Accepted for publication in *Journal of Cognitive Neuroscience*, 1996

The end of the line for a brain-damaged model of unilateral neglect

Michael C. Mozer
Department of Computer Science &
Institute of Cognitive Science
University of Colorado
Boulder, CO 80309–0430

Peter W. Halligan
Rivermead Rehabilitation Centre
Abingdon Road
Oxford OX1 4XD
United Kingdom

John C. Marshall
Neuropsychology Unit
University Department of Clinical Neurology
The Radcliffe Infirmary
Woodstock Road
Oxford OX2 6HE
United Kingdom

Abstract

For over a century, it has been known that damage to the right hemisphere of the brain can cause patients to be unaware of the contralesional side of space. This condition, known as *unilateral neglect*, represents a collection of clinically related spatial disorders characterized by the failure in free vision to respond, explore, or orient to stimuli predominantly located on the side of space opposite the damaged hemisphere. Recent studies using the simple task of *line bisection*, a conventional diagnostic test, have proved surprisingly revealing with respect to the spatial and attentional impairments involved in neglect. In line bisection, the patient is asked to mark the midpoint of a thin horizontal line on a sheet of paper. Neglect patients generally transect far to the right of the center. Extensive studies of line bisection have been conducted, manipulating—among other factors—line length, orientation, and position. We have simulated the pattern of results using an existing computational model of visual perception and selective attention called MORSEL (Mozer, 1991). MORSEL has already been used to model data in a related disorder, neglect dyslexia (Mozer & Behrmann, 1990). In this earlier work, MORSEL was "lesioned" in accordance with the damage we suppose to have occurred in the brains of neglect patients. The same model and lesion can simulate the detailed pattern of performance on line bisection, including the following observations: (1) no consistent across-subject bias is found in normals; (2) transection displacements are proportional to line length in neglect patients; (3) variability of displacements are proportional to line length, in both normals and patients; (4) position of the lines with respect to the body or the page on which they are drawn has little effect; (5) for lines drawn at different orientations, displacements are proportional to the cosine of the orientation angle. MORSEL fails to account for one observation: across patients, the variability of displacements for a particular line length is roughly proportional to mean displacement. Nonetheless, the overall fit of the model is sufficiently good that we believe MORSEL can be used as a diagnostic tool to characterize the specific nature of a patient's deficit, and thereby has potential down the line in therapy.

INTRODUCTION

Neglect

Unilateral neglect or *hemi-inattention* describes a collection of behavioral symptoms in which patients appear to ignore, forget, or turn away from contralesional space (Heilman, Watson, & Valenstein, 1993). Neglect after right sided lesions are more frequent, long lasting, and severe than after equivalent lesions of the left hemisphere. The disorder can compromise visual, auditory, tactile, and olfactory modalities and may involve personal, peripersonal, extrapersonal, and "imaginal" space (Halligan & Marshall, 1993a). Unilateral neglect is far from a unitary phenomena and has been shown to fractionate into a number of dissociable components in terms of sensory modality, spatial domain, laterality of response, motor output, and stimulus content (Barbieri & De Renzi, 1989). Furthermore, the lateralized behaviors observed cannot be explained in terms of concomitant sensorimotor deficits because neglect is manifested in free vision and under conditions of testing that do not necessarily require the use of motorically impaired limbs. Visual neglect has been repeatedly shown to double dissociate from hemiplegia, hemianopia, and hemianasthesia. Furthermore, the lesions that produce neglect are not restricted to primary sensory or motor cortex.

Unlike many brain-injured patients who rapidly learn to compensate for a visual field deficit or plegia of the upper limb, patients with unilateral neglect often continue to behave as if one half space has ceased to exist in any meaningful form (Mesulam, 1985). As a result, patients with neglect commonly collide with objects on the left side of space, fail to eat from the left side of the plate, and may dress only one side of the body. When copying or drawing, patients tend to confine their drawings to the right side of the page; the drawings themselves often include adequate representations of the right side of the object with the left side either entirely omitted or grossly distorted.

A related condition, commonly found in association with recovery from florid neglect, is *extinction*. In extinction, a patient can report seeing or feeling an object on the side opposite the brain damage when one stimulus is presented, but fails to detect the same stimulus when a second stimulus is simultaneously presented to the unaffected side. Extinction is regarded by some researchers (e.g., Heilman, Watson, & Valenstein, 1993) as a milder form of neglect which is revealed in the course of partial recovery.

Line bisection

Of all the clinical tests used to diagnose visual neglect, asking the patient to divide a horizontal line at its midpoint is by far the simplest and most widely employed (Bisiach, Bulgarelli, Sterzi, & Vallar, 1983; Bisiach, Capitani, Colombo, & Spinnler, 1976; Halligan & Marshall, 1988; Schenkenberg, Bradford, & Ajax, 1980; Tegner & Levander, 1991). As a clinical test, line bisection can be traced back some hundred years to the practice of the first clinicians who documented spatial asymmetries of perception in patients with unilateral brain damage (Axenfeld, 1915; Poppelreuter, 1917).

Typically in a line bisection task, patients are shown a thin line centered on a sheet of paper, which is centered with respect to the midsaggital plane. They are instructed to mark the midpoint of the line with a rapid, ballistic movement. Patients with left-sided neglect behave as if they have failed to notice the full left extent of the line, generally transecting the line far to the right of center. Performance can be quantified by measuring the *transection displacement*, the deviation of the patient's mark from the true center. We follow the convention that a positive displacement is a shift to the right of center, indicating left neglect; a negative displacement is a shift to the left of center, indicating right neglect.

Halligan and Marshall have studied line bisection extensively manipulating line length (Halligan & Marshall, 1988, 1989c; Marshall & Halligan, 1989, 1990a), orientation (Burnett-Stuart, Halligan, & Marshall, 1991; Halligan & Marshall, 1993b; Marshall & Halligan, 1990b), position (Halligan & Marshall, 1989b; Marshall & Halligan, 1990a), and the context in which the line is embedded (Halligan, Manning, & Marshall, 1991; Halligan & Marshall, 1989a, 1991a; Manning, Halligan, & Marshall, 1990). All of these variables have been shown to have reliable effects on the measured severity of neglect.

Computational modeling of neglect

We have replicated the pattern of results of neglect patients on the line bisection task in a computational model of two-dimensional object recognition and spatial attention (Mozer, 1991). The model, called MORSEL, was originally developed with two objectives in mind: (1) to build a computational mechanism that could analyze complex scenes consisting of multiple visual stimuli presented simultaneously, and (2) to account for a broad spectrum of psychological data, including perceptual errors that arise when several stimuli appear simultaneously in the visual field, facilitatory effects of context and redundant information, and attentional phenomena. The architecture and details of MORSEL arose from constraints imposed by these two objectives.

Mozer and Behrmann (1990) "lesioned" MORSEL in accordance with the damage that was hypothesized to occur in the brains of neglect patients. The lesioned model was then used to simulate some puzzling aspects of the performance of patients with neglect dyslexia—a reading disorder associated with neglect. For example, MORSEL could account for the data of Behrmann et al. (1990) indicating that patients, when shown a pair of words such as *cow* and *sun*, would generally report only the word on the right—a form of extinction. However, when the two words were related in that they could be combined to form a lexical item, e.g., *cow* and *boy*, extinction was less likely.

In the present work, we have used MORSEL, along with the lesion proposed in the earlier neglect modeling work, to simulate a set of findings on the line bisection task. The model, with no extensions or parameter adjustment, is successful in explaining a wide variety of effects. Interestingly, MORSEL makes an incorrect prediction for one minor aspect of the data, revealing a probable flaw in the model—most likely a technical, not conceptual, flaw. Nonetheless, the model makes interesting and testable predictions, and has the potential to be used as a diagnostic tool to better characterize the specific nature of a patient's deficit.

DESCRIPTION OF MORSEL

At this point, readers might wish for a detailed description of the data that we model. We ask the reader's indulgence, however, and begin with an overview of MORSEL. Following this overview, we can then present the data side-by-side with simulation experiments. Readers who are not fully motivated now may wish to first skim the "Data and Simulations" section to get a sense for the results.

Input to MORSEL

MORSEL was designed primarily to model letter and word perception. In contrast to earlier models of word perception (e.g., McClelland & Rumelhart, 1981), MORSEL has the capability of processing several items simultaneously. MORSEL has a visual field or *retina* on which stimuli are presented. The retina is a 36×36 array of cells. In each cell is a set of five feature detectors. These detectors register, or *become active* in response to, the presence of certain features in their region of the retina. Four of the detectors register oriented lines and the fifth detector registers the termination of a line. Figure 1 shows a sample input to MORSEL representing the phrase BAD ARTIST. (Note that we show only as many rows of the retina as are necessary in our figures.) Each letter produces a 3×3 pattern of activity on the retina.

This primitive activity pattern is interpreted by a connectionist network called BLIRNET, which attempts to identify the letters and words appearing on the retina. Figure 2 shows a sketch of MORSEL, with the BLIRNET module mapping the retinal activity pattern to a representation of letters and words. We envision a potentially large collection of processing modules that analyze various properties of the retinal array. For the line bisection task, for example, a module is required that identifies the center of mass of the stimulus activity.

MORSEL has two means of performing attentional selection: a late-selection component called the *pull-out net* and an early-selection component called the *attentional mechanism* or AM for short. The pull-out net acts on the outputs of BLIRNET to select items that have a certain lexical status or semantic content. The AM acts on the input to BLIRNET to select certain spatial regions for processing. BLIRNET, and other processing modules that operate on the retinal representation, is limited in capacity, and hence requires the AM's direction to avoid processing too much information at once.

We omit details of BLIRNET and the pull-out net, as they are not relevant for the current simulations. The interested reader can consult Mozer (1991). The simulations of line bisection, however, are dependent on the behavior of the AM, and thus we describe this component in some depth.

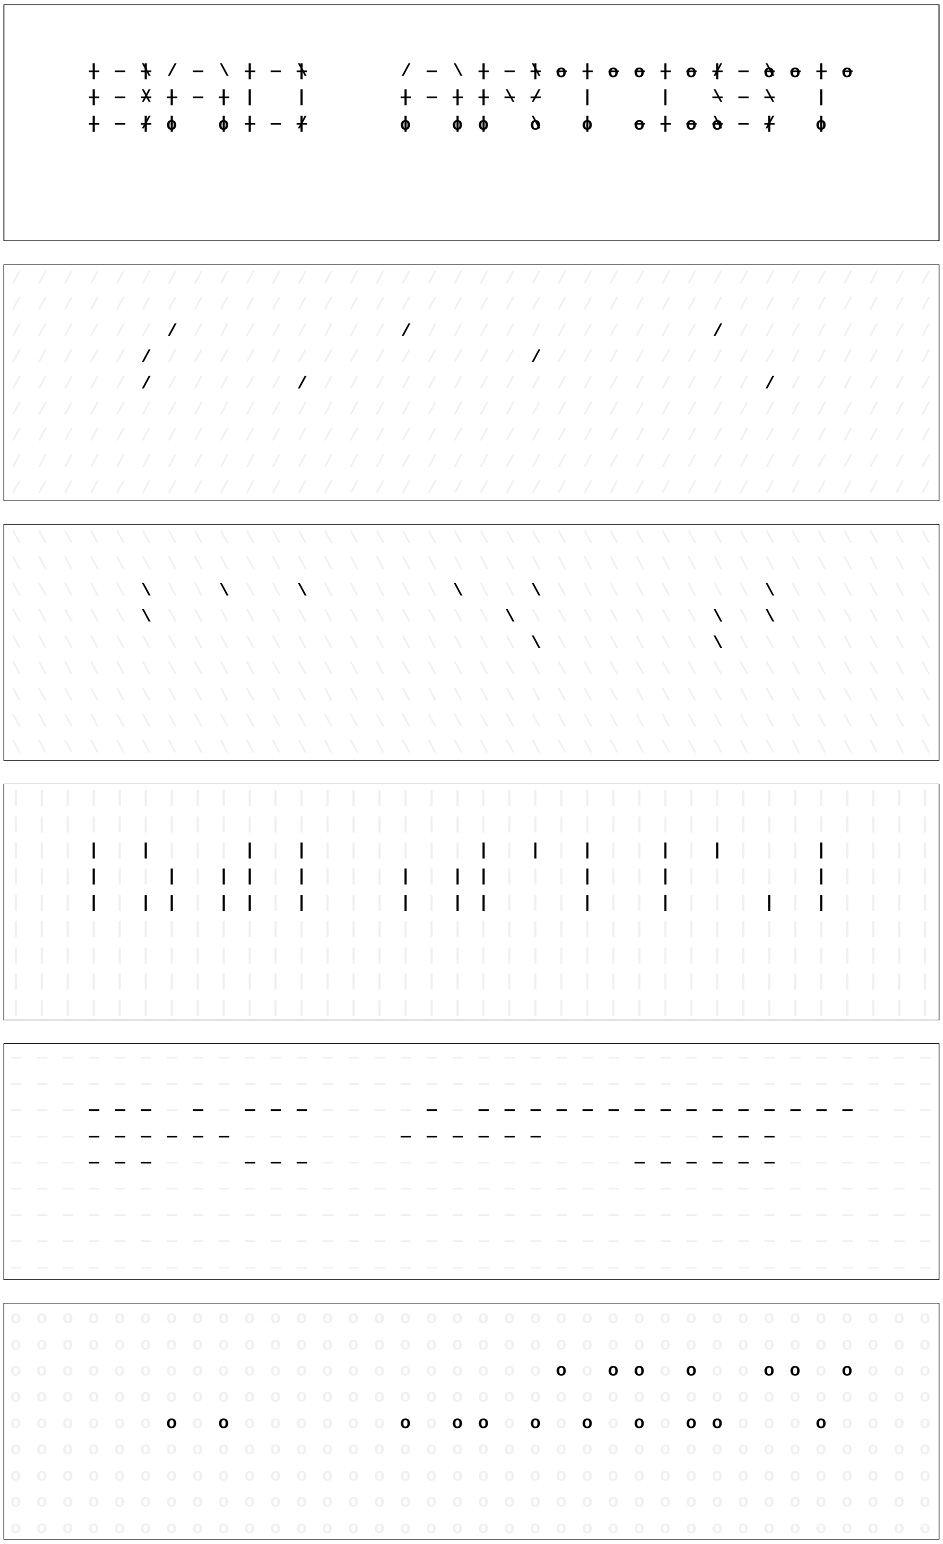

Figure 1: The top array shows the set of active feature detectors for a sample input, the phrase BAD ARTIST, presented on MORSEL's retina. The arrays below show activity by detector type. Each character in an array represents the activity of a single detector. The symbols –, /, |, and \ correspond to the 0°, 45°, 90°, 135° line detectors, and the symbol **o** corresponds to line terminator detectors. A darkened symbol indicates that the detector is active.

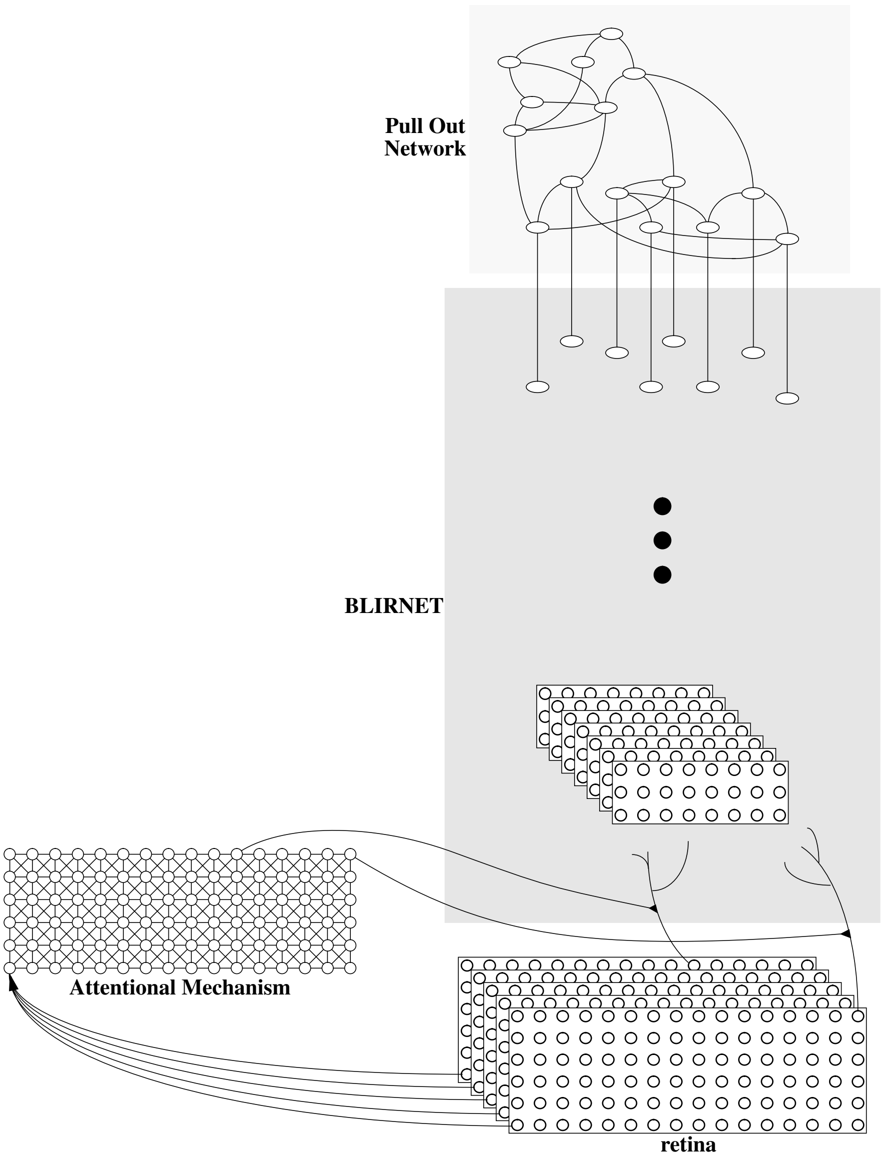

Figure 2: A sketch of MORSEL. The retinal representation is shown on the bottom right. Each circle corresponds to a detector or processing unit. The sheets of units correspond to a spatiotopic array of units all having the same feature type. The ovals at the top of the BLIRNET module are units that represent letters and words.

The Attentional Mechanism

The AM receives advice 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 within its bounds relative to those outside. This causes preferential treatment in BLIRNET and other processing modules. However, the AM does not act as an all-or-none filter. Information from unattended regions of the retina undergoes some degree of analysis by BLIRNET and other processing modules. This partial processing of unattended information distinguishes the AM from other early-selection filtering mechanisms that have been proposed (e.g., Koch & Ullman, 1985; Laberge & Brown, 1989), although this is not relevant in the current work.

The AM 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. The connections in Figure 2 from the retina into the AM provide this input. 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 2, 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 retina into BLIRNET and other processing modules. Specifically, the activity level of a unit in a given location in the retinotopic array is transmitted to the processing modules with a probability that is monotonically related to the activity of the AM unit the the corresponding array 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.

Each unit in the AM gets bottom-up input from the detectors in the corresponding location in all of the 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 array, with a bias towards locations indicated by bottom-up and top-down inputs. Details of the AM selection process are provided in Appendix 1. Figure 3 shows an example of the AM in operation. Two blobs of activity are presented to the AM via its external inputs, and the AM settles on the left blob.[1]

Damaging MORSEL to Produce Neglect

To model data from neglect dyslexia, Mozer and Behrmann (1990) proposed a particular form of lesion to the model—damaging the bottom-up connections to the AM from the retinal feature

[1]Mozer and Behrmann (1990) assumed a slight amount of blurring in the bottom-up input to the AM: Each retinal activation provided input not only to the corresponding location in the AM but also to the immediately adjacent locations, with a relative strength of 2%. This small amount of spread is unlikely to affect processing, but we have preserved it to maintain consistency with the original model.

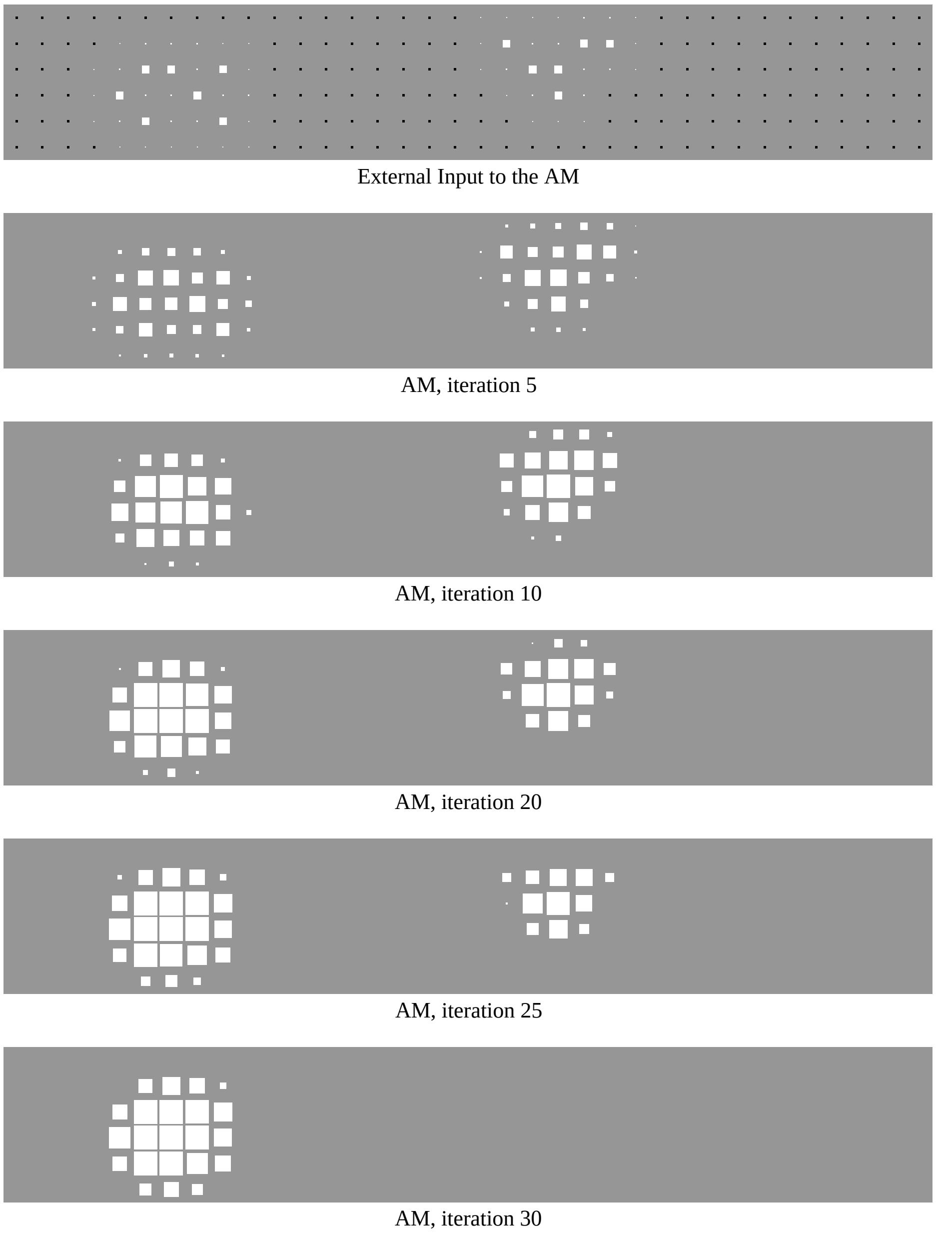

Figure 3: An example of the AM selecting a single location. The top panel shows the external input to the AM, arranged in a two-dimensional grid. The area of each white square is proportional to the magnitude of the external input to that location on the grid; the black dots indicate AM grid locations that receive no input. The lower five panels show activity in the AM as it settles over time. The area of each white square in these figures is proportional to the activity of the AM unit in the homologous position of the array. "Iteration" refers to the number of discrete time steps that have transpired.

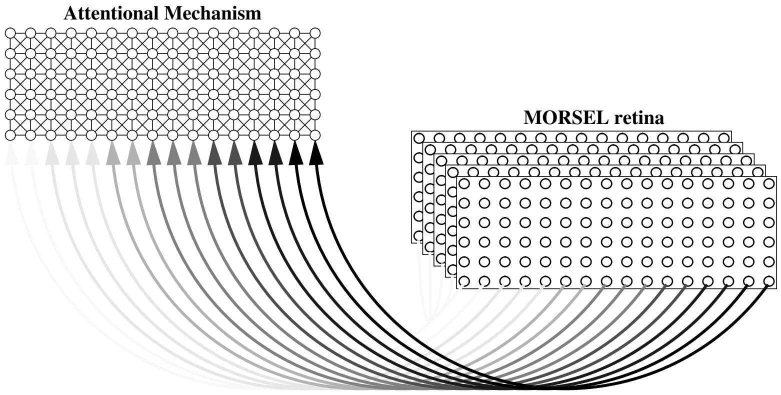

Figure 4: A sketch of the AM and some of its inputs from MORSEL's retina. Each detector on the retina connects to the homologous unit in the AM. In neglect, we propose that there is graded damage to these connections, causing retinal detectors to be less effective in activating the AM. The damage is depicted by the fainter connections toward the left side of the field.

arrays. 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 will throughout this article). Figure 4 depicts the damaged connections into the AM. The graded damage is important; complete destruction of the connections in the left field and fully intact connections in the right field would yield a qualitatively different sort of behavior, e.g., Mozer and Behrmann achieved what might be interpreted as object-based neglect via the graded damage. It is also important to contrast this proposal for lesioning the model with two alternatives. First, one might damage the visual recognition system (BLIRNET) itself. However, this would lead to blindness, and is inconsistent with the view of neglect as an attentional phenomenon and with the neuroanatomical lesion sites that give rise to neglect. Second, one might lesion the AM directly, either changing the activation dynamics or connectivity of the units such that damaged units integrated activity more slowly or had a weakened influence on the activity of other units. We conjecture that these types of lesions would yield a behavioral effect similar to the proposed lesion for the simulation studies reported in this article.

The damage affects the probability that features present on the retina 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 mustered, features will be analyzed normally by BLIRNET and other processing modules.

Mozer and Behrmann showed that this form of damage was compatible with both early, peripheral effects observed in neglect dyslexia because the disruption directly affects a low-level representation. However, the damage was also compatible with late, higher-order effects in neglect dyslexia by virtue of the fact that the pull-out net is able to reconstruct elements of a string—via lexical and semantic knowledge—that are attenuated by the AM.

MODELING LINE BISECTION

Line bisection is a simple cognitive task; it does not require visual object recognition and does not invoke higher-order knowledge. For this reason, much of MORSEL, which deals with shape recognition and the top-down influence of semantic and lexical knowledge, is irrelevant. Only the early selection attentional system—the AM—is pertinent. Thus, in terms of MORSEL, line bisection is an extremely pure and primitive cognitive task.

In brief, we model the line bisection task by presenting the line stimulus on the retina, allowing the AM to settle, and then supposing a motor process that places the transection point at the center of the attended region. We assume that the motor process thus veridically bisects the attended region; there is no motor neglect or motor bias.

Stimuli

Figure 5 shows the representation we constructed for a line stimulus appearing on MORSEL's retina. The line is made up of two rows of horizontal segments with terminators at the extremities.[2]

The number of features active in a given cell of MORSEL's retina determines the amount of input fed to the corresponding cell of the AM. Figure 6 shows the external input to the unlesioned AM for line stimuli of varying lengths, and Figure 7 shows the external input for line stimuli of various orientations.

To establish a correspondence between distances on the retina and physical distances in the world, we arbitrarily assume that each cell of the retina corresponds to an 8.47 millimeter (= 1/3 inch) distance in the stimulus display. Consequently, the line lengths of 3 to 33 cells, depicted in Figure 6, correspond to stimulus lines of 25 to 279 mm.

Lesioning MORSEL

All parameters and details of MORSEL—including connection strengths, interconnectivity patterns, and time constants—were left unchanged from the earlier work of Mozer (1991) and Mozer and Behrmann (1990). However, MORSEL has several remaining degrees of freedom that relate to the nature of the attentional deficit, which presumably differs from patient to patient.

Specifically, four parameters determine a function relating the horizontal position of a feature on MORSEL's retina to the probability that the feature will be transmitted to the corresponding cell of the AM. This function is shown in Figure 8. It is a piecewise linear curve with a flat segment, followed by a segment with positive slope, followed by another flat segment. The left extreme of the curve represents the left edge of the retina, the right extreme of the curve represents the right edge of the retina. The probability that the AM will register a feature is low in the left field, and it is monotonically nondecreasing further to the right.

[2] The exact representation does not appear to be critical. We also experimented with lines having thicknesses of one and three cells, and the qualitative pattern of results was unchanged.

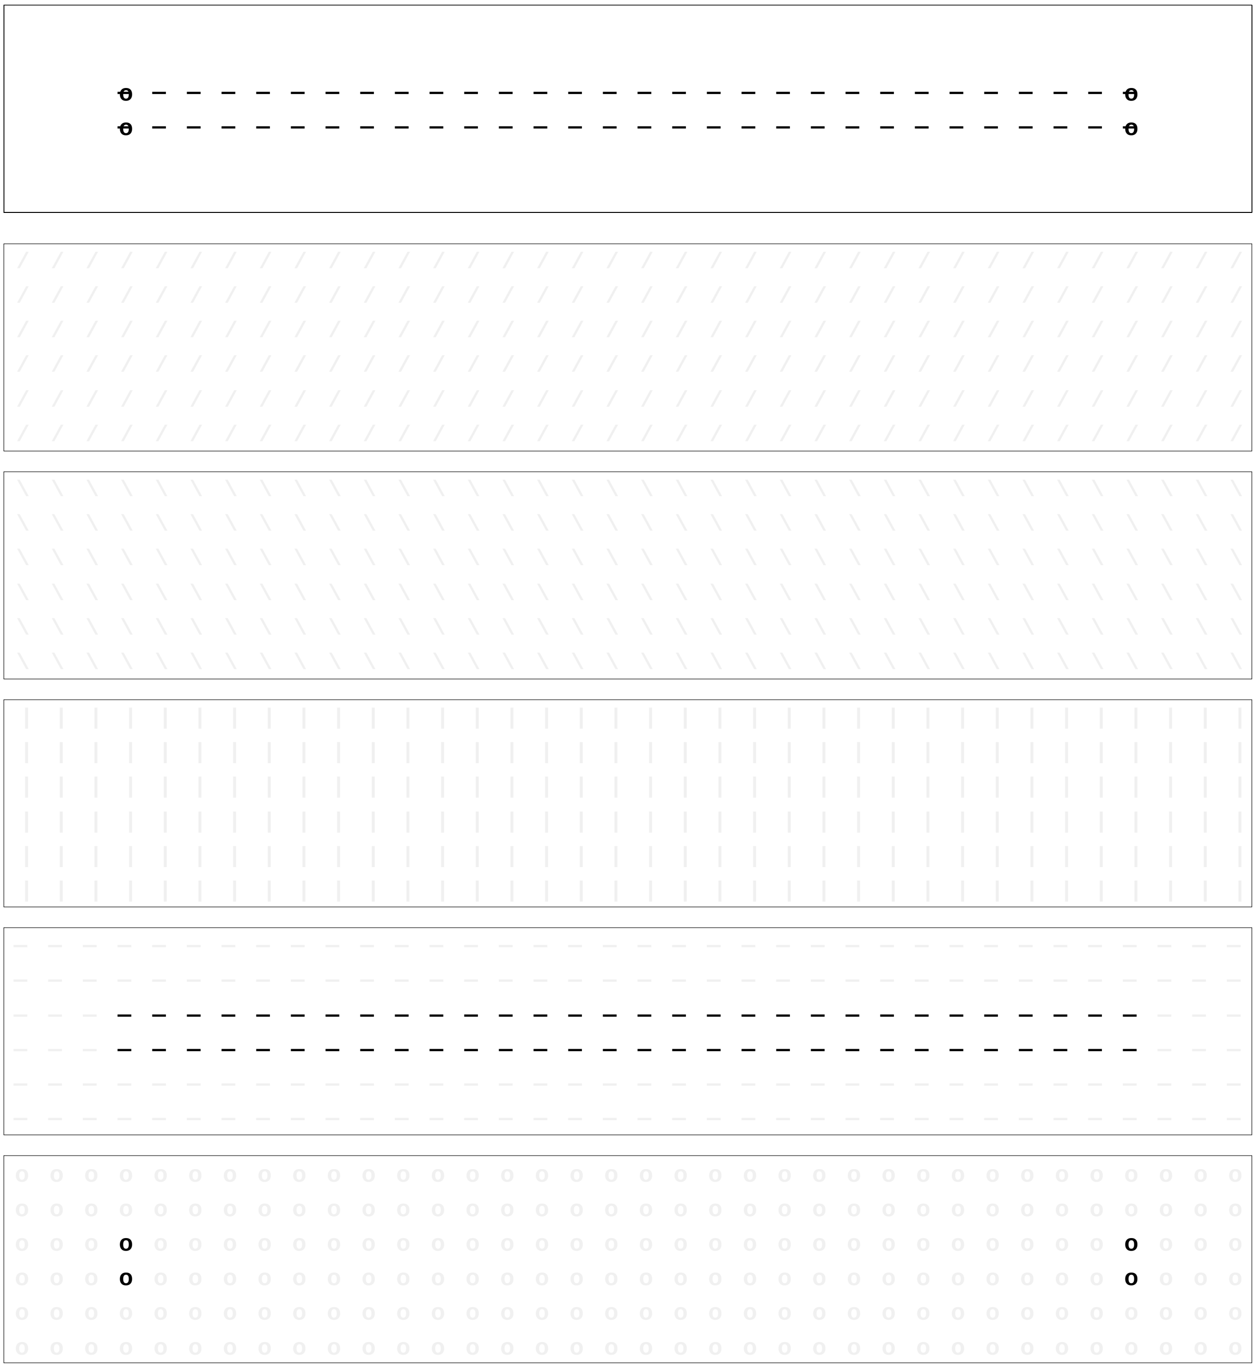

Figure 5: The pattern of feature activity produced by a line on MORSEL's retina.

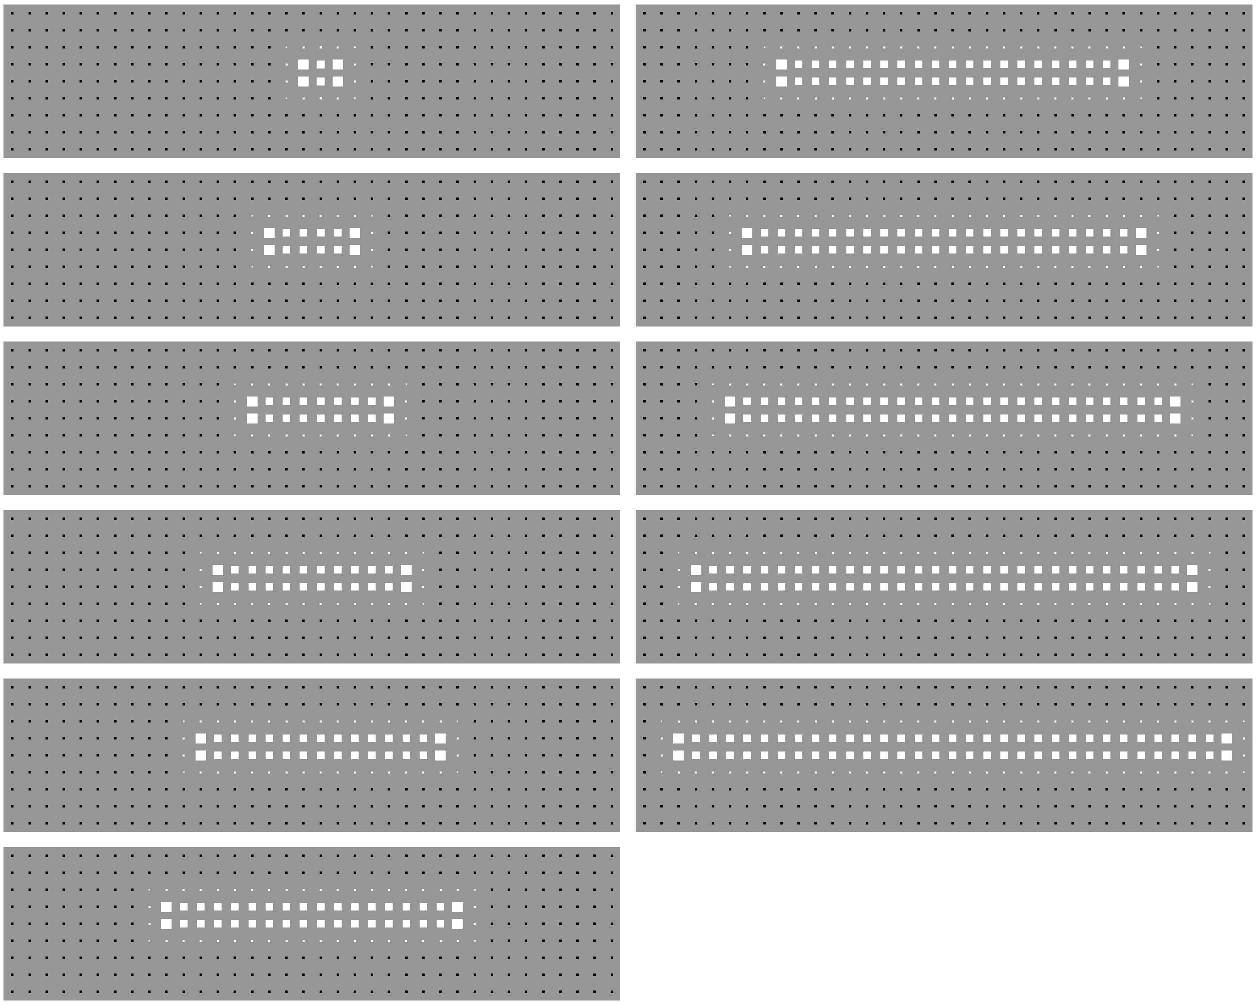

Figure 6: External input to the unlesioned AM for line stimuli of varying lengths. The size of a white square is proportional to the magnitude of the external input. The black dots indicate AM grid locations that receive no external input. Because both horizontal and terminator detectors are active at the endpoints of a line, external input is twice as large at the endpoint locations. Note that a feature at a given location on the retina provides input not just to the corresponding location on the AM but also provides a small amount of input to adjacent locations. This spillover results in the halo around the stimulus locations.

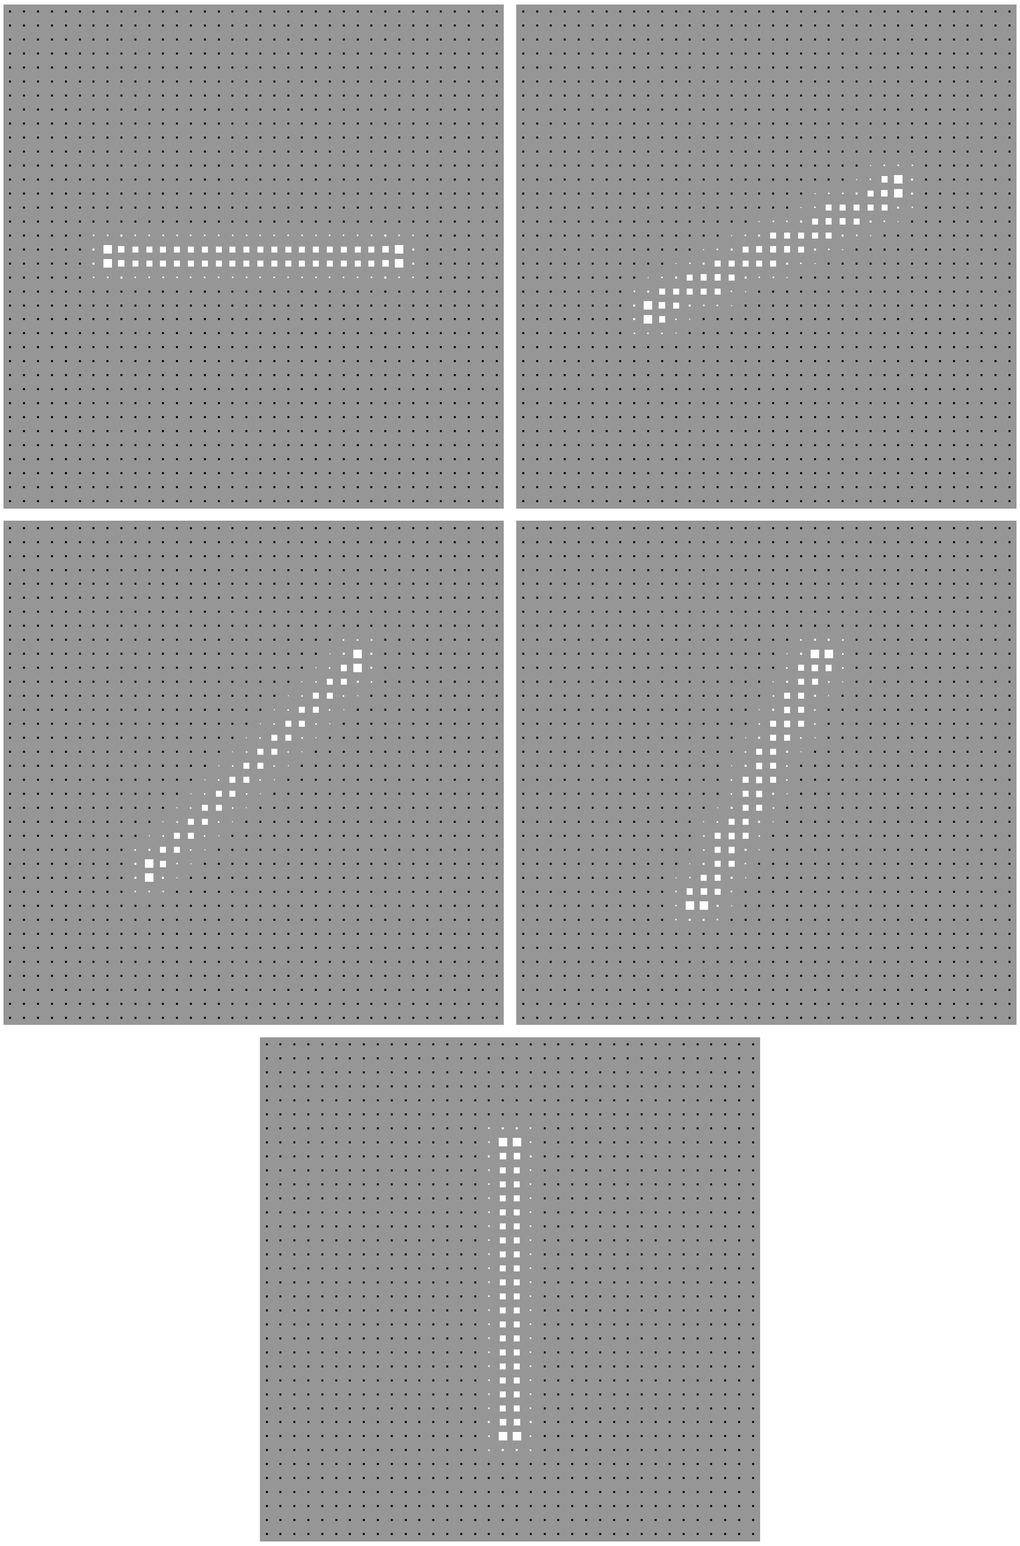

Figure 7: External input to the unlesioned AM for line stimuli of various orientations.

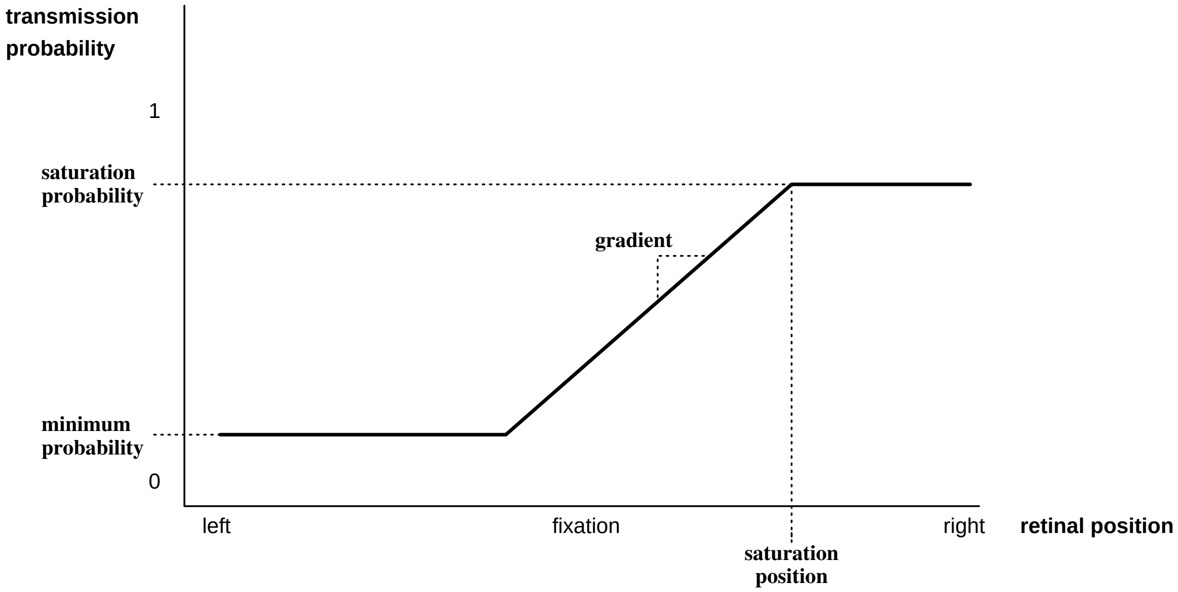

Figure 8: The transmission probability curve representing the damage to the attentional system in MORSEL. This function relates the position of a feature on MORSEL's retina to the probability that the feature will be detected by the corresponding cell of the AM. The function is for a left neglect patient; the probability that the AM will register a feature is low in the left field, and it is monotonically nondecreasing further to the right.

The curve is characterized by four parameters: (1) the probability of feature transmission on the right end of the retina (*saturation probability*), (2) the horizontal position on the retina where the probability reaches aymptote (*saturation position*), (3) the minimum transmission probability (*minimum probability*), and (4) the slope of the curve (*gradient*). Mozer and Behrmann (1990) used this same general type of curve, without the explicit parameterization, in their earlier neglect simulations. The basic notion of an attentional gradient is due to Kinsbourne (1993). The parameterization of the curve allows a variety of transmission functions, including forms corresponding to normals (e.g., a minimum probability close to 1 and a gradient of 0), a homogeneous slope across the entire field (e.g., a shallow gradient and a saturation position at the far right edge), and a sharp discontinuity at the hemifield crossing (a very large gradient and a saturation position just to the right of fixation).

This curve representing damage is cleaner and sharper than one would expect from a biological system. Minimally, it should be curvilinear with no discontinuities in slope. However, minor changes to the form of the curve have little effect on the model's behavior, and the parameterization of curve is quite flexible in that it allows a wide variety of possible forms of damage.

It turns out that the behavior of the AM is robust over various transmission probability curves, except when extreme values of the parameters are selected. Interestingly, some sets of parameters do lead to odd and incorrect predictions. This does not invalidate MORSEL, but rather places constraints on the range of parameter values and their interrelationships that can occur in patients. This in itself is a useful contribution of the modeling effort. We say more about this issue below.

Rather than picking a single set of parameters somewhat arbitrarily, the simulations reported below explored a collection of parameter sets. Because each patient presumably has a slightly different form of deficit, and because we are modeling mean performance across patients, it seems sensible to construct an ensemble of lesioned models, each with a slightly different transmission probability curve. Our ensemble was formed by considering saturation probabilities of .9 or 1.0, gradients of .01 or .02, saturation positions of 50%, 75%, or 100% from the left edge of the retina, and minimum probabilities of .2 or .4. The Cartesian product of these parameter values was formed, yielding 24 different transmission probability curves.

In most of our simulations, we modeled ten replications of each of the 24 transmission probability curves, yielding a total of 240 simulated patients. The results reported below are means over these simulated patients. Except where otherwise noted, we ran ten trials per simulated patient per condition. This is the amount of data collected on real patients in the corresponding studies.

Methodology

When a stimulus is presented on MORSEL's retina, processing proceeds as follows. Activation is transmitted from the retina to the AM according to the transmission probability function. This triggers the AM to select a contiguous spatial region on the retina. The AM iterates toward a stable state over time. Examples of this in the normal and lesioned model are shown in Figures 9 and 10, respectively. The response of the unlesioned model is symmetric about the center of the line stimulus, whereas the response of the lesioned model is shifted toward the right side of the field. While Figure 10 shows the response on a single trial, of a version of the AM with a particular transmission probability function and a particular sampling of activation transmitted to the AM based on this function, it is quite typical, as will be evident from results presented in subsequent sections.[3]

Once the AM has settled, the center of mass of the blob of activity is computed using the formula:

$$p = \frac{\sum_x \sum_y x a_{xy}}{\sum_x \sum_y a_{xy}}$$

where a_{xy} is the activity of AM unit at location (x, y). If p_0 is the horizontal position of the true center of the line, then the transection displacement is $p - p_0$. This displacement can scaled to units of millimeters, using the conversion factor described above.

There are several reasonable procedures for translating the AM pattern of activity into a transection response from the model. For instance, one could determine the leftmost and rightmost point of activity in the AM, call them p_l and p_r, and the transection point could be computed as the average, $p = (p_l + p_r)/2$. We experimented with several such alternative procedures. Fortunately, the results of simulations were quite robust to the read out procedure.

[3]Examining the external input to the AM in Figure 10, one notices that the actual input strength varies from location to location, and some locations to the right of center actually have stronger input than some locations to the left of center. This is because the Figure shows a particular sampling from retinal input based on the transmission probability function, rather than the *expected* input strength, which would increase monotonically from left to right.

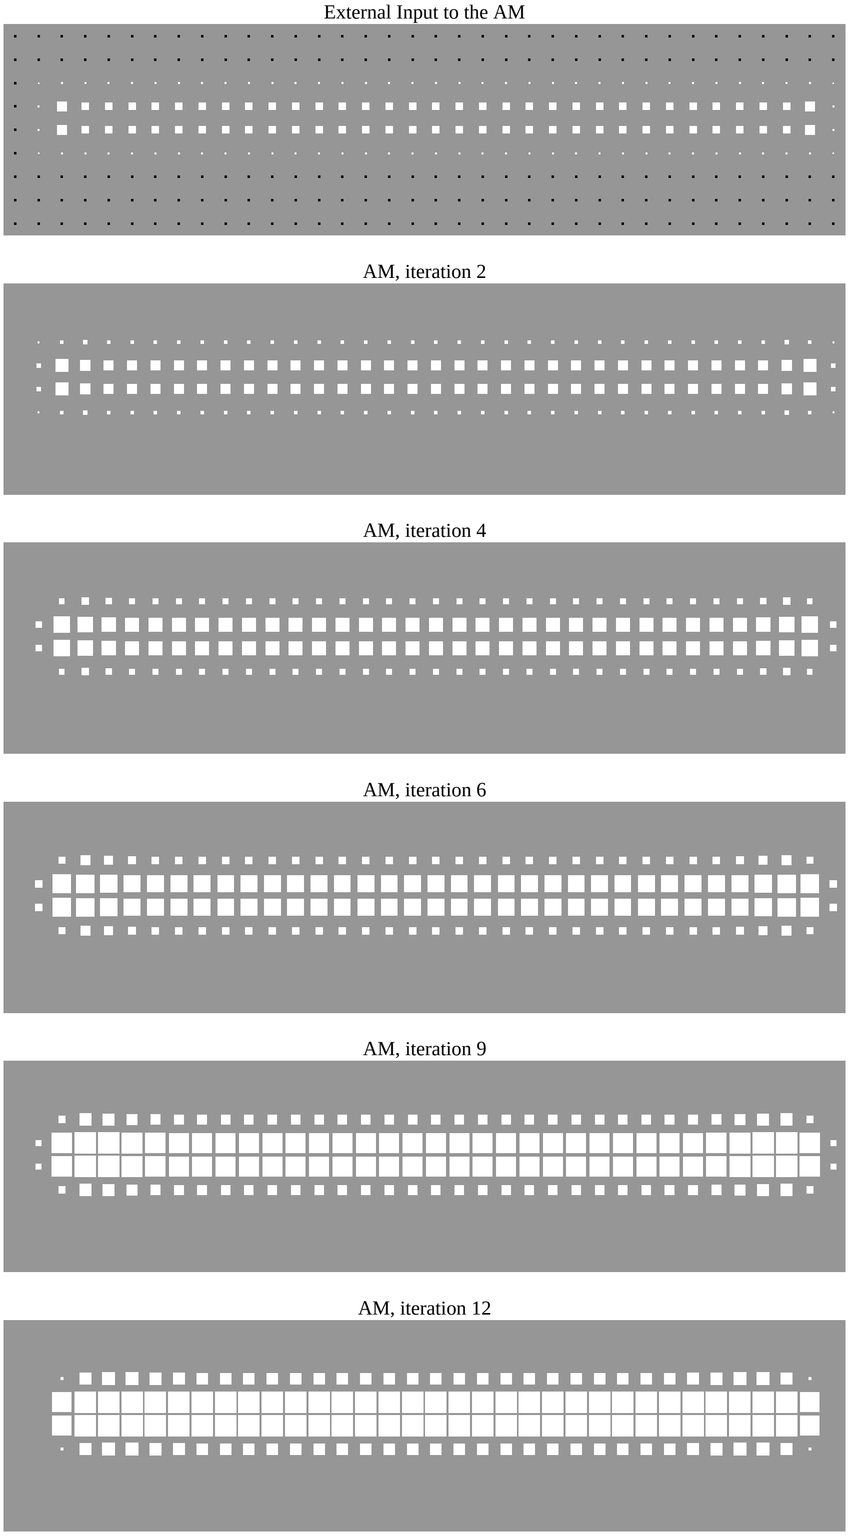

Figure 9: The unlesioned AM responding to a line stimulus. The top panel shows the external input to the AM. In this simulation, the AM selects all locations containing features of the line, in addition to some neighboring locations, and its response is symmetric about the center of the line.

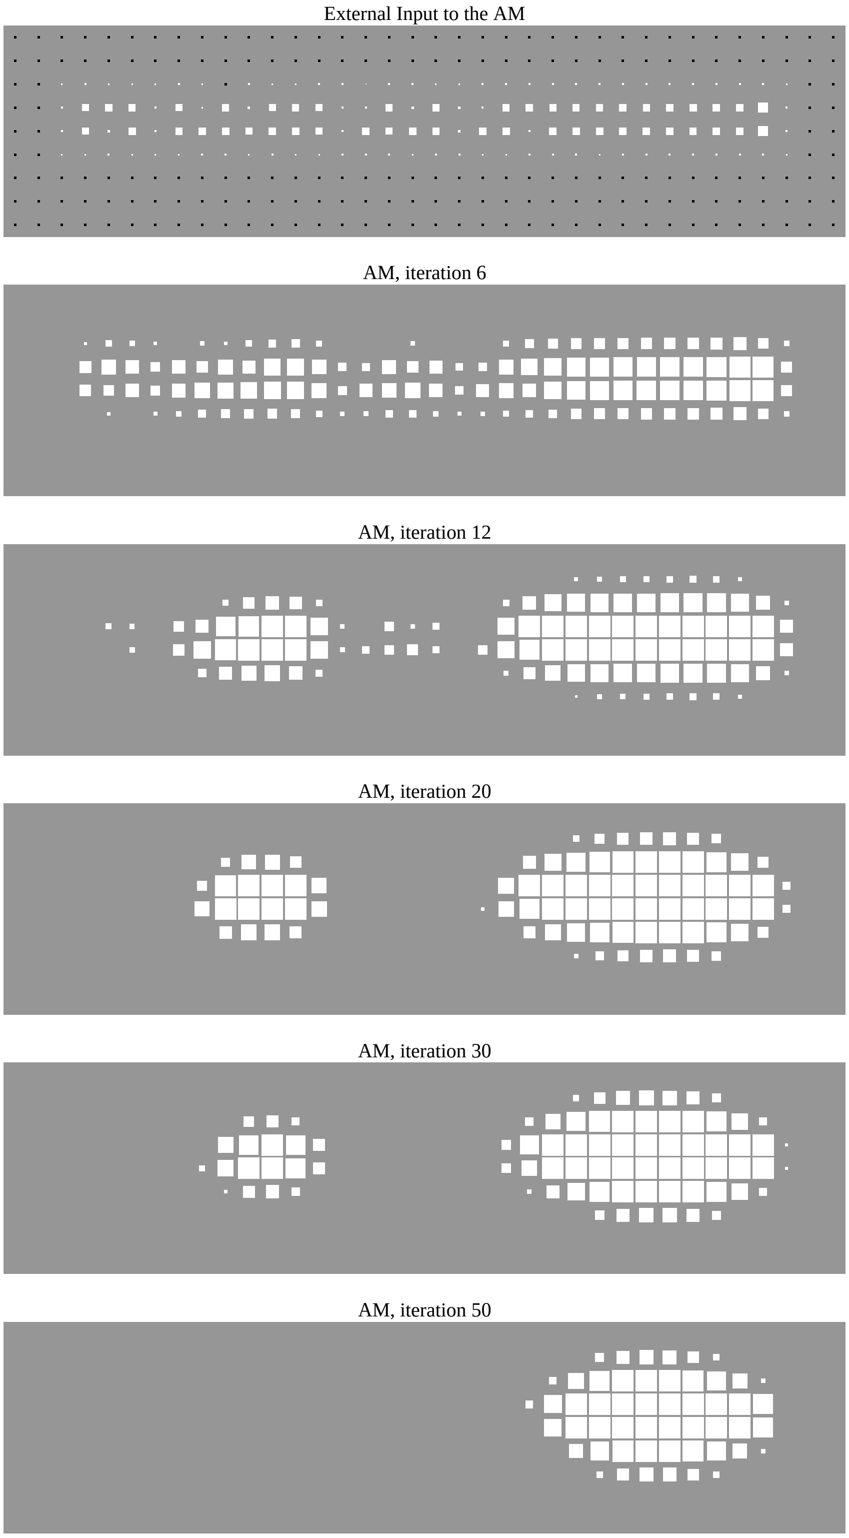

Figure 10: The lesioned version of the AM attending to a line stimulus. In contrast to Figure 9, the external input to the AM is degraded, particularly in the left region of the retina. Rather than processing the line as a single entity, the AM drops some portion of the line, resulting in a competition between the remaining portions of the line. The right portion wins out as its external input is stronger. The final state of the AM is not symmetric about the center of the stimulus line.

DATA AND SIMULATION ACCOUNTS

In this section, we describe the basic phenomena of line bisection, along with the model's explanation of the phenomena. A great deal of data has been collected on this task. We have attempted to sift through the data (including unpublished data) and characterize the consistent and robust phenomena, as well as curious but statistically reliable quirks.

1. Normal performance

Studies of normal performance on line bisection have been carried out using college-aged subjects, elderly controls, and neurological patients without neglect (e.g., Bisiach, Capitani, Colombo, & Spinnler, 1976; Bradshaw, Nettleton, Nathan, & Wilson, 1985; Halligan & Marshall, 1989c; Heilman, Bowers, & Watson, 1984; Manning, Halligan, & Marshall, 1990). Some studies have shown overall mean leftward transection displacements, others have shown overall mean rightward displacements, while other have shown no overall bias. The mean displacements are quite small, less than 1–2% of the total line length in magnitude.

The variability across studies is not mysterious. It seems clear that individual subjects do show overall biases, roughly 3–4% of total line length. About half of the subjects transect to the left and half to the right. Consequently, with small sample sizes, one should expect small positive or negative means across subjects.

Two additional facts have emerged. First, the mean displacement for a given subject—either to the left or to the right—increases with the line length, although the effect is small because the displacements are small. Second, within-subject variability of transections is positively correlated with line length. For example, Manning, Halligan, & Marshall (1990) found a linear correlation between line length and standard deviation of responses to be .97; Halligan & Marshall (1989) found a correlation of .88.

Simulations

In the unlesioned version of MORSEL, the transmission probability curve is uniform across the visual field—there is no attentional bias or gradient. To model normal performance, we must nonetheless specify the transmission probability. If the probability is 1.0 everywhere in the field, then the model will always select all features of the line (see Figure 9) and will thus always bisect the line correctly. Some noise must be introduced into performance to model the variability among normals. One way to do this would be to assume noise in the motor response process, but it's also possible to introduce noise in perception by setting the transmission probability to, say, .9 uniformly across the field. (Mozer, 1991, also used a transmission probability of .9 to model normals.)

Figure 11 shows the mean transection displacement as a function of line length for the unlesioned model with uniform transmission probability of .9. The mean displacements are not reliably different from zero, regardless of the line length. However, there is a clearly monotonic relationship between the line length and the variability of the responses across simulated subjects: The standard deviation of transection displacements, which are linearly related to the error bars shown in the

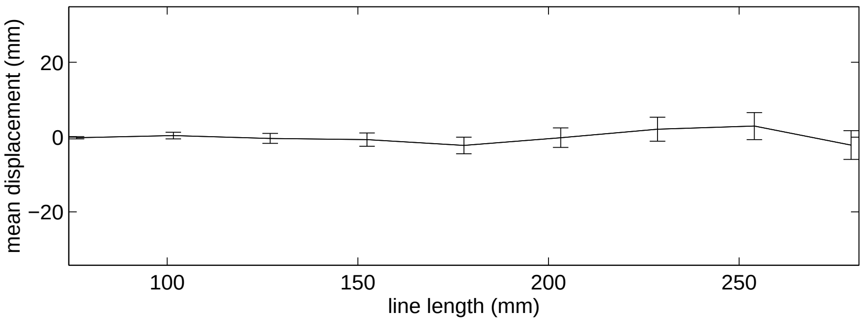

Figure 11: Simulation of unlesioned MORSEL on line bisection task as a function of line length. The bars indicate the standard error of the mean. Two hundred and forty simulated subjects were run, with ten trials at each line length for each subject. The ten trials were averaged, resulting in one data point at each line length for each subject.

Figure, get larger as the line length increases. This replicates the general pattern of normal human performance.

2. Effect of line length

In neglect patients, transection displacement is monotonically related to line length (Bisiach, Bulgarelli, Sterzi, & Vallar, 1983; Halligan & Marshall, 1988; Nichelli, Rinaldi, & Cubelli, 1989; Riddoch & Humphreys, 1983). Figure 12a shows a typical curve for a patient. For short lines of about 50 mm, the patient is fairly accurate, but as the line length increases, the transection point shifts farther to the right of center. Roughly, the displacement magnitude is proportional to the line length. The proportionality constant varies from one patient to another, but all show this general pattern.

Data from forty-three patients with unilateral right hemisphere damage were reanalyzed to quantify this observation. Twenty six of these patients were tested with lines varying from 18 to 180 mm in 18 mm increments; sixteen were tested with lines of 25 to 279 mm in 25.4 mm increments. Ten trials were run for each line length. We examined the shape of the curves relating line length to transection displacement to determine the linearity of the relationship. Because subjects perform very well at the shortest lengths, and hence there is a floor effect, we removed 50 mm and shorter lines from the analyses.

Of the forty-three patients, forty two showed a strong linear relationship between line length and transection displacement. The mean percent of variance accounted for with a linear regression was 90.3%. With a quadratic term included in the regression, the percent variance accounted for rose only to 92.2%.

The one exception, patient MN, showed a significant quadratic component in his response function. The curve was negatively accelerated: in proportion to line length, MN showed less neglect for long lines than for short lines. MN was studied over fifteen sessions to assess the reliability of his performance. Averaging MN's responses over the fifteen sessions, a linear regression accounted for only 85.8% of the variance, a quadratic regression accounted for 97.6%.

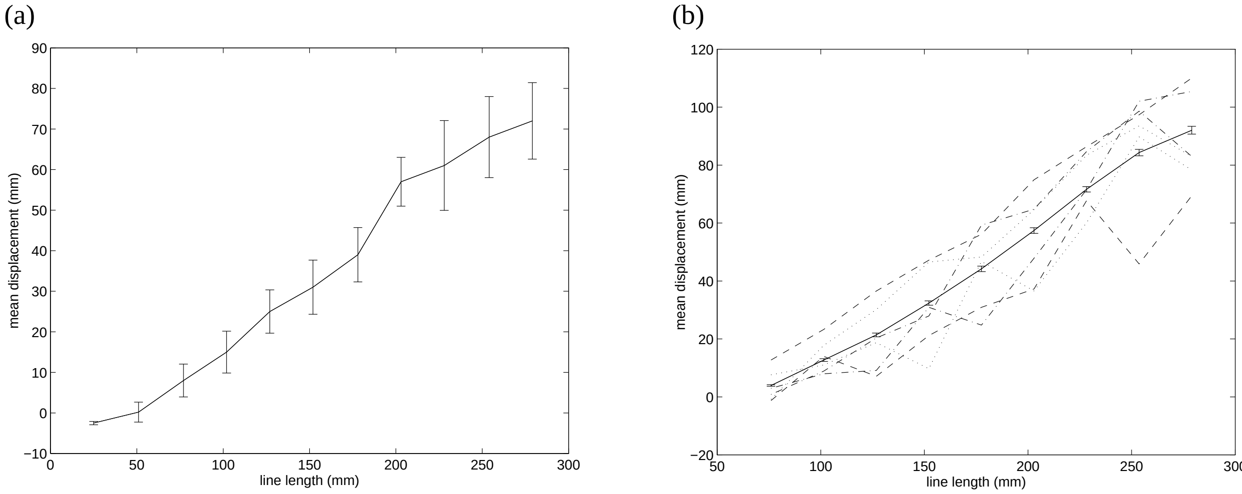

Figure 12: (a) Mean transection displacement as a function of line length for patient PP from Marshall & Halligan (1989). Each point is the average over 10 trials. The length of the vertical bar extending in each direction from the point is +/− 1 standard error of the mean. (b) Simulation of lesioned MORSEL on line bisection task as a function of line length. The dotted and dashed lines show typical curves from six simulated patients. The solid line with the error bars shows the mean across simulated patients.

Simulations

Varying-length lines were presented to the lesioned MORSEL. Figure 12b shows mean transection displacement as a function of line length averaged over the 240 simulated patients, as well as the data for six individual simulated patients who represent the sort of variability found in the data.

The mean curve is clearly linear. To examine the nature of the curve for individual simulated patients, we computed the proportion of variance accounted for by a linear regression and a quadratic regression (i.e., having both first and second order terms) for each simulated patient. Averaged over simulated patients, the linear fit accounted for 90.3% of the variance, and the quadratic fit 91.7%. This closely replicates the corresponding statistics from the human patients—90.3% for the linear fit and 92.2% for the quadratic.

Over all simulated patients and line lengths, the mean transection displacement was 23% of the line length with a standard deviation of 12%. These figures match typical patient performance; for example, patient PP (Figure 12a) showed a mean displacement 22% of the line length and with a standard deviation of 13%.

Note that we omitted very short lines—having lengths of 25 or 51 mm—from Figure 12b and the regression analyses. This is because most of the simulated patients performed nearly perfectly at these lengths, resulting in a floor effect (see Table 1) and corrupting the linearity of the relationship between line length and transection displacement. However, these very short lines were also omitted from the the human performance statistics for the same reason; hence, the the model and human data are comparable.

Table 1: Mean transection displacement for various line lengths—lesioned model

line length (mm)	*mean displacement*
25	0.1
51	1.1
76	3.9
102	12.7
127	21.4
152	32.4
178	44.2
203	57.4
229	71.6
254	84.3
279	92.0

3. Cross over of displacements at short lengths

Typically, neglect patients transsect to the right of the true midpoint. However, for very short line lengths—less than about 50 mm—some patients consistently bisect to the *left* of the midpoint (Halligan & Marshall, 1988; Marshall & Halligan, 1989). Thus, the left neglect observed for long lines becomes a sort of right neglect for the shortest lines. When this occurs, the transection displacements *cross over* the objective center of the lines. Linear extrapolation from the displacements for longer lines often predicts this counterintuitive phenomenon.

Account in terms of the model

Although this phenomenon is on the surface puzzling, we offer a simple explanation. The phenomenon might be viewed as emerging from a combination of two assumptions: (1) some normals consistently err to the left on line bisection; and (2) patients exhibit little or no neglect for short lines. The evidence for the first point was given earlier, and the leftward bias exhibited by some normals may be present even after brain damage. If this leftward bias is additive with large rightward shifts due to neglect, it will not be detected. However, if neglect plays no significant role for short lines, then the leftward bias might again dominate. Thus, rather than viewing the cross over phenomenon as fundamentally tied to neglect, one could sensibly attribute cross over to the *absence* of neglect for short lines.

Cross over in most human patients occurs only for line lengths of 25 and 51 mm. As Table 1 shows, the lesioned MORSEL produces little or no neglect at these line lengths, allowing even a slight leftward bias to push the responses to the left of the midpoint.

To quantify this argument, consider the combination of two independent influences on performance: left neglect which causes deviations to the right of the midpoint that are proportional to line length, and a fixed bias to deviate to the left of the midpoint for all line lengths. One can express this mathematically as

$$d = .5nl - b$$

where d is the observed transection displacement, l is the line length, n quantifies the severity of neglect (0 = none, 1 = complete), and b is the fixed leftward bias. When the line length is sufficiently short, the contribution of the first term will be overwhelmed by the fixed leftward bias, and the displacement will end up negative, i.e., to the left of the midpoint.

Note that we make no assumptions concerning what system or systems produce the leftward bias. It is unrelated to brain damage, as normals exhibit these same biases. It could stem from perceptual, motor, or perhaps even spatial reasoning systems.

4. Response variation as a function of line length

At each line length, experimenters typically collect multiple responses from patients. The responses are averaged to factor out some sources of variability in performance. However, this variability is itself interesting. One can quantify the variability as the standard deviation of transection displacements over the trials for a given line length. This standard deviation tends to be proportional to line length. For example, patient AL was tested over six sessions and correlations between standard deviations and line length ranged from .55 to .89 (Halligan & Marshall, 1991b). Patients DF, MN, CC, BS, TB, EF, TR, and TM, reported in Halligan and Marshall (1993b), had correlations of .88, .94, .64, .96, .89, .92, .92, and .92. Patient PP's correlation was .93 (Marshall & Halligan, 1989; see Figure 12a). Patients PB (Halligan & Marshall, 1988), JH, and PS (Halligan & Marshall, 1989c) showed correlations of .84, .84, and .76 respectively.

Qualitatively, the distribution of responses is unimodal, with 5–25% of the transections being made to the left of center. The vast majority of these left-sided responses occurred at the shorter line lengths. In the case of the serial assessment of AL, the mean proportion of left-sided transections comprised 21% of the total responses. For the other neglect patients described above—DF, PP, MN, CC, BS, TB, EF, TR, TM, PB, JH, PS—the percentage of left sided responses was as follows: 24, 12, 0, 5, 5, 9, 8, 5, 17, 8, and 17.

Simulations

We examined the distribution of transection displacements across simulated patients. The mean transection displacement for each patient and line length was computed over ten trials and then a histogram of displacements over the set of patients was constructed. Figure 13 shows this distribution for 229 mm lines using from data of 16 human patients (left) and 240 simulated patients (right). In both the human and simulation data, the distribution of responses is roughly unimodal.[4]

The correlation coefficient between line length and standard deviation of transection displacements was computed for each simulated patient. Line lengths of 25 and 51 mm were included in this simulations, because they were also used to compute the human patient correlations.

A mean correlation of .58 was obtained for the 240 simulated patients. This is perhaps lower than

[4]This pattern is clearer for the simulation data, as there are fifteen times as many patients composing the distribution. Note also that for the human data, patients were classified as normal if their displacements were less than 10 mm, causing the truncation of the distribution at the left end.

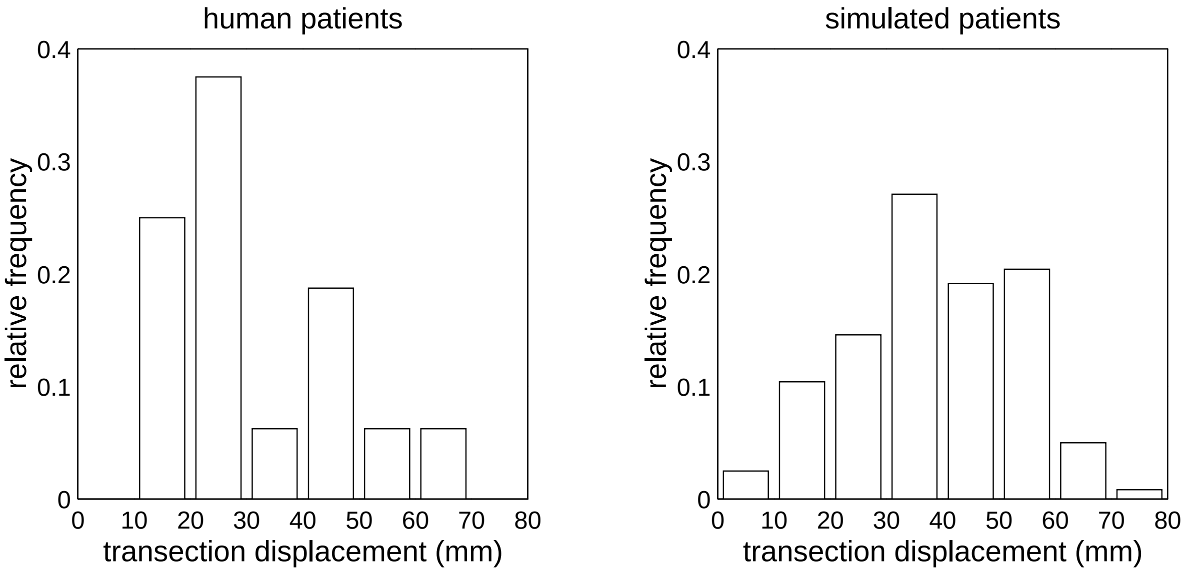

Figure 13: Distribution of mean transection displacements for 16 human patients (left) and 240 simulated patients (right) for 229 mm lines, averaged over ten trials per patient.

the correlations reported for humans (in the range .55–.95), but clearly nonzero and of the correct sign.[5]

5. Response variation as a function of displacement

In the last section, we characterized the variability of responses as a function of the length of lines. One can also consider the variability of responses as a function of mean displacement, for a given line length. Roughly, the standard deviation of transection displacements is proportional to the displacement: If a subject accurately bisects a line on average, then the variation from one attempt to the next will tend to be small; however, if the subject makes large rightward displacements, then the variability is generally larger. To quantify this effect, we reanalyzed the data of twenty patients transecting 180 mm lines, with ten trials per patient. The correlation between mean displacement and standard deviation of the displacements was .51.

[5]One parameter of the model that affects this correlation is the settling criterion—the total change in the AM activity pattern must drop below this criterion in order for a response to be read out. In the simulations reported here, a criterion of .0001 was used. If this criterion is increased to .001, the correlation jumps to .75. This is because increasing the criterion terminates the settling process sooner, and the AM response tends to be more dependent on stimulus conditions (e.g., line length) earlier in the settling process. Of course, it is not valid to adjust parameter settings differently for each simulation, and if all other simulations were run again with a settling criterion of .001, other results would no doubt change. We mention this experiment only to indicate that the magnitude of the effect is sensitive to changes in minor parameter values, not to essential properties of the model.

Simulations

For a particular line length, we examined the relation between a simulated patient's mean transection displacement and the variability of the responses that constituted this mean. While the human patients showed a positive correlation, the simulated patients showed exactly the opposite pattern. For example, with 177 mm lines, the correlation was –.84, indicating that patients who produced larger displacements tended to show *less* variability in their responses. MORSEL thus fails to model this aspect of the data.

While it would be extremely difficult to point to a specific component of MORSEL that is responsible for this failure, it is possible to describe a general characteristic that the model appears to possess, and how this characteristic is responsible for producing the observed simulation data.

Neglect in MORSEL corresponds to the failure of features registered on the retina to be detected by the AM, specifically features of the left side of the stimulus. The severity of neglect will be related to the degree to which the AM fails to detect the presence of these features. This is depicted schematically in the two graphs of Figure 14. Both graphs show the expected transection displacement as a function of the proportion of features on the left side of the retina that are successfully transmitted to the AM. When all features are transmitted, there will be no neglect, and the transection displacement will be zero; when the AM fails to detect features on the left side of the retina, transection displacements will be large. Thus, there should be a monotonic relationship between feature transmission probability and transection displacements. However, the relationship can be either positively or negatively accelerated, corresponding to the left and right curves in the figure, respectively.

For a particular lesion to MORSEL, the exact set and number of features transmitted to the AM on any trial will vary. This is because feature transmission operates probabilistically (see Figure 8). Consequently, a particular lesion will produce a *range* of different values along the x-axis of the curves in Figure 14, as depicted by the two dotted vertical lines. This range in feature transmission rates will result in a range of transmission displacements, as depicted by the two dotted horizontal lines. One can clearly see that with the curve on the left, larger transection displacements will occur with smaller variability, whereas with the curve on the right, larger transection displacements will occur with larger variability. The curve on the left corresponds to the behavior of the model; the curve on the right is what is required to fit the patient data.

While these curves are descriptive, translating them into a computation mechanism is not simple. The curve on the left emerges from the dynamics of the AM.[6] The difference between the curves—the sign of their second derivatives—is subtle indeed, and it will be a challenge to develop a revised version of the AM dynamics that achieves the desired behavior.

[6]It further does not give the whole story, because it says nothing about why shorter lines result in smaller displacements. We could simply relabel the y-axis "*relative* transection displacement" to deal with this issue, without influencing the discussion above.

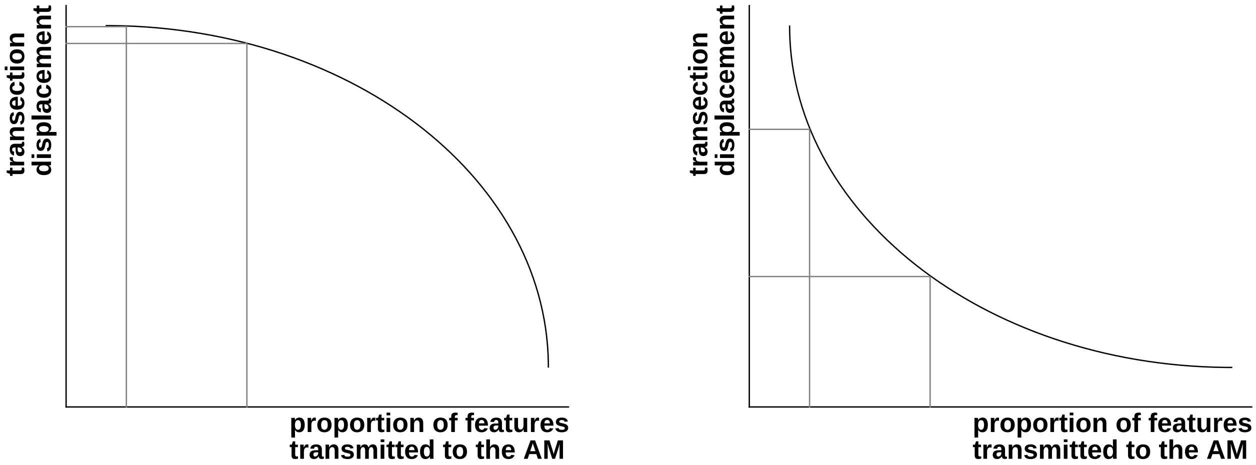

Figure 14: Two possible curves relating the expected transection displacement as a function of the proportion of features on the left side of the retina that are successfully transmitted to the AM. The larger this proportion is, then the less neglect that should be expected.

6. Effect of stimulus position

In the studies described above, stimulus lines were presented centered on a sheet of paper, which was also centered with respect to the midsaggital plane. One patient, PS, was studied not only in this condition, but also with a set of lines whose left endpoints were anchored 1.5 mm from the left edge of the sheet of paper, and another set all of whose right endpoints were anchored 1.5 mm from the right edge of the sheet of paper (Marshall & Halligan, 1990a). Let us call these three conditions the *centered*, *left aligned*, and *right aligned* presentations. A variety of line lengths were tested, ranging from 12.5 mm to 279 mm in length. Note that the full extent of right-aligned lines 127 mm and shorter were entirely within right hemispace, and likewise for short left aligned lines lying in the left hemispace.

Figure 15a shows PS's performance for centered, left aligned, and right aligned stimuli. In all three conditions, transection displacements are linearly related to line length. The average magnitude of neglect is no greater for left-aligned than for centered stimuli; right-aligned stimuli appear to produce less neglect. The slope of the linear relationship between line length and mean displacement can also be used as a measure of the severity of neglect; the larger the slope, the larger the displacement as a fraction of line length. The slope was .406 for centered, .369 for left aligned, and .299 for right aligned presentations. All three of these slopes are quite large relative to slopes obtained for normals, which are in the neighborhood of .03. Although there appears to be somewhat less neglect for the right aligned presentations by this measure, one cannot ascertain the statistical reliability of this conclusion from the one patient's data.[7]

If neglect occurs with respect to a reference frame that is centered on the sheet of paper or centered on the body midline, one would expect relatively severe neglect for left-aligned stimuli and relatively mild neglect for right-aligned stimuli. The support for this pattern of results is not strong. Instead,

[7]Indeed, as a rough test, the slopes of the lines formed by consecutive pairs of data points in Figure 15a were compared for the centered and right-aligned conditions, yet no reliable difference was found ($F(1,7) < 1$).

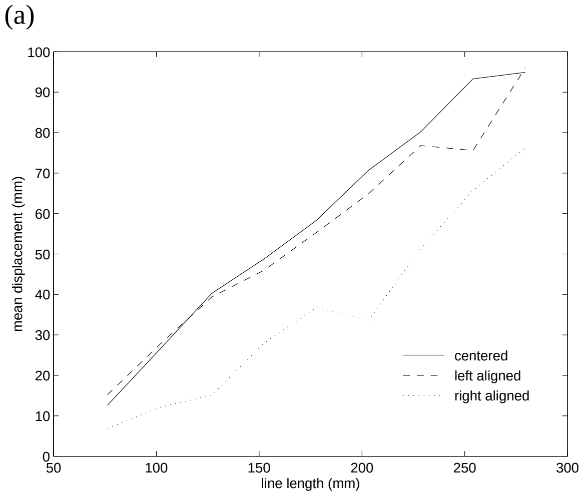

Figure 15: (a) Mean transection displacement as a function of line length for centered, left-aligned, and right-aligned stimulus presentations for patient PS, from Marshall & Halligan (1990). (b) Mean transection displacement as a function of line length for centered, left-aligned, and right-aligned stimulus presentations for one simulated patient. The transmission probability curve of this patient is characterized by a saturation probability of .9, a saturation position at the right edge of the retina, a gradient of .02, and a minimum transmission probability of .4. Other simulated patients with these parameters showed a similar pattern of results, indicating a systematic relationship between the parameters and the observed performance.

the data appears to provide evidence for neglect that occurs with respect to the reference frame of the stimulus itself.

Simulations

In the simulation experiments reported thus far, stimulus lines were centered on MORSEL's retina. Here, we add two conditions: lines whose endpoints are anchored either to the left or to the right side of the retina. Short lines thus lie entirely in one hemifield or the other, and the longer the line, the farther it extends into the opposite hemifield.

Figure 16 shows the transection displacement as a function of line length averaged over the 240 simulated patients, for centered, left aligned, and right aligned conditions. The centered-condition data are the same as those in Figure 12b. The clear result is that the placement of the lines makes little difference in the pattern of neglect. Although the transmission probability varies as a function of absolute position of the stimulus, what appears to be critical to the AM's performance is the *relative* strength of the left and right extremes of a stimulus line. Thus, the AM produces object-centered neglect—neglect with respect to the left and right ends of the object—rather than retinotopic neglect—neglect with respect to the position on the retina. This behavior of the model is somewhat unintuitive because the deficit itself is retinotopic. Mozer and Behrmann (1990) similarly found that the retinal position of a stimulus word in the damaged MORSEL had a relatively minor influence on reading performance.

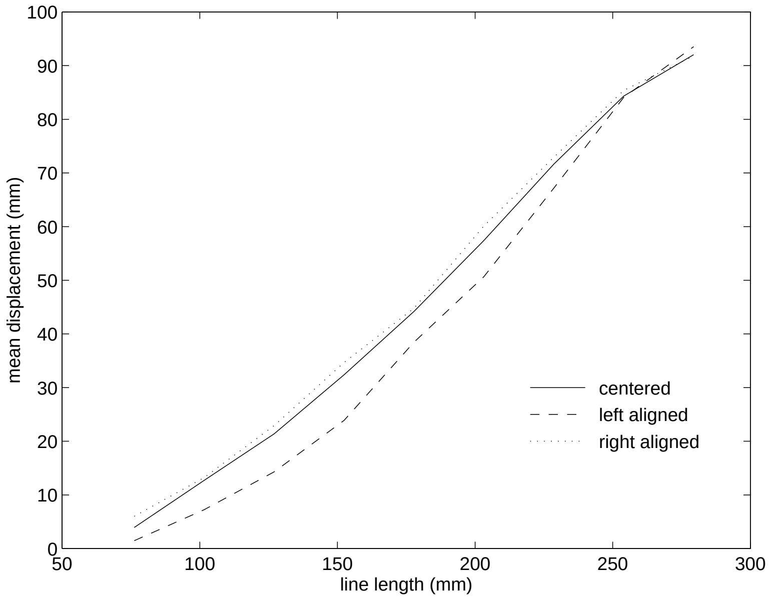

Figure 16: Mean transection displacement produced by simulations of lesioned MORSEL as a function of line length for centered, left-aligned, and right-aligned stimulus presentations.

Table 2: Linear regression coefficient for different line placement conditions

	left aligned	*centered*	*right aligned*
patient PS	.369	.406	.299
simulated patients	.479	.455	.449

This simulation result is consistent with the data obtained from PS, the only human patient who has been studied in these conditions (Figure 15a) in two important respects. First, there is severe neglect in all three conditions. Second, the slope of the functions in all three conditions are of the same order of magnitude, and are an order of magnitude larger than those obtained for normals. Because the simulation data represents a mean over 240 patients, whereas the human data are those of a single patient, a precise fit should not be expected. However, it is not difficult to find individual simulated patients whose data shows a remarkably similar pattern to that of PS (Figure 15b).

7. Effect of stimulus orientation

In the studies described above, lines were presented horizontally on the page and horizontally with respect to the patient. Marshall and Halligan (1990b) and Burnett-Stuart, Halligan, and Marshall (1991) have explored performance as a function of stimulus orientation. Patients were presented with stimulus lines of a fixed length (180 mm) drawn at various orientations on the paper.

Figure 17a shows the mean transection displacement as a function of orientation for one patient for whom a large corpus of data was collected. Positive displacements are to the right of center, regardless of orientation. The angle 0° is horizontal; the angles 90° and −90° are vertical.[8] This

[8]Note that x° lines mirrored around the x-axis become $-x^{\circ}$ lines.

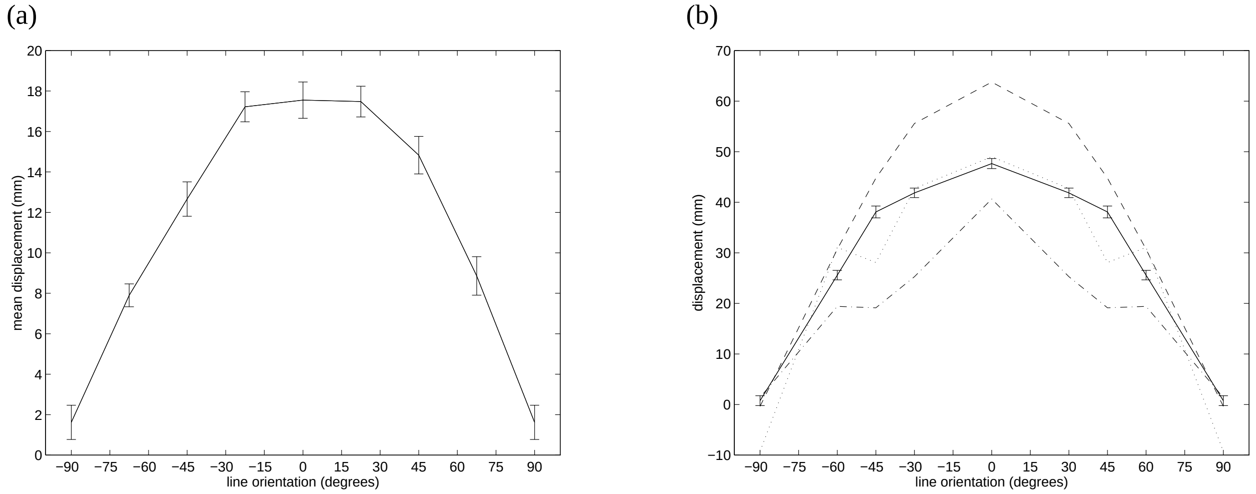

Figure 17: (a) Mean transection displacement as a function of line orientation for patient MN, from Marshall & Halligan (1990). Each point is the average over 60 trials. The vertical bars indicate +/− 1 standard error of the mean. (b) Simulation of lesioned MORSEL on line bisection task as a function of line orientation. The dotted and dashed lines show curves from three simulated patients. The solid line with the error bars shows the mean across simulated patients.

particular patient shows relatively mild neglect; his transection displacements for horizontal lines are about one third as large as are typically observed. Nonethless, it is worth presenting this data because it clearly shows the pattern of results confirmed with statistical analyses: transection displacements are linearly related to the cosine of the orientation. That is, the severity of neglect is determined by the length of the projection of the line onto the horizontal axis (Burnett-Stuart et al., 1991).

Let us formalize this characterization of the results. The data obviously cannot be modeled by a straight line; it appears that a quadratic or cosine fit is necessary. Nonetheless, one can ask whether the data can be modeled by a piecewise linear function. That is, suppose we examine the data for orientations 0–90°. In Burnett-Stuart et al.'s study of six patients, the fraction of the variance accounted for by a linear fit of these data is only 74.1%, whereas a cosine fit accounts for 88.8% of the variance.

Interpreting the data is complicated somewhat by the fact that for some patients, neglect is not entirely a function of lateral extent, but also altitudinal extent (Butter et al., 1989; Rapcsak, Cimino, & Heilman, 1988). For these patients, one can conceptualize the axis of the attentional deficit as lying not along the horizontal, but rotated slightly clockwise (resulting in both left and upper field neglect) or slightly counterclockwise (resulting in both left and lower field neglect). Accordingly, for oriented lines, neglect is a function of the cosine of the angle, but there is a constant phase shift that reflects the fundamental axis of the deficit.[9]

[9] The cosine regression computed by Burnett-Stuart et al. includes a phase shift parameter to model the altitudinal component of neglect.

Table 3: Percent of variance accounted for in oriented line simulation

	linear regression	*cosine regression*
human patients (n=6)	74%	89%
simulated patients (n=240)	65%	72%

Simulations

Stimuli used in this simulation experiment were fixed length lines, corresponding to 180 mm lines viewed by human patients, at five orientations: 0°, 30°, 45°, 60°, and 90° (Figure 7). Because of quantization effects on MORSEL's retina, it was tricky to generate the non-horizontal and non-vertical lines. We matched the diagonal lines on the total number of input features and the average number of neighbors for each feature, which are the factors that seem to most affect the AM response.

Each of the 240 simulated patients was presented with ten trials of each of the five different line orientations, and the mean transection displacement at each orientation was computed. Figure 17b shows the transection displacement as a function of line orientation. The dotted and dashed curves show typical responses of individual simulated patients, and the solid curve is the mean of the group, with error bars indicating the spread of the distribution.

Because the response of MORSEL for $-x°$ lines must necessarily mirror the response for $x°$ lines, the Figure has been extended to include the orientations between –90° and 0°. Clearly, the mean displacement curve is bowed, nonmonotonic in line orientation, and has the general shape of a cosine function.

For each simulated patient, we performed a linear and cosine regression determine which type of curve best fit the individual's data for line orientations between 0° and 90°. The mean percent of variance accounted for in the two regression analyses is presented in Table 3. As it did for the human patients, the cosine regression yielded a reliably better fit for the simulated patients ($F(1, 239) = 74, p < .001$). The explanation for this finding is simple: As we observed earlier, the critical factor in determining the severity of neglect in MORSEL is the difference in the transmission probability between the left and right ends of the stimuli. For oriented lines of a fixed length, this difference is monotonically related to the length of the projection of the stimulus line onto the horizontal axis, which in turn is proportional to the cosine of the line orientation.

One discrepancy between simulations and humans is that the simulated patients seem to show somewhat more variance in their responses than do the humans, which leads to slightly worse fits for both the linear and cosine regressions, as well as a weaker difference between the two. A partial explanation for this discrepancy may have to do with quantization effects on MORSEL's retina. In our preliminary simulation studies, we found that the responses of MORSEL were sensitive to small changes in the stimulus-induced activity pattern on the retina. A finer quantization of the retinal array should make the activity of any individual cell less critical and might reduce response variability. Another partial explanation for the discrepancy is the fact that the human experiments used 22.5° and 67.5° lines whereas the simulation used 30° and 60° lines. It was much simpler to construct 30° and 60° lines on MORSEL's retina due to quantization effects, but this created a uniform spacing of orientations, which may have masked differences between linear and cosine regressions.

We could model the performance of patients whose deficit was not strictly along the horizontal axis. This would be easy to do by rotating the attentional gradient on the AM's map, corresponding to the sort of gradient that we suppose for these patients.

DISCUSSION

MORSEL is an existing computational model whose architecture, dynamics, and parameters have been specified in earlier simulation studies. Even the form of the damage that occurs in unilateral neglect is dictated by earlier modeling work. Thus, MORSEL was not in any way designed to account for the line bisection data modeled here. MORSEL makes strong predictions and the data were reexamined to determine whether the model and data were consistent.

We have shown that MORSEL can achieve a remarkable fit to the data. MORSEL can model the following phenomena: (1) no consistent across-subject bias is found in normals; (2) transection displacements are proportional to line length in neglect patients; (3) variability of displacements are proportional to line length, in both normals and patients; (4) position of the lines with respect to the body or the page on which they are drawn has little effect; (5) for lines drawn at different orientations, displacements are proportional to the cosine of the orientation angle. MORSEL also fails to account for one observation: across patients, the variability of displacements for a particular line length is roughly proportional to mean displacement.

What have we learned from this modeling exercise? It has been valuable for a variety of reasons, on which we elaborate.

- In determining what data to model, we were forced to reexamine and reanalyze the patient data, combining results across multiple studies, separating robust and reliable patterns from quirks in individual patient performance. Prior to this effort, the main phenomena related to line bisection had never been distilled and summarized. We thus have a better understanding of the data.

- MORSEL is able to provide a comprehensive account of the varied corpus of data in a unified framework. We have pointed to essential properties of the model that are responsible for the observed effects. Phenomena that could be seen as puzzling without a theoretical framework—such as the finding that neglect is related to the total line length, despite the fact that the patient does not apprehend the entire line—have straightforward, mechanical explanations in terms of the model.

- MORSEL has been able to explain a large collection of new data with almost no additional assumptions or mechanisms. This increases our confidence in the correctness and importance of the model, and gives us further motivation to continue expanding the breadth and coverage of MORSEL.

- The one phenomenon that MORSEL mispredicts—the relation between the mean and variability of transection displacements—points to a possible problem with the model. We suspect that fixing this problem involves a fairly minor change to the dynamics of the AM, but it is difficult to determine at this point; making any change to the model requires verification that the change does not affect the model's behavior on any other data. However, we are confident that the change can be made without repercussion, as no other data previously modeled depends on the aspect of the AM that we have pointed to as the culprit of the misprediction. Indeed, whether or not the effort of fixing the model is warranted depends on whether we expect to run across data in the future that will require this correction. At this point, we suspect not.

- MORSEL provides an accurate portrayal of the average performance of groups of patients, yet it also possesses the flexibility—via the transmission probability curve that specifies the nature of the attentional deficit—to model individual differences. If we have confidence in the model's validity, we can use it prescriptively to *characterize the attentional deficit* of individual patients in terms of their transmission probability curves. We have done this in one case above to show that the stimulus position effects observed for patient PS (left aligned and centered stimuli producing comparable neglect, while right aligned stimuli producing less) could be modeled by the selection of a particular transmission probability curve. We can do a similar sort of exploration in the model's parameter space to try to understand the behavior of the only patient of forty three studied who shows a significant quadratic component in his performance on varying-length lines. Thus, MORSEL could well be used as a diagnostic tool to characterize the specific nature of a patient's deficit.

  A related and important application of MORSEL is in characterizing the sort of attentional deficits that are possible following brain damage. Exploring the model's parameter space, we have found that certain transmission probability curves yield patterns of data that are quite bizarre and have never been observed in patients, and thus are unlikely to occur in nature. Figure 18 shows two examples of this which result from shallow attentional gradients lying over just a small portion of the visual field. If we dichotomize the space of transmission probability curves into those that are and are not consistent with the human data, and if we have information about the site of a patient's lesion, we may gain insight into the underlying cortical organization and neural structures that give rise to certain types of deficits but not others.

- In the course of our modeling efforts, MORSEL has suggested experiments that should be run to help settle theoretical issues and to test additional predictions of the model. We are currently exploring the literature to determine if the relevant experiments have been conducted, and if not, we plan to conduct them in the future. The experiments will, for example, help disentangle perceptual and motor components of neglect, explore line thickness as a manipulation in bisection studies, explore the effects of brief, masked stimulus presentations, and collect more data on the effect of stimulus position and different frames of reference. For the masking studies, MORSEL makes the surprising prediction that under some conditions masking could alleviate neglect. In the last set of studies, MORSEL makes the strong prediction that, averaged over many patients, stimulus position should have little effect on performance, although there may be considerable individual differences. Assuming that these predictions are borne out, this may not be the end of the line for MORSEL after all!

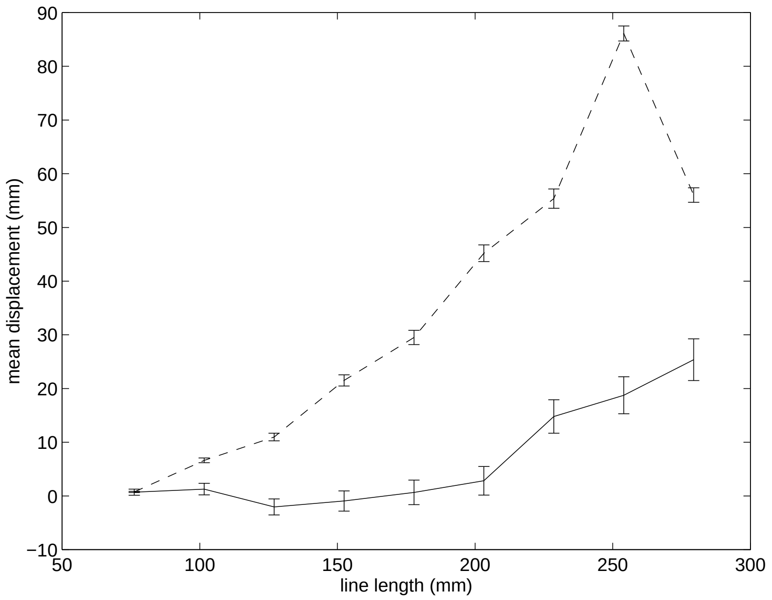

Figure 18: Simulation of lesioned MORSEL on line bisection task as a function of line length. The two lines show performance for two different transmission probability curves. The solid line—which corresponds to a saturation probability of .9, gradient of .01, saturation position on the right edge of the retina, and a minimum probability of .8—shows no neglect except for the longest lines. The dashed line—which corresponds to a saturation probability of 1.0, a gradient of .01, saturation position in the center of the retina, and a minimum probability of .8—shows proportionately less neglect for the longest line length than for the next shorter line length. Neither of these patterns are observed in the human patient data, although a much less extreme version of the dashed line is occasionally found.

APPENDIX 1: AM DYNAMICS

The task 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 a *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 behavior 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 AMbe set in order to best satisfy a number of possibly conflicting constraints? The two primary constraints here are that the AMshould focus on locations suggested by the bottom-up and top-down inputs, and the AMshould 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 *steepest 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 used in the AM is:

$$H = \sum_{\substack{(x,y)\in \\ \text{ALL}}} ext_{xy}\, a_{xy} \;-\; \frac{\mu}{4} \sum_{\substack{(x,y)\in \\ \text{ALL}}} \; \sum_{\substack{(i,j)\in \\ \text{NEIGH}_{xy}}} (a_{ij} - a_{xy})^2 \;+\; \frac{\theta}{2} \sum_{\substack{(x,y)\in \\ \text{ACTIVE}}} (\gamma \bar{a} - a_{xy})^2 \, ,$$

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 eight 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—

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

[10] We should note that many other Harmony functions would suffice equally well if not better than the one presented here. Mozer and Behrmann (1990) experimented with several different functions, and the qualitative system behavior was unaffected by the details of the Harmony function.

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 *gamma* serves as a discounting factor: with *gamma* 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 update rule for a_{xy} is:

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

Further, a_{xy} is prevented from going outside the range [0,1] by capping activity at these limits.[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 $\frac{1}{8}$, θ at $\frac{1}{2}$, and γ at .11 times the total external input.

Acknowledgements

This research was supported by NSF Presidential Young Investigator award IRI–9058450 and grant 90–21 from the James S. McDonnell Foundation to MM, grants from the British Medical Research Council to PH and JM, and an Oxford McDonnell-Pew Centre Travel Grant. Our gracious thanks to Tim Shallice and Alphonso Caramazza for setting up the Trieste meeting which served as the catalyst for this work, and to Lesley Court for facilitating our communications across the Atlantic.

References

Axenfeld, D. (1915). Hemianopische Gesichtsfeldstroungen nach Schadelschussen. *Klin. Monatsble. Angerheilkd.*, *55*, 126–143.

[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.

Barbieri, C., & De Renzi, E. (1989). Patterns of neglect dissociation. *Behavioural Neurology*, *2*, 13–24.

Behrmann, M., Moscovitch, M., Black, S. E., & Mozer, M. C. (1990). Perceptual and conceptual mechanisms in neglect dyslexia: Two contrasting case studies. *Brain*, *113*, 1163–1183.

Bisiach, E., Bulgarelli, C., Sterzi, R., & Vallar, G. (1983). Line bisection and cognitive plasticity of unilateral neglect of space. *Brain and Cognition*, *2*, 32–38.

Bisiach, E., Capitani, E., Colombo, A., & Spinnler, H. (1976). Halving a horizontal segment: A study on hemisphere-damaged patients with focal cerebral lesions. *Achives Suisses de Neurologie, Neurochirgurie et de Psychiatrie*, *118*, 199–206.

Bradshaw, J. L., Nettleton, N. C., Nathan, G., & WIlson, L. E. (1985). Bisecting rods and lines: Effects of horzontal and vertical posture on left side underestimation by normal subjects. *Neuropsychologia*, *23*, 421–436.

Burnett-Stuart, G., Halligan, P. W., & Marshall, J. C. (1991). A Newtonian model of perceptual distortion in visuo-spatial neglect. *NeuroReport*, *2*, 255–257.

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., Evans, J., Kirsch, N., & Kewans, D. (1989). Altitudinal neglect following traumatic brain injury; a case report. *Cortex*, *25*, 135–146.

Halligan, P. W., Manning, L., & Marshall, J. C. (1991). Hemispheric activation versus spatio-motor cueing in visual neglect: a case study. *Neuropsychologia*, *29*, 165–175.

Halligan, P. W., & Marshall, J. C. (1988). How long is a piece of string? A study of line bisection in a case of visual neglect. *Cortex*, *24*, 321–328.

Halligan, P. W., & Marshall, J. C. (1989a). Laterality of motor response in visuo-spatial neglect: a case study. *Neuropsychologia*, *27*, 1301–1307.

Halligan, P. W., & Marshall, J. C. (1989b). Perceptuo-cueing and perceptuo-motor compatibility in visuospatial neglect: a single case study. *Cognitive Neuropsychology*, *6*, 423–435.

Halligan, P. W., & Marshall, J. C. (1989c). Line bisection in visuo-spatial neglect: Disproof of a conjecture. *Cortex*, *25*, 517–522.

Halligan, P. W., & Marshall, J. C. (1991a). Left neglect for near but not far space in man. *Nature*, *350*, 498–500.

Halligan, P. W., & Marshall, J. C. (1991b). Recovery and regression in visuo-spatial neglect: a case study of learning in line bisection. *Brain Injury*, *5*, 23–31.

Halligan, P. W., & Marshall, J. C. (1993a). The history and clinical presentation of visual neglect. In I. H. Robertson & J. C. Marshall (Eds.), *Unilateral neglect: Clinical and experimental studies*. London: Lawrence Erlbaum.

Halligan, P. W., & Marshall, J. C. (1993b). The bisection of horizontal and radial lines: A case study of normal controls and ten patients with left visuo-spatial neglect. *International Journal of Neuroscience*, *70*, 149–169.

Heilman, K. M., Bowers, D., & Watson, R. T. (1984). Pseudo-neglect in a patient with partial callosal disconnection. *Brain*, *107*, 519–532.

Heilman, K. M., Watson, R. T., & Valenstein, E. (1993). Neglect and related disorders. In K. M. Heilman & E. Valenstein (Eds.), *Clinical neuropsychology* 2nd ed. (pp. 279–336). New York: Oxford University Press.

Kinsbourne, M. (1993). Orientational bias model of unilateral neglect: Evidence from attentional gradients within hemispace. In I. H. Robertson & J. C. Marshall (Eds.), *Unilateral neglect: Clinical and experimental studies.* Hove, UK: Lawrence Erlbaum.

Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology*, *4*, 219–227.

LaBerge, D., & Brown, V. (1989). Theory of attentional operations in shape identification. *Psychological Review*, *96*, 101–124.

Manning, L., Halligan, P. W., & Marshall, J. C. (1990). Individual variation in line bisection: A study of normal subjects with application to the interpretation of visual neglect. *Neuropsychologia*, *28*, 647–655.

Marshall, J. C., & Halligan, P. W. (1989). When right goes left: An investigation of line bisection in a case of visual neglect. *Cortex*, *25*, 503–515.

Marshall, J. C., & Halligan, P. W. (1990a). Line bisection in a case of visual neglect: Psychophysical studies with implications for theory. *Cognitive Neuropsychology*, *7*, 107–130.

Marshall, J. C., & Halligan, P. W. (1990b). The psychophysics of visuo-spatial neglect: a new orientation. *Medical Science Research*, *18*, 429–430.

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

Mozer, M. C. (1991). *The perception of multiple objects: A connectionist approach.* Cambridge, MA: MIT Press/Bradford Books.

Mozer, M. C., & Behrmann, M. (1990). On the interaction of selective attention and lexical knowledge: A connectionist account of neglect dyslexia. *Cognitive Neuroscience*, *2*, 96–123.

Nichelli, P., Rinaldi, M., & Cubelli, R. (1989). Selective spatial attention and length representation in normal subjects and in patients with unilateral spatial neglect. *Brain and Cognition*, *9*, 57.

Poppelreuter, W. (1917). Die Storungan der Niederan und hoheren Seheistungen durch Nerletzung des Okzipitalhirns. In *Die Psychischen Schadigungen durch Kopfschuss in Kreige (1914/1916). Vol. 1.* Leipzig: Voss.

Rapcsak, S. Z., Cimino, C. R., & Heilman, K. M. (1988). Altitudinal neglect. *Neurology*, *38*, 277–281.

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

Schenkenberg, T., Bradford, D. C., & Ajax, E. T. (1980). Line bisection and unilateral visual neglect in patients with neurological impairment. *Neurology*, *31*, 509–517.

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.

Tegner, R., & Levander, H. (1991). The influence of stimulus properties on visual neglect. *Journal of Neurology, Neurosurgery, and Psychiatry*, *54*, 882–887.