

Thesis

University of Stirling 07 1708

Daniel P. F. Sturdy

---

**THE REPRESENTATION OF LETTER  
STRINGS: PSYCHOLOGICAL EVIDENCE AND  
COMPUTATIONAL MODELS**

---

Submitted for the degree of Ph.D.

December, 1990



---

# TABLE OF CONTENTS

---

ACKNOWLEDGEMENTS	i
ABSTRACT	ii
CHAPTER 1: INTRODUCTION	1
1. 1 Overview	1
1. 2 The issue	3
1. 3 The representation of relational information	19
1. 4 Psychological and neuropsychological evidence	34
1. 5 Experimental methods and predictions	47
CHAPTER 2: WHOLE-PART RELATIONS IN THE TRANSFER OF TRAINING	50
2. 1 Introduction	50
2. 2 Experiment 1: Fixed and Moved-part transfer	51
2. 3 Simulation 1: Interactive Activation	64
2. 4 Simulation 2: MAPPER	65
2. 5 General discussion	67
CHAPTER 3: LOCAL AND NON-LOCAL RELATIONSHIPS	75
3. 1 Introduction	75
3. 2 Experiment 2: Fixed-part transfer in local and non-local parts	77
3. 3 Experiment 3: Part-part learning	83

3. 4 Experiment 4: Part-whole learning in local and non-local parts	88
3. 5 Experiment 5: Moved-part transfer in local and non-local parts	93
<b>CHAPTER 4: MIGRATION EXPERIMENTS</b>	<b>99</b>
4. 1 Introduction	99
4. 2 Experiment 6: Positional information	99
4. 3 Experiment 7: Across-string migrations	111
4. 4 Experiment 8: Detecting and correcting spelling mistakes	117
<b>CHAPTER 5: GENERAL DISCUSSION</b>	<b>130</b>
5. 1 Summary of findings	130
5. 2 Theoretical implications	132
5. 3 Modelling graphemic representational space	135
<b>BIBLIOGRAPHY</b>	<b>145</b>
<b>APPENDICES</b>	<b>173</b>

---

# ACKNOWLEDGEMENTS

---

My sincerest thanks go to my supervisor, Bill Phillips, without whom none of this would have been possible.

Sincere thanks for a stimulating workplace go to all the members of the Centre for Cognitive and Computational Neuroscience at Stirling, especially, and alphabetically, to Roland Baddeley, Lorna Begg, Peter Cahusac, Steven Dakin, Peter Hancock, Ian Patterson, and Peter Richardson.

Thanks go to many for technical support and advice, most memorably to Lindsay Wilson, Gerry Orchard, Bruce Sutherland, Bob Lavery, and Randal Macdonald.

Special thanks to family and flatmates, to Fran, Lise, Rachel, Martin, Conor, Paddy, and Heather.

For Barbara.

---

## ABSTRACT

---

Two ways of representing the spatial arrangement of letters in letter-strings are distinguished. In part-whole representations, the relationship of a letter to the letter-string as a whole is encoded. In part-part representations, the relationships of a letter to other letters in the string are encoded. Computational models of word perception typically use the former, but part-part representations are a very general feature of some neurocomputational models. Experiments are reported that examine for nonword and word wholes the representations used to encode their constituent parts; the first five experiments use measures of facilitation to infer encoding type, the next three primarily use error measures.

Experiment 1 shows that when a part of a recently learned letter-string is maintained in a briefly-presented test string, the test string is more accurately reported, showing perceptual transfer of training. No significant difference in the amount of transfer is found between maintaining the part in the same position (fixed-part) in the string and maintaining the part in a different position (moved-part) in the string. It is argued that this confirms part-part theories because transfer was obtained when only inter-letter relationships are maintained. Experiment 1 simulated on two implementations of part-whole theories shows that they fail to produce the obtained pattern of performance. This indicates that part-whole relational encoding is not a major part of the representations mediating these transfer effects. Experiment 2 replicates the fixed-part transfer and shows that it is restricted to parts made of adjacent letters. Experiments 3 and 4 use a prototype-extraction paradigm to show that novel parts made of adjacent letters are easier to learn than parts made of non-adjacent letters. Experiment 5 replicates the moved-part transfer and shows that it is restricted to parts made of adjacent letters. These results show that the major inter-letter relationships encoded are between neighbouring letters.

These first five results are taken as showing that pre-processing of the image to provide position-in-the-string information is not important for the representations that produce transfer. It is suggested that modelling the input to the graphemic input lexicon as the Primal Sketch of the image is more appropriate. In particular, realistic early vision algorithms such as MIRAGE appear to be potentially capable of modelling the results obtained.

Experiment 6 shows that reports of letters in nonwords have gradients of positional accuracy, with most positional errors occurring close to the correct position. Experiment 7 finds that migrations into the report of the second of two briefly-presented nonwords from the first nonword do not always maintain position though many do. Experiment 8 involved the presentation of mis-spelled words preceded by nonwords that either encouraged the detection of the mis-spelling or its lexicalisation. Lexicalisation responses involve the migration of a letter from the preceding string. These occur when primed by the lexicalisation letter in the same, but not in moved, positions in the first string, but only when presented in the context of neighbouring letters. Detection of mis-spelling shows both facilitation and inhibition. Facilitation is obtained with the part in moved positions in the source string but not in the same position, in which case inhibition is found. Facilitation is also obtained by prior presentation of the misspelled word or prior presentation of the correctly spelled word. These results are interpreted as showing that facilitation is obtained when the facilitating part of the preceding string either fully or minimally activates a representation of the word mis-spelled on second presentation. Partial activation of the word produces inhibition.

The results suggest that part-whole encoding is used for letters in familiar wholes, while part-part encoding is used for letters in unfamiliar wholes. This conclusion is used to motivate a model of the organisation and access of graphemic representations in which the concept of scale plays an important role. The model is extended to other tasks involving visually presented words and nonwords and a brief account of the major findings attempted. Finally some extensions of the model to the domain of object perception are outlined.

---

# INTRODUCTION

---

## 1. 1 Overview

Information is present in the environment at a variety of scales<sup>1</sup>. To survive in natural environments it is necessary to make decisions about perceptual events rapidly and efficiently. The information available for making these decisions is often either incomplete or more detailed than necessary for a correct decision. Useful increases in efficiency can be achieved by processing the information initially at a coarse scale and continuing to a fine-grained analysis only when necessary. This is particularly true of information about spatial layout, the positions of items within an image (Watt, 1988). Two questions then arise. How are spatial positions computed and represented, *within* any of a variety of scales? And how is spatial position information combined *across* scales? This is the most general formulation of the issue at hand. I will treat the former as a question of inter-item relations, the latter as a question of within-item relations, referring to them loosely as part-part and part-whole relationships respectively. The issue then becomes: When are part-part and when are part-whole relationships used in processing, and what is the relationship between them?

---

<sup>1</sup> Key words will be printed in bold and defined, where appropriate, in footnotes. Other footnotes offer developments of interesting points that are not crucial to the argument. By scale, I mean both spatial scale, as measured by the types of information extracted by band-pass filters of varying sizes, loosely correlated with spatial frequency, and a symbolic (representational) scale, as measured by the types of information present in different sizes of symbol structures, perhaps correlated with temporal frequency information. In language, for example, letters, morphemes, words, sentences, paragraphs, and texts, are all structures at different scales.

The issue phrased in this way is very generally concerned with spatial relations, but it is differences between different types of objects that are interesting in determining which spatial relations are relevant. For some objects the information present at coarse scales differs qualitatively from the information present at fine scales. Some objects have global symmetry, for example. For other objects the coarse information is more fully specified by the fine information. Both cases are true of verbal objects. Speech, for example, has global patterns of intonation and stress determined by the whole discourse rather than by the words that constitute the sentence, while the overall shape of a written word, for example, is more fully determined by the shape of the word's individual letters. Thus written words are good candidate visual objects for investigating how representations combine information over different scales. Words and letters provide an a priori clear distinction between coarse and fine scales, because words are made out of letters. Words are perceptual wholes at a coarse scale, and letters are parts<sup>2</sup> of word-wholes, but letters are perceptual wholes at a fine scale. This simple distinction raises a number of possibilities: the research reported in this thesis tests theories of the representation and processing of the relationships (a) between letter-wholes within word-wholes; (b) between letter-parts within word-wholes; (c) between letter-parts and word-wholes; and (d) between letter-wholes and word-wholes.

Several, neurally motivated<sup>3</sup>, computational models of word and object perception have recently been developed. Some make strong claims about the representation of relationships between parts and wholes. These claims can be tested by asking how the models deal with changed relationships among the parts of familiar objects/words. Some of the models can be tested by running simulations of their performance to provide quantitative data for comparison with the performance of humans using similar stimuli under various experimental conditions. The conditions most likely to test the theories vary, independently, part-whole and part-part relationships. The paradigms used to provide these test conditions

---

<sup>2</sup> Defining parts and wholes is problematic. In this thesis, individual words or nonwords will be treated as the wholes. Letters are generally considered to be the relevant parts but clearly letter groups, especially syllables and morphemes, are also candidate parts. Different languages may differ in the relative salience of these candidate parts.

<sup>3</sup> In the tradition of connectionist, or parallel distributed processing models.



are based on the following assumption: to the extent that the representations produced by a target stimulus are functionally similar to those elicited by a recently seen or well-learned stimulus, transfer to the processing of that subsequent stimulus is likely to occur. In other words perceptual transfer from one representation (or set of representations) to another is likely depending on the similarity of those representations. The amount and the nature of the transfer of processing can then be used to make inferences about the nature of the underlying similarities between the representations.

The aims of this research are summarised as follows: (a) to investigate experimentally facilitation and interference in conditions in which part-whole and part-part relationships vary independently; (b) to derive theoretically, and in some cases computationally, predictions from models of relational representations about their performance under the same conditions, and to compare the predictions with the experimental data.

The structure of the rest of this chapter is this: §1. 2 sets out the issue; §1. 3 looks in detail at theories of the representation of relational information, with examples taken from psychological and computational models of word and object recognition; §1. 4 examines the psychological, and neuropsychological evidence for and against the models; §1. 5 derives testable predictions from the different representational schemes and outlines the rationale for the experimental methods chosen for the research reported in this thesis. The experimental results are presented in Chapters 2, 3, and 4; Chapter 5 assesses the implications of the results for the models discussed in Chapter 1, and offers a, speculative, theoretical development that could be implemented by one of the models.

## 1. 2 The issue

The question of how spatial relations are represented is posed here within the problem domain known, by analogy with pattern recognition and object recognition, as *visual word recognition*. Research is characterised by such questions as: How are written words recognised? What is the nature of the representations that mediate word recognition? How are (interpreted) word representations computed from (uninterpreted) sensory

representations?<sup>4</sup> A bewildering variety of “word recognition” tasks have been used, from the ecologically valid tasks of naming written words aloud or writing them down, through tasks such as proof-reading for spelling mistakes, which are commonly performed, to tasks such as categorising a stimulus as a word or nonword (lexical decision), which are rarely, if ever, naturally encountered. Seidenberg (e.g., 1985), in particular, has emphasised performance differences between paradigms (between naming and lexical decision, for example), but these differences are often ignored.

The treatment offered here takes a different approach from Seidenberg's. Instead of a gross distinction between paradigms, I will use a distinction in terms of task demands. **Whole-word** tasks require that the whole stimulus be processed; lexical decision and naming, are examples. **Letter-level** tasks, on the other hand, require only that part of the stimulus be processed. Examples are the letter search task, where a target letter is presented and subjects decide whether that letter was present in a preceding, subsequent, or simultaneous letter-string. One intensively studied task is the two-alternative forced-choice (2AFC) discrimination, in which subjects decide which of two letters was present in a preceding letter-string. These two types of task differ in that in the whole-stimulus task the letters are processed *as parts of a whole*, while in the part-stimulus task the letters are processed as *wholes within a larger whole*.<sup>5</sup> Two examples of the effects of this distinction are presented, the second is highly speculative.

(i) Strong effects of word frequency are well established for whole-word tasks, such as lexical decision (e.g., Gernsbacher, 1984; Gordon, 1983; 1985), but not for letter-level

---

<sup>4</sup> These questions seem straightforward enough (if unanswered), but they may be misleading, because they assume that we know what is meant by the term “recognition”. Recognition is usually evidenced by successful performance of a task designed to test ... word recognition.

<sup>5</sup> To my knowledge this distinction has not been made before in these terms. An analogous distinction is made for the spelling-sound mapping between addressed and assembled phonology (Norris & Brown, 1985; Patterson & Morton, 1985). In the former, phonology is derived via word-wholes, the constituents of which are treated as parts. In the latter, phonology is derived via sub-lexical clusters, treated as wholes within a larger whole. Kimchi and Palmer (1982; 1985) make a similar distinction between processing parts of an overall pattern either as texture or as shapes in their own right.

tasks: 2AFC, for example, shows no effect of word frequency (Gunther, Gfroerer, & Weiss, 1984). Letter frequency, on the other hand, (measured by letter-in-position frequencies) has no effect on whole-word tasks such as lexical decision (Coltheart, Davelaar, Jonasson, & Besner, 1977; Gernsbacher, 1984), but strong effects on 2AFC (McClelland, 1976; McClelland & Johnston, 1977; Rumelhart & McClelland, 1982) and letter-in-word search (Mason, 1975).

(ii) A double dissociation has sometimes been claimed between impaired word-level processing and impaired letter-level processing, but the dissociation may be more subtle than that. *Letter-by-letter readers*, although a heterogeneous group, are generally impaired at whole-word processing (hence their description, and alternative classification as *word-form dyslexics*; e.g., Warrington & Shallice, 1980; Patterson & Kay, 1982). Because they usually read nonwords as accurately as words, it is claimed that word-wholeness has no influence on their performance. Recent evidence from Bub (1990)<sup>6</sup> shows that this is not always the case, because performance in 2AFC can show the normal word-nonword superiority. Processing letters as wholes within word-wholes may be intact in the presence of impaired processing of letters as parts of word-wholes. The converse disorder has not been named but appears to be present in some deep dyslexic patients (e.g., Howard, 1987; Van Lancker, 1990). These reports indicate relative preservation of performance on whole-word tasks compared to whole-letter tasks<sup>7</sup>. Critical tests, such as 2AFC and letter-in-word search, have not been done, but the evidence suggests relatively better preserved processing of letters as *parts* of words than as *wholes* within words. Processing parts of wholes should be dependent on the quality of the whole to a greater extent than processing wholes within word-wholes; this may explain these patients' complete inability to read nonwords and their difficulty with visually disrupted words.

This distinction will be used as the basis for a discussion of the processing of spatial

---

<sup>6</sup> And cf. related evidence from Coslett and Saffran (1989) and Shallice & Saffran (1986) showing preserved categorisation in the absence of explicit naming.

<sup>7</sup> The whole-letter task is the cross-case letter matching of a letter target to one of a group of four letters; this is very similar to the letter-in-word search task.

relations; the different types of representation which could be implicated in the processing of spatial relations are considered next. Clearly models of object/word perception are only relevant to the question of how spatial relations are processed if the same representations plausibly underlie both the processing of identity (or recognition) and the processing of spatial locations.

### 1. 2. 1 Spatial relations

To see how spatial relations apply to word recognition consider first Harris and Coltheart's (1986) definition of word recognition and the lexicon:

We have to learn to recognise the sounds, *spellings* and meanings of individual words, and to store this information in such a way that we can call upon it when we encounter ... written words. Since dictionaries also contain information about the orthography (i.e., spelling), phonology (i.e., pronunciation) and semantics (i.e., meaning) of words, terms such as ... 'mental lexicon' ... have been used to refer to the internalised system of knowledge we use when we perceive ... words (p. 135; my italics).

The idea that word recognition involves recognising a word's spelling is crucial because it assumes that words are not recognised by their whole-word visual characteristics, and thus that word representations are built out of representations of letter-wholes. Watt (1988) is more explicit:

[In] reading... the necessary information is the set and sequence of character identities. Spatial layout information in excess of the sequence (treating the space between words as a character) is not required (p. 92).

There is no doubt that increasing ability at spelling is a large part of the development of being able to read. The use of alphabetic rather than logographic scripts means that changes in the ordering of letters (one aspect of spelling) play a role as large as the role of word order within a sentence in determining symbol meanings. This possibility is not excluded from logographic scripts in which each symbol is constructed from a number of radicals, for example, but it is not clear that spatial relations between these elements are structured in such an organised way as in alphabetic scripts. To give an example from English, the words

TAR, ART, and RAT are all made out of the same letters and their status as different words depends solely on the different orderings of their letters. The reason for this is clear: logographic scripts make no attempt at representing sounds, while alphabetic scripts try to represent visually the phonemes of spoken language. To do this, temporal order information, which specifies spoken words, is converted directly into spatial order information, "*when* is recoded as *where*" (Mason, Pilkington, & Brandau, 1981). There is some evidence that reading ability correlates with the ability to deal with "where" information, both in simple perceptual tasks (Mason, 1980) and in tasks in which symbol order determines the required response (Mason, Pilkington, & Brandau, 1981).

However, spatial order, and therefore spelling, can be treated in two different ways. In one, each letter is identified as occupying a particular position within the word. In the other, each letter is identified as coming before or after another letter or letters. These are very different ways of representing spatial relations: Watt (in press) refers to them as a "*direct code for position*" and a "*relational code for position*". The direct code represents directly the spatial relation of a letter to a whole, or overall frame; the indirect code represents the relations between letters in a word.

A similar distinction is made in theories of the representation of order in other modalities and tasks<sup>8</sup>. Detailed computational models of the representation of order information in short-term memory have been developed (discussed in McNicol & Heathcote, 1986); a characteristic distinction in these models is between *position-dependent* codes and *context-sensitive* codes, i.e., between models which *tag* items with order information and models which represent inter-item *associations*. In tag models (e.g., Ratcliff, 1987) position is encoded as an absolute figure on some scale, in associative models (e.g., Murdock, 1983) inter-item relations are encoded.

The situation is more complicated than this, because McNicol and Heathcote (1986) distinguish both tag and associative models from a third possibility, which they describe as

---

<sup>8</sup> Order information has been extensively examined in short-term memory tasks, typically with auditory presentation. This research is not covered in detail here, though relevant findings are briefly discussed in subsequent sections.

"recoding information into chunks". The difference is that tag and associative models both represent item and order information independently, whereas recoding into larger-size chunks entails the joint treatment, or dependency, of at least some item and order information. In the terms of §1. 2, tag models are encoding whole-letter to larger-whole relationships, associative models are encoding whole-letter to whole-letter relationships, and recoding models are encoding letter-part to larger-whole relationships. McNicol and Heathcote (1986) find that dependency models better cope with short-term memory research results, but only with verbal stimuli which are easily recoded into larger chunks. This finding suggests that spatial relations between letter-wholes are not the whole story because processing wholes can be larger than the individual letter.

As suggested in §1. 1, because words are wholes at a coarse scale of representation they can potentially be recognised on the basis of word-whole features that are not present in the constituent letter features. Word length is a plausible example of a useful whole-word feature<sup>9</sup>. Clearly the possibility that words can be treated as wholes has strong implications for spatial relational encoding: "it is possible that the visual system deals differently with the spatial arrangement of object parts, and the relative position of whole objects" (Duncan, 1987, p. 42). In particular, letters within a word could be represented either as wholes within a larger whole, as in the tag model, or as wholes in relation to other wholes, or as parts of a whole. In the latter case word-based coordinates are imposed on the representations of the parts, but where does the word-based frame come from? Is it part of the internalised store of knowledge about visual symbols, or part of a low-level visual description of the image? Questions about which representations are implicated by different ways of coding spatial relations are discussed next.

---

<sup>9</sup> Word length is useful because, given that the number of different basic letter-symbols is limited, it allows massive increase in the number of unique letter-combinations. If words were all the same length, say four letters, then the number of possible letter-combinations would be limited, to 456,976 in this case. Letter-order is important because if different orderings of the same letters are treated as identical, then the number of unique letter-strings is greatly reduced, to 15,000 in this example. Even so, the subset of phonotactically legal members of the larger letter-combination set is much smaller than 456,976, and an order of magnitude smaller than the number of words in the language. This limitation is overcome by allowing words that differ in length as well as spatial ordering.

## 1. 2. 2 Representational domains

Some initial distinctions need to be introduced. As a first step, the nature of the problem requires definition. The treatment offered here is to model word recognition as an input-output mapping (c.f., Allport, 1987). The input is the grey-level representation over the retina, caused by the stimulus. At least four potential target outputs are available: semantics (the meaning of the stimulus), input and output phonological descriptions (how the stimulus sounds, and is pronounced), and graphic output descriptions (how the stimulus is written down). These descriptions will be treated uniformly, as output descriptions into which grey-level representations have to be mapped.

Most models of object and word recognition assume that this mapping is too complex to be achieved in one step<sup>10</sup>, and thus that some intermediate representations have to be computed. These intermediate representations are commonly divided into two categories, referred to as sensory representations and object- or word-based representations. Sensory representations are often assumed to represent position in coordinates determined by some aspect of the viewer (e.g., retina, head, body), whereas object representations are assumed to represent position in coordinates that are independent of the viewer.

The distinction between sensory and object representations is adopted here, though it should be emphasised that it is only a heuristic which hides many complexities. Phillips (1974; 1983) summarises the main empirical differences between the two domains: (1) *Capacity*: very much smaller for object than sensory representations; (2) *Durability*: fractions of a second for sensory representations, fractions of a minute for object representations; (3) *Pattern complexity*: much smaller effects on sensory than object representations; (4) *Masking*: much larger effects on sensory than object representations; (5) *Spatial restriction*: sensory representations appear to be tied to positions, object representations generalise over position.

---

<sup>10</sup> With the notable exception of "direct" perceptual theorists (e.g., Gibson, 1979). I am not, however, aware of detailed applications of this approach to word recognition.

Some examples of this distinction are presented in Table 1.1.

Models	Sensory Domain	Object/Word Domain
MAPPER (Hinton, 1981a, b, and c)	Retinotopic frame	Object-centred frame
Marr (e.g., 1982)	Primal Sketch & 2 1/2D Sketch	3-D model representations
Phillips (1983)	Sensory memory	Short-term visual memory
Monk (1985)	Retinal coordinates	Word-centred coordinates
Four-frames (Feldman, 1985)	Retinotopic frame	Stable-feature frame & World knowledge formulary
Dynamic-link (von der Malsburg & Bienenstock, 1986)	Layer 1	Layer 2
Howard (1987)	Visual analysis	Abstract Grapheme units & Word recognition units
De Yoe and van Essen (1988)	Sensory cues	Inferred object attributes
Humphreys and Bruce (1989)	Image properties	Object properties
Phillips, Hancock, Wilson, & Smith (1989)	Sensory data	Object descriptions
Kosslyn, Flynn, Amsterdam, and Wang (1990)	Visual buffer & Pre-processing	Pattern activation & Object representations

**Table 1.1.** Examples of the distinction between sensory and object-based representations.

The suggestion that object representations generalise over position is important here, because one interpretation of it suggests that object representations “refer to structure and position separately” (Phillips, 1983), allowing the representation of structural information to generalise over position. This interpretation is discussed in §1. 2. 2. 2, but first I discuss my treatment of sensory representation. Section §1. 2. 2. 2 discusses object representations in general, while §1. 2. 2. 3 discusses specifically verbal object representations. I describe the sensory domain as providing a visual input description (VIP), because I treat it as the input to the mapping into the various target outputs. I treat the word-based domain as an intermediate representation through which the mapping from VIP into a target output is performed. Loosely conceived, this intermediate representation approximates to the level of hidden units in a three layer back-propagation neural network. I will refer to it as the graphemic representational level.



### 1. 2. 2. 1 *The visual input description: MIRAGE*

This level of representation constructs a description of the useful deviations from randomness in the grey-level image, perhaps in terms of primitive features such as oriented edges, corners, line stops, their connections, size, and colour. This description has come to be referred to as the Primal Sketch (after Marr, 1982). MIRAGE is a low-level algorithm designed to provide a Primal Sketch that is robust and information rich (for its development see Watt & Morgan, 1985; Watt, 1987; 1988). I take this as a state-of-the-art model of the Primal Sketch, because it is motivated by psychophysical evidence and computational analysis. The MIRAGE process that constructs the Primal Sketch uses multiple spatial filters arranged at a variety of spatial scales, but its details are not at issue here. Its most important implications for subsequent processing are outlined by Watt (1988) as follows:

1. At the largest scale in operation at a particular moment it [MIRAGE] computes spatial positions.
2. At finer scales, if present, a statistical representation is applied.
3. Time permitting, the largest spatial filters are switched out, adding finer detail to the representation of spatial position. Filter switches are all-or-none.
4. The starting and terminating values of the largest spatial scale to be in operation can be set in advance by a high-level government (p.140).

This description will be treated as the database from which object representations are constructed in the course of task-specific mappings into target output representations. Some of the implications for the word recognition literature, and the usefulness of treating the VIP as the MIRAGE Primal Sketch, are discussed in Chapter 5. For the moment, the salient points of the VIP include those identified by Phillips (1983) for non-verbal objects that have been generally supported by research with letter-strings (under the guise of work on "iconic" or sensory memory):

(1) *Capacity*: the cued partial-report superiority effect (Sperling, 1960) shows that the amount of information extracted from briefly presented displays is far larger than can be

reported, presumably because of limitations on graphemic representational capacity.

(2) *Duration*: partial-report superiority lasts only for 250 ms (Sperling, 1960)<sup>11</sup>.

(3) *Masking*: partial-report superiority is abolished by backward noise-masking the displays (Averbach & Coriell, 1961; Sperling, 1960). However, it is clear that graphemic representations can also be masked (e.g., Michaels & Turvey, 1979; Taylor & Chabot, 1978) because masks have different effects on words and nonwords, but this can be explained by distinguishing between two different masking effects (cf. Ganz, 1975; Richman & Simon, 1989).

*Integration masking* is produced by noise-masks, is maximal at very short target onset-mask intervals, can be produced by pattern-masks when the mask is much brighter than the target, and is also produced by forward-masking (Turvey, 1973); the suggestion is that this reflects poor temporal resolution of briefly presented stimuli, through the integration of their sensory representations.

*Interruption masking* is much stronger with pattern-masks than with bright and dim noise-masks, but is only found at longer target-mask intervals, maximally with 40-50 ms between target and mask onset (Turvey, 1973); the suggestion is that these masks interrupt the attainment of a representation that takes about 50 ms to construct. Evidence from another paradigm suggests that text can be read normally with 50 ms masked presentations (Rayner, Inhoff, Morrison, Slowiaczek, & Bertera, 1981), and thus that graphemic representations are constructed within 50 ms of stimulus onset. This means that the two masking effects are operating on different representations, and that only graphemic representations show interruption masking. This then explains why different masks differentially effect words and nonwords.

---

<sup>11</sup> The advantage in sequential letter matching for same-case (A-A) over changed-case (A-a) letters can last for seconds (e.g., Kroll et al., 1970; Parks et al., 1972; Posner & Keele, 1967; Posner et al., 1969), but it is not clear whether this represents persistence of VIP representations or of different graphemic representations. The greatly reduced stimulus complexity in the letter-matching task may enable an explanation in terms of greater clarity of VIP representations to explain the difference in duration estimates without invoking graphemic representations to explain this result.

### 1. 2. 2. 2 *Object representations: what is represented?*

The claim (in §1. 2. 2) that structural information and position information are represented separately<sup>12</sup> has important implications. Representations of structural information, structural descriptions, are assumed by Phillips (1983) to:

represent the major structural features of objects, such as what parts they have and how these relate... To be useful such representations must include more than can be seen from any particular view (e.g., the unseen sides of objects) (p. 296).

Evidence that such structural descriptions are used comes primarily from mental rotation experiments (Hinton, 1979; Hinton & Parsons, 1981), but not exclusively so (Neisser & Kerr, 1973; Phillips, Hobbs, & Pratt, 1978). However, the notion of structural descriptions raises an interesting problem<sup>13</sup>. By definition, structural descriptions represent spatial relations: Hinton (1981a) describes hierarchical structural descriptions as containing

---

<sup>12</sup> The separate representation of structure and position is also apparent in the distinction between "what" and "where" processing drawn by some neurophysiologists (e.g., Mishkin, 1982). This distinction is based on the findings that inferotemporal cortex shows selective responses to object identities, irrespective of their position (e.g., Gross, Rocha-Miranda, and Bender, 1972), and that posterior parietal cortex is specialised for the processing of spatial relations (e.g., Wurtz, Goldberg, and Robinson, 1982).

<sup>13</sup> A side issue raised by this distinction is that because structural descriptions of novel visual inputs can be constructed (Phillips, 1983), achieving a structural description is not the same as object recognition. One solution is to distinguish *structural* representations from *class* representations (Marr, 1982). Neuropsychological evidence is pertinent: some agnosics can recognise objects, but only when the view is straightforward (Warrington & Taylor, 1973; Humphreys & Riddoch, 1984); others cannot recognize even these simple views, but can match different views of the same objects they fail to recognize (Warrington, 1975). This distinction refines the notion of recognition: representing an object's structure can be thought of as object recognition or description; classifying an object can be thought of as object identification (cf. Kosslyn et al., 1990); the evidence suggests that these dissociate. An alternative interpretation, however, is possible in terms of the distinction introduced in §1. 2: if the cross-views matching task requires processing the parts of the object as distinct wholes in order to achieve matches, then patients deficient on this task may be showing intact processing of wholes with a deficit in processing object parts as wholes within a larger whole; the recognition deficit may be the converse, impaired processing of object parts as parts of a whole with intact processing of parts as wholes within a larger whole.

a node for each object that is linked to lower-level nodes for its parts. These lower-level nodes, in turn, are linked to nodes for their parts, and so on until a level of primitive entities like edge-segments is reached. Each node in a structural description has its own object-based frame of reference, and each link between two nodes is labelled with the spatial relationship between their two object-based frames (p. 1092).

But if the nodes in the structural description continue down to edge-segments, why do they not also continue up to whole scenes? If scenes, or collections of objects are represented in this way, though, each object or object part's position in the scene is also represented by the structural description. If positions are encoded in structural descriptions, then why are structure and position represented separately? Phillips (1983) proposes that object descriptions include the separate representation of two types of position information, egocentric and exocentric position. Egocentric positions are encoded relative to the viewer's body, and are useful for reaching for, or moving towards, objects. Exocentric positions are encoded in terms of the spatial relations of the objects in the scene, and are useful for generalising over changes in view (e.g., Rieser, 1989).

So if both structural descriptions and exocentric positions encode the relations of objects in the scene, then why duplicate the encoding? A good answer cannot appeal to a particular size at which structural descriptions stop and exocentric position coding takes over because of the arbitrary definition of an 'object' with respect to size. One possibility is that the two descriptions encode position in different ways: in structural descriptions an item's position is defined in relation to the larger item of which it is a part, whereas exocentric coding of spatial relationships relates each item to any of the other items, irrespective of whether the items are parts of each other. Hinton (1981a) reaches a slightly different conclusion based on a re-analysis of what a representation of an object's structure is representing.

There is little evidence that the whole of a complex structural description is actively represented at the same time. It may well be that our attention flits between levels and that at each moment, we only focus on one node, i.e. we impose the object-based frame appropriate for this node and form a Gestalt for it... One important difference between a hierarchical structural description and a hierarchy of object-based feature units is that each link between nodes in the structural description is labelled with an explicit spatial relationship, whereas there are no

explicit representations of the spatial relationships between the various object-based features. An object-based feature is activated by the *combination* of a particular feature type with a particular relationship to the global object-based frame of reference... Higher-level feature units can be activated directly by combinations of lower-level ones. They do not need to check the relationships between these lower-level features, because the relationships are implicitly encoded by which of the lower-level units are active (p. 1092).

His definition of a Gestalt, "a coherent organisation of the parts of a figure into a perceptual whole which transcends the individual parts", raises a problem for the idea that only one Gestalt is active at any moment: "How can there be a Gestalt for the whole without Gestalts for the parts also being present?" His solution:

There are two quite different ways of binding together the shape and other properties<sup>14</sup> of a particular instance in a network of neuronlike units. When an instance is perceived as a Gestalt, the method of simultaneity can be used.<sup>15</sup> This allows the very same active units to represent the shape of an instance whatever its other properties. When an instance is seen as a constituent of a larger Gestalt, however, the multi-dimensional method is used.<sup>16</sup> This allows many constituents to be coded at once, and it allows the effects of each constituent to depend on its particular parameter values relative to the whole. The representation of an instance when it is seen as a Gestalt is therefore quite different from its representation when it is seen as a constituent of some larger whole. The Gestalt for the whole does not in any way involve the Gestalts for its parts (Hinton, 1981, p. 1093).

Both this solution and that offered by Phillips (1983) make the claim that object representations include a description of object structure in which parts are represented in relation to the whole object; they differ, however, in their claims about the representation of the position of the whole: Hinton suggests that wholes are represented in terms of an unspecified coordinate system, presumably either viewer-centred or scene-based; Phillips suggests that wholes are represented both in relation to the viewer and in relation to other

---

<sup>14</sup> Including position.

<sup>15</sup> This is the simultaneous activation of separate representations of the shape and the position of an instance.

<sup>16</sup> This is the use of units that encode the conjunction of particular shapes in particular positions.

wholes. The description of the solution in these terms makes clear its similarity to the suggestion offered in §1. 2, that letters can be processed either as parts of a whole or as wholes within a whole, and that the representations in each case have different properties. This leads to consideration of the structure of the graphemic representational domain within which the word and letter processing takes place.

### 1. 2. 2. 3 *Graphemic representations: the visual input lexicon?*

This level of description computes from the VIP a representation that specifies which words or word-like structures are currently being foveated. The word-identification domain is taken here to cover skilled readers' complete store of knowledge about visual linguistic symbols. The term "*visual input lexicon*" is sometimes used with this meaning but more often is restricted to the lexical (whole-word) aspects of that knowledge. To include knowledge about sublexical regularities the phrase *graphemic representations* is used here. Many models of the graphemic representational domain have been developed. Some of the most influential of these will be briefly reviewed.

(a) **Logogens.** Perhaps the most influential model of all has been Morton's (1969; 1970; 1979) *logogen model*, which is also the simplest model possible. Each known word is represented by a unit in the visual input lexicon, that fires when sufficient evidence for its presence has accumulated to overcome a threshold firing level. Although logogen units were not intended to be equivalent to single neurons, many have treated them as such, by analogy to neurophysiologists' hypothetical grandmother cells. The problems posed by grandmother cell encoding, such as the inability to treat similar items as similar, are not addressed directly by the model, simply because it makes no detailed claims about the nature of the visual evidence which is collected by a logogen. This leaves two questions which are directly addressed, both of which are relevant: (i) how performance on word recognition tasks is affected by absolute word frequency and by relative frequency, where the latter is reflected in facilitation of performance for repeated words (priming); (ii) how to define the "wholes" involved in word recognition, where the argument has been between word-wholes and

morpheme wholes<sup>17</sup>.

(b) **Multi-level representations.** Many models have been developed to address directly the problem of what information is extracted from the image to activate word-recognition units. The usual solution is to propose a hierarchy of levels, each of which processes combinations of the items explicitly represented at the level immediately below (e.g., Drewnowski & Healy, 1977; Estes, 1975; 1977; Gibson, 1971; LaBerge & Samuels, 1974; Massaro, 1975; Rumelhart, 1977; Smith, 1971). Three levels are commonly identified: the word level, the letter level, and the letter-feature level. The letter-feature level is not properly part of the graphemic representational domain since the features may be common to all visual domains. In effect, then, these models reduce to two levels, a lexical and a sub-lexical level.

Typical sublexical units represent single letters, described, for example, as *letter detectors* (Johnston & McClelland, 1980; McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982), *the letter-form system* (Patterson & Kay, 1982), *letter-based early orthographic processes* (Warrington & Shallice, 1980), *orthographic analysis* (Morton & Patterson, 1980), and *preliminary letter identification* (Rayner, McConkie, & Zola, 1980). Sometimes the sublexical representations include letter groups (graphemes, which map on to single phonemes, and larger groups such as syllables) as well as single letters (Patterson & Morton, 1985).

A rather different solution is to maintain the levels approach but to allow input to the higher levels directly from VIP without going through the lower levels (e.g., Howard, 1987). This is also how Morton (1970) described input to the logogens: if the stimulus is the word cat " ... the output from the visual analysis might include the attributes <three letter word>, <tall letter at the end>, <initial c >, <final t >, and so on " (p. 206).

---

<sup>17</sup> This argument remains unresolved. There is some evidence of morphemic effects for some (root) morphemes in full report (e.g., Murrell & Morton, 1974), less so for lexical decision (Fowler, Napps, & Feldman, 1985; Stanners et al., 1979; Taft, 1979; 1987; Taft & Forster, 1976), and little for other morphemes, such as prefixes (e.g., Smith et al., 1984; Lima, 1987).

(c) **Fuzzy representations.** Allowing heterogeneous sublexical units at the lower level begins to blur the distinction between levels. Two developments of this have been offered. In one, it is proposed that a preliminary orthographic, or letter, level maps into a level containing representations of differently sized items, essentially word-sized but also morpheme, syllable, and bigram sized. In this case, the sizes of the units are determined by their usefulness or regularity for performing a mapping into an output representation. In other words, they are the representations developed by the hidden-units in a three-layer network. These have been used to model the mapping into phonology (e.g., Seidenberg & McClelland, 1989; Sejnowski & Rosenberg, 1987) and into semantics (Hinton & Shallice, 1989).

A more radical approach has been to blur the distinctions between differently-sized units still further by proposing that a single visual word-form system processes VIP information at a variety of different scales (Shallice & McCarthy, 1985; Shallice & Warrington, 1980; Shallice, Warrington, & McCarthy, 1983; Smith & Spoehr, 1974; Spoehr & Smith, 1973; 1975). This approach is also taken by the LW model (Golden, 1986) of letter-in-word perception. LW collapses the feature, letter and word levels into a single visual-feature-in-position "graphemic" level, over which distributed representations of letters or words are sustainable.

How best to interpret these different positions? The strictest approach is to analyse the hypothesized levels of representation in terms of the information they contribute to the tasks they help perform. This is a computational question. To map directly from VIP to the output target representations in one step (using current neurocomputational technology) is possible only if the regularities in the mappings are all first-order statistical regularities. Clearly this is not true of spelling-sound mappings because, for example, the pronunciation of any letter is very often dependent on the identity of neighbouring letters, and sometimes dependent on the identity of neighbouring syllables. To model second-order statistical regularities a three-layer network is needed, using, for example, the back-propagation algorithm to learn the mappings. Whether three-layer networks can learn third- and higher-order regularities with different learning algorithms is unclear as yet. If they can, then a single, multi-scaled, visual word-form system is sufficient.



## 1. 3 The representation of relational information

The purpose of this section is to analyse the different methods of encoding relational information used by various computational models. A discussion of the possible ways of representing relational information is presented in §1. 3. 1. This includes an analysis of the computational requirements of some of the possibilities, and how current connectionist models might deal with the requirements. Part-whole relational representations are discussed in more detail in §1. 3. 2. Part-part relational encoding theories are discussed in §1. 3. 3, though these schemes have not been implemented in as much detail as part-whole representational schemes.

### 1. 3. 1 Computational implications

The distinction between sensory and object representations emphasises that it is easier to map into output representations from object representations than from sensory representations. This is because object representations generalise over variations in viewpoint so that different images of the same objects or words can be recognised as the same irrespective of their retinal position, orientation, and size. Word representations have to generalise over variations because a word remains the same word, irrespective of the case, script, font, colour, and context in which it is presented. Some models obtain generalisation by mapping into object-centred coordinates (e.g., Hinton, 1981a; b; and c)

Models that use viewer-centred coordinates in the sensory domain and object-centred coordinates in the object domain propose a qualitative difference between the domains. Mapping from viewer-centred to object-centred descriptions imposes a severe computational problem. Relaxation networks of two types have been used to provide solutions. The simplest, the generalised Hough transform (Ballard, 1981) succeeds in performing the mapping, but at the expense of reduplicating a network for every perceptible object. The dynamic Hough transform, as implemented in the MAPPER model (Hinton, 1981), performs the mapping, and uses only one network for all objects, but can only successfully

map one object at a time. The problems these models encounter may be due to the way they encode position. Both retinal and object-centred frames impose a set of coordinates upon the possible representations within each domain: position information is represented in absolute terms as  $x, y$  coordinates in relation to a frame, a direct code for position. In effect this type of code represents the position of parts in relation to a whole, where the whole is defined as the retina or retinotopic map in one case, and as an object or scene in the other. The difficulty of the mapping problem is a function of the complexity of the representations assigned to the parts.<sup>18</sup>

There are five plausible alternatives to mapping from part-whole representations into a different domain of part-whole representations: to map (i) part-whole representations into a representation that encodes relationships between parts; (ii) from one domain that uses part-part relational representations into another domain using part-part representations; (iii) from part-part representations into a part-whole representational domain; (iv) from part-part representations into a combination of part-part and part-whole representations; (v) from part-whole representations into a combination of part-part and part-whole representations;

Watt (in press) argues strongly against the notion that the retinotopic maps found in visual cortex are anything like real maps with  $x, y$  coordinates. Instead he suggests that the only positional information directly available from early visual descriptions is in the relations between statistically significant parts of the image. Direct codes, he argues, provide a coordinate description of position but have to be computed from the relational code. MIRAGE (Watt, 1988) uses relational position to generate visual descriptions of the image. Using part-part relations in the sensory domain rules out possibilities (i) and (v) above. The remaining possibilities are (ii), (iii), and (iv), that the sensory domain maps into part-part object/word representations or into part-whole object/word representations, or into some combination of the two. The first two cases will be treated as the part-part and the part-whole hypotheses.

Some models assume that the relational information in the visual input description is

---

<sup>18</sup> The same argument applies to models that map from part-whole viewer-centred sensory representations into part-whole viewer-centred object representations (see e.g., Rock, 1973; 1983).

adequate for the mapping into object representations. The object representations in these models are "frame-independent descriptions" (Corballis, 1988), or "topological categories" (Roberts, 1965; Minsky, 1975, both cited in Hinton, 1981a), and are treated here as part-part relational descriptions. Some part-part relational models assume that once an object representation has been mapped into, an object-centred part-whole description becomes available for further processing of the visual input description (see e.g., Corballis, 1988; also proposed by Barrow, Tenenbaum, Bolles, & Wolf, 1977, cited in Hinton, 1981a)<sup>19</sup>. Other models assume that object-centred information is accessed simultaneously with activation of a viewer-centred object representation (e.g., Kosslyn et al, 1990; Feldman, 1985). In the part-part hypothesis the distinction between sensory and object/word domains is blurred by the use of the same positional information in both domains. The part-whole hypothesis implies a much stronger distinction.

To summarise: object/word perception involves mapping from VIP descriptions to object/graphemic representations. Representations at the object/graphemic representational level are either viewer or object-centred. Object-centred representations always involve part-whole encoding of position. Viewer-centred representations are sometimes based on part-whole encoding of position, but more usually involve part-part encoding of position. Some models assume that both types of positional encoding are simultaneously activated, others that part-part representations are activated first and form the basis for the derivation of part-whole representations.

### 1. 3. 2 Part-whole theories


In this scheme the object/word-centred frame is typically composed of a fixed number of slots, each labelled by specific coordinates, into which parts are slotted. The number of slots, and thus the size of the frame, is fixed but potentially unlimited: any number of slots could be added or subtracted at will. The relationships of the parts to the whole are encoded by their coordinates within the frame. The position information is implicit at this level, but

---

<sup>19</sup> Compare with *verification* models of word recognition (e.g. Becker, 1976; Grossberg & Stone, 1986; Paap, Newsome, McDonald, & Schvaneveldt, 1982).

could be made explicitly available to a further level of representation. Suggestions from the word recognition literature are of ordinal position representations of the form 1 = leftmost, to  $n$  = rightmost (Seymour, 1979), or 1 = first, to  $n$  = last ordinal position (Monk, 1985).

(i) **Interactive Activation.** The best-known example of the ordinal position representational scheme is the Interactive Activation model of word recognition (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982), henceforth referred to as IAM. IAM can recognise 1,179 four letter words; it does so by a two-way cascade process. Processing at one stage does not have to be completed before processing at the next stage begins. This leads to a variety of effects, such as the enhancement of weak signals, error correction, and, in particular, the facilitation of processing at one stage (letters) by processing at a subsequent stage (words). The system is shown in Figure 1.1.



**Figure 1. 1.** Simplified sketch of the Interactive Activation Model. "Units within the same rectangle stand for incompatible alternative hypotheses about an input pattern and all are mutually inhibitory. The bidirectional excitatory connections between levels are indicated for one word and its constituents." (Adapted from "Putting Knowledge in its Place: A Scheme for Programming Parallel Processing Structures on the Fly" by J.L. McClelland, 1985, *Cognitive Science*, 9, p. 115.)

IAM produces a word-centred description which is mapped into from early vision by an unspecified preprocessing system that normalises the input in terms of case, size, orientation, and retinal position. IAM has three levels of units, each at different levels of abstraction: visual features-in-letters, letters-in-position, and words. Each level forms a separate representation, though they are simultaneously active and interact with each other. Activation initially propagates from the feature detectors through letter detectors to the word units. As described by Monk (1985),

There is also potentially inhibition within a level, and top-down activation or inhibition. In their implementation of the model, which recognises four-letter words from a pre-processed input, lateral inhibition is limited to the word and letter units, and top-down influences to the excitation of letter in the active word units. Each of the letter-recognition units has a specific position in the word, thus there are detectors for "A" as the initial letter, "A" as the second letter and so on... Similarly, the feature-detecting units are specific to some position in a letter at some position in the string... The input ... is simply a vector indicating which of the 64 features are present, there being 16 possible features in each of the four-letter positions (p.622).

Each unit has a momentary activation value, a resting value (which for the word units depends on that unit's frequency, i.e., it is higher for high-frequency words), maximum and minimum values, and a decay rate. All excitatory and inhibitory influences on one unit summate algebraically to give that unit's activation level. When IAM is presented with an input stimulus, it works as follows. A number of feature-in-position units become activated. Activated feature-in-position units excite all letter units containing that particular feature in that particular position, and inhibit all letter units for that position which do not contain that feature. When these letter-in-position units become active they inhibit units representing different letters in the same position, excite the word units that contain that letter in that position, and inhibit all other word units. The word units start to compete with each other, and feed back excitation to the supportive letter-in-position units in the level below. Over a number of relaxation cycles, the system settles down into a state in which one word unit, four letter-in-position units, and a collection of feature-in-position units are active. This mutually supportive assembly of active units constitutes the word-centred description of the input stimulus. It is word-centred because units at the levels lower than the word are

encoded relative to their position in the word.

The central phenomenon which IAM was designed to simulate is the Word Superiority Effect, the finding that 2AFC is more accurate for letters in words and pseudowords than for letters on their own or in nonwords (e.g., Adams, 1979; Aderman & Smith, 1971; Reicher, 1969; Wheeler, 1970). As was pointed out in §1. 2, this task can be seen as requiring the processing of letters as wholes within word-wholes; this Word Superiority Effect<sup>20</sup> is evidence that whole-within-whole processing is facilitated by the presence of a familiar word-whole. This point will be important in Chapter 5.

As mentioned earlier this coding scheme works with pre-processed input<sup>21</sup>; for feature-in-position units to be activated by the same input, the image must be normalised for variation in size, orientation, and, crucially, labelled with ordinal position within the string. Johnston & McClelland (1980) are particularly explicit about the latter assumption: a letter position pre-processor encodes the stimulus as a sequence of unanalysed visual blobs, labelled with position in the sequence, and then passes this description to the first level for feature analysis. Labelling unanalysed blobs, however, cannot be the whole story because the input still requires normalisation. How this is done is unspecified, but Rumelhart & McClelland (1982) suggest that the canonical normalisation performed by MAPPER (Hinton, 1981b; see §1. 3. 2. 3) could perform the required computations, perhaps with a degree of uncertainty so that input was slightly "smeared" across the feature detectors. While MAPPER does normalise the input it is not at all clear how position labelling is done: presumably it requires a description of the whole blob, so that the labelling process knows where to start and stop, and subsequent segmentation of this blob into its parts. Processing from coarse to fine in this way is reminiscent of MIRAGE, which raises the possibilities (a) that this level of processing is capable of performing complex operations on the input that might reduce the computational burden on the graphemic network, and (b) that the coarse blob information could be input to the graphemic network together with, or even before, the

---

<sup>20</sup> This effect should be distinguished sharply from word superiority on word-whole tasks which involve processing letters as parts of the whole; this distinction, though, is not usually made.

<sup>21</sup> In this respect all three levels of representation can be considered as part of the graphemic representational domain (see §1. 2. 2. 3).

fine blob information.

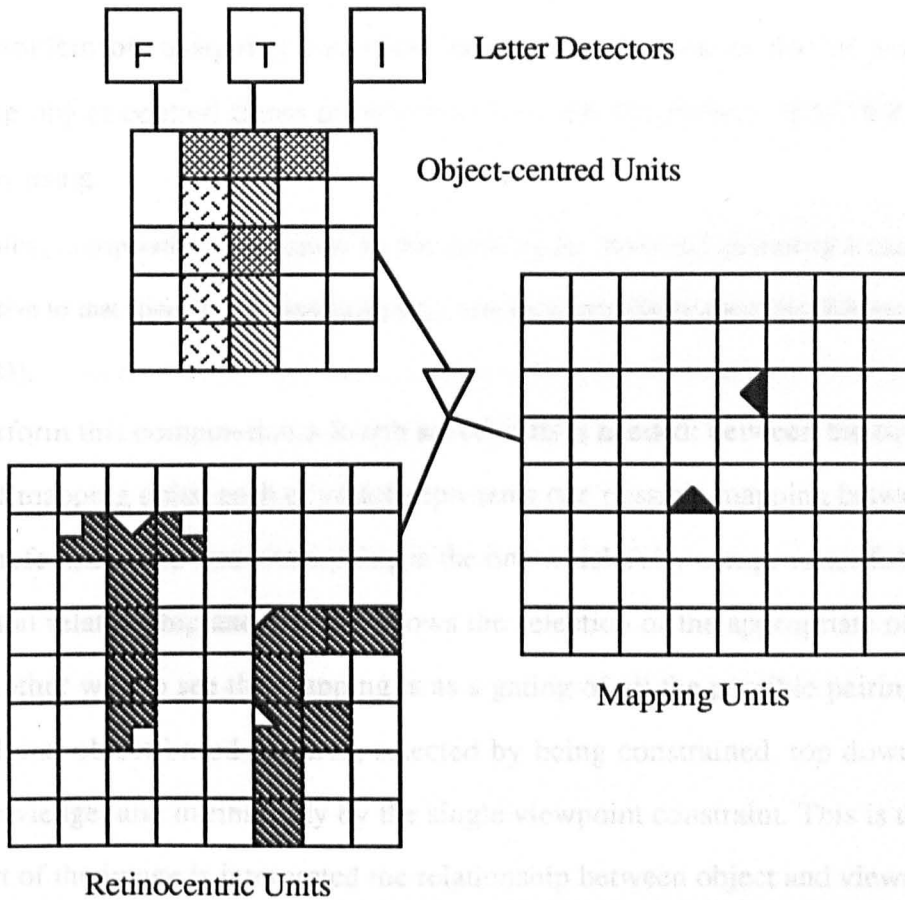
(ii) *Other word perception models.* Many other models of word recognition also use this coding scheme, including models that map from spelling to sound with connectionist (Lacouture, 1989; Van Orden, 1987), and non-connectionist algorithms (Brown, 1987), though the latter also maps bigrams and trigrams-in-position. Other models that use this coding scheme are a serial simulation of letter-in-word perception, the Elementary Perceiver and Memorizer (EPAM, Richman & Simon, 1989)<sup>22</sup>, and the connectionist LW model of letter-in-word perception (Golden, 1986).

A more complex example of part-whole encoding is the development of IAM as PABLO (McClelland, 1986). Representation is of part-whole relations combined with some part-part relational information in the VIP. Each letter is represented as a set of four 2-letter combinations: before or after a space, and before or after any other letter. Position is coarsely coded by the overlapping units activated by the letters within a letter string, with explicit information about which letters are at the beginning and end of the string. PABLO, however, replicates the set of coarse-coded letter-cluster units for each position within a word, which reduces the possibilities of within-word interactions.

(iii) *MAPPER.* An example of this coding scheme applied to both object and word perception is Hinton's connectionist model of canonical recognition, and its computer simulation, MAPPER (Hinton, 1981a; b; and c; Hinton & Lang, 1985). MAPPER can recognise prespecified simple patterns ("objects" or letters) irrespective of their "retinal" input location. It does so by computing an object-centred, or canonical, description of the input. The architecture of MAPPER is shown in Figure 1.2, and is described below.

---

<sup>22</sup> EPAM differs in that it first processes strings at the word-level, and only moves to sublexical levels when that fails. This idea is taken up again in Chapter 5.



**Figure 1. 2.** Simplified sketch of the MAPPER model, applied to letter perception (adapted from Hinton & Lang, 1985, p. 253).

The system is composed of four sets of units. Each top level unit represents one of the prespecified patterns. These are shape recognition units, equivalent to “pictogens” (Warren and Morton, 1982), or letter-detectors. Below this level are two separate arrays of units representing object-based and retinally-based coordinate frames. Units in these arrays represent particular relations of the pattern's parts to that particular frame. Units in the retinally-based frame are called *retinocentric units* because they encode particular feature types (line segments and junctions) in particular positions on the retina. Their activity thus depends on the spatial relationships between objects and retina in a direct manner. In terms of hierarchical models of visual perception they represent the highest level of processing of features at which retinotopic information is still encoded, roughly equivalent to Marr's 2 1/2 D sketch. Units in the object-centred frame represent the spatial structure of the pattern's parts relative to a frame which is intrinsic to the pattern itself. Such descriptions are canonical because they are not altered by changes in the retinal description of the pattern.



The problem of assigning canonical descriptions is seen as that of assigning the appropriate object-centred frame on which to base the description. MAPPER solves the problem by using

parallel,...cooperative computation so that choosing the frame and generating a description relative to that frame occur simultaneously, with each influencing the other (Hinton, 1981b, p.683).

To perform this computation a fourth set of units is needed: between the two frames is an array of mapping units, each of which represents one possible mapping between the two frames of reference. The correct mapping is the one which fully compensates for the object-retina spatial relationship and thereby allows the selection of the appropriate object-based frame. Another way to see the mapping is as a gating of all the possible pairings between the retinal and object-based features, selected by being constrained, top down by stored object knowledge, and intrinsically by the single viewpoint constraint. This is that as soon as one part of the image is interpreted the relationship between object and viewer becomes highly constrained because every retinal point is formed from exactly the same viewpoint.

MAPPER successfully performs position-independent recognition of any of the objects that it knows; it can also do orientation-independent letter recognition. Its limitations include the inability to recognise new objects, to perform any size generalisation, and to recognise more than one object at a time. Theoretical expansions of the system attempt to account for other perceptual phenomena<sup>23</sup>. When human subjects are presented with brief displays of three differently coloured letters and two black digits, and asked to report what they have seen, they occasionally report one of the presented letters, but in the colour of another letter, an "*illusory conjunction*" (Treisman and Schmidt, 1982). MAPPER, modified for letter recognition, makes similar errors when pattern masked (Hinton & Lang, 1985). The mask

---

<sup>23</sup> One modification of MAPPER adds a set of *scene-based units*, receiving input from the object-based units, gated by another set of mapping units. These scene-based units perform the work of a spatial working memory for scenes. Such a working memory is assumed (Hinton, 1981b) to hold recently presented Gestalts, and thus to mediate anorthoscopic perception (in which an object is correctly perceived despite being presented one piece at a time through a keyhole, e.g., Hochberg, 1968) and context effects in perception (e.g., Palmer, 1975a).

removes the retinotopic position information, allowing conjunctions to be made between the dominant letter identity in the object-centred units and the dominant viewpoint in the mapping units, conjunctions which are sometimes incorrect.

*(iv) Extensions to more complex object representations.* A more complex object-centred representational scheme explicitly encodes the disposition of parts relative to wholes, in terms of their three-dimensional distance from any of the major axes (such as elongation, or symmetry) intrinsic to the object whole. This is done hierarchically, by decomposing each object into a succession of parts, each of which is treated as a whole for the next decomposition into parts. As each part becomes a whole it is assigned an axis, which is taken as the basis for the coordinates at that level. Any "whole" therefore will incorporate not only a description of its parts relative to itself, but also a description of how its major axis relates to the larger whole of which it is, in turn, a part.<sup>24</sup> A clearly worked out example is the model of object recognition developed by Marr (e.g., Marr and Nishihara, 1978; Marr, 1982).

In this model, object recognition consists of the decomposition ("segmentation") of complex objects into a hierarchy of parts each with their own central axis and associated cylindrical coordinate system, and the matching of this hierarchical description to a "catalogue" of descriptions stored in memory. Information about both the spatial arrangement and lengths of the axes is assumed to be involved in discriminating between similar objects. Since the coordinates are based on the object and not on the retina or the environment, this process produces an object-centred description. The model works most naturally for objects composed of elongated parts, each describable as a generalised (or variable diameter) cone.

This scheme differs from the previous one in that it produces a hierarchical structural description. Larger parts cannot be formed out of smaller parts without first checking the spatial relationships between the smaller parts. In the previous scheme, smaller parts at the

---

<sup>24</sup> The idea of having many local (not necessarily hierarchical) origins to describe spatial positions can be found in Atneave (1954), and Palmer (1975b).

object based level directly contribute to the formation of larger parts because the relational information is intrinsic to their description relative to the object-based frame (Hinton, 1981a).

### 1. 3. 3 Part-part theories

(i) *Neighbourhood Relations*. The simplest version of part-part encoding is a system which describes the position of each part in terms of the parts to which it is a neighbour, in other words based on the relationship "next to". This is the natural method of description for parts spaced out in one dimension only, such as letter-strings, but becomes more difficult to use in more than one dimension. Theories of order information in short-term memory distinguish between inter-item relations that are non-directional (e.g., Murdock, 1983), and those that specify directions between pairs (e.g., Pike, 1984).

One simple example of the second kind of arrangement is found in Biederman's theory of object recognition (e.g., 1987). Instead of simply using generalised cones as the basic parts, Biederman proposes a wider range of basic shapes which he calls *geons* (geometric ions).<sup>25</sup> This includes simple shapes such as cones, wedges, spheres, and deformations of the basic shapes that do not introduce concave (i.e. pointing into the object) discontinuities. The importance of concave discontinuities is that they are powerful cues for segmenting an object, or its occluding contour, into different parts (e.g., Hoffman & Richards, 1984). Biederman proposes that geons are identified from the image by various "non accidental" or "landmark" properties of the edges in the image, and the nature and arrangement of the geons is then used to match structural models of objects in memory. The important point here is that it is very simple spatial arrangements of geons that are computed, such as relative sizes, orientation, place of attachment, all of which are derived from the relationships between neighbouring geons.

Biederman proposes that geons are the object equivalents of phonemes and letters, the

---

<sup>25</sup> A similar proposal is found in Pentland (1986), where the shapes are referred to as superquadric components.

spatial or temporal arrangements of which go to make up words. It has been suggested that the basic elements of words are not letters, because the number of possible relationships among them is so large, but groups of three letters. These triplets, proposed by Wickelgren (1969), and known as Wickelgraphs, are a form of context-sensitive encoding: the representation of any letter depends on the letters that immediately precede and follow it. The representation of "aTe" is different from that of "eTs". Thus simple neighbourhood relations are built into this encoding scheme. The most interesting features of this scheme are as follows: (a) "the unordered set of codes is sufficient to reconstruct the ordered components of the word" (Mozer, 1987), which means that each set of codes activated by a word is enough to identify it uniquely; (b) within limits, this scheme allows the simultaneous representation of two words without interference; and (c) the number of letter-clusters needed to account for all occurring clusters is not large, and the 1,000 most common account for 50% of all occurring letter-clusters (Mozer, 1987). Strings of Wickelgraphs, however, are sometimes dealt with in the standard part-whole manner, that is as a string coded in terms of ordinal positions (e.g., the PABLO model discussed earlier).

Other uses of Wickelgren-type encoding can be found in Cohen and Grossberg's (1986) model of word recognition, and in Mozer's (1987) model, **BLIRNET**, which "Builds Location-Independent Representations of multiple words". **BLIRNET** uses low-level visual features at its bottom level. Four subsequent layers recode the information in the preceding layer in terms of more complex combinations ("relative spatial arrangements") of lower-level features, each of which generalises over increasingly larger regions of retinal space. This solution to the position-invariance problem is based on that offered by the Neocognitron (Fukushima & Miyake, 1982). The fifth layer encodes 750 feature-types, which learn to map on to a letter level containing 6000 letter-cluster feature types. These clusters are of three letters (including a space); the three letters do not have to be directly adjacent: they generalise over a single intervening letter in any position. Essentially trigrams are represented, except that the trigram can contain an extra unidentified letter among its identified letters. **BLIRNET** uses letter-cluster information as the input to a Pull-out net that uses competitive interactions between candidate words to reduce noise and select a single

word response. Words are represented in the net as excitatory connections between consistent letter-cluster units, not as single units.

A more complex version is found in the Seidenberg and McClelland (1989) three-layer back-propagation model of the mapping into phonology. Each input unit represents a triplet of 10 possible first letters, 10 possible middle letters, and 10 last letters; each orthographic input triple activates a set of units, usually numbering about 20. The connections between these units and the hidden units that perform the mapping into phonology are not pre-programmed. The back-propagation algorithm develops its optimum interpretation of the input given the output requirements. In practice this means that a set of orthographic input units sometimes activate a unit that represents a whole word, but more usually activate sets of units, the combination of which constitutes a whole word. The distinction between parts and wholes is blurred because the hidden units organise their own representations.

(ii) *Hierarchical Neighbourhood Relations.* A complex elaboration of the previous scheme, but hierarchical by its very nature, is a scheme whereby parts are again related to their neighbours, but parts can be of any degree of complexity or size. This produces a highly redundant, rich description of the input. Each part can be described in relation to any of a hierarchy of differently-sized, neighbouring parts. Each part can also contribute to many differently sized groups of parts, described in relation to other parts on their own and to other groups of all sizes. Clearly the problem of assigning a definition to the term "part" becomes particularly acute for this scheme. Potentially the smallest discriminable feature could be used as a "part", but the number of relations a part of this size enters into could become intractable. Such a scheme does not seem to have been concretely specified in the literature on letter and word perception: it has been put forward as a way of modelling the acquisition of new concepts (Hayes-Roth & Hayes-Roth, 1977), in terms of a concept's properties being represented to each other in all possible combinations.

This scheme emerges directly from the model of object recognition proposed by von der Malsburg and co-workers as an example of how a modified version of connectionism, dynamical connectionism might work (e.g., Bienenstock & von der Malsburg, 1987; von

der Malsburg, 1981; 1985; von der Malsburg & Bienenstock, 1986).<sup>26</sup> Because dynamic connections embody relationships, a network of active connections can be seen as a labelled graph, or structural description. Pattern recognition can then be treated as a problem of graph matching: i.e. matching the input labelled graph with one stored in memory:

The labels are local features extracted from the image, such as grey-level and color, edge elements, descriptors of texture etc. The links of the graph embody the neighbourhood relationships between these local features in the image. The matching algorithm thus operates not on the raw data ... but on a relational graph derived from the data. Such a relational description is designed to be intrinsically invariant (Bienenstock & von der Malsburg, 1987; p.122).

The importance of neighbourhood relations is clear: "The refined plasticity of correlation theory is ... analogous to measurement of the probabilities of letters to be adjacent" (von der Malsburg, 1981; p.33). It is also clear that von der Malsburg intends that this process operate at a variety of scales: "A way to take advantage of the topology" (loosely, the relational graph structure) "is a simple 'divide-and-conquer' strategy... : the map is first roughly outlined, and then progressively refined" (Bienenstock & von der Malsburg, 1987; p.125). Von der Malsburg (1981) presents the reasons for this strategy:

In our cultural world we form symbols of a higher order by the juxtaposition (in time or space) of symbols of a lower order, e.g. words out of letters or phonemes. According to the localization theory neurons are the basic symbols in the brain. Their position is fixed and

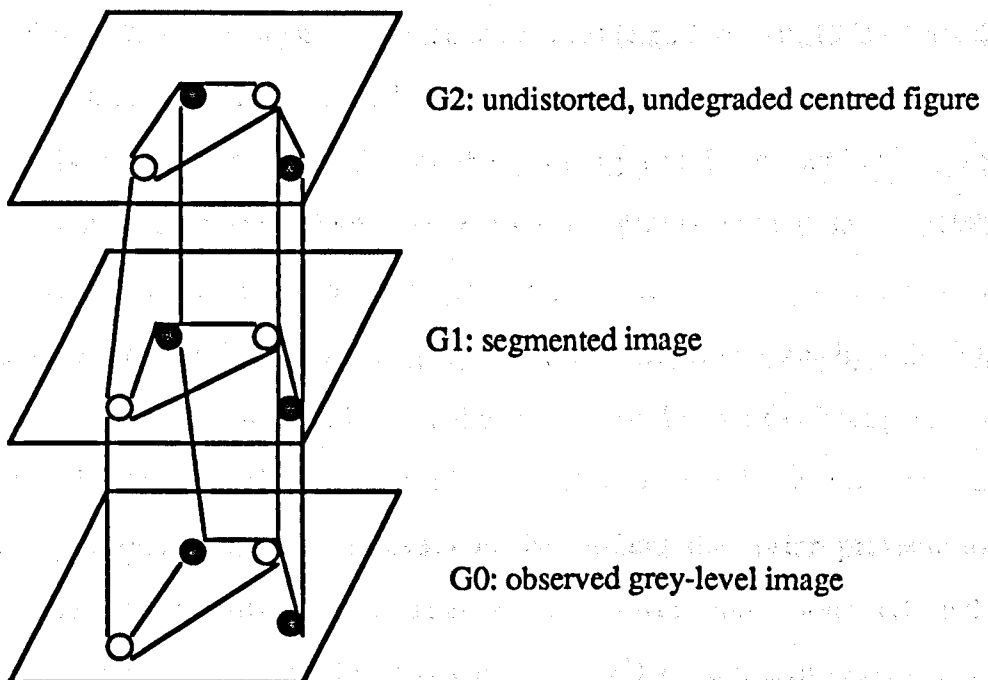
---

<sup>26</sup> The new elements of dynamical connectionism are as follows. In conventional connectionist models, the weights between units serve only as the underlying structure of LTM. Dynamical connectionism proposes that they play a more active role by switching on and off on the time scale of processing (similar ideas in Feldman, 1982; Hinton & Plaut, 1987). Connections that are "on" operate at their maximum strength; connections that are fully "off" operate at minimum strength; between these two extremes is a resting value to which the connection strength returns as it becomes inactive. Maximum strength is attained when the neurons on either side of the connection (the pre- and post-synaptic neurons) are simultaneously active. Connections detect coincidences, or correlations, between the firing patterns of the neurons they relate. Feedback between connection strength and the rates and synchronicity of firing patterns creates assemblies of mutually supportive connections, disconnected or decorrelated from all other possible assemblies on the same set of neurons.

cannot be used to form groupings. Another code is required to represent association of cells into patterns forming symbols on a higher level.

When we analyse a complex visual scene it is important to break it down into patterns which are simple enough so that we can hope to recognise them, e.g. identify them with objects we saw before. A single pattern in turn has to be broken down into subpatterns, possibly through several stages, e.g. man-arm-hand-finger-joint ... It should be possible to group neurons into such a hierarchy of patterns in a flexible way, without the introduction of new hardware for new patterns (p. 8)

Figure 1. 3 shows how this process is hypothesized to work for object recognition. Cells in G1 encode their neighbourhood relations by signal correlations. These are propagated to G2 "where they activate a connection pattern which encodes the topology of G1" (von der Malsburg, 1985). Cells in G2 are position-invariant and feature-specific: the activated connections between them signal the relationships between the features in the image.



**Figure 1. 3.** Simplified sketch of dynamical connectionist network for pattern recognition (adapted from Bienenstock, 1987, Relational Models in Natural and Artificial Vision, British Psychological Society Annual Conference.)

## 1.4 Psychological and neuropsychological evidence

The central paradigm used to investigate the encoding of relational information is priming. The logic of the phenomenon is that performance of a task on a letter-string stimulus is sometimes affected by prior performance of the same task or another task on another letter-string. To the extent that performance on a letter-string is unaffected by a prior performance of another letter-string the representations used in the performance of the task are assumed to be independent of one another. The extent to which performance is changed by prior performance is taken as a measure of the similarity of the representations used for one letter-string to the representations used for the other letter-string. Some independent measures of the similarity of the two letter-strings to each other are taken and correlated with the effects on performance.

It is well established that performance of many tasks on a particular letter-string is affected by prior performance of the same task, or a different task, on the same letter-string. This is called *repetition priming* and is well documented (e.g., Monsell, 1985). Identifying which representations are implicated by repetition priming is more controversial. In particular it is sometimes claimed that the priming has its effect on "episodic" representations of the context in which a letter-string was processed (e.g., Jacoby, 1983). If this is the case then inferences about graphemic representations would have to be more tentative (but see Monsell, 1985, for strong arguments against the episodic hypothesis).

Prior presentation of semantically related words is also known to affect performance (e.g., Forster, 1981), but this effect is still less clearly useful for inferences about graphemic representations. More relevant is the finding that prior presentation of orthographically similar words also effects performance. This phenomenon is referred to as *orthographic priming* (Humphreys, Evett, & Quinlan, 1990), or *similarity priming* (Rueckl, 1990). It is also clearly established that all forms of priming can involve both facilitation and inhibition of performance, depending on the task requirements (for examples of repetition producing interference see e.g., Kanwisher & Potter, 1990; Humphreys, Besner, & Quinlan, 1988). In the next sections the relevant similarity priming literature is reviewed.



This forms part of a more general attempt to find evidence for part-whole and part-part relational encoding.

### 1. 4. 1 Similarity priming

Similarity priming has been reported for the naming task (Feustel, Shiffrin, & Salasoo, 1983; Masson & Freedman, 1990), and for deciding whether the string is a word or not (lexical decision: Balota & Chumbley, 1984; Besner & McCann, 1987; Besner, Dennis, & Davelaar, 1985; Fowler, Napps, & Feldman, 1985; Monsell, 1985). As with repetition priming (Humphreys et al., 1988), similarity priming can be inhibitory when no mask is presented between prime and target (Colombo, 1986). Nonword and word pronunciations can be biased and delayed by prior naming of similar regular or irregular words (Kay & Marcel, 1981; Taraban & McClelland, 1987). Similarity priming is also sometimes found for full report (Murrell & Morton, 1974; Rueckl, 1990) but not always (Humphreys, Evett, Quinlan, & Besner, 1987). Pseudoword similarity priming is found for lexical decision (Fowler, Napps, & Feldman, 1985; Colombo, 1986) and full report (Whittlesea, 1987; Rueckl, 1990).

Similarity priming can also occur when the prime is not responded to, and not consciously perceived (subthreshold priming). This is found for naming (Manso de Zuniga, Quinlan, & Humphreys, 1988), naming and fixation duration with parafoveal preview (Rayner, McConkie, & Zola, 1980; Rayner, Well, Pollatsek, & Bertera, 1982), and full report (Evett & Humphreys, 1981; Humphreys, Besner, & Quinlan, 1988; Humphreys, Evett, & Quinlan, 1990) but not for lexical decision (Forster, 1987; Forster & Davis, 1984; Manso de Zuniga et al., 1988). This latter negative result emphasises that each task requires analysis in terms of its processing demands before inferences about particular representations can be made safely. In particular, lexical decision specifically requires descriptions of the whole string for its successful performance. Whole descriptions may be primed by similar descriptions only when the similarity becomes clear. This is so for above-threshold similar items which are described as wholes, but not for subthreshold items which cannot be clearly described as related wholes. There is also good evidence from similarity

priming that lexical decisions can be performed on phonological as well as graphemic representations (Seidenberg & McClelland, 1989, present this argument in detail). Simultaneous presentation of orthographically similar non-rhymes ("couch-touch") inhibits lexical decisions, whereas rhyme pairs produce facilitation (Meyer, Schvaneveldt, & Ruddy, 1974). Since facilitation is also obtained with differently-spelled rhymes such as "cake-break" (Hillinger, 1980), phonological representations are implicated. One mapping into phonology is pulling the mapping of the other word into the same phonological representation, slowing the development of stable representations. When phonological recoding for lexical decisions is made unnecessary by using random letter-strings instead of pseudowords as the nonwords, facilitation of both rhyming and non-rhyming orthographically similar words is obtained (Shulman, Hornak, & Sanders, 1978).

In general similarity priming has been used to investigate issues other than part-whole relations such as the effect of morphemic structure on priming, for example. These studies rarely defined prime-target relationships except in terms of the absolute number of letters shared by prime and target. Recent work by Humphreys and colleagues has begun to specify more precisely the nature of the relevant similarities that produce priming. This has established that priming increases non-linearly with the number of shared letters in the prime, and is larger when the first and last letters are shared (e.g., Humphreys, Evett, & Quinlan, 1990). Both observations are consistent with the idea that relationships between letters, and between letters and end-of-string spaces are important. More directly, Experiment 6 of Humphreys, Evett, and Quinlan (1990) compared priming between conditions in which letter-relationships were maintained (e.g., "bvK"- "BLACK") and those in which absolute positions in the string, part-whole relationships, were *also* maintained (e.g., "btvuk"- "BLACK"). No additional priming obtained when positions in the string were also maintained.

Taken together, these results provide strong evidence that the representations that mediate this similarity priming make use of part-part relationships. The generality of this result is severely limited by the following experimental factors: the prime letter-string was presented below threshold, was masked, and was a nonsense letter-string, whereas the target was always a word, also presented below threshold. Additionally the range of

relationships manipulated is restricted to end-letters, though their Experiment 5 also found significant priming when two internal letters had their relationship maintained across prime and target. Whether their results apply to other situations is clearly an important question.

Duncan (1987, Experiment 6) found much weaker effects in a different paradigm. When target words (e.g., STAB) are flanked by two unattended (peripheral) primes, 2AFC accuracy on the target is not improved relative to a neutral control either by (a) having the target wholly present as one of the two primes, (b) having the target present in the primes, first half in one prime, second half in the other (e.g., STUX and ICAB), or (c) having the target present in a prime but with its part-whole relations disrupted (e.g., ABST). Although none of these prime conditions were better than controls, all of them were better than when the same three conditions primed the alternative word offered for the choice discrimination. In summary, priming the letters of the alternative is equally harmful, whether the letters are all in the same word and in correct positions, all in the same word and in incorrect positions, or in correct positions across two words. Priming under these conditions appears not to be position-specific, though the data show a tendency for the moved-position priming of the alternative to be less harmful than the fixed-position cross-string priming.

One other situation in which relative and absolute position information has been compared is that in which successive letter-strings are presented for a same-different judgement. The information used for same-different matching is not retinal position, because changing the spacing between letters at test does not effect performance (Bjork & Murray, 1977; Estes, 1982; Ratcliff, 1987), nor is the information completely precise: there is a gradient of positional uncertainty so that when adjacent items in the study string are interchanged in the test string, subjects find it difficult to respond that they are different, but this becomes easier the greater the distance between interchanged letters (Angolillio-Bent & Rips, 1982; Proctor & Healy, 1985; Ratcliff, 1981; 1987; Ratcliff & Hacker, 1985). Comparing conditions in which more relative positions are maintained than ordinal positions and vice versa, should allow their relative contributions to be assessed. Ratcliff (1987) interprets his data for this comparison as evidence for the dominance of ordinal positions, but this interpretation is questionable, and no statistics were presented to support it. When

the data are re-analysed<sup>27</sup> in terms of the number of neighbourhood relations (and the relationship of a letter to a preceding or following space is included), the larger the number of relations maintained the harder it is to detect the difference, though responses are fastest when the most and least relations are maintained and slowest in the middle. This could reflect confident false and correct judgments, with a less confident stage between; speed-accuracy trade-offs were not presented for this data. For increasing number of ordinal positions maintained there is no clear effect on accuracy but responses generally become slower.

#### 1. 4. 2 Neighbourhood effects

An effect that is similar to similarity priming can also be obtained, but in this case the similar letter-strings do not need to be presented as stimuli. Performance of tasks on particular letter-strings can be influenced by the total number of words that are similar to the letter-string. Similarity is defined here in terms of number of common letters. The most used measure is that of neighbourhood size, or Coltheart's *N* factor, the number of words derivable from a given word by changing one letter (Coltheart, Davelaar, Jonasson, & Besner, 1977). For this reason this effect is usually referred to as the neighbourhood effect; it is treated as further evidence for interactions between similar representations. As examples of the neighbourhood effect it has been reported that words with larger *N* are named faster (Grainger, 1988; Gunther & Greese, 1985, both cited in Segui & Grainger, 1990), and in lexical decision nonword rejection is slower for nonwords with larger numbers of neighbours (Coltheart, Davelaar, Jonasson, & Besner, 1977; Gunther & Greese, 1985; Scheerer, 1987).

The simple neighbourhood effect, however, interacts with word frequency in a complex manner. In general neighbourhood effects are limited to, or larger for, low-frequency words and nonwords. This has been most consistently demonstrated for the naming task where it applies (i) to naming latencies (Andrews, 1989), (ii) to the priming of words in low-

---

<sup>27</sup> This re-analysis is my own.

frequency but not high-frequency phrases (Osgood & Hoosain, 1974) and the priming of word-bodies selectively in low-frequency words (Bowey, 1990), (iii) to the more widespread biasing of pronunciation by primes for nonword than word targets (Rosson, 1983; Taraban & McClelland, 1987), and (iv) to the effect of the regularity of a word-body's pronunciation on the naming latency of a word containing that body (Rosson, 1985; Parkin, 1982; Parkin & Underwood, 1983; Seidenberg et al, 1984; Waters & Seidenberg, 1985). Words with high-frequency neighbours produce slower lexical decisions than those without (Davelaar, Coltheart, Besner, & Jonasson, 1978; Grainger, O'Regan, Jacobs, & Segui, 1989) and longer full report thresholds (Havens & Foote, 1963). Mis-spelled high-frequency words are more slowly rejected in lexical decisions than mis-spelled low-frequency words (O'Connor & Forster, 1981).

Similarity priming shows strong effects of neighbourhood frequency. For lexical decision similarity priming, low-frequency primes inhibit high-frequency orthographically related target neighbours, while high-frequency primes have no effect on low-frequency target neighbours (Segui & Grainger, 1990; similar results, and also with pseudowords as the low-frequency words, in Colombo, 1986). It was said earlier that lexical decisions generally show no subthreshold similarity priming. However, high-frequency subthreshold primes inhibit lexical decisions to lower frequency neighbours (Segui & Grainger, 1990; though no consistent effect in Forster et al, 1987), whereas low-frequency primes have no effect on higher frequency neighbours; presumably this reflects that processing has gone further for the high-frequency prime, but not far enough for it to become unambiguous. Similarly, full report subthreshold similarity priming is larger with word than nonword primes (Humphreys, Evett, & Taylor, 1982). Facilitation is obtained by related nonword primes for short word targets with very few neighbours and for eight-letter but not four-letter word targets (Forster, Davis, Schoknecht, & Carter, 1987), because number of neighbours decreases with increasing word length. These last two findings suggest that low-frequency (nonword) primes facilitate low-frequency targets.

Taken together the neighbourhood frequency effects strongly suggest important differences between the representations that mediate these priming effects for high-frequency words and those that mediate priming for low-frequency words. Assuming that

the same *type* of representation is involved for both frequencies of word, the results suggest that similar low-frequency words are more similarly represented than are equally similar high-frequency words. In other words, low-frequency word representations overlap more in graphemic representational space. One possibility is that low-frequency word representations involve more overlapping part-part relationships than high-frequency word representations. Possibly, then, a difference between nonwords and high-frequency words might be expected to emerge in terms of the encoding of part-part and part-whole relationships. It seems reasonable to assume that the "whole" of a high-frequency word is better, or more clearly represented than the whole of a low-frequency word, and thus that higher-frequency word representations make more use of part-whole relationships.

### 1. 4. 3 Positional information

The most straightforward approach to the investigation of the encoding of position information is to analyse errors in the report of singly presented pattern-masked letter-strings. This paradigm shows that for nonwords position information decays independently (Townsend, 1973) and faster than identity information (Long, 1980), and that mislocation errors account for the majority of all errors, especially under masked conditions (Mewhort & Campbell, 1978; Mewhort, Campbell, Marchetti, & Campbell, 1981). How to interpret this basic finding is not so clear. One possibility which is often suggested is to take this as evidence for the distinction between VIP and graphemic representations and to argue that position information is available from the VIP but not from the graphemic representations of nonwords. At the very least, this paradigm provides no evidence for the use of accurate absolute position information in nonwords.

Moreover, accuracy in various paradigms is not equivalent across all letter positions. It is well established that accuracy is higher for the end letters of pattern-masked letter-strings, is usually higher for the central, fixated letter, and is higher for letters to the left of fixation for full report, and to the right of fixation for partial report (Bryden, 1966; Crovitz & Schiffman, 1965; Hershenson, 1969; Mason, 1975; Merikle, 1974; Merikle, Coltheart, & Lowe, 1971; Merikle & Coltheart, 1972; Merikle & Glick, 1976; Estes, Allmeyer, & Reder,

1976; Shaw, 1969; Taylor & Brown, 1972; Townsend, Taylor, & Brown, 1971; Winnick & Bruder, 1968; Wolford & Hollingsworth, 1974). Whether these effects are all attributable to the nature of the VIP is not fully established, but since increasing the visual similarity of the display increases the number of errors overall (Morrison & Butler, 1986), explanations in terms of VIP are plausible.

(a) The superiority for right-side letters (Bryden, Mondor, Loken, Ingleton, & Bergstrom, 1990) may reflect greater ease of constructing representations in the left-hemisphere from left rather than right-hemisphere VIP.

(b) The central letter advantage is selectively reduced by post-stimulus pattern-masks, and because it is a function of fixation position rather than within-string position it is attributable to foveal-peripheral acuity differences.

(c) The end-letters advantage has been attributed to the differential effects of lateral masking from surrounding letters, and is also found in the output from MIRAGE (Watt, Bock, Thimbleby, & Wilkins, in press). Since the end-letters advantage is selectively reduced by pre-stimulus masks (which, it was argued in §1.2.2.1, are selectively harmful to VIP representations), and also by more spatially extensive post-stimulus masks (Mewhort & Campbell, 1978), the lateral masking explanation seems plausible<sup>28</sup>.

When tasks demand word processing at the letter-in-word level, similar position effects to those found in nonwords are usually reported (e.g., LaBerge, 1983; Mason, 1975; Rumelhart & McClelland, 1982), but not when tasks demand whole-word processing (LaBerge, 1983).

Mislocation errors are clearly a useful tool for investigating position information. They can be thought of as "intra-string migration errors" to bring out their similarity to cross-string migration errors. Mislocation errors can be reduced in two ways. Firstly, attending to letters selectively reduces their number (Morrison & Butler, 1986) presumably because of the role of spatial attention in maintaining positional information. Similarly,

---

<sup>28</sup> One anomaly about the VIP explanation for the end-letters advantage is that it appears to be specific to strings of letters or digits because rows of geometric shapes show an advantage only for the central stimuli (Hammond & Green, 1982). This discrepancy is not clearly accounted for but may depend on differences in internal masking.

advance attention to the critical letter-position reduces or abolishes the normal 2AFC word-nonword superiority, both with extra spacing between the letters to facilitate their separate perception (Johnston & McClelland, 1974), and with normal spacing (Johnston, 1981), by impairing words and improving nonwords and single letters (e.g., McClelland & Johnston, 1977). Secondly, words protect against mislocation errors. Errors are less common overall, and individual errors tend to be smaller in terms of number of positions moved (Mewhort & Campbell, 1978). The reverse of this effect is that word targets act as attractors to similar nonwords such that nonwords that are anagrams of words (BCAK) tend to be reported as those words (BACK; Johnston, Hale, & Van Santen, 1983), though Duncan (1987) found no effect when *all* positions are disrupted (ABST-STAB). Even in nonwords order-inverted reports of letter-pairs correlate with bigram frequency (Estes et al., 1976).

The distinction between words and nonwords can be taken as further evidence for the distinction developed in the previous section between the encoding of familiar items, which may be able to call on part-whole relations, and the encoding of unfamiliar combinations of letters, which is only able to use relations between letters. The second distinction suggested by this section is between tasks which can make use of whole-word level descriptions and tasks which demand letter-level descriptions. The evidence reviewed in this paragraph suggests that word representations are under attentional control in that they can be processed in response to task demands at either level whereas nonwords can only be processed at the letter-level, a level at which attention to particular spatial positions can boost performance. These two manifestations of attentional flexibility, for words between a coarse and a finer representation, and for nonwords the boost supplied by attention to single letter positions, suggests that graphemic representations may be usefully considered to differ in terms of spatial scale (coarse to fine), by analogy with the scale-space processing of VIP by MIRAGE.

It would be nice if these distinctions were supported by neuropsychological evidence. The two disorders most relevant to the issue of the encoding of parts are neglect dyslexia and attentional dyslexia. Typical neglect dyslexics neglect the left side of space, read only the right halves of lines of text (e.g., Gilliat & Pratt, 1952; Diller & Weinberg, 1977),



and of isolated words. When reading text these patients do not neglect specific words. Rotating text 90 degrees improves reading dramatically (Kinsbourne & Warrington, 1962; Ellis, Flude, & Young, 1987), so that the vertical lines on the left of the page are now read as well as the lines on the right. The data are compatible with the idea that left-right horizontal control of attention is selectively compromised in some forms of neglect.<sup>29</sup> Attentional scanning from left-right at a variety of scales appears to be impaired, though the effect of scale is under-researched. Possibly control at fine scales is more vulnerable, because one patient has been reported who neglects isolated words but not text (Patterson & Wilson, 1988), and the converse has not been reported.

Ellis et al. (1987) claim that neglect dyslexics errors tend to preserve word length because letters are encoded in their positions in the word, as 1st, 2nd, ..nth. (i.e. part-whole encoding). Their evidence is weak: taken together with the errors reported by Costello and Warrington (1987), of 169 errors on three and four-letter words, 78 preserve length, 7 are deletions, and 84 are additions.<sup>30</sup> More than 50% of these errors do not preserve length, and not all of the deletions and additions involve just one letter. More plausibly, coarse information is available from the neglected letters; the errors might, therefore, be expected to preserve word-shape.

Analysing the Costello and Warrington corpus produces suggestive, but not strong, evidence: of 92 three- and four-letter word errors, 54 preserved word shape, strictly defined in terms of ascenders and descenders, and 15 of the errors that did not involved the addition of "st" which may be a particularly frequent beginning to words. The general pattern of errors in both patients are better explained in relation to fixation point. Assuming that they fixate on the centre of words, to the right of fixation they are always correct; errors are made to the left of fixation, usually the leftmost letter only (which is deleted or replaced by one or more letters), especially for three- and four-letter words, but also often the next letter in.

---

<sup>29</sup> Such a deficit is analogous to the selective impairments to voluntary eye-movements and attention shifts in the vertical direction (Posner, Rafal, Choate, & Vaughan, 1985). See also recent evidence suggesting that focussing attention brings about a left-right perceptual organisation "that predominates over that provided by other egocentric reference axes" (Nicoletti & Umiltà, 1989).

<sup>30</sup> This re-analysis is my own.

This effect is shown most clearly for five-letter words, because fixation point for four-letter words is between the second and third letters: 90 errors on the leftmost, 40 on the second letter. Words longer than five letters do not show such a clear pattern, presumably because normal readers often need to refixate words longer than the foveal span of roughly six letters. Errors are less common on the leftside nearest-to-fixation letter because it is constrained by the letters to the right, an influence which declines for more leftward letters, and also because the first letter is more informative than the second so that any coarse information about the leftmost position, such as an ascender-descender description, constrains the identity of the second letter to a smaller range of probabilities.

If words and nonwords can be dissociated in terms of their representations then neglect might be expected to be more severe for nonwords and, more crucially, to preserve word length more for words than nonwords. Ellis et al (1987) do not report either pattern in their patient using pronounceable pseudowords for the comparison. For digit-strings, however, the prediction appears to be upheld: as compared to word errors the errors on digit-strings show no tendency to preserve word length. In addition several reports now exist of the dissociation between word and unpronounceable-nonword neglect with severer neglect for nonwords (Friedrich, Walker, & Posner, 1985; Sieroff & Michel, 1987; Sieroff, Pollatsek, & Posner, 1988). Unfortunately these reports do not assess error types in more detail across words and nonwords.

**Attentional dyslexics** (Shallice & Warrington, 1977b) are impaired selectively at reporting the identity of elements from multi-element displays, but only if all the elements are of the same visual category (letters, digits, dots etc.). Attentional dyslexics, therefore, make large numbers of migration errors. These are examined in detail in the next section.

In multi-word pattern-masked displays a word that was not present but could be made by a letter migrating from one word to another, e.g., "sand-lane", is sometimes reported, e.g., "land-lane" (Allport, 1977). This is the standard letter migration error. The strong claim has been made that migrations respect position-in-the-string and are therefore evidence for part-whole encoding (Allport, 1977; Duncan, 1987; McClelland & Mozer, 1986; Shallice & McGill, 1978; Mozer, 1983). A similar finding comes from the word search

paradigm: search is slowed when the distractor strings share some letters with the target, and especially so when the shared letters are in the same position within the word (Flowers & Lohr, 1985).

The claim of position specificity, however, has not been rigorously tested, and indeed may be the result of the very limited word sets used in migration experiments. Moreover, in a slightly different paradigm (successive rather than simultaneous presentation, and subthreshold presentation of the first string) no tendency for migrations to respect position was found (Humphreys, Evett, & Quinlan, 1990). One other difference between this paradigm and the more usual migration experiments was that the first string was a nonword and the second a word. Another anomaly is that parafoveal preview apparently never produces migrations to the report of a foveal word (McConkie, Zola, Blanchard, & Wolverton, 1982). Most migration experiments have only used word stimuli; one that compared nonwords and words did not report data on the position specificity of migrations (Treisman & Souther, 1986).

A strong neighbourhood effect is found on migrations, similar to the neighbourhood effects discussed earlier but one from which slightly different conclusions are usually drawn; this is the effect of context on migration errors. This effect, the "*surround*" or "*context*" effect, is that migrations are more likely between two words the more similar those two words are to each other, in terms of number of common letters in position. Shared context displays also reduce overall accuracy (McClelland & Mozer, 1986; Shallice & McGill, 1978; Mozer, 1983). This effect is usually interpreted in terms of competing word representations, activated by letters-in-position, but since context conditions also maintain larger numbers of inter-letter relationships, another possible implication of this neighbourhood effect is that relations between parts do contribute to migrations, contrary to the claim made on behalf of the position specificity of migrations. This implication, however, would be unwarranted if the context effect only applies to words and not nonwords. Indeed, a reverse of the context effect is found for intrusions from subthreshold similar nonword primes to word targets (Humphreys, Evett, & Quinlan, 1990).

If the context effect applies only to words an interpretation in terms of some feature of word representations not shared by nonword representations, such as word shape, might be

possible, though interactive networks which feed back activity from word to letter representations, and do not use overall word shape information, have been shown to be capable of generating the context effect (e.g., McClelland, 1986). McClelland and Mozer (1986) found that the surround effect does not apply to letters embedded in digit strings, which suggests that it is not similarity defined in purely physical terms that mediates the context effect. This does not provide an alternative definition of similarity, except that it is limited to similarity within letter and word representations.

Migrations are less likely into lower-case than upper-case words, and tend to preserve the ascender-descender characterisation of the word (McClelland & Mozer, 1986), which suggests that word shape information acts as a constraint on potential migrations. However, because the potential migratory letters in this experiment were not presented in the same case as they would have migrated into, the visual characteristics of the opposite case of a presented letter constrain the likelihood of migrations; this suggests that migrations are the result of interactions between VIP and abstract graphemic representations; similarly migrations are not reduced between different case words and letters (Shallice & McGill, 1978; McClelland & Mozer, 1986). Thus migrations should interact with the regularities in graphemic representations. But in a search task neither lexicality nor pronounceability protects against migrations, and lexical and pronounceable items are not much more likely to be formed after migrations; the effects obtained were small, and disappeared in a full report version of the task (Treisman & Souther, 1986). Almost certainly the negative result is due to lack of sensitivity of full report and search measures because contradictory results are obtained in the cued-report paradigm: more migration errors on pseudoword stimuli, and the results more likely to be words than nonwords (McClelland & Mozer, 1986).

Evidence is also accumulating that migrations are subject to semantic influences, essentially that they are more likely when they fit semantically with the context (or prime) in which they are presented (Shallice & McGill, 1978; Strain & Cowie, 1989 unpubl.; van der Velde, van der Heijden, & Schreuder, 1989). This effect suggests that interactions at the whole-word level of representation are responsible for generating migrations.

#### 1. 4. 4 Summary

The evidence reviewed above addresses the issue of the representation of the relationships between letters, and between letters and word-wholes. Some of this evidence has been used to support the idea that letters are encoded in terms of their position in relation to the whole of which they form a part; detailed examination of this claim suggests that, at the least, the claim has been overstated. Evidence from similarity priming makes clear that relationships between letters are also involved in some experimental contexts. Which contexts and which relationships remain unclear.

A further line of evidence was introduced, the neighbourhood frequency effect, which can be taken as support for a clear distinction between the representations and processing of words and nonwords. This raises the possibility that it is particularly nonwords that are represented in terms of part-part relationships. This possibility was then explored with reference to the literature on positional errors in the report of singly-presented letter-strings. Not only was the distinction upheld but no evidence was found that letters are accurately encoded in terms of their absolute positions in the string. Neuropsychological evidence from neglect dyslexics appears to contradict this viewpoint, but there is at least a suggestion in the literature that neglect affects words and some nonwords differently. No information is available from these studies as to how this distinction bears on the question of the encoding of spatial relations. Finally, the pattern of errors made in reporting multi-string displays is subject to the same qualifications. A priori the errors appear to be evidence for part-whole encoding, but this claim has not been subjected to rigorous testing. Further, the distinction between words and nonwords in terms of migration errors has not been explored in depth, and appears to be unresolved. The research reported in this thesis addresses some of these unresolved issues.

### 1. 5 Experimental methods and predictions

The research used letter-strings and words as experimental stimuli for the following reasons. Firstly ready definitions are available for what counts as "parts" and "wholes".

Both syllables and single letters potentially act as parts in clearly defined ways. Their relationships to the whole string are also transparent. Similar control over experimental conditions is not possible with objects and object parts as stimuli. Secondly, some detailed computational models have been developed using word recognition as their domain of application. This enables detailed analysis of the predictions they make under various experimental conditions. Thirdly, order information is a form of spatial information that is especially salient for word and letter-string stimuli, and can be defined either in part-part or part-whole terms.

One direct way to compare part-part and part-whole descriptions is to examine how experience with one letter-string transfers to the processing of another letter-string depending on which relationships are maintained across the two letter-strings. This paradigm, similarity priming between nonwords, allows comparison between letter-string pairs that maintain relational order information with letter-string pairs that maintain both relational and ordinal order information. The paradigm also allows different types of relationships between parts to be compared for their efficacy in producing transfer. Similarity priming between nonwords is used to investigate transfer between two processing experiences with letter-strings.

A slightly different paradigm that has been used to investigate this same general issue is prototype extraction. Here, processing experiences with particular regularities are built up through repeated exposure to letter-strings containing those regularities. By manipulating the regularities present in the repeated experiences the ease of extracting different relational regularities can be assessed. These two paradigms, similarity priming and prototype extraction, are used in the first five experiments. These are described in Chapter 2 and Chapter 3.

Chapter 4 describes experiments that use slightly different paradigms. As a baseline for this investigation the errors made in reporting the positions of letters in singly presented letter-strings are analysed in detail. Then transfer of experience between two letter-strings is examined again, but looking at errors made in reporting the second string. The errors of most interest are letters reported that were present in the first but not the second letter-string,

migration errors. This allows measures to be taken of the probability that migrations maintain position-specificity, and whether this probability differs for words and nonwords.

Simply put, part-whole theories predict that the relationships between parts are not as important as the relationships between parts and the whole, and thus that the part-whole relationships should dominate processing. Thus when part-whole relationships are disrupted but part-part relationships maintained, processing should also be disrupted. Similarly ease of learning should be more dependent on maintaining part-whole relationships than on the maintenance of any additional part-part relationships. Finally errors made during processing should show evidence of part-whole encoding as in the incorrect report of a letter from a preceding letter-string as being present in a subsequent letter-string and in the same position. These predictions are tested in the experiments reported in the following three chapters.

---

# WHOLE-PART RELATIONS IN THE TRANSFER OF TRAINING

---

## 2. 1 Introduction

As discussed in Chapter 1 similarity priming is potentially a useful tool for examining the encoding of relational information. The most direct application of the paradigm to this issue was reported by Humphreys, Evett, and Quinlan (1990) using subthreshold primes and word targets. The relevant results of their investigations are summarised next. (i) Similarity priming increases, non-linearly, as the number of shared letters increases, when the shared letters are in the same positions in both source and target string (Experiment 1). (ii) Similarity priming does not obtain when shared letters maintain neither fixed nor relative positions (Experiment 2). (iii) Similarity priming obtains when end letters are maintained as end letters but in different absolute positions on the display. This priming increases when the end letters' immediately interior neighbours are also maintained. Priming is also obtained when internal letters are maintained as an internal cluster (Experiments 4 and 5). (iv) Similarity priming is not increased when absolute as well as relative positions are maintained, at least for end-letters (Experiment 6). What conclusions can be drawn from these results?

The first two results appear to indicate that fixed position priming is more powerful than moved position priming (Experiments 1 and 2), whereas the second two results indicate the reverse (Experiments 4, 5 and 6). This contradiction is resolved by developing a definition of the parts that are either fixed or moved. When the part is a single letter no moved-part



transfer is obtained, but when the part is two adjacent letters, or two adjacent letters and an end space, moved-part transfer is obtained. Single letter parts do not produce fixed-part transfer either: when only a single letter is maintained in fixed position no transfer is obtained unless it is the first letter in which case it is a part made of a single letter *plus* the preceding space.

However the demonstration of moved-part transfer can be criticised as inadequate. In Experiment 4, for example, it was found that BLCK primed BLACK, and this was taken as evidence of moved-part transfer because the Interactive Activation coding scheme (McClelland & Rumelhart, 1981) predicts no transfer under these conditions. This claim is difficult to assess because Interactive Activation has no provision for processing words of different lengths, and in any case, a simple assignment of a positional slot to each letter would allow priming from B/1, L/2 in both strings. This criticism is even stronger when addressing other part-whole theories because a coding scheme that labelled the slots as First....Last (Monk, 1985) would also predict transfer, especially if the interior letters were coded relative to the end letters. More generally, the degree of movement over which transfer is maintained is not large enough to be convincing evidence of moved-part transfer. The paradigm developed (independently) for the research reported below used a more convincing test of moved-part transfer.

## 2. 2 Experiment 1: Fixed and Moved-part transfer

The experiment reported here develops work by Begg (unpublished, 1988). Using the transfer paradigm, Begg asked whether there was any perceptual facilitation when a part of an original whole was presented in the same relationship to a new whole, and when a part was presented in a new relationship to a new whole. The stimuli she used were six-letter pronounceable nonwords, such as SOLMEP, which were learned for ten seconds. These nonwords can be easily broken down into two constituent syllable parts, SOL and MEP. Half her subjects learned SOLMEP on its own, half in addition learned SOL or MEP. Test stimuli were presented for 120 ms, and pattern masked for 140 ms, after which

subjects were asked to type in what they had seen. The data were scored in two ways: number of letters correct over all six letters, and number of letters correct over the three letters that came from the learned letter string. For example, in the moved-part condition **PIKSOL** is a test stimulus if **SOLMEP** had been learned; if any transfer took place it would be most likely to show up as improved performance on the moved part itself (**SOL**), rather than on the new, unfamiliar part (**PIK**). Accordingly the score on the moved part itself may be of most interest. (These three relevant letters, whether as a moved or as a fixed part, are the "*focal trigram*".)

Her results show evidence of both fixed-part transfer and moved-part transfer, but only when both the whole and the part are learned separately. When only the whole is learnt there was no significant transfer, but the data show trends in that direction. Possibly with longer learning than ten seconds those trends might become significant. In the moved-part condition the position of the parts in relation to the whole is changed whereas some of the relative positions of the parts are maintained. Begg's data thus suggest that the position of letters relative to each other is encoded.

Her data show that in absolute terms performance is better on the fixed-part than the moved-part conditions. A further statistical analysis of her data shows that when both sets of results are summed (i.e., over subjects who learned only the six letter word, and subjects who learned both the six and the three letter word) the difference between the fixed-part and the moved-part conditions is significant,  $t(7) = 2.83, p < .05$ . This difference applies to the results scored over all six letters. It is possible that part-whole encoding is used in addition to part-part encoding, allowing the difference in performance to be attributed to the maintained absolute positions of the fixed-part. Alternatively, it may still be attributable to part-part encoding, because more of the overlapping relationships (as in the Wickelgren encoding scheme) are maintained in the fixed than in the moved-part condition.

However there are several limitations in Begg's data. The main problem is that each learnt word was tested for each condition more than once. This may have allowed learning of the stimuli to continue during the test phase, so that, for example, the moved-part conditions would no longer fit the requirement that they provide a new relationship to a new whole. This may mean that superior performance on any of the experimental conditions as

compared to the control can be explained as priming effects at the level of individual letters, through the internal representations of the letters being in a primed state from recently having been activated. The experiment reported in this chapter attempted to control for these limitations and examine in more detail the relative size of fixed- and moved-part transfer.

## 2. 2. 1 Method

*Stimuli and Design:* All stimuli for this experiment were unpronounceable letter strings<sup>1</sup>. These were generated pseudo-randomly from the whole alphabet, with the proviso that easily pronounced letter strings were not allowed (see Appendix A for examples). Letters were sampled with replacement. Stimuli were strings of either six or three letters, presented in upper case with no spaces between the letters. Each string was bordered on either side by an indented arrow, two spaces away from each end letter. There was a distance of 4.5 cm between the 2 indented arrows. The masking stimulus was a row of six "eights" (i.e. "888888"). Each letter was 5 mm high and 3 mm wide, presented in green on a black background. Viewing distance was not controlled, but the distance of the keyboard from the screen constrains it to roughly 30 cm, so the total display subtended 8.5 degrees of visual angle, and each letter roughly 0.9 degrees of visual angle.

The design differed in several ways from Begg's experiment. To examine in more detail the contribution of learning a part of the letter string as an additional whole, whole-part transfer was directly, and separately, tested. An additional conservative control condition, referred to as the **letter-prime control**, was also introduced. This was a string similar to that of the fixed-part conditions: three of the letters from the learned word were maintained, the same three that were tested in the fixed-part trials, but their order was jumbled. If improved performance on the fixed-part condition is simply due to priming of the

---

<sup>1</sup> The stimuli used in the Begg experiment were pronounceable nonwords, in which the part boundaries always coincided with the syllable boundaries. This has the effect of making each part more "whole-like", and thus may reduce the likelihood of each part being perceived as a part of a larger whole. It also increases the possibility that any transfer obtained is due to phonological recoding.

representations for specific letters, without any information about the relative positions of each letter, then performance on the letter-prime control would be better than performance on the ordinary, random control condition and as good as on the moved-part condition. Thus any differences between the experimental and the letter-prime control conditions, if found, can be treated as stronger evidence than differences from the ordinary control.

A within-subjects design was used with repeated measures, two learning conditions and six test conditions. The independent variable was the relation of the test conditions to the learned letter string and the dependent variable was accuracy of response, scored over all six letters and, for some of the conditions, over the relevant three letters, the focal trigram. Half the letter-strings that were learned were three letters long, half were six letters long. In the descriptive terminology used by Humphreys, Evett, and Quinlan (1990) the learned six-letter string is represented as 123456, and the test strings for the learned six-letter string are as follows. (In this terminology, "d" represents a different, randomly selected letter.)

1. **The Prime condition:** the same letters as originally learned: 123456.
2. **Fixed-part transfer:** a letter string made up of three of the letters of the original string in the same positions, but with the other three letters changed. On half the trials it was the first three letters that were kept constant: 123ddd; on the other half it was the last three: ddd456.
3. **Moved-part transfer:** a letter string made up of three of the letters of the original string but in different positions, and with the other three letters changed. On half the trials the first three letters were kept constant but moved to the last three positions: ddd123; on the other half the last three letters were constant but moved to the first three positions: 456ddd.
4. **Letter-prime control:** a six-letter string in which three of the letters of the original string were kept constant but their absolute positions altered. Half the trials used the first three letters of the learned word: 312ddd; the other half used the last three: ddd645.
5. **Six-letter control:** six different letters, randomly chosen without replacement: dddddd.
6. **Part on its own (part-to-whole transfer):** a three-letter string made up of three of

the letters of the original string: half the trials had the first three letters: 123; half had the last three: 456.

### 7. Three-letter control: ddd.

*Subjects:* 10 members of Stirling University, a mixture of undergraduates and staff; seven were members of the Psychology Department. Six were male and four female. Ages ranged from 19 to 50. All were voluntary participants in the study, and had normal or corrected to normal vision.

*Apparatus:* An Apple IIe microcomputer was used to generate and present the visual stimuli, and to record and score subjects' responses. The computer was set up in a darkened sound-proof cubicle and subjects were seated in front of the screen, within easy reach of the keyboard. Lighting was provided by a standard 60 watt reading lamp.

*Procedure:* Subjects were seated in the experimental cubicle in front of the VDU and introduced to the apparatus; the computer chose randomly whether to present first the part-whole or the whole-part condition, generated the appropriate stimuli, and presented a letter string to be learned. Stimulus durations for the letter strings and the mask were chosen by the experimenter and typed in. Instructions were presented on the screen, telling subjects that they would have to learn a letter string, and then be tested on other strings, some of which would be identical to the original string, some variations on the original, and some completely different.

Six-letter strings were presented for three minutes of learning, three-letter strings for one minute. Subjects were asked before the experiment began actively to test themselves by looking away and writing down the letter string during this learning period. When the learning time was finished subjects were asked to type in the string they had learned. For each six-letter string learned there then followed 11 test trials, chosen at random from the conditions described in the Design section. For the conditions described as having two sub-conditions, both sub-conditions were presented for each learned letter-string. For each three-letter string that was learned the test conditions were randomly chosen from the following six conditions: Prime; Whole-to-part transfer, first three letters; Whole-to-part

transfer, last three letters; Letter-prime control, first three letters; Six letter control; Three letter control.

Each test stimulus was presented for 100 ms, followed immediately by the pattern mask for another 100 ms. After presentation of each test stimulus subjects were asked to type in the letters they had seen, and then to proceed (by pressing Return and then Space Bar) at their own pace to the next test trial.

Subjects learned eight letter-strings, four of six letters and four of three letters, alternately one of each. The six-letter learned strings were followed by 11 test conditions each, the three-letter learned strings by 6 test conditions each. Thus 68 test trials were given to each subject. Before test trials began subjects were given 10 practice trials with the same stimulus durations and procedure as on the experimental trials. The experimental session lasted for about fifty minutes, and was followed by a debriefing and the opportunity to ask questions about the experiment.

## 2. 2. 2 Results

The results for the two phases of the experiment (learning 123456 and learning 123 respectively) are presented separately.

### Part-to-part and Part-to-whole transfer

(a) In absolute terms the prime condition produced more transfer than the fixed-part condition which produced more than the moved-part condition. Percentage correct over all six letters was 36% for the Control, as compared to 90% for the Prime, 53% for the Fixed-part, and 49% for the Moved-part conditions. Averaging over the focal trigrams the percentage correct is 82% for the Fixed-part and 56% for the Moved-part conditions. As a percentage of the amount of transfer produced by the prime condition, over all six letters the Fixed-part condition produces 31% and the Moved-part condition 24% facilitation. The average of the two focal trigrams transfer, separately compared with focal trigram control scores, is 54% for the Fixed-part and 43% for the Moved-part conditions. Table 2. 1

summarises the results for each condition.

Condition	Mean		
	All Letters	First three	Last three
Prime	5.4 (0.9)	2.9 (0.3)	2.5 (0.7)
Control	2.2 (0.6)	1.9 (0.5)	0.2 (0.2)
<b>First three letters as focal trigram</b>			
Fixed-part	3.0 (0.8)	2.5 (0.6)	-
Moved-part	2.9 (1.4)	-	1.0 (0.9)
Letter-prime Control	2.3 (0.5)	1.9 (0.5)	-
<b>Last three letters as focal trigram</b>			
Fixed-part	3.3 (1.6)	-	2.4 (1.0)
Moved-part	2.9 (1.2)	2.4 (0.7)	-
Letter-prime Control	2.3 (0.8)	-	0.3 (0.2)
<b>Part-to-whole conditions</b>			
First three letters	2.8 (0.5)	-	-
Last three letters	2.8 (0.2)	-	-
Control	2.5 (0.5)	-	-

*Note: The scores in the All Letters scored column are out of six for the first eight rows, out of three for the bottom three rows; the scores for the Part-whole conditions are all out of three. Standard deviations of the scores are presented in brackets. Focal trigram refers to the three letters taken from the learned letter string; in the Moved-part conditions the first focal trigram moves to the last three letters in the string, and the second focal trigram to the first three letters. A "-" sign means that this cell does not exist for the particular condition. Chance is 0.2 for the six-letter scores, and 0.1 for the three-letter scores.*

**Table 2. 1.** Mean Number of Letters in Position Correctly Reported as a Function of Transfer Condition and Scoring Method: Part-to-part and Part-to-whole Transfer, Experiment 1.

The difference between performance on the prime condition and the control condition is highly significant,  $t(9) = 14.96$ ,  $p < .001$ . Both halves of the prime condition are highly significantly different from the relevant halves of the control string,  $t(9) = 8.14$ ,  $p < .001$ ,

and  $t(9) = 12.1, p < .001$ . These results show that there is transfer to the prime condition, i.e., that learning successfully took place. All t-tests presented in this section refer to one-tail t-tests, because the direction of transfer has already been predicted from Begg's results.

There are no major differences between the letter-prime control conditions and the ordinary six-letter control. None of the comparisons, over focal trigram or over all six letters, averaged together over absolute position or taken separately, were significant.

(b) The finding of central importance is that **moved-part transfer** is obtained. The two Moved-part sub-conditions (ddd123 and 456ddd) were added together and averaged to control for absolute position effects. For the scores over all six letters, performance is significantly better than on the Control condition and better than on the two letter-prime control sub-conditions added together and averaged,  $t(9) = 2.7, p < .05$  and  $t(9) = 2.22, p < .05$  respectively. For the focal trigram scores, the difference from the two letter-prime control scores is significant,  $t(9) = 2.59, p < .05$ , and the difference from the Control condition is highly significant,  $t(9) = 3.25, p < .01$ .

(c) Significant **fixed-part transfer** is also obtained. Performance on the two Fixed-part conditions (123ddd and ddd456) added together and averaged is significantly better than on the Control,  $t(9) = 3.83, p < .01$ , and the letter-prime control conditions (added together and averaged),  $t(9) = 8.95, p < .001$ . For the focal trigram scores, all the comparisons give the same outcome: the two Fixed-part conditions together are highly significantly better than the Control condition,  $t(9) = 4.78, p < .001$ . They are also highly significantly different from the letter-prime control conditions,  $t(9) = 4.48, p < .001$ .

(d) There is a small absolute difference between performance on the Fixed-part and on the Moved-part conditions, but this does not reach significance, either over all six letters,  $t(9) = 1.01, p = .17$ , or over the focal trigram,  $t(9) = 1.4, p = .10$ .

The Fixed and Moved-part conditions were also analysed for effects of absolute position: the same pattern emerges for both conditions. In the Fixed-part conditions when scored over all six letters there is no difference between the two sub-conditions,  $t(9) = 0.8, p = 0.22$ ; but for the scores on the focal trigram this difference becomes highly significant,  $t$



(9) = 5.48,  $p < .001$ . When scored over all six letters, both sub-conditions are individually significantly different from Control,  $t(9) = 4.29$ ,  $p < .001$  and  $t(9) = 2.81$ ,  $p < .01$  respectively. Both sub-conditions can be compared on focal trigram scores with the relevant three letters of the Control: the first three letters for the Fixed-part condition 123ddd, and the last three for the Fixed-part condition ddd456. Both sub-conditions show highly significant differences from their controls,  $t(9) = 4.13$ ,  $p < .01$  and  $t(9) = 4.08$ ,  $p < .01$  respectively.

Looking at the difference between the two Moved-part sub-conditions, the same pattern is found: no difference when looked at over all six letters,  $t(9) = 0.1$ ,  $p = 0.46$ , but a highly significant difference when looked at over the focal trigram,  $t(9) = 6.4$ ,  $p < .001$ . Comparing each sub-condition individually with the Control, both over all six letters and with the relevant three-letter part of the Control divided into its two constituent parts produces the following results. The first half Moved-part condition (ddd123) is significantly different over all six letters,  $t(9) = 2.2$ ,  $p < .05$ , and significantly different over the focal trigram  $t(9) = 3.07$ ,  $p < .01$ . The second half Moved-part condition (456ddd) is also significantly different from the Control on both scoring measures,  $t(9) = 2.94$ ,  $p < .01$ , and  $t(9) = 3.04$ ,  $p < .01$  respectively.

(e) There is evidence for significant **part-to-whole transfer**, i.e., when 123 or 456 are presented on their own at test as wholes. The comparisons are made with the three-letter control rather than with the six-letter control. The difference between the first part-to-whole transfer condition and the Control just misses significance,  $t(9) = 1.69$ ,  $p = 0.052$ . The other part-to-whole condition, however, shows a significant difference from the Control,  $t(9) = 1.82$ ,  $p < .05$ . There is no difference between the two part-to-whole conditions,  $t(9) = 0.14$ ,  $p = 0.45$ , suggesting that it does not make any difference to part-to-whole transfer whether the part comes from the beginning or the end of the original whole.

### Whole-to-part transfer

The results are summarised in Table 2. 2. There is significant evidence of direct transfer: the comparison between the Prime and the three-letter Control conditions is

significant,  $t(9) = 2.54, p < .05$ .

Condition	Mean		
	All letters	First three	Last three
Prime	2.8 (0.4)	--	--
3-letter Control	2.5 (0.3)	--	--
6-letter Control	2.6 (0.6)	2.2 (0.4)	0.3 (0.4)
<b>First three letters as focal trigram</b>			
Whole-to-part	3.0 (0.6)	2.6 (0.6)	-
Letter-Prime Control	2.2 (0.6)	2.0 (0.6)	-
<b>Second three letters as focal trigram</b>			
Whole-to-part	3.5 (1.4)	-	1.4 (1.1)

*Note: The scores in the All Letters Scored column are out of three for the first three rows; the scores for the other rows are all out of six. Standard deviations of the scores are presented in brackets. Focal trigram refers to the three letters taken from the learned letter string; in the Whole-part conditions the first focal trigram is the first three letters in the learned string, and the second focal trigram is the last three letters. A "--" sign means that the scores for these cells were not analysed. A "-" sign means that this cell does not exist for the particular condition. Chance is 0.2 for the six-letter scores, and 0.1 for the three-letter scores.*

**Table 2. 2.** Mean Number of Letters Correctly Reported as a Function of Transfer Condition and Scoring Method: Whole-to-part Transfer, Experiment 1.

(a) There is significant evidence of whole-to-part transfer. This comes from the comparisons between the two whole-to-part conditions, and the two controls, standard and letter-prime. The first-half whole-to-part condition (learning 123) is significantly different from the Control, over all six letters,  $t(9) = 2.35, p < .05$ . On the focal trigram analysis this comparison is also significant,  $t(9) = 2.02, p < .05$ . It is also highly significantly different from the letter-prime control condition on both scoring measures,  $t(9) = 6.13, p < .01$ , and  $t(9) = 3.86, p < .01$ , respectively.

The second-half whole-to-part condition (learning 456) is also significantly different from the control condition on both scoring measures,  $t(9) = 2.27, p < .05$ , and  $t(9) =$

3.38;  $p < .01$ , and highly significantly different from the letter-prime control over all six letters,  $t(9) = 4.1$ ,  $p < .01$ .

(b) When scored over all six letters there is no difference between the two whole-to-part conditions,  $t(9) = 1.69$ ,  $p = 0.065$ , but on the focal trigram analysis there is a highly significant difference,  $t(9) = 5.71$ ,  $p < .001$ . There is no difference between the two Control conditions,  $t(9) = 1.43$ ,  $p = 0.095$  (over all six letters), and  $t(9) = 1.36$ ,  $p = .105$  (over the focal trigram, i.e., the first half of the control compared with the letter-prime control) showing that the improved performance on the whole-part as compared to the control conditions is not simply a matter of those particular letters being primed independently.

### 2. 2. 3 Summary

The results of this experiment, in summary, show that: there is significant evidence of fixed-part transfer; significant evidence of moved-part transfer; no significant difference in the amount of fixed and moved-part transfer, though the fixed-part transfer is greater in absolute terms; significant evidence of part-to-whole transfer (123456 to 123 or 456); significant evidence of whole-to-part transfer (123 to 123ddd or ddd123), no difference when the part is at the beginning or the end of the new whole, and no evidence of priming of letters independent of their position. The importance of the significant moved-part transfer is that it is not predicted by part-whole theories of the representation of relational information.

For moved-part transfer to be found there must be some encoding of the relative positions of the letters. If the letters were only encoded in their positions relative to the whole (as positions in the string, for example), then when those positions are changed the entire description would be changed and no transfer between the two would take place. That transfer does occur is evidence that more than part-whole descriptions are used. What the moved-part condition has in common with the fixed-part and the prime conditions, but not with the control or letter-prime control conditions, is that some of the positions of the letters

relative to each other are maintained. This is the only difference between the moved-part and the letter-prime control conditions. Since the transfer is attributable to the letters maintaining (some of) their relative positions, there must be a description that encodes the inter-relationships between parts of a whole that is distinct from a description of the relationships between the part and the whole.

Both part-whole and part-part theories predict fixed-part transfer. However its relative size (compared to the whole-word priming and moved-part facilitations) is of considerable interest. In Begg's experiment the fixed-part transfer was significantly larger than the moved-part transfer, whereas in Experiment 1 the difference was not statistically significant. Begg's first finding, that fixed-part is greater than moved-part transfer, is weakened by the limitations in her experimental design; in any case it only applies as a summary statistic over her two groups of subjects, one of which showed transfer from a six-letter word, the other from a six and a three-letter word learned together. In Experiment 1 the difference between the fixed and moved-part condition was not statistically significant. This is important because Begg suggested the possibility that her original finding could be interpreted as support for the idea that a part-whole code is used in addition to part-part representations. Even if her finding was valid, however, this conclusion is not inevitable: in fixed-part conditions some part-part coding systems maintain more relationships constant than in the moved-part conditions. For example, if the scheme encodes the first letter of the string and its relationship to the blank on its immediate left, then this particular relationship is not maintained in the moved-part condition when the first three letters of the learned word move to the last three positions. Thus, even if there is a difference in performance under the two conditions, this can be handled by part-part theories without additional part-whole coding.

The remaining positive findings to be discussed are the significant part-to-whole, and whole-to-part transfer. Both part-to-whole and whole-to-part transfer show that transfer is obtainable across different sizes of letter strings (and therefore, object frames or "wholes"). For transfer to occur across different sizes of object frames, either there must be a coarse-coded representation of the position of the part relative to the whole (a description such as "first letter out of six" would not suffice), or there must be encoding of the relationships between parts.

The whole-to-part transfer results show that there is transfer when the whole becomes a part both at the beginning and at the end (i.e., 123ddd or ddd123) of the new six letter string. This evidence suggests that any codes that use ordinal position to describe the relationships of the letters in a string do not make any necessary contribution to the transfer obtained. If 123 is described, for example, as "1-first, 2-second, 3-third", then there would only be transfer to 123ddd and none to ddd123. The results do not rule out the possibility that these codes are used, but if they are, they are not the only codes being used.

The same arguments apply to the part-to-whole transfer results: there was no difference in performance when the first three letters of a learned six-letter string were presented from when the last three letters were presented. An ordinal coding scheme predicts transfer for 123 and none for 456. The transfer for 123 was large in absolute terms, but just missed significance, whereas for 456 it was just as large, and significant.

The last finding that needs to be discussed is that there was no difference in performance between the control and the letter-prime control conditions. This suggests that no letter priming contributed to the observed transfer, i.e., that letter identity information without positional specification did not play a role. Only a scheme that does not use letter-identity-without-position information but still allows moved-part transfer can explain this pattern of results. One such scheme is to encode the positions of parts in relation to other parts; if relative encoding of position is used then when the relative positions of the parts are changed, as they are in the letter-prime control condition, no transfer is possible. This raises the questions: which relationships between which letters are encoded? And what defines a part that can be encoded relative to other parts? These questions are taken up in the General Discussion to this chapter and in the experiments reported in Chapter 3.

To confirm the claim that part-whole theories cannot deal with the moved-part transfer results, two simulation experiments that test two part-whole models under similar conditions are reported. In the first (Simulation 1) the Interactive Activation Model was tested for fixed- and moved-part transfer; in the second (Simulation 2) MAPPER was tested with the same conditions.

## 2.3 Simulation 1: Interactive Activation

The Interactive Activation Model, available as a CMU package, was tested on an Acorn Archimedes 310. The experimental paradigm was the same as in Experiment 1: measuring the effects of a prime string on subsequent related target strings. A prime word was presented for 16 processing cycles, followed by a pattern mask for 4 cycles; the test word, which bore one of four relationships to the prime word, was then presented for another 20 cycles. Because Interactive Activation only uses four-letter words, the parts for each experimental condition were defined as two neighbouring letters. On each cycle the model's accuracy of two-alternative forced-choice (2AFC) identification was recorded for one letter position. As in standard procedures (e.g., Reicher, 1969) the two alternative letters both were consistent with words in the model's vocabulary.

The model was tested with four four-letter word and four four-letter nonword primes; each word and each nonword was tested in a different one of the four positions to control for absolute position. All the word stimuli were taken from the list of words in the model's vocabulary. The experimental conditions were as follows; examples are from the actual word and nonword stimuli used in the experiment:

0. Prime:	PORT PMLQ
1. Full Prime target:	PORT PMLQ
2. Fixed-part:	POND PMSF
3. Moved-part:	CAPO WDPM
4. Control:	GAME OXVU

The results, grouped together into four blocks of five cycles each, appear in Table 2. 3. Significance tests were not performed because the results are clear-cut: fixed-part transfer is 100%, moved-part transfer is 0%. This confirms the analysis presented in Chapter 1, that part-whole encoding of this type is not capable of supporting generalisation when familiar parts of one whole appear in unfamiliar relationships to other wholes.

Condition	% Correct/ Block			
	1	2	3	4
<b>1. Words</b>				
Prime	81	89	95	98
Fixed-part	78	89	95	97
Moved-part	50	52	59	74
Control	50	52	59	74
<b>2. Nonwords</b>				
Prime	68	74	83	90
Fixed-part	68	74	83	91
Moved-part	50	53	60	72
Control	50	53	60	72

**Table 2. 3.** Identification Accuracy as a Function of Transfer Condition, and Cycle Block:  
Experiment 1. 2.

## 2. 4 Simulation 2: MAPPER

A simulation of MAPPER<sup>2</sup> was tested on an Acorn Archimedes 410; non-quantitative work with the simulation suggested that no moved-part transfer was evident in its performance. This work used variations among the patterns that it already knows, but to make more direct comparisons with Experiments 1 and 1. 2 the following paradigm was set up. By modifying the program it is possible to obtain read-outs of the activity levels in any of the processing units. The units in the object-centred plane were concentrated on because of their analogy with the letter-level units of IAM. MAPPER was presented on its retinocentric plane with patterns made up of six active elements; the activity levels of three of the object units were recorded. These three units maintained exactly the spatial relations of three of the active elements in the input pattern but could not maintain absolute positions. This is because the retinocentric plane is a 10 x 10 array while the object-unit array is 5 x 5.

<sup>2</sup> Written by Peter Cahusac at Stirling.

In order to avoid edge effects, which MAPPER is known to produce because of the nature of its algorithm, the patterns were presented as close to the centre of the retinocentric array as possible.

The same test pattern was presented for all experimental conditions, and the activities of the same three units were recorded. The first four cycles of activity were used for scoring; MAPPER changes its activity levels very rapidly, even when its computations are slowed by forcing them to take more into account of the most recent activity levels (inertia), and the most striking changes were observed within the first four cycles. The test pattern was preceded for eight cycles of activity by one of four priming patterns; activity in the mapping and object units did not decay in between presentation of each prime and test pair. The conditions were as follows:

0. Prime: the same pattern as on the test.
1. Control: an unrelated pattern that shared neither absolute positions nor relational positions with the test.
2. Fixed-part: a pattern in which three of the test pattern elements were present in the same absolute and relative positions and the other three elements are unrelated.
3. Moved-part: a pattern in which three of the test elements were present in the same relative positions to each other but in different absolute positions.

A total of four different test patterns, with corresponding priming conditions were presented. In general the results showed much more variability than in the IAM simulation, to the extent that the control pattern sometimes produced more priming than the full prime pattern. This variability is not surprising: MAPPER tries to construct an object-centred description even when it knows nothing about the patterns presented, and the biases in the algorithm that perform the mapping will favour some patterns over others. Nevertheless, averaged over all four test patterns, activity levels were higher after the fixed-part prime than after the moved-part prime: 0.7727 vs. 0.5384. The variation in activity for both conditions is from below 0.2 to above 1.1. The prime scores averaged slightly higher than the fixed-part prime at 0.885; the control scores averaged 0.687, only slightly below the fixed-part prime but larger than the moved-part prime, but showed enormous variability both



within and between patterns. The results appear to indicate a difference between the fixed and the moved-part prime conditions, but it is not clear that this paradigm was sensitive enough to pick up differences, because the fixed-part priming is only marginal.

## 2. 5 General discussion

Some general points of interpretation need to be made; these points will apply equally to all the experiments in the thesis. The experiments reported in this chapter were designed to test theories of the representation of relational information in visually presented letter-strings. Two problems immediately present themselves.

(i) The discussion in Chapter 1 assumes that relational information is represented in part of the specifically visual description of the input. Clearly it is possible that the same relational information could be represented by a non-visual description, a phonological recoding of the input, for example.

(ii) Further, it is assumed that the type of visual description which includes the representation of relational information is at the object/word level, and at the graphemic representational level in particular. If the type of relational information implicated in producing the transfer results reported in Experiment 1 is more easily attributed to a visual description other than graphemic representations, then inferences about the nature of graphemic representations become problematic. There are two obvious candidate alternatives: first that relational information is represented in a separate representational domain altogether; and second, that the relational information implicated is a function of a level of representation prior to graphemic representations, the Primal Sketch, or VIP in terms of Chapter 1.

As to the first problem, the candidate non-visual descriptions are all those which are possibly involved in the performance of the experimental task. This candidate list must include the graphic representations used to produce the typed letter-string responses. It may also include phonological descriptions since there is evidence that visual letter-strings are

automatically (i.e., unavoidably, uncontrollably) recoded into phonology (e.g., Dennis & Newstead, 1981; McCutchen and Perfetti, 1982) though this evidence applies to words rather than to nonword letter-strings. The same applies to semantic representations since the Stroop effect is evidence of obligatory semantic recoding, but only for word stimuli. Because the stimuli were all difficult to recode into phonology and meaningless there is no obvious reason to suppose that the transfer results depend exclusively on either phonological or semantic representations.

One clue to determining the involvement of graphic representations is provided by analysing the results in terms of left-right effects. This is because tasks that require full report of the stimulus in writing show a marked decrease in accuracy from the leftmost positions to the rightmost, presumably because the left positions benefit from being reported before the right positions. Inspection of Table 1. 2 reveals that performance is indeed better on the first three than the last three letters for all conditions, and thus that the left-right effect is operating. However the amount of priming is not less for the last three letters, indeed it is rather larger: the full prime condition produces a facilitation equivalent to one extra letter reported for the first three letters (2.9 - 1.9), and 2.3 letters for the last three letters (2.5 - 0.2). This is also true of the fixed-part (0.6 and 1.1 facilitation) and moved-part (0.5 and 0.8 facilitation) conditions. Thus whatever produces the transfer is not compromised by the left-right effect. Clearly this does not rule out the involvement of graphic representations in the generation of transfer effects, indeed it suggests an interaction between graphic and some other representations, but it demonstrates that the transfer is not wholly dependent on the left-right effect which is a defining symptom of graphic representations.

The second problem is more difficult to deal with. As discussed in Chapter 1, there are good reasons for distinguishing between the representation of object structure and object position (e.g., Hinton, 1981a; Phillips, 1983). It seems plausible that Hinton's "method of simultaneity" involves the binding together of the identity and spatial position of a recognised whole. If this is the case, then the question of the representation of the relationships of the parts and the whole can be distinguished from the representation of the position of the whole. It then becomes possible for a part-whole theory (such as Hinton's

“multi-dimensional method”) to be developed such that the precision of the representation of part-whole relationships is a function of the precision of the representation of the whole and its position. This idea is not incorporated into the part-whole theories whose simulations were reported in §2.3 and §2.4, but it is a natural development of Hinton’s (1981a) distinction between methods of representing positions. However the relations between parts and wholes are represented, this differs from the representation of the spatial relations of wholes.

This leaves the problem of whether to attribute the transfer results to the VIP, the low-level visual description, or to graphemic representations (the visual input lexicon). Taking MIRAGE as the model of the VIP has the implication that the VIP uses part-part relational descriptions of position. The results of Experiment 1 indicate that part-part relational representations are *also* involved in the processing of letter-strings, and thus, by definition, are a component of graphemic representations. The problem of distinguishing between VIP and graphemic representations, introduced in Chapter 1, is made acute by this result. Furthermore, as discussed in the next section, it is possible to use the type of coding used by MIRAGE to model accurately the results obtained in Experiment 1.

To summarise the discussion so far: before the results can be used to test theories of how relational information is coded in graphemic representations it is necessary to rule out other types of representation as possible causes of the results obtained. Phonological and semantic codes appear unlikely to have played much part. Graphic representations mediating the responses are definitely involved. Spatial representations independent of identity appear not to be heavily involved, but the role of non-graphemic, low-level visual descriptions is very difficult to assess. Whichever representations are implicated in the experimental results, however, the results apply to all of them equally well. There is no evidence, from a task that probably calls on VIP, graphemic, and graphic representations, that part-whole relations in any of these representational domains play much of a role in mediating the transfer obtained.

The next problem in interpreting the results in terms of theories of relational representations is to find a part-part representational scheme that can produce nearly

equivalent fixed and moved-part transfer. Before the predictions made by different schemes with can be compared with the experimental results, the appropriate measure from the results must be obtained. As percentages of the full transfer in the prime condition the fixed and moved-part conditions produced 31% and 24% transfer respectively over all six letters and 54% and 43% respectively over the focal trigrams.<sup>3</sup> Because predictions are being derived from the schemes about the encoding of relationships throughout whole letter-strings the facilitation over all six letters is used for comparison. The schemes looked at in the following discussion are all variants of the type of context-sensitive encoding proposed by Wickelgren (1969).

The first example is a Wickelgren-type scheme explored in Begg (1988). In this scheme

---

<sup>3</sup> The differences in these estimates can be used as the basis for two quite complex arguments about the non-independence of letter position processing.

The first argument is that transfer to the processing of a maintained part of a letter-string might be expected to facilitate processing of the non-maintained part, given the assumptions that the processing resources available for the task are (a) limited, and (b) stretched to maximum. Inspection of Table 2. 1 shows that this is not the case. First fill in the blanks in the table with the scores for the non-maintained part, obtained by subtracting the focal trigram score from the score over all six letters. For example, in the Moved-part, first three letters as focal trigram condition, the score for the non-maintained part is  $2.9 - 1.0 = 1.9$ . This score can then be compared with the relevant control focal trigram score, 1.9 in this case. Most comparisons of this sort reveal very small differences at the most, suggesting that the non-maintained letters are not facilitated at all. This absence of facilitation can be interpreted in two ways: (a) as evidence for the independence of letter positions in processing, because the maintained part does not influence processing of the non-maintained part, and (b) as evidence for the non-independence of letter position processing, because the expected facilitation, given the argument about processing resources, is counteracted by the disruption of relationships between the maintained and the non-maintained part across the source and the test string.

The second argument concerns the relative size of the facilitation for the two scoring methods, over all six letters and over the focal trigram. Under the assumption of independent processing of position the percentage of facilitation should be twice as great for the focal trigram scores compared to the all-letters scores. This is because the all-letters scores include the (non-facilitated) scores for three letters not present in the prime condition (where all letters are facilitated) and also not present in the focal trigram scores. The fact that the relative transfer is less than twice as great for the focal trigram scores is evidence for the non-independent processing of letter positions

all neighbouring parts are encoded in relation to each other and parts can be of any size from single letters upwards. The following example will serve to illustrate all the representational schemes examined in this section: a six-letter letter-string is represented as \*123456\*, where "\*" stands for the immediately preceding or following space and each number represents a letter position. The symbol "1-2" means either an explicit representation of the fact that "1" is to the left of "2" or simply a representation of the fact that they are neighbours. In the extended Wickelgren-type scheme under discussion a total of 84 relations exist. The fixed-part conditions of Experiment 1 are described as the following string: \*123ddd\*, where "d" refers to a changed letter. Under these conditions ten of the Wickelgren-type relations are maintained, as follows: \*-1; 1-2; 2-3; \*1-23; \*1-2; 12-3; \*12-3; \*-12; \*-123; 1-23. Maintaining 10 out of 84 relations should allow 12% of the transfer in the full prime condition, assuming that transfer depends on the similarity of representations in terms of relational descriptions. For the moved-part conditions only 5% of transfer is predicted because only 4/84 relations are maintained: 1-2; 2-3; 12-3; 1-23. This scheme, then, predicts much larger transfer in the fixed than the moved-part conditions, and transfer at a lower percentage than is found in Experiment 1.

The simple modification of the Wickelgren-type encoding scheme to allow only parts of one size to be represented in relation to neighbouring parts of the same size produces the following predictions. Of the 16 total relations, four (25%) are maintained in the fixed-part condition: \*-1; 1-2; 2-3; \*1-23. In the moved-part condition only two (13%) are maintained: 1-2; 2-3. Again the difference between the two conditions is much larger than the obtained difference.

A more complex version of the Wickelgren-type scheme allows any part to be represented in relation to any other *unspecified*, or unidentified, part. In this scheme the fixed-part condition maintains 49/84 (58%) of relations, while the moved-part condition maintains 11/84 (13%). Restricting this type of encoding to parts of the same size produces 8/16 (50%) transfer for fixed-parts and 4/16 (25%) for moved-parts. In both cases the differences are still too large.

Many other context-sensitive schemes are possible; McClelland (1986) in PABLO uses units that represent coarsely coded letters. This means that each letter is represented in a manner that depends on that letter's immediate neighbours. The letter "A", for example, can be represented by one or more of the following units: \*A; dA; Ad; A\*, depending on where it appears in a letter-string. Using this scheme the fixed-parts maintain 6/12 (50%) of relations, while the moved-parts maintain 4/12 (33%). The same scheme adapted to use trigram rather than bigram units (dAd, for example) predicts 2/6 (33%) transfer for fixed-parts and 1/6 (17%) transfer for moved-parts. Neither of these schemes offers much similarity to the obtained results.

The final coding scheme discussed here is derived from MIRAGE and is offered as speculation rather than in terms of precise predictions. MIRAGE produces descriptions of images at a variety of scales. If presented with a letter-string it is natural to assume that a coarse-scale description of the input would be centred on the fixation point, the centre of the letter-string. Assuming that a description at this scale of resolution does not provide enough information for accurate identification of letters in a nonword letter-string, descriptions at successively finer scales of resolution need to be made available. If the description at the coarsest scale is of a blob surrounding the whole letter-string, i.e., [\*123456\*], then one possible way to construct finer blobs is to divide the larger blob into two equal blobs centred around fixation, i.e., [\*123] [456\*]. These two blobs can then be subdivided in turn: [\*1] [23] [45] [6\*]. These four blobs then become six: [1] [2] [3] [4] [5] [6]; and these six "letter" blobs could in turn be subdivided into letter-feature blobs, like, for example, [/] [-] [N], to represent "A".

If this simplistic coding scheme is used clearly it is arbitrarily easy to set the scales so that the relationships represented within each blob fit the pattern of transfer results obtained in Experiment 1. Using the scheme outlined above, with two letter-feature blobs per letter blob, the fixed-part conditions maintain 12/25 (48%) of relations, while the moved-part conditions maintain 10/25 (40%) of relations. These particular figures are rather high, but they assume that the time needed to run through the particular range of scales is available for processing the input. Under the experimental conditions used in Experiment 1 this is

implausible. It may be the case that processing can continue to move from coarse to fine scales after the removal of the stimulus, but not as accurately as when the stimulus is present in the image<sup>4</sup>. This loss of accuracy would then be reflected in the obtained results being lower than predicted for ideal conditions. Since the processing at finer scales in the absence of continued visual information must depend on the processing already done at the coarser scales, it is an easy step to suggest that words are represented at the coarser scales in such a way that the coarse representations of familiar words are more helpful to continued processing at finer scales than are the coarse representations of unfamiliar words and nonwords.

This idea appears to have a natural affinity with the distinction presented in §1. 2 between processing letters as parts of word-wholes and as wholes within word-wholes: the process of refining the analysis of a coarse description would be analogous to switching from processing the coarse whole as made up of unidentified parts to processing the parts as wholes within the overall whole. One implication of these two suggestions is that nonword letter-strings do not receive much processing at the whole/coarse level, simply because they are not good perceptual wholes; thus, part-whole processing of letter-strings is a misleading term because letter-strings have to be processed as a collection of wholes within a larger whole. This allows two further questions: (i) What is the nature of the parts of the larger whole that are also processed as wholes; is it letters, or letter-clusters? (ii) Does processing letters as parts have different implications from processing letters as wholes for models of the representation of spatial relations? The first question is examined in Chapter 3, the second in Chapter 4.

To conclude: it appears possible that the transfer results of Experiment 1 can be dealt with by a model of visual processing that includes no verbal knowledge, and thus that the results implicate VIP rather than graphemic representations. Nevertheless, this model has

---

<sup>4</sup> The evidence of Rayner et al., 1981, that normal reading is possible with 50 ms masked word presentations, is suggestive of this point because overall fixation durations with these stimuli were not much changed.

the implication that the organisation of knowledge and processing within the graphemic representational domain is similar to the organisation of the VIP. The problem of distinguishing between the two domains becomes acute. This problem, and some of the speculations of the preceding paragraph, are returned to in Chapter 5.



---

## LOCAL AND NON-LOCAL RELATIONSHIPS

---

### 3. 1 Introduction

The four experiments reported in this chapter examine the first issue raised at the end of Chapter 2. Experiment 1 shows that part-part relationships are important, but says nothing about the nature of the parts that do matter. Assuming that letters are the relevant parts, the results show that letter inter-relationships are important in producing transfer under some conditions. An alternative is to consider the three neighbouring letters, the focal trigram, as the relevant part.<sup>1</sup> In this case the interpretation changes slightly because the results only show that trigram parts in particular relationships to the whole are not the *only* relations represented. To generate the transfer, the trigram part *may* be represented in relation to the whole, together with a relational description of the letters that make up the trigram part. Letter-clusters made of groups of three neighbouring letters were used in Experiment 1. Experiments 2, 3, 4, and 5 look also at letter-clusters made of non-neighbouring letters.

Part-whole theories of relational encoding make a clear, testable prediction: if encoding is simply of the relationship of the letter parts to the wholes, then differences in the relationships of a letter-cluster's constituents to each other that do not also differ in the relationships of the letter parts to the whole will not make any difference. In particular, the number and nature of a letter's neighbours will make no difference as long as the

---

<sup>1</sup> To distinguish single-letter parts from letter-cluster parts, the latter will be referred to throughout the introduction as "trigram parts", but throughout the rest of the chapter simply as "parts".

neighbours also maintain consistent part-whole relationships. If, however, in the more complex picture outlined above, letter-clusters, rather than letters, are the (trigram) parts represented in relation to a whole, then a difference in performance for different trigram part's inter-letter relations might only reflect differences in the relational coding of a letter-cluster's parts to each other and say nothing about trigram part-whole relationships. Non-locally defined trigram parts, might, for example, be particularly difficult to learn in the source string. In this case, then, when a letter-cluster's relationship to the whole changes there ought to be an interaction with the type of letter-cluster's letter-interrelationships. This is because the trigram part made of locally-related letters, being easier to learn, is more easily related to the whole than the trigram parts made of non-local letters. Stronger trigram part-whole representations might be more interfered with by changed part-whole relationships than weaker trigram part-whole representations

Thus these experiments look at two experimental conditions drawn up to differ only in the relationships among their parts. Treating letters as parts, the definition of interrelationships between parts used in Experiment 1 is one based on "neighbourliness", "degree of localness", or local relations. A non-locally defined relationship is one in which the letters that make up the relationship are not next-neighbours to each other, but are interspersed with other, irrelevant, letters. For example in the letter string 123456, 123 has locally defined relations, whereas 1-3-5 has non-locally defined relations. Treating letter-clusters as parts, 123 is a local part, 1-3-5 is a non-local part. These conditions allow examination of Mozer's (1987) suggestion that parts made up of at least some nonadjacent letters can be represented. Experiment 2 compares local and non-local trigram parts in terms of the amounts of fixed-part transfer they support; Experiments 3 and 4 compare the relative ease of learning local and non-local parts in prototype-extraction paradigms. Experiment 5 compares local and non-local trigram parts for moved-part transfer; the particular point of interest is whether changing a letter-cluster's relationships to the whole has equal effects on locally and non-locally defined trigram parts.

## 3. 2 Experiment 2: Fixed-part transfer in local and non-local parts

Experiment 2 uses the same paradigm as Experiment 1 to examine whether there are any differences in the amount of transfer obtained with fixed-parts, when those parts are either locally or non-locally (non-adjacently) defined. A theory of inter-item relational encoding based purely on neighbouring letters predicts no non-local transfer at all, whereas a theory that allows some generalisation over nonadjacent letters only predicts a large difference in performance. Any difference in performance on the local and non-local conditions is evidence against part-whole relational encoding theories, because they do not take into account the relationships among the items that make up a trigram part.

### 3. 2. 1 Method

*Stimuli and Design:* The stimuli were all strings of six consonants selected randomly without replacement (see Appendix B for examples). A within-subject, repeated measures design was used. The independent variable was whether the parts were local or non-local, the dependent variable was accuracy of response, scored over all six letters and over the three letters that were kept constant, the focal trigram.

*Subjects and Apparatus:* Nine undergraduates at Stirling University; five females and four males, participating as a course requirement. Ages ranged from 17 to 28. All had normal or corrected to normal vision. An Apple II microcomputer was used, as described for Experiment 1.

*Procedure:* Subjects were seated in a darkened, sound-proofed cubicle in front of the computer. All instructions were presented on the computer screen. Subjects were told that they would have to learn a letter string made up of six consonants, and then type in the string that they had learned, and subsequently respond to a series of letter strings that might be

similar or identical to the original letter string, again by typing in what they had seen. The letter string that was to be learned was presented for 3 minutes. Subsequent test strings were flashed on the screen for 100 ms, followed immediately by a mask made up of a row of eights for another 100 ms. Subjects were told to respond as quickly as possible but slow responses were not penalised.

Each subject learned a letter string and was then given eight repetitions of four test trial conditions, before learning another, unrelated, letter string and being tested with another 32 test trials. The test trial conditions were as follows:

1. **Prime:** the same consonant string as originally learned: 123456.
2. **Control:** six different consonants, randomly chosen: dddddd.
3. **Local fixed-part:** a consonant string made up of three of the letters that had been learned and three different letters. The three letters kept constant in this condition were neighbouring letters: half the time it was the first, second, and third; and on the other half the fourth, fifth, and sixth to balance out absolute position effects: 123ddd and ddd456.
4. **Nonlocal fixed-part:** a string made up as in the local condition, except that the three letters kept constant were not direct neighbours: either the first, third, and fifth, or the second, fourth, and sixth letters: 1d3d5d and d2d4d6.

### 3. 2. 2 Results

The results, scored as in Experiment 1, are summarised in Table 3. 1. As a percentage of the full amount of transfer obtained in the prime condition, the local fixed-part condition produces 22% transfer over all six letters and 56% over the focal trigram. The non-local fixed-part condition produces 13% transfer over all six letters.

The large difference between the scores for the prime and the control conditions shows that learning successfully took place. When scored over all six letters this difference is highly significant,  $t(8) = 4.96, p < .001$ . (As before all t-tests are one-tailed). The prime condition is also significantly better than both the local and non-local conditions,  $t(8) =$

4.14,  $p < .01$ , and  $t(8) = 4.77$ ,  $p < .001$ , respectively.

Condition	Letters Scored			
	All 6	1st F.T	2nd F.T	F.T. mean
Prime	4.4 (1.5)	2.4	1.7	2.1
Control	2.1 (0.6)	1.8	0.3	1.1
Local	2.6 (0.6)	2.3	0.7	1.5
Non-local	2.4 (0.5)	1.6	1.1	1.4

*Note: F.T. = Focal trigram. The scores in the All 6 Letters Scored column are out of six. Standard deviations are in brackets. The F.T scores are out of three. The focal trigram is the three letters taken from the learned letter string. In the Prime, Control, and Local-part conditions the two focal trigrams are the first three letters and the last three letters. In the Non-local conditions the two focal trigrams are the 1st, 3rd, 5th letters, and the 2nd, 4th, 6th letters. The scores in the All 6 column for the local and non-local conditions are the means of both the subconditions; the F.T scores for these conditions are specific to each subcondition. Chance is 0.2 for the six-letter scores, and 0.1 for the three-letter scores.*

**Table 3. 1.** Mean Number of Letters in Position Correctly Reported as a Function of Transfer Condition and Scoring Method: Experiment 2.

The critical finding is that there is a significant difference in performance on locally and non-locally defined parts of a whole: when scored over the focal trigram the local condition scores are significantly better than the non-local scores,  $t(8) = 2.45$ ,  $p < .05$ . Since all letter positions are included equally often in both conditions, this difference cannot be due to effects of absolute position.

There is evidence of significant local fixed-part transfer. Performance on the local condition is significantly better than on the control condition when the data are scored over all six letters,  $t(8) = 2.72$ ,  $p < .05$ . The mean local focal trigram score is significantly better than that of the control condition divided by two to make an appropriate comparison,  $t(8) = 3.04$ ,  $p < .01$ . When the local condition is broken down into its two sub-conditions (123ddd and ddd456), the local condition 123ddd does not differ significantly (scored over all six letters) from the control score divided by two,  $t(8) = 1.69$ ,  $p = 0.065$ . The

local condition ddd456 is significantly better, however,  $t(8) = 2.66, p < .05$ . Both these comparisons when looked at over the focal trigram scores (i.e., compared with the control scores over either the first or the second three letters) are significant,  $t(8) = 2.37, p < .05$ , and  $t(8) = 2.87, p < .05$ , respectively. The two local sub-conditions do not differ significantly from each other when scored over all six letters,  $t(8) = 0.72, p = .245$ , though on the focal trigram scores this comparison shows a highly significant difference in favour of condition 123ddd,  $t(8) = 11.09, p < .001$ .

The non-local condition overall was not significantly different from the control condition, either over all six letters, or over the focal trigram,  $t(8) = 1.46, p = .09$ , and  $t(8) = 1.47, p = .09$ , respectively. Neither of the non-local sub-conditions differ significantly from the control over all six letters. The two sub-conditions do not differ from each other when scored over all six letters,  $t(8) = .18, p = .43$ , but when scored over the focal trigram, condition 1d3d5d is significantly better than d2d4d6,  $t(8) = 3.29, p < .01$ .

### 3. 2. 3 Discussion

The results of this experiment can be summarised briefly: significant local, fixed-part transfer; no non-local fixed-part transfer; and a significant difference between performance on the two conditions. Thus there is a significant difference in performance on locally and non-locally defined parts of a whole, even when the definition of the part is the only distinguishing feature of the two conditions. Both conditions had the same absolute amount of similarity with the learned letter-string, in each case three out of the six letters were maintained, and across the two sub-conditions combined each letter position occurred equally often in each condition. The difference in performance shows that the internal description of the learned letter-string is not based on representations which give equal importance to each part, and treat each part as an independent entity in a particular relationship to the whole of which it forms a part. Only a theory which takes into account the relationships among parts can explain this result. Part-part encoding theory allows a

difference in performance whenever there is a difference in the number of part-part relationships that are kept constant.

In this experiment the local condition maintained the relationships among the items 1, 2, and 3 that make up the local part every time that the local condition was presented as a test trial for that particular learned letter string. On the other hand every time the non-local condition was presented its three items were in different relationships, 1d, 3d, and 5d, where "d" changed each time. Differential learning during test trials is thus one possible explanation for the difference in performance. However this explanation is not plausible, given that each test trial was only presented for 100 ms, and that brief inspection of the raw data shows no improvement in performance as the test trials continue. A more likely alternative explanation is that the difference is attributable to the construction of the internal description during the learning period. Either a part that is locally defined within the whole is more salient within the description (perhaps more richly or complexly described, or else more explicitly), or it is easier to learn such a part. These two are not mutually exclusive: they are different ways to express the same hypothesis. The second hypothesis, that locally defined parts are easier to learn than non-locally defined parts is tested in Experiments 3 and 4.

The description of the non-local part, on these results, is not one on which transfer can be based: there was no evidence of non-local transfer. In other words a letter-string in which the letters of the trigram part are interspersed with novel letters is perceived no more accurately than a letter-string made up completely of new letters. This suggests that relationships among neighbouring letters form the basis of the description of the learned letter-string. For the letter-string 123456 the part 123 has the following neighbourly relations, looking only at relations between individual letters for simplicity, and representing a space as "\*": \*-1; 1-2; 2-3; 3-4. In the local test condition, 123ddd, three of these four relationships are maintained. For the same letter string the part 1-3-5 has the following neighbourly relations: \*-1; 1-2; 2-3; 3-4; 4-5; 5-6. But in the non-local test condition 1d3d5d only the first out of the six relationships is maintained, the rest are changed. Thus a neighbouring-items relational description does not provide any basis for transfer of

processing to a non-local part, but can explain the perceptual transfer for local parts.

The attribution of the transfer results to graphemic representations is subject to the same arguments as presented in the general discussion of Experiment 1. The same counter-arguments also apply. The effect of left-right response biases, for example, does not explain away the transfer. From Table 3.1 it can be seen that the transfer in the prime condition is considerably greater over the second three letters (1.4) than over the first three letters (0.6). This was also the pattern of results found in Experiment 1. However the transfer for the local part is roughly the same over both halves of the test strings (0.5 and 0.4). These results would be difficult to explain if performance was completely dependent on left-right response bias effects.



### 3. 3 Experiment 3: Part-part learning

Experiment 2 shows that the representation that allows transfer involves an encoding of the local relationships between letters. It should, therefore, be easier to learn parts that are defined in terms of local rather than non-local relationships. Experiment 3 tests this prediction, using a modified prototype extraction procedure. The prototype extraction paradigm introduced by Posner & Keele (1968) has already been applied to the learning of nonword letter-strings (Whittlesea, 1987) and connectionist models of the sort described in the introduction are capable of simulating much of the data derived from human subjects using this paradigm (Knapp & Anderson, 1984). A number of these experiments have shown that when subjects are presented with a series of patterns (exemplars) derived from underlying prototypes, they respond to the subsequent presentation of the actual prototypes as if they had seen them during the learning trials, and classify them more accurately than some of the exemplars, including ones actually seen during the learning trials (e.g., Franks & Bransford, 1971; Posner & Keele, 1968; 1970).

Whittlesea's (1987) experiments looked at a range of exemplars derived from five-letter prototypes, but did not look specifically at differences in performance when the exemplars were locally rather than non-locally related to the prototype. For this experiment, two underlying six-letter prototypes were constructed for each subject, one for the local and one for the non-local conditions. Each prototype was made up of two parts, each of three letters. Subjects worked their way through a series of trials, each of which contained one or the other parts from one of the underlying prototypes, the other three letters randomly varying. In the local condition the two parts that were repeatedly shown were made up of three neighbouring letters from the underlying prototype. In the non-local condition each part was made up of three non-neighbouring letters from the underlying non-local prototype. Thus each condition had 50% consistency with its underlying prototype throughout the trials, but one maintained simple neighbourhood relations where the other violated them. The experimental hypothesis is that subjects are better at extracting underlying regularities when those regularities are locally rather than non-locally defined, and thus that they learn parts

from the locally defined prototype faster and better than from the non-locally defined prototype.

### 3. 3. 1 Method

*Stimuli and Design:* All stimuli were randomly generated letter-strings made up of six letters selected from the whole alphabet, with replacement. They were identical to those described in Experiment 1 (see Appendix C for examples). A within-subjects design was used with repeated measures and four conditions. The conditions were as follows:

1. **Local**, first three letters: 123ddd.
2. **Local**, last three letters: ddd456.
3. **Nonlocal** first, third, and fifth letters: 1d3d5d (the underlining makes clear that the letters in the non-local parts were different from those used in the local parts).
4. **Nonlocal**, second, fourth, and sixth letters: d2d4d6.

Each subject was presented with 64 trials, 16 of each of the four conditions. The 64 trials were divided into four blocks of 16 trials, each block made up of four trials of each condition. These blocks were themselves broken down into four sub-blocks, each with one trial of each condition. Within each sub-block the order of the conditions was randomised. The independent variable was the local vs non-local nature of the parts of the stimulus, and the dependent variable the number of letters correctly reported and in the correct positions. The underlying prototypes were different for the local and the non-local conditions, so that each subject learned two prototypes altogether.

*Subjects and Apparatus:* Ten members of Stirling University, six males and four females; six were undergraduates, four were staff members. Ages ranged from 17 to 50. All were voluntary participants in the study and all had normal or corrected to normal vision. An Apple II microcomputer was used, as before.

*Procedure:* Each subject was seated in an experimental cubicle in front of the computer.

Instructions were displayed on the screen. The rate of stimulus presentations was controlled by the subjects, using the space bar to proceed to the next trial when they were ready. Each stimulus was presented for 100 ms and followed immediately by a pattern mask made up out of a row of six "eights" for another 100 ms. After the mask left the screen subjects were asked to type in what they had seen. After pressing the return key they were then able to view the next stimulus by pressing the space bar.

### 3. 3. 2 Results

The data were scored on two measures: (1) the number of letters correctly reported in their correct position over the whole letter-string; (2) the number of letters correctly reported and in their correct positions over the focal trigram. The results are summarised in Table 3. 2, as mean number of letters correct per letter-string for each condition.

Condition	Block			
	1	2	3	4
<b>a) All Six Letters</b>				
Local, first	1.9	2.2	2.3	2.5
Local, second	1.8	2.2	2.3	2.4
Nonlocal, first	1.7	1.9	1.9	2.0
Nonlocal, second	1.8	2.2	2.0	2.1
<b>b) Focal trigram</b>				
Local, first	1.6	1.9	2.0	2.1
Local, second	0.5	0.8	0.6	0.7
Nonlocal, first	1.0	1.2	1.2	1.2
Nonlocal, second	0.8	0.9	1.0	1.0

*Note: The scores in the first four rows are out of six; the other scores are out of three. Focal trigram refers to each condition's three letters kept constant throughout the learning trials; in the Local-part conditions the two focal trigrams are the first three letters and the last three letters; in the Nonlocal conditions the two focal trigrams are the 1st, 3rd, 5th letters, and the 2nd, 4th, 6th letters. Chance is 0.2 for the six-letter scores, and 0.1 for the three-letter scores.*

**Table 3. 2.** Mean Number of Letters in Position Correctly Reported, as a Function of Learning Condition, Block, and Scoring Method: Experiment 3.

Performance on Block 4 is better than performance on Block 1. This suggests that learning has taken place. This is confirmed by  $t$ -tests of the difference in performance between Block 1 and Block 4. Analysed over all six letters, the difference between the local condition (the two local sub-conditions added together and averaged) on Block 1 and the local condition on Block 4 is significant,  $t(9) = 3.31, p < .01$ , on a one-tail  $t$ -test. The difference between the non-local condition (the two non-local sub-conditions added together and averaged) on Block 1 and on Block 4, however, is not significant,  $t(9) = 1.64, p = 0.07$ , suggesting that subjects did not learn the non-local part. The focal trigram analysis confirms these findings: the local condition shows a significant improvement in performance on Block 4 compared to Block 1,  $t(9) = 2.32, p < .01$ ; the same comparison for the non-local condition is not significant,  $t(9) = 1.58, p = .075$ .

The second analysis looks at differences between the local and non-local conditions, summing over both the sub-conditions. Across all four blocks this difference is significant for both the six letter and the focal trigram scores,  $t(9) = 1.99, p < .05$ , and  $t(9) = 1.93, p < .05$ . Across blocks 2 to 4 both analyses are again significant,  $t(9) = 2.34, p < .05$ , and  $t(9) = 2.33, p < .05$ .

The third analysis looks at the difference between the two sub-conditions of the two main conditions. Separate  $t$ -tests were performed for both accuracy measures, summing over all four blocks. Within both local and non-local conditions no significant differences were found on the first accuracy measure (all six letters), but on the second measure (the focal trigram) both local and non-local sub-conditions were significantly different from each other,  $t(9) = 17.07, p < .01$ , and  $t(9) = 24.49, p < .01$  respectively.

### 3. 3. 3 Discussion

The results show three things: (1) that learning of the local part but not the non-local part took place; (2) that performance on the local conditions is significantly better than on the non-local conditions; (3) that on the focal trigram scores there is a difference between the two sub-conditions of both the local and non-local main conditions. Each result will be

discussed in turn

The first result confirms the usefulness of this pseudo-prototype extraction paradigm and replicates the finding of Experiment 2 that under some conditions subjects do not learn a non-locally defined part in such a way as to support any transfer of that learning. The implications for theories of relational encoding of Experiment 2 are now reinforced: these results show that when a series of letter-strings have an underlying regularity in some of their constituent items, subjects have difficulty detecting that regularity, unless the items are direct neighbours to one another. This strengthens the conclusion that the internal representation of letter strings must partly consist of a description of the letters relative to each other.

The second finding is the crucial one: however the results are scored subjects do better when the part that is kept constant has its internal constituents in neighbourhood relations than when the constituents are interleaved with randomly varying other letters. The only difference between the two conditions is the relationship of the constituents of the parts to each other, because the same amount of consistency of absolute position is present in both conditions. Part-whole encoding theory is unable to account for this finding.

The third result shows that over all trials subjects responded more accurately to constant letters in the left half of a string. Local condition 123ddd was easier than ddd456, and non-local condition 1d3d5d was easier than d2d4d6, but this was only found for the responses scored over the focal trigram, not over all six letters. This effect of left-right scanning most probably emerges in the output from graphic representations directing the responses from left to right. This difference does not show up when all six letters are scored because, irrespective of the position of the focal trigram, the first three letters are facilitated.

Several improvements could be made to the basic experiment. The learning that took place is not very impressive. Even after 16 trials performance is only at 50% accuracy. If subjects had more time to learn the stimuli, or if they were easier to learn (for example if they were easily pronounceable), then performance would improve. Although the experiment is in the prototype extraction paradigm, the underlying prototypes were themselves never presented or tested. The next experiment will do so.

### 3. 4 Experiment 4: Part-whole learning in local and nonlocal parts

Experiment 4 incorporates the two suggested improvements to Experiment 3. Firstly, in an attempt to make sure that subjects' learning improved, half the trials are five seconds duration. These are called the learning trials, and were blocked, as were the remainder, the test trials. Test trials were each 100 ms in duration. Blocks of learning and of test trials were alternated. Subjects were expected to perform at 100% accuracy on the learning trials. Secondly, although this was also expected to contribute to improved learning, during the blocks of fast trials subjects were tested on the two underlying prototypes, randomly mixed in with the other four test conditions. It is also possible, however, that seeing the prototype particularly boosts part-whole learning and thus reduces the difference between the two conditions that was found in Experiment 3.

The reason for the former modification is a pilot study that failed. This study, with 15 subjects, included both prototypes as test stimuli four times each, but with only 36 other trials in which learning could take place, each of which was a 100 ms masked presentation. The results were negative: performance was no better on the last trials than on the first. No learning had taken place, and so no local vs. non-local difference was possible.

#### 3. 4. 1 Method

*Stimuli and Design:* The stimuli were identical in type to those used in Experiment 3 (see Appendix D for examples). A within-subjects design was used with repeated measures and six experimental conditions. The conditions were as follows :

1. **Local**, first three letters: 123ddd.
2. **Local**, last three letters: ddd456.
3. **Nonlocal**, first, third, and fifth letters: 1d3d5d.
4. **Nonlocal**, second, fourth, and sixth letters: d2d4d6.

5. Local prototype: 123456.

6. Nonlocal prototype: 123456.

Each subject learned two prototypes, one locally defined, one non-locally. They did this by cycling through conditions 1- 4 eight times, i.e., each experimental period was made up of eight blocks, each one of which contained, in random order, the four non-prototype conditions. Alternate blocks were designated as "slow" and "fast"; during the slow blocks each stimulus was presented without a mask for 5 seconds; during the fast blocks each stimulus was presented for 100 ms, followed immediately by a pattern mask for 100 ms. The slow blocks can be conceived of as learning trials, alternating with blocks of test trials. Additionally, during each of the test blocks both of the underlying prototypes were presented as additional test conditions (conditions 5 and 6 above). In total each subject was presented with 40 trials, 16 slow, and 24 fast.

*Subjects and Apparatus:* Twelve psychology undergraduates at Stirling University, 5 males and 7 females, participating as part of a course requirement. Ages ranged from 17 to 36. All had normal or corrected to normal vision. The stimuli were presented on an Apple II microcomputer as before.

*Procedure:* Each subject was seated in a darkened and sound-proofed cubicle in front of the computer. Instructions were displayed on the screen. Subjects were warned when the fast trials were about to begin and again when the slow trials began. In all other respects the procedure was identical to that of Experiment 3, described above.

### 3. 4. 2 Results

Results are presented in the same way as those of Experiment 3. Table 3. 3 shows the mean scores per letter-string for each condition. By comparing performance on block 1 and block 4, the amount of learning can be analysed. The difference for the local condition (the two local sub-conditions added together and averaged) between block 1 and block 4,

analysed over all six letters is significant,  $t(11) = 1.77, p < .05$  on a one-tail  $t$ -test. For the non-local condition (again the two sub-conditions added together and averaged) this difference is not significant,  $t(11) = 0.37, p = .36$ .

Condition	Block			
	1	2	3	4
a) All Six Letters				
Local, first	2.8	2.9	3.0	3.1
Local, second	2.2	2.8	2.8	2.6
Non-local, first	2.4	2.8	2.3	2.3
Non-local, second	2.6	2.3	2.5	2.8
Local prototype	3.1	3.5	4.0	3.7
Nonlocal prototype	2.5	2.5	3.0	3.1
b) Focal trigram				
Local, first	2.5	2.4	2.8	2.7
Local, second	0.6	1.1	0.9	1.0
Non-local, first	1.6	1.7	1.3	1.5
Non-local, second	1.0	1.0	1.3	1.3

*Note: The scores in the first six rows are out of six; the other scores are out of three. Focal trigram refers to each condition's three letters kept constant throughout the learning trials; in the Local-part conditions the two focal trigrams are the first three letters and the last three letters; in the Nonlocal conditions the two focal trigrams are the 1st, 3rd, 5th letters, and the 2nd, 4th, 6th letters. Chance is 0.2 for the six-letter scores, and 0.1 for the three-letter scores.*

**Table 3. 3.** Mean Number of Letters in Position Correctly Reported as a Function of Learning Condition, Block, and Scoring Method: Experiment 4.

The main analysis looks at the differences between the two main conditions, local, and non-local, both over part test trials and over prototype test trials. The two sub-conditions of each main condition are summed together for this analysis. The difference between the local and non-local conditions over all four blocks added together is significant when analysed over all six letters and over the focal trigram,  $t(11) = 2.21, p < .05$ , and  $t(11) = 3.6, p < .01$  respectively. Again over all four blocks, the difference between the local and non-local



prototypes is also significant,  $t(11) = 2.53, p < .01$ .

Over blocks 2, 3, and 4 added together the difference between the local and non-local conditions is highly significant for both scoring measures,  $t(11) = 2.73, p < .01$ , and  $t(11) = 4.38, p < .01$  respectively. The difference between the local and non-local prototypes is also significant,  $t(11) = 2.51, p < .05$ .

This analysis can be continued by looking at the difference between the local and the non-local conditions for each block. Over all six letters the conditions are significantly different on blocks 2 and 3 only: block 2,  $t(11) = 1.86, p < .05$ , and block 3,  $t(11) = 2.01, p < .05$ . Over the focal trigram they are significantly different in blocks 2, 3, and 4: block 2,  $t(11) = 3.34, p < .01$ ; block 3,  $t(11) = 2.64, p < .05$ ; and block 4,  $t(11) = 1.88, p < .05$ . The prototypes are significantly different on blocks 2 and 3 only: block 2,  $t(11) = 2.67, p < .05$ ; and block 3,  $t(11) = 2.07, p < .05$ .

Comparison of the local prototype with the local conditions on block 4 shows that performance on the prototype is significantly better:  $t(11) = 2.06, p < .05$ ; performance on the non-local prototype is also significantly better than on the non-local conditions:  $t(11) = 2.27, p < .05$ .

The final analysis looks at whether there are any differences between the two sub-conditions of the main conditions in each block. Looked at over all six letters, the difference between the two local sub-conditions is only significant in block 1,  $t(11) = 1.94, p < .05$ ; and the difference between the two non-local sub-conditions is only significant in block 2,  $t(11) = 2.03, p < .05$ . Looked at over the focal trigram, most of the differences are significant: the local sub-conditions are significantly different in each block: block 1,  $t(11) = 12.12, p < .001$ ; block 2,  $t(11) = 4.88, p < .001$ ; block 3,  $t(11) = 5.32, p < .001$ ; and block 4,  $t(11) = 4.84, p < .001$ . The non-local sub-conditions are different in blocks 1 and 2 only: block 1,  $t(11) = 3.8, p < .001$ ; and block 2,  $t(11) = 3.95, p < .001$ . This pattern of results is very similar to that found in Experiment 3, with many differences on the focal trigram scores, but very few over all six letters.

### 3. 4. 3 Discussion

The results of this experiment can be summarised as follows: taken over all four blocks of trials a highly significant difference between the local and non-local conditions emerges. Superior performance on the local parts is also found when just the last three blocks are taken together, and when blocks 2 and 3 are analysed on their own (block 2 also shows this difference on the focal trigram analysis). In other words, the difference is most readily found in blocks 2 and 3, but by block 4 is becoming less apparent. This may be because a ceiling is being reached, or because learning of the local condition reaches a temporary plateau, or because of a sudden increase in the learning of the non-local parts. There was, however, no significant improvement in performance on block 4 over block 1 for the non-local condition, which again suggests that no non-local learning took place, although it is possible that comparison of the first and last presentations would show a difference hidden by averaging within each block. Because each block includes four trials of each condition, learning could have taken place within the first block itself. Nevertheless, there was a significant increase for the local condition, suggesting that learning of the local parts continues throughout the trials.

Exactly the same pattern is found for performance on the underlying prototypes, although at a higher level of accuracy than on the separate parts. This shows that although the prototypes were never seen as wholes during the learning trials, performance on them is more accurate than on the sub-conditions of which they are composed. Within the two conditions very few differences are found when scored over all six letters, but on the focal trigram measures performance on the left-hand part is consistently superior to performance on the right-hand part. This pattern of results replicates that found by Experiment 3.

These results are as predicted, and serve to replicate and extend (to performance on actual prototypes) the findings of Experiment 3. Comparison with Experiment 3 shows that this paradigm, alternating fast and slow trials, produces more learning: in block 4 performance reaches a maximum 52% accuracy over all six letters, compared with 41% in Experiment 3, and 89% over the focal trigram, compared with 68% in Experiment 3.

### 3. 5 Experiment 5: Moved-part transfer in local and non-local parts

This experiment has two main purposes: firstly to replicate the finding in Experiment 1 of significant moved-part transfer, but with a different learning method. Instead of the passive reading and internal recital of the letter-string to be learned, a more interactive procedure is used. In this procedure subjects look at the string for 10 seconds at a time, then type it in, and move on to the next 10 second presentation, of which there are six in all. A minute of this more interactive learning should be the equivalent of a considerably longer period of passive learning. Secondly, this experiment seeks to examine any differences between performance on locally and non-locally defined parts. In particular it tests the prediction that no non-local moved-part transfer will be found. Part-whole theories of encoding predict no moved-part transfer, and part-part encoding theories predict no, or little, encoding of non-locally defined parts, so the strong prediction can be made that there will be no non-local moved-part transfer.

There is a problem designing non-local moved-part stimuli: in the locally defined part, for example 123ddd, each letter moves three places in the moved-part condition, ddd123, but because the non-local part of the letter string is, for example, the first, third, and fifth letters, it is impossible for each of them to move three places. In the moved-part condition used in this experiment, each letter moves instead only one place; i.e., 1d3d5d becomes d1d3d5. This might make the non-local moved-part condition easier than it otherwise would be, and thus make any difference that is found in favour of local moved-part transfer even more compelling. Theoretically, the different number of spaces moved in each condition makes no difference: even a small amount of movement by the part produces a new set of relationships, both to the whole and to the other letters in the string, and it is the relationships that are at issue.

### 3. 5. 1 Method

**Stimuli and Design:** All stimuli were strings of six letters selected pseudo-randomly without replacement, constrained to be difficult to pronounce (see Appendix E for examples). A within-subjects, repeated measures design was used, with five experimental conditions. The independent variable was whether the parts that were moved were locally or non-locally defined in the learned word; and the dependent variable was accuracy of response, scored over all six letters and just over the focal trigram.

**Subjects and Apparatus:** Thirteen psychology undergraduates, six males and seven females; all took part as a course requirement. Ages ranged from 17 to 28. All had normal or corrected to normal vision. An Apple II microcomputer was used, as before.

**Procedure:** The procedure differed from that of Experiment 2 in only one way: instead of subjects learning the base words by looking at them for three minutes, a more interactive method was used, partly to prevent subject boredom and partly as a comparison technique. In this method the base words were presented six times for 10 seconds each time and the subject asked to type it in after every presentation. After these six presentations the test trials began; there were five test conditions, as follows:

1. Control: dddddd
2. Local moved-part, first half: ddd123
3. Local moved-part, second half: 456ddd
4. Non-local moved-part, 1st, 3rd, and 5th letters: d1d3d5
5. Non-local moved-part, 2nd, 4th, and 6th letters: 2d4d6d

Each subject was presented in all with eight base words, each of which was followed once by the five test conditions, a total of 40 test trials per subject.

### 3. 5. 2 Results

The results for each of the moved-part test conditions are summarised in Table 3. 4. as accuracy means per letter string for each condition, over all six letters and, where relevant, over the focal trigram.

Condition	Mean	
	All six letters	Focal trigram
Local, first	1.9 (0.9)	0.5
Local, second	1.8 (0.7)	1.5
Non-local, first	1.2 (0.7)	0.4
Non-local, second	1.3 (0.8)	0.8
Control	1.5 (1.1)	-

*Note: Standard deviations in brackets. Focal trigram refers to the three letters in the learned letter string that are maintained in the test conditions; in the Local-part conditions the two focal trigrams are the first three letters and the last three letters; in the Nonlocal conditions the two focal trigrams are the 1st, 3rd, 5th letters, and the 2nd, 4th, 6th letters. A "-" sign means that the scores for these cells were not analysed. Chance is 0.2 for the six-letter scores, and 0.1 for the three-letter scores.*

**Table 3. 4.** Mean Number of Letters in Position Correctly Reported as a Function of Learning Condition, and Scoring Method: Experiment 5.

The two local conditions added together and averaged are highly significantly different from the two non-local conditions added together and averaged,  $t(12) = 5.7, p < .01$ ; additionally the two local conditions taken together are significantly different from the control condition,  $t(12) = 1.92, p < .05$ , whereas the two non-local conditions taken together are not. Of the conditions compared individually, and scored over all six letters, to the control condition, only the first-half local condition shows any significant difference,  $t(12) = 2.22, p < .05$ . There are no differences between the two local conditions, or between the two non-local conditions. Overall performance is lower for the non-local conditions,

though not significantly so, than on the control condition. In other words non-local moved parts show no sign of being learned, because they are not recognised any more accurately than unlearned stimuli.

Looking at the scores taken over the focal trigram confirms the highly significant difference between the two local conditions taken together from the two non-local conditions taken together,  $t(12) = 4.41, p < .01$ . On this scoring method the two local sub-conditions are highly significantly different from one another,  $t(12) = 6.55, p < .001$ , as are the two non-local sub-conditions,  $t(12) = 4.44, p < .001$ .

### 3. 5. 3 Discussion

The main finding of Experiment 5 is that there is evidence of significant local moved-part transfer, and no evidence of any non-local moved-part transfer. Indeed, overall performance in the non-local condition is lower, though not significantly so, than on the control condition. In other words non-locally defined parts show no sign of being learned, because they are not recognised any more accurately than unlearned stimuli. The significance of this finding is as follows: the achievement of local moved-part transfer argues against the theory that parts are encoded in relation to wholes; but the lack of non-local moved-part transfer puts a severe restriction on the generality of the alternative theory, part-part relational encoding, even suggesting that only neighbouring items at one level of complexity (letters in this case) are encoded in terms of one another.

Both local and non-local conditions have the same absolute amount of similarity with the learned word: in each case three out of the six letters were maintained. The differences in performance in Experiment 5 confirms the inference from Experiment 2, that the internal description of the learned letter string is not based on this information. The difference in performance found between these conditions rules out theories which give equal importance to each part, treating each part as an independent entity in a particular relationship to the whole of which it forms a part. The relationships among parts need to be taken into account to explain this result.

One purpose of this experiment was to replicate the significant moved-part transfer of Experiment 1. Table 3. 5 compares the two experiments on performance on the local moved-part (averaged over the two moved-part sub-conditions) and control conditions, and on the amount of transfer obtained.

Condition	Experiment 1		Experiment 5	
	All six letters	Focal trigram	All six letters	Focal trigram
Moved-part	2.9	1.7	1.9	1.0
Control	2.2	1.1	1.5	0.8
Transfer	0.7	0.6	0.4	0.2

*Note: the focal trigram control scores are obtained by dividing the control score by two.*

**Table 3. 5.** Comparison of amount of transfer obtained between Experiments 1 and 5.

It is clear that performance in Experiment 5 is lower overall than in Experiment 1. For the moved-part condition alone this could be attributed to reduced learning time, but because performance on the control is also lower, it may be because the stimuli were made up of consonants only, rather than including some vowels. On the argument that the stimuli in Experiment 5 were less pronounceable than those in Experiment 1, the small difference in transfer between the two experiments suggests that phonological recoding has, at most, a small influence on the transfer produced.

The critical comparison for the question, discussed in §3. 1, of possible differences between letter-interrelationships and letter-cluster relationships, is the comparison across experiments, of the difference between fixed and moved-part conditions, for local and non-local parts. The data for this comparison are presented in Table 3. 6. The prediction was that local conditions would show more disruption than non-local conditions from the changed part-whole relationships in the moved-part conditions, *if* it was the case that local letter-clusters were easier to learn and so more likely to have their part-whole relationships represented. Assessing this prediction is complicated by the fact that comparisons have to be across experiments, but the following figures can be offered. In the local conditions, the

average amount of disruption (the fixed-part transfer scores minus the moved-part transfer scores) for the data from Experiment 2 and Experiment 5 is 0.15. For Experiment 1, the figure is 0.55. The comparison disruption figure for the non-local conditions, from Experiments 2 and 5, is 0.5. Whichever figure for the local conditions is used, the prediction is disconfirmed: disruption is not larger for the local than the non-local conditions. Indeed, on the figures derived from Experiments 2 and 5 alone, the disruption is larger for the non-local conditions.

Condition	Local Score	Local Transfer	Non-local score	Non-local Transfer
<b>(Experiment 2)</b>				
Fixed-part: all six	<b>2.6</b> (3.2)	<b>0.5</b> (1.0)	<b>2.4</b>	<b>0.3</b>
Control: all six	<b>2.1</b> (2.2)		<b>2.1</b>	
Fixed-part: focal trigram	<b>1.5</b> (2.5)	<b>0.4</b> (1.4)	<b>1.4</b>	<b>0.3</b>
Control: focal trigram	<b>1.1</b> (1.1)		<b>1.1</b>	
<b>(Experiment 5)</b>				
Moved-part: all six	<i>1.9</i> (2.9)	<i>0.4</i> (0.7)	<i>1.3</i>	<i>-0.2</i>
Control: all six	<i>1.5</i> (2.2)		<i>1.5</i>	
Moved-part:focal trigram	<i>1.0</i> (1.7)	<i>0.2</i> (0.6)	<i>0.6</i>	<i>-0.2</i>
Control: focal trigram	<i>0.8</i> (1.1)		<i>0.8</i>	

*Notes: The scores written in Bold are from Experiment 2; the scores in brackets are from Experiment 1; the scores in italics are from Experiment 5. The focal trigram control scores are obtained by dividing the control score by two.*

**Table 3. 6.** Comparison of amount of local vs. non-local transfer obtained in fixed and moved-part conditions: Experiments 1, 2, and 5.

This result indicates that there is no evidence that learned letter-clusters are represented in terms of their relationship to the whole. What this result says nothing about is whether learning a letter-cluster increases the likelihood, or amount of encoding of the relationships between *its* letter-parts and itself as a whole. One reason to suspect that it might is that letter-clusters are more likely to be familiar, or at least less likely to be completely novel, than six letter nonwords, and familiar items may represent relational information rather differently from unfamiliar items. The latter claim is examined in Chapter 4.



# MIGRATION EXPERIMENTS

---

## 4. 1 Introduction

As discussed in Chapter 1, migration errors are often taken as evidence of part-whole coding because of their reported tendency to maintain position. This evidence is not strong because of the constraints in the stimuli used in most migration experiments, but the interesting possibility derived from the results of the preceding five experiments is that positional encoding differs for words and nonwords. Part-whole encoding may be used for words but not nonwords; thus migrations between words might maintain position while migrations between nonwords are not position-specific. Three experiments are reported that analyse the positional specificity of migrations in words and nonwords. Experiment 6 looks at within-string migrations in reports of briefly presented letter-strings. Experiment 7 looks at across-string migrations between two briefly presented letter-strings. In Experiment 8 migrations of target letters into word-like letter-strings were either encouraged or discouraged by preceding source letter-strings.

## 4. 2 Experiment 6: Positional information

This experiment is designed to explore some of the parameters of positional information by examining the positional errors made in reporting briefly presented nonword letter-strings. Positional errors in report are loosely analogous to migration errors. The difference

is that classical migration errors are transpositions from one letter string to another. Positional errors in report of single letter-strings can be seen as within-string migrations. Correct report of the position of an identified letter represents in these terms a migration of zero positions, or a fixed-position migration. This is an unnatural way to consider identification accuracy until the full report task is thought of as requiring the correct combination of two sources of information, letter identity and letter position.

Previous work on mislocation errors (S1. 4. 3) suggests that "information ... concerning the location of displayed letter can be described by an uncertainty gradient around the true location" (Estes, Allmeyer, & Reder, 1976). This gradient describes a measure of *positional uncertainty*; this gradient is known to have the properties that uncertainty, or gradient variance, increases with eccentricity from fixation point, and that transposition responses drift from the periphery towards the centre, an overall skew in the gradient towards fixation point (Estes et al., 1976). This experiment attempts to derive uncertainty gradients for each position in a briefly presented pattern-masked letter-string.

## 4. 2. 1 Method

**Stimuli and Design:** Stimuli were strings of seven letters selected randomly without replacement; none were easily pronounceable. No experimental conditions were drawn up. All subjects saw the same letter-strings but order of presentation was randomised.

**Subjects and Apparatus:** Five psychology undergraduates at Stirling University, three male and two female; aged between 20 and 30, with normal or corrected to normal vision. Participation fulfilled a course requirement. The stimuli were created and presented, and the results analysed on an Archimedes 310 microcomputer.

**Procedure:** Subjects were seated in a sound-proofed experimental cubicle, illuminated by a reading lamp providing enough light to read the keyboard and screen of the computer. Instructions presented on the screen informed subjects that they would be briefly shown strings of seven letters and asked after each one to type in what they had seen. Before

presentation of the test trials subjects were given 10 practice trials each. Testing took place on two separate occasions for each subject, 50 test trials on each occasion.

Each trial was preceded by a fixation point appearing between two flanking arrows in an otherwise blank screen. Each letter-string was presented for 100 ms, followed immediately by a mask (a string of seven "X"s) for another 100 ms. Subjects were encouraged to respond as quickly as possible but slow responses were not penalised. Because the stimulus string was made up of seven letters without any repetitions, it was made impossible for subjects to type in the same letter more than once. Subjects were forced to type in exactly seven letters by making it impossible for them to proceed to the next trial without the full number. Deletions and corrections could be made at any time while responses were still on the screen. Subjects proceeded through the trials at their own speed by pressing the Spacebar. No feedback was given during the experimental sessions, which took twenty to thirty minutes. The second session was followed by a debriefing, in which the purpose of the experiment was explained and any questions answered.

## 4. 2. 2 Results and Discussion

Because this was essentially an exploratory experiment the results are presented descriptively. They are discussed in three sections: (i) differences between identity-alone, or Item information (identity correct and position incorrect), and identity-and-position, or Position information (identity *and* position correct)<sup>1</sup>; (ii) the spread of positional errors (uncertainty gradients) for each letter position; (iii) analysis of the results in terms of single letters and clusters of letters, providing measures of Order information, defined here as the probabilities of getting letter-sequences correct, both in and out of absolute position.

---

<sup>1</sup> The terms Item and Order information are the traditional ones (see e.g., Sperling & Melchner, 1976) but here Order information is restricted to inter-letter positional information.

(i) *Item vs. Position information.* The overall results<sup>2</sup>, summarised in Table 4.1, show: (a) the standard left-right effect in report of letter-strings, together with the end-letter effect whereby both end letters are more accurately identified; and (b) that Item information can be available when Position information is absent.

Letter Position	1	2	3	4	5	6	7	Mean
Item	95.4	72.0	59.0	44.6	31.8	32.2	36.0	53.0
Position	94.2	56.6	35.8	13.4	8.4	7.2	13.8	32.8
Difference	1.2	15.4	23.2	31.2	23.4	25.0	22.2	20.2

**Table 4. 1.** Percentage correct for each letter: Position (identity plus position) and Item (identity irrespective of position), and the difference between the two scores.

The left-right and the end-letter effects are apparent both in the Item results, and the Position results for which they are more usually reported (e.g., Merikle & Coltheart, 1972). The left-right effect is much stronger in the Position scores, and the end-letter effect correspondingly weaker. Table 4. 1 (bottom row) shows that over all seven letters there is a mean 20% difference between Item and Position scores. The strong inference from this result is that letters in nonword letter-strings are not simply represented by the *conjunction* of their identity and their position, contrary to part-whole theory. To confirm this inference, 1147 (32.8%) of all responses are correct in Position, and *another* 698 (19.9%) are correct in Item. The remaining 1655 (47%) are completely incorrect.

What are the variables that affect the probability of correctly reporting a letter identity? Using the end-letter scores as anchors and rounding to the nearest 5%, the data can be fit by three simplifying assumptions: (i) the left-right effect produces 10% accuracy decrements for each successive position to the right; (ii) each letter position other than the end-letters receives a 15% decrement, a *mid-string confusability effect*; (iii) the central letter (position

<sup>2</sup> Because subjects could not repeat letters it is probable that these scores, and those for the next two experiments, are underestimates: if subjects typed in the correct letter for position 4 as their response to position 3, the effective probability of them getting position 4 correct is zero. Nevertheless, the results are qualitatively similar to those of other experiments.

4) receives an additional 5% decrement, while its right-hand neighbour receives an additional 10% decrement, a *central letter confusability effect*. The overall pattern of scores after these three modelling assumptions becomes a good match to the observed results:

1	2	3	4	5	6	7
95	70	60	45	30	30	35

These figures are undoubtedly very dependent on experimental details, but the three postulated effects may be quite general. (i) If the left-right effect indicates a reduction in accuracy that is constant for each successive letter position, this implicates a constant decay rate as they are recoded into graphic output, given the assumption of parallel letter input. Unfortunately it is not possible to tell whether the decay results from the serial recoding of each letter into an output code, or from being held successively longer in a graphic output buffer after parallel recoding into output representations. (ii) The assumption of parallel letter input is supported by the mid-string confusability effect because it indicates the harmful impact of surrounding letters on any one letter (lessened for the end-letters because they have less surrounding letters), implying that processing any one letter is strongly influenced by that letter's neighbours. The interfering effect of neighbouring letters has been shown often (e.g., Bjork & Murray, 1977; Eriksen & Eriksen, 1974; Santee & Egeth, 1980), and is also shown by MIRAGE (Watt, 1988; Watt et al., in press), which implicates VIP representations in this effect. (iii) The central confusability effect may simply be an exaggeration of the mid-string confusability effect for the central, foveated letters<sup>3</sup>.

The difference between the Item and Position scores provides a positional index for each letter position: the proportion of times on which identity information is available when position information is also available. Table 4. 1 shows that this happens more for the central positions than the end-letter positions, and more for the right than the left positions. To remove the overall effects of accuracy for the different positions, the Position scores were transformed into percentages of the Item scores for each position.

---

<sup>3</sup> I am not aware of previous reports of this effect. It appears to implicate relational coding of letter positions because it suggests that letters in the centre of a string are more strongly masked by surrounding letters than letters to the right or left of centre. Whether MIRAGE simulates this effect is not established.

1	2	3	4	5	6	7
98.7	78.6	60.7	30.0	26.4	22.4	38.3

There is a clear-cut left-right effect, modified by the end-letter effect<sup>4</sup>. Modelling the data in the same way as before, the same constant left-right effect of 10% emerges, but the confusability effects are slightly different: the data are fit most neatly by assuming a constant 10% decrement for all mid-string letters, as compared to 15% before. This leaves a central confusability effect for positions 3, 4, 5, and 6 of 10%, 30%, 25%, and 20% respectively. As for the Item scores the central confusability effect is strongest for the central letter and position 4, but extends more widely to all other mid-string positions except position 2.

This result indicates (a) that Item information is more robust than Position information, and (b) that Position information varies non-linearly with position in the string. In §1.4.3 it was suggested that the relative fragility of positional information reflects masking of the VIP, assuming that position information is more usefully represented in the VIP than in the graphemic description. This cannot be the whole case, however: on the argument that the probability of correctly reporting letter identities depends to a significant degree on the relative clarity of VIP descriptions, it might be expected that the positional information index would closely mimic the Item scores if the VIP is the main source of position information. Although the relationship between the two is close, it is not exact; in particular the central letter scores much lower on the positional index than expected from its Item score. Thus it seems necessary to consider interactions between VIP and graphemic descriptions and within graphemic descriptions themselves. In particular the non-linearities in the position index imply that the positions of identified letters are represented in relation to each other. Two ways to consider this implication offer themselves: (a) as suggesting that neighbouring letter identities are temporarily bound to each other, by a mechanism similar to the dynamic links discussed in §1.3.3.2; or (b) as suggesting that letter identities are represented as

---

<sup>4</sup> This implies that an explanation of the results in terms of response constraints is not easily available: there is less positional uncertainty for the 7th letter than the 6th, which could be because there are less positions available for subjects to put their responses in for the 7th than the 6th letter; but there is more positional uncertainty for the 6th than the 5th letter.

parts of letter-clusters, implicitly coding their relative positions<sup>5</sup>. Both deal simply with the effects of position-in-the-string; in the neighbouring letters version the neighbourhood relations are less clearly/strongly specified for less clearly identified items, and in the letter-cluster version the less clearly identified letters are represented by a wider, more ambiguous set of letter-clusters. The difference between the proposals is *not* that only one uses dynamic links, because the letter-cluster version can use dynamic connections to bind letter-clusters, but more that in the former version the crucial representational item is the single letter. This leads to the prediction that if the position of a single letter is incorrectly encoded it is as likely to be completely wrong as it is to be partially wrong. This is examined next.

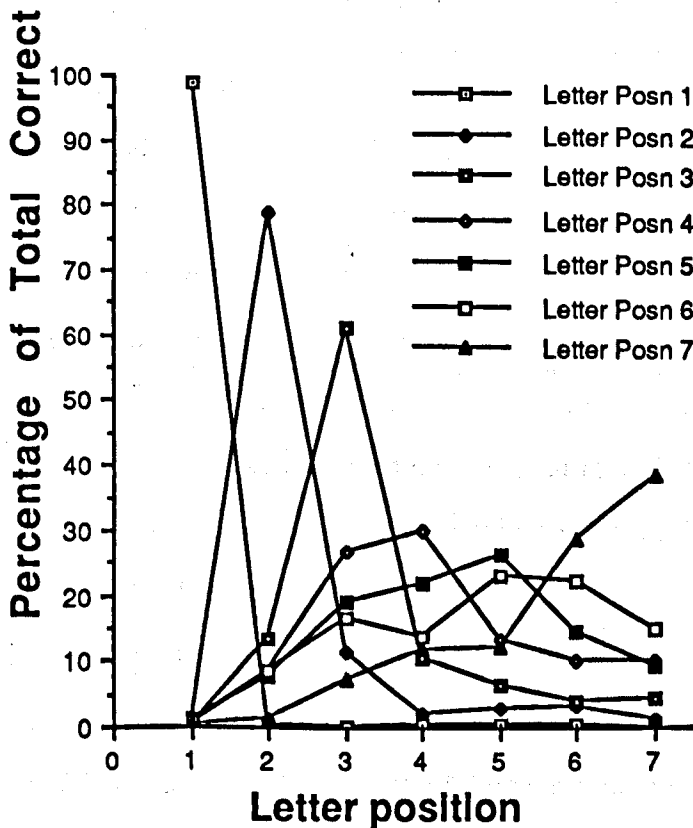
(ii) *Position uncertainty gradients.* For each letter-position, the number of times that that letter's identity was reported in each of the seven letter positions provides a measure of the gradient of positional uncertainty. This is presented in Table 4. 2. where *target position* refers to the position in which a response was made and *source position* to the correct position for that letter. Target position represents positions *to* which migrations are made; source position represents positions *from* which migrations are made.

Source Position	Target Position						
	1	2	3	4	5	6	7
1	471	2	0	2	1	1	0
2	4	283	41	7	10	11	4
3	1	40	179	31	19	12	13
4	1	19	60	67	30	23	23
5	2	12	30	35	42	23	15
6	1	14	27	22	37	36	24
7	1	2	13	21	22	52	69

**Table 4. 2.** Number of correct reports for each letter position in each position. The main diagonal contains the number of correct Position reports; these are also part of the data of Table 4. 1.

<sup>5</sup> Logically there is no reason why these proposals should apply to graphemic representations any more than they do to graphic output representations. On the other hand, it is plausible that both domains use similar representational structures, and possible that the propagation of identity information from graphemic to graphic representations is more robust than the propagation of position information.

These results are shown graphically in Figure 4. 1, where, in order to remove the effects of absolute accuracy for each position, the data-points for *one* position-in-the-string are the number of Item responses *in each* letter position as percentages of the overall Item score for that letter. Each presented letter-position is represented by its own position-response curve; the normalisation makes the area under each curve equal, but does not change the gradient. The curves for each letter-position are superimposed to compare their tuning; each position can be identified by finding its peak response: usually this corresponds to the correct position for that letter.



**Figure 4.1.** Migrations from each letter position to each letter position (including migrations of zero movement) as a percentage of overall correct responses for each letter-position.

The clearest implication of these results is that Item information, as evidenced by reports of correct identity but incorrect position, is not simply correct identity information in the absence of *all* position information. For all letter positions the incorrect position responses show a gradient of error: positional errors are more likely to be a nearby than a distant letter



position. Position information cannot, therefore, be encoded in an all-or-none manner. Figure 4. 1 shows that the gradients differ across letter-positions. Positions 1, 2, and 3 (and 7 to a lesser extent), are tightly tuned: positional errors are extremely likely to be on immediately adjacent letters, and very unlikely to be further away. Positions 4, 5, and 6 show much broader tuning: positional errors are equally frequent over a wide range of positions. Indeed, for Letter-Position 6 responses in position 5 are more frequent than in the correct position. The first three positions show a much more marked discontinuity between correct and incorrect positional responses. Position 7 appears to be intermediate between the two types of curves.

Very similar results have been reported for the report of auditorily presented strings of six letters (Jahnke, Davis, & Bower, 1989), except that the most marked transition is between the intermediate letters, which are all broadly-tuned, and the end-letters, which are more sharply tuned; the final letter shows much sharper tuning with auditory than with visual presentation. Typically the final letter in a serial position curve shows a much stronger *recency effect* than the last-letter effect with visual presentation; this suggests that Item and Position curves are much more similar to each other with auditory presentation. In turn this may be because of the constraints imposed by sequential presentation on the encoding of serial order.

Interpretation of the obtained results is not straightforward: if inter-letter relationships are used to encode position, and if the strength of the representation of an inter-letter relationship depends on the strength of the representations of the inter-related letters, then it is difficult to explain the sharp transition between the gradients of positional uncertainty for positions 3 and 4. (This may, however, be how temporal order is represented). One solution would be to suggest that non-neighbouring letters are sometimes represented as neighbours, thereby producing position errors, but the results of Experiments 2-5 make this doubtful. Some additional complexity is needed; this may be obtained by the idea that letter positions are encoded implicitly in letter-cluster representations. If the letters are encoded as overlapping bigram and trigram clusters, and if each cluster decays equivalently before output, then the transition from position 3 to 4 can be explained as position 3 being in the

bigram pairs 2-3 and 3-4, but also in the trigram 1-2-3 which clearly provides superiority over the longer delayed trigrams of which position 4 is a part. This interpretation suggests that it should be possible to find direct evidence for the role of letter-clusters.

(iii) *Letter groups.* The analysis of between-letter interactions continues by looking at responses for different sizes of letter-group. Figure 4. 2 shows the number of correct Position reports separately for each size of letter-group. The figures are summed over all available letter positions; the different sized-letter groups are scored independently of each other, such that the single-letter correct responses were not included if they were part of a correct bigram group. The number of correct seven-letter groups (three) is the number of times the whole string was correctly reported. The results are presented in absolute numbers and after a logarithmic transform of the data.

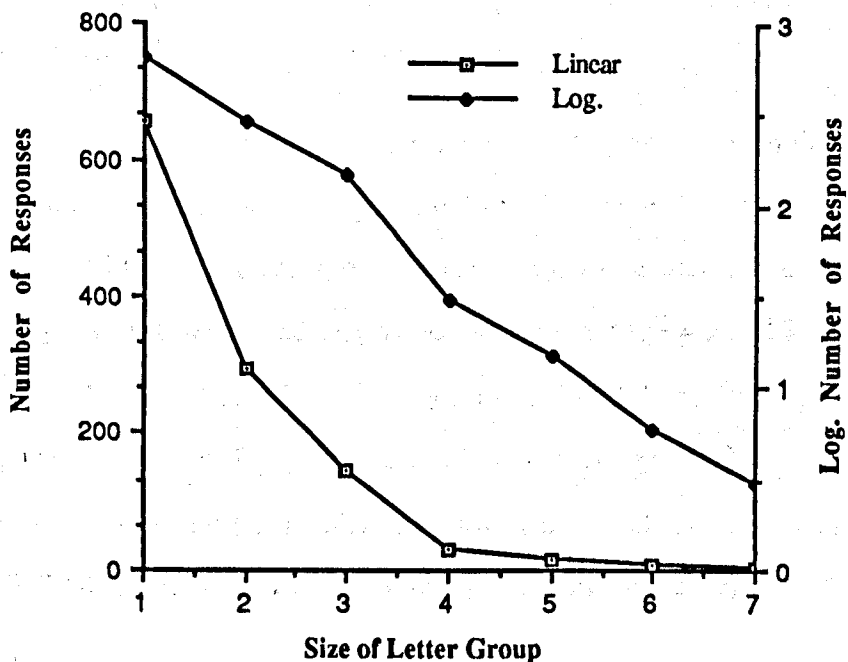


Figure 4. 2. Absolute and logarithmic number of correct reports for letter-groups of all possible different sizes. Absolute scale on the left, logarithmic on the right.

There is a log-linear relationship between number of correct (Position) reports and size of letter-group (correlation coefficient  $R = 0.994$ ). This strongly implies that there is no privileged size of letter-group, at least when the letter-groups are not orthographically

rédundant: neither bigrams (Humphreys et al., 1990) nor trigrams (Mozer, 1987; McClelland, 1986), show signs of being more easily processed. Whether there is evidence that letter-groups are processed as the sum of their single-letter components is not so clear, though the absence of a linear relationship is suggestive.

One approach to this question is to ask for each letter-position whether accuracy is affected by whether the preceding or subsequent letter-position was responded to correctly or incorrectly. The results of this analysis are presented in Table 4.3.


Letter Position	1	2	3	4	5	6	7
Preceding letter correct	-	274	146	38	15	9	10
Preceding letter incorrect	-	9	33	29	27	27	59
Subsequent letter correct	274	146	38	15	9	10	-
Subsequent letter incorrect	197	137	141	52	33	26	-

**Table 4. 3.** Number of correct reports for each letter position as a function of whether the preceding or subsequent letter was either correct or incorrect.

This analysis shows that positions 3 and 4 are not affected by accuracy on the preceding position; accuracy on positions 2 and 3 is significantly higher if the preceding position was correct ( $\chi^2 = 138.1$  and  $71.3$  respectively,  $p < .01$  for both), while accuracy on positions 6 and 7 is significantly lower if the preceding position was correct ( $\chi^2 = 9$  and  $34.8$ ,  $p < .01$ ). Accuracy on position 2 is unaffected by accuracy on position 3, but accuracy on the first position is significantly higher if position 2 is also correct ( $\chi^2 = 12.6$ ,  $p < .01$ ), and accuracy on positions 3, 4, 5, and 6 is significantly lower if the subsequent position is also correct ( $\chi^2 = 59.3$ ,  $20.4$ ,  $13.7$ , and  $7.1$  respectively,  $p < .01$  in each case)<sup>6</sup>.

<sup>6</sup> A cruder way to assess the independent-letters hypothesis is to calculate probabilities for each size of letter-group. If letters contribute independently, the probability for any particular size group should be a multiple of the single-letter probability. This analysis is done on the inclusive scores because exclusive probabilities implicitly take into account for any sized letter-group whether the adjacent letters were also correct and thus fail to reflect independent letter probabilities. The probability of getting any letter correct is 0.32, so the probability of getting a letter-pair correct should be 0.10; the observed probability is 0.164. Although this difference is small it is multiplied for larger letter combinations which suggests that individual letters do not contribute independently to the observed probabilities of correctly reporting letter-groups.

Another approach is to compare position information across the different size letter-groups. Position response curves for each size of letter-group are shown in Figure 4. 3. The results are presented logarithmically to compress them into a smaller range. Each group size shows essentially the same pattern: more frequent responses in the correct position, tailing off over the full range of possible movements by an approximately logarithmic function. The apparent deviations for the single-letter group from the log-linear function are probably because of the greater likelihood of making positional errors that differ by three or four positions by chance: a letter in the centre of the string can only move a maximum of three or four positions either side of centre.



**Figure 4. 3.** Position-specificity for different size letter-groups, presented as the logarithm of the number of reports in positions that differ by increasing amounts from the correct position. Correct position responses are represented as zero positions moved.

It is clear that some bigrams and trigrams move position; if this result is not due to chance, it is strong evidence for the encoding of relationships between the letters of bigrams and trigrams. The observed probability<sup>7</sup> of a single letter moving one position is 0.063; if

<sup>7</sup> As measured by the observed number of times a single letter moves one place divided by the number of times a single letter could have moved one place; for one seven-letter string there are 12 possible ways in which a single letter and a bigram can move one place.

letters are processed independently the probability of two adjacent letters moving one position is 0.004; the observed probability is 0.007, which is marginally higher. Whether this implicates part-whole coding within letter-clusters, or the encoding of letter-interrelationships within a letter-cluster is not clear. Nevertheless Experiment 6 provides further evidence against the idea of part-whole encoding of positional information when the parts are single letters and the wholes are nonword letter-strings.

### 4. 3 Experiment 7: Across-string migrations

This experiment is designed to extend the findings of the previous experiment to a different situation. Instead of within-string migrations, though, Experiment 7 looks at migrations between strings. The question of interest is whether, as found in some circumstances with word stimuli, migrations between nonwords maintain position as they move from source string to target string. As discussed in Chapter 1 migrations between words do apparently maintain position, but the issue has not been examined in detail for nonwords. Pilot studies indicated that sequential presentation (cf. Intraub, 1985; Treisman & Souther, 1986) produced similar number of migrations to the more usual simultaneous presentation, so sequential presentation of two letter-strings was used; the task was to read the first string, the *source*, but respond only to the second, the *target*.

#### 4. 3. 1 Method

**Stimuli and Design:** Letter-strings were seven-letter nonword strings, selected from consonants only, with the constraint that the same letter did not appear in both strings. A within-subjects, repeated measures design was used.

**Subjects and Apparatus:** Eight undergraduates, four male and four female. Other details as for Experiment 6.

**Procedure:** The only difference in procedure from Experiment 6 is that two letter-strings were successively presented, the first for 360 ms, masked for 100 ms, the second for 120 ms, also masked for 100 ms. This presentation sequence has been called *four-field masking* (e.g., Evett & Humphreys, 1981). Stimulus durations were chosen on the basis of pilot work determining optimum presentation times for maximum numbers of migrations. Subjects were told to attend to both strings but to type in only the second string. Each subject received 20 practice trials unless they asked for more, in which case an extra 10 were presented; a total of 50 experimental trials were then presented to each subject.

### 4. 3. 2 Results and Discussion

Position (i.e., identity-*and*-position) scores are summarised in Table 4. 4; the main interest of the analysis is in the error data, which thus comprise over 70% of the total reports.

1	2	3	4	5	6	7	Mean
83	47	28.8	14.3	7.5	8.3	12	28.7

*Note: chance is 4.8%.*

**Table 4. 4.** Accuracy of report for each position, presented as percentage accuracy.

Comparison with the results of Experiment 6 suggests two differences. When the data are analysed in the same simplifying manner as for Experiment 6, using the end-letter scores as anchors, the left-right effect is stronger, most neatly modelled by constant 12% rather than 10% decrements, and the central confusability effect slightly different: as before this is most marked for position 4, but extends more strongly to the left than to the right. Assuming a constant mid-string decrement of 15% again, the additional decrements are 13%, 16%, 19%, and 15% for positions 2, 3, 4, and 5 respectively. This result suggests that the interfering effect of the source is stronger on the left than the right letters of the target, presumably reflecting stronger competition between the more clearly identified letters

on the left of both strings.<sup>8</sup>

Analysing the migration errors in terms of where in the source they come from, supports this interpretation: the central letter, position 4, and position 3 produce most migrations (101), while position 5 only produces 79. There is a strong tendency, however, for migrations to move towards the right of the target: positions 2, 3, and 4 receive an average of 78 migrations, whereas positions 5, 6, and 7 receive an average of 118. Figure 4.4 shows both these effects more clearly.

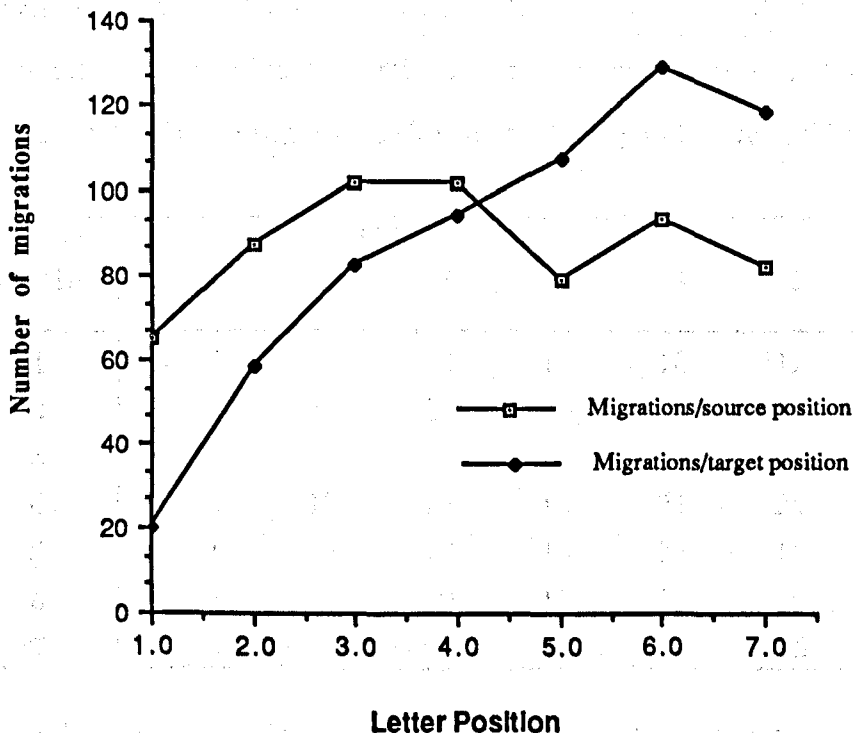


Figure 4.4. Number of migrations from each position in the source string to each position in the target string.

<sup>8</sup> Note that this explanation is ambiguous: it was assumed earlier that the left-right effect depended on recoding into graphic output representations rather than on graphemic representational clarity. This means that the left letters of the source are exerting their interference either (a) because the source was automatically recoded, and preferentially for the left letters, into a representational domain shared by the recoding of the target, presumably the graphic domain, or an intermediate phonological domain, or (b) because the source was recoded, again preferentially for the left letters, into a domain not shared by the recoding of the target but that feeds back activation to a domain shared by the coding of the target. Possibly the source has time to be recoded phonologically, which feeds back activation to the graphemic representations, thereby interfering with the target which does not have time to be recoded phonologically.

The overall tendency of migrations is clearly from the left of the source to the right of the target; the inference from this result is that competition between letters across strings is not position-specific, and thus that letters in nonwords are not represented in part-whole relations.

The position specificity of the migrations is analysed next. Table 4. 5 shows the numbers of migrations from the source, for each position in both source and target strings. Target position represents positions *to* which migrations are made; source position represents positions *from* which migrations are made. The successive diagonals above and below the central diagonal provide the numbers of migrations that move different numbers of positions. The totals of these figures are shown graphically in Figure 4. 5 below.

Source Position	Target Position						
	1	2	3	4	5	6	7
1	7	3	8	10	10	12	15
2	1	15	14	9	20	16	12
3	3	10	24	20	11	21	12
4	2	10	12	18	21	24	14
5	1	5	10	15	23	14	11
6	2	10	9	10	18	26	18
7	4	5	5	12	4	16	36


**Table 4. 5.** Number of single-letter migrations for each target position from each source position. The central diagonal column contains the number of same-position migrations.

The majority of the migration errors do not maintain position: the total number of single-letter migrations is 608. Of these, 149 (25%) are same-position migrations; the remaining 459 (75%) are moved-position migrations. Of all the errors, 7% are same-position migrations, 23% are moved-position migrations. It is possible, however, that this simply reflects a scoring bias: for example, there are only seven ways for a fixed-position migration to be produced, but 12 ways in which a migration of one position can be produced. This is because letter-position 2 in the source can migrate either to position 1 or position 3 in the target and be scored as a migration of one position. To differing degrees the same is true of



other sizes of migration: migrations of six positions, for example, can only be produced in two ways, from position 1 to position 7, or vice versa.

When the migrations are analysed in terms of the number of migrations possible for each particular size of movement, migrations that maintain position are more frequent than migrations of any particular size of movement, but migrations to different positions are still obtained. Figure 4. 5 shows the migrations scored in both ways, as absolute numbers and as percentages of possible migrations.



**Figure 4. 5.** Number of migrations that move each possible distance, presented in absolute terms and as a percentage of the number of possible ways of producing a migration of that particular amount of movement.

Even on the percentage scoring method, a large number of migrations between letter-strings do not maintain position. Table 4. 5 suggests that this may be more true of some source positions than others. For positions 1 to 4 the migrations are overwhelmingly to different positions: only 22% (64/290) maintain position; for positions 5 to 7 exactly 50% (85/169) maintain position. Two different causes of migrations may be at work: in one,

competition between letter identities for the left letters, as suggested earlier, produces migrations which do not respect position; in the other, competition between letters-in-position for the right letters produces migrations which do respect position.

This provides an interesting contrast to Experiment 6, in which the left letters were more tightly bound to position-in-the-string than the right letters. This contrast offers the following possibility: assume that the brief presentation times of Experiment 6 and the fact that the stimulus was an unfamiliar whole prevent the construction of an unambiguous representation of the stimulus as a whole. This means that the representation of position must be done through the relations between letter-wholes. The left-right superiority, however, means that the relations and the identities are more clearly represented for the left letters, which helps explain their tighter position response curves. The sharp difference between the left trigram and the other letter positions suggests that in addition the left trigram has a partial representation as a whole, within which its letters are represented in part-whole relations. The longer processing durations available for the source string in Experiment 7 means that the right trigram now attains the status of a whole, and the part-whole coding of its letters explains the fixed-position migrations from the right trigram; in the left trigram, however, the letter parts have attained clear enough representations to be treated as separate wholes represented in relation to each other; their status as wholes now allows them to escape their part-whole coding and migrate to different positions when the actual letters for those positions are only poorly represented as wholes themselves. Only 39 bigram migrations were obtained, 18 (46%) of these maintained position, but again the tendency to maintain position is stronger for the right letters (13/19) than the left letters (5/20). This interpretation is pursued in the next experiment, which looks at different types of wholes.

# ERRATA

A mistake was made in the calculation of the three Newman-Keuls analyses of the results of Experiment 8. The result of this mistake is that several apparently significant results are actually non-significant. The discussion of the results, in this chapter and in Chapter 5, should, therefore, be treated with caution. The patterns discerned in the results are not changed by the mistake, but the differences between conditions are smaller than claimed.

The results are, after correction, as follows:

(a) Focal letter accuracy scores: the WordPrime condition is significantly better than the Full-Prime condition, and both are significantly better than all the other conditions. No other significant differences.

(b) Accuracy over the other six letters: no significant differences.

(c) Lexicalisation of the focal letter: the Full-Prime condition is significantly worse than the following three conditions: Prime-Fixed Context, Migration-Fixed Context, and WordPrime. No other significant differences.

## 4. 4 Experiment 8: Detecting and correcting spelling mistakes

This experiment is designed to test whether the moved-position priming and migration effects reported in the preceding experiments apply to word-like stimuli as well as to nonword letter-strings. The paradigm used to explore this issue is the same as used in Experiment 6, the successive presentation of two different letter-strings. The second letter-string, the target, is usually a word presented with one letter, the *focal letter*, misspelled, SHBJE<sub>T</sub>, for example. The task of reporting the target is intermediate between letter-as-part and letter-as-whole tasks; it requires full report of the whole string (part-whole processing), but also focusses attention on individual letters within the string (whole-whole processing). Different preceding strings, or sources, might encourage processing the target letters either as wholes or as parts of the whole string. Correctly reporting the misspelling can be taken as evidence of successful letter-as-whole processing, while reporting the misspelled letter as the letter it *should* have been is suggestive of letter-as-part processing.

Preceding the target with a source containing the misspelled focal letter, XHVCFDS, for example, might make it easier to detect the misspelled focal letter as compared to a control string which does not contain the focal letter. Detection of the misspelling might be primed by appropriate sources, so that focal letter accuracy would increase. Priming can be compared across conditions that maintain different relationships to the target. In particular the amount of priming obtained from the focal letter in the same position across both strings can be compared with the priming obtained when the focal letter is in different positions in both strings. Priming in the latter case is equivalent to moved-position priming. There is evidence, however, that moved-position priming of this sort does not obtain (Humphreys et al., 1990). This may be because this condition removes all relative position as well as absolute position similarity between the two strings.

Two other conditions of interest, therefore, are those in which relative positions are either maintained or disrupted, in both the fixed and moved absolute position situations. The conditions that maintain relative positions, referred to as **Context** conditions, contain three

of the target's neighbouring letters in the source, one letter either side of the focal letter, SHBdddd, for example. Conditions that disrupt relative positions, NoContext conditions, contain the same three letters but not as neighbours, dHdBdSd, for example. An additional condition maintaining all the target letters in the source is included for comparison; in this case the two strings are identical and serve as a measure of repetition priming for word-like stimuli.

In the priming conditions the focal letter is always the same as the misspelled letter in the target. In another set of conditions the sources contain the letter that would correct the misspelling if it migrated from the source to the target, SUBJECT, for example. These conditions are designed to elicit migrations, or lexicalisations of the target, and are manipulated in the same way as the priming conditions described above. The final condition is a Word Prime condition, SUBJECT preceded by SUBJECT; this was included (a) to provide a measure of repetition priming for words to compare with the same measure for word-like strings, and (b) to prevent subjects from inferring that all the sources were nonwords.

#### 4. 4. 1 Method

*Stimuli and Design:* A within-subjects repeated measures design was used: the independent variable was the relationship between the two letter-strings presented; the dependent variable was accuracy of response, scored over a variety of measures described under Results.

The stimuli used were of two sorts, one for the source letter-string, one for the target letter-string. The source was a nonword in a particular relationship to the letters present in the the target. Twelve different relationships were constructed. In each case the target was either a word misspelled by one letter or a real word. For each word-like target the source had one of three relationships. The neutral relationship is the control condition in which the source contained none of the letters present in the target. In the Prime conditions the source contained the misspelled letter of the target. In the Migration conditions the source contained

the letter that would correct the misspelling in the target if it migrated across the strings.

Within the Prime and Migration conditions the focal letter that was either a repeat of the misspelled letter in the target or its correction, was sometimes in the same position in both strings and sometimes in different positions. These two manipulations are referred to as fixed position and moved position. Again within both Prime and Migration conditions the critical letter in the source was sometimes surrounded by the two letters that surrounded it in the target and sometimes surrounded by other letters not repeated in the target. These two manipulations are referred to as Context and NoContext.

Three positions of the target were manipulated: the second, fourth, and sixth letters. For each of these three positions, 12 conditions were constructed. As an example the word SUBJECT is shown with all conditions based around changes in the second letter.

<u>Condition</u>	<u>Source String</u>	<u>Target String</u>
1. Prime, fixed, context	SHBdddd	SHBJECT
2. Prime, fixed, no context	dHdBdSd	SHBJECT
3. Prime, moved, context	ddddSHB	SHBJECT
4. Prime, moved, no context	dBdSdHd	SHBJECT
5. Full prime	SHBJECT	SHBJECT
6. Control	dddddd	SHBJECT
7. Migration, fixed, context	SUBdddd	SHBJECT
8. Migration, fixed, no context	dUdBdSd	SHBJECT
9. Migration, moved, context	ddddSUB	SHBJECT
10. Migration, moved, no context	dBdSdUd	SHBJECT
11. Full migration	SUBJECT	SHBJECT
12. Word prime	SUBJECT	SUBJECT

**Subjects and Apparatus:** Ten undergraduates, 5 male and 5 female. All subjects were unpaid volunteers. Other details as for Experiment 6.

**Procedure:** Generally the procedure was as for Experiments 6 and 7, described more fully in Experiment 6. The mask was made up of a random collection of non-alphabetic symbols (e.g. \$£%@&\*). Each subject was presented with trials derived from 12 base words, 144 trials in all. The base-words used for each subject were randomly selected from a group of 20. Before testing began, each subject received 20 practice trials, or more if requested. The procedure for each trial was as follows: first a fixation point between two flanking arrows was presented, then a 100 ms mask, then the source letter-string for 360 ms, then the mask for another 100 ms, then the target letter-string for 120 ms, and finally the mask for another 100 ms. This was followed by presentation of a row of seven dashes each marking the position of one of the presented characters. Each dash was replaced with the letter the subject typed in on the keyboard. Subjects had unlimited time to complete and alter their responses, and proceeded to the next trial by pressing the Spacebar, at their own speed. No feedback was given, but a debriefing was supplied at the end of the experimental session.

## 4. 4. 2 Results

Overall accuracy scores for each position are shown in Table 4. 6 for the Word-prime condition. In comparison with the positional accuracies reported for Experiments 6 and 7 these data show that the left-right effect is greatly attenuated for primed word stimuli but is still present (cf. Estes et al., 1976). The central confusability effect resembles neither that for the single presentation letter-string results (Experiment 6) nor that for the double presentation letter-string results (Experiment 7) in that it is approximately equally distributed across all central positions.

1	2	3	4	5	6	7	Mean
99.1	96.7	94.2	92.5	92.5	90.8	92.5	94.1

**Table 4. 6.** Percentage accuracy scores for each letter position in the Word-prime condition, Experiment 8.

The data for the first cross-position analyses are summarised in Table 4. 7; the Migration analysis is of the number of times the letter misspelled in the target was corrected. It is thus a measure of lexicalisation as much as a true measure of migrations, and is referred to as such.

Condition	Number correct in position	Accuracy scores on focal letters	Lexicalisation scores on focal letters
1. Prime, Fixed-Context	5.7	11 (1.5)	90 (2.4)
2. Prime, Fixed-NoContext	5.3	16 (2.1)	82 (2.8)
3. Prime, Moved-Context	5.6	27 (3.0)	78 (2.9)
4. Prime, Moved-NoContext	5.4	19 (1.8)	79 (2.3)
5. Full Prime	5.8	52 (3.7)	52 (3.5)
6. Control	5.5	17 (2.3)	86 (2.4)
7. Migration, Fixed-Context	5.5	7 (0.7)	100 (1.6)
8. Migration, Fixed-NoContext	5.3	20 (2.2)	82 (2.7)
9. Migration, Moved-Context	5.3	15 (2.4)	78 (3.1)
10. Migration, Moved-NoContext	5.3	17 (1.6)	77 (2.5)
11. Full Migration	5.4	23 (2.4)	81 (3.4)
12. WordPrime	6.6	111 (1.3)	111 (1.3)

**Table 4. 7.** Accuracy and Lexicalisation scores for all conditions over all seven letters, and over the focal letter for the second Prime and the Migration analysis. The accuracy scores over all seven letters are presented as the number of letters correct per trial; the accuracy and migration focal letter scores are presented as absolute totals. Standard deviations of the scores are in brackets alongside the obtained scores. Experiment 8.

A one-way analysis of variance for the overall accuracy scores revealed a significant effect of treatments ( $F(11,108) = 3.6, p < 0.01$ ), but a planned set of linear contrasts revealed no significant differences<sup>9</sup> A Repetitions\*Positions\*Conditions\*Subjects analysis

<sup>9</sup> The following comparisons were made: the two Fixed-prime conditions vs. control; the two Moved-Prime conditions vs. control; the Full Prime condition vs. control; the two Fixed-Prime conditions vs. the two Moved-Prime conditions; and the five Prime conditions vs. the five Migration conditions.



of variance performed on the accuracy scores over the focal letters revealed significant main effects of both position and conditions ( $F(2,18) = 4.95, p < 0.05$ , and  $F(11,99) = 41.71, p < 0.01$ ), but the planned linear contrasts showed no significant differences. Analysis of the reports scored in terms of migrations from the source to the target was also unsuccessful. A Subjects \* Treatments \* Position analysis of variance found a significant treatment effect ( $F(11,95) = 6.95, p < .01$ ), but no significant planned linear contrasts.

Because of these non-significant, planned comparisons, and because post-hoc other comparisons looked interesting, the results were re-analysed by Newman-Keuls tests. This analysis is presented in three parts: focal letter accuracy, accuracy scores over the remaining six letters, and lexicalisation of the focal letter.

*(a) Focal letter accuracy scores.* To obtain an error term for the focal letter accuracy Newman-Keuls test, a one-way analysis of variance of the data gave a significant effect of treatment ( $F(11,108) = 16.5, p < 0.01$ ) and an error term of 4.9. The results of the focal letter analysis are presented in Table 4. 8; each star shows a significant difference between the condition identified by the row and the condition in the column; the row and column numbers are identified on the right of the tables; the results are also presented graphically: conditions which not significantly different from each other are other overlined by a common line.

The conditions which are significantly better than control are: WordPrime; Full Prime (SHBJECT); and Prime-Moved Context (ddddSHB). The Full Migration condition approaches significance. Migration-Fixed Context (SUBdddd) is significantly worse than control and Prime-Fixed Context (SHBdddd) approaches significance. The WordPrime result is evidence of normal repetition priming, which is significantly larger than the nonword repetition priming in the Full-Prime condition (as in e.g., Rueckl, 1990). WordPrime results will not be considered further because the conditions of most interest are those that allow or encourage the processing of letters within words.

	1	2	3	4	5	6	7	8	9	10	11	12	Condition
1	-		*	*	*	*	*	*	*	*	*	*	M-FC
2		-					*	*	*	*	*	*	P-FC
3			-							*	*	*	M-MC
4				-						*	*	*	P-FNC
5					-					*	*	*	M-MNC
6						-				*	*	*	Control
7							-			*	*	*	P-MNC
8								-		*	*	*	M-FNC
9									-	*	*	*	Full-M
10										-	*	*	P-MC
11											-	*	Full-P
12												-	WordPrime

M-FC	P-FC	M-MC	P-MNC	Full-M	P-MC	Full-Prime	WordPrime
(SUBddd)	(SHBddd)	(dddSUB)	(dBdSdUd)	(SUBJECT)	(dddSHB)	(SHBJECT)	(SUBJECT)
		M-MNC	M-FNC				
		(dBdSdUd)	(dUdBdSd)				
		P-FNC					
		(dHdBdSd)					
		<b>Control</b>					
		(dkkkk)					

*Note: The conditions are abbreviated as follows: P: Prime condition; FNC: Fixed position, no context; FC: Fixed position, with context; MNC: Moved position, no context; MC: Moved position, with context; M: Migration condition.*

**Table 4. 8.** Significant differences between conditions for accuracy of focal letter report: Newman-Keuls test.

As a preliminary observation it is difficult to see how Interactive Activation could deal with these results. Prime-Fixed Context (SHBddd) should produce priming but the observed effect is inhibitory. This observation is equally damaging to all models in which letters have independent, position-dependent contributions to processing. Moreover, it is only Context, and never NoContext, conditions that differ from control, either significantly or with a tendency to do so. As in Humphreys et al. (1990), the presence of the focal letter alone has no effect. Under these conditions the representation and processing of isolated letters in the source does not effect subsequent processing of the target. The results also show that moved-position priming is obtained. This extends the previous finding with nonwords to word-like strings. The fixed-part priming differs from that obtained with

nonwords in that here it is inhibitory. Both SHBddd and SUBddd produced inhibition, with a tendency for SUBddd to produce more.

(b) *Accuracy over the other six letters.* The focal accuracy scores subtracted from the overall accuracy scores gives data on accuracy over the other six letters. For this Newman-Keuls test a one-way analysis of variance performed on the data revealed no significant treatment effect ( $F(11,108) = 0.91$ ) and a Mean Square error term of 40.7. The Newman-Keuls test results are presented in Table 4. 9.

	1	2	3	4	5	6	7	8	9	10	11	12	Condition
1	-								*	*	*	*	M-FNC
2		-							*	*	*	*	Full-M
3			-							*	*	*	P-FNC
4				-						*	*	*	M-MNC
5					-					*	*	*	M-MC
6						-					*	*	P-MNC
7							-				*	*	Full-P
8								-			*	*	P-MC
9									-		*	*	Control
10										-		*	M-FC
11											-		P-FC
12												-	WordPrime

Full-Mig (SUBJECT)	M-MC (dddSUB)	P-MNC (dBdSdUd)	Control (dkkkd)	M-FC (SUBddd)	P-FC (SHBddd)	WordPrime (SUBJECT)
M-FNC (dUdBdSd)	M-MNC (dBdSdUd)	Full-Prime (SHJECT)				
P-FNC (dHdBdSd)	P-MC (dkkSHB)					

*Note: The conditions are abbreviated as follows: P: Prime condition; FNC: Fixed position, no context; FC: Fixed position, with context; MNC: Moved position, no context; MC: Moved position, with context; M: Migration condition.*

**Table 4. 9.** Significant differences between conditions for accuracy of report on the six-letters other than the focal letter: Newman-Keuls Test.

The results of the analysis of six-letter accuracy differ in several ways. Prime-Fixed

Context is significantly better than control and Migration-Fixed Context is associated with Prime-Fixed Context but is not significantly different from control. Full Migration and Migration-Fixed NoContext are significantly worse than control, and are most strongly associated with the following three conditions, none of which differ significantly from control: Prime-Fixed NoContext, Migration-Moved NoContext, and Migration-Moved Context. No clear pattern emerges, but the conditions with the least structured letter groups that have some letters in common across source and target, the NoContext conditions, all tend to inhibit performance relative to control.

The two conditions better than control on focal letter accuracy do not differ from control on six-letter accuracy, and are slightly associated with worse performance than control: Full-Prime and Prime-Moved Context. The two conditions tending to be worse than control for the focal letter (Prime-Fixed Context and Migration-Fixed Context) are also the two conditions tending to be better than control on six-letter accuracy, though only Prime-Fixed Context is significantly better in this analysis and significantly worse in the former analysis.

(c) *Lexicalisation of the focal letter (migration analysis).* Table 4. 10 presents the results of the Newman-Keuls re-analysis of the migration data. Only Full-Prime produces less migrations than control, and only WordPrime and Migration-Fixed Context produce more than control. The four Prime and Migration Moved conditions are all associated with less migrations than control, but are not significantly worse; Prime-Fixed Context is associated with more migrations than control but not significantly.

	1	2	3	4	5	6	7	8	9	10	11	12	Condition
1	-	*	*	*	*	*	*	*	*	*	*	*	Full-P
2		-								*	*	*	M-MNC
3			-							*	*	*	P-MC
4				-						*	*	*	M-MC
5					-					*	*	*	P-MNC
6						-					*	*	Full-M
7							-				*	*	M-FNC
8								-			*	*	P-FNC
9									-		*	*	Control
10										-	*	*	P-FC
11											-	*	M-FC
12												-	WordPrime

Full-Prime (SHBJEKT)	M-MC (dddSUB)	Control (dddddd)	P-FC (SHBddd)	M-FC (SUBddd)	WordPrime (SUBJECT)
	M-MNC (dBdSdUd)	Full-Mig (SUBJECT)			
	P-MNC (dBdSdUd)	M-FNC (dUdBdSd)			
	P-MC (dddSHB)	P-FNC (dHdBdSd)			

*Note: The conditions are abbreviated as follows: P: Prime condition; FNC: Fixed position, no context; FC: Fixed position, with context; MNC: Moved position, no context; MC: Moved position, with context; M: Migration condition.*

**Table 4. 10.** Significant differences between conditions for number of migrations: Newman-Keuls test.

Fixed-position migrations into word-like strings are more likely than moved-position migrations. This finding supports the claim (Mozer, 1983; McClelland & Mozer, 1986; Shallice & McGill, 1978; Treisman & Souther, 1986) that migrations tend to maintain position. The results also replicate the context effect, such that presentation of the surrounding letters of a particular letter (Prime-Fixed Context) can induce that letter to be reported even though it was not actually presented. The difference between the conditions where the focal letter is absent and where it is present suggests that having the letter actually present in the source increases the likelihood of it migrating, though since there is no difference between these conditions when the context is moved this only applies when the letter is present in the source in the correct position. That there is also an absence of context effects for the moved position Prime and Migration conditions suggests strongly that for migrations to word-like strings, position-in-string is crucial.

#### 4. 4. 4 Discussion

The most striking result is that the regular trigrams (SdB and SUB, where d represents a changing letter) tend to produce lexicalisations and inhibit focal letter accuracy

when in fixed position but that the SdB trigram facilitates accuracy when in moved position. One interpretation of this result is in terms of neighbourhood frequency. SdBddd and SUBddd can be considered as low-frequency neighbours of the higher frequency word SUBJECT. Segui and Grainger (1990) showed that low-frequency neighbours inhibit processing of high-frequency neighbours. They explain the inhibitory effect through the idea of a "lexical" representational space, with more frequent words as attractor regions into which less frequent neighbours are drawn; in order to process the low-frequency neighbour accurately, the high-frequency neighbour has to be actively inhibited. But why should inhibition of SUBJECT disrupt processing the d in SdBJECT, the target?

One possibility is that when letters require processing as wholes, as for the focal letter, being part of a familiar word-whole confers an advantage over letters in nonwords (as evidenced by the word superiority effect on 2AFC); thus disruption of the whole, SUBJECT, lessens the ability of the whole word to facilitate processing its letters as wholes. This formulation allows a re-interpretation of the Segui and Grainger data: a high-frequency word will dominate a lower frequency neighbour at the whole-word level; to overcome the high-frequency dominance, low-frequency words tend to be processed more at the letter-whole level where their distinctiveness is more apparent; changing the scale of processing within graphemic representational space is harmful to (neighbouring) high-frequency words, which only show their dominance at the word-whole level.

If the inhibition produced by SdBddd and SUBddd on SUBJECT explains the focal letter accuracy inhibition, then the fact that both tend to produce lexicalisations, apparently showing the continued influence of SUBJECT, requires explanation. Here it is necessary to distinguish more carefully between different scales. The attractor effect on SdBddd and SUBddd moves their processing away from the whole-word to sublexical levels; the largest scale at which they can be unambiguously represented is the trigram level, where both SUB and SdB are regularities, both orthographically and in the experiment. When the target SdBJECT is presented, its representation at the whole-word level is inhibited by the change in scale caused by the processing of the source. The effect of the

inhibition is to reduce the speed at which the target letters start to be processed as wholes; this means more opportunity for the target to be processed at the trigram level. At this level a representation of the trigrams in the source remains strongly activated. When the source trigram is SUB this interferes with the development of an SdB representation for the target letters, thereby increasing the likelihood of lexicalisations from the SUB trigram.

The explanation is slightly different when the source trigram is SdB: in order to disambiguate SdB from its high-frequency competitor SUB, processing of the SdB trigram continues down to the letter-whole level. This change of scale is even more inhibitory to SdBJECT, and more specifically to SdBJECT than the inhibition produced by the unrelated control string. The inhibitory effect means that SdBJECT only has the opportunity to be processed at the whole-word level, at which level the common letter d is not apparent. This means that (a) accuracy on the focal letter is reduced, (b) accuracy on the other six letters is enhanced, because they are implicit in the whole-word description, and (c) that lexicalisations tend to occur, though their likelihood is counteracted by the absence of any active representation of U and the relative inhibition of the U in the SUB trigram. Because the inhibition must be stronger from the SdB source than from the control, a strong implication of this interpretation is that the relative inhibitory effect is very specific to neighbours. Neighbours seem to be defined in terms of position-in-the-string because the SUB trigram in a different position produces none of these effects. This implies that the neighbourhood effects are based on representations in which parts are encoded relative to wholes.

How does this interpretation deal with the moved-part priming from ddddSdB? Facilitation can occur because no inhibition of the target is caused by the change in scale of processing. This must mean that the representation of the trigram SdB is not tied to position-in-the-string, so that the representation of the target at trigram and letter levels can interact with the representation of the source at the same levels. In turn this implies that representations of wholes within wholes are more independent of position than representations of parts of wholes.

When does neighbourliness change from being a distinction that requires exaggeration

to one that requires generalising over? This can be examined in the results from the word-like source string, **SdBJECT**, and the word source **SUBJECT**. **SdBJECT** differs from **SdBddd** in that it primes focal letter accuracy, and reduces lexicalisations; **SUBJECT** has broadly similar effects, though the focal letter effect is marginal, and lexicalisations do not differ at all from control. Stimuli that differ by one letter, where one is a misspelling of the other, appear to be treated as the same: **SdBJECT** as source primes **SUBJECT** which enables the target **SdBJECT** to be more rapidly processed at the letter-as-whole level. **SUBJECT** also primes processing of the target **SdBJECT**, but not as specifically as **SdBJECT** does. Monk and Hulme (1983) report an analogous effect whereby misspellings that delete a letter are more difficult to detect than letter substitutions. Letter substitutions, as in this experiment, cue the place of the misspelling while deletions only cue the fact that a misspelling is present somewhere. **SdBJECT** receives the benefit of being very similar to **SUBJECT** but also specifies where and what the subsequent deviation from **SUBJECT** will be. Thus accuracy on the other six letters is higher for **SdBJECT** than for **SUBJECT**. The latter may prime subjects to expect a misspelling and thus lead to pseudo-corrections of other letters, while the former primes subjects to expect a misspelling in the relevant position, thereby reducing pseudo-corrections.

Several complex developments of the original distinction between parts of wholes and wholes within wholes have been used to interpret the results of Experiment 8; in particular the idea that processing principally takes place at the largest scale at which the input is disambiguated, but that stimuli which are better, more familiar wholes at one scale are better able to change the scale at which they are processed. Perhaps the most speculative development is the suggestion that the scale at which one stimulus is processed affects the processing of subsequent, similar stimuli, but has no effect on stimuli further away in multidimensional graphemic representational space. These ideas are discussed in more detail in Chapter 5.



---

## GENERAL DISCUSSION

---

In Chapter 1 two main aims of this research were outlined: to investigate the nature of the internal representations that encode relational information in letter-strings, and to relate the results of the investigation to current, computationally motivated, models of word recognition. Three issues were distinguished as being particularly important. Firstly, the nature of the representations used in the immediate access to language-specific descriptions that embody knowledge about visually presented words, in particular whether the visual input is preprocessed into position-specific slots prior to access. Secondly, and more generally, the relationship between low-level visual descriptions, not specialised for visual language, and representations embodying graphemic knowledge. Thirdly the organisation of the graphemic representational space, in particular whether and in what ways familiar and unfamiliar letter-strings are differently represented. Before returning to these issues in §5. 2, the experimental findings are briefly reviewed in §5. 1; §5. 3 builds the implications into a speculative model, discusses some other evidence for the position adopted, and finally provides in a brief conclusion a sketch of the more general relevance of the findings.

### 5. 1 Summary of findings

Experiment 1 showed that when a trigram part of a recently learned nonword letter-string was maintained in a briefly-presented test string, the test string showed perceptual transfer from the learned letter-string in that it was more accurately reported. No difference in the amount of transfer was found between maintaining the part in the same position (fixed-part) in the string and maintaining the part in a different position (moved-part) in the string. Transfer was also

obtained from three-letter to six-letter strings and vice versa. Experiment 2 replicated the finding of fixed-part transfer but only for parts made of adjacent rather than non-adjacent letters. Experiments 3 and 4 used a prototype-extraction paradigm and found that novel parts made of adjacent letters, and novel wholes made of those parts, were easier to learn than parts made of non-adjacent letters and the wholes from those parts. Experiment 5 replicated the finding of moved-part transfer and found that it was restricted to parts made of adjacent letters. It was argued that these results largely constrain the encoding of inter-letter relationships to neighbouring letters.

Experiments 6-8 explored migration errors in a variety of contexts. Experiment 6 found that reports of letters in nonwords showed gradients of positional accuracy, with most positional errors occurring close to the correct position. The range over which positional errors were made was different for different absolute positions in the string. Experiment 7 found that intrusions into the report of the second of two briefly-presented nonwords from the first nonword did not invariably maintain position though large numbers of them did. Experiment 8 presented misspelled words preceded by nonwords that either encouraged the detection of the misspelling or its correction. Correction involved the migration of a letter from the source and was found only to occur significantly more than control when the letter was in the same position in both strings and surrounded by contextual letters. Evidence was obtained about the size of parts important for processing: migrations were slightly more likely when only the context and not the target migration letter was present in the source. This suggests that processing elements larger than the single letter and capable of generalising over variations are responsible for at least some of the processing. Detection of misspelling showed both facilitation and inhibition. Facilitation was obtained when the facilitating part was in a different position in the source as compared to the target. When the part was in the same position inhibition was found. This result was interpreted in terms of the scales used to disambiguate the input and generate a task-relevant output.

Two strong conclusions follow: (i) the coding of the position of letters in nonwords does not only involve the representation of part-whole relationships; (ii) the coding of the position of

letters in words or familiar wholes strongly involves the representation of part-whole relationships. The results do not allow the inferences that relational coding in nonwords does not use part-whole relationships at all, and that relational coding in words does not use part-part relationships at all, but they do suggest an important difference between familiar and unfamiliar wholes. This conclusion exactly parallels McNicol and Heathcote's (1986) summary of research into the coding of (temporal) order information in auditorily presented sequences of letters: with sequences of letters that do not allow *grouping*, the characteristic error in report is to confuse two adjacent letters, while with sequences that do allow grouping the most characteristic error is to confuse two items in similar positions within different groups (e.g., Ryan, 1969; Wickelgren, 1967).

These conclusions can be extended into the thesis that parts of familiar wholes have their position represented in the form of part-whole relationships, while parts of unfamiliar wholes have their positions represented in the form of part-part relationships. This formulation leaves many questions unanswered: Are familiar wholes represented in relation to other familiar wholes? Are familiar wholes represented in relation to the larger wholes of which they are a part? Under what conditions does a particular item act as a part or a whole? What is it that determines grouping, or the representation of an item as a familiar whole?

## 5. 2 Theoretical implications

The results of Experiments 1-7 provide evidence that preprocessing of the image to provide position-in-the-string information is not inevitable. This conclusion is not easily accommodated by part-whole theories of relational encoding. As discussed in Chapter 1, models of the processing of letter-strings, such as IAM (McClelland & Rumelhart, 1981) assume a stage of preprocessing before access to graphemic representations and after the first stages of low-level visual analysis. There is now a varied body of evidence to suggest that the part-whole theory as exemplified in the position-specific slots of IAM is incorrect, at least partially. The results of the present experiments (1-7), together with similar findings of part-part transfer by Humphreys et

al (1990), constitute direct evidence. Additional evidence comes from two sources. When words are presented with one letter incorrect, misspellings which delete a letter are harder to detect than letter substitutions (Monk & Hulme, 1983). A strict IAM coding of the input would predict the opposite. A strong prediction of the IAM coding scheme is that processing repeated tokens of the same letter type should present no problem because letter tokens are repeated for each letter-in-the-string position. In fact when presented with nonword strings containing repeated letters, subjects tend to underestimate the number of repetitions of particular letters (Mozer, 1989), suggesting that the type/token distinction is not as clear as proposed by IAM. At the very least IAM needs modification of its feature and letter levels, to allow relational encoding between letters in nonword letter-strings. The solution offered in PABLO (McClelland, 1986) will not work, because the coarse-coded letter-level description is still tied to position-in-the-string. BLIRNET (Mozer, 1987), which maps relationally encoded letter-clusters into a lexical network, has not been extensively tested but appears to offer a partial solution, as does the Seidenberg & McClelland (1989) back-propagation network which maps letter-clusters into hidden units.

On the other hand there are at least two sources of evidence that support part-whole encoding theory. The first is evidence from a number of paradigms that positional-frequencies of individual letters in different positions in words can strongly influence processing (e.g., Mason, 1975). However the same paradigms very often provide evidence for the influence of inter-letter transitional probabilities independent of the influence of letter-position frequencies (e.g., Massaro, Venezky, & Taylor, 1979). Since both types of probability, one position-dependent, the other position-independent, affect both words and pseudowords it is very difficult to make strong inferences without more detailed research. Two forms of error also provide support for the part-whole encoding idea: the errors of some neglect dyslexics and migration errors. However, the neglect dyslexia data are not complete enough to motivate firm conclusions, and, in any case, the claimed position specificity is not impressive.

Experiment 8 shows that part-whole encoding does play a role in the processing of word-like stimuli; the evidence is the inhibition of focal letter accuracy and corresponding increase in

lexicalisations produced by fixed-position parts. The inhibition is taken to be the result of interference from a strong attractor on letter-string processing. Initial activation of the attractor effect appears to depend on letters or letter-groups in specific positions in the string. The full story is more complex because of the moved-part facilitation and because very word-like, very similar or identical, source strings also facilitate detection of the spelling mistake in the target. The word representation is able to facilitate processing at some sublexical level, such that a single spelling mistake is detected more fluently. Moreover, at this sublexical, or fine grain, level of processing, the representations that can facilitate processing are not tied to positions in the string. Migrations implicate part-whole processing but more strongly for words than nonwords.

On the other hand priming is obtainable when letter-letter but not letter-whole relationships are maintained for word-like as well as for nonword letter-strings. The moved-part priming remains difficult to interpret here and in Experiments 1-5: it strongly implies that familiar trigrams are not always encoded in their relationships to the whole string. Whether this means that familiar trigrams are encoded in relation to other trigrams is not possible to say. Nor do the results allow inferences about the encoding of the trigram's parts because both part-whole and part-part descriptions within the trigram would produce moved-part priming as long as the trigram itself is not encoded in part-whole terms. The tentative solution offered to this problem is to treat parts and wholes differently depending on the task requirements. If both letters and trigrams can be treated either as parts of a whole or as wholes within a larger whole then the following possibility arises. When items at any scale are processed as parts they are represented in terms of their part-whole relations; this possibility may be less open for unfamiliar word-wholes such as the letter-strings used in Experiments 1-7 simply because they are less well processed as wholes. When items are processed as wholes they are represented in terms of their relations to other wholes at the same scale, when necessary. Otherwise they are treated as wholes in isolation.<sup>1</sup> It remains unclear whether whole-whole relationships are directly

---

<sup>1</sup> In terms of the possibilities outlined in Chapter 1, wholes represented in this way may, in fact be represented in egocentric coordinates. This type of coding of position is obviously less relevant in word than object recognition.

encoded, or are represented through the dynamic linking of their representations. Either way the method used is general across stimulus objects which explains why an interpretation of the results of Experiment 1 in terms of MIRAGE was possible. The implication is that flexible representations are needed, depending on the task and on the type of input. Neither BLIRNET nor Seidenberg and McClelland's (1989) back-propagation network offer the required flexibility.

## 5. 3 Modelling graphemic representational space

### 5. 3. 1 Mapping into graphemic representations

The first step to a solution to the problem of providing flexible processing within multiple scales of representation may be to use a more flexible visual input. This is exactly what MIRAGE provides: a visual description that scans from coarse to fine scales over time. There is not room for a full discussion of the ways in which MIRAGE could be used to model the VIP into graphemic representations, but some striking possibilities are provided by fixation and reading span data.

(a) Brief fixations separated by saccades provide VIP in discontinuous chunks. Skilled readers' perceptual span<sup>2</sup>, approximates to 18 letters, 3 to the left of fixation, and 15 to the right (McConkie & Rayner, 1975; 1976). The width of the span is less for difficult fixated words (Henderson & Ferreira, 1990), which shows that it is directly under control from higher levels of representation, as does the reversal of the span asymmetry for languages written right to left (Pollatsek, Bolozky, Well, & Rayner, 1981). Fixation duration is also heavily dependent on word-frequency and familiarity. Average fixation duration is 250 ms (Rayner & Pollatsek, 1987), but fixations vary with text difficulty (Just & Carpenter, 1980), and readers' experience (Taylor, Frackenpohl, Pettee, 1960). Both findings can be accommodated by assuming that

---

<sup>2</sup> The optimum size for normal reading of a window of unmutilated text surrounded by rows of Xs.

MIRAGE coarse filters are switched out until a description is achieved that is adequate for an unambiguous graphemic representation.

(b) Only the six letters centred around fixation provide accurate letter-identity information (McConkie & Zola, 1987). The 12 letters, roughly two words, to the right of fixation provide insufficient information for letter identification, and therefore need re-foveation. Saccade length and fixation duration can be affected by parafoveal presentation of space information alone (Morris, Rayner, & Pollatsek, 1990), by parafoveal words' end letter identities, and, more controversially, by parafoveal word shape (Rayner, 1975; Rayner, McConkie, & Ehrlich, 1978; Rayner, Well, Pollatsek, & Bertera, 1982), which suggests that parafoveal information is of the coarseness of the MIRAGE filters that provide length, some word shape, and end-letter information (Watt et al., in press).

(c) Since average word length is 4.5 characters, or 6 characters with a space on either side, and average saccade length is also 6 characters (Rayner & Pollatsek, 1987), it is a useful simplifying assumption that the middle of words are normally fixated (O'Regan, Levy-Schoen, Pynte, & Brugailiere, 1984). This makes the point that further processing of the input to normalise it in terms of position is unnecessary given the flexibility of the input itself. Fixations tend exactly to bring the centre of the next word into the centre of the fovea (O'Regan et al, 1984), independent of word frequency and nonwordness, and even for longer words for which the optimum viewing position is to the left of centre (O'Regan & Levy-Schoen, 1987); the same is found for pairs of line targets (Coren & Hoenig, 1972; Findlay, 1982). The switching out of filters is implicated in that delaying saccades to 500 ms for words (Coeffe & O'Regan, 1987) and to over 150 ms for line targets (Findlay & Harris, in press), increases their accuracy.

### 5. 3. 2 Word recognition

Perhaps the strongest implication of 5. 3. 1 for models of the graphemic representational space is that processing takes place at a variety of scales, beginning with coarse information and

continuing to fine-grained information when necessary. This idea has been expressed before, but not strongly, and not, therefore, subjected to direct experimental scrutiny. EPAM, for example, the serial model of word recognition (Richman & Simon, 1989) discussed in §1.3.2, processes input at the word-level if possible and only moves to the letter-level if no adequate word representation is available, when processing nonwords, for example. This is one candidate case where word-wholes are not used in processing. Another is the set of tasks described in §1.2 as involving treating letters in letter-strings as wholes within larger wholes. These tasks contrast with tasks such as lexical decision that require processing the entire letter-string as a whole and may allow the constituent letters to be treated as parts of the word-whole. Humphreys & Bruce (1989) make a similar case for a distinction between lexical decision and naming:

Words are visually processed at various "levels". Lexical decisions can be based on descriptions coded across the whole-word. More local descriptions of sub-word segments are also derived, but more slowly. Such segments can be transcribed into phonology, and primarily affect naming tasks (p. 234).

Thus the claim is that stimuli are initially processed at a coarse, word-scale, level, and then if necessary at a sublexical (e.g., letter) level. This contrasts with IAM, for example, where processing begins at the sublexical level, though the interactions between levels allow the word-level to influence the letter-level. This interaction generates the word superiority effect on 2AFC, because feedback is stronger when word-level representations are active. If processing begins at the word-level, the implication of the 2AFC word superiority effect is that the presence of a word-whole representation enables processing to switch more accurately, or more quickly, to the letter-level scale.

Three points require brief documentation: (i) the claim that words differ in the "goodness" of their coarse representations, (ii) that words are better than nonwords at enabling processing to switch to sublexical levels, and (iii) that word-level information is available *before* sublexical information.

**(i) Whole-word representations:** there is much evidence to show that stimuli differ in



their whole-word or whole-string representations. (1) Word-nonword advantages are found in a variety of tasks that require whole-string computations: letter-string-letter-string search (Staller & Lappin, 1981); full report thresholds (Osgood & Hoosain, 1974); full report accuracy (Cattell, 1886; Neisser, 1967); delayed full report, or recall (Miller, Bruner, & Postman, 1954); string matching (Eichelman, 1970; Pollatsek, Well, & Schindler, 1975), and string naming (Theios & Muise, 1977). (2) Moreover words differ among themselves on these tasks, such that more frequent words show best performance (e.g., full report accuracy: Johnston, 1978; McClelland & Johnston, 1977; naming times: Forster & Chambers, 1973; Frederiksen & Kroll, 1976). (3) In general, lower-case print is read faster than upper case (Woodworth, 1938; Smith, 1969; Fisher, 1975) which may be because lower-case word-shapes are more distinctive than upper-case, because of the patterns produced by different letter-features. (4) Pseudowords also show advantages over nonwords on most tasks, which suggests that they, too, are better represented than nonwords at a coarser scale than the letter-level (e.g., recall: Miller et al, 1954; matching: Staller & Lappin, 1981). (5) There is also evidence that wholes at other scales differ in their goodness as well: (a) letter constituents of bigrams only show repetition priming when part of a low rather than high-frequency bigram (Greenberg & Vellutino, 1988); and (b) similarly word constituents of two-word phrases only show repetition priming when parts of low-frequency phrases (Osgood & Hoosain, 1974).

(ii) *Word superiority at letter-level tasks*: again this claim is well documented: words and pseudowords both provide better performance than nonwords on 2AFC (e.g., Aderman & Smith, 1981; Carr, Davidson, & Hawkins, 1978; Reicher, 1969; Wheeler, 1970), and on letter-word search (Krueger, 1970; Mason, 1975). Evidence that the word superiority depends on the availability of coarse-scale representations is that word performance on 2AFC improves with increasing word length, at least from two to four letters, while nonword performance does not change (Samuel, van Santen, & Johnston, 1982); no superiority obtains for letters that are also words ("I", "A"; Samuel et al, 1982; Wheeler, 1970). This picture is complicated by the finding of word *inferiority* at some tasks, notably at target letter cancellation in passages of text (e.g., Drewnowski & Healy, 1977; Healy, 1976; 1980), where the typical

finding is of an increased number of errors on more frequent words. However this task differs from 2AFC in that (a) it presents passages rather than single words, and that (b) errors increase when reading for meaning (Smith & Groat, 1979), whereas repetition priming of 2AFC is not influenced by prior reading of the prime word in a coherent context (Ratcliff, McKoon, & Verwoerd, 1989). These differences make it plausible to treat word superiority and word inferiority as being effects at different levels, and to see the word inferiority results as more evidence for differences in the goodness of whole-word representations.

*(iii) Availability of coarse and fine information:* the claim that coarse information is more quickly available is on less firm ground. However a number of different lines of evidence point to this conclusion. (1) Word-word search is faster than letter-word search, with successive presentation (Johnson, 1975; 1977; Johnson, Turner-Lyga, & Pettegrew, 1986; LaBerge, 1983; Marmurek, 1977; Sloboda, 1976; 1977). When words are processed as words, by instruction, reaction times to single-letter probes are faster than when processed as strings of letters (LaBerge, 1983). The word-nonword advantage on 2AFC can be abolished by manipulations which allow increased processing time (longer stimulus-mask delays; Massaro & Klitzke, 1979) or which specify in advance the critical letter or letter-position (e.g., Smith & Haviland, 1972; Johnston & McClelland, 1974; Thompson & Massaro, 1973)<sup>3</sup>. (2) In general word length has no effect on words in whole-string tasks but large effects on nonwords (matching: Eichelman, 1970; lexical decision: Young & Ellis, 1985; Frederiksen & Kroll, 1976). Moreover when doing letter-word matching with prespecified positions in the string, the nature of the other letters in the string affects words but not nonwords (Johnson, 1986; Johnson & Marmurek, 1978; Marmurek, 1986). (3) There is some evidence that learning to read aloud proceeds from reading at a whole-word level to reading via sublexical levels when required (Harris & Coltheart, 1986). (4) In reading aloud, effects of sublexical regularities on naming times are found, but only for more slowly read, low-frequency words (Andrews, 1982;

---

<sup>3</sup> It is controversial whether specification of the letter or the position is crucial: precueing letters but not position maintains the word superiority sometimes (Reicher, 1969; Spector & Purcell, 1977) but not always (Estes, 1975; Holender, 1979).

Parkin, 1984; Seidenberg, Waters, Barnes, & Tanenhaus, 1984; Taraban & McClelland, 1987). Moreover, when reading nonwords aloud the effects of regularities at levels coarser than the single-letter level are apparent, suggesting that these coarser effects are obligatory (e.g., Glushko, 1979; Kay & Bishop, 1987; Taraban & McClelland, 1987).

What are the implications of these points for models of graphemic representations? It might be possible to save IAM by simply turning it upside-down so that processing begins at the word-level, with coarsely defined information, and proceeds to sublexical levels. Flexibility in terms of task demands is provided by allowing read-out from any of the levels (it was assumed that letter-level read-out was crucial), but the model still suffers two weaknesses: no intermediate levels of regularity, corresponding to bigrams or trigrams, implicated by the results of all the experiments reported herein, and no way of processing letters in nonwords independent of their positions in the string.

Golden's LW model offers a partial solution by collapsing feature, letter, and word levels into a single graphemic network, in which word representations are distributed across their constituent parts. This appears to be helpful, except that the network is coded with multiple token representations for different positions in the string, provided by the same preprocessing assumed by IAM.

A rather different solution is to treat graphemic representations as being embodied in a single network within which processing can occur at a variety of scales. At the coarse scale no independent information is explicitly available about the constituents of what is represented; letters are treated as parts of wholes, to be made explicit if necessary by other domains of representation. At finer scales different sizes of item become wholes and are explicitly represented; the larger wholes are now no longer made explicit. The scale of processing within this network can be set, either quite specifically to particular neighbourhoods or regions of representational space (as in Experiment 8), or more generally, as evidenced by the findings that expecting nonwords reduces the 2AFC word-nonword advantage (Aderman & Smith, 1971), and expecting pseudowords increases the pseudoword-nonword advantage (Carr et al, 1978), and the similar finding that search for words is slower in nonwords than in words, and

search for nonwords slower in words than in nonwords (Treisman & Souther, 1986).

Such a network, if self-organised, would provide an information-theoretic optimum solution: largest amount of processing space devoted to the most probable occurrences, but it presents a learning problem. If, as seems plausible, letter representations are learned from, or abstracted out of, word representations, rather than the other way round, then an interesting possibility arises. Letter representations should be more able to generalise over variations in the VIP than word representations. There is some evidence to support this: responses to letters are almost completely orientation-invariant (Corballis & Nagourney, 1978; Eley, 1982; Hock & Tromley, 1978; White, 1980), whereas words, and particularly pseudowords, show strong orientation effects (Koriat & Norman, 1984; 1985; 1989; Navon, 1978). This explains why on letter-whole tasks, like 2AFC, visual disruptions such as cAsE-mIxInG equally impairs words, pseudowords, and nonwords (McClelland, 1976; Adams, 1979). For whole-string tasks the effects of visual manipulations should depend on frequency. Increasing frequency of occurrence means experience in a wider range of visual forms which will allow better generalisation. Low-frequency words, pseudowords, and nonwords are more vulnerable than high-frequency words. Nonwords are more disrupted than pseudowords and words by case-mixing in lexical decision and naming (Besner, Davelaar, Alcott, & Parry, 1984; Besner & Johnston, 1987; Besner & McCann, 1987). Case-mixing and handwriting particularly disrupt naming low-frequency words (Besner & McCann, 1987; Manso de Zuniga, 1988).

However, if the task is one in which nonwords processed at the letter-level normally show equivalent performance to words, then pseudowords could show more disruption because of their obligatory processing at coarser, more easily disrupted levels. As evidence for this claim, vertical presentation disrupts pseudoword but not nonword naming<sup>4</sup> (Bryden, 1970), and matching pseudowords is more disrupted by case-mixing than matching nonwords (Pollatsek, Well, & Schindler, 1975; Taylor, Miller, & Juola, 1977), when matching can be performed on initial letter identity information.

---

<sup>4</sup> Similar presentational manipulations normally interfere with word and pseudoword but not nonword naming (Mewhort, 1974).

The effects of different masks can be interpreted in a similar way. The more similar the representation of the mask to that of the target, the greater the disruption: word masks more than letter-fragment more than light-flash masks (e.g., Taylor & Chabot, 1978). Letter-fragment masks increase word-nonword 2AFC superiority (Johnston & McClelland, 1973; 1980) by being more damaging to nonwords than words. This may be because they interfere selectively with coarse-scale representations, emphasising the more complete specification of constituents in words than nonwords. Flash-masks impair representations equally at the whole and part levels. Word-masks affect both coarse-scale descriptions and some of the descriptions at finer-levels, leaving less difference between word and nonword representations<sup>5</sup>.

### 5. 3. 3 Conclusions and extensions

A model of the graphemic representational domain is required in which the parts of a familiar whole are tied to their part-whole relationships but the wholes within a larger whole can be represented in relation to one another. Processing must be sensitive to the degree to which a stimulus is a good whole at a coarse scale, to task demands, and to dynamic changes in the scale of currently active representations. It is possible that dynamical connectionism (5. 3. 2) can provide such a model. Von der Malsburg (1985) provides a suggestion that the type of scale-space that seems desirable might be a natural consequence of processing with dynamic links:

A typical dynamic process would start in an initial state in which a number of cells are active and communicate by a matrix  $W$ . Most of the elements of  $W$  vanish because the corresponding synapses do not exist. The existing synapses may be in their resting state or they may already have been modulated in strength (e.g. by externally induced correlations). This initial structure of  $W$  (and, of course, also of the activity distribution) will now evolve dynamically, until a stable state in the form of a decomposition of the set of cells (i.e. of the matrix) into blocks is reached.

---

<sup>5</sup> This explanation is exactly opposite to that usually offered for masking effects (e.g., McClelland & Rumelhart, 1981).

It is important that under realistic conditions the blocks formed correspond to a useful segmentation into subsymbols... The system can find segments which correspond to coherent objects in the external world... Feedback between the signal correlations and synaptic modulation soon leads to a clear-cut block structure. As a consequence, cells responding to parts of different objects are temporally anti-correlated and thus are prevented from interfering with each other during pattern recognition.

The blocks formed in a first stage of organization may be unstable and may decompose into smaller blocks. This can go on through a number of stages until blocks of a certain minimal size are formed... If the system is regulated such that connections between blocks are only weakened and not ruptured completely, a hierarchical system of blocks and sub-blocks may be formed... Such a system is ideally suited to form hierarchically structured semantic symbols (p.705).

A working simulation of this system at this level of complexity has still not been produced, let alone applied to word recognition. It makes the point, however, that similar representational structures may underlie the processing of objects and faces as well as words. There is much evidence for commonalities between the three. Objects and faces both show superiority effects comparable to the word superiority effect, i.e. processing parts of a face compared to processing parts of a jumbled face (e.g., Homa, Haver, & Schwartz, 1976; Purcell & Stewart, 1988; Weisstein & Harris, 1974; McClelland, 1978). Faces as well as words are severely disrupted by being shown upside-down (e.g., Diamond & Carey, 1980), and when inverted their features contribute independently and serially to recognition latencies (Sergent, 1984), as do the letters of inverted words (Koriat & Norman, 1984).

Moreover there are disorders of object and face recognition which seem to be analogous to letter-by-letter reading (e.g., Humphreys & Riddoch, 1987), or at least to implicate problems in maintaining representations of wholes while switching to processing parts as wholes. Interactions between parts and wholes of objects (or local and global processing) have also been widely studied. One classic finding is the superiority of global over local information when the two are set against each other in large letter stimuli constructed as a pattern of many repetitions of much smaller letters (e.g., Navon, 1977), though this effect now seems to be

dependent on a complex variety of factors such as visual angle (Kinchla & Wolfe, 1979), attention to a particular level (Grice, Canham, & Boroughs, 1983), and number and size of parts (Kimchi & Palmer, 1982).

To my knowledge, the full application of an analysis of these findings in terms of spatial scale and the way attention can, at least partly, determine the scale of processing has yet to be published. One particularly suggestive result is that performing global discriminations enhances responses to low spatial frequency patterns, while doing local discriminations enhances responses to high spatial frequency patterns (Shulman & Wilson, 1987a and b). The idea of scale-space, within which processing treats constituents of an identified whole as parts of that whole, but treats identified wholes in relation to one another, may prove useful in the development of computational models of word, face, and object recognition.

# BIBLIOGRAPHY

- Adams, M.J. (1979). Models of word recognition. *Cognitive Psychology*, 11, 133-176.
- Aderman, D. & Smith, E., (1971). Expectancy as a determinant of functional units in perceptual recognition. *Cognitive Psychology*, 2, 117-129.
- Allport, D.A. (1977). On knowing the meaning of words we are unable to report: The effects of visual masking. In S. Dornic (Ed.), *Attention and Performance*, VI. Hillsdale, N.J.: Erlbaum.
- Allport, D.A. (1987). Selection for action: Some behavioural and neurophysiological considerations of attention and action. In H. Heuer & A.F. Sanders (Eds), *Perspectives on perception and action*. Hillsdale, N.J.: Erlbaum.
- Andrews, S. (1982). Phonological recoding: Is the regularity effect consistent? *Memory and Cognition*, 10, 565-575.
- Andrews, S. (1989). Frequency and neighbourhood effects on lexical access: Activation or search? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 802-814.
- Angolillio-Bent, J.S. & Rips, L.J. (1982). Order information in multiple element comparison. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 392-406.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, 61, 183-193.
- Averbach, E. & Coriell, A.S. (1961). Short-term memory in vision. *Bell System Technical Journal*, 40, 309-328.
- Ballard, D.H. (1981). Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition*, 13, 111-122.
- Balota, D.A. & Chumbley, J.I. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance*, 10, 340-357.
- Baron, J. & Thurston, I. (1973). An analysis of the word superiority effect. *Cognitive Psychology*, 4, 207-228.



- Becker, C.A. (1976). Allocation of attention during visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 556-566.
- Begg, L.E.J. (1988). *Transfer effects in word perception and theories of object-based relational descriptions*. Unpublished thesis, University of Stirling.
- Besner, D., Dennis, I., & Davelaar, E. (1985). Reading without phonology. *Quarterly Journal of Experimental Psychology*, 37A, 477-491.
- Besner, D., Davelaar, E., Alcott, D., & Parry, P. (1984). Wholistic reading of alphabetic print: Evidence from the FDM and the FBI. In L. Henderson (Ed.), *Orthographies and reading*. London: Erlbaum.
- Besner, D. & Johnston, J.C. (1987). Reading and the mental lexicon: On the interaction of visual, orthographic, phonological, and lexical information. In W. Marslen-Wilson (Ed.), *Lexical processes and representation*. Cambridge, Mass.: MIT Press.
- Besner, D. & McCann, R.S. (1987). Word frequency and pattern distortion in visual word identification and production: An examination of four classes of models. In M. Coltheart (Ed.), *Attention and Performance, XII*. London: Lawrence Erlbaum Associates Ltd.
- Besner, D. & Swan, M. (1982). Models of lexical access in visual word recognition. *Quarterly Journal of Experimental Psychology*, 34A, 313-325.
- Biederman, I. (1987). Recognition by components: A theory of human image understanding. *Psychological Review*, 94, 115-145.
- Bienenstock, E. & von der Malsburg, C. (1987). A neural network for invariant pattern recognition. *Europhysics Letters*, 4, 121-126.
- Bjork, E. & Murray, J.T. (1977). On the nature of input channels in visual processing. *Psychological Review*, 84, 472-484.
- Bowey, J.A. (1990). Orthographic onsets and rimes as functional units in reading. *Memory and Cognition*, 18, 419-427.
- Brown, G. (1987). Resolving inconsistency: A computational model of word naming. *Journal of Memory and Language*, 26, 1-23.
- Bryden, M.P. (1966). Accuracy and order of report in tachistoscopic recognition. *Canadian Journal of Psychology*, 20, 262-272.

- Bryden, M.P. (1970). Left-right differences in tachistoscopic recognition as a function of familiarity and pattern orientation. *Journal of Experimental Psychology*, 84, 120-122.
- Bryden, M.P., Monder, T.A., Loken, M., Ingleton, M.A., & Bergstrom, K. (1990). Locus of information in words and the right visual field effect. *Brain and Cognition*, 14, 44-58.
- Bub, D. (1990). *Implicit word recognition in acquired dyslexia*. Paper presented to Symposium on Consciousness and Cognition: Neuropsychological Perspectives, St. Andrews.
- Carr, T.H., Davidson, B.J., & Hawkins, H.L. (1978). Perceptual flexibility in word recognition: Strategies affect orthographic computation but not lexical access. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 674-690.
- Cattell, J.M. (1886). The time taken up by cerebral operations. *Mind*, 11, 220-242; 377-392; 524-538.
- Coeffe, C. & O'Regan, J.K. (1987). Reducing the influence of non-target stimuli on saccade accuracy: Predictability and latency effects. *Vision Research*, 27, 227-240
- Cohen, M. & Grossberg, S. (1986). Neural dynamics of speech and language coding: Developmental programs, perceptual grouping, and competition for short-term memory. *Human Neurobiology*, 5, 1-22.
- Coltheart, M., Davelaar, E., Jonasson, J.T., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention and Performance*, VI. Hillsdale, N.J.: Erlbaum.
- Colombo, L. (1986). Activation and inhibition with orthographically similar words. *Journal of Experimental Psychology: Human Perception and Performance*, 12, 226-234.
- Corballis, M.C. (1988). Recognition of disoriented shapes. *Psychological Review*, 95, 115-123.
- Corballis, M.C. & Nagourney, B.A. (1978). Latency to categorise disoriented characters as letters or digits. *Canadian Journal of Psychology*, 32, 186-188.
- Coren, S. & Hoenig, P. (1972). Effect of non-target stimuli upon the length of voluntary saccades. *Perceptual and Motor Skills*, 34, 499-508.

- Coslett, H.B. & Saffran, E. (1989). Evidence for preserved reading in "Pure alexia". *Brain*, *112*, 1091-1110.
- Costello, A.de L. & Warrington, E.K.(1987). The dissociation of visuospatial neglect and neglect dyslexia. *Journal of Neurology, Neurosurgery and Psychiatry*, *50*, 1110-1116.
- Cowie, R. (1985). Reading errors as clues to the nature of reading. In A.W. Ellis (Ed.), *Progress in the psychology of language, 1*. London: Erlbaum.
- Crovitz, H.F. & Schiffman, H.R. (1965). Visual field and the letter span. *Journal of Experimental Psychology*, *70*, 218-223.
- Davelaar, E., Coltheart, M., Besner, D., & Jonasson, J.T. (1978). Phonological recoding and lexical access. *Memory and Cognition*, *6*, 391-402.
- Dennis, I. & Newstead, S.E. (1981). Is phonological recoding under strategic control? *Memory and Cognition*, *9*, 472-477.
- De Yoe, E.A. & Van Essen, D.C. (1988). Concurrent processing streams in monkey visual cortex. *Trends in Neurosciences*, *11*, 219-226.
- Diamond, R. & Carey, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, *115*, 107-117.
- Diller, L. & Weinberg, J.P. (1977). Hemi-inattention in rehabilitation: The evolution of a rational remediation program. In E.A. Weinstein & R.P. Friedland (Eds.), *Hemi-inattention and hemisphere specialisation. (Advances in neurology, 18)*. New York: Raven Press.
- Drewnowski, A. & Healy, A.F. (1977). Detection errors on *the* and *and*: Evidence for reading units larger than the word. *Memory and Cognition*, *5*, 636-647.
- Duncan, J. (1987). Attention and reading: Wholes and parts in shape recognition. In M. Coltheart (Ed.), *Attention and Performance, XII*. London: Erlbaum.
- Eichelman, W.H. (1970). Familiarity effects in the simultaneous matching task. *Journal of Experimental Psychology*, *86*, 275-282.
- Eley, M.G. (1982). Identifying rotated letter-like symbols. *Memory and Cognition*, *10*, 25-32.

- Ellis, A.W., Flude, B.M., & Young, A.W. (1987). "Neglect dyslexia" and the early visual processing of letters in words. *Cognitive Neuropsychology*, 4, 439-464.
- Eriksen, B.A. & Eriksen, C.W. (1974). Effects of noise letters upon the identification of a target letter in a non-search task. *Perception and Psychophysics*, 16, 143-149.
- Estes, W.K. (1975). The locus of inferential and perceptual processes in letter identification. *Journal of Experimental Psychology*, 104, 122-145.
- Estes, W.K. (1977). On the interaction of perception and memory in reading. In D. LaBerge and S.J. Samuels (Eds), *Basic processes in reading: Perception and comprehension*. Hillsdale, N.J.: Erlbaum.
- Estes, W.K. (1982). Similarity-related channel interactions in visual processing. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 353-362.
- Estes, W.K., Allmeyer, D.H., & Reder, S.M. (1976). Serial position functions for letter identification at brief and extended exposure durations. *Perception and Psychophysics*, 19, 1-15.
- Evett, L.J. & Humphreys, G.W., (1981). The use of abstract graphemic information in lexical access. *Quarterly Journal of Experimental Psychology*, 33A, 325-350.
- Feldman, J.A. (1982). Dynamic connections in neural networks. *Biological Cybernetics*, 46, 27-39.
- Feldman, J.A. (1985). Four frames suffice: A provisional model of vision and space. *Behavioral and Brain Sciences*, 8, 265-289.
- Feustel, T.C., Shiffrin, R.M., & Salasoo, A. (1983). Episodic and lexical contributions to the repetition effect in word identification. *Journal of Experimental Psychology: General*, 112, 309-346.
- Findlay, J.M. (1982). Global visual processing for saccadic eye movements. *Vision Research*, 22, 1033-1045.
- Findlay, J.M. & Harris, L.R. (in preparation), Saccadic eye movements to single and double targets.
- Fisher, D.F. (1975). Reading and visual search. *Memory and Cognition*, 3, 188-196.

- Flowers, J.H. & Lohr, D.L. (1985). How does familiarity affect visual search for letter strings? *Perception and Psychophysics*, *37*, 557-566.
- Forster, K.I. (1981). Priming and the effects of sentence and lexical contexts on naming time: Evidence for autonomous lexical processing. *Quarterly Journal of Experimental Psychology*, *33A*, 465-496.
- Forster, K.I. (1987). Form-priming with masked primes: The best match hypothesis. In M. Coltheart (Ed.), *Attention and Performance, XII*. London: Erlbaum.
- Forster, K.I. & Chambers, S.M. (1973). Lexical access and naming time. *Journal of Verbal Learning and Verbal Behavior*, *12*, 627-635.
- Forster, K.I. & Davis, C. (1984). Repetition priming and frequency attenuation in lexical access. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 680-698.
- Forster, K.I., Davis, C., Schoknecht, C., & Carter, R. (1987). Masked priming with graphemically related forms: Repetition or partial activation? *Quarterly Journal of Experimental Psychology*, *39A*, 211-252.
- Fowler, C.A., Napps, S.E., & Feldman, L.B. (1985). Relations among regular and irregular morphologically related words in the lexicon as revealed by repetition priming. *Memory and Cognition*, *13*, 241-255.
- Franks, J.J. & Bransford, J.D. (1971). Abstraction of visual patterns. *Journal of Experimental Psychology*, *90*, 65-74.
- Frederiksen, J.R. & Kroll, J.F. (1976). Spelling and sound: Approaches to the internal lexicon. *Journal of Experimental Psychology: Human Perception and Performance*, *2*, 361-379.
- 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.
- Fukushima, K. & Miyake, S. (1982). Neocognitron: A new algorithm for pattern recognition tolerant of deformations and shifts in position. *Pattern Recognition*, *15*, 455-469.
- Ganz, L. (1975). Temporal factors in visual perception. In E.C. Carterette & M.P. Friedman (Eds), *Handbook of perception: Vol. V. Seeing*. New York: Academic Press.

- Gernsbacher, M.A. (1984). Resolving 20 years of inconsistent interactions between lexical familiarity and orthography, concreteness, and polysemy. *Journal of Experimental Psychology: General*, *113*, 256-281.
- Gibson, E.J. (1971). Perceptual learning and the theory of word perception. *Cognitive Psychology*, *2*, 351-368.
- Gibson, J.J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- Gilliatt, R.W. & Pratt, R.T.C. (1952). Disorders of perception and performance in a case of right-sided cerebral thrombosis. *Journal of Neurology Neurosurgery and Psychiatry*, *15*, 264-271.
- Glushko, R.J. (1979). The organisation and activation of orthographic knowledge in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance*, *5*, 674-691.
- Golden, R.M. (1986). A developmental neural model of visual word perception. *Cognitive Science*, *10*, 241-276.
- Gordon, B. (1983). Lexical access and lexical decision: Mechanisms of frequency selectivity. *Journal of Verbal Learning and Verbal Behavior*, *22*, 24-44.
- 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.
- Grainger, J., O'Regan, J.K., Jacobs, A.M., & Segui, J. (1989). On the role of competing word units in visual word recognition: The neighbourhood frequency effect. *Perception and Psychophysics*, *45*, 189-195.
- Greenberg, S.N. & Vellutino, F.R. (1988). Evidence for processing of constituent single- and multi-letter codes: Support for multilevel coding in word perception. *Memory and Cognition*, *16*, 54-63.
- Grice, G.R., Canham, L., & Boroughs, J.M. (1983). Forest before trees? It depends where you look. *Perception and Psychophysics*, *33*, 121-128.
- Gross, C.G., RochaMiranda, E.E., & Bender, D.B. (1972). Visual properties of neurons in inferotemporal cortex of the macaque. *Journal of Neurophysiology*, *35*, 96-111.

- Gunther, H., Gfroerer, S., & Weiss, L. (1984). Inflection, frequency, and the word superiority effect. *Psychological Research*, 46, 261-281.
- Hammond, E.J., & Green, D.W. (1982). Detecting targets in letter and non-letter arrays. *Canadian Journal of Psychology*, 36, 67-82.
- Harris, M. & Coltheart, M. (1986). *Language Processing in Children and Adults*. Routledge and Kegan Paul Ltd.
- Havens, L.L. & Foote, W.E. (1963). The effect of competition on visual duration threshold and its independence of stimulus frequency. *Journal of Experimental Psychology*, 65, 6-11.
- Hawkins, H.L., Reicher, G.M., Rogers, M., & Peterson, L. (1976). Flexible coding in word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 380-385.
- Hayes-Roth, B. & Hayes-Roth, F. (1977). Concept learning and the recognition and classification of exemplars. *Journal of Verbal Learning and Verbal Behavior*, 16, 321-338.
- Healy, A.F. (1976). Detection errors on the word *the*: Evidence for reading units larger than letters. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 235-242.
- Healy, A.F. (1980). Proofreading errors on the word *the*: New evidence on reading units. *Journal of Experimental Psychology: Human Perception and Performance*, 6, 45-57.
- Henderson, J.M. & Ferreira, F. (1990). Effects of foveal processing difficulty on the perceptual span in reading: Implications for attention and eye movement control. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 417-429.
- Hillinger, M.L. (1980). Priming effects with phonemically similar words: The encoding-bias hypothesis reconsidered. *Memory and Cognition*, 8, 115-123.
- Hinton, G.E. (1979). Some demonstrations of the effects of structural descriptions in mental imagery. *Cognitive Science*, 3, 231-250.
- Hinton, G.E. (1981a). Shape representations in parallel systems. In vol. 2, *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, Vancouver, Canada. Los Altos: Kaufman.

- Hinton, G.E. (1981b). A parallel computation that assigns canonical object-based frames of reference. In vol. 2, *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, Vancouver, Canada. Los Altos: Kaufman.
- Hinton, G.E. (1981c). The role of spatial working memory in shape perception. In *Proceedings of the Third Annual Conference of the Cognitive Science Society*.
- Hinton, G.E. & Parsons, L.M. (1981). Frames of reference and mental imagery. In A.D. Baddeley & J. Long (Eds) *Attention and Performance, IX*, Hillsdale, N.J.: Erlbaum.
- Hinton, G.E. & Lang, K.J. (1985). Shape recognition and illusory conjunctions. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 252-259.
- Hinton, G.E. & Plaut, D.C. (1987). Using fast weights to deblur old memories. *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*.
- Hinton, G.E. & Shallice, T. (1989). *Lesioning a connectionist network: Investigations of acquired dyslexia*. University of Toronto, technical report: CRGTR-89-3.
- Hochberg, J. (1968). Parts and wholes: A response to Arnheim. *New Ideas in Psychology*, 4, 285-293.
- Hock, H.S. & Tromley, C.L. (1978). Mental rotation and perceptual uprightness. *Perception and Psychophysics*, 24, 529-533.
- Hoffman, D.D. & Richards, W.A. (1984). Parts of recognition. *Cognition* 18, 65-96.
- Holender, D. (1979). Identification of letters in words and of single letters with pre- and postknowledge vs. postknowledge of the alternatives. *Perception and Psychophysics*, 25, 313-318.
- Homa, D., Haver, B., & Schwartz, T. (1976). Perceptibility of schematic face stimuli: Evidence for a perceptual Gestalt. *Memory and Cognition*, 4, 176-185.
- Howard, D. (1987). Reading without letters? In M. Coltheart, G. Sartori, & R. Job (Eds), *The cognitive neuropsychology of language*. London: Erlbaum.
- Humphreys, G.W. & Bruce, V. (1989). *Visual cognition: Computational, experimental, and neuropsychological perspectives*. London: Erlbaum.



- Humphreys, G.W., Besner, D. & Quinlan, P.T. (1988). Event perception and the word repetition effect. *Journal of Experimental Psychology: General*, *117*, 51-67.
- Humphreys, G.W., Evett, L.J., & Quinlan, P.T. (1990). Orthographic processing in visual word identification. *Cognitive Psychology*, *22*, 517-560.
- Humphreys, G.W., Evett, L.J., Quinlan, P.T., & Besner, D. (1987). Orthographic priming: Qualitative differences between priming from identified and unidentified primes. In M. Coltheart (Ed.), *Attention and Performance, XII*. London: Erlbaum.
- Humphreys, G.W., Evett, L.J. & Taylor, D.E. (1982). Automatic phonological priming in visual word recognition. *Memory and Cognition*, *10*, 576-590.
- Humphreys, G.W. & Riddoch, M.J. (1984). Routes to object constancy: Implications from neurological impairments of object constancy. *Quarterly Journal of Experimental Psychology*, *36A*, 385-415.
- Humphreys, G.W. & Riddoch, M.J. (1987). *To see but not to see: A case study of visual agnosia*. London: Erlbaum.
- Intraub, H. (1985). Visual dissociation: An illusory conjunction of pictures and form. *Journal of Experimental Psychology: Human Perception and Performance*, *11*, 431-442.
- Jacoby, L.L. (1983). Perceptual enhancement: Persistent effects of an experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*, 21-38.
- Jahnke, J.C., Davis, S.T., & Bower, R.E. (1989). Position and order information in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 859-867.
- Johnson, N.F. (1975). On the function of letters in word identification: Some data and a preliminary model. *Journal of Verbal Learning and Verbal Behavior*, *14*, 17-29.
- Johnson, N.F. (1977). A pattern-unit model of word identification. In D. LaBerge and S.J. Samuels (Eds), *Basic processes in reading: Perception and comprehension*. Hillsdale, N.J.: Erlbaum.
- Johnson, N.F. (1986). On looking at letters within words: Do we "see" them in memory? *Journal of Memory and Language*, *25*, 558-570.

- Johnson, N.F. & Marmurek, H.C.C. (1978). Identification of words and letters within words. *American Journal of Psychology*, *91*, 401-415.
- Johnson, N.F., Turner-Lyga, M., & Pettegrew, B.S. (1986). Part-whole relationships in the processing of small visual patterns. *Memory and Cognition*, *14*, 5-16.
- Johnston, J.C. (1978). A test of the sophisticated guessing theory of word perception. *Cognitive Psychology*, *10*, 123-154.
- Johnston, J.C. (1981). The effects of advance precuing of alternatives on the perception of letters alone and in words. *Journal of Experimental Psychology: Human Perception and Performance*, *7*, 560-572.
- Johnston, J.C. & McClelland, J.L. (1973). Visual factors in word perception. *Perception and Psychophysics*, *14*, 365-370.
- Johnston, J.C. & McClelland, J.L. (1974). Perception of letters in words: Seek not and ye shall find. *Science*, *184*, 1192-1194.
- Johnston, J.C. & McClelland, J.L. (1980). Experimental tests of a hierarchical model of word recognition. *Journal of Verbal Learning and Verbal Behavior*, *19*, 503-524.
- Johnston, J.C., Hale, B.L., & Van Santen, J.P.H. (1983). *Resolving letter position uncertainty in words*. Technical report TM-83-11221-19. Murray Hill, NJ: Bell Labs.
- Just, M., & Carpenter, P.A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review*, *87*, 329-354.
- Kanwisher, N.G. & Potter, M.C. (1990). Repetition blindness: Levels of processing. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 30-47.
- Kay, J. & Bishop, D. (1987). Anatomical differences between nose, palm, and foot, or, the body in question: Further dissection of the processes of sub-lexical spelling-sound translation. In M. Coltheart (Ed.), *Attention and Performance, XII*. London: Erlbaum.
- Kay, J. & Marcel, A.J. (1981). One process, not two, in reading aloud: Lexical analogies do the work of non-lexical rules. *Quarterly Journal of Experimental Psychology*, *33A*, 397-414.
- Kimchi, R. & Palmer, S.E. (1982). Form and texture in hierarchically constructed patterns. *Journal of Experimental Psychology: Human Perception and Performance*, *8*, 521-535.

- Kimchi, R. & Palmer, S.E. (1985). Separability and integrality of global and local levels of hierarchical patterns. *Journal of Experimental Psychology: Human Perception and Performance*, 11, 673-688.
- Kinchla, R.A., & Wolfe, J.M. (1979). The order of visual processing: "Top-down," "bottom-up" or "middle-out". *Perception and Psychophysics*, 25, 225-231.
- Kinsbourne, M. & Warrington, E.K. (1962). A variety of reading disability associated with right hemisphere lesions. *Journal of Neurology, Neurosurgery and Psychiatry*, 25, 339-344.
- Knapp, A.G. & Anderson, J.A. (1984). Theory of categorization based on distributed memory storage. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 616-637.
- Koriat, A. & Norman, J. (1984). What is rotated in mental rotation? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 421-434.
- Koriat, A. & Norman, J. (1985). Reading rotated words. *Journal of Experimental Psychology: Human Perception and Performance*, 11, 490-508.
- Koriat, A. & Norman, J. (1989). Why is word recognition impaired by disorientation while the identification of single letters is not? *Journal of Experimental Psychology: Human Perception and Performance*, 15, 153-163.
- 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.
- Kroll, N.E., Parks, T., Parkinson, S.R., Bieber, S.L., & Johnson, A.L. (1970). Short-term memory while shadowing: Recall of visually and aurally presented letters. *Journal of Experimental Psychology*, 85, 220-224.
- Krueger, L.E. (1970). Search time in a redundant visual display. *Journal of Experimental Psychology*, 83, 391-399.
- LaBerge, D. (1983). Spatial extent of attention to letters and words. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 371-379.
- LaBerge, D. and Samuels, S.J. (1974). Toward a theory of automatic information processing in reading. *Cognitive Psychology*, 6, 293-323.

- Lacouture, Y. (1989). From mean square error to reaction time: A connectionist model of word recognition. In D. Touretzky, G.E. Hinton, T. Sejnowski (Eds.), *Proceedings of the 1988 Connectionist Models Summer School*. Morgan Kaufman.
- Lima, S.D. (1987). Morphological analysis in sentence reading. *Journal of Memory and Language*, 26, 84-99.
- Long, G.M. (1980). Iconic memory: A review and critique of the study of short term visual storage. *Psychological Bulletin*, 88, 785-820.
- von der Malsburg, C. (1981). *The correlation theory of brain function*. Internal report 81-2, Max-Planck-Institute for Biophysical Chemistry.
- von der Malsburg, C. (1985). Nervous structures with dynamic links. *Ber. Bunsenges Phys. Chem.*, 89, 703-710.
- von der Malsburg, C. & Bienenstock, E. (1986). Statistical coding and short-term synaptic plasticity: A scheme for knowledge representation in the brain. In E. Bienenstock et al. (Eds.), *Disordered systems and biological organization*. Berlin: Springer-Verlag.
- Manso de Zuniga, C.M. (1988). *The effects of handwriting and repetition on visual word recognition*. Unpublished Ph.D. thesis, London University.
- Manso de Zuniga, C.M., Quinlan, P.T., & Humphreys, G.W. (1988). *Task constraints on priming with masked primes*. Paper presented to the Experimental Psychology Society, Reading, March.
- Marmurek, H.C.C. (1977). Processing letters in words at different levels. *Memory and Cognition*, 5, 67-72.
- Marmurek, H.C.C. (1986). Whole and part comparisons of words and nonwords. *Memory and Cognition*, 14, 113-123.
- Marr, D. (1977). Analysis of occluding contour. *Proceedings of the Royal Society of London*, B197, 441-475.
- Marr, D. (1982). *Vision*. San Francisco: W.H. Freeman.
- 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.

- Mason, M. (1975). Reading ability and letter search time: Effects of orthographic structure defined by single-letter positional frequency. *Journal of Experimental Psychology: General*, 104, 146-166.
- Mason, M. (1980). Reading ability and the encoding of item and location information. *Journal of Experimental Psychology: Human Perception and Performance*, 6, 89-98.
- Mason, M., Pilkington, C., Brandau, R. (1981). From print to sound: Reading ability and order information. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 580-591.
- Massaro, D.W. (1975). Primary and secondary recognition in reading. In D.W Massaro (Ed.), *Understanding language: An information processing analysis of speech perception, reading, and psycholinguistics*. New York: Academic Press.
- Massaro, D.W. & Klitzke, D. (1979). The role of lateral masking and orthographic structure in letter and word recognition. *Acta Psychologica*, 43, 413-426.
- Massaro, D.W., Venezky, R.L., & Taylor, G.A. (1989). Orthographic regularity, positional frequency, and visual processing of letter strings. *Journal of Experimental Psychology: general*, 108, 107-124.
- Masson, M.E.J. & Freedman, L. (1990). Fluent identification of repeated words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 355-373.
- McClelland, J.L.(1976). Preliminary letter recognition in the perception of words and non-words. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 80-91.
- McClelland, J.L.(1978). Perception and masking of wholes and parts. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 210-223.
- McClelland, J.L. (1985). Putting knowledge in its place: A scheme for programming parallel processing structures on the fly. *Cognitive Science*, 9, 113-146.
- McClelland, J.L.(1986). The programmable blackboard model of reading. In J.L. McClelland & D.E. Rumelhart (Eds), *Parallel distributed processing: Explorations in the microstructure of cognition, II*. Cambridge, Mass.: MIT press.
- McClelland, J.L.& Johnston, J.C. (1977). The role of familiar units in perception of words and nonwords. *Perception and Psychophysics*, 22, 249-261.

- McClelland, J.L. & Mozer, M. (1986). Perceptual interactions in two-word displays: Familiarity and similarity effects. *Journal of Experimental Psychology: Human Perception and Performance*, 12, 18-35.
- McClelland, J.L. & Rumelhart, D.E. (1981). An interactive activation model of context effects in letter perception. 1. An account of basic findings. *Psychological Review*, 88, 375-407.
- McClelland, J.L. & Rumelhart, D.E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114, 159-188.
- McConkie, G.W. & Rayner, K. (1975). The span of the effective stimulus during a fixation in reading. *Perception and Psychophysics*, 17, 578-586.
- McConkie, G.W. & Rayner, K. (1976). Asymmetry of the perceptual span in reading. *Bulletin of the Psychonomic Society*, 8, 365-368.
- McConkie, G.W. & Zola, D. (1987). Visual attention during eye fixations while reading. In M. Coltheart (Ed.), *Attention and Performance, XII*. London: Erlbaum.
- McConkie, G.W. & Zola, D., Blanchard, H.E.W., & Wolverton, D.S. (1982). Perceiving words during reading: Lack of facilitation from prior peripheral exposure. *Perception and Psychophysics*, 32, 271-281.
- McCutchen, D. & Perfetti, C.A. (1982). The visual tongue-twister effect: Phonological activation in silent reading. *Journal of Verbal Learning and Verbal Behavior*, 21, 672-687.
- McNicol, D. & Heathcote, A. (1986). Representation of order information: An analysis of grouping effects in short term memory. *Journal of Experimental Psychology: General*, 115, 76-95.
- Merikle, P.M. (1974). Selective backward masking with an unpredictable mask. *Journal of Experimental Psychology*, 103, 589-591.
- Merikle, P.M., & Coltheart, M. (1972). Selective forward masking. *Canadian Journal of Psychology*, 26, 296-279.
- Merikle, P.M., Coltheart, M., & Lowe, D.G. (1971). On the selective effects of a pattern masking stimulus. *Canadian Journal of Psychology*, 25, 264-279.

- Merikle, P.M. & Glick, M.J. (1976). Processing order in visual perception. *Quarterly Journal of Experimental Psychology*, 28, 17-26.
- Mewhort, D.J.K. (1967). Familiarity of letter sequences, response uncertainty, and the tachistoscopic recognition experiment. *Canadian Journal of Psychology*, 21, 309-321.
- Mewhort, D.J.K. & Campbell, A.J. (1978). Processing spatial information and the selective-masking effect. *Perception and Psychophysics*, 24, 93-101.
- Mewhort, D.J.K., Campbell, A.J., Marchetti, F.M., & Campbell, J.I.D. (1981). Identification, localisation and "iconic memory": An evaluation of the bar-probe task. *Memory and Cognition*, 9, 50-67.
- Meyer, D.E., Schvaneveldt, R.W., & Ruddy, M.G. (1974). Functions of graphemic and phonemic codes in word recognition. *Memory and Cognition*, 2, 309-321.
- Michaels, C.F. & Turvey, M.T. (1979). Central sources of masking: Indexing structures supporting seeing at a single, brief glance. *Psychological Research*, 41, 1-61.
- Miller, G.A., Bruner, J.S., & Postman, L. (1954). Familiarity of letter sequences and tachistoscopic identification. *Journal of General Psychology*, 50, 129-139.
- Mishkin, M. (1982). A memory system in the monkey. *Philosophical Transactions of the Royal Society of London*, B298, 85-96.
- Monk, A.F. (1985). Coordinate systems in visual word recognition. *Quarterly Journal of Experimental Psychology*, 37A, 613-625.
- Monk, A.F. & Hulme, C. (1983). Errors in proofreading: Evidence for the use of word shape in word recognition. *Memory and Cognition*, 11, 16-23.
- Monsell, S. (1985). Repetition and the lexicon. In A.W. Ellis (Ed.), *Progress in the psychology of language (Vol. 2)*. London: Erlbaum.
- Morris, R.K., Rayner, K., & Pollatsek, A. (1990). Eye movement guidance in reading: The role of parafoveal letter and space information. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 268-281.
- Morrison, I.R. & Butler, B.E. (1986). A set intersection model of letter recognition in the spatial probe task. *Canadian Journal of Psychology*, 40, 136-160.

- Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, 76, 165-178.
- Morton, J. (1970). A functional model for memory. In D.A. Norman (ed.). *Models of Human Memory*. Academic Press: New York.
- Morton, J. (1979). Some experiments on facilitation in word and picture recognition and their relevance for the evolution of a theoretical position. In P.A. Kolers, M.E. Wrolstad, & H. Bouma (Eds), *Processing of visible language, 1*. New York: Plenum.
- Morton, J. & Patterson, K.E. (1980). A new attempt at an interpretation, or, an attempt at a new intepretation. In M. Coltheart, K.E. Patterson, & J.C. Marshall (Eds), *Deep dyslexia*. London: Routledge & Kegan Paul.
- Mozer, M.C. (1983). Letter migration in word perception. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 531-546.
- Mozer, M.C. (1987). Early parallel processing in reading: A connectionist approach. In M. Coltheart (Ed.), *Attention and Performance, XII*. London: Erlbaum.
- Mozer, M.C. (1988) *A connectionist model of selective attention in visual perception*. Technical report, CRGTR-88-4.
- Mozer, M.C. (1989). Types and tokens in visual letter perception. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 287-303.
- Murdock, B.B. (1983). A distributed memory model for serial-order information. *Psychological Review*, 90, 316-338.
- Murrell, G.A. & Morton, J. (1974). Word recognition and morphemic structure. *Journal of Experimental Psychology*, 102, 963-968.
- Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, 9, 353-383.
- Navon, D. (1978). Perception of misoriented words and letter strings. *Canadian Journal of Psychology*, 32, 129-140.
- Neisser, U. (1967). *Cognitive psychology*. New York: AppletonCenturyCrofts.
- Neisser, V. & Kerr, N. (1973). Spatial and mnemonic properties of visual images. *Cognitive Psychology*, 5, 138-150.



- Nicoletti, R. & Umiltà, C. (1989). Splitting visual space with attention. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 164-169.
- Norris, D. & Brown, G. (1985). Race models and analogy theories: A dead heat? A reply to Seidenberg. *Cognition*, *20*, 155-168.
- O'Connor, R.E. & Forster, K.I. (1981). Criterion bias and search sequence bias in word recognition. *Memory and Cognition*, *9*, 78-92.
- Osgood, C.E. & Hoosain, R. (1974). Salience of the word as a unit in the perception of language. *Perception and Psychophysics*, *15*, 168-192.
- O'Regan, J.K. & Levy-Schoen, A. (1987). Eye movement strategy and tactics in word recognition and reading. In M. Coltheart (Ed.), *Attention and Performance*, *XII*. London: Erlbaum.
- O'Regan, J.K., Levy-Schoen, A., Pynte, J., & Brugailere, B. (1984). Convenient fixation location within isolated words of different lengths and structures. *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 250-257.
- Paap, K.R., Newsome, S.L., McDonald, J.E., & Schvaneveldt, R.W. (1982). An activation-verification model for letter and word recognition: The word superiority effect. *Psychological Review*, *89*, 573-594.
- Palmer, S.E. (1975a). The effects of contextual scenes on the identification of objects. *Memory and Cognition*, *3*, 519-526.
- Palmer, S.E. (1975b). Visual perception and world knowledge. Notes on a model of sensory-cognitive interaction. In D.A. Norman, D.E. Rumelhart, & the LNR Research Group (Eds), *Explorations in cognition*. San Francisco: W.H. Freeman.
- Parkin, A.J. (1982). Phonological recoding in lexical decision: Effects of spelling-to-sound regularity depend on how regularity is defined. *Memory and Cognition*, *10*, 43-53.
- Parkin, A.J. (1984). Redefining the regularity effect. *Memory and Cognition*, *12*, 287-292.
- Parkin, A.J. & Underwood, G. (1983). Orthographic vs. phonological irregularity in lexical decision. *Memory and Cognition*, *11*, 351-355.

- Parks, T.E., Kroll, N.E., Salzberg, P.M., & Parkinson, S.R. (1972). Persistence of visual memory as indicated by decision time in a matching task. *Journal of Experimental Psychology*, *92*, 437-438.
- Patterson, K.E. & Kay, J. (1982). Letter-by-letter reading; psychological descriptions of a neurological syndrome. *Quarterly Journal of Experimental Psychology*, *34A*, 411-441.
- Patterson, K.E. & Morton, J. (1985). From orthography to phonology: An attempt at an old interpretation. In K.E. Patterson, J.C. Marshall, & M. Coltheart (Eds), *Surface dyslexia: Neuropsychological and cognitive studies of phonological reading*. London: Erlbaum.
- Patterson, K.E. & Wilson, B. (1988). *A rose is a nose: A deficit in initial letter identification*. Paper presented to the Experimental Psychology Society, London.
- Pentland, A. (1986). Perceptual organisation and the representation of natural form. *Artificial Intelligence*, *28*, 293-331.
- Phillips, W.A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception and Psychophysics*, *16*, 283-290.
- Phillips, W.A. (1983). Short-term visual memory. *Philosophical Transactions of the Royal Society of London*, *B302*, 295-309.
- Phillips, W.A., Hobbs, S.B., & Pratt, F.R. (1978). Intellectual realism in children's drawings of cubes. *Cognition*, *6*, 15-33.
- Phillips, W.A., Hancock, P.J.B., Wilson, N.J., & Smith, L.S. (1989). On the acquisition of object concepts from sensory data. In R. Eckmiller & C. von der Malsburg (Eds.), *Neural computers*. London: Springer-Verlag.
- Pike, R. (1984). Comparison of convolution and matrix distributed memory systems for associative recall and recognition. *Psychological Review*, *91*, 281-294.
- Pollatsek, A., Well, A.D., & Schindler, R.M. (1975). Familiarity affects visual processing of words. *Journal of Experimental Psychology: Human Perception and Performance*, *1*, 328-338.
- Pollatsek, A., Bolozky, S., Well, A.D., & Rayner, K. (1981). Asymmetries in the perceptual span for Israeli readers. *Brain and Language*, *14*, 174-180.

- Posner, M.I. & Keele, S.W. (1967). Decay of visual information from a single letter. *Science*, *158*, 137-139.
- Posner, M.I. & Keele, S.W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353-363.
- Posner, M.I. & Keele, S.W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology*, *83*, 304-308.
- Posner, M.I., Boies, S.J., Eichelman, W.H., & Taylor, R.L. (1969). Retention of name and visual codes of single letters. *Journal of Experimental Psychology*, *79*, 1-16.
- Posner, M.I., Rafal, R.D., Choate, L.S., & Vaughan, J. (1985). Inhibition of return: Neural basis and function. *Cognitive Neuropsychology*, *2*, 211-228.
- Proctor, R.W. & Healy, A.F. (1985). Order-relevant and order-irrelevant decision rules in multiletter matching. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *11*, 519-537.
- Purcell, D.G. & Stewart, A.L. (1988). The face-detection effect: Configuration enhances detection. *Perception and Psychophysics*, *43*, 355-366.
- Ratcliff, R. (1981). A theory of order relations in perceptual matching. *Psychological Review*, *88*, 552-572.
- Ratcliff, R. (1987). Order information and distributed memory models. *Proceedings of the Cognitive Science Society*, 474-486.
- Ratcliff, R. & Hacker, M.J. (1985). Speed and accuracy of same and different responses in perceptual matching. *Perception and Psychophysics*, *30*, 303-307.
- Ratcliff, R., McKoon, G., & Verwoerd, M. (1989). A bias interpretation of facilitation in perceptual identification. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *15*, 378-387.
- Rayner, K. (1975). The perceptual span and peripheral cues in reading. *Cognitive Psychology*, *7*, 65-81.
- Rayner, K. & Pollatsek, A. (1987). Eye movements during reading: A tutorial review. In M. Coltheart (Ed.), *Attention and Performance, XII*. London: Erlbaum.

- Rayner, K., McConkie, G.W., & Ehrlich, S. (1978). Eye movements and integrating information across fixations. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 529-544.
- Rayner, K., McConkie, G.W., & Zola, D. (1980). Integrating information across eye movements. *Cognitive Psychology*, 12, 206-226.
- Rayner, K., Well, A.D., Pollatsek, A., & Bertera, J.H. (1982). The availability of useful information to the right of fixation in reading. *Perception and Psychophysics*, 31, 537-550.
- Rayner, K., Inhoff, A.W., Morrison, R.E., Slowiaczek, M.L., & Bertera, J.H. (1981). Masking of foveal and parafoveal vision during eye fixations in reading. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 167-179.
- Reicher, G.M. (1969). Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology*, 81, 275-280.
- Richman, H.B., and Simon, H.A. (1989). Context effects in letter perception: Comparison of two theories. *Psychological Review*, 96, 417-432.
- Rieser, J.J. (1989). Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 15, 1157-1165.
- Rock, I. (1973). *Orientation and form*. New York: Academic Press.
- Rock, I. (1983). *The logic of perception*. Cambridge Mass: Bradford MIT Press.
- Rosson, M.B. (1983). From SOFA to LOUCH: Lexical contributions to pseudoword pronunciation. *Memory and Cognition*, 11, 152-160.
- Rosson, M.B. (1985). The interaction of pronunciation rules and lexical representations in reading aloud. *Memory and Cognition*, 13, 90-99.
- Rueckl, J.G. (1990). Similarity effects in word and pseudoword repetition priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 374-391.
- Rumelhart, D.E. (1977). Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and Performance, VI*. Hillsdale, N.J.: Erlbaum.

- Rumelhart, D.E. & McClelland, J.L. (1982). An interactive activation model of context effects in letter perception. 2. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, 89, 60-94.
- Ryan, J. (1969). Temporal grouping, rehearsal and short-term memory. *Quarterly Journal of Experimental Psychology*, 21, 148-155.
- Samuel, A.G., van Santen, J.P.H., & Johnston, J.C. (1982). Length effects in word perception: We is better than I but worse than you or them. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 91-105.
- Santee, C. & Egeth, H.E. (1980). Selective attention in speeded classification and comparison of multidimensional stimuli. *Perception and Psychophysics*, 28, 191-204.
- Scheerer, E. (1987). Visual word recognition in German. In D.A. Allport, D. Mackay, W. Prinz, & E. Scheerer (Eds); *Language perception and production: Shared mechanisms in listening, speaking, reading, and writing*. London: Academic Press.
- Segui, J. & Grainger, J. (1990). Priming word recognition with orthographic neighbors: Effects of relative prime-target frequency. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 65-76.
- Seidenberg, M.S. (1985). The time course of information activation and utilization in visual word recognition. In D. Besner, T.G Waller, & G.E. MacKinnon (Eds), *Reading research: Advances in theory and practice*, 5. New York: Academic Press.
- Seidenberg, M.S. & McClelland, J.L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96, 523-568.
- Seidenberg, M.S., Waters, G.S., Barnes, M.A., & Tanenhaus, M.K. (1984). When does irregular spelling or pronunciation influence word recognition? *Journal of Verbal Learning and Verbal Behavior*, 23, 383-404.
- Sejnowski, T.J. & Rosenberg, C.R. (1987). Parallel networks that learn to pronounce English text. *Complex Systems*, 1, 145-168.
- Sergent, J. (1984). An investigation into component and configural processes underlying face perception. *British Journal of Psychology*, 75, 221-242.
- Shallice, T. & McCarthy, R. (1985). Phonological reading: From patterns of impairments to possible procedures. In K.E. Patterson, J.C. Marshall, & M. Coltheart (Eds), *Surface*

- dyslexia: Neuropsychological and cognitive studies of phonological reading*. London: Erlbaum.
- Shallice, T. & McGill, J. (1978). The origin of mixed errors. In J. Requin (Ed.), *Attention and Performance, VII*, Hillsdale, N.J.: Erlbaum.
- Shallice, T. & Saffran, E. (1986). Lexical processing in the absence of explicit word identification: Evidence from a letter-by-letter reader. *Cognitive Neuropsychology*, 3, 429-458.
- Shallice, T. & Warrington, E.K. (1977). The possible role of selective attention in acquired dyslexia. *Neuropsychologia*, 15, 31-41.
- Shallice, T. & Warrington, E.K. (1980). Single and multiple component central dyslexia syndromes. In M. Coltheart, K.E. Patterson, & J.C. Marshall (Eds), *Deep dyslexia*. London: Routledge & Kegan Paul.
- Shallice, T., Warrington, E.K., & McCarthy, R. (1983). Reading without semantics. *Quarterly Journal of Experimental Psychology*, 35A, 111-138.
- Shaw, P. (1969). Processing of tachistoscopic displays with controlled order of characters and spaces. *Perception and Psychophysics*, 6, 257-266.
- Shulman, G.L. & Wilson, J. (1987a). Spatial frequency and selective attention to local and global information. *Perception*, 16, 89-101.
- Shulman, G.L. & Wilson, J. (1987b). Spatial frequency and spatial attention. *Perception*, 16, 103-111.
- Shulman, H.G., Hornak, R., & Sanden, E. (1978). The effects of graphemic, phonetic, and semantic relationships on access to lexical structures. *Memory and Cognition*, 6, 115-123.
- Sieroff, E. & Michel, F. (1987). Verbal visual extinction in right/left hemisphere lesion patients and the problem of lexical access. *Neuropsychologia*, 25, 907-918.
- Sieroff, E., Pollatsek, A., & Posner, M.I. (1988). Recognition of visual letter strings following injury to the posterior visual spatial attention system. *Cognitive Neuropsychology*, 5, 427-449.

- Sloboda, J.A. (1976). Decision times for word and letter search: A wholistic word identification model examined. *Journal of Verbal Learning and Verbal Behavior*, 15, 93-101.
- Sloboda, J.A. (1977). The locus of the word priority effect in a target detection task. *Memory and Cognition*, 5, 371-376.
- Smith, F. (1969). Familiarity of configuration vs. discriminability of features in the visual identification of words. *Psychonomic Science*, 14, 261-262.
- Smith, F. (1971). *Understanding Reading*. New York: Holt, Rinehart, & Winston.
- Smith, E.E. & Haviland, S.E. (1972). Why words are perceived more accurately than nonwords: Inference vs. unitization. *Journal of Experimental Psychology*, 92, 59-64.
- Smith, E.E. & Spoehr, K.T. (1974). The perception of printed English: A theoretical perspective. In B.H. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition*. Hillsdale, N.J.: Erlbaum.
- Smith, P.T. & Groat, A. (1979). Spelling patterns, letter cancellation and the processing of text. In P.A. Kolars, M.E. Wrolstad and H. Bouma (Eds), *Processing of visible language (Vol. 1)*. New York: Plenum Press.
- Smith, P.T., Meredith, T., Pattison, H.M., & Sterling, C. (1984). The representation of internal word structure in English. In L. Henderson (Ed.), *Orthographies and reading*. London: Erlbaum.
- Spector, A. & Purcell, D.G. (1977). The word-superiority effect: A comparison between restricted and unrestricted alternative set. *Perception and Psychophysics*, 21, 323-328.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, 74 (11, Whole no. 498).
- Sperling, G. & Melchner, M.J. (1976). Estimating item and order information. *Journal of Mathematical Psychology*, 2, 192-221.
- Spoehr, K.T. & Smith, E.E. (1973). The role of syllables in perceptual processing. *Cognitive Psychology*, 5, 71-89.

- Spoehr, K.T. & Smith, E.E. (1975). The role of orthographic and phonotactic rules in perceiving letter patterns. *Journal of Experimental Psychology: Human Perception and Performance*, 104, 21-34.
- Staller, J.D. & Lappin., J.S. (1981). Visual detection of multi-letter patterns. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 1258-1272.
- Stanners, R.F., Neiser, J.J., Herson, W.P., & Hall, R. (1979). Memory representations for morphologically related words. *Journal of Verbal Learning and Verbal Behavior*, 18, 399-412.
- Strain, E. & Cowie, R. (1989). *The effects of meaning and display complexity on migration errors*. Unpublished manuscript.
- Taft, M. (1979). Recognition of affixed words and the word frequency effect. *Memory and Cognition*, 7, 263-272.
- Taft, M. (1987). Morphographic processing: The BOSS re-emerges. In M. Coltheart (Ed.), *Attention and Performance, XII*. London: Erlbaum.
- Taft, M. & Forster, K.I. (1976). Lexical storage and retrieval of polymorphemic and polysyllable words. *Journal of Verbal Learning and Verbal Behavior*, 15, 607-620.
- Taraban, R. & McClelland, J.L. (1987). Conspiracy effects in word pronunciation. *Journal of Memory and Language*, 26, 608-631.
- Taylor, S.G. & Brown, D.R. (1972). Lateral visual masking: Supraretinal effects when viewing linear arrays with unlimited viewing time. *Perception and Psychophysics*, 12, 97-99.
- Taylor, D.A. & Chabot, R.J. (1978). Differential backward masking of words and letters by masks of varying orthographic structure. *Memory and Cognition*, 6, 629-635.
- Taylor, S., Frackenpohl, H., & Pettee, J. (1960). Grade level norms for the components of the fundamental reading skill. *Bulletin No. 3*. Huntingdon, N.J.: Educational Developmental Laboratories, Inc.
- Taylor, G.A., Miller, T.J., & Juola, J.F. (1977). Isolating visual units in the perception of words and nonwords. *Perception and Psychophysics*, 21, 377-386.



- Theios, J. & Muise, J.G. (1977). The word identification process in reading. In N.J. Castellan, D.B. Pisoni, & G.R. Potts (Eds), *Cognitive theory*. Hillsdale, N.J.: Erlbaum.
- Thompson, M.C. & Massaro, D.W. (1973). Visual information and redundancy in reading. *Journal of Experimental Psychology*, 98, 49-54.
- Townsend, V.M. (1973). Loss of spatial and identity information following a tachistoscopic exposure. *Journal of Experimental Psychology*, 98, 113-118.
- Townsend, J.J., Taylor, S.G., & Brown, D.R. (1971). Lateral masking for letters with unlimited viewing time. *Perception and Psychophysics*, 10, 375-378.
- Treisman, A. & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, 14, 107-141.
- Treisman, A. & Souther, J. (1986). Illusory words: The roles of attention and of top-down constraints in conjoining letters to form words. *Journal of Experimental Psychology: Human Perception and Performance*, 12, 3-17.
- Turvey, M.T. (1973). On peripheral and central processes in vision: Inferences from an information processing analysis of masking with patterned stimuli. *Psychological Review*, 80, 1-52.
- Van Lancker, D. (1990). *Reading and writing without letters: A case of deep dysgraphia attributed to right hemisphere function*. Paper presented at INS Innsbruck.
- Van Orden, G.C. (1987). A ROWS is a ROSE: Spelling, sound, and reading. *Memory and Cognition*, 15, 181-198.
- Van der Velde, F., van der Heijen, A.H.C., & Schreuder, R. (1989). Context-dependent migrations in visual word perception. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 133-141.
- Warren, C.E.J. & Morton, J. (1982). The effects of priming on picture recognition. *British Journal of Psychology*, 73, 117-130.
- Warrington, E.K. & Shallice, T. (1980). Word-form dyslexia. *Brain*, 103, 99-112.
- Warrington, E.K. & Taylor, A.M. (1973). The contribution of the right parietal lobe to object recognition. *Cortex*, 9, 152-164.

- Warrington, E.K. & Taylor, A.M. (1978). Two categorical stages of object recognition. *Perception*, 7, 695-705.
- Waters, G.S. & Seidenberg, M.S. (1985). Spelling-sound effects in reading: Time course and decision criteria. *Memory and Cognition*, 13, 557-572.
- Watt, R.J. (1987). Scanning from coarse to fine spatial scales in the human vision system after the onset of a stimulus. *Journal of the Optical Society of America*, 4A, 2006-2021.
- Watt, R.J. (1988). *Visual processing: Computational, psychophysical and cognitive research*. London: Lawrence Erlbaum Associates Ltd.
- Watt, R.J. (1990, in press). Visual analysis and representation of spatial relations. *Mind and Language*, 5.
- Watt, R.J. & Morgan, M.J. (1985). A theory of the primitive spatial code in human vision. *Vision Research*, 24, 1387-1397.
- Watt, R.J., Bock, J., Thimbleby, H., & Wilkins, A. (submitted) *Visible aspects of text*.
- Weisstein, N. & Harris, C.S. (1974). Visual detection of line segments: An object-superiority effect. *Science*, 186, 752-755.
- Whaley, C.P. (1978). Word-nonword classification time. *Journal of Verbal Learning and Verbal Behavior*, 17, 143-154.
- Wheeler, D.D. (1970). Processes in word recognition. *Cognitive Psychology*, 1, 59-85.
- White, M.J. (1980). Naming and categorization of tilted alphanumeric characters do not require mental rotation. *Bulletin of Psychonomic Society*, 15, 153-156.
- Whittlesea, B.W.A. (1987). Preservation of specific experiences in the representation of general knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 3-17.
- Wickelgren, W.A. (1967). Rehearsal, grouping and hierarchical organisation of serial position curves. *Quarterly Journal of Experimental Psychology*, 19, 97-102.
- Wickelgren, W.A. (1969). Context-sensitive coding, associative memory and serial order in (speech) behavior. *Psychological Review*, 76, 1-15.

- Winnick, W.A. & Bruder, G.E. (1968). Signal detection approach to the study of retinal locus in tachistoscopic recognition. *Journal of Experimental Psychology*, 78, 528-531.
- Wolford, G.L. & Hollingsworth, S. (1974). Lateral masking in visual information processing. *Perception and Psychophysics*, 16, 315-320.
- Woodworth, R.S. (1938). *Experimental psychology*. New York: Holt, Rinehart & Winston.
- Wurtz, R.H., Goldberg, M.E., & Robinson, D.L. (1982). Brain mechanisms of visual attention. *Scientific American*, 246, 124-135.
- Young, A.W. & Ellis, A.W. (1985). Different methods of lexical access for words presented in the left and right visual hemifields. *Brain and Language*, 24, 326-58.

**APPENDIX A****Examples of stimulus set used in Experiment 1***Prime words*

QFHBMQ

RDUSPL

ZSHRTF

EFRODM

OBMQNZ

YHOQKU

GJOEMA

AJBPSO

IQRTL B

VNTWRE

**APPENDIX B****Examples of stimulus set used in Experiment 2***Prime words*

DCRJNT

CLYBFZ

MBVNDS

QNWDRK

RMYKQD

HMNWX P

VPLTZR

KRDWNQ

LRVXJK

NDMVGC

## APPENDIX C

## Examples of stimulus set used in Experiment 3

<i>Local regularities</i>	<i>Non-local regularities</i>
NHDASB	PSGOXM
NHDYDK	PLGSXU
NHDZAL	PZGTXP
NHDCKY	PKGSXO
NHDGXU	PMGNXO
GLQKFM	CGLDGQ
VLUKFM	CGSDAQ
HOCKFM	IGLDZQ
WZNKFM	VGODLQ
ZXQKFM	RGMDMQ

## APPENDIX D

## Examples of stimulus set used in Experiment 4

<i>Local prototypes</i>	<i>Non-local prototypes</i>
LQKDZM	KUDOXK
QIGRQL	KNRTPC
SJFGPD	GYYNPS
YEQSNS	ABKHXW
JHMFVC	TNRMZS

## APPENDIX E

## Examples of stimulus set used in Experiment 5

*Prime words*

DJSOVT

JDSWRF

KPFESQ

CXIBGV

NLOESW

## APPENDIX F

## Stimulus set used in Experiment 8

COUNTRY

DUSTBIN

CONIFER

SUBJECT

ANGUISH

WHISPER

SPECIFY

CHARITY

PHOEBUS

CLOSURE

POSTURE

PROBLEM

ROUTINE

HANDFUL

TROUBLE

SPECIFY

WHISTLE

MEDICAL

SYRINGE

STRANGE