or semantic properties. Thus, attention is not a single, localizable component of MORSEL. Selection operates at two distinct stages, and, as we have demonstrated, these stages can interact with one another in rather unintuitive ways. We eagerly await future research in cognitive neuroscience and neuropsychology that examines further the notion of attentional selection occurring at multiple levels in the human information processing system.

Appendix 1: PO net dynamics

The task of the PO net is to select a set of letter cluster units that can be assembled to form a unique letter string and are consistent with the activations produced by BLIRNET. If the letter string read out by the PO net is an English word, the PO net must also select a set of semlex units that represent the semantic/lexical entry.

Letter-cluster unit connections

Two letter clusters are said to be *neighbors* if they can be aligned so as to overlap on two letters or delimiters ("*"). Some examples of neighbors are: MON and ONE (overlap on O and N), **M and *MO (* and M), M_NE and V_NE (N and E), and E_** and F_** (* and *). Two neighbors are said to be *compatible* if, when aligned, they do not conflict in any letter position. The first two examples above are compatible, the second two incompatible.

Based on this classification of compatible and incompatible neighbors, four types of connections between letter-cluster units are warranted: (1) *excitatory* — between compatible neighbors; (2) *inhibitory* — between incompatible neighbors; (3) *-excitatory* — a special case of an excitatory connection where both letter clusters contain delimiters and the presence of one cluster necessitates the presence of another, e.g., *MO implies **M and **_O. Note that these connections are not symmetric: neither **M nor **_O alone implies *MO. (4) *-inhibitory* — a special case of an inhibitory connection where both letter clusters contain delimiters. For these pairs, the presence of one cluster precludes the presence of the other, e.g., *MO and **_A, Y** and E**. These connections are symmetric.

Each connection type has associated with it a different weight. The excitatory connections have positive weights, inhibitory negative. The *-connections have weights of a greater magnitude. The values used in our simulations, as well as other weight parameters of the PO net described below, are listed in Table 9.

Semlex unit connections

As stated in the text, the semlex representation is intended to be a distributed encoding of word meanings. Because of the difficulty in devising a complete distributed semantic representation, the PO net instead uses a semi-distributed representation in which each word meaning is associated with a distinct pool of units. These units are not shared by different words. In our simulations, this is effectively equivalent to a *lexical* representation because simulations involved few if any synonyms.

The number of semlex units associated with each word in MORSEL's lexicon was twice the number of letters in the word. Each of these units was connected to five randomly selected letter clusters of the word, with the restriction that all letter clusters had approximately the same number of semlex connections. Because the number of letter clusters in an l-letter word is $3l+2$ and the total number of semlex-letter cluster connections is $10l$, each letter cluster unit of a word is on average connected to slightly over three of the word's semlex units. This particular scheme was selected because, unlike other schemes we considered, it made the PO net fairly neutral with regard to word length; there was no bias towards either shorter or longer words.

Table 9

PO net Connection Strengths	
Connection Type	Value
excitatory	.06
inhibitory	−.18
*-excitatory	.24
*-inhibitory	−.24
letter cluster to semlex	.10
semlex to letter cluster excitatory	.10
semlex to letter cluster inhibitory	−.001
semlex to semlex inhibitory	−.05
feedforward (ω_F)	.0005
global suppression (ω_G)	−.14

The connections between letter cluster and semlex units are symmetric and excitatory. In addition, each semlex unit slightly inhibits all letter cluster units to which it is not connected. Semlex units also inhibit all semlex units that are associated with different words. It is this inhibition that forces the PO net to select a pattern of activity in the semlex units corresponding to a single word. (See Table 9 for values of these parameters.)

PO net activation function

Initially, the PO net receives feedforward excitation from the letter cluster units of BLIRNET. Interactions then take place within the PO net and it gradually iterates towards a stable state. PO units were given the same dynamical properties as units in McClelland and Rumelhart's (1981) interactive-activation model. Units are continuous-valued in the range [-0.2,1.0]. Information coming in to each unit is summed algebraically, weighted by the connection strengths, to yield a "net input":

$$net_i = \sum_{\substack{j \in \\ ACTIVE}} w_{ij}p_j + \omega_F b_i + \omega_G \bar{p} ,$$

where ACTIVE is the set of all PO units with positive activity at the current time, w_{ij} is the strength of connection to PO unit i from PO unit j, p_j is the activity of PO unit j, b_i is the activity of letter-cluster unit i of BLIRNET (if i is a semlex unit, then b_i is zero), and ω_F is the strength of feedforward connections from BLIRNET to the PO net. The final term, $\omega_G \bar{p}$, applies only to the letter-cluster units and is explained below.

The activation value of each PO unit is updated by the net input according to the rule:

$$\Delta p_i = \begin{cases} net_i [1.0 - p_i] & \text{if } net_i > 0 \\ net_i [p_i - (-0.2)] & \text{otherwise.} \end{cases}$$

If the net input is positive, activation is pushed towards the maximum value of 1.0; if negative, activation is pushed towards the minimum value of -0.2. The effect of the net input is scaled down as the unit approaches its maximum or minimum activation level.

The network as described thus far is inadequate. The problem is as follows. Many letter clusters compete and cooperate directly with one another, in particular, the clusters representing ends of words and the clusters sharing letters. Often, however, these interactions are not enough. For instance, suppose two words are presented, LINE and FACT, and that clusters of LINE are more active initially. Clusters like **F and CT* of FACT experience direct competition from the corresponding clusters of LINE, and are therefore suppressed, but the inner clusters of FACT such as FAC and F_CT do not. The pull out process thus yields LINE along with the inner clusters of FACT. To get around this problem, some type of "global inhibition" is useful.

The mechanism we opted for inhibits each letter-cluster unit in proportion to the average activity of all clusters above threshold, which can be computed as follows:

$$\bar{p} = \frac{1}{|ACTIVE_L|} \sum_{\substack{i \in \\ ACTIVE_L}} p_i ,$$

where $ACTIVE_L$ is the set of all letter cluster units with positive activity at the current time. The equation for net_i incorporates this term, weighted by the parameter ω_G. This scheme allows the set of letter cluster units whose activity grows the fastest to shut off the other units. Activity grows fastest for units that have many active compatible neighbors.

Appendix 2: AM dynamics

The goal of the AM is to construct a "spotlight" of activity that highlights a single item appearing on MORSEL's retina. Defining an item to be a set of features in close proximity, the spotlight should form *contiguous* region on the retina consistent with the bottom-up and top-down inputs to the AM.

In connectionism, the standard method of transforming this description of the target behaviour of the AM into a network architecture is to view the AM's task as an *optimization* problem: to what activity value should each unit in the AM be set in order to best satisfy a number of possibly conflicting constraints? The two primary constraints here are that the AM should focus on locations suggested by the bottom-up and top-down inputs, and the AM should focus on a single item.

The first step in tackling such an optimization problem is to define a *Harmony* function (Smolensky, 1986) that computes the goodness of a given pattern of activity over the entire AM (the AM *state*). This goodness is a scalar quantity indicating how well the AM state satisfies the optimization problem. The maxima of the Harmony function correspond to desired states of the AM.

Given a Harmony function, H, one can ask how the activity of the AM unit at a retinal location (x,y), denoted a_{xy}, should be updated over time to increase Harmony and eventually reach states of maximal Harmony. The simplest rule, called *gradient ascent*, is to update a_{xy} in proportion to the derivative $\partial H / \partial a_{xy}$. If $\partial H / \partial a_{xy}$ is positive, then increasing a_{xy} will increase H; thus a_{xy} should be increased. If $\partial H / \partial a_{xy}$ is negative, then decreasing a_{xy} will increase H; thus a_{xy} should be decreased.

Returning to the problem faced by the AM, devising a Harmony function that computes whether the pattern of activity is contiguous is quite difficult. Instead of constructing a function that explicitly rewards contiguity, we have combined several heuristics that together generally achieve convex, contiguous patterns of activity.[10] The Harmony function we use is:

$$H = \sum_{\substack{(x,y) \\ \in ALL}} ext_{xy} a_{xy} - \frac{\mu}{2} \sum_{\substack{(x,y) \\ \in ALL}} \sum_{\substack{(i,j) \in \\ NEIGH_{xy}}} (a_{ij} - a_{xy})^2 + \frac{\theta}{2} \sum_{\substack{(x,y) \in \\ ACTIVE}} (\gamma \bar{a} - a_{xy})^2 \, ,$$

where *ALL* is the set of all retinal locations, ext_{xy} is the net external (bottom up and top down) activity to the AM at location (x,y), $NEIGH_{xy}$ is the set of 8 locations immediately adjacent to (x,y) — the *neighbors*, *ACTIVE* is the set of locations of all units with positive activity, $\bar{a}$ is the mean activity of all units with positive activity —

$$\bar{a} = \frac{1}{|ACTIVE|} \sum_{\substack{(x,y) \in \\ ACTIVE}} a_{xy} \, ,$$

and μ, θ, and γ are weighting parameters.

The first term encourages each unit to be consistent with the external bias. The second term encourages each unit to be as close as possible to its neighbors (so that if a unit is off and the neighbors are on, the unit will tend to turn on, and vice versa). The third term encourages units below the mean activity in the network to shut off, and units above the mean activity to turn on. The constant γ serves as a discounting factor: with γ less than 1, units need not be quite as active as the mean in order to be supported. Instead of using the average activity over *all* units, it is necessary to compute the average over the *active* units. Otherwise, the effect of the third term is to limit the total activity in the network, i.e., the number of units that can turn on at once. This is not suitable because we wish to allow large or small spotlights depending on the external input. (The same type of scheme was used to limit activity in the PO net, as described in Appendix 1.)

---

[10] We should note that many other Harmony functions would suffice equally well if not better than the one we devised. In fact, we experimented with several different functions, and the qualitative system behaviour was unaffected by the details of the Harmony function.

The update rule for a_{xy} is:

$$\Delta a_{xy} = \frac{\partial H}{\partial a_{xy}} = ext_{xy} + \mu \sum_{\substack{(i,j) \in \\ NEIGH_{xy}}} (a_{ij} - a_{xy}) - \theta (\gamma \bar{a} - a_{xy}).$$

Further, a_{xy} is prevented from going outside the range [0,1] by capping activity at these limits. The activation function is essentially the same as described in Mozer (1988b), changing $\bar{a}$ to represent average activity of only the active units.[11]

To explain the activation function intuitively, consider the time course of activation. Initially, the activity of all AM units is reset to zero. Activation then feeds into each unit in proportion to its external bias (first term in the activation function). Units with active neighbors will grow the fastest because of neighborhood support (second term). As activity progresses, high-support neighborhoods will have activity above the mean; they will therefore be pushed even higher, while low-support neighborhoods will experience the opposite tendency (third term).

In all simulations, μ was fixed at 1 and, with informal experimentation, θ was chosen to be .5 and γ to be .11 times the total external input.

---

[11] To follow the objective function exactly, the third term should actually be zero if a_{xy} is currently inactive. However, including this term at all times prevents oscillation in the network and does not otherwise appear to affect the quality of the solution.

Appendix 3: Details of AM simulations

In this Appendix, we describe the stimuli used as input to the AM and the simulation methodology.

Input assumptions

In the font we have designed, letters presented on MORSEL's retina each occupy a 3×3 region of the input map. Letters within a word are presented in horizontally adjacent positions. Thus, a three-letter word subtends a 3×9 retinal region. Two three-letter words, with a single space between them (a 3×3 gap), subtend a 3×21 region.

The featural activations arising on MORSEL's retina at a given location serve as the bottom-up input to the corresponding location of the AM. The input is thus nonzero only at locations where letters are present. To simplify our simulations, rather than presenting real words on MORSEL's retina and using the resulting featural activations as input to the AM, we assumed, for a stimulus string occupying a given retinal region, a uniform distribution of input within that region — an external input of .01 at each location. We assumed an additional input of .01 along the outer border of the region, representing an input from a boundary contour system (e.g., Grossberg & Mingolla, 1985).[12] Finally, we assumed a bit of blurring: Each retinal activation provided not only bottom-up input to the corresponding location in the AM but also to the horizontally, vertically, and diagonally neighboring locations. This activation strength was only .0002, much smaller than the direct input.

With the input as described, two three-letter words presented simultaneously produce exactly the same pattern of bottom-up input. Without some degree of randomness, the AM has no means of breaking symmetry and selecting one word or the other. Thus, for simulations of the normal model, we assumed that each bottom-up input is transmitted to the AM with only 90% probability. This causes the strength of a word to vary from one trial to the next.

The basic claim of MORSEL is that neglect dyslexia results from graded damage to the bottom-up AM inputs, most severe on the left and least on the right. One way of expressing this damage is in terms of the probability of transmitting an input to the AM. Rather than a uniform probability close to 1, we assumed in the damaged model that the probability varies with lateral retinal position: at the left edge of the retina, the probability was 48% and increased by 2% for each successive pixel location to the right, with a maximum of 90%. Thus, words presented in the "standard" position (starting 6 pixels from the left end of the retina; this was the position used in most simulations) had a transmission probability of 60% for their left edge, and the probability reached the 90% ceiling by the sixth letter position (21 pixels from the left of the retina).

Simulation methodology

Simulation experiments were conducted for two simultaneous three-letter words and single 5-7 letter words, presented in the standard position or shifted one or two letter positions to the right. For each simulation, 1000 replications were run. On each replication, every source of bottom-up input was considered independently and was fed to the AM in accordance with the probabilistic transmission function. Thus, on each replication the AM detected a slightly different subset of the inputs.

The AM was then allowed to run until equilibrium was reached, that is, until all units settled on stable activation values. The total attention to each letter position was then measured by averaging the activities of the nine AM units in the region corresponding to a given letter. If this average activity was greater than .5, the letter was considered to have been attended. The attentional state for the stimulus was then determined by combining the individual letter results. By the .5 activity criterion, there were occasional responses that didn't fit into one of the expected attentional states, for example, attending to positions 2 and 4-6 of a six-letter word but not position 3. We placed such responses into the closest reasonable category; here, the state of attending to positions 2-6.

---

[12] Such an input seems of critical importance in determining the focus of attention. Attention should turn to to changes in the visual environment, not homogeneous regions.

Appendix 4: Details of BLIRNET and PO net simulations

About six thousand letter clusters are required to represent the most common words of English. However, running a simulation with this number of clusters is computationally infeasible. If each cluster is connected to, say, two hundred other clusters, the total number of connections will exceed 1.2 million, and this estimate completely ignores the cost of the semlex units, which is major factor if the simulation includes many lexical items. Constructing a full-scale PO net is wasteful, too: for a given stimulus, most of the units will not come into play in determining the PO net's response. Thus, rather than constructing one gigantic PO net to handle all simulations, we constructed a specialized PO net for each stimulus item. This smaller net contained only the letter-cluster and semlex units that seemed relevant for the particular stimulus.

In this Appendix, we describe the procedure used to select letter-cluster and semlex units for inclusion in the PO net simulation, the rules used for determining the BLIRNET activation levels of these units, and finally, the PO net simulation methodology.

Selection of alternative responses

For each stimulus, we generated a set of *alternative responses* — strings that had enough in common with the stimulus to be plausible responses. For the single word stimuli, the alternative responses included:

- the stimulus word itself (e.g., PARISH);
- all right segments and left segments of the stimulus with three or more letters (e.g., ARISH, PARIS, RISH, PARI, ISH, PAR);
- all words in the Kučera and Francis (1967) corpus ending with the last three letters of the stimulus and having the same length as the stimulus, plus or minus one letter (e.g., ENGLISH, BRITISH, JEWISH, FINISH, SPANISH, FURNISH, IRISH, POLISH, FOOLISH, TURKISH, ABOLISH, ANGUISH, DANISH, RADISH, RELISH). If more than fifteen such words existed, the fifteen with the highest word frequency counts were selected.[13]
- all words in the Kučera and Francis corpus beginning with the first three letters of the stimulus and having the same length as the stimulus, plus or minus one letter (e.g., PARTY, PARTS, PARENTS, PARKER, PARTIES, PARTLY, PARKED, PARTNER, PARKING, PARADE, PARKS, PARLOR, PARENT, PARTIAL). If more than fifteen such words existed, the fifteen with the highest word frequency counts were selected.
- six pseudowords having the same final three letters and overall length as the stimulus (e.g., SUNISH, COWISH, PEAISH, OFFISH, INKISH, EARISH).

For the two-word stimuli used in the extinction experiment, the alternative responses were determined by combining the two words into a single string (e.g., SUN and FLY to SUNFLY) and using the above criteria, in addition to:

- the individual three-letter stimuli (e.g., SUN, FLY);
- all three-letter words in the Kučera and Francis corpus ending with the last two letters of either stimulus (e.g., RUN, GUN, FUN, HUN, NUN, BUN, PUN, SLY, PLY);
- all three-letter words beginning with the first two letters of either stimulus (e.g., SUM, SUE, SUB, FLA, FLU).

The net constructed for a given stimulus included all letter-cluster units composing each of the alternative responses as well as a set of semlex units for each alternative response that was an English word. This allowed the PO net to potentially read out any of the alternative responses. Table 10 presents the average number of alternative responses generated for stimuli in each of the four simulation experiments we conducted, as well as the average number of units and connections contained in the PO net.

---

[13] Note that by including only alternative responses that had approximately the same length as the stimulus, we artificially limited the model to responses that preserve stimulus word length. Neglect dyslexia patients do in fact show a preservation of word length (Behrman et al., 1989; Ellis et al., 1987), but clearly not because all the words they know are of the same length as the stimulus. We believe that a fuller implementation of MORSEL should include a processing module similar to BLIRNET that computes word shape information instead of word identity information. The word shape and identity information could then be integrated by the PO net to select responses that were consistent with both, thereby allowing a preservation of word length even in neglect dyslexia patients.

Table 10

PO Net Statistics						
Experiment	Number of alternative responses			Avg. number of units		Avg. number of connections
	Average	Minimum	Maximum	Letter cluster	Semlex	
Retinal position	35	29	42	324	281	7545
Word/pseudoword	28	20	35	216	245	4958
Right-embedded morpheme (REM)	32	22	42	303	254	6779
Related/unrelated morpheme	50	38	62	393	412	13637

Rules for determining BLIRNET activations

Once the set of letter cluster units has been selected for a given stimulus, the BLIRNET activation of each unit must be determined. As we explained in the main text, we did not actually simulate BLIRNET. Instead, we used a simple algorithm to obtain activations similar to what BLIRNET would have produced had a full-scale simulation been conducted. Given an input stimulus and a focus of attention produced by the AM, this algorithm worked as follows for a particular letter cluster.

The letter cluster is compared to every subsequence of the stimulus by aligning the cluster in every possible way with the stimulus. For a given alignment, each of the three characters of the cluster (i.e., letters or delimiters — the "don't care" underscores were ignored) is matched against the corresponding character of the stimulus. If the cluster character is a letter, the match score is $\alpha\rho$; if the character is a delimiter, the match score is $1-\alpha(1-\rho)$. α is the level of attention to the corresponding character of the stimulus; α is 1 if the character is attended or .368 if unattended. ρ lies in the range [0,1] and is a measure of the featural similarity of the stimulus and cluster characters. If the characters are identical — a perfect match, ρ is 1; to the extent that the characters are physically similar (as measured by the dot product of their feature vectors), ρ is greater than zero. Thus, the physical appearance of the letters comes into play in determining BLIRNET activations.

To summarize, the match will be close to 1 for a letter of a cluster if the corresponding stimulus letter is attended and is physically similar to the cluster's letter. For delimiters of the cluster, however, the match will be close to 1 *either* if the corresponding stimulus position contains a blank space (ρ is 1) or if the corresponding stimulus position is unattended (α is small). The reason for the second condition is that if the position is unattended, few features are transmitted through BLIRNET; consequently, it will appear as if the position is blank.

A cluster character is not only matched against the corresponding stimulus character but also against the left and right neighbors of the stimulus character. The scores obtained for the neighbor matches are multiplied by .5, and the largest of these two scores and the original match score is selected as the overall character match. The reason for matching neighbors is that BLIRNET confuses exact letter positions and often produces partial activations of clusters with letters in a slightly incorrect order.

The geometric mean of the overall character matches is computed to obtain an *overall cluster match*; i.e., $(m_1 m_2 m_3)^{1/3}$, where m_i is the match for character i. Character matches involving a delimiter in the cluster and a blank space in the stimulus are ignored in computing overall cluster match. This overall cluster match is computed for each possible alignment of the cluster and the stimulus, and the activation level of the cluster is simply the sum of the overall cluster matches over all alignments.

What this procedure boils down to is simply that a cluster is assigned an activity level of 1.0 if the cluster is contained in the attended portion of the word; the cluster is assigned an activity level of .05 ($=.368^3$) if it contained in the unattended portion; the cluster is assigned an activity level intermediate between .05 and 1.0 if it crosses the boundary between the unattended and attended portions. If a cluster does not match the stimulus exactly because some letters are different or the letters are in a slightly different order, the cluster still attains some degree of activation.

To obtain different responses on each run, gaussian noise with mean zero and standard deviation .10 was added to the activity level of each cluster, and the activities were thresholded to lie in the range [0,1].

PO net simulation methodology

To obtain reliable simulation results, each stimulus was tested with alternative sets of semlex unit connections and random fluctuations in the BLIRNET activities. To elaborate, for each stimulus we reconstructed the PO net ten times, each time with the semlex units connected to a different random subset of their associated letter clusters. For each version of the network thus constructed, we allowed the net to settle ten times, each time starting with a different pattern of noise added to the BLIRNET activations. In total, then, every stimulus item was presented one hundred times.

A measure of the strength of a particular response was computed according to the formula:

$$strength = \frac{1}{2}\left[\frac{t}{n}+\frac{t}{T}\right] \ ,$$

where t is the summed activity of target clusters — those composing the response, n is the number of target clusters, and T is the summed activity all clusters. Cluster activities were thresholded to lie in the range [0,1]. The first term in the formula represents the average activity of the target clusters and approaches 1 as the clusters of the response increase in activity. The second term represents the activity of the target clusters relative to nontarget clusters. The strength ranges from 0 to 1 and reaches 1 only if all target clusters are fully active and no nontarget cluster is active.

On each run, the PO net was allowed to run until it reached equilibrium (usually within 50 processing iterations) and the response with the greatest strength was taken as MORSEL's selection. Generally, this response had strength 1. It was necessary to use only the letter-cluster activity in determining MORSEL's selection; use of the semlex units would have precluded nonword responses.

This simulation procedure was carried out for each stimulus and each attentional state. The individual stimulus results were then averaged to produce a distribution of responses conditional upon a particular attentional state. These conditional probabilities could then be combined with the relative probabilities of different attentional states to obtain an overall distribution of responses.

References

Barbut, D., & Gazzaniga, M. (1987). Disturbances in conceptual space involving language and speech. *Brain, 110*, 1487-1496.

Baxter, D., & Warrington, E. K. (1983). Neglect dysgraphia. *Journal of Neurology, Neurosurgery, & Psychiatry, 46*, 1073-1078.

Behrmann, M., Moscovitch, M., Black, S. E., & Mozer, M. C. (1989). *Perceptual and conceptual factors in neglect: Two contrasting case studies.* Manuscript submitted for publication.

Bisiach, E., Perani, D., Vallar, G., & Berti, A. (1986). Unilateral neglect: Personal and extra-personal. *Neuropsychologia, 24*, 759-767.

Bisiach, E., & Vallar, G. (1988). Hemineglect in humans. In F. Boller & J. Grafman (Eds.), *Handbook of neuropsychology, Volume 1* (pp. 195-222). North Holland, Amsterdam: Elsevier Science Publishers, BV.

Broadbent, D. E. (1958). *Perception and communication.* London: Pergamon.

Brunn, J. L., & Farah, M. J. (1989). *The relation between spatial attention and reading: Evidence from the neglect syndrome.* Manuscript submitted for publication.

Butter, C. M. (1987). Varieties of attention and disturbances of attention: A neuropsychological analysis. In M. Jeannerod (Ed.), *Neurophysiological and neuropsychological aspects of spatial neglect* (pp. 1-24). Amsterdam: North Holland.

Butter, C. M., Mark, V., & Heilman, K. M. (1989). An experimental analysis of factors underlying neglect in line bisection. *Journal of Neurology, Neurosurgery, & Psychiatry.* In Press

Caplan, B. (1987). Assessment of unilateral neglect: A new reading test. *Journal of Experimental and Clinical Neuropsychology, 9*, 359-364.

Carr, T. H., Davidson, B. J., & Hawkins, H. L. (1978). Perceptual flexibility in word recognition: Strategies affect orthographic computation but not lexical access. *Journal of Experimental Psychology: Human Perception and Performance, 4*, 674-690.

Deutsch, J. A., & Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review, 70*, 80-90.

Ellis, A. W., Flude, B., & Young, A. W. (1987). Neglect dyslexia and the early visual processing of letters in words and nonwords. *Cognitive Neuropsychology, 4*, 439-464.

Farah, M. J., Brunn, J. L., Wong, A. B., Wallace, M. A., & Carpenter, P. A. (1989). *Frames of reference for allocating attention to space: Evidence from the neglect syndrome.* Manuscript in preparation.

Grossberg, S., & Mingolla, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception & Psychophysics, 38*, 141-171.

Heilman, K. M., Watson, R. T., & Valenstein, E. (1985). Neglect and related disorders. In K. M. Heilman & E. Valenstein (Eds.), *Clinical neuropsychology* (2nd ed.). New York: Oxford University Press.

Hillis, A. E., & Caramazza, A. (1989). *The effects of attentional deficits on reading and spelling* (Technical Report 44). Baltimore, MD: Cognitive Neuropsychology Laboratory, Johns Hopkins University.

Hinton, G., & Shallice, T. (1989). *Lesioning a connectionist network: Investigations of acquired dyslexia* (Technical Report CRG-TR-89-3). Toronto, Canada: University of Toronto, Department of Computer Science, Connectionist Research Group.

Hinton, G. E. (1981). A parallel computation that assigns canonical object-based frames of reference. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence* (pp. 683-685).

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

Karnath, H. O. (1988). Deficits of attention in acute and recovered hemi-neglect. *Neuropsychologia, 26*, 27-43.

Kinsbourne, M. (1987). Mechanisms of unilateral neglect. In M. Jeannerod (Ed.), *Neurophysiological and neuropsychological aspects of spatial neglect* (pp. 69-86). Amsterdam: North Holland.

Kolb, B., & Whishaw, I. Q. (1985). *Fundamentals of human neuropsychology*. New York: W. H. Freeman.

Kučera, H., & Francis, W. N. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.

LaBerge, D., & Samuels, S. J. (1974). Toward a theory of automatic information processing in reading. *Cognitive Psychology, 6*, 293-323.

Ladavas, E. (1987). Is the hemispatial deficit produced by right parietal lobe damage associated with retinal or gravitational coordinates?. *Brain, 110*, 167-180.

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

McClelland, J. L., & Johnston, J. C. (1977). The role of familiar units in perception of words and nonwords. *Perception and Psychophysics, 22*, 249-261.

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

Mesulam, M-M. (1981). A cortical network for directed attention and unilateral neglect. *Annals of Neurology, 10*, 309-325.

Mesulam, M-M. (1985). *Principles of behavioral neurology*. Philadelphia, PA: F. A. Davis Company.

Mewhort, D. J., Marchetti, F. M., Gurnsey, R., & Campbell, A. J. (1984). Information persistence: A dual-buffer model for initial visual processing. In H. Bouma & D. G. Bouwhuis (Eds.), *Attention and Performance X* (pp. 287-498). Hillsdale, NJ: Erlbaum.

Mozer, M. C. (1987). Early parallel processing in reading: A connectionist approach. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 83-104). Hillsdale, NJ: Erlbaum.

Mozer, M. C. (1988). A connectionist model of selective attention in visual perception. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society* (pp. 195-201). Hillsdale, NJ: Erlbaum.

Mozer, M. C. (1988). *The perception of multiple objects: A parallel, distributed processing approach* (ICS Technical Report 8803). La Jolla: University of California, San Diego, Institute for Cognitive Science.

Mozer, M. C. (1989). Types and tokens in visual letter perception. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 287-303.

Norman, D. A. (1968). Toward a theory of memory and attention. *Psychological Review, 75*, 522-536.

Ogden, J. A. (1985). Anterior-posterior interhemispheric differences in the loci of lesions producing visual hemineglect. *Brain and Cognition, 4*, 59-75.

Pashler, H., & Badgio, P. C. (1985). Visual attention and stimulus identification. *Journal of Experimental Psychology: Human Perception and Performance, 11*, 105-121.

Patterson, K. E., Seidenberg, M. S., & McClelland, J. L. (1989). *Connections and disconnections: Acquired dyslexia in a computational model of the reading process.* Manuscript submitted for publication.

Posner, M. (1988). Structures and functions of selective attention. In T. Boll & B. Bryant (Eds.), *Clinical neuropsychology and brain function: Research, measurement, and practice* (pp. 179-202). The Master Lecture Series, American Psychological Association.

Posner, M., & Petersen, S. E. (1989). *The attention system of the human brain* (Technical Report 89-3). Eugene, OR: Institute of Cognitive and Decision Sciences, University of Oregon.

Prince, A., & Pinker, S. (1988). Wickelphone ambiguity. *Cognition*, 189-190.

Riddoch, M. J., & Humphreys, G. W. (1983). The effect of cuing on unilateral neglect. *Neuropsychologia, 21*, 589-599.

Rumelhart, D. E., & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception: Part II. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review, 89*, 60-84.

Seidenberg, M. (1988). Cognitive neuropsychology and language: The state of the art. *Cognitive Neuropsychology, 5*, 403-426.

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

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

Sieroff, E. (1989). *Perception of visual letter strings in a case of left neglect: Manipulation of the word form.* Manuscript submitted for publication.

Sieroff, E., & Michel, F. (1987). Verbal visual extinction in right/left hemisphere lesion patients and the problem of lexical access. *Neuropsychologia, 25*, 907-918.

Sieroff, E., Pollatsek, A., & Posner, M. I. (1988). Recognition of visual letter strings following injury to the posterior visual spatial attention system. *Cognitive Neuropsychology, 5*, 451-472.

Smolensky, P. (1986). Information processing in dynamical systems: Foundations of Harmony Theory. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume I: Foundations* (pp. 194-281). Cambridge, MA: MIT Press/Bradford Books.

Smolensky, P. (February 1987). *On variable binding and the representation of symbolic structures in connectionist systems* (Technical Report CU-CS-355-87). Boulder, CO: Department of Computer Science, University of Colorado, Boulder.

Touretzky, D. S., & Hinton, G. E. (1988). A distributed connectionist production system. *Cognitive Science, 12*, 423-466.

Treisman, A. M. (1969). Strategies and models of selective attention. *Psychological Review, 76,* 282-299.

Uhr, L. (1987). *Highly parallel, hierarchical, recognition cone perceptual structure* (Technical Report 688). Madison, WI: Computer Sciences Department, University of Wisconsin.

Warrington, E. K., & Shallice, T. (1980). Word-form dyslexia. *Brain, 103,* 99-112.

Zemel, R. S., Mozer, M. C., & Hinton, G. E. (1989). TRAFFIC: A model of object recognition based on transformations of feature instances. In D. S. Touretzky, G. E. Hinton, & T. J. Sejnowski (Eds.), *Proceedings of the 1988 Connectionist Models Summer School* (pp. 452-461). San Mateo, CA: Morgan Kaufmann.