

Using Anatomical Information to Enrich the Connectionist Modelling of Normal and Impaired Visual Word Recognition

Richard Shillcock (rcs@cogsci.ed.ac.uk)
Padraic Monaghan (pmon@cogsci.ed.ac.uk)
Centre for Cognitive Science, University of Edinburgh
2 Buccleuch Place, Edinburgh, EH8 9LW, U.K.

Abstract

We argue that the connectionist modelling of visual word recognition can be made more explicit and more accurate by the incorporation of information concerning the initial projection of the visual field to the visual cortex. We show that this initial projection involves the precise splitting of the visual field into two hemifields, even in the case of the foveal projection: when a single word is fixated, that part of the word to the right of the fixation point initially goes to the left hemisphere and that part to the left initially goes to the right hemisphere. We present a number of reasons why this initial splitting should be assumed to persist into the higher cognitive processing concerned with visual word recognition. We explore three different phenomena – the processing priority given to the exterior letters of words, the optimum and preferred viewing positions for word recognition, and the core phenomena of dyslexia – and show that in each case a model based on a split architecture makes the correct predictions and captures the human data in a parsimonious and natural way. We conclude that anatomical information concerning the initial visual projection can enrich the modelling of visual word recognition.

Modelling and Anatomical Information

The relative success of the connectionist cognitive modelling of visual word recognition partly reflects the fact that such models embody some of the style of processing of the neuronal level of description: many simple units working in parallel, graded activity, distributed representations. To this extent, these models reflect cortical physiology and anatomy. In this paper we explore what benefits might accrue to the modelling of visual word recognition by incorporating the gross anatomical structure of the visual system – specifically, the precise splitting and contralateral projection of the visual field. Current models of visual word recognition assume that the relevant processing has equal and simultaneous access to all of the visually presented word. We make the case that the initial splitting and contralateral projection of the visual field – including the foveal projection – persists into the higher cognitive processes of word recognition. We explore this conclusion with respect to existing data on normal visual word recognition, and also demonstrate that the impairment of a split connectionist architecture provides a principled and parsimonious explanation of dyslexic data that enriches existing connectionist lesioning accounts.

The Foveal Projection is Split

The visual field is split in its initial projection to the primary visual cortex. This precise splitting is true even of the foveal projection: when an isolated word is read, the left

of the word is initially projected to the right hemisphere (RH) and the right to the left hemisphere (LH). Sugishita, Hamilton, Sakuma, and Hemmi (1994) present compelling data to this effect, with a commissurotomed subject reading single words straddling the fixation point and controlled by eye-tracking equipment. They demonstrate that any sharing of visual information between the initial projection to the two hemispheres is very limited, and adjacent letters may typically be projected to different hemispheres¹.

This fact of peripheral visual processing has no role in conventional modelling of visual word recognition as it is assumed either that this initial foveal splitting does not happen or that such splitting is completely transcended by the time that the processing directly relevant to word recognition occurs. See Shillcock & Monaghan (*submitted*) and Brysbaert (1994) for further discussion. Below, we present four arguments to the effect that this conventional view is too strong a position to sustain, and that word recognition is influenced to some degree by foveal splitting.

Intra- and Interhemispheric Processing

(a) Our first argument against ignoring foveal splitting in visual word recognition concerns the nature of information and information-flow in the brain. Recurrent connectivity is pervasive in the visual system, and in the rest of the brain. Visual information does not make its way unidirectionally from more peripheral, cruder representations/areas-of-the-brain to more central, more sophisticated ones, and ultimately to those mediating word recognition. Rather, conscious word recognition depends upon the orchestration of activity in all the relevant areas, including the precisely split primary visual cortex.

Although it is possible to claim that the higher visual areas contain receptive fields that span the two hemifields, and that the primarily homotopic arrangement of callosal neurons can be interpreted as providing each hemisphere with information from the entire visual field (see, e.g., Berlucchi & Antonini, 1990), we claim that conclusions about the functional architecture of visual word recognition cannot be directly drawn from observations of receptive fields alone. For instance, the parietal lobule (in the macaque), concerned with visual attention, has been shown to have generally large and bilateral receptive fields, with an average diameter of 44°. However, the centres of the receptive fields (the retinal location at which peak response occurs) tend to be located in the contralateral visual field (Andersen, Asanua, Essick, & Siegel, 1990). The necessary data on how the functional architecture of word recognition

¹ Note that these arguments apply to each eye.

finds expression in the physiology and anatomy of the human brain is as yet underdetermined, and we argue that the cautious assumption is to accept that the initial splitting of the visual field may not be completely transcended by a single passing of information across the corpus callosum to join the relevant information from the other hemifield. Indeed, a recent Hebbian account suggests that bi-hemispheric cell-assemblies mediate word recognition (Pulvermüller, *submitted*).

(b) Laterality studies constitute the second argument against the current assumption that the two hemispheres delay any and all processing relevant to word recognition until complete and equal sharing of all information has occurred. Laterality studies typically involve the presentation of letters or words to a single hemifield, and many such studies have shown different processing predispositions when presentation is exclusively to one or other hemisphere (see, e.g., Hellige, 1995; Mohr, Pulvermüller, & Zaidel, 1994). These differences would not occur if the two hemispheres postponed lexical processing – letter and word recognition – until visual information in the two hemifields had been completely shared.

(c) In neglect dyslexia, damage to one hemisphere (typically the RH) can cause asymmetric (contralesional) impairment of lexical representations (see, e.g., Ellis, Flude & Young, 1987). The complete, conscious awareness of a single visually presented word seems to require both hemispheres to be intact and may be disrupted in spatially predictable ways by the (often not very specific) impairment of one hemisphere (see also Chatterjee, 1995).

(d) In some cases of neglect, damage to one hemisphere causes contralesional errors in oral spelling (Hillis & Caramazza, 1990). Even the high-level cognitive processing involved in oral spelling, with no external visual stimulus, seems to show spatially predictable asymmetries.

Thus these four different types of data suggest that the initial splitting of the visual field into two hemifields may be perpetuated in some form in the higher cognitive processing immediately relevant to visual word recognition. The information from the ipsilateral hemifield may arrive later, or it may be accorded a lesser status for some part of the processing prior to word recognition. A range of options is possible, short of the complete and equal sharing of information from the two hemifields before any letter or word recognition.

Implications for Word Recognition

Below, we investigate some of the consequences of foveal splitting for modelling visual word recognition. We assume a degree of intrahemispheric processing of the information from the contralateral hemifield. We report three different sets of visual word recognition data and demonstrate that in each case the human data emerges naturally from the requirements of a split processing architecture.

Processing exterior letters of words

The first and last letters of visually presented words seem to be accorded some sort of processing priority. Thus, the “exterior” letters are better than the interior letters at priming the whole word (Forster & Gartlan, 1975) and subjects can

better identify speeded postmasked presentations of pairs of legitimate exterior letters (like *d k*, as in *dark*) compared with non-legitimate pairs (like *z k*) (Jordan, 1990).

To model this behaviour we employ a simple network as shown in Figure 1. The four-letter input words are presented across any and all sets of four adjacent input nodes. Three input positions straddled the “fixation point” between the two hemifields, two fell wholly in one or other half of the model. Individual letters are coded in terms of eight binary units representing visual features. The model was trained on the 60 most frequent four-letter words in the CELEX database, and it successfully learned to reproduce the words at the output units for all five positions in which the word occurred at the input units. The network was judged to have learned the task when mean squared error for each word was less than 0.5. This occurred after approximately 200 epochs of training.

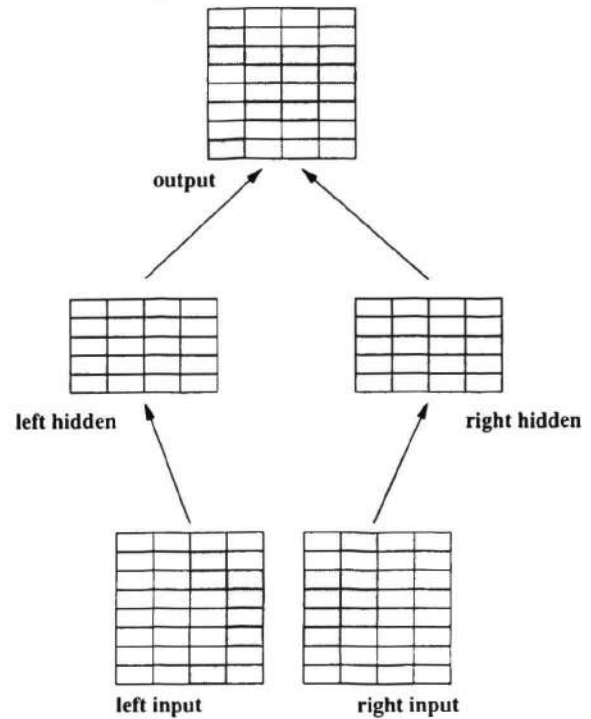


Figure 1: The split model used to simulate the exterior and interior letters effect

The model was then tested using as inputs the exterior and interior letters of the words, respectively, and replacing the missing letters with a mask, produced by setting each input feature to half its maximal value. (In the experiments being simulated, word length is typically conveyed by making the mask exactly overlap the four word positions.)

As Figures 2 and 3 show, the exterior letters effect emerges, whether the nodes measured at output correspond to the whole word ($t(59) = 7.82, p < 0.001$) or just to the letters presented at input ($t(59) = 10.12, p < 0.02$). This robust effect is directly due to the split architecture, as shown by parallel simulations carried out with a non-split architecture with equivalent computational resources (i.e. the same number of units and connections). In this set of simulations, the behaviour of the non-split network did not reliably produce the exterior letters effect. An analysis of variance (ANOVA) comparing the behaviour of the two

networks gave a significant interaction between the factors of architecture (split vs. non-split) and input (exterior vs. interior letters: for the whole word, $F(1, 236) = 28.23, p < 0.001$; for the presented letters, $F(1, 236) = 31.44, p < 0.001$).




Figure 2: Mean error on whole word at output, with presentation of the exterior letters and the interior letters for the split and non-split architectures.




Figure 3: Mean error at output for only the presented letters, for the exterior and interior letters and for the two models.

Subsequent exploration of the behaviour of the two networks showed that the split network invariably gives the exterior letters effect, and that the non-split network only occasionally adopts a strategy of prioritising the processing of the exterior letters. The effect is due to the split disproportionately disrupting attempts to define interior

letters in terms of their immediate contexts. Interestingly, the effect also emerges in our split model for those presentations solely to one half of the model; this behaviour raises a novel principle in theorising about hemispheric interaction, that when a single hemisphere attempts a task, it may be convenient for it to recapitulate the splitting that typically occurs when the two hemispheres collaborate.

We conclude that the exterior letters effect – a phenomenon that has remained largely outside the purview of current models of word recognition – is a processing signature of a split processor.

Optimum and Preferred Viewing Position

Subjects' optimal or preferred viewing position for single words is either in the centre of the word, or slightly to the left of the centre (O'Regan, 1990). This behaviour may be directly observed using eye-tracking, or it may be inferred from the speed and accuracy of word recognition when the word is presented with different letter positions corresponding to the initial fixation point; performance is by definition best at the optimal fixation point. This behaviour is conventionally accounted for by saying that the processor is attempting to position a limited window of high-fidelity processing over the most informative part of the word (see, e.g., Legge, Klitz & Tjan, 1997). We will now explore a different account based on the hypothesis that this fixation behaviour is optimal for a split architecture.

We will first consider a level of representation dealing with whole words. It is a common observation that the first halves of English words tend to be more informative than their last halves (see Yannakoudakis and Hutton (1992) for comparable phonological statistics): syllable onsets contain a greater range of segment combinations than do codas, and (in English) the ends of words often contain suffixes. A simple consequence of this fact is that a processor that was trying to provide each hemisphere with the best information about the identity of the single word straddling its fixation point – simultaneously minimising entropy in each hemisphere – would optimally position its fixation point in the centre of the word or to the left of centre.

We now consider the lower level of individual letter recognition. Letter identity and location must be determined during word recognition. We may reconceptualise this problem as one of how to align the words known to the processor in such a way that the ("vertical") distribution of letters at any one position is such as to minimise entropy at that position. Lower entropy at a certain position (signifying a skewed distribution of letters at that position) is taken to entail less processing, and greater reliance on known probabilities. Figure 4 illustrates three options: left-justified (the strategy suggested by Legge et al., 1997), right-justified (for purposes of comparison), and middle-justified² (the strategy favoured by a split processor). (Note that our model's central split gives us an absolute point around which to organise such justification, allowing us to be less specific about the existence of letter "slots".) The relevant entropies were calculated from all of the words in

² For odd numbers of letters, the midpoint is placed to the left.

CELEX, multiplied by their frequencies, and plotted (see Shillcock & Monaghan, *submitted*).

The plots revealed that there was no difference in total entropy between the left- and middle-justified conditions ($t(19) = -0.787$, *n.s.*), but that both differed significantly from the right-justified condition, which showed greater entropy (respectively, $t(19) = -7.424$, $p < 0.001$; $t(19) = -4.124$, $p < 0.001$). As the plots also made clear, for a split architecture there is a natural advantage in the middle-justified condition in that the distribution of entropy by letter position is symmetrical, compared with the skewed left-justified condition. It is preferable for the labour to be shared by the two hemispheres, and for the positions with highest entropy to fall in the foveal centre, where discrimination is best.

We conclude that, if fixation position is determined so as to facilitate individual letter recognition, there is no inherent advantage in the left-justified strategy suggested by Legge et al. as approximating to optimal behaviour, compared with the middle-justified condition. In a split processor with significant intrahemispheric processing of hemifield-specific information, the observed optimum human fixation behaviour optimises both letter-in-word and word recognition.

(a)

a	s	p	a	r	a	g	u	s		
c	a	r	p	e	t					
f	l	o	o	r						
t	r	a	i	n	s					
p	e	n								
i	t									

(b)

	a	s	p	a	r	a	g	u	s	
				c	a	r	p	e	t	
					f	l	o	o	r	
				t	r	a	i	n	s	
							p	e	n	
								i	t	

(c)

	a	s	p	a	r	a	g	u	s	
			c	a	r	p	e	t		
			f	l	o	o	r			
			t	r	a	i	n	s		
				p	e	n				
				i	t					

Figure 4: Examples of (a) left-justified, (b) right-justified, and (c) middle-justified.

Impaired Visual Word Recognition

We now briefly review some of the central behaviours associated with the different subtypes of dyslexia and show that dysfunction in a split processor can provide a principled basis for many of these behaviours. The simplest dysfunction that we consider is the maladaptive coordination of the two halves of the model, each with their hemifield-specific information. One consequence of such a dysfunction is that more hypotheses and/or more specific hypotheses

may be developed intrahemispherically – before effective interhemispheric communication can occur – than would normally be the case. Problematic hemispheric interaction and atypical lateralisation have long been linked to dyslexia in a variety of ways (see, e.g., Davidson & Saron, 1992).

Surface Dyslexia The regularisation errors of surface dyslexia (*pint* pronounced as /pint/) (Marshall & Newcombe, 1973) naturally emerge as failures of access to the whole of the orthographic word. The relevant simulation has been presented by Plaut and McClelland (1993) in which they show that the appropriate pronunciation of the vowel in an exception word like *pint* requires the coordination of information from the orthographic input onset, nucleus and coda. In contrast, their model allows onset, nucleus and coda individually to determine their own pronunciation in regular words like *mill*. Substituting a phonological output for the orthographic one in Figure 1, we can see how dysfunctional coordination of the two halves of the model would cause regularisation errors based on the two halves of the input word: the first half of *pint* would be pronounced as in *pink* and *pill*, and the last half as in *hint* or *lint*.

Existing connectionist accounts of surface dyslexia rely on reducing the computational resources available to the processor (in the developmental case) (see, e.g. Harm & Seidenberg, 1996) or else characterise it in terms of the abilities of a surviving orthography-to-phonology pathway that had a division of labour with a now-impaired orthography-to-semantics pathway (in the acquired case) (see, e.g., Plaut, McClelland, Seidenberg & Patterson, 1996). Our own account may be seen as a special instance of the reduced resources account. It is a parsimonious account of surface dyslexia because hemispheric coordination is particularly vulnerable to non-specific insult or developmental abnormality. It is more motivated than the division of labour account; after all, even a small network model such as that presented by Plaut and McClelland (1993) can succeed in pronouncing exception words in a general model of pronunciation. Finally, the account we offer applies to both acquired and developmental surface dyslexia.

Phonological Dyslexia Phonological dyslexia is characterised by disproportionately poor performance in the pronunciation of nonwords, such as *nust* (Déroutés & Beauvois, 1979). A split processor allows a degree of more or less autonomous processing to occur in each half of the model, based on the visual information in the relevant (contralateral) hemifield. If we allow autonomous semantic processing to occur in each half of the model in this way, then a natural account of phonological dyslexia emerges. One way that a split model can adapt to poor coordination of the visual/orthographic input is to allow semantic processing to develop further in each half of the model, with the part of the word in each hemifield encouraging the partial activation of the semantics of the words possibly involved. The left half of *chair* might activate the semantics of *chase*, *chair*, *choir*, ... and the right half *air*, *chair*, *flair*, The intersection of these two sets is the presented word, *chair*. We might expect that mapping a word's phonology onto the semantic coordination of its two halves might be less

problematic than mapping a word's phonology onto the visual coordination of its two orthographic halves: callosal connectivity is more extensive for the "higher" cortical regions; orthographic-phonological mapping may be more exacting, requiring detailed temporal processing of the phonological representation in conjunction with position-specific processing of the orthographic representation.

A processor that relies on a semantics-to-phonology mapping cannot cope with nonwords or unknown words. The best that can be done is to be aware of similarities between known words and the two halves of the nonword, an observation which phonological dyslexics can often readily make.

This outline of an account of phonological dyslexia may complement existing accounts which rely on an impairment of the phonological representations (Harm & Seidenberg, 1996; Brown, 1997; Plaut, McClelland, Seidenberg & Patterson, 1996), although we might note that the originally observed phonological dyslexic (LB) seemed to present with no such phonological difficulties. It should also be noted that callosal agenesis – a close human equivalent to one end of the range of split models we have begun to explore – seems to involve a reading impairment that closely resembles phonological dyslexia (Temple, Jeeves & Vilarroya, 1990). The correlation between phonological impairment and phonological dyslexia may not be causal.

Deep Dyslexia Hinton and Shallice (1991) and Plaut and Shallice (1993) provide comprehensive coverage of the data from deep dyslexia, using a model that maps from orthography to semantics (and ultimately to phonology). The introduction of a split architecture into their connectionist modelling of deep dyslexia provides a principled and specific basis for the disintegration of orthographic and/or semantic processing. Relatively autonomous semantic processing in the two halves of the model provides an account of one very specific behaviour produced by deep dyslexics: when a confusion occurs over a word, it is typically the case that a less concrete word is replaced by a more concrete word, as in *plan* → *flan*. It is generally assumed that the RH is concerned with the representation of more concrete words. By allowing more independent processing to occur in the RH, we allow for the fact that more concrete interpretations of the data may become more influential in later processing. This account provides a basis for a longstanding suggestion (Coltheart, 1980) that deep dyslexia reflects the behaviour of a RH reading system.

In general, in this necessarily brief examination of the different subtypes of dyslexia, we can see that split connectionist architectures are able to provide a principled and specific basis for the emergence of a range of relevant types of reading impairment. A unified treatment of the subtypes of dyslexia is possible, based on maladaptive hemispheric coordination. At the same time, such models may be expected to inherit the general advantages, already demonstrated in the literature, of connectionist models in simulating other aspects of reading dysfunction.

Conclusions

We began by considering how the initial projection to the visual cortex is precisely split, and we then explored the implications of assuming that this split is not completely transcended by the time that visual word recognition occurs. We have shown that this assumption leads directly to natural accounts of a set of data drawn from different parts of the word recognition literature. The exterior letters effect is given an explanation that is not available to current connectionist models of visual word recognition. Similarly, the data on the optimum and preferred viewing position are given a coherent reinterpretation from the perspective of a split architecture. We have shown, empirically, that right-justification and middle-justification are no different in their ability to minimise entropy in individual letter positions. Finally, we have shown that a split model has the potential to provide motivated and specific explanations of a variety of behaviours found within the different subtypes of dyslexia.

The exact details of the functional architecture of word recognition are still far from being adequately mapped onto the known physiology and anatomy of the human visual system. However, the disparate data we have discussed above represent a growing body of circumstantial evidence that the most accurate and illuminating way to model word recognition is by means of a class of split architectures. This general model makes specific and testable predictions concerning a range of word recognition paradigms. For instance, if the processing in the two hemispheres is more autonomous than is allowed for in current models of visual word recognition, then a variety of confusion errors between the different hypotheses being entertained in the two hemispheres might be expected to occur. Just such errors have been reported in the literature: Mozer (1985) presented subjects very briefly with pairs of words such as *line* and *lace* and observed that "letter migration" errors occurred, in which subjects reported seeing *lice* or *lane*. From the perspective of the model discussed above, these errors might be better described as "hemifield mismatch errors".

We conclude that split connectionist models of visual word recognition will largely inherit the demonstrable advantages of existing connectionist models of visual word recognition, and will augment their coverage of the data in a parsimonious manner. Models of visual word recognition should reflect the way in which the brain makes sense of its immediate data about the visual world, something that is reflected in the large-scale anatomy of the visual system. At the same time, this modelling can inform our theories of the hemispheres' more general representational principles and strategies.

Acknowledgements

This research was supported partly by MRC (UK) grant G9421014N.

References

- Andersen, R.A., Asanuma, C., Essick, G.K., & Siegel, R.M. (1990). Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *Journal of Comparative Neurology*, 296, 65–113.

- Brown, G.D.A. (1997). Connectionism, phonology, reading and regularity in developmental dyslexia. *Brain and Language*, 59, 207–235.
- Brysaert, M. (1994). Interhemispheric transfer and the processing of foveally presented stimuli. *Behavioural Brain Research*, 64, 151–161.
- Chatterjee, A. (1995). Cross-over, completion and confabulation in unilateral spatial neglect. *Brain*, 118, 455–465.
- Coltheart, M. (1980). Deep dyslexia: a right-hemisphere hypothesis. In (M. Coltheart, K. Patterson & J.C. Marshall, eds.) *Deep Dyslexia*. Routledge and Kegan Paul.
- Davidson, R.J. & Saron, C.D. (1992). Evoked potential measures of interhemispheric transfer time in reading disabled and normal boys. *Developmental Neuropsychology*, Vol. 8(2–3), 261–277.
- Déroutésné, J. & Beauvois, H.F. (1979). Phonological processing in reading. *Journal of Neurology, Neurosurgery and Psychiatry*, 42, 1125–32.
- Ellis, A.W., Flude, B. & Young, A.W. (1987). Neglect dyslexia and the early visual processing of letters in words and nonwords. *Cognitive Neuropsychology*, 4, 439–464.
- Forster, K.I. & Gartlan, G. (1975). *Hash coding and search processes in lexical access*. Paper presented at the Second Experimental Psychology Conference, University of Sydney.
- Harm, M.W. & Seidenberg, M.S. (1996). Computational bases of two types of developmental dyslexia. *Proceedings of the 18th Cognitive Science Society Conference*, Hillsdale, NJ: Law. Erl. Assoc..
- Hellige (1995). Coordinating the different processing biases of the left and right cerebral hemispheres. In F. Kitterle, (Ed.) *Hemispheric communication: Mechanisms and models*. Law. Erl. Assoc., Hove.
- Hillis, A.E. & Caramazza, A. (1990). The effects of attentional deficits on reading and spelling. In A. Caramazza (Ed.) *Cognitive Neuropsychology and Neurolinguistics: Advances in Models of Cognitive Function and Impairment*. Law. Erl. Assoc., Hove, UK.
- Hinton, G.E. & Shallice T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review*, 98, 74–95.
- Jordan, T.R. (1990). Presenting words without interior letters: Superiority over single letters and influence of postmark boundaries. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 893–909.
- Legge, G.E., Klitz, T.S. & Tjan, B.S. (1997). Mr. Chips: An ideal-observer model of reading. *Psychological Review*, 104, 524–553.
- Marshall, J.C. & Newcombe, F. (1973). Patterns of paralexia: A psycholinguistic approach. *Journal of Psycholinguistic Research*, 2, 175–199.
- Mohr, B., Pulvermüller, F. & Zaidel, E. (1994). Lexical decision after left, right and bilateral presentation of function words, content words and non-words: Evidence for interhemispheric interaction. *Neuropsychologia*, 32, 105–124.
- Mozer, M.C. (1983). Letter migration in word perception. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 531–546.
- O'Regan, J.K. (1990). Eye movements and reading. In E. Kowler (Ed.) *Eye movements and their role in visual and cognitive processes*, North Holland: Elsevier Science Pub..
- Plaut, D.C. & McClelland, J.L. (1993). Generalization with componential attractors: word and nonword reading in an attractor network. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society*, Hillsdale, NJ: Law. Erl. Assoc..
- Plaut, D.C., McClelland, J.L., Seidenberg, M.S. & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56–115.
- Plaut, D. C., Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, 10, 377–500.
- Pulvermüller, F. (submitted). Words in the brain's language. *Behavioral and Brain Sciences*.
- Shillcock, R.C. & Monaghan, P. (submitted). Inter- and intrahemispheric processing and the modelling of visual word recognition. *Psychological Review*.
- Sugishita, M., Hamilton, C.R., Sakuma, I. & Hemmi, I. (1994). Hemispheric representation of the central retina of commissurotomy subjects. *Neuropsychologia*, 32, 399–415.
- Temple, C.M., Jeeves, M.A. & Vilarroya, O.O. (1990). Reading in callosal agenesis. *Brain and Language*, 39, 235–253.
- Yannakoudakis, E.J. & Hutton, P.J. (1992). An Assessment of *N*-phoneme statistics in phoneme guessing algorithms which aim to incorporate phonotactic constraints. *Speech Communication*, 11, 581–602.