COGNITIVE NEUROPSYCHOLOGY, 1990, 7 (5/6) 391–445

Levels of Representation, Co-ordinate Frames, and Unilateral Neglect

Alfonso Caramazza
Cognitive Science Center, The Johns Hopkins University, Baltimore, U.S.A.

Argye E. Hillis
HealthSouth Rehabilitation Corp., Baltimore, U.S.A.

We describe the performance of a brain-damaged subject, NG, who made reading errors only on the right half of words. This problem persisted even when the subject had demonstrated accurate recognition of the letters in a stimulus through naming all the letters. Furthermore, the spatially determined reading impairment was unaffected by topographic transformations of stimuli: identical performance was obtained for stimuli presented in horizontal, vertical, and mirror-reversed form. The same pattern of errors was also obtained in all forms of spelling tasks: written spelling, oral spelling, and backward oral spelling. The performance of the subject is interpreted in the context of a multi-stage model of the word recognition process. It is concluded that the locus of the deficit responsible for NG's reading impairment is at a stage of processing where word-centred grapheme representations are computed. The spatially determined pattern of performance reported for NG, as well as other patterns observed for other brain-damaged subjects, are interpreted as providing support for the proposed multi-stage model of word recognition. The more general implications of the reported results for models of visual processing and attention are also considered.

Requests for reprints should be addressed to Alfonso Caramazza, Cognitive Science Centre, The Johns Hopkins University, Baltimore, MD 21218, U.S.A.

The research reported here was supported in part by N.I.H. grant NS22201 and by grants from the Seaver Institute and the McDonnell/Pew Program in Cognitive Neuroscience. This support is gratefully acknowledged. We would especially like to thank NG for her patience and good cheer throughout our interminable experiments and probings. We are grateful to the members of the Cognitive Neuropsychology Laboratory, and in particular Brenda Rapp, for their helpful suggestions at various points of the research reported here. We thank Howard Egeth, Glyn Humphreys, Brenda Rapp, Jane Riddoch, Tim Shallice, Eric Siéroff, Paolo Viviani, and Steven Yantis for comments on an earlier version of this paper.

INTRODUCTION

In this paper we investigate several aspects of the reading process that are illuminated by consideration of the reading performance of patients presenting with a striking perceptual/cognitive disorder that disproportionately affects a spatially defined part of visually presented objects—unilateral visual neglect. We will argue that the clinical category of unilateral neglect consists of a heterogeneous set of patients with deficits at different levels of the perceptual system. One level of deficit involves damage to the perceptual mechanisms that compute retino-centric representations of visual inputs; another level of deficit involves damage to mechanisms that compute a viewer- or stimulus-centred representation of the font-specific and orientation-specific letter shapes, and a third level of deficit involves damage to mechanisms that compute word-centred (or object-centred) representations of the abstract letter identities that comprise a word. This claim about the various types of possible impairments is based on strong assumptions about the types and structure of representations that are computed in the course of recognising a string of letters as a word of the language. Our discussion of these issues is organised as follows: we begin by identifying the relevant computational problems that must be solved in reading—in the specific case considered here, the computational problems associated with the recognition of a string of letters as a word of the language; we then describe experimental results obtained with a brain-damaged subject which severely constrain plausible claims about the structure of the representations and the processes that underlie the recognition of written words; we conclude with a discussion of various problems and pseudoproblems that have arisen in the context of considerations of reading and unilateral neglect.

A Computational-level Analysis of Written Word Recognition: Levels of Representation and Co-ordinate Frames

Within an information processing framework, the computational goal in word recognition consists of determining the types and structure of representations that are computed in the course of mapping a visual stimulus onto a lexical-orthographic representation. We take it that this process is not dissimilar in its general form to that involved in visual object recognition. Following Marr (1982; Marr & Nishihara, 1978), we assume that the latter process entails computing several different types of representations, prior to actual recognition: beginning with a description of the visual array in terms of perceptual primitives, blobs, edges, and bars—the primal sketch; proceeding to a description of the visible surfaces of objects in terms of local surface orientation and distance—the 2½-D sketch; and, finally, an abstract, canonical description of the object—the 3-D model.[1] In this framework, the descriptions computed at each level of processing differ not only in terms of the types of representations that are computed, but also in terms of the co-ordinate system within which the respective representations are defined. Our working hypothesis is that the early stages of reading, up to word identification, are similarly organised (see Monk, 1985, for a discussion of this parallel). And, in fact, in the absence of evidence to the contrary, we will assume that the representational spaces involved in visual word and object recognition are the same. On this view, the visual processes involved in word recognition can be characterised in terms of three stages of analysis: (1) the computation of a retino-centric feature map; (2) the computation of a viewer- or stimulus-centred letter shape map; and (3) the computation of a word-centred grapheme description.[2,3] The grapheme descriptions computed at the latter level of analysis serve to activate lexical-orthographic representations for word identification (see Fig. 1).

The multi-stage model of word recognition adopted here makes a number of specific assumptions about the representations computed at

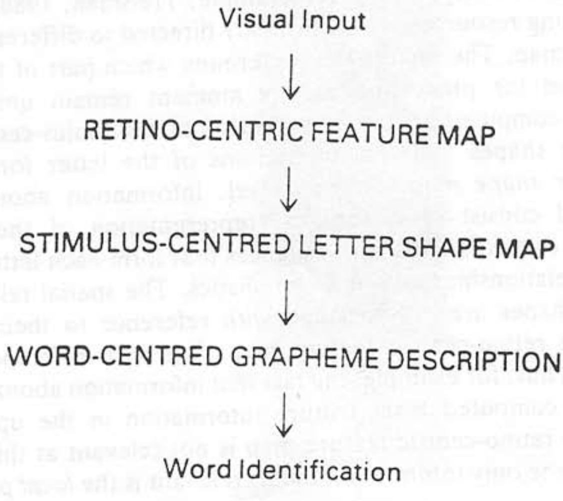

FIG. 1 Levels of representation in visual word recognition.

[1]See Hildreth and Ullman (1989) and Kosslyn, Flynn, Amsterdam, and Wang (1990) for recent accounts of the computational problems in visual perception.

[2]We follow Venezsky (1970) and Cummings (1988) in using the term *grapheme* to refer to the basic units of analysis in orthographic representations—abstract (case- and font-independent) letter identities.

[3]It should be noted that the term "word" in "word-centred co-ordinate system" is used to refer to a letter string, bounded by perceptually salient features, that only potentially forms a word in the linguistic sense. More on this follows.

each level of the recognition process. The major assumptions may be rendered explicit by considering the computational problem that must be solved at each of the hypothesised levels of processing. At the first level, the problem is that of extracting directly from surface-reflected light intensities the relevant discontinuities that define edges in the image (within the limits of visual acuity). It is assumed that this process is spatially parallel across the entire visual field, in the sense that edge information may be computed simultaneously at every location of the retinal image. The representation computed at this stage of processing consists of a retino-centric description of the edges in a retinally projected image—a *feature map*. Thus, for example, feature information extracted from a word presented in the upper right quadrant of the visual field will be represented in the upper right quadrant at the retino-centric feature level.

The computational problem at the second level of analysis involves recovering from a feature map the shape properties of contours and the spatial relations that obtain among parts of an image (stimulus). This process is assumed to be only locally parallel, in the sense that processing is spatially and temporally inhomogeneous across the feature map computed at the first level of analysis (see, for example, Treisman, 1988). In other words, processing resources are *sequentially* directed to different subparts of the feature map. The factors that determine which part of the feature map is selected for processing at any moment remain unclear. The representation computed at this level consists of a stimulus-centred[4] description of the shapes and spatial relations of the letter forms in the image—a *letter shape map*. At this level, information about a word stimulus would consist of a veridical representation of the spatially arranged lines, curves, angles, and distances that form each letter, as well as the spatial relationships among letter shapes. The spatial relationships among letter shapes are *not* specified with reference to their absolute position in the retino-centric feature map, but with reference to the stimulus itself. Thus, for example, the fact that information about a specific letter shape is computed from feature information in the upper right quadrant of the retino-centric feature map is not relevant at this level of representation: the only information that is relevant is the *local* position of

[4] We have chosen to use the term "stimulus-centred" instead of the more common "viewer-centred" in order to emphasise the stimulus-bound aspect of the representation computed at this level of processing. That is, we wish to emphasise the fact that we can think of the several levels of processing considered here as ordered in terms of progressive abstraction from the physical stimulus: as a process that abstracts away from irrelevant physical detail to a representation of only the perceptually relevant information. However, at this point we are agnostic on the issue of whether processing resources are directed to a *stimulus* or to a *region of space that contains a stimulus*. At the level of detail we are working the distinction is not significant.

the letter shape relative to other letter shapes in the stimulus. For example, the *leftmost* letter shape of a stimulus in the *upper right* quadrant of the retino-centric feature map would be represented on the *left* of the stimulus-centred letter shape representation. In general, then, the information that is retained at this level of representation is the local topographic relationship (e.g. vertical arrangement) of the letter shapes that comprise a stimulus.

At the third level of analysis, the computational problem consists in computing from simple shape properties the abstract letter identities—case-, font-, and orientation-independent letter representations—that comprise a letter string. It is not clear whether this process should also be assumed to be locally parallel, in the sense that simultaneous processing is restricted to a spatially defined subset of the shape description, or whether it should be thought to be strictly serial for each segregated shape (for discussion of this issue see Duncan, 1987; Egeth, Jonides, & Wall, 1972; Pashler & Badgio, 1987; Treisman & Gelade, 1980; and see Allport, 1989, for a general review of the issues and results). The representation computed at this level consists of a word-centred description of the graphemes and their relative spatial position in a word—a *grapheme description*. At this level of analysis there is no difference among the representations for the stimuli *CHAIR*, *chair*, *ChAIr*, *cHAiR*, CHAIR and so forth; in each case the representation consists of the grapheme sequence [<c>, <h>, <a>, <i>, <r>].[5] It is also assumed that the computed grapheme description is "normalised" to an orientation-invariant, canonical format, with the horizontal plane as the major axis. On this assumption, the grapheme representations for the stimuli shown in Fig. 2 would in each case have the form <chair> with the grapheme <a> occupying the central position of the grapheme description, the graphemes <c> and <h> would be to the left of centre and the graphemes <i> and <r> would be to the right of centre in the horizontal plane. Figure 3 summarises the hypothesised progression of abstraction from physical parameters to a canonical grapheme representation. The representation computed at the word-centred grapheme level is used by lexical processing mechanisms in word identification.

Various proposals have been offered concerning the processing structure of the lexical access stage of word identification (see, e.g., Coltheart, 1981; McClelland & Rumelhart, 1981; Taft, 1985). The proposal we will entertain here shares important similarities with some of these, but also differs from them in important respects. For present purposes, the aspects of the proposed lexical access procedure that need to be made explicit are the following:

[5] Letters in triangular brackets indicate graphemes.

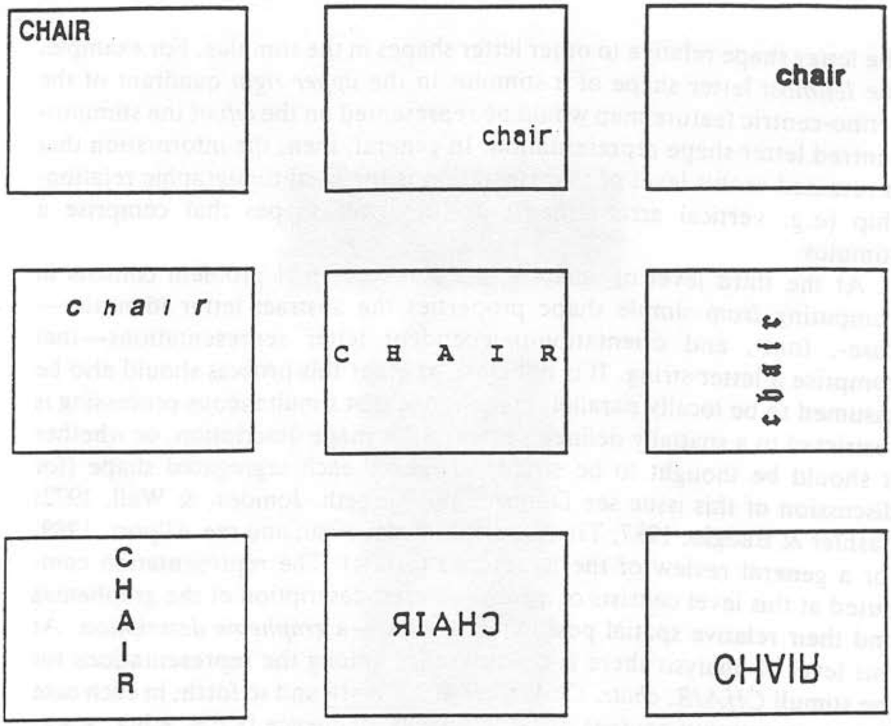

FIG. 2 Various word formats for which the graphemic description would be identical.

1. The unit of representation that serves as input to the lexical access system is a grapheme string;

2. input grapheme strings activate in parallel all access units in the lexicon that share graphemes with the input representation (e.g. the input grapheme representation <pots> activates, to varying degrees, the access units POTS, POT, PIT, NOT, TOPS, PINS, POTATO, and so forth);

3. access units in the lexicon consist of the known words *and* morphemes of the language (e.g. WALKING, WALK-, -ING, and so on);[6]

4. the activation level of an access unit in the lexicon is proportional to the degree of similarity between input grapheme string and access unit, where similarity is indexed by the degree of grapheme overlap defined in terms of both identity and relative position of graphemes (e.g. for the input

[6] Although we assume that there are both whole-word and morpheme access units in the lexical access procedure, in this report we will only refer to word access units. This choice is not intended to reflect a theoretical choice: it is adopted only for convenience. The aspects of the word identification process discussed in this report are largely unaffected by the distinction between word and morpheme access units. In the Discussion we will return to the issue of the nature of the units of analysis computed at the level of grapheme representations.

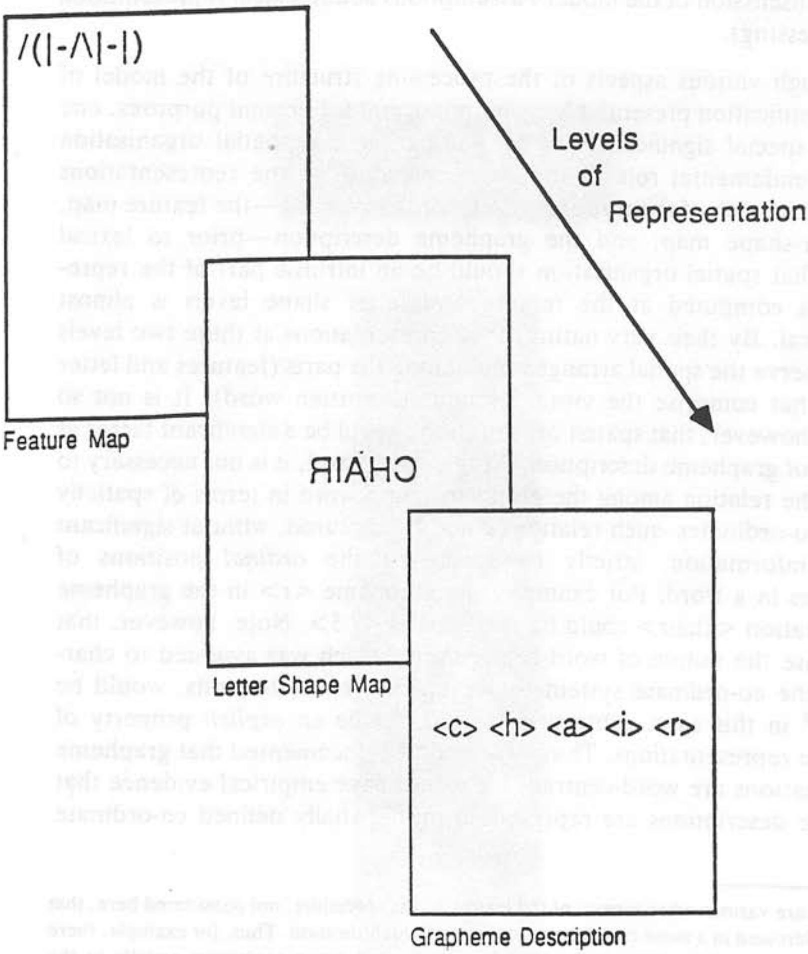

FIG. 3 Schematic illustration of levels of representation in visual word recognition for a mirror-reversed word projected to the upper right quadrant. The first level of analysis is the retino-centric feature map (upper left panel), followed by a stimulus-centred letter shape map (middle panel), and, finally, the word-centred grapheme description (lower right panel).

grapheme string <pots>, the access unit in the lexicon, POTS, would be activated more than PATS, which would be activated more than PAT, which would be activated more than STEP, and so forth; furthermore, the input grapheme string <pots> would maximally activate POTS and much less STOP, TOPS, and SPOT because even though all four access units have the same graphemes, those of the latter three do not have the same spatial arrangement/order as those of the input representation);

5. access units differ in terms of the amount of activation required to reach identification threshold—threshold settings for access units are

determined by word frequency, context, and, perhaps, other factors[7] (Morton, 1969; see Caramazza, Laudanna, & Romani, 1988, for a more detailed discussion of the model's assumptions about lexical representation and processing).

Although various aspects of the processing structure of the model of word identification presented here are important for present purposes, one point of special significance is the assumption that spatial organisation plays a fundamental role in the characterisation of the representations computed at each of the three levels of word processing—the feature map, the letter-shape map, and the grapheme description—prior to lexical access. That spatial organisation should be an intrinsic part of the representations computed at the feature and letter shape levels is almost tautological. By their very nature, the representations at these two levels must preserve the spatial arrangement among the parts (features and letter shapes) that comprise the visual stimulus (a written word). It is not so obvious, however, that spatial organisation should be a significant factor at the level of grapheme description. At this latter level, it is not necessary to capture the relation among the graphemes in a word in terms of spatially defined co-ordinates; such relations could be captured, without significant loss of information, strictly by specifying the *ordinal* positions of graphemes in a word. For example, the grapheme <r> in the grapheme representation <chair> could be specified as <r/5>. Note, however, that in this case the notion of word-centredness, which was assumed to characterise the co-ordinate system for grapheme representations, would be vacuous:[8] in this case, centredness could not be an *explicit* property of grapheme representations. Thus, if it could be documented that grapheme representations are word-centred, we would have empirical evidence that grapheme descriptions are represented in a spatially defined co-ordinate system.

[7]There are various other aspects of the lexical access procedure, not considered here, that must be addressed in a more complete model of word identification. Thus, for example, there is the issue of whether or not the graphemes in an input string contribute equally to the activation of an access unit in the lexicon: perhaps the initial letters contribute more than the final letters, or the initial and final letters could contribute more than the medial letters, and so forth. As another example, there is the issue of whether there might not be facilitory or inhibitory links among access units (see, e.g. Caramazza et al., 1988; McClelland & Rumelhart, 1981). Although these and other aspects of the lexical access procedure are important, they are not considered here because they do not affect the specific issues we wish to address.

[8]To be sure, it is possible to articulate an alternative representation format with ordinal positions for graphemes that gives the notion of word-centredness a non-vacuous reading. For example, the graphemic representation for chair could be specified as [<c/−2>, <h/−1>, <a/0>, <i/1>, <r/2>]. However, as will become apparent shortly, this solution fails to provide a principled account for the experimental results reported here.

The assumption that visual word identification involves a series of processing stages which compute different types of representations, each specified in a spatially defined co-ordinate system, constrains plausible claims about the form of impairment that could result from selective damage to any of the stages of processing. One obvious implication is that, although spatially defined impairments may be observed following damage to any of the three stages of processing, the units of representation affected would be different in the three cases—edges for the feature map, shapes for the letter shape map, and graphemes for the grapheme description. On this reasoning, the spatial specificity of a processing deficit would not, on its own, provide evidence as to the locus of damage responsible for the observed impairment. To determine the locus of damage we must have evidence about the type of representations affected.

In this paper, we describe the reading and spelling performance of a brain-damaged subject, NG, who presents with the clinical picture of unilateral neglect. Various aspects of her performance have been reported elsewhere (Hillis & Caramazza, 1990) and some of the results described here were summarized in Caramazza and Hillis (1990). In reading and spelling, NG only made errors in processing the right half of written words, regardless of the length of the word and independently of the topographic arrangement of the stimulus in reading, and the form of output in spelling. It is argued that these results, as well as others described here, cannot be explained by a deficit to low-level visual processing mechanisms. Instead, we must assume that the deficit is at the grapheme level. Furthermore, given that the impairment consistently involved the *"right" half* of words, irrespective of task type, we must assume that grapheme descriptions are represented in a word-centred co-ordinate system. The implications of these conclusions for models of word recognition are considered, as are the more general implications for claims about visual processing and object recognition and implications for the nature of unilateral neglect.

CASE REPORT

Social and Medical History

NG is a 79-year-old woman who completed 8th grade in a parochial school. She reports premorbid left-hand dominance for all tasks except writing; she was taught to use her right hand for writing at school. NG has always been active in the community, and was reportedly able to read and write well. She lived alone between the time of her husband's death and her stroke, but currently lives with her son, and attends a senior centre several days a week.

NG had a stroke in the fall of 1986, resulting in right hemiparesis of the arm and leg, which persists. There was no reduction of her visual field, nor were there signs of aphasia or dysarthria. A C.T.-scan revealed a large area of infarction in the left parietal white matter and a smaller area in the left anterior basal ganglia, adjacent to the head of the caudate. She has since had several T.I.A.'s that did not cause any lasting changes in her neurological status.

NG initially showed right "neglect" in virtually all tasks—she walked into people on her right side, failed to eat food on her right, made right-sided errors in line cancellation, copying, and matching tasks, and omitted words on the right side of the page in reading. She has since shown improvement in many of these areas. Thus, for example, she no longer has right-sided problems in eating, and no longer makes errors in simple left-to-right stimulus matching tasks. She also rarely omits whole words on the right in reading sentences or paragraphs. Despite these improvements in responding to stimuli presented in the right hemispace, her reading and spelling of individual words have not changed to any significant extent during the three years we have studied her performance. She continues to make errors in which the right-most letters are replaced with incorrect letters, both in reading single words—e.g. *park* read as "part"—and sentences—e.g. she read *The quick brown fox jumps over the lazy dog* as "The quiet brown fox jumped over the lazy doctor."

NG's pattern of spelling errors has also remained highly stable over the testing period. She has continued to make spelling errors on the right-most part of words in writing to dictation, written naming, and spontaneous writing. For example, her written name in response to a picture of a church was *churc* and to a picture of a guitar was *guiton*. She wrote *Hou much was the postal to mailed the packy* to convey "How much was the postage to mail the package".

Cognitive Evaluation

On retesting at 24 months post-stroke, NG made no errors in responding to sequential commands or repeating sentences on the Boston Diagnostic Aphasia Examination (Goodglass & Kaplan, 1972). Verbal description of the "Cookie Theft" picture was normal in all respects. On the Modified Token Test (De Renzi & Faglioni, 1978) she received a normal score of 33.5/36. Unlike performance on a previous administration, she made no errors that might be attributed to failure to attend to the tokens on the right. Performance on the Weschler Memory Scale (Weschler, 1972) was within normal limits for her age, except in visual reproduction (score = 5/15; errors on right) and mental control (errors in counting by 3's and completing the alphabet). Forward digit span was 7; backward span was 3. She showed persisting right-sided tactile and visual extinction with double simultaneous stimulation. In a line cancellation task she crossed out 40/46 lines (improved from 13/46); all errors were on the extreme right. She copied a flower normally, but omitted the right side of each figure in copying a scene (see Appendix). Impaired performance in line bisection, reading, and spelling will now be described in detail.

EXPERIMENTAL METHOD, RESULTS, AND DISCUSSION

Reading and Spelling

Methods

The data reported were collected between 4 and 24 months after NG's stroke. The most recent studies, involving mirror-reversed reading, backward oral spelling, line bisection, and identifying the effects of affixes on reading performance, took place between 18 and 24 months post-stroke.

NG was presented with a set of words in a variety of formats and tasks: normal (horizontal) oral reading, vertical reading, recognition of aural spelling, mirror-reversed reading, delayed copying, oral spelling, written spelling, and backwards oral spelling. A total of 976 words and 208 nonwords were presented individually (through a window card) for oral reading. The stimuli included 374 items from the Johns Hopkins University Dyslexia Battery (Goodman & Caramazza, 1986) and 602 items from the Johns Hopkins Morphology Battery (Badecker & Caramazza, 1987). The words in these batteries vary across various dimensions of experimental interest: grammatical class, frequency, concreteness, length, phonological transparency (for pronunciation), and affix type (prefix versus suffix, opaque versus transparent suffix, and so forth). The stimuli from the J.H.U. Dyslexia Battery were printed in large block print; the stimuli from the Morphology Battery were printed in lower case, medium font (12pt Times Roman). Neither case nor print size significantly affected NG's performance. Stimuli for the other tasks were drawn from this pool. Specifically, for "vertical reading", a subset of 300 words and 60 nonwords, printed vertically on the page in large block letters, were presented individually. For recognition of aural spelling, the same subset of 300 words and 60 nonwords were spelled aloud to her, and she was asked to name the word or nonword. Stimuli for mirror-reversed reading and written and oral spelling consisted of the same subset, plus additional items from the J.H.U. Dyslexia Battery, for a total of 626 words and 74 nonhomophonic nonwords, plus 34 pseudohomophones (e.g. *hunnee*) for reading tasks only. Stimuli for mirror-reversed reading were made by photocopying the words printed in block letters onto transparencies, and flipping the transparencies onto a white paper background.[9] Individual words were exposed through a window card. Delayed copying involved exposing an individual word or nonword briefly (from a subset of 84 words and 40 nonwords from the main list), and immediately asking her to write it. Each task was trained with as many practice items as needed until NG produced consistently appropriate, if not accurate, responses. At least 6 weeks intervened between tasks that involved the same stimuli.

Detailed descriptions of the effects of various stimulus parameters such as frequency and grammatical class on NG's reading and spelling performance, compared to performance of two other brain-damaged subjects with spatially specific reading problems and an age-matched control subject, are reported in Hillis & Caramazza (1990). Here, we will only briefly summarise the results obtained for NG, along with a qualitative analysis of her performance. Since performance in mirror-reversed reading and backward spelling have not previously been reported in any detail, we will provide somewhat more information for these tasks. The principal focus in this report, as in Caramazza & Hillis (in press), will be on the distribution of reading and spelling errors as a function of position within a word for different types of tasks.

Results: Normal (Horizontal) Reading

NG made reading errors in response to 222/976 (22.7%) words and 118/208 (56.7%) nonwords. The only stimulus parameters that were found to affect significantly her reading accuracy were word frequency and lexicality (i.e. word versus nonword status). She correctly read 141/145 (97.2%) high-frequency words compared to 123/145 (84.8%) low-frequency words matched for length and word class ($X^2_1 = 12.21$, $P < 0.001$). Furthermore, the mean frequency of correctly read words was significantly higher than the mean frequency of incorrectly read words (103.0 versus 43.4: $P < 0.001$ by 2-tailed t-test). NG correctly read 75/84 (89.3%) words versus 20/68 (29.4%) nonwords matched for length in letters. She also read functors more accurately than open-class words (100% versus 89.3%), but this effect can be accounted for by the higher frequency and perhaps shorter length of functors. Although there were no significant differences with respect to word length when this variable was tested with 14 words of each length from 4 to 8 letters, the mean length of correctly read words was significantly shorter than the mean length of incorrectly read words (6.02 letters versus 6.92: $P < 0.0001$ by 2-tailed t-test). However, the latter effect may simply reflect the fact that NG's error rate in reading was substantially higher for suffixed words than for

[9] Credit for this technique belongs to B. Wilson and K. Patterson.

unsuffixed words of the same length and surface frequency (56% errors on 85 suffixed words vs. 18% on 85 unaffixed controls; $X^2_1 = 34.51$; $P < 0.0001$). On a list of 46 suffixed, 46 unaffixed, and 46 prefixed words, her error rate was higher for suffixed words (61% errors) than for matched unaffixed words (24% errors) or prefixed words (11% errors; $X^2_2 = 28.45$; $P < 0.001$). Thus, the result that correctly read words were significantly shorter than incorrectly read words can be accounted for by the fact that suffixed words were longer and less accurately read than unaffixed words. Finally, it should be noted that the seeming discrepancy between the overall error rate of about 23% and the much smaller error rates for some word lists (e.g. the frequency list) merely reflects the fact that the latter lists did not contain suffixed words. The implications of affixing on reading performance are discussed below.

All of NG's reading errors, both for words and nonwords, involved the right part of the stimulus. Examples of these errors are reported in Table 1. As illustrated by these examples, the majority of her responses were approximately the same length (in letters) as the stimulus. The correlation between stimulus length and response length was highly significant ($r = 0.67$, $P < 0.0001$). The mean length of the stimuli was 6.94 (s.d. = 1.7) letters, and the mean length of the responses was 6.75 (s.d. = 2.0) letters.

In response to nonwords, NG tended to produce words that shared at least the initial half of the letters with the stimulus (e.g. *aftes* → "after", *faunch* → "fault". However, her 90 correct responses demonstrated that she understood the task and had no specific deficit in converting print to sound.[10]

NG showed precisely the same pattern of reading errors in two other tasks: when words were presented tachistoscopically in her *left* visual field, and even when she first named all the letters in a word correctly. For example, when instructed to name each letter before saying the word, NG

[10]On the assumption that the graphemes at the right end of the stimuli could not be processed normally, it is unclear how she read nonwords correctly. One possibility is that when information about letters on the left is either (1) sufficient to reject the string as a word, or (2) alone corresponds to a whole word, the letter string could be parsed into separate substrings, so that the corresponding graphemic descriptions could be individually re-centred for further processing. So, for example, the left half of the nonword *teybull* (*teyb*) would be sufficient to reject the stimulus as a word, so that the grapheme description be parsed into separate representations—<tey> and <bull>—that can be centred individually and processed further. In fact, she read essentially all two-syllable nonwords correctly or incorrectly as though they were composed of two short words (e.g. *teybull* → "tea-bull" (correct); *haygrid* → "hay grit", *haytrid* → "hay trial"; *mushrume* → "mush and rum") or one real word (e.g. *hunnee* → "human"). Also consistent with this notion of parsing and re-centring is the observation that NG read compound words made of two 4-letter words (e.g. *bookmark*) as accurately as 4-letter words, and much more accurately than 8-letter monomorphemic words.

TABLE 1
Examples of NG's Errors in Normal (Horizontal) Reading (Stimulus → Response)

Word Stimuli		
humid → human	hound → house	stripe → strip
sprinter → sprinkle	dumb → dump	study → stud
though → thoughts	emotionally → emotional	hazardous → hazard
Nonword Stimuli		
petch → petcher	dring → drill	stould → stoutly

read *journal* as "j-o-u-r-n-a-l . . . journey" and read *fing* as "f-i-n-g . . . fine". In reading 150 words in this fashion, she made 39 errors (26%—comparable to her error rate in normal reading), of which all but one were restricted to the end of the word. This result establishes that NG's right-sided errors cannot be explained as simply arising from a low-level visual perceptual disorder. Tachistoscopic reading was not studied systematically, but her spontaneous reading of words in a lexical decision task, in which stimuli were presented for 100msec. to the left of the fixation point (without a mask), revealed very similar types of errors (e.g. *allow* → "allot"; *dollrb* → "dollar", *pulsr* → "pulse"). In this lexical decision task, she correctly rejected all but 5% of nonwords in which letters on the left half of the stimulus violated graphotactic constraints (e.g. *fkirt*), but failed to reject any of the nonwords in which the violation occurred on the right half of the stimulus (e.g. *suggesb*). Performance was essentially the same in a lexical decision task in which nonwords were created by omitting the final (right) or initial (left) letter of words, and in which stimuli were presented for 200msec. randomly to the left, right, or centred at the fixation point (without a mask). Her spontaneous oral reading in this task revealed some word completions on the right (e.g. *prett* → "pretty"; *golde* → "golden"), but also many letter substitutions on the right (*bowle* → "bowls", *expect* → "express").

We have noted that NG read many words correctly. An important question we may ask is: do NG's correct responses reflect occasional normal processing of information at the right end of words, or do they simply reflect default responses that are correct by chance? It is difficult to answer this question directly. In order to do so we would need to have an estimate of how much information on the right part of the word NG was able to process, and we would also need an estimate of the responses that would be possible (and their relative subjective frequencies—Gordon, 1985) given that amount of information. In the absence of clear information about NG's effective response set for a given stimulus, and in the absence of a direct procedure for estimating the degree of usable information on the right part of a word that may be available to her, we must resort to indirect means in order to determine how correct responses are produced. For these purposes we relied on the following procedure.

NG was asked to read two lists of words, each comprised of 30 nouns. One of the lists consisted of 15 regular plural and 15 singular forms. The other list consisted of the corresponding singular or regular plural forms of the words in the first list. All plural forms had the suffix -s (rather than -es). It was hypothesised that if correct responses were the result of occasional normal processing of the full stimulus, then the expectation is that NG should show above-chance discrimination of the presence/absence of the plural suffix. Considering only items for which the stem was read correctly, NG was at chance level in producing the correct form of the word. She correctly identified, as indicated by her reading response, the presence/absence of the plural suffix for 46% (12/26) of the words on list 1 and 51.8% (14/27) of the words on list 2. For 28/30 items, she produced the same response to the word whether or not the stimulus was plural. Thus, she read both *dollar* and *dollars* as "dollars" and read both *fabric* and *fabrics* as "fabric". The exceptions were: *offenses* → "offensive" versus *offense* → "offend" and *planets* → planet versus *planet* → "plant". Most often, the form of the word produced by NG was the form with the higher surface frequency; thus, she made 22 deletions of the plural suffix, 6 suffix additions, and 6 substitutions of suffixes or letter sequences (e.g. *pursuit* → "pursue") on the 60 words. These results indicate that NG had virtually zero usable graphemic information at the ends of words. The conclusion invited by the results is that correct reading performance is the product of default responses that are correct by chance.

We have claimed that NG's errors virtually always involved the right end of the stimulus. A quantitative analysis of the distribution or errors as a function of letter position within a word provided the basis for a more precise characterisation of the spatial nature of the subject's reading impairment. Words of each length were scored separately. The letter of the stimulus word in each position from the left (e.g. first letter = 1; last letter of a 4-letter word = 4, last letter of a 5-letter word = 5) was scored. If the letter did not appear in the response, or was substituted with another letter, 1 error was recorded for that position (e.g. *stripe* → "strip" was scored as 1 error in position 5, and *swam* → "swan" was scored 1 error in position 4). Transposed letters were scored as 0.5 error in each position. Thus, *quite* → "quiet" was scored as 0.5 error in positions 4 and 5. When one or more letters were added on the end of the word, 1 error was scored in the last position. To illustrate, *though* read as "thoughts" was scored as 1 error in position 6.

The upper panel of Table 2 presents the distribution of reading errors as a function of position within a word, separately for words of different lengths. It is immediately apparent upon inspection that NG's errors occurred almost exclusively on the right end of words, irrespective of their length. It is equally apparent, however, that for longer stimuli, NG is able to get more letters correct (in absolute terms). Thus, for example, the fourth position of 4-letter words engendered 15% errors, whereas the fourth position of 9-letter words engendered no errors at all. More generally, it seems that the point in a word at which NG makes reading errors shifts leftward for longer words, with errors seemingly occurring only on the *right half* of words. This is, in fact, a correct description of the results, as may be seen from the lower panel in Table 2, where the percentage of errors at each letter position are arranged by reference to the centre of words for each word length. When so arranged, several facts become clear: (1) the vast predominance of errors occurred on the right half of words irrespective of their length; (2) errors occurred at equal rates as a function of absolute distance from the centre of a word; and (3), errors increased "linearly" as a function of distance from the centre of the word.

Discussion

There are several aspects of NG's reading performance that are relevant

TABLE 2
Rate of Reading Errors as a Function of Letter Position in Words of Different Lengths

Left Aligned

Word Length	N	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*
		Position in Word:								
4	141	0	0	4	15					
5	219	0	0	1	8	18				
6	204	0	0	0	4	8	25			
7	82	0	0	1	4	5	16	31		
8	88	0	0	0	1	3	18	21	34	
9	5	0	0	0	0	2	10	23	28	37

Centred

Word Length									*Word Centre* x								
4						0		0		4		15					
5					0		0		1		8		18				
6				0		0		1		4		8		25			
7			0		0		1		4		5		16		31		
8		0		0		0		1		3		18		21		34	
9	0		0		0		0		2		10		23		28		37

for understanding the processing structure of the word recognition system, and for determining the locus (or loci) of functional deficit responsible for the observed reading impairment. These include: the effects of word frequency and of lexicality on reading accuracy, the high correlation between stimulus and response length, and the spatial specificity of the impairment. However, these results on their own do not allow an unambiguous decision regarding the specific locus of deficit within the word recognition system. To be sure, the spatial specificity of the impairment restricts the possible locus of impairment to a stage of processing prior to lexical access. This constraint is not sufficiently specific, however. There are at least three stages of processing in the proposed model of word recognition—the feature map, the letter shape map, and the grapheme description—which if damaged could result in the observed pattern of impairment. However, the fact that NG continued to make right-sided errors with tachistoscopic presentation of stimuli to the left visual field rules out the feature level of representation as a possible locus of deficit. This still leaves as possible loci of deficit the letter shape and the grapheme levels. Selective damage to the right part of either of the latter two levels of representation could result in the observed difficulty in processing the right end of words. To distinguish between these alternative hypotheses we need information regarding NG's reading performance with topographically transformed stimuli or other lexical processing tasks that differentially affect the letter shape and the grapheme levels of representation. This is explored next.

Results: Reading Topographically Nonstandard Text, Spaced Letters, and Recognition of Aural Spelling

In reading topographically nonstandard text—vertically presented and mirror-reversed words—NG made precisely the same types of errors as she had in reading topographically standard (horizontal) text: errors virtually only occurred at the end of words. The majority of errors in these tasks involved substitutions of letters on the right end of words—e.g. *thousand* (presented in mirror-reversed form) was read as "thought". Some letter omissions on the right were also noted—e.g. *dewt* (in mirror-reversed form) was read as "dew". Additional examples of errors in these tasks are reported in Table 3. It is important to note that, in the case of mirror-reversed reading, these word-end errors occurred in response to letters *physically (absolutely and relatively) on the left* ("good") part of the word. NG correctly read 75% (225/300) of vertically printed words, and 70% (440/626) of mirror-reversed words (compared to 72% of normally printed words). Accuracy of responses to nonwords was also similar across formats: 27% (16/60) in vertical reading and 23% (25/108) in mirror-reversed reading.

TABLE 3
Examples of Reading Topographically Nonstandard Text

Errors in Vertical Reading		
blending → blemish	vivid → vivian	rang → ran
motionless → motel	discovery → discover	habitual → habit
strist → strip	neithem → neither	sipter → sip
Errors in Mirror-reversed Reading		
common → comet	joint → joint	regulated → regular
greenish → greenery	discovery → disco	dashes → dash
cring → crime	vigid → vigor	dring → drink
Errors in Recognition of Aural Spelling		
earns → earring	sparrow → space	village → villa
basis → bass	requirement → require	planet → plane
dring → drink	fing → fine	womar → woman

TABLE 4
Rate of Errors as a Function of Letter Position in Words of Different Lengths for Vertical and Mirror-reversed Reading and Recognition of Aurally Spelled Words. (x = Word Centre)

Word Length (Number of Letters)								
	Vertical Reading (x)							
4	0	1	2	18				
5	0	0	1	10	24			
6	0	0	2	11	21	34		
7	0	0	1	10	11	21	38	
8	0	0	4	10	11	23	31	39
	Mirror-reversed Reading (x)							
4	0	2	12	22				
5	0	1	3	11	21			
6	0	1	6	10	21	33		
7	0	0	7	18	27	31	41	
8	0	2	7	14	22	30	39	42
	Recognition of Aural Spelling (x)							
4	0	0	0	16				
5	0	0	1	10	22			
6	0	0	5	13	19	31		
7	0	0	0	6	14	26	36	
8	0	0	3	6	6	19	36	39

NG also made the very same types of errors when words were spelled aloud to her. Responses included 130 completion errors and 34 responses that omitted letters in word-final positions. For example, NG's response to the stimulus "e-x-c-e-s-s" was "exceed", and her response to "p-l-a-n-e-t" was "plane". Table 3 includes additional examples.

Table 4 presents the rate of errors at each grapheme position in vertically arrayed, mirror-reversed, and aurally spelled words, displayed as a function of the distance (in grapheme positions) from the centre of the word. Although the absolute rate of errors in these tasks was higher than in reading normal (left to right) print, it is clear that NG shows the same pattern of errors with respect to their distribution across grapheme positions in the word (and, for mirror-reversed words, the reverse of the pattern of errors across letter positions in the physical stimulus). Although the data from these tasks are a bit more "noisy" than those from normal reading, due in part to the fact that there were a lesser number of stimuli of each word length in these tasks, it is quite clear that reading errors are concentrated on the *right half* of words irrespective of the topographic arrangement of the stimulus or the form of the sensory signal (provided it concerns letters).

An additional list of 108 6-letter words were administered in all 4 reading tasks—the horizontal, vertical, mirror-reversed, and aural-spelling tasks—at 24 months post stroke. This allowed us to compare directly the levels of performance across tasks and, at the same time, to compare the distribution of errors as a function of position in a word. Overall accuracy rates were as follows: 79% for normal reading, 70% for recognition of aural spelling, 67% for vertical reading, and 60% for mirror-reversed reading. (The somewhat higher accuracy in normal reading can be accounted for by the fact that she never failed to attempt reading a normally printed word, whereas she declined to attempt reading 10–15 of the items in each of the other formats.) Table 5 illustrates that in all 4 tasks

TABLE 5
Distribution of Errors as a Function of Letter Position
(Given in % of Total Errors)

Task	*Letter Position in 6-letter Words*					
	1	*2*	*3*	*4*	*5*	*6*
Regular Reading	0	0	0	9	27	64
Vertical Reading	0	0	2	10	35	52
Naming of Orally Spelled Words	0	0	7	20	28	45
Mirror-reversed Reading	0	1	9	16	29	46

errors occurred virtually only on the right half of words, and increased at comparable rates as a function of distance from the centre of the word.

NG was also asked to read a list of 82 words (5–8 letters in length) in 4 formats on 4 different days: once with no spaces between letters (e.g. *cough*), once with 2 spaces between each pair of letters (*c o u g h*), once with 3 spaces between each pair of letters (*c o u g h*). The spacing had no effect on her reading accuracy: 76%, 74%, and 77% correct responses for no spaces, 2 spaces, and 3 spaces between letters, respectively.

Discussion

There are three crucial features of the results reported in this section: (1) NG made reading errors virtually only on the right half of words irrespective of the form of input—vertical, mirror-reversed, or aural spelling; (2) she produced quantitatively similar error patterns across tasks, both in terms of overall error rates and in terms of the distribution of errors as a function of position within a word (see Tables 4 and 5); and (3) her performance was unaffected by different spacings of the letters in a word. These results undermine unambiguously the possibility that NG's disorder could result from damage to the letter shape level. This conclusion follows from the assumption that information at this level of processing is represented in a stimulus-centred co-ordinate system. In such a co-ordinate system, the absolute location of a stimulus in space is not represented, but the absolute positions of its component parts in relation to each other are. In other words, a letter shape map veridically represents the within-stimulus spatial relations among letter shapes. Thus, for example, the letter-shape representation of a vertically presented word would preserve the vertical orientation of the stimulus. The implication of the foregoing is that since the letter-shape representations for horizontal, vertical, and mirror-reversed words are not the same, there can be no single form of deficit to this level of representation that could account for the qualitative and quantitative similarity of performance across reading tasks. This conclusion is further supported by the fact that the same pattern of spatially defined errors was observed for the visual presentation tasks and for the aural spelling recognition task, despite the fact that the latter task does not involve processing at the level of the letter shape map. Finally, the absence of a spacing effect on NG's reading performance allows the inference that the locus of impairment must be at a level of representation where relative and not absolute distances are encoded. We must conclude, therefore, that the spatially specific reading impairment considered here cannot be assumed to result from damage to low-level visual processing mechanisms.

Having ruled out damage to low-level visual processing mechanisms as the basis for the reported performance, the only remaining possibility is that the deficit concerns processing mechanisms at the grapheme level. At this level of processing, the information represented consists of a canonical description of letter identities (graphemes) and their relative order. Information about orientation of stimuli and absolute distances among letters are not represented at this level of processing. This means that the representations of a word presented in horizontal, vertical, or mirror-reversed form would be identical—in each case a canonical word-centred grapheme description. Consequently, damage to the processing mechanisms that operate at this level of representation would have identical effects on all reading tasks, irrespective of the topographic arrangement of stimuli—precisely the reported result.

Additional evidence for the hypothesis presented here accrues from the analysis of NG's spelling performance. Since one of the processing stages in spelling a word involves computing a grapheme-level representation of its orthographic structure (see Caramazza & Miceli, 1989; Caramazza, Miceli, Villa, & Romani, 1987; Ellis, 1988; for detailed discussion), damage to mechanisms operating at this level of representation should result in a spelling disorder (see Fig. 4). And, on the assumption that reading and spelling depend on the same processing mechanisms in com-

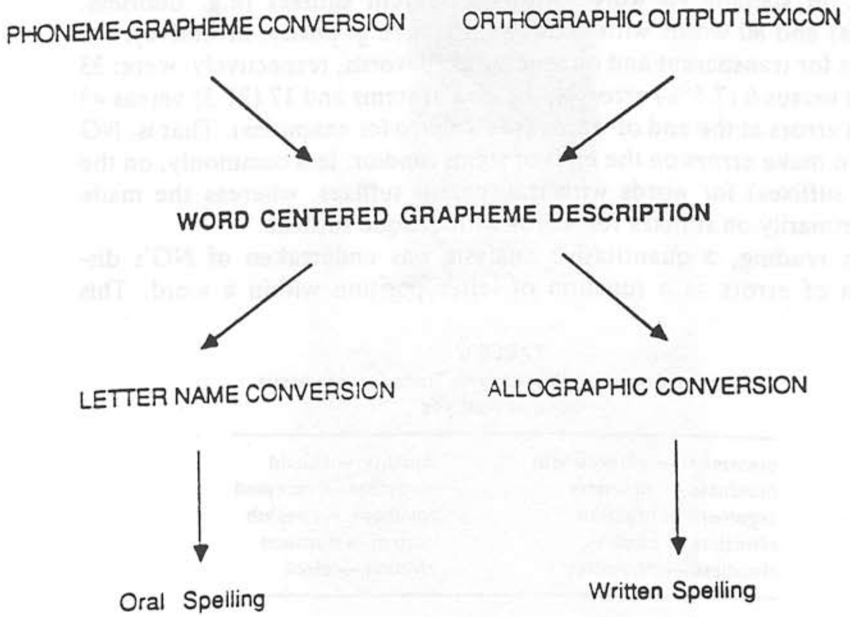

FIG. 4 The grapheme description and output processes in spelling.

puting grapheme-level representations, we would expect that damage to the grapheme level would result in qualitatively similar patterns of reading and spelling impairments.

Results: Writing to Dictation

NG's performance in writing to dictation was not significantly influenced by lexicality (word versus nonword), lexical parameters (frequency, concreteness, word class), or orthographic regularity. She correctly spelled 33% of the words and 21% of the nonwords matched for length in letters ($X^2_1 = 3.0$; n.s.). However, her spelling deteriorated steadily as a function of length, from 50% correct for 4-letter words to 7% correct for 8-letter words on a list counterbalanced for word length in letters, phonemes per word, and frequency ($X^2_1 = 5.6$; Mantel-Haentsel test for linear trend, $P < 0.03$).

Nearly all of NG's written spelling errors consisted of substitution, deletion, or insertion of letters at the end of the word or nonword. Most of the misspellings resulted in phonemically implausible nonwords, such as "advise" spelled as *advisd*, "sneeze" as *sneed* and /fout/ as *fol*. As in reading, her responses in spelling included some omissions of letters on the right (e.g. "broom" → *broo*). Suffixed words elicited some errors at the end of one or both component morphemes, such as "bottomless" spelled as *bottless* and "brightness" spelled as *brignesss*. These errors occurred more commonly in response to words with highly productive, transparent suffixes. In spelling 80 words with transparent suffixes (e.g. odorless, darkness) and 80 words with opaque suffixes (e.g. purity, inventive), the contrasts for transparent and opaque suffixed words, respectively, were: 33 (41.3%) versus 6 (7.5%) errors at the end of stems and 17 (21.3) versus 43 (53.8%) errors at the end of words (see Table 6 for examples). That is, NG tended to make errors on the ends of stems (and/or, less commonly, on the ends of suffixes) for words with transparent suffixes, whereas she made errors primarily on suffixes for words with opaque suffixes.

As in reading, a quantitative analysis was undertaken of NG's distribution of errors as a function of letter position within a word. This

TABLE 6
Spelling Errors on Words with Transparent versus Opaque Suffixes

placement → palacesment	equality → equald
blindness → blineness	exception → excepted
argument → argment	roughage → roughth
effortless → effuless	normal → normant
cloudless → clouless s	election → eleck

TABLE 7
Rate of Spelling Errors at Each Position of Words

Length in Letters	*Position in Word* 1	2	3	4	5	6	7
Left Aligned							
4	0	2	13	25			
5	0	0	6	20	29		
6	0	0	5	15	26	39	
7	0	0	3	5	15	28	51

							Word Centre x						
Centred Words													
4				0		2		13		25			
5			0		0		6		20		29		
6		0		0		5		15		26		39	
7	0		0		3		5		15		28		51

analysis revealed that virtually all of her errors occurred at the end of words (except for the noted contrast for transparent versus opaque suffixed words). The upper panel in Table 7 reports the results of this analysis separately for 4- to 7-letter words. It is apparent upon inspection of these data that the point in a word at which NG made spelling errors moved rightward for longer words. When the data on the distribution of errors as a function of position within a word are arranged relative to the word's centre, it is observed that errors virtually only occurred on the right half of a word, irrespective of word length (see lower panel in Table 7).

Discussion

The features of NG's spelling performance that merit close attention include the relatively high rate of spelling errors when compared to reading performance, the absence of frequency and lexicality effects, the fact that errors almost invariably resulted in nonword responses, the effect of morphological structure of stimuli, and the spatial-specificity of the impairment. Although all these features of NG's performance are important, for the moment we will focus on the spatial character of her impairment. The other features of NG's spelling performance will be considered in the General Discussion.

The fact that, as in reading, NG's spelling errors were virtually restricted to the right half of words, and that, also as in reading, errors increased "linearly" as a function of distance in number of grapheme positions from the centre of a word, invites the inference that damage to a processing stage common to reading and spelling is responsible for the dyslexic and dysgraphic performance. The hypothesis we will entertain here is that this common processing stage is the grapheme description. On this hypothesis, we would expect qualitatively and quantitatively similar patterns of performance in all types of spelling tasks, irrespective of the form (written versus oral spelling) or order of output (forward versus backward oral spelling).

Results: Oral Spelling, Backward Oral Spelling, and Delayed Copying

NG's performance was essentially identical across written and oral spelling and delayed copying (for words read correctly); she respectively made 64%, 67%, and 66% errors in the three spelling tasks. The quality of errors was identical across tasks. That is, errors were restricted to the rightmost part of a word, ordinarily resulting in nonword responses. The performance in the delayed copying task contrasts with direct copying, which she performed essentially flawlessly (97% correct). Furthermore, she made precisely the same types of errors in backwards oral spelling. To illustrate, "compare" was spelled backwards as "d-r-a-p-m-o-c" (compard) and "sheets" was spelled backwards as "e-e-h-s" (shee). In backward spelling of unaffixed words 80.7% of her errors were restricted to the end of words. In spelling suffixed words, 69.7% of her errors were restricted to the end of the word, and an additional 18.2% occurred at the end of one or both morphemes (e.g. "listed" was spelled as "d-e-s-i-l" (lised) and "sickness" was spelled backwards as "e-e-n-c-i-s" (sicnee). Because these word-end errors occurred at the *beginning* of the response in backward spelling, they cannot be ascribed to a general problem maintaining "attention" to the end of a stimulus or response (see also Baxter & Warrington, 1983). Examples of errors in these spelling tasks are shown in Table 8.

Table 9 reports the distribution of errors made by NG at various positions of a word for the same set of 108 6-letter words used in all spelling (and reading) tasks. Scoring of the position of spelling errors followed scoring procedures described for reading of the same set. It is apparent upon inspection that the distributions of errors are highly similar across tasks, and that errors essentially only occurred on the right half of a word in all tasks. Overall accuracy rates for these stimuli were also similar across tasks given during the same time period: 50% for written spelling, 48% for oral spelling, 44% for delayed copying, and 37% for backward oral spelling.

TABLE 8
Examples of Errors in Various Spelling Tasks

Errors in Written Spelling		
floor → floore	sneeze → sneed	cloud → clou
unit → unite	jury → jurd	faith → fait
skart → skarr	remmun → remmey	chench → chen
Errors in Oral Spelling		
career → carred	sneeze → sneed	ground → grou
poodle → poodler	afraid → afrain	period → perio
achieme → achiemd	emplain → emplained	spond → spone
Errors in Backward Spelling		
absorb → absown	sky → skik	church → chur
garbage → garbsi	oyster → oyste	sample → sampl
Errors in Delayed Copying		
square → squard	afraid → afrain	method → meth
turkey → turket	fabric → fabrict	starve → starv

Discussion

The most important aspect of the results reported in this section is the highly consistent, spatially specific pattern of errors in all spelling tasks. The fact that spelling errors virtually only occurred on the right end of a word, irrespective of the form and order of output (written versus oral spelling, and forward versus backward spelling) strongly argues for a deficit to a common level of representation in the spelling process. We have suggested that the only level of representation common to all spelling tasks is the grapheme description. Thus, we are led to conclude that the underlying cause for NG's spelling impairment is damage to the stage of processing where grapheme representations are computed. In fact, the hypothesis entertained here is that this level of representation is common to the reading and spelling processes. Thus, it is our expectation that reading and spelling performance should, in relevant respects, be qualitatively and quantitatively similar across all reading and spelling tasks. This expectation was borne out, as may be seen from a comparison of the results reported in Tables 5 and 9. These present, respectively, the distributions of reading and spelling errors as a function of within-word position.

These results further buttress the already strong evidence for the hypothesis that damage to a common level of representation—the grapheme description—is responsible for the co-occurring impairments in reading and spelling. And, since the deficit has a distinctly spatial character, involving only the right half of words, we are further led to conclude that information at the grapheme level is represented in a word-centred co-ordinate system. As noted earlier, it would stretch credulity beyond imagination to suppose that the remarkably similar patterns of performance in reading and spelling could result from damage to distinct processing mechanisms. It would seem, then, that the evidence from reading and spelling performance converge on a common account of the nature of NG's deficit, and, therefore, on crucial aspects of the processing structure of word recognition and production. Before turning to a discussion of these issues, we present further evidence about NG's reading performance that is relevant to the problem of word identification.

TABLE 9
Distribution of Errors as a Function of Letter Position
(Given in % of Total Errors)

	Letter Position in 6-letter Words					
Task	*1*	*2*	*3*	*4*	*5*	*6*
Written Spelling	0	0	6	19	31	46
Oral Spelling	0	1	2	7	29	60
Backward Spelling	0	0	1	12	34	53
Delayed Copying	0	1	1	13	28	57

The Effects of Adding a Suffix or a Prefix on Reading Performance

We have argued that NG is unable to process normally the right half of the word-centred grapheme description that serves as input to lexical access mechanisms. The evidence adduced for the hypothesis that the deficit involves the right half of a grapheme representation is rather compelling. Less clear is the precise nature of the object that is centred at the grapheme description level of representation. Although we have called this object "word", by this term we have not meant the corresponding linguistic object identified by that term. Instead, the term "word" has been used loosely as a label for the perceptually defined object that is represented at the grapheme description level. At this level of processing, the information available to the processing system consists of grapheme strings which may or may not turn out to correspond to words of the language. However, the grapheme string selected for representation at the grapheme description level is not arbitrary. It consists, instead, of the perceptual system's computation of a *potential* word, as signalled by perceptually significant features such as spacing. Thus, one characterisation of the object represented at the grapheme description level is that it consists of a grapheme string bounded by spaces. Another perceptual parameter that might be used by the visual processing system to segregate potential words is information about possible grapheme sequences in the language. Thus, for example, a string of identical graphemes, say a series of Xs, which violate graphotactic constraints of well-formedness, might be sufficiently perceptually salient to serve as the basis for defining the boundaries of a potential word. The experiments that follow were designed to evaluate the hypothesis that spacing and letter repetition constitute salient perceptual features for defining boundaries of potential words. We consider first the effects of spacing.

On the assumption that spacing constitutes one of the perceptually salient features used by the visual processing system to define potential words, and hence the representation computed at the grapheme description level, we would expect that the effective length of a grapheme string processed at this level of representation is determined by the spacing feature. The implications of this assumption for NG's reading performance are straightforward. Consider the case where NG is asked to read the word *lead*. Given the empirically established fact that NG is unable to process normally the right half of a word, she will (with some probability) have difficulty processing the *ad* part of the stimulus. However, given that the *le* part was processed normally, this information may be sufficient to activate the lexical representations LEAD, LEAN, LEFT, LEARN, LET, and so forth. Everything else being equal, the response she produces from the set of lexical representations that were activated, will be the one with the lowest activation threshold, say, "left" in this case. Consider now the case in which the word to be read is *leading*. In the latter case, the probability that she would produce the correct stem "lead" should be much greater than in the case where the stimulus was *lead*. The reason for this expectation is that, given the assumption of word-centredness, the addition of the suffix *ing* to the stem *lead* has the effect of shifting more of the graphemes in the stem (lead) to the left half of the grapheme description. Consequently, responses such as "left", "let", "letting", "learn", "learned", and so forth, should not be produced. Instead, responses should contain at least the graphemes in the left half of the grapheme description—e.g., "leading", "leads", "leader", or "leaden". Thus, the effect of adding a suffix, whether legal or illegal (e.g. *leadest*), should be to increase the probability of correctly reading the stem of the word if not the word itself.

The expectations for the addition of a legal or illegal prefix are the opposite of those described for the addition of a suffix—it should lead to a diminution of the number of correctly read stems. Thus, for example, if the stimulus were *relead*, likely responses would include "relive", "relay", "relief", "relate", "relent", "religion", "reliant", and so forth. The basis for this expectation is that adding a prefix has the effect of shifting stem graphemes to the difficult-to-process right half of the word-centred grapheme description.

Several experiments were designed to assess the effects on reading performance of adding letters to the beginning or end of a word. For all lists, words were presented individually without time limits. For lists with "illegal" affixes (e.g. prespend, leadest), NG was told that some of the stimuli were not real words, and that she should try to pronounce the entire letter string. For lists that contained words with a series of identical letters attached—a series of Xs—she was asked to ignore the series of identical letters.

Lists 1 and 2: Prefix and Suffix Effects

Methods. Forty-one words, 3 to 7 letters in length, that could be legally prefixed and suffixed (e.g. lead → misleading) were used in this experiment. Each word was presented in 4 forms: unaffixed (e.g. lead), prefixed (mislead), suffixed (leading), and both prefixed and suffixed (misleading), for a total of 164 stimuli. Four experimental blocks were constructed with each block containing roughly one fourth of each of the 4 types of unaffixed and affixed words.

Another set of 54 words, 3 to 6 letters in length and counterbalanced for word class, frequency, and length in letters, was presented on 3 occasions. One third of the words on each occasion were presented with illegal prefixes (e.g. malnoise), one third were presented with illegal suffixes (e.g. noiseful), and one third were presented without affixes (e.g. noise).

Results and Discussion. Table 10 presents a summary of the results on the 2 lists. The main finding from list 1 was that NG was twice as likely to make an error on a suffix when there was also a prefix. She made 50% suffix errors (substitutions or deletions) in the suffixed condition, compared to 94% suffix errors in the prefix plus suffix condition ($X^2_1 = 16.1$, $P < 0.0001$). For instance, she read *readable* correctly and *unreadable* as "unreading"; and she read *writes* correctly and *rewrites* as "rewritten". There was also a slight tendency to produce more stem errors when the stem was prefixed than when it was unaffixed (e.g. *loyal* → "loyalty" versus *disloyal* → "dislodge" and *content* → "content" versus *discontent* → "discontinue"), but the effect was small (10% versus 2%) due to the low rate of stem errors. This latter tendency was confirmed by results from list 2. NG made significantly more errors on the same stems when they were prefixed than when they were unaffixed (38.9% versus 11.1% errors on prefixed and unaffixed words, respectively; $X^2_1 = 9.68$, $P < 0.01$). That is, she was less likely to read the end of a word correctly if a prefix was present. To illustrate, she read *resist* correctly, but read *deresist* as

TABLE 10
Prefix and Suffix Effects

	Number of Errors (% of Stimuli in Parentheses)	
	Errors on Stem	*Errors on Suffix*[a]
List 1		
prefixed	4/41 (10)	n/a[b]
unaffixed	1/41 (2)	n/a
suffixed	1/41 (2)	20/40 (50)
prefixed + suffixed	5/41 (12)	34/36 (94)
List 2		
prefixed	21/54 (39)	n/a
unaffixed	6/54 (11)	n/a
suffixed	2/54 (4)	22/54 (41)

[a] We scored only those suffixes for which the corresponding stem was read correctly.

[b] There were three suffix insertions on these words.

"dereal"; and she read *jury* correctly, but read *misjury* as "misjudge".

The overall pattern of results confirms the expectations derived from the hypothesised structure of the representations at the grapheme description level. Specifically, the results confirm the expectation that grapheme descriptions are word-centred in the sense that a string of graphemes bounded by spaces, whether or not it constitutes a word, is the object that is centred at this level of representation.

Lists 3 and 4: Suffix Effects

Method. The suffix effect reported for list 1 may not have been significant only because NG's error rate on the unaffixed words was not high. Two new lists were constructed with the goal of inducing a high error rate on unaffixed words, so that any position effect of adding a suffix would be evident. Both lists consisted of pairs (and 1 triplet) of words that share at least the initial 4 letters (e.g. content and contest). List 3 included 96 words that could legally be affixed with a suffix of 3 or more letters (e.g. steepest and steering). This list was presented twice; each time half of the words were suffixed, so that each word was presented once with and once without a suffix. The same procedure was used with List 4, which consisted of 287 words that were presented once with and once without an *illegal* suffix.

Results and Discussion. NG made significantly fewer errors on word stems when a legal *or* illegal suffix was present. On list 3, she made 25%

TABLE 11
Legal and Illegal Suffix Effects

	Number of Errors (% of Stimuli in Parentheses)	
	Errors on Stem	*Errors on Suffix*[a]
List 3		
unaffixed	24/96 (25)	(8 suffix insertions)
legally suffixed	4/96 (4)	66/96 (67)
List 4		
unaffixed	98/287 (34)	(11 suffix insertions)
illegally suffixed	21/287 (7)	156/287 (54.3)

[a]We scored only those suffixes for which the corresponding stem was read correctly.

stem errors in response to unaffixed words, compared to only 4.2% stem errors in response to the same words when they were suffixed ($X^2_1 = 15.10$, $P < 0.001$). For example, she read *access* as "accept" and read *accessibility* as "accessible". Similarly, on list 4, she made 43.1% stem errors in the unaffixed condition, compared to only 7.3% stem errors when each word was illegally suffixed ($X^2_1 = 61.23$, $P < 0.0001$). Table 11 summarises the results obtained with Lists 3 and 4.

Another important result from these data was that on both lists NG made approximately twice as many errors on the lower-frequency stems than on the higher-frequency stems in each pair of words with the same first 4 letters. On the combined lists, she made 80 errors on the lower-frequency words, compared to 42 errors on the higher-frequency words, when they were unaffixed: and she made 17 errors on the lower-frequency stems compared to 8 errors on the higher-frequency stems when they were suffixed. To illustrate this effect, she read *almost* correctly, and read *almond* as "almost": and she read *review* correctly, and *revive* as "review". Table 12 demonstrates NG's tendency to produce high-frequency words in response to both high- and low-frequency stimuli. A further confirmation of this result was that the mean frequency of correctly read stems was significantly higher than the mean frequency of incorrectly read stems: 51.1 versus 13.0 ($t = 2.108$, 95 d.f., $P = < 0.05$ by two-tailed t-test) for list 3, and 57.8 versus 14.4 ($t = 2.605$, 286 d.f., $P = < 0.01$ by two-tailed t-test) for list 4.

NG's reading performance for lists 3 and 4 convincingly shows that the object that is centred at the grapheme description level consists of a grapheme string bounded by spaces. One implication that follows from this conclusion is that the representation computed at the level of the grapheme description does not contain morphological structure.

TABLE 12
Stem Errors as a Function of Stem Frequency

	Number Errors on Highest Frequency Stem (% in Parentheses)	Number Errors on Lower Frequency Stem
List 3		
unaffixed	8/48 (17)	16/48 (33)
legally suffixed	1/48 (2)	3/48 (6)
List 4		
unaffixed	34/143 (24)	64/144 (44)
illegally suffixed	7/143 (5)	14/143 (10)

Having established that spacing is a sufficient condition for defining the boundaries of the representation that is computed at the grapheme description level, we turn next to a consideration of more subtle perceptual features—repeated identical letters.

List 5: Effects of Identical Letter "Affixes"

Method. A list of 42 words was presented in 6 conditions, distributed equally across 6 forms. Each word was presented once without any affix (e.g. tempt), once with a string of 4 Xs as a "prefix", once with a string of 4 Xs as a "suffix", once with a legal suffix (e.g. temptation), once with a legal suffix and a 4-X prefix, and once with a legal suffix followed by a 4-X suffix. For example, each form of list 5 contained one of the following: *problem*, *xxxxproblem*, *problemxxxx*, *problematic*, *xxxxproblematic*, or *problematicxxxx*.

Results and Discussion. The results are summarised in Table 13. Adding a string of Xs to the beginning of the word had no effect at all on

TABLE 13
Effects of Graphotactically Ill-formed Suffixes

Affix	Number (%) Errors on Stem	Suffix Errors
No affix	5/42 (12)	3/42 (7) (insertions)
4-X prefix	5/42 (12)	2/42 (5) (insertions)
4-X suffix	1/42 (2)	11/42 (26) (insertions)
Legal suffix	0/42 (0)	26/42 (62)
Legal suffix + 4-X prefix	0/42 (0)	26/42 (62)
Legal suffix + 4-X suffix	0/42 (0)	20/42 (48)

reading. Rates of suffix and stem errors were precisely the same for 4-X prefix and no prefix conditions, for both suffixed and unsuffixed words (e.g. *problem* versus *xxxxproblem* and *problematic* versus *xxxxproblematic*). On the other hand, a string of Xs on the end of a word resulted in fewer stem errors (2% versus 12%) and more suffix insertions (26% versus 7%), for 4-X suffixed words and unaffixed words, respectively. For example, *material* was read as "mature", whereas *materialxxxx* was read as "materialise". In reading legally suffixed words, there were no stem errors, but the addition of a string of Xs at the end had the effect of reducing the number of suffix substitutions and deletions. To illustrate, *childishly* was read as "childless", whereas *childishlyxxxx* was correctly read as "childishly". Thus, repeated-X suffixes, which she was instructed to ignore, were almost as effective as graphotactically appropriate suffixes, which she was instructed to read, in reducing stem errors (2% versus 0% stem errors for 4-X and legal suffixes, respectively).

The results of this experiment are important for two reasons. First, the fact that a 4-X prefix can effectively be ignored shows that perceptual features—visual identity, in this case—may be used to segregate grapheme strings into potential words for centring (or re-centring) at the grapheme description level of representation. Second, the fact that a 4-X suffix functioned as effectively as a regular suffix to shift the centre of a grapheme representation rightward shows that NG cannot process adequately even "low-level" information, such as identity, when this information falls on the right part of the grapheme representation. The latter result is consistent with our earlier observation that there seems to be no graphemically usable information on the right end of NG's word-centred grapheme representations.

Line Bisection

Thus far, we have focused entirely on NG's reading and spelling performance. However, this emphasis must not be assumed to reflect a belief that NG's impairment is language specific. We have already noted that she made errors in various tasks not involving words or letter strings. Interestingly, this impairment in processing other visual materials is, in important respects, of the same form as that documented for words. This contention is supported by NG's performance in line-bisection tasks.

Method. NG was presented with 5 horizontal lines of each of the following lengths: 4, 6, 8, and 10 inches, for a total of 20 trials. Each line was centred on an 11-inch wide by 8½-inch high sheet of paper. Lines of various lengths were presented randomly, and were placed directly in front of her on a flat surface. She was instructed to mark the centre of the line.

TABLE 14
Line Bisection as a Function of Line Length

	Standard Task Administration		After Marking an "X" at Each End	
Length	Mean Displacement (in eighths of an inch)	Proportion of Line	Mean Displacement	Proportion of Line
4"	6.2	0.10	5.0	0.08
6"	9.4	0.10	6.6	0.07
8"	12.0	0.09	10.6	0.08
10"	15.0	0.09	11.2	0.07

This procedure was repeated, with the additional instruction to mark each end of the line with an "x" before marking the midline.

Results and Discussion. NG's performance on these tasks revealed a marked leftward bias in indicating the centre of a line, suggesting inability to process the right part of the line adequately. The mean deviations from the actual centre of the line for each length on the first task are shown in Table 14. Her mean responses deviated in each case by 9–10% of the stimulus. The leftward bias in marking the centre of a line was not affected significantly by first having to mark the ends of the line before bisecting it, even though she never erred in marking each end of the line (Table 14). The overall mean deviation for the standard line bisection task was 1.33ins. (s.d. = 0.83ins.), and the overall mean deviation for line bisection after placing an x at each end was 1.04ins. (s.d. = 0.5ins.) ($t = -1.332$, 38 d.f., $P = 188$, n.s. by 2-tailed t-test).

The line bisection results show that NG's spatially specific processing deficit is not restricted to processing letter strings. This is an important result. It allows the inference that the deficit concerns visual processing in general, and not just letters. Consequently, conclusions reached from the detailed analysis of reading may be generalised to object recognition. The implications of this generalisation are now discussed more fully.[11]

[11] The results of the line bisection tasks, although seemingly not inconsistent with the claim that NG has a spatially specific processing deficit at the grapheme level, do raise an important issue. There are computational reasons for expecting that a grapheme level representation must be computed in the process of recognising or producing a written word. However, it is not immediately obvious that there is any need to compute such an abstract level of representation for perception of lines. It might have been assumed that the representation computed at the level of the stimulus-centred 2½-D sketch (corresponding to the letter shape level in word recognition) should have been sufficient to support normal performance in line bisection. And, since the latter level of representation is not impaired in NG, as indicated by the fact that her errors in reading words in different topographic orientations always involved

GENERAL DISCUSSION

The results we have reported severely constrain plausible claims about the possible locus of damage to NG's visual word recognition system. In particular, the results converge in support of two specific claims: (1) the damage is in a stage of processing where word-centred grapheme representations (or object-centred 3-D model descriptions) are computed; and, (2) the deficit concerns only the right half of these representations. The evidence adduced in support of these conclusions is the following:

1. NG's reading errors involved the right part of words and nonwords, irrespective of their topographic arrangement (horizontal, vertical, or mirror-reversed) or the sensory modality of input (visually vs. aurally presented stimuli). Her word recognition and reading impairment persisted even when she successfully named all the letters in the stimulus, and when the stimulus was presented tachistoscopically to her normal left visual field. Furthermore, reading performance was unaffected by the spacing between letters in a word. These results rule out as the locus of deficit damage to retino-centric feature level and the stimulus-centred letter shape level. The first of the two levels is ruled out by the fact that the impairment was not sensory-modality specific, by the fact that the impairment persisted despite accurate reading of all the letters in the stimulus, and by the fact that the impairment was independent of the topographic location and arrangement of the stimulus. The second level is ruled out by the fact that her impairment remained unchanged under topographic transformations of the stimulus.

2. NG's distribution of reading errors as a function of letter position within a word (or nonword) were restricted to the right half of the word (nonword), irrespective of length, topographic orientation, and modality of input. Furthermore, the probability of an error increased "linearly" as a function of distance from the centre of the word. These facts unambiguously establish that the deficit concerns the right half of grapheme representations.

3. NG's spelling performance was qualitatively identical to her reading performance: she made spelling errors only on the right half of words and

---

the end of words, we would have expected normal performance in the line bisection task. How, then, do we account for the line bisection results? The interpretation we would like to offer for the line bisection results is based on two assumptions: (1) that the computation of the 3-D model level of visual representation is mandatory even for perceptually simple forms such as line segments: and (2) that low-level perceptual processes are impenetrable to cognitive operations (Fodor, 1983). Specifically, the contention is that although the representation at the level of the 2½-D sketch is normal, this representation cannot be queried by cognitively driven operations such as that of deciding the mid-point of a line.

nonwords, irrespective of modality (written versus oral spelling) and order of output (forward versus backward spelling). This fact provides strong converging evidence for the hypothesis that the deficit responsible for NG's reading and spelling disorder concerns a graphemic level of representation—the only level of representation that may plausibly be assumed to be common to all the tasks for which a spatially specific reading and spelling deficit was observed. This fact also supports the hypothesis that the deficit concerns only the right half of grapheme representations.

4. NG's systematic leftward shift in the line bisection task and her omission errors of the right part of objects in a drawing task indicate a deficit in processing the right part of *all* visually based representations, whether words or objects. This result sanctions generalisations from NG's word recognition performance to visual recognition, more generally.[12]

The fact that it is possible to explicate highly detailed features of NG's impairment by proposing damage to a stage of processing where word-centred grapheme representations are computed may not only be taken as providing evidence for the proposed level of representation, but also for the model that contains that level of representation. However, the validity of the model depends also on its ability to explicate other features of the subject's performance besides the spatially determined impairment on which we have focused thus far. In addition, the model must provide a principled account for the performance of other patients reported in the literature. A consideration of these results allows a more stringent assessment of the validity of the overall architecture of the proposed model of word recognition, as well as a greater articulation of the representational and processing structure of the hypothesised stages of processing. We begin with a discussion of results relevant to the word-centred grapheme level of representation and proceed to discuss the stimulus-centred letter shape and the retino-centric feature levels.

[12]The claim that a common deficit is responsible for NG's spatially determined impairment in processing words and objects would seem to be at variance with the results reported by Costello and Warrington (1987) for patient JOH. This patient presented with greater difficulties in processing the *left* part of words, but the *rightmost* object of series of objects and the *right* part of lines. However, JOH had bilateral parietal damage with a dense right homonymous hemianopia. Furthermore, inspection of the reading errors made by JOH revealed that he also made a nontrivial number of errors on the *right* part of words (e.g. *keep* → "knee"; *England* → "angle"). It would seem, then, that the reported dissociation between "object neglect" and "word neglect" may simply be the result of the interaction between processing mechanisms for the recognition of words and objects with different types of spatially specific deficits to the right *and* left representational spaces.

The Word-centred Grapheme Level: Further Evidence

Although NG's performance in reading and spelling words and nonwords was qualitatively identical with respect to the spatially determined nature of the impairment, it was strikingly different in terms of overall levels of performance, and in terms of the effects of (nonspatial) characteristics of stimuli on performance. Four factors are important for consideration here: (1) the fact that the overall accuracy level in reading words was in the order of 70% to 90%, depending on various factors such as frequency, whereas overall accuracy in reading nonwords was around 25% to 35%, and accuracy in spelling both words and nonwords was in the range of 20% to 40%; (2) the fact that while reading performance was affected by word frequency, spelling performance was not; (3) the fact that error responses in reading both words and nonwords were almost always words, whereas error responses in spelling both words and nonwords were almost invariably nonwords; and (4) the fact that in reading errors occurred at the end of words irrespective of their morphological structure, but in spelling errors occurred not only at the end of words but, in transparently suffixed words, also at the end of stems (e.g. brightness → "brignesss"). Factors 2 and 4 are summarised in Table 15.

These dissociations could, of course, merely reflect the co-occurrence of deficits to different components of the reading and spelling systems, resulting in different patterns of performance for the two tasks. Alternatively, and more interestingly, the observed pattern of performance could reflect characteristics of the processing structure at the level of the damaged word-centred grapheme description that is assumed to be responsible for the spatially determined impairment in this subject. In the latter case, we would have further evidence for the hypothesised model of word recognition (and production).

TABLE 15
Percentage Accuracy and Predominant Error Type in Reading and Writing Words and Nonwords

	Overall Accuracy	*Types of Errors*
Reading Words	89%	Visually similar words (e.g. *hound* → "house")
Reading Nonwords	29%	Visually similar words —Lexicalisations (e.g. *dring* → "drill")
Spelling Words	33%	Nonwords (e.g. sneeze → *sneed*)
Spelling Nonwords	21%	Nonwords (e.g. spond → *spone*)

An explanation of the observed difference in levels of accuracy between words and nonwords follows directly from two assumptions about the structure of the lexical access procedure outlined in the proposed model of word recognition: (1) the assumption that a word-centred grapheme representation activates in parallel all lexical representations in proportion to the degree of grapheme overlap between the input and lexical representations: and (2) the assumption that the lexical representation that receives the most activation, above a prespecified minimal level, or the one to reach threshold first, will be processed further and produced as a response. As discussed in detail in an earlier section of this paper, the implications of these processing assumptions is that even though the right part of the grapheme representation of a word stimulus cannot be processed normally, the usable left part may be sufficient in many cases to activate the correct lexical representation, resulting in a correct response. For nonwords, this situation does not obtain, so that default correct responses are unlikely. Consequently, the expectation is that reading accuracy for words should be superior to reading accuracy for nonwords—the obtained result.

Similar considerations account for the discrepancy in accuracy levels between reading and spelling despite the assumed commonality in the level of deficit responsible for the impairments in these tasks. The reason for the observed discrepancy in performance levels is that different constraints are at play in the use of the grapheme representations in reading and spelling: in reading the grapheme representation must be processed in parallel for lexical access, whereas in spelling the graphemes in the grapheme representation must be processed sequentially for conversion into specific letter shapes or letter names: in reading the output of the operations applied to the grapheme representation is a single object, a lexical representation, whereas in spelling the output of the operations applied to the grapheme representation is a set of independent letter shapes or letter names. Consequently, whereas in reading the normally processed part of the grapheme representation can constrain the activation of a lexical entry in the orthographic lexicon, in spelling, the normally processed part of the grapheme representation cannot constrain the independent operations of the allographic conversion mechanism which operates over individual graphemes. Thus, in the case of spelling there is no default mechanism to constrain possible *letter* responses. If the information on the right part of the word-centred grapheme representation is not usable, spelling responses for the graphemes in this part of the representation can only consist of random letter errors.[13] Thus, word spelling accuracy is expected

[13]Or, spelling responses for the graphemes in the right (impaired) part of the representation might consist of letters corresponding to the most available graphemes—for example, *x*, *z*, and *q* were rarely, if ever, produced at the end of spelling responses, whereas *s* and *t* were very common.

to be poorer than word reading accuracy. Furthermore, since lexical status is not a factor in the sequential processing of graphemes for output processes in spelling, the expectation is that words and nonwords should be spelled with comparable levels of accuracy (Caramazza et al., 1987). The results obtained for NG are consistent with these expectations.

The remaining three features of NG's performance also receive ready explanations in light of these considerations. Thus, the fact that word frequency affected reading but not spelling performance is explained by the fact that frequency is a factor in lexical access in reading—frequency affects threshold settings for activating lexical entries—but plays no role in the sequential processing of graphemes in spelling. Consequently, we expect better reading performance for high- than low-frequency words, but no effect of word frequency in spelling. In the case of the difference in error responses for reading and spelling, this is explained by the fact that the default responses of the lexical access procedure in reading involve words, whereas the default responses in spelling can only be individual letters.

The differential effect of morphological structure on reading and spelling also follows from assumptions we have made about the structure of the processes that are applied to grapheme representations in the two tasks. In the case of reading, the factors that determine the unit of representation that serves as input to the lexical access procedures (at least on a first pass), are strictly perceptual—spacing and identity, for example. The representation that is submitted for further processing at the lexical access stage is a letter string, and *not* a morpheme or word. Consequently, spatially determined errors will occur at the end of the letter string, whether it is a word, a morpheme, or a nonword. In spelling, by contrast, the representations computed prior to the grapheme level consist of lexical units—words and/or morphemes. If the production system can output morphemes as well as words, then, we might expect errors at the end of stems (or roots) and the end of suffixes when the unit of output is the morpheme. The fact that NG made errors both at the end of roots and suffixes supports the hypothesis of morphological compositionality in output (see Badecker, Hillis, & Caramazza, 1990; Miceli & Caramazza, 1988; for further discussion of this hypothesis). For present purposes, the important point is that the contrast in the spatial distribution of errors for spelling and reading follows directly from computationally motivated assumptions about lexical access in reading and lexical production in spelling.

In short, then, the proposed model of word recognition (and production) can account not only for NG's spatially determined impairment in reading (and spelling) but also for the pattern of response accuracy, the effects of frequency, and the effect of response type observed in her reading and spelling performance.

In the effort to provide an account for NG's contrasting performance on some aspects of reading and spelling, an important property of the processing structure at the grapheme level was made explicit: whether a representation at a specific level of processing is processed in parallel or serially need not be an intrinsic property of the representation level itself, but could be determined by the interactions that are possible between that representation level and the mechanisms that operate over the information represented at that level of processing. Thus, in the case of reading, the grapheme representation is processed in parallel because the information relevant to the lexical access mechanism is the whole grapheme representation. This, on its own, does not imply that representations at this level must necessarily be processed in parallel, but confers a high degree of plausibility to such a possibility. By contrast, in the case of spelling, graphemes at the word-centred grapheme level are processed serially because the mechanisms that operate over representations at this level of processing (allographic conversion and letter name conversion) are concerned with the computation of individual letter forms, and not words. This computational distinction between different forms of processing at the grapheme level of representation—parallel for lexical access in reading, serial for letter form conversion in writing—is supported by the reported results.

The Word-centred Grapheme Level: Spatial or Ordinal Co-ordinates

Throughout this report we have assumed, without explicit justification, that order information among graphemes at the word-centred grapheme level is encoded in a spatially defined co-ordinate system (see also Hillis & Caramazza, 1989; 1990). This contention is on the face of it implausible, even seemingly contradictory: graphemes are, by hypothesis, abstract objects encoding font-, case-, size-, and orientation-independent letter information. And, yet, we have reported results that unambiguously support the thesis that NG's impairment concerns the right half of a word-centred representation.

The evidence we have presented undermines the hypothesis that the deficit concerns an ordinally determined part of a grapheme representation—the distribution of NG's reading and spelling errors indicates that errors are restricted to the *right half* of a word (or nonword), irrespective of stimulus length. Thus, the deficit cannot be specified simply in terms of ordinal positions defined in terms of either the beginning or end of a word. An alternative possibility is that grapheme order is specified in terms of ordinal position from the centre of the word. Thus, for example, the order of the graphemes for the word stimulus *chair* could be specified as [<c/-2>, <h/-1>, <a/0>, <i/1>, <r/2>]. This form of representation for the order of the graphemes in a word can successfully capture the fact that NG's impairment concerns a part of the stimulus specified by reference to the mid-point of the word. However, this representational format fails to provide a motivated account for the fact that the impairment concerns the right part of the word. Of course, it could be objected that we have prejudged the issue by labelling the impaired part as the "right" part. However, this choice is not unmotivated. Recall that NG was not only impaired in processing words, but also in line bisection, object drawing, and other perceptual tasks—tasks for which a characterisation of the relations among the parts of the stimulus could not be given in any simple way in terms of ordinal relations. It is, therefore, much more parsimonious to assume that a common, spatially determined deficit is responsible for the observed processing impairment for words and objects. Nonetheless, we recognise that there are nontrivial theoretical questions introduced by the joint assumptions that representations at the level of the word-centred grapheme description consist of abstract letter identities and that the order among these objects is encoded in a spatially defined co-ordinate frame.

Finally, there is the matter of the mechanism that centres the grapheme string in the spatially defined co-ordinate system. One possibility is that the centre at the grapheme level is simply the same as the one at the letter shape level. On this view, the centre of the string of shapes at the stimulus-centred, letter shape level serves as the anchor point about which to place other graphemes at the grapheme level. This way of establishing a centre for the grapheme string does not preclude the possibility that the grapheme string may be re-centred on the basis of grapheme level information (see section on the effects of affixation).

The Stimulus-centred Letter Shape Level and the Retino-centric Feature Level

NG's reading performance is in many respects similar to that of other brain-damaged subjects who have been classified as "neglect dyslexics" (Behrmann, Moscovitch, Black, & Mozer, in press; Brunn & Farah, Note 1; Costello & Warrington, 1987; Ellis, Flude, & Young, 1987; Friedrich, Walker, & Posner, 1985; Kinsbourne & Warrington, 1962; Riddoch, Humphreys, Cleton, & Fery, this issue; Warrington & Zangwill, 1957; Young, Newcombe, & Ellis, this issue (Part 2, 1991); see Shallice, 1988, for review). Like these other patients, NG's reading impairment has a spatially determined character—errors occur only at one end (the left or the right) of the word, the right end in her case. However, there is a crucial difference between NG's performance and that of other cases described in the literature. Unlike these other cases, NG's reading impairment was invariant under topographic transformations of the stimulus (horizontal, vertical, and mirror-reversed letter strings), and she also showed a similar spatially determined impairment in all forms of spelling. The subjects reported by Ellis et al., Behrmann et al., and Riddoch et al. made errors on the same relative side of space for standardly presented and mirror-reversed words—that is, these subjects made errors on the *left part of the physical stimulus* whether or not it corresponded to the beginning of the word. For example, VB (Ellis et al.) misread the standardly presented word *yellow* as "pillow", but she misread the mirror-reversed stimulus *plant* as "plane". Furthermore, VB showed a dissociation between preserved oral spelling ability and impaired written spelling, the latter characterised by stroke errors and a tendency to write down the right side of the page (Ellis et al., 1987). Thus, there can be no doubt about the fact that the locus of impairment in VB cannot be the same as that hypothesised for NG. In fact, for VB, we must assume that the impairment involves a stage of processing prior to the word-centred grapheme level: either the stimulus-centred letter shape level or the retino-centric feature level, or, more likely, both. Although the results reported by Ellis et al. do not exclude the possibility that VB has an impairment at the level of the retino-centric feature map, they do establish clearly that she had an impairment at the level of the stimulus-centred letter shape map. The evidence for this contention is that VB made errors on the left part of stimuli even when these were presented tachistoscopically in the "intact" right visual field, and even when she had successfully read a digit to the immediate left of the word. The latter results rule out as the *only* locus of deficit the retino-centric feature level and implicate damage to the stimulus-centred letter shape level as at least one factor responsible for VB's reading impairment[14] (see also Behrmann et al., in press; Kinsbourne & Warrington, 1962; Riddoch et al., this issue; Young et al., this issue (Part 2, 1991)).

[14] Ellis et al. (1987) interpret their result with mirror-reversed words as providing evidence against the hypothesis that a word-centred representation level is computed in the course of visual word recognition. They argue (p. 459) that ". . . if neglect operated on word-centred co-ordinates, then initial letters would be neglected, regardless of the orientation of the word." Because this result did not obtain with their subject, VB, they conclude that order information in an abstract letter representation is encoded in terms of ordinal spatial positions (Seymour, 1979). This conclusion is inadequate for several reasons. First, it is empirically inadequate as demonstrated by the results reported for NG—there are patients whose performance is invariant under topographic transformations of stimuli (see also Hillis & Caramazza, 1990). Second, it is inadequate because it fails to account in a motivated way for the right-of-centre/left-of-centre character of the reading difficulties recorded for various "neglect" patients—as noted here, a simple representation of order information cannot capture the centredness character of the impairment of many of the patients studied, including VB. And, third, it is forced to make an unprincipled distinction between perceptual processes involved in reading words and numbers—the latter resulted most often in deletion errors whereas word errors often resulted in letter substitutions. This dissociation follows naturally from the assumption (see earlier) that word reading is constrained by possible lexical responses, whereas number reading is not subject to similar constraints.

Further evidence for spatially determined selective damage to the stimulus-centred letter shape level and the retino-centric feature level is reported by Rapp & Caramazza (this issue). These authors investigated the performance of a brain-damaged subject, HR, who presented with clinical symptoms of "letter-by-letter" reading—she identified each letter in a word either audibly or subvocally before attempting to pronounce it. Detailed investigation of the subject's ability to identify letters in a visual display documented a low-level visual processing deficit. The deficit was characterized by an increasingly severe left-to-right processing limitation for absolute spatial positions of letters. Furthermore, a relative spatial position effect was observed for horizontally arrayed stimuli, independent of absolute spatial position, with greater processing difficulties for letters on the right relative to other letters in the array. This effect of relative spatial position was not observed for vertically arrayed stimuli. These results were interpreted as reflecting spatially-determined damage for positions on the right part of representations at the level of the retino-centric feature map *and* stimulus-centred letter shape map (see Rapp & Caramazza, this issue (Part 2, 1991), for a detailed account of the possible relationship between the hypothesized form of damage in this case and her letter-by-letter reading performance).

The performance by patient MO (Riddoch et al., 1990) may be explained as resulting from selective damage to the retino-centric feature map. MO's left-sided errors in reading words occurred only when the initial letters were presented in his left visual field (perhaps within the portion of the field affected by his hemianopia). That is, MO made errors only in response to words presented tachistoscopically, which would prevent compensation by eye movements for his impairment in processing the left half of the retino-centric representation. Furthermore, when words were presented for sufficiently brief durations to prevent refixating them (250msec.), he made as many errors on the second letter of 5-letter words as on the first letter of 4-letter words (which were presented in the same retinal position). And, when 4-letter words were presented further to the right of fixation, MO made fewer errors on the first letter of each word. Together, these results are consistent with the hypothesis that MO's spatially specific deficit disrupted word recognition at the level of the retino-centric feature map.

The performance of patient TB (Patterson & Wilson, 1990) could also be accounted for by assuming damage to the retino-centric level of representation such that damage is restricted to a small segment of the feature map just at and to the left of fixation (as can result from a scotoma in one eye, at least when the other eye is nonfunctional as in TB's case). If we assume that he fixates just to the left of the second letter of 4-letter words, then he would have impaired perception of the first letter, resulting in errors like *rose* read as "nose". Further, if he were to fixate on the second letter of 5-letter words he would have impaired perception of the first and second letters, resulting in errors like *xlead* → "read" and *spear* → "wear". Adding letters, symbols, or words on the left side would shift his fixation (O'Reagan & Lévy-Schoen, 1987; O'Reagan, Lévy-Schoen, Pynte, & Brugaillère, 1984), such that more of the left letters of the word would fall within the spared retino-centric feature map, and would be read correctly—just the pattern of results reported. On longer words, only 1 or 2 letters at and to the left of fixation would be affected, so that the available information would often be sufficient to activate the correct representation in the orthographic input lexicon. For example, perception of *elephant as el_hant*, would activate the lexical representation for elephant more than any other representation. In general, the longer the word, the fewer representations would share the correctly perceived graphemes, a fact that would account for his better performance on longer words. The reported performance on stimuli like *cashland*, which elicited responses that substituted the initial letter of the first word only, would also be predicted on this account (but *not* on the McClelland & Rumelhart [1981] account as argued by Patterson & Wilson, 1990), since the entire second word would fall on the intact retino-centric map, assuming he fixated toward the beginning or in the middle of the letter string. Further, such damage would explain his perception of nonword letter strings, if they are normally fixated and perceived in a fashion similar to words. A scotoma at and to the left of fixation would also result in impaired perception of individual letters (and vertically printed words, if each letter is fixated individually). Upper-case letters might have been perceived by TB more adequately than lower-case because they are generally larger (so that more feature information would fall in the intact field). Finally, the hypothesised deficit would explain his performance in text reading—errors only on the left side of individual short words—since each word would be fixated individually. Of course, the deficit should also affect visual perception of stimuli other than alphanumeric characters, to the extent that the stimuli are comparable in size and discriminability to letters and numbers. TB's perception of such stimuli could not be tested because of his deteriorating vision. (Thus, his medical history is in line with the hypothesised impairment, since deterioration of vision to only light/dark vision, as in his left eye, would be expected in the case of macular degeneration, the most common cause of scotomas.)

In some cases, it is not possible to determine the locus of impairment with respect to the level of representation in word recognition from the data reported. For example, RNR (Warrington, this issue (Part 2, 1991) made errors that were very similar to those of NG in terms of the increase in rate of errors as a function of the distance from the centre of the word on the right side. This pattern of errors would be expected to result from right-sided damage at the level of the stimulus-centred letter-shape map *or* damage at the level of the word-centred grapheme description. Other features of RNR's reading performance are also consistent with hypothesised damage at either level of visual word recognition. Thus, RNR made errors in reading words presented for brief durations to his left (intact) visual field at a rate comparable to his error rate in untimed reading. Like NG, RNR also made some right-sided errors in vertical reading that appeared to be similar to his right-sided errors in normal (horizontal) reading. However, RNR also made many other kinds of errors in the vertical reading task (thus, less than 50% of his errors on this task occurred only on the end of the word). Similarly, his spelling errors were not restricted to the right side, presumably due, at least in part, to a premorbidly low level of spelling skill. Because of the inconclusive data from vertical reading and spelling, and absence of data on reading mirror-reversed words or recognising aural spelling, it is not possible to ascertain with any confidence the locus of damage responsible for RNR's right-sided errors in reading.

In summary, the fact that contrasting patterns of deficits to the word-centred grapheme level—NG (see also RB and HH reported in Hillis & Caramazza, 1990, and ML and DH reported in Hillis & Caramazza, 1989)—the stimulus-centred letter shape level—VB, HR, JB, SP and patients reported by Kinsbourne and Warrington (1962)—and the retinocentric feature levels, HR, MO, and TB (and perhaps JB, SP, and VB), as well as cases reported in Behrmann et al. (in press) provide strong evidence for the multi-level model of visual word recognition proposed in this paper.

On the Interpretation of Neglect: Real and Imagined Problems

A Disorder of Attention or of Representation? Thus far the focus of our discussion has been entirely on the implications of the spatially specific nature of NG's deficit for claims about the processing structure of written word recognition and production. We have avoided all discussion of whether the cause of the "neglect" impairment recorded for NG is to be specified in terms of a deficit to attentional mechanisms (e.g., Heilman, Bowers, Valenstein, & Watson, 1985; Kinsbourne, 1970; Posner, Cohen, & Rafal, 1982; Riddoch & Humphreys, 1987) or to the information represented at some stage of processing (e.g., Bisiach, Luzzatti, & Perani, 1979; DeRenzi, Faglioni, & Scotti, 1970). Although it is not possible to give a definitive answer to this question, there are several considerations that make the attentional account highly problematic. The default conclusion will be that it is more reasonable to assume that, at least for NG, the underlying cause of the observed spatially determined impairment is damage to the grapheme representation computed at the word-centred grapheme level.

The considerations that led to this conclusion are as follows. We have shown that NG's processing at the retino-centric feature level and the stimulus-centred shape level is essentially intact. This implies that the attentional mechanisms that direct processing resources to parts of the feature map and to parts of the letter shape map must be functioning sufficiently well to guarantee "normal" processing at these levels of representation. Were this not the case, we would have observed spatially determined deficits at the level of retinotopic and stimulus-centred representations. However, this was not the case. NG's spatially determined impairment concerned only the right half of a word-centred grapheme representation. This pattern of performance rules out a general deficit in directing attention to representational spaces. Consequently, either we assume that there are distinct attentional mechanisms operating at different levels of representation, or we reject the hypothesis that the deficit concerns an attentional mechanism. To save the attentional hypothesis, we could, of course, abandon the claim that there is a single, general mechanism for directing resources at all levels of representation. However, the price of this move would be to give up the possibility of having a representation-independent characterisation of attentional mechanisms. Alternatively, we could assume that the underlying cause of NG's deficit is that it directly concerns the right half of representational space at the level of word-centred grapheme (or, object-centred, 3-D model) descriptions.

It must be emphasised at this point that the hypothesis advanced here is *not* that the proposed deficit concerns stored representations of words or operations of the lexical system in general. These hypotheses are clearly inconsistent with the reported results: NG's impairment involved words *and* nonwords in qualitatively similar ways. Furthermore, NG's performance with suffixed and pseudosuffixed letter strings shows that she was able to access lexical stems which she was unable to read correctly when these were presented unsuffixed. These results show that the deficit responsible for NG's spatially determined impairment in word recognition cannot concern stored lexical representations, nor can it concern the operations of the lexical access process itself. Instead, as proposed earlier, the deficit must concern the grapheme representations that serve as input to the lexical access process. Of course, this does not exclude the possibility that other patients' spatially determined impairment in word processing may in fact be the result of a deficit to the lexical system. However, no detailed evidence or argument has been presented in favour of such a possibility.

Finally, note that the conclusion reached here is not to be understood as a general claim about the underlying cause of "neglect" in general, as has been suggested by Bisiach et al. (1979) and De Renzi et al. (1970). The claim advanced here concerns only the spatially determined impairment described for NG: it is not a claim about all patients who may be classified clinically as presenting with "neglect". The specificity of the conclusion does not make it any less important, for it establishes that at least some "neglect" symptoms may result from direct damage to the information in a spatially defined part of a representational space. This position leaves open the possibility that the clinical symptoms of "neglect" in other patients may result from damage to other perceptual or attentional mechanisms (see Riddoch et al., this issue, for putative cases of attentional deficits).[15]

On the Preservation of Positional Encoding. Ellis et al. (1987: patient VB; see also Riddoch et al., this issue: patient JB; Young et al., this issue (Part 2, 1991): patient SP) have proposed that "neglect" may affect information differentially about the identity of the letters and the length of a word. Specifically, they suggest that a patient may be unable to recover information about the identity of the letters at the beginning or end of a word while retaining the ability to encode the overall length of the word. The evidence adduced in support of this claim is the striking correlation between stimulus and response length in reading errors. Thus, all three patients discussed by these authors, VB, JB, and SP, produced reading errors that were the same or nearly the same length as the misread stimulus word. This correlation was also observed for NG, and represents a significant factor to be accounted for by theories of visual word recognition. We do not think, however, that the proposed account of independent encoding of letter identity and word length provides the correct characterisation for the observed correlation. The reasons for this contention follow.

One objection is theoretical in nature: no motivated account has been offered of how the independently encoded information about word length and letter identities are integrated in the process of word recognition.

[15]We do not think that Riddoch et al.'s conclusion, that the performance of their two patients results from deficits to attentional mechanisms as opposed to representational spaces or computed representations, is supported by the evidence they cite. Thus, the fact that JB (with left-sided processing difficulties) continued to neglect the (physical) left side of words for mirror-reversed stimuli does not imply that the deficit is to attentional mechanisms: it could just as easily be due to a deficit to the left-of-centre part of the computed stimulus-centred shape representation. The fact that JB's pattern of performance rules out a deficit to the word-centred grapheme level does not imply that the form of damage must be to an attentional mechanism (although it could be!). The situation with case MO is equally indeterminate. As argued earlier, MO's performance suggests a deficit in processing at the level of the retino-centric feature map. The fact that this patient's "neglect" could be eliminated by allowing him to fixate position 1 of the letter string suggests that the observed reading impairment may result from his hemianopia (though not necessarily!). In short, the available evidence does not allow an unambiguous conclusion on whether the reading impairment of JB and MO is most plausibly explicated by assuming a representational or an attentional deficit at one or another level of the visual word recognition process.

Thus, let us suppose that information about word length and letter identities has been encoded for a visually presented word. How are these two types of information represented? Are they represented in the same spatial representation? How are the two sorts of information used in the word recognition process? In the absence of answers to questions such as these, we must remain sceptical about the value of the proposed claim.

Empirical objections can also be raised against the claim of independent encoding of word length and letter identities. Let us suppose that the basis for the observed correlation between the length of error responses and stimuli was that the patients were able to encode word length correctly despite their inability to process some of the letters at the right or left half of words adequately. Were this to be the case, we would expect the patients in question to perform normally in line bisection or length estimation tasks where the only relevant information for performing the task concerns the length of the line. However, as documented here, NG systematically underestimated length in line bisection tasks. Consequently, it does not seem likely that the observed correlation between response and stimulus length for NG results from the independent, correct encoding of word length.

The observed correlation between the length of error responses and stimuli receives a ready explanation within the model of visual word recognition we have proposed. Briefly, the correlation between stimulus and response length is a direct consequence of the assumption that graphemes are represented in a word-centred co-ordinate system, and thus necessarily encode positional information. The implicitly encoded position information is part of the information needed for lexical access. Thus, it is important that the grapheme <p> in the stimulus *pot* is the first grapheme and not the third as in *top*. And, on the assumption that lexical representations are activated in proportion to their similarity to the input grapheme representation, this guarantees that the lexical representations receiving the maximum activation will be those that share graphemes at the same relative positions. A direct consequence of these assumptions is that the lexical representations with maximum activation will have the same length as the stimulus. Now, when a spatially determined part of a grapheme representation is damaged (say part of the right half, as is the case for NG), the undamaged part of the grapheme representation will activate lexical representations that are (approximately) the same length as the stimulus.

To illustrate this claim consider the following example. Suppose that the patient is shown the stimulus *canter*. On the assumption that the right half of the grapheme representation of this stimulus is damaged, the only fully usable information is <can>. Why, then, does the patient not just produce "can" instead of the more likely responses "cannot", "cannon", "canned", "candle", "canopy", and so forth? The reason is to be found in the fact that the undamaged graphemes <c>, <a>, and <n> occupy particular positions in the representational space of the grapheme description: the grapheme <c> occupies the third position to the left of centre, the grapheme <a> occupies the second position to the left of centre, and the grapheme <n> occupies the first position to the left of centre. When these graphemes occupy the indicated positions they will maximally activate lexical representations with graphemes in the same positions relative to their centre—the lexical entries CANNON, CANDID, CANTOR, and so forth, and not CAN, CANE, CANNIBAL, CANTALOUPE, and so forth. Consequently, everything else being equal, the patient will produce a response that is the same length as the stimulus even though its length has not been encoded explicitly. Thus, the model of visual word recognition we have proposed can account for the observed correlation between stimulus and error response length without having to make unmotivated assumptions about the independent encoding of stimulus length.

Neighbourhood Effects. The model of visual word recognition we have proposed predicts that patients with damage at any of the three levels of visual processing will show significant effects of the density of the orthographic neighbourhood of words (as reported by Riddoch et al., this issue; Patterson & Wilson, this issue). That is, on the assumptions that (1) the representation at the grapheme description level activates in parallel (proportionally to the degree of graphemic overlap) all entries in the orthographic input lexicon with which it shares graphemes in specific positions, and (2) the lexical representation that is activated for further processing depends upon the degree of activation it receives from the grapheme representation and on its own threshold level of activation, we would expect that words with many "neighbours" would have a lower chance of being read correctly than words with few or no "neighbours" (particularly if at least one of the neighbours is higher in frequency than the stimulus, as we reported for NG in lists 3 and 4).

Along the same lines, we can account for inconsistencies in reported word length effects in reading by patients who make errors restricted to the left or right side of words. We have reported that NG made more errors on longer words, and that the mean length of correctly read words was significantly shorter than the mean length of words that were read incorrectly. Ellis et al. (1987) reported similar findings for VB. By contrast, Costello and Warrington (1987) reported that JOH made more errors on shorter words: and Patterson and Wilson (this issue) reported that their patient made errors *only* on short words. The discrepancy among the reported effects of word length is just as likely to reflect differences in the nature of the stimuli used by different investigators as it is to reflect differences in the nature of damage to different components of the visual word recognition process. Thus, for example, most of the longer words incorrectly read by NG were suffixed words. In the case of these stimuli, there is always at least one word in the lexicon that shares with the stimulus the left half of the letters (thus, for example, NG might well read *accommodated* as "accommodate", "accommodates", "accommodation", "accommodations", and so on, depending on the relative accessibility of these representations). On the other hand, a patient with left "neglect" would be very likely to read *accommodated* correctly, since no other word in the lexicon ends in *-odated*. Thus, an interpretation of word length effects cannot be undertaken without considering the neighbourhood structure of the stimuli used by different experimenters.

On Supposedly Contradictory Results. It has been claimed that there are contradictory results regarding the stages at which neglect is supposed to occur—the retinotopic encoding stage or some later stage of processing (Behrmann et al., in press; Mozer & Behrmann, Note 2). However, these supposedly contradictory results can only be taken as such if one assumes that neglect is a unitary phenomenon that can only affect a specific stage of processing. Once we abandon this theoretically unmotivated restriction about possible loci at which the spatially determined deficit can occur, the seemingly contradictory results referred to by these authors are no longer paradoxical—they simply reflect damage to different levels of the visual word recognition system.

The Lexicality Effect: Automatic and Top-down Processing? One final issue we consider briefly concerns the interpretation that has been offered for the lexicality effect in reading—better performance in reading words than nonwords—that has been observed in various cases of neglect (Brunn & Farah, Note 1; Sieroff, Pollatsek, & Posner, 1988). This result has led some authors (e.g. Sieroff et al., 1988) to conclude that words are processed automatically—without visual attention—whereas nonwords require attention for normal processing (see also Brunn & Farah, 1989; Mozer & Behrmann, Note 2). This conclusion has been reached in the context of a model that fails to distinguish between the several levels of representation within the pre-lexical access part of the visual word recognition system. Consequently, the interpretation of results has been based on the simple distinction between pre-lexical, letter processes (common to word and nonwords),[1n] and lexical processes (exclusively for words). In

---

[1n]That is, no account is offered of the levels of representation that are needed to solve the computational problems in word recognition. Such a step would have involved making explicit assumptions about the nature and organisation of the information computed at each stage of processing of the word recognition system. With such an account on hand, it would have been possible to evaluate the plausibility of claims about the presumed role of attentional mechanisms at some level of processing. The absence of such an explicit account makes statements such as (Siéroff et al., 1988; p. 427) ". . . spatial attention is unnecessary for access to the lexical network that produces a visual word form" not especially informative—we simply have a juxtaposition of two poorly understood concepts: spatial attention and lexical network.

this context, and on the assumption that neglect concerns the pre-lexical level of processing, the better performance in processing words versus nonwords invites the inference that words and nonwords are processed differently at this level of the word recognition system—hence the conclusion that words are processed "automatically" and that nonwords require "attention" for normal processing. This conclusion is problematic. Even ignoring problems of interpretation of the notion of "automatic processing", it is not necessary to invoke such a notion in order to account for the superior performance with words. As discussed earlier, better performance with words may only reflect the fact that, for word stimuli, the set of possible word responses (given partial information about the grapheme string) may be highly restricted, leading to many correct responses by default. Furthermore, since no hypothesis is offered about the structure of relevant visual processes at pre-lexical access stages, the claim that processing at these stages may be automatic or attentive, depending on the type of stimulus (word versus nonword), is hardly informative.

Some aspects of the performance of some patients with spatially determined reading impairments have given rise to claims about the role of "top-down" processing in reading (e.g. Behrmann et al., in press; Mozer & Behrmann, Note 2). It has been argued that words are read better because of the influence of top-down lexical effects. The motivation for this claim has concerned the superiority of word over nonword reading performance in "neglect" patients. If by "top-down" effects is meant no more than the fact that the availability of lexical representations in the course of reading constrains possible responses, leading to superior performance for words over nonwords, then this notion is superfluous—we have shown that one can instantiate lexical constraints in a strictly bottom-up model such as the one we have proposed in this paper. If by "top-down" effects is meant that lexical knowledge affects pre-lexical perceptual processes, then the evidence cited in its support is not adequate. To defend the latter claim it would have to be shown that the letter shapes and grapheme descriptions computed for word stimuli are *different* from those computed for nonwords. No such demonstration has been provided. And, since it is possible to explicate the word/nonword contrast in reading performance in a model with strictly bottom-up processing architecture (as shown here), we must reject as premature the introduction of overly powerful notions such as "top-down" processing.

Conclusion

In this report we have presented a relatively detailed multi-stage model of visual word recognition (and spelling). Explicit assumptions were made about the processing structure of early stages of visual form analysis, about the processing structure of the mechanisms of lexical access, and about the interaction between these two sets of processes. This model received considerable support through the extensive analysis of NG's reading and spelling performance. The model could account for all the main features of NG's performance: (1) the spatially determined reading deficit, involving only the right half of words, regardless of the topographic arrangement of stimuli; (2) the spatially determined spelling deficit, involving only the right half of words, regardless of the form of output (written and oral spelling and backward oral spelling); (3) the better performance reading words compared to nonwords, and the comparably poor performance in spelling words and nonwords; (4) the presence of a word frequency effect in reading but not spelling; and (5) the prefix and suffix effects on reading performance. The fact that highly detailed aspects of NG's performance could be given a clear interpretation in the context of the proposed model provides empirical support for the model. The proposed model is further supported by the fact that it could also account for the pattern of spatially determined reading deficits in other brain-damaged subjects described in the literature.

Although we have succeeded in providing a detailed account of a wide range of facts, there remain a number of problems which we have either not been able to address satisfactorily or completely omitted from consideration because they are too difficult to yield to analysis at this time. One puzzle concerns the assumption that graphemic information is represented in a spatially defined co-ordinate system. As already noted, this assumption is not unproblematic—it remains unclear how to reconcile the assumption of abstractness of graphemic information with the assumption that this information is arrayed in a spatially defined co-ordinate frame. A long-standing, important puzzle concerns the correlation between side of damage to representational space and side of hemispheric damage. The significance of this correlation remains a mystery despite the many efforts to clarify it (see DeRenzi, 1982; and papers in Jeannerod, 1987; for discussions of this problem). And, finally, there is the poorly understood relationship between the spatial-specificity of the deficit, especially its phenomenological character, and the unity of consciousness (see Bisiach, in press, for discussion). It is to be hoped that future research will find means of addressing these important theoretical issues. For now, we can do no more than to note that significant progress on these issues is unlikely to be forthcoming without the formulation of computationally explicit accounts of the processing stages that subserve visual word and object recognition. It is our hope that the work reported here has contributed to the elucidation of at least some of the general prc erties of the visual word recognition system.

Manuscript received 7 March 1990
Revised manuscript received 15 June 1990

REFERENCES

Allport (1989). Visual attention. In M. I. Posner (Ed.), *Foundations of cognitive science*. Cambridge, Mass.: M.I.T. Press.

Badecker, W. & Caramazza, A. (1987). *The Johns Hopkins Morphology Battery*. Baltimore, Maryland: The Johns Hopkins University.

Badecker, W., Hillis, A. E., & Caramazza, A. (1990). Lexical morphology and its role in the writing process: Evidence from a case of acquired dysgraphia. *Cognition*, *35*, 205–244.

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

Behrmann, M., Moscovitch, M., Black, S. E., & Mozer, M. C. (in press). Perceptual and conceptual factors in neglect: Personal and extrapersonal. *Brain*.

Bisiach, E. (in press). The (haunted) brain and consciousness. In A. J. Marcel & E. Bisiach (Eds.), *Consciousness and contemporary science*. Oxford: Oxford University Press.

Bisiach, E., Luzzatti, C., & Perani, D. (1979). Unilateral neglect, representational schema, and consciousness. *Brain*, *102*, 609–618.

Bisiach, E. & Vallar, G. (1989). Hemineglect in humans. In F. Boller & J. Graffman (Eds.), *Handbook of neuropsychology, Vol. 1*. North Holland: Elsevier Science Publishers, B.V., 195–222.

Caramazza, A. & Hillis, A. E. (1990). Internal spatial representation of written words: Evidence from unilateral neglect. *Nature*, *346*, 267–269.

Caramazza, A., Laudanna, A., & Romani, C. (1988). Lexical access and inflectional morphology. *Cognition*, *28*, 297–332.

Caramazza, A., Miceli, M., Villa, A., & Romani, C. (1987). The role of the graphemic buffer in spelling: Evidence from a case of acquired dysgraphia. *Cognition*, *26*, 59–85.

Caramazza, A. & Miceli, M. (1989). Orthographic structure, the graphemic buffer, and the spelling process. In C. von Euler, I. Lundberg, & G. Lennerstrand (Eds.), *Brain and reading*. London: MacMillan/Wenner-Gren International Symposium Series, 257–268.

Coltheart, M. (1981). Disorders of reading and their implications for models of normal reading. *Visible Language*, *15*, 245–286.

Costello, A. D. & Warrington, E. K. (1987). The dissociation of visuospatial neglect and neglect dyslexia. *Journal of Neurology, Neurosurgery, and Psychiatry*, *50*, 1110–1116.

Cummings, D. W. (1988). *American English spelling*. Baltimore: The Johns Hopkins University Press.

De Renzi, E. (1982). *Disorders of space exploration and cognition*. New York: Wiley.

De Renzi, E. & Faglioni, P. (1978). Normative data and screening power of a shortened version of the Token Test. *Cortex*, *14*, 41–49.

De Renzi, E., Faglioni, P., & Scotti, G. (1970). Hemispheric contribution to the exploration of space through the visual and tactile modality. *Cortex*, *6*, 191–203.

Duncan, J. (1987). Attention and reading: Wholes and parts in shape recognition. In M. Coltheart (Ed.), *Attention and performance*, *Vol. 12*. London: Lawrence Erlbaum Associates Ltd.

Egeth, H., Jonides, J., & Wall, S. (1972). Parallel processing of multi-element displays. *Cognitive Psychology*, *3*, 647–698.

Ellis, A. W. (1988). Normal writing processes and peripheral acquired dysgraphias. *Language and Cognitive Processes*, *3*, 99–127.

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

Ellis, A. W., Flude, B., & Young, A. (1987). "Afferent dysgraphia" in a patient and in normal subjects. *Cognitive Neuropsychology*, *4*, 465–486.

Fodor (1983). *The modularity of mind*. Cambridge, Mass.: M.I.T. Press.

Friedrich. F. J.. Walker. J. A.. & Posner. M. I. (1985). Effects of parietal lesions on visual matching: Implications for reading errors. *Cognitive Neuropsychology*. 2. 253–264.
Goodglass. H. & Kaplan. E. (1972). *The Boston Diagnostic Aphasia Examination*. Philadelphia. Penn.: Lea & Febiger.
Goodman. R. & Caramazza. A. (1986). *The Johns Hopkins Dyslexia Battery*. Baltimore. Maryland: The Johns Hopkins University.
Gordon. B. (1985). Subjective frequency and the lexical decision latency function: Implications for mechanisms of lexical access. *Journal of Memory and Language*. 24. 631–645.
Heilman. K. M.. Bowers. D.. Valenstein. E.. & Watson. R. T. (1985). Hemispace and hemispatial neglect. In M. Jeannerod (Ed.). *Neurophysiological and neuropsychological aspects of spatial neglect*. North Holland: Elsevier Science Publishers.
Hildreth. E. C. & Ullman. S. (1989). The computational study of vision. In M. I. Posner (Ed.). *Foundations of cognitive science*. Cambridge. Mass.: M.I.T. Press.
Hillis. A. E. & Caramazza. A. (1989). The graphemic buffer and attentional mechanisms. *Brain and Language*. 36. 208–235.
Hillis. A. E. & Caramazza. A. (1990). The effects of attentional deficits on reading and spelling. In A. Caramazza (Ed.). *Cognitive neuropsychology and neurolinguistics: Advances in models of language processing and impairment*. Hillsdale. N.J.: Lawrence Erlbaum Associates Inc.
Jeannerod. M. (Ed.) (1987). *Neuropsychological and physiological aspects of spatial neglect*. New York: Elsevier Science Publishers. 69–86.
Kinsbourne. M. (1970). A model for the mechanism of unilateral neglect of space. *Transactions of the American Neurological Association*. 95. 143.
Kinsbourne. M. & Warrington. E. (1962). A variety of reading disability associated with right hemisphere lesions. *Journal of Neurology. Neurosurgery. and Psychiatry*. 25. 334–339.
Kosslyn. S. M.. Flynn. R. A.. Amsterdam. J. B.. & Wang. G. (1990). Components of high-level vision: A cognitive neuroscience analysis and accounts of neurological syndromes. *Cognition*. 34. 203–277.
Marr. D. (1982). *Vision*. New York: W. H. Freeman & Co.
Marr. D. & Nishihara. H. K. (1978). Representation and recognition of the spatial organisation of three-dimensional shapes. *Proceedings of the Royal Society of London*. B200. 269–294.
McClelland. J. L. & Rumelhart. D. W. (1981). An interactive activation model of context effects in letter perception: Part I. An account of basic findings. *Psychological Review*. 88. 375–407.
Miceli. G. & Caramazza. A. (1988). Dissociation of inflectional and derivational morphology. *Brain and Language*. 35. 24–65.
Monk. A. F. (1985). Co-ordinate systems in visual word recognition. *Quarterly Journal of Experimental Psychology*. 37(A). 613–625.
Morton. J. (1969). The interaction of information in word recognition. *Psychological Review*. 76. 165–178.
O'Reagan. J. K. & Lévy-Schoen. A. (1987). Eye-movement strategy and tactics in word recognition and reading. In M. Coltheart (Ed.). *Attention and performance XIII: The psychology of reading*. Hillsdale. N.J.: Lawrence Erlbaum Associates Inc.
O'Reagan. J. K.. Lévy-Schoen. A.. Pynte. J.. & Brugaillère. B. (1984). Convenient fixation location within isolated words of different length and structure. *Journal of Experimental Psychology: Human Perception and Performance*. 9. 912–922.
Pashler. H. & Badgio. P. C. (1987). Attentional issues in the identification of alphanumeric characters. In M. Coltheart (Ed.). *Attention and performance XIII: The psychology of reading*. Hillsdale. N.J.: Lawrence Erlbaum Associates Inc.

Patterson, K. & Wilson, B. (1990). A ROSE is a ROSE or a NOSE: A deficit in initial letter identification. *Cognitive Neuropsychology*, *7*, 447–477.

Posner, M. I., Cohen, Y., & Rafal, R. D. (1982). Neural systems control of spatial orienting. *Philosophical Transactions of the Royal Society, London*, *298B*, 60–70.

Rapp, B. & Caramazza, A. (1991). Spatially determined deficits in letter and word processing. *Cognitive Neuropsychology*, this issue (Part 2).

Riddoch, M. J. & Humphreys, G. W. (1987). Perceptual and action systems in unilateral visual neglect. In M. Jeannerod (Ed.), *Neuropsychological and physiological aspects of spatial neglect*. New York: Elsevier Science Publishers, 151–181.

Riddoch, M. J., Humphreys, G. W., Cleton, P., & Fery, P. (1990). Interaction of attentional and lexical processes in neglect dyslexia. *Cognitive Neuropsychology*, *7*, 479–517.

Seymour, P. H. K. (1979). *Human visual cognition*. West Drayton: Collier Macmillan.

Siéroff, E., Pollatsek, A., & Posner, M. I. (1988). Recognition of visual letter strings following injury to the posterior visual spatial attentional system. *Cognitive Neuropsychology*, *5*, 427–450.

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

Taft, M. (1985). The decoding of words in lexical access: A review of the morphological approach. In D. Besner, T. G. Waller, & G. E. MacKinnon (Eds.), *Reading research: Advances in theory and practice, Vol. 5*. London: Academic Press.

Treisman, A. (1988). Features and objects: The fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology*, *40A*, 201–237.

Treisman, A. & Gelade, G. (1980). A feature-integration theory of attention. *Cognition and Psychology*, *12*, 97–136.

Venezky, R. L. (1970). *The structure of English orthography*. The Hague: Mouton.

Warrington, E. (1991). Right neglect dyslexia: A single case study. *Cognitive Neuropsychology*, this issue (Part 2).

Warrington, E. K. & Zangwill, O. L. (1957). A study of dyslexia. *Journal of Neurology, Neurosurgery, and Psychiatry*, *20*, 208–215.

Weschler, D. (1972). *The Weschler Memory Scale*. New York: The Psychological Corp.

Young, A. W., Newcombe, F., & Ellis, A. W. (1991). Different impairments contribute to neglect dyslexia. *Cognitive Neuropsychology*, this issue (Part 2).

REFERENCE NOTES

1. Brunn, J. L. & Farah, M. J. (1988). *The role of spatial attention in word reading: Further evidence from neglect patients*. Paper presented at the 26th Annual Meeting of the Academy of Aphasia: Montreal.
2. Mozer, M. C. & Behrmann, M. (1989). *On the interaction of selective attention and lexical knowledge: A connectionist account of neglect dyslexia*. Unpublished manuscript. Department of Computer Science, University of Colorado at Boulder.

APPENDIX A

NG's Direct Copying