

Towards a unified process model for graphemic buffer disorder and deep dysgraphia

David W. Glasspool

University College London, UK, and Cancer Research UK, London, UK

Tim Shallice

University College London, UK, and SISSA, Trieste, Italy

Lisa Cipolotti

University College London, UK, and National Hospital for Neurology and Neurosurgery, London, UK

Models based on the competitive queuing (CQ) approach can explain many of the effects on dysgraphic patients' spelling attributed to disruption of the "graphemic output buffer". Situating such a model in the wider spelling system, however, raises the question of what happens when input to the buffer (e.g., from a semantic system) is degraded while the buffer remains intact. We present a preliminary exploration of predictions following from the CQ approach. We show that the CQ account of the graphemic buffer predicts and explains the finding that deep dysgraphic patients generally show features of graphemic buffer disorder, as disrupted input from a damaged semantic system has an inevitable effect upon the functioning of the buffer. The approach also explains the most salient differences between the two syndromes, which are seen as consequences of the difference between an intact sequence generation system operating on degraded input versus a damaged sequencing system operating on intact input.

INTRODUCTION

There have been two implicit approaches to connectionist computational models among cognitive neuroscientists. Some have viewed connectionist models as challenging older information-processing accounts. Others have viewed them as unpacking information-processing accounts at a deeper theoretical level.

Traditionally, neuropsychological reasoning has tended to concentrate on determining what distinct processing systems, or modules, exist and their functional organization. Shallice (1988) emphasizes the use of dissociation data in such inferences. But modellers taking the first approach to connectionism have argued that some dissociation data may imply different types of damage to a single system rather than distinct

Correspondence should be addressed to D. Glasspool, Advanced Computation Laboratory, Cancer Research UK, 44 Lincoln's Inn Fields, London WC2A 3PX, UK (Email: david.glasspool@cancer.org.uk).

This work was supported in part by a grant from the McDonnell–Pew programme in cognitive neuroscience. We are grateful to George Houghton, Argye Hillis, Jon Machtynger, Richard Cooper, Gordon Brown, Andrew Ellis, and an anonymous referee for cogent comments on the modelling work and on earlier versions of this report.

processing systems (e.g., Plaut, 1995), undercutting a major theoretical tool in neuropsychology.

The second approach to connectionism has been less threatening to information-processing accounts. On this view connectionist modelling is a tool for developing concrete theories of the internal operation of the modules inferred by more traditional techniques (e.g., Burgess & Hitch, 1992). Computational modelling allows explanations for effects of damage to cognitive systems to operate at a finer level of detail, encompassing the effects of damage on the internal operation of a module as well as on the disrupted functioning of the system as a whole.

Shallice (1988) took a cautious approach in arguing that dissociation data are the most relevant data for neuropsychological reasoning, and that the two other types of information that have most often been taken as neuropsychological evidence—association of deficits and error data—both provide only weak constraints for theory. Caramazza and McCloskey (1991) and McCloskey and Caramazza (1991) defend the use of these types of evidence, as long as both the data and the putative theory are characterized at a sufficiently detailed level. Computational modelling enables concrete theorizing at this fine level of detail and in particular provides a framework for the appropriate use of error data in cognitive neuropsychology.

In this paper we develop a connectionist computational model of the second type, which draws on both these types of data—it is inspired by detailed consideration of patient errors, and it is used to account for an association of symptoms across two disorders.

Dysgraphic patients with “graphemic buffer disorder” (GBD) produce characteristic errors in their spelling, which indicate problems with the serial output of a sequence of letters in the final output stage of the spelling process. A second group of patients with so-called “deep dysgraphia” produce spelling errors indicative of a higher level locus of damage, with semantic content affecting errors. The two syndromes have generally been seen as quite distinct. However, as Cipolotti, Bird, Glasspool, and Shallice (2004) point out, they do in fact share many features in common.

In this paper we argue that this association of features is predicted on a straightforward assumption about the internal operation of one component of the spelling system, the graphemic output buffer (GOB).

In previous work we have modelled the operation of the GOB and have simulated spelling problems thought to be due to disruption to this component (GBD). In the work described here we take a first step towards situating our GOB model in a wider model of the spelling system by adding input from a semantic spelling route. Such a model allows us to compare the effect of degraded input to the GOB (following damage to the semantic components of the model) versus the effect of a damaged GOB operating on intact input. A number of gross features are shared by the error patterns produced by these two manipulations. Several aspects of the mechanisms leading to errors are, however, modified by the change in locus of damage, leading to both major and minor differences in the error patterns generated. There is a correspondence with the association of features and the salient differences between deep dysgraphia and GBD that allows us to propose a unified theoretical account for the two disorders.

We start by discussing the theoretical and empirical background to the work, and we propose a theoretical account for the similarities and differences between two key groups of GOB patients. We describe a computational model based on this account and report the results of a number of simulations comparing the model with a range of patients. Finally we discuss both successes and limitations of the modelling enterprise and relate them to aspects of the underlying theoretical approach.

THEORETICAL AND EMPIRICAL BACKGROUND

Two spelling syndromes

The standard information-processing model of spelling that was developed in the 1980s involved a number of routes—2 or 3 in different

versions—by which information was transmitted from a semantic system and/or auditory input (or output) lexicon system to the GOB (Ellis, 1984; Margolin, 1984; Shallice, 1988). Information was then thought to be held in the buffer before being translated letter by letter into allographs and then motor stroke representations (Ellis, 1982). In the context of a model of this type, Caramazza, Miceli, and their colleagues, in a series of papers, described two patients with an acquired dysgraphia whose deficit was held to be at the level of the GOB (FV: Miceli, Silveri, & Caramazza, 1985; and LB: Caramazza, Miceli, Villa, & Romani, 1987; Caramazza & Miceli, 1990). Since the initial descriptions of these patients other patients with a similar type of difficulty have been described (e.g., SE: Posteraro, Zinelli, & Mazzucchi, 1988; ML, DH: Hillis & Caramazza, 1989; JES: Aliminosa, McCloskey, Goodman-Schulman, & Sokol, 1993; HE: McCloskey, Badecker, Goodman-Schulman, & Aliminosa, 1994; JH: Kay & Hanley, 1994; AM: De Partz, 1995; AS: Jonsdottir, Shallice, & Wise, 1996; SFI: Miceli, Benvegna, Capasso, & Caramazza, 1995).

The spelling disorders of these patients have four main types of property (see Caramazza et al., 1987, and Shallice, Glasspool, & Houghton, 1995, for review):

1. There is a tendency for spelling errors to increase with word length.
2. Spelling accuracy is not affected by semantic or syntactic variables such as concreteness or part of speech, or variables relating to the mapping from phonology to orthography, and is generally unaffected by word frequency.
3. Nonwords are spelled in a roughly similar fashion to words but less accurately length for length.
4. Errors typically correspond to one or two operations of the following type: substitution, deletion, transposition, and insertion of single letters.

Characteristics 2 and 3 indicate that the disorder occurs no higher than the stage at which the outputs from the semantic and phonological

systems to the graphemic output system converge. The other properties correspond to what one would expect from an impairment to a buffer in which individual letter representations are held as in the standard model.

More recently a second set of patients have been described whose errors also show effects of word length and include substitutions, deletions, transpositions, and insertions—Properties 1 and 4 above. However, Properties 2 and 3 do not apply; unlike the first type of patient, semantic variables such as concreteness affect spelling performance, as does word frequency, and nonwords are typically very poorly spelled (Cipolotti et al., 2004). This second set of patients seem to combine failure of the GOB with a deep dysgraphia (Bub & Kertesz, 1982); the dysgraphia syndrome with analogous properties to those of deep dyslexia.

Additionally, there are a number of differences in detailed aspects of the error pattern compared with the “classic” GOB disorder described above. Patients in this second set show spelling error rates that increase monotonically from start to end of word, whereas patients in the first set show a bowed serial curve with lower error rates at the start and end of words. The second set of patients show a dominance of letter deletion errors whereas the first set often produce a high proportion of letter substitutions. Finally, this second set of patients produce large numbers of a new type of error not reported for the first set—the fragment (Ward & Romani, 1998), responses two or more letters shorter than the target word, often preserving earlier portions of the word. This set of patients include HR (Katz, 1991), BA (Ward & Romani, 1998), TH and PB (Schiller, Greenhall, Shelton, & Caramazza, 2001), and DA (Cipolotti et al., 2004).

For convenience we label these two groups of patients as displaying “GOB disorder Type A” and “GOB disorder Type B”, respectively. There are sufficient differences between the groups that one might be tempted to propose that they correspond to distinct functional syndromes. However, two factors suggest that the picture may not be so straightforward.

First, while most patients fall into one group or the other, a number of patients exist who appear to be somewhat intermediate between the two groups. The most common type of intermediate patient are those apparently of Type A who show a significant effect of word frequency (AS, JH, DH, JES, HE). A few patients apparently of Type B produce no semantic errors—HR and AZO (Miceli, Capasso, Ivella, & Caramazza, 1997)—and a few patients have anomalous serial position curves (FM, Tainturier & Caramazza, 1996, and AM are apparently Type B but have bowed serial position curves; FV is Type A but has a flat serial position curve; GSI, Miceli, Capasso, Benvegnu, & Caramazza, 2004, is apparently of Type A but has a monotonic serial error curve). One patient appears to combine both the last two types of discrepancy: BH has both Type A characteristics (bowed serial position curve, no semantic errors) and Type B characteristics (frequency and imageability effects and fragment errors).

Second, it is not clear that the two groups are functionally distinct. In many respects “prototypic” Type B disorder is similar to Type A disorder with the addition of some extra features—“deep dysgraphic” type semantic errors, sensitivity to word frequency and semantic content, and fragment errors. According to the standard information-processing model of spelling, GOB disorder was held to be essentially unrelated to deep dysgraphia, which was usually seen as involving complete damage to the phonological route or routes and also partial damage to the semantic system or to its direct connections to the GOB. However, Cipolotti et al. (2004) review a number of accounts of the second type of GOB syndrome in the literature. Where data were given, Cipolotti et al. conclude that the relevant patients show GOB Type B errors as well as semantic and frequency effects. The suggestion that deep dysgraphia includes aspects of GOB disorder implies a much more intimate link between the two syndromes.

In this paper we propose a theoretical account for GOB disorder Types A and B and their relationship to deep dysgraphia. We treat “prototypic” Type A and B disorders as relating to

orthogonal types of damage, so that the two disorders represent end-points of a potential continuum of intermediate cases where both types of damage are present. We propose that Type A disorder results from disrupted operation of the GOB caused by damage to the GOB itself, upstream systems (including semantic systems) remaining intact. Type B disorder also results from disrupted operation of the GOB, but in this case the disruption is caused by damage to upstream semantic systems, which then feed degraded input to an intact GOB.

However, while Type B disorder in most respects simply adds features to Type A, a few features are altered. The mix of different error types is changed, deletion errors becoming more prominent, and the serial position curve for errors, which is bowed in Type A disorder, becomes a monotonic rising curve in Type B. If an account of the type that we suggest is to be successful it must explain these differences in terms of differences in the operation of the GOB when presented with degraded input compared with its operation when it is itself damaged. This is a situation where “classical” neuropsychological reasoning, based on identifying processing modules, which are then treated as “black boxes”, cannot help us. In order to provide an account of this type it is necessary to theorize about the internal operation of a module—the GOB—and the effects of different types of disruption.

The general theoretical position outlined above is compatible with our process model of the GOB (Glasspool, 1998; Glasspool & Houghton, 2005; Houghton, Glasspool, & Shallice, 1994; Shallice et al., 1995) and with our approach to modelling the semantic system (Plaut & Shallice, 1993). It is also compatible with the suggestion of Ward and Romani (1998) that the Type A features of spelling errors in their Type B patient BA might result from damage to a GOB model of the sort that we have proposed.

Our purpose in this paper, then, is to begin to explore the consequences of placing a process model of the GOB in the context of other parts of the spelling system and to test the theoretical account proposed above by exploring the effect of

degraded input on the operation of the GOB, on both gross and detailed error patterns. Our aim is to lay the foundations for a parsimonious theoretical explanation of the two types of GOB disorder, A and B, including the similarities and differences between them and their relation to deep dysgraphia.

Competitive queuing accounts of the graphemic output buffer

Spelling, like other forms of language production, is an inherently serial process. While generation of serially ordered sequences of responses is basic to human and animal behaviour, the processes underlying serial behaviour are not well understood (Houghton & Hartley, 1995; Lashley, 1951). One effect of the recent rise of connectionist models has been to force modellers to address the problem of serial behaviour from first principles, rather than relying on the serial constructs available in traditional computing paradigms to side-step the issue. The variety and complexity of the solutions that connectionists have found to the problem of serial behaviour is testament to the difficulty of achieving this.

One class of models that has been successful in explaining a number of common features of human automatic serial behaviour, in the fields of language and of short-term memory for example, are those termed competitive queuing (CQ) models (Glasspool, 1998, 2005; Houghton, 1990). A number of successful models in a range of psychological paradigms have been based on this approach, including verbal short-term memory (e.g., Burgess & Hitch, 1992, 1996), typing (Rumelhart & Norman, 1982), spelling (Houghton et al., 1994; Shallice et al., 1995), speech production (e.g., Hartley & Houghton, 1996), and action plans (e.g., Cooper & Shallice, 2000).

Although there are important differences between implemented CQ models, the following three features are typically present:

1. *A set of response representations, which can become refractory:* These constitute a finite pool of

representations of distinct responses or actions from which individual sequences are generated. The responses are potentially refractory in that when an item is produced as part of a sequence, it becomes temporarily unavailable for further use.

2. *Parallel activation of response representations and activation gradients:* The representations of responses in a target sequence are activated in parallel at the beginning of sequence generation, but with a gradient of activation over them such that the sooner a response is to be produced the more active it is. The set of active response representations forms the “competitive queue”, as the representations compete for output on the basis of their activation level. The relative activation levels may be static throughout production of a sequence or may change over time.
3. *A competitive output mechanism:* This mechanism has to resolve the response competition in the queue by selecting the currently most active response representation. This process of selection triggers the subsequent inhibition of the chosen representation.

The process of generating a sequence of responses in such a model involves an *activating mechanism* generating a gradient of activations over some subset of item representations. The competitive output mechanism then repeatedly selects for output the most active item. As each item is output it becomes refractory and hence temporarily unavailable. In this way the activated items are output in order of their relative activation values, from the most to the least active.

CQ models rely on the transient inhibition of item representations after they appear in a sequence to prevent perseveration of responses. Sequences with items occurring more than once can nonetheless be generated, but this requires that an item representation used earlier in a sequence is sufficiently strongly excited at a later point that the transient inhibition is overcome. Generally this entails allowing the pattern of excitation applied to sequence items to change during

the course of sequence production (Glasspool, 1998; Houghton, 1994). While sequences with repeated items separated by other items can be represented, immediate doubling of an item presents a particular difficulty and requires a special-purpose mechanism. We have argued (Houghton et al., 1994; Shallice et al., 1995) that the apparent marked status of double items in sequences, and in particular the types of error made on doubled items (geminate)—for example, in spelling (Caramazza & Miceli, 1990) and typing (Rumelhart & Norman, 1982), in which the “doubled” status appears to be represented separately from the identity of the item being doubled—constitute strong evidence for special treatment of doubled items in serial behaviour. The CQ model provides a principled explanation of their special status. Moreover, errors involving doubles are strongly suggestive of the errors found in implemented CQ systems with special-purpose doubling mechanisms added to them.

Our previous work on modelling neurological disorders of spelling using the CQ modelling paradigm has been concerned with GOB disorder Type A. Within a symbolic-stages model framework (e.g., Caramazza & Miceli, 1990) this disorder can be characterized as resulting from an impairment to a buffer that holds graphemic representations prior to their production in writing or spelling aloud. However, current accounts using the symbolic-stages approach of the Type A disorder (e.g., Caramazza & Miceli, 1990) make claims about the form of information represented within the spelling system without providing an account of the mechanism by which spelling is achieved from these representations. While it is important to have representational theories, certain features of an error pattern may be due to the mechanism that is damaged rather than the information that it is processing. Indeed we believe this to be the case for errors arising from damage to the “graphemic buffer”. It is important, therefore, to consider both representation and mechanism.

Our approach has been to view the GOB as a CQ system for generating sequences of letter identities when receiving input representing the

identity of the word to be spelled. It is assumed that these letter identities are processed by downstream systems to produce written or spoken output. The GOB is therefore viewed primarily in the role of a sequence generator. On this approach the Type A disorder was explained as resulting from an increased uncertainty in the selection of the “winning” letter from the set of letter activations available at each moment during spelling production. This might, for example, be due to disruption of the selection system itself, disruption to or noise in the system generating the activation gradient, or a reduction in the general level of activation of letter representations in the face of a fixed amount of background noise. Our theoretical approach has been to explain what we consider to be the most salient features of Type A GBD as resulting directly from the characteristic breakdown pattern of the CQ sequencing process under random perturbation of “winning item” selection. This pattern includes a clear effect of word length on error rate, bowed serial error curves, and a characteristic set of error types including ordering errors such as the exchange of two letters within a word, letter substitution, insertion and deletion errors, and errors involving movement of a gemination operator. Within this framework other regularities in the GBD error pattern are most naturally interpreted as reflecting features of the activating system that establishes the activation levels of queued letters (Glasspool & Houghton, 1997). Regularities within this system lead to a characteristic pattern of vulnerability, upon which the detailed CQ error pattern is superimposed. This pattern of vulnerability would include, for example, the propensity for consonants to exchange with consonants and vowels with vowels in some patients’ errors.

In common with most models of the CQ type, this work has been based on fully localist single-layer connectionist networks (Glasspool, 1998; Glasspool & Houghton, 2005; Glasspool et al., 1995; Houghton et al., 1994; Shallice et al., 1995). These papers describe a number of different versions of the basic localist model. Our approach has been to start with the basic CQ

sequence generation mechanism and then to elaborate the activating mechanism in order to attempt to simulate Type A GBD in more detail. For instance, the earlier models do not include explicit consonant/vowel representations, while the later models do.

However, from a computational perspective the form of localist representations and learning algorithm used in this class of CQ models has some disadvantages. A separate representation for the temporal context of each letter must be created for each word—the generation of timing information that drives sequential recall is integral to the representation of the sequence, and there are no opportunities for generalization across words based on temporal position. Additionally the lack of shared representations at the letter level of the model means that there are no cross-word generalizations in the spelling properties of the model; thus frequency of letter combinations and letter sequential dependencies are not represented in any way. No efficiencies of generalization are possible in the storage of several very similar words. Each must be stored separately. Moreover the problem of combining information sources, as when phonological and lexical information are combined in a two-route spelling model, is difficult (Glasspool et al., 1995).

In Glasspool (1998) and Glasspool, Shallice, and Cipolotti (1999) we describe an extension to the standard fully localist CQ architecture to allow CQ sequence production in a multilayer connectionist network trained with a backpropagation-type rule. This type of model does not have the theoretical inadequacies of the earlier localist versions. However, replacing localist context representations by distributed representations and word-specific sets of connections by a shared set of connections used by all words required hundreds of times more training epochs; it was computationally much more demanding. The model therefore had to be more restricted in its scope with omission, for instance, of consideration of geminates.

The overall philosophy of this approach is not to view these distributed models as inherently more adequate than the localist models, but only

more adequate in particular respects. The modeling philosophy is to produce a set of models, each of which captures salient parts of an overall ideal model. If the models each work well with respect to the particular set of empirical phenomena most related to their key assumptions then an attempt would be made to combine them into a macromodel. We view these models as attempts to explore the possibilities of elaborating in various ways upon this central theme and also to test our hypothesis that a core set of features of GBD are due to the dynamic behaviour of the CQ sequencing mechanism rather than to the details of any particular implementation.

Predictions of the CQ GOB account

It is reasonable to assume that output from a damaged semantic pathway impinging on the GOB will be degraded relative to its usual form. What, though, will be the effect of such degraded input on the operation of the buffer? On the standard information-processing model of spelling, it is not possible to make any prediction.

Our view of the GOB as a CQ mechanism, however, predicts that degraded input will affect the generation of letter sequences in a number of ways. If degraded input leads to suboptimal activation of “queued” letter representations, then in the face of a constant level of background noise we would expect the likelihood of errors in sequence generation to be higher and the types of error to be broadly similar to those seen in GOB disorder. We would, however, expect differences in the detailed form of the errors produced. For example, our recent spelling models have, in common with many other CQ models, reacted to subthreshold activation of letters by omitting a response and have interpreted very low overall activation across letter representations as indicating that the end of a word has been reached. We would therefore expect omission of letters and premature ending of words to be more likely when disruption is due to low input activation level rather than increased uncertainty in letter selection.

When the semantic pathway is damaged, one would expect the degree of degradation in the representation of particular words to vary depending on semantic factors, such as semantic category, and on the robustness of the semantic representation, which can be expected to vary with concreteness or imageability. The degree of GOB disruption would be expected to reflect these variations, superimposing these effects on the pattern of errors produced by the disrupted GOB. The error pattern following from damage to the GOB in the presence of an optimal input from an intact semantic system, however, is dependent only on the degree of disruption internal to the CQ system and would not be expected to vary with the semantic category or the lexical status of the target word.

Both the localist and distributed versions of the CQ GOB model make an interesting prediction with respect to word frequency. The CQ approach to sequence generation requires only that the correct item is most active at a particular point in a sequence; the exact activation level, and the state of other items, is not important. CQ systems are thus less demanding than usual in their learning algorithms, and our models consequently have the property that once a word has been learned sufficiently that it can be correctly recalled (or once a specified level of robustness has been achieved in its recall), further exposure to the word will not lead to increased robustness of representations. Hence, providing all words in a training corpus have been learned at least to the minimum level required for correct recall, we would predict that errors caused by disruption to the GOB will not be sensitive to word frequency. The semantic-to-orthographic mapping may be modelled by a more conventional connectionist network, which we would expect to be sensitive to word frequency.¹

Our general prediction is thus that damage to systems or pathways impinging on the GOB will

lead to error patterns with the broad characteristics of GOB disorder—increase of error rate with word length, presence of ordering errors, separation of geminate features from letter identity in errors, and possibly preservation of consonant/vowel status. However, there would be a greater rate of omission errors, particularly at the end of words due to the early ending of the production of letters, than is seen when sequence generation is disrupted by uncertainty in winning-letter selection. We would expect the somewhat different balance of error mechanisms at work to give rise to altered serial position curves for errors. Moreover, these effects will be superimposed on a pattern related to the vulnerability of the damaged system or pathway—in the case of the semantic system, we would expect that these will include effects related to semantic class and concreteness and word frequency. This is suggestive of the characteristics of Type B disorder. Disruption to the GOB, on the other hand, while the semantic orthographic system is undamaged, should lead to classic Type A error patterns with no effect of semantic class or word frequency. These predictions suggest that a rather simple model might provide a parsimonious unified explanation for the features of Type A and B dysgraphias and for the relationship between them.

In order to test these predictions, and to determine the detailed form of the error pattern resulting from suboptimal input to a CQ sequence generation system, we describe here a series of simulations with an outline computational model combining the GOB with input from lexical semantic representations, based on an extension to our multilayer network model of the GOB (Glasspool, 1998; Glasspool et al., 1999). In the current model we only consider inputs to the CQ system from semantic representations. This is done for two reasons. First, the overall

¹ In the model presented here a standard backpropagation network is used to model the semantic pathway. A more elaborate model of this pathway might include attractor dynamics (Plaut & Shallice, 1993), and it is not clear what effect this might have on the sensitivity of the semantic route to word presentation frequency. The incremental way in which such attractors emerge during training suggests that word frequency would affect robustness of representation in such a network, however, so we do not expect that the present simplified model would behave qualitatively differently in this respect.

modelling philosophy, as discussed earlier, is to increase the complexity of the basic CQ model in a step-wise fashion. Thus we increase the complexity by adding to existing models a simulation of one "route" to the graphemic buffer rather than two or more and do this by incorporating a route that carries representations that have no inherent temporal structure in them so that two different temporal structures do not need to be combined theoretically. Secondly, as the phonological route or routes, which are not simulated, are generally held to be grossly damaged in deep dysgraphia so the semantic route, which is simulated, is the route that should be critical for an account of the determining properties of deep dysgraphia.

However, we make one exception to this concentration on semantic aspects. The initial creation of the output graphemic representations will be primarily driven from sublexical phonology (Lennox & Siegel, 1994). Thus certain phonological characteristics may be expected to be implicitly represented in the structure of output graphemic representations. We therefore make the most limited assumption of this type, as proposed by Caramazza and Miceli (1990; see also Cubelli, 1991), that consonant and vowel status are represented in the output.

THE MODEL

The architecture of the model is shown in Figure 1. The following informal description presumes a familiarity with the principles of connectionist models; a more formal treatment is given in the Appendix. Full technical details of the GOB model on which it is based are given by Glasspool (1998).

The model comprises two subsystems representing the semantic activating system and the GOB (dotted outlines in Figure 1).

The *semantic representation* field represents the input to the semantic spelling pathway and specifies the word to be spelled in terms of a pattern of activation distributed over nodes representing

semantic features. The pattern is held steady during the spelling of a word.

The immediate input to the GOB system cannot be the semantic representation itself since in the full system its input can also be based on translation from the phonological system. In fact, during normal development of the spelling system spelling is likely to derive initially from phonology (Lennox & Siegel, 1994). One would therefore expect to find representations intermediate between semantic features and letters even if these simply represent the hidden layer of a phonology-to-spelling system. Moreover, without a joint input from the two routes there would be no transfer of spelling knowledge between them. For these reasons the immediate input to the GOB on the present model is what we call the *word identity field*. It would correspond to the graphemic output lexicon on a symbolic approach and is held to represent the identity of the to-be-spelled word independent of semantic or phonological content. The semantic representation field projects via *Hidden Layer 1* to the word identity field. The semantic activating subsystem is thus a straightforward multilayer connectionist network mapping from semantic representation to word identity.

The GOB subsystem has two input fields of nodes. One is the word identity field discussed above. Activation of the other input field, *serial position*, changes in a stereotyped way through a series of patterns representing successive serial letter positions in the word. This is equivalent to similar fields in other CQ sequencing models (e.g., Burgess & Hitch, 1992; Houghton, 1990). The two input fields project via *Hidden Layer 2* onto an output field representing individual letter identities.

The parallel activation pattern over letter nodes is resolved into a sequence of letter identities by a competitive queuing output stage. The most active letter node is determined at each step during production, and this constitutes the output of the model. Following output of a letter the corresponding letter identity node becomes refractory and is briefly unavailable for further output. In our previous models this function has been performed by

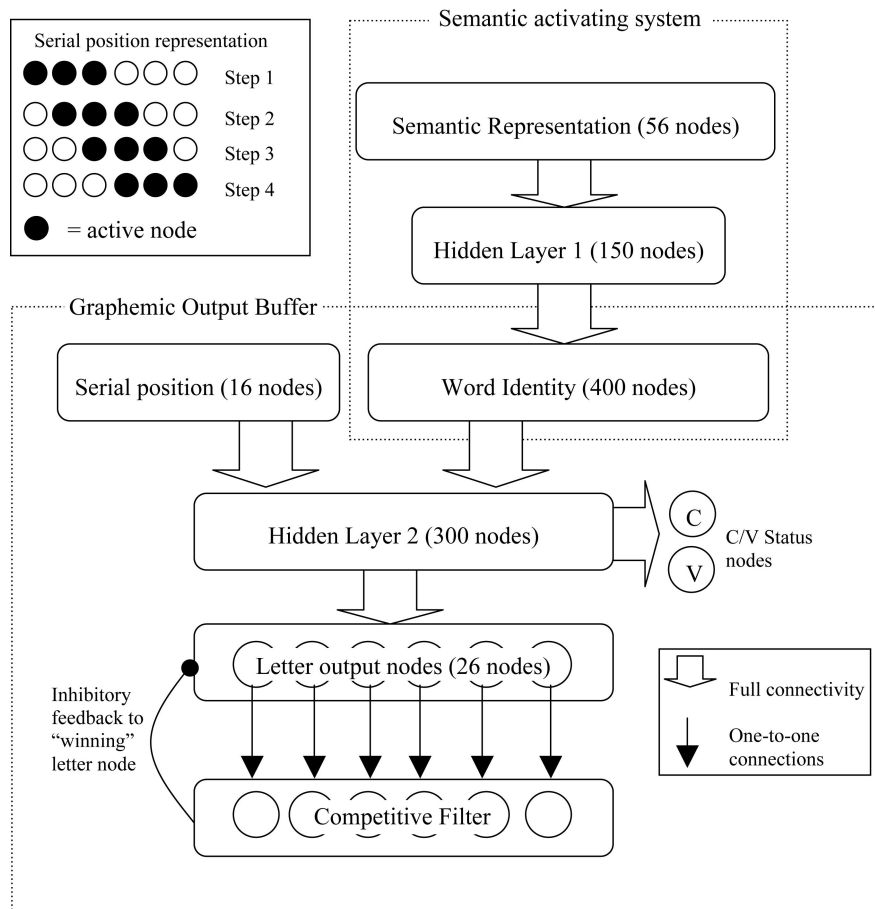


Figure 1. Architecture of the model. See text for details.

a “competitive filter”—a field with strong inhibitory connections between nodes creating a winner-takes-all competition (Houghton, 1990). In the present model a competitive filter is simulated by a peak-picking algorithm. The GOB network learns to associate a particular serial position in a particular word with a pattern of excitation on the output field, which, in conjunction with the “select-inhibit” dynamics of the CQ approach, results in the corresponding letter node being most active for that serial position during spelling.

Two threshold levels are free parameters of the model. If on a particular time step no letter activation exceeds the response threshold T_r no

response is made at that time step. If no letter exceeds the lower stopping threshold T_s , then the letter sequence for the current word is ended.

As in previous work (Glasspool, 1998; Glasspool et al., 1999) we add a pair of nodes at the output layer to indicate the consonant/vowel (CV) status of each letter. During training the network is required to activate the C or V node in parallel with the appropriate letter node. The CV nodes are ignored during testing (this has some similarities with the approach of Christiansen, 1997). This is to encourage the network to represent C/V status of letters; we make no theoretical claims for the way this is

achieved in the current model. We interpret this as an example of a common phenomenon in serial behaviour—the imposition of domain-specific serial constraints or biases. Such constraints on the class of elements participating in an error also operate in models of speech output and verbal short-term memory, for example (see Glasspool & Houghton, 1997, for discussion). The C and V nodes are not included in the select/inhibit output process mediated by the competitive filter.

For reasons given above, we have not extended the multilayer version of our spelling model to include a mechanism for doubled letters, although we do not anticipate that this would in principle be more difficult than for the single-layer localist models in which such mechanisms have been realized (Houghton et al., 1994; Rumelhart & Norman, 1982; Shallice et al., 1995). We do not aim to account for doubled letters in the present model but assume that the general account provided in previous models will hold.

Representations and training

The representations used for input and output to the various elements of the model, the training procedures used, and the details of the training corpus are as follows.

Input representations

The input to the semantic route model (*semantic representation* in Figure 1) is a semantic representation of the word to be spelled. For the purposes of the current model we have not attempted to encode the semantics of each word individually. Each word is instead assigned a randomly generated vector of 56 semantic features, each of which may be “on” or “off”. A total of 28 of the semantic features are labelled “concrete” features, and high-concreteness words each have 14 of these features active, such that every possible pair of words differ on at least 5 features. The other 28 features are labelled “abstract”, and low-concreteness words each have 5 of these features active such that each pair differs on at least 3. High- and low-concreteness words thus differ in the

density of their semantic representations. This method of representing abstraction is similar, for example, to that used by Plaut and Shallice (1993) in their explicit representation of semantic information.

The field of nodes forming the output of the semantic system and the input to the GOB (the *word identity* field in Figure 1) may be interpreted as corresponding to the output graphemic lexicon in the standard model. In a fuller model this field would represent the place where phonological lexical and semantic lexical information interact prior to the GOB. Representations here are assumed to be independent of semantic content, but some degree of overlap is nonetheless assumed between orthographically similar words to allow for some generalization across such words in spelling knowledge. We have based the chosen representation on the assumption that this field might represent orthographic words as the blend of a small number of idealized quasi-morphemes. A sparse distributed representation was therefore selected for this layer with 400 vectors defined over 400 nodes such that each vector has 4 active nodes. The scheme used ensures that each vector overlaps (shares active nodes) with six others, of which two each overlap by one active node, two overlap by two nodes, and two overlap by three. Vectors are randomly assigned to words.

CQ models require a means of *roughly* representing item position within a larger unit. This is the job of the *serial position* units. Using binary valued inputs, a suitable representation for serial position is that introduced by Burgess and Hitch (1992, 1996; see inset, Figure 1). A set of patterns is generated by shifting a “window” of active units across a field of inactive units. Each position is uniquely represented, but there is some overlap with other positions, the overlap being greater the closer positions are together. In the model a window of 8 active units is shifted across a field of 16 units to produce eight vectors that can be used to identify serial positions in the word to be spelled.

Output representations

Since the CQ approach depends on the localized inhibition of a single output item at each step in

sequence production, a local representation of letter identity is used at the output layer (*letter output nodes* on Figure 1) as in previous CQ models. A set of 26 nodes represents letter identities. An additional pair of output nodes represents the consonant/vowel status of letters during training.

Training corpus

A corpus of 400 words was used, 100 each of length 4, 5, 6, and 7 letters. Words were selected from the MRC psycholinguistic database (Coltheart, 1981; Quinlan, 1993). While word frequency and concreteness are simulated within the model (as detailed elsewhere in this section), the orthographic structure of words may vary systematically with these factors. We therefore used actual words that were either high or low on frequency or concreteness or both in the corpus. Half the words of each length were low frequency (mean Kucera–Francis frequency 3.8), the other half high frequency (mean frequency 190). In each of the resulting eight sets of words half were high concreteness (mean concreteness value in the MRC psycholinguistic database 536) and half low concreteness (mean MRC concreteness value 313). Frequency values were balanced across high- and low-concreteness conditions and concreteness values across high- and low-frequency conditions, and both were equalized across different word lengths, as far as possible.

Training procedure

Other than the one-to-one connections between output nodes and the competitive filter, all connections in the network are modifiable. For connections in the semantic route model a standard cross-entropy backpropagation training procedure was used (see Appendix). This section of the network was trained in isolation until the maximum error on any word identity node was below 1%, which occurred after 885 epochs. In each training epoch every high-frequency word was presented to the network, but each low-frequency word was presented with a probability of .3. After training, weights on this section of the network were frozen, and the entire network

from the input layers to the letter identity field was trained in the following fashion.

Conventionally, backpropagation algorithms assume that the activation level of every output unit is equally important for the generation of a correct output. What is critical in the present situation is that at each time step the correct letter is the most active one. The CQ mechanism itself ensures that activation levels of units other than the most active one have no effect on immediate behaviour. Moreover subsequent behaviour needs to be influenced by the current activation levels of the competitors. For the remainder of the model, which implements a CQ output stage, a more minimal procedure than standard backpropagation is therefore used. A “lazy” learning rule is used in an approach that is similar in some respects to that of Jordan (1986). This learning procedure, proposed by Glasspool (1998), is “lazy” in that learning only occurs for the individual letters explicitly involved when an error occurs in sequence generation. The error signal used to drive learning in the network is calculated for each letter node at each time step essentially according to the following rules (further details are given in the Appendix):

- If the letter that should be most active in the current position within the current word—the target letter—is not activated above the preset response threshold T_r , the error signal is such that this letter is reinforced, making it likely to be more active in this position in subsequent trials. If it is activated above threshold, no error signal is generated for the target letter.
- If the most active letter is not the target letter, then the error signal for the most active letter is such that this letter is punished, making it likely to be less active in this position in subsequent trials.
- No error signal is generated for any other letters, regardless of their activation level.

Training is continued for two letter positions beyond the end of the word, in which the model is trained to reduce all letter activation below the stopping threshold T_s . A small margin is incorporated in comparisons with thresholds in these rules

(see Appendix), meaning that the model is trained with slightly more stringent thresholds than those that are used during recall. This is similar to the “learning margin” used in our previous CQ models and prevents undue fragility of spelling performance simply due to operation of the model too close to its thresholds. The error values thus calculated are used in a cross-entropy backpropagation learning algorithm (see the Appendix for details). In each training epoch every word is presented to the network once, in random order.

While our preferred explanation for the lack of word frequency effects in GBD is as outlined above—the clearly defined limits on learning in CQ systems resulting in no overlearning for frequent items—there is a practical difficulty for this explanation in the present model. Pilot studies indicated the pragmatic computational need in a network of this complexity for a “momentum” term in the learning rule (see Appendix) in order to learn a substantial number of words in a tractable time on our equipment. The use of a momentum term is a standard technique in networks with backpropagation learning rules and relates purely to speeding up learning. However, the momentum term also entails that weight changes in the network during training are artificially affected by presentation frequency. As this is a pragmatic feature of the implementation intended to accelerate learning we do not place any theoretical weight on it. However, for practical purposes in the majority of simulations reported here the momentum term is present. To avoid the artifactual effect of presentation frequency in the GOB section of the network we therefore only expose the semantic pathway section of the model to varying word frequency—high- and low-frequency words are presented equally often to the GOB section. Because this is an unrealistic assumption a version of the model was also tested in which no momentum term is used, and in which both subsystems are exposed to varying word frequency. Of necessity this is a smaller scale model; however, it is sufficient to test our predictions regarding word frequency effects. It is discussed in the section on frequency effects below.

Simulation procedure

Type A disorder

On the model, Type A disorder is the consequence of disruption to the sequence generation mechanism of the GOB itself. It has been argued elsewhere (e.g., Houghton et al., 1994) that one obvious manipulation exists whereby nonspecific damage to the operation of a CQ system may be simulated—the addition of random noise to the activation levels of competing items. This has the effect of rendering the competitive process nondeterministic and corresponds to a loss of positional specificity in the sequencing process. Accordingly, random disruption of the output competition is how we simulate damage to the GOB component of our model. However, simply adding noise directly to item activation levels can lead to instability when recall is halted using a threshold on activation levels. While other mechanisms are possible for halting sequence generation at the end of a word, none is so intuitively satisfying as stopping when no letter remains with superthreshold activation (Glasspool, 1998). We follow Glasspool (1998) and Glasspool and Houghton (2005) in adding noise at the level of the competitive filter rather than at the letter output field (see Appendix). This may be thought of as targeting disruption directly at the process of response selection.

Type B disorder

On the model, Type B disorder results from degraded input to the GOB. Because we are interested in the relationship between this disorder and deep dysgraphia we have included parts of the lexical semantic system in the model, and we simulate Type B disorder by lesioning connections in the projections from the semantic representation field via Hidden Layer 1 to the word identity field. Bullinaria and Chater (1995) show that, unless a very large number of hidden units can be used, a lesion to such a projection in a multi-layer feedforward network should always be simulated by uniform reduction in strength of all connections in the projection rather than random cropping of connections. As the number of

connections in a projection increases towards the high numbers found in biological neural systems, the results obtained by random cropping asymptotically approach those of uniform weight reduction. Conversely in the limited simulations that are tractable on current equipment input-output mappings are distributed unevenly across connections and hidden units, and cropping and weight reduction produce divergent results. Accordingly we simulate damage to the semantic pathway by reducing connection weights by a uniform proportion.

SIMULATIONS

The model was trained as described above until 20 epochs were achieved with all words spelled correctly. This was after 708 epochs. The intact model, when fully trained, produces correct spellings for all 400 words in its corpus (the present model is not intended to model the spelling of novel words; see the Discussion section for comments on this process). From this single intact model four different damaged models were produced, which were used to provide the results reported in this section.

A Type A model was produced by adding random noise to activation levels of all nodes in the competitive filter of the intact model. The noise level was adjusted to give a performance of approximately 50% correct spelling for six-letter words, which is comparable to our previous simulation work and with the performance of several well-known Type A patients (e.g., AS and LB). This was achieved with a noise magnitude of ± 0.04 .

A Type B model was produced by setting a much lower background level of noise in the competitive filter nodes (± 0.005 , sufficient to cause less than 0.05% errors with intact weights) to simulate an intact GOB. Damage to the semantic system was then simulated by degrading weights. All weights in the projections from the semantic representation field via Hidden Layer 1 to the word representation field were scaled by a uniform factor of 0.58, which was found to give

similar overall spelling performance to DA (4% correct spelling on seven-letter words, average 15% correct over four- to six-letter words).

One concern is that the Type B model is operating with a considerably more severe deficit than the Type A model, which raises the possibility that differences in error patterns between the two models could be due simply to the difference in severity of damage rather than locus of damage. We therefore produced two control models—the “Type B control” model is a Type B model with a lower level of impairment (weights scaled by 0.67) to match its performance approximately to the Type A primary model, and the “Type A control” model is a Type A model with a higher noise magnitude (± 0.09) to match its performance to the Type B primary model. In the results reported below the primary models are given solid lines in graphs, and the control models broken lines. All results were averaged over 500 runs of the models over the test corpus (i.e., 200,000 attempted spellings for each model). We illustrate the performance of the models by comparing with LB (Caramazza et al., 1987) and AS (Jonsdottir et al., 1996), Type A patients, and with DA (Cipolotti et al., 2004) and BA (Ward & Romani, 1998), Type B patients. Of these, AS, DA, and BA are English speakers, and LB is an Italian speaker.

Word length effect

Figure 2 shows overall spelling performance of each model plotted against word length. The damage introduced to the models produces spelling errors, and these show a clear effect of word length with fewer errors on shorter words. Strong sequence length effects are a standard finding with CQ models and occur for at least three reasons. First, more items in a sequence present more opportunities for error. Second, there is a tendency for recall to be better (due to relatively less confusability in item position) near the start and end of a sequence, and these end effects have a relatively larger impact on shorter sequences. Third, in some models the gradient of activation levels of the sequenced items is steeper in shorter sequences




Figure 2. Overall spelling performance of the models for each word length, compared with Type A patient AS and Type B patient DA.

(this makes the difference in activation between adjacent items larger, and so successive items become more distinctive). Given the type of representation that we use at the sequence position layer we assume that the first two of these factors are primarily responsible for the effect in the present model.

The primary models produce curves with similar gradients to their controls. The curves appear more related to severity than to locus of damage in the models. Figure 2 also plots comparable data for Type A patients AS and LB and Type B patients DA and BA. The Type A modelled curves are similar in shape and gradient to the Type A patients. The Type B curves are similar to patient DA but patient BA clearly has a less severe deficit.

Concreteness and frequency

Figure 3 shows overall performance of each of the damaged models on (a) the high- and low-frequency words in the test corpus and (b) the high- and low-concreteness words (200 words in each category).

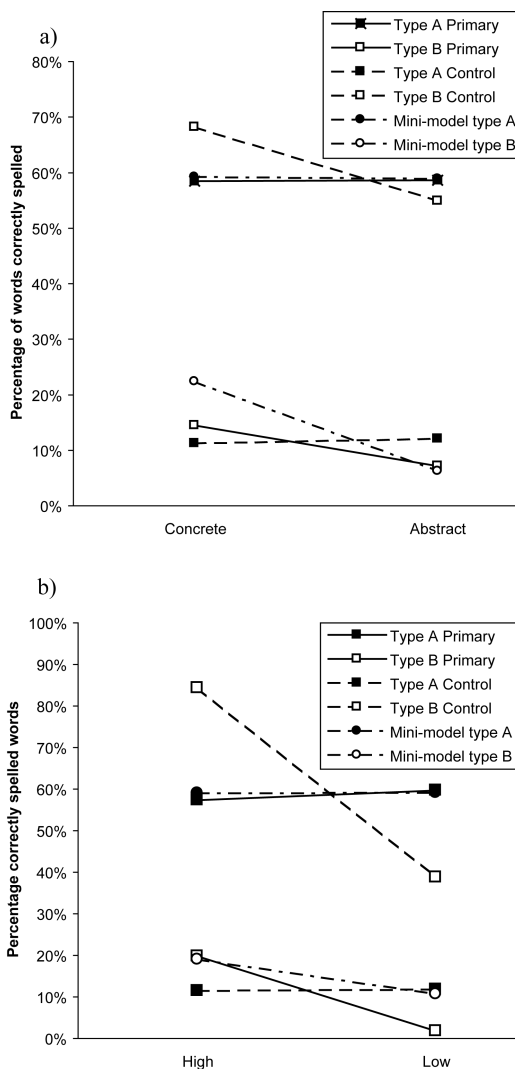


Figure 3. The performance of the models on high- and low-concreteness (a) and high- and low-frequency (b) words. The "mini" models vary word frequency during training of the GOB component and do not include a momentum term in the backpropagation learning procedures. They are trained on only 100 words.

The Type A models show no effect of either concreteness or frequency. This is expected as the intact semantic spelling route in these models produces essentially perfect activation patterns at the word representation field regardless of word frequency or concreteness. The GOB subsystem

thus operates with normal input regardless of either variable, and errors occur only as a result of sequencing inaccuracies introduced by noise in the competitive queuing process. By contrast, the Type B models show effects of both frequency and concreteness. Here errors result from sub-optimal input to an intact GOB, and as expected the effects of semantic route damage are greater for low-concreteness and low-frequency words, which produce more abnormal word representation patterns and hence higher error rates.

The effect of concreteness on the Type B models is clear in Figure 3(a), but the magnitude of the effect is less than that for patient DA (Cipolotti et al., 2004, Table 5). However no attempt has been made in the present model to accurately model the difference in semantic representation between high- and low-concreteness words. Rather, our representational scheme for word semantics aims to capture the qualitative difference between the two groups. Again for Figure 3(b), while the qualitative effect of word frequency is clear and accords with the effect in patient DA (Cipolotti et al., 2004, Table 7), the magnitude is rather low. However, we have represented word frequency extremely crudely in the model with the aim of capturing only the qualitative effect.

As discussed above, it was not practical to explore our preferred theoretical explanation for the lack of a frequency effect in Type A disorder in the full-scale model due to the need to have a momentum term in the learning rule in order to make training tractable. We have therefore not exposed the GOB model to word frequency differences during training. This is a highly unrealistic assumption. To test our preferred explanation we produced a small-scale version of the model, which is identical to the standard model except for the following features. A corpus of 100 words was used, with 25 words of each of four, five, six, and seven letters. A total of 12 words at each length were high frequency, 13 were low frequency, 12 words were of high concreteness, and 13 were of low concreteness, arranged so that all four permutations were covered at each word length as before. The words used were selected

from the larger corpus. Hidden Layer 2 was reduced to 50 nodes. The momentum term in the learning rules for both semantic and GOB sections of the model (see Appendix) was set to 0, and word frequency was represented in GOB training as well as semantic route training, using the same technique. The semantic network was trained to criterion on the full set of 400 input representations after 9,545 epochs. However, to reduce training time for the second and more computationally demanding simulation, only the first 12 low-concreteness and first 13 high-concreteness word identity representations at each word length were fed to the graphemic buffer network. The GOB network was trained until 200 consecutive correct epochs were achieved (the tighter criterion being due to the probabilistic occurrence of low-frequency words in each epoch). Without the momentum term a much larger number of training epochs were required (76,743). Type A and Type B primary models were produced (with a noise level of ± 0.0025 and with a scaling factor of 0.65 alongside background noise at ± 0.0005 , respectively). The results are shown in Figure 3 as “mini-model” (dashed lines).

Despite the fact that the GOB network has been exposed to high-frequency words three times as often as low-frequency words, there is no difference in the Type A error rate between high- and low-frequency words. By contrast the Type B model shows a clear effect of frequency. This is just what would be expected on our explanation in terms of the particular requirements of learning in CQ systems.

Serial position

Figure 4 shows the number of single-error responses produced at each serial position for each of the full-scale models. These figures are normalized across word lengths using the scheme of Wing and Baddeley (1980), which distributes errors for all word lengths to five notional serial positions.

The Type A models (Figure 4a) produce curves which are bowed, as typically are those for Type A patients (see Caramazza et al., 1987; Jonsdottir

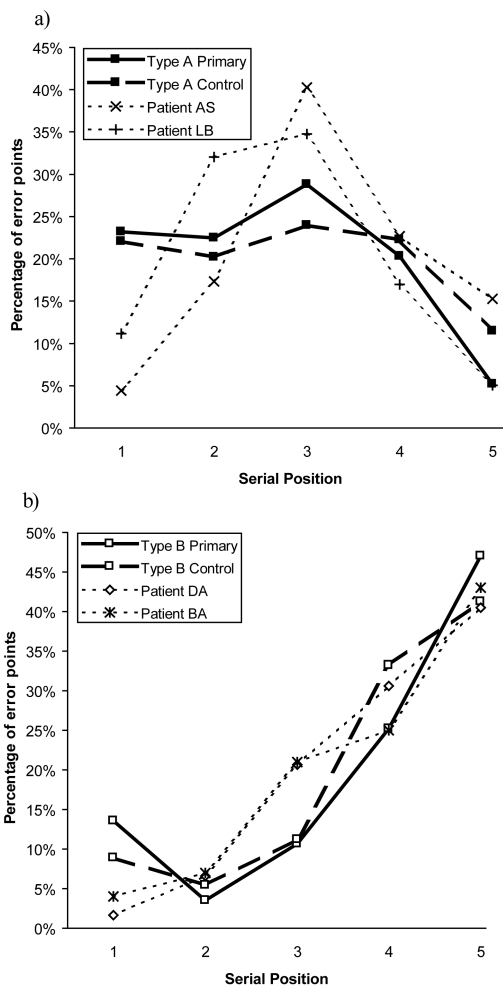


Figure 4. The incidence of errors occurring in each serial position in responses containing a single identifiable error, (a) for Type A models and patients and (b) for Type B models and patients. Position of error is determined according to the scoring scheme of Caramazza and Miceli (1990) and is normalized across word lengths by the method of Wing and Baddeley (1980).

et al., 1996). Serial curves showing primacy and recency effects—relatively preserved performance on items near the start and end of sequences—are typical of damaged CQ models. A number of factors contribute:

1. As sequencing progresses earlier sequence items are inhibited and thus removed (to some extent, depending on the degree of

inhibition) from the pool of available competitors. This tends to reduce errors late in a sequence.

2. In some models sequence items have higher absolute activation nearer the start of a sequence (this is not the case in the present model). This can lead to better separation of consecutive items and hence fewer errors.
3. In models with dynamic cueing of serial position (as in the present model) “end effects” arise where serial positions near the start and end of the sequence overlap in their representations with fewer adjacent positions than those in the middle of the sequence.

In the present Type A model Factors 1 and 3 are likely to contribute to the observed serial position effect.

The Type B models (Figure 4b), however, produce a monotonically increasing serial error curve. We interpret this as resulting from the different mix of error mechanisms that predominate when letter activations are reduced in strength. While the profiles of all error types change to some extent, errors involving omission of letters or the early stopping of sequencing due to subthreshold letter activation predominate in this situation, and the large number of late deletion errors is the main contributor to the altered serial error curve.

Figure 5 repeats this analysis for all errors, plotting the total percentage of incorrect letters at each serial position, again normalized across word lengths. This is the way that the serial position curves of errors made by Type B patients have typically been represented. Again bowed curves are produced for Type A, but monotonic increasing curves for Type B. The Type B patients who have been described show serial position curves of this type (see Cipolotti et al., 2004). The bowing of the Type A curves in Figure 5 is very shallow, however. We assume that the difference in degree of bowing between Figures 4 and 5 reflects differences in the distribution of complex errors (i.e., errors involving more than a single transposition, insertion, deletion, or substitution).




Figure 5. The percentage of letters produced in each serial position that were incorrect, normalized across word lengths by the method of Wing and Baddeley (1980).

Error types

Table 1 gives examples of the errors produced by the Type A and Type B models, and Figure 6 shows the proportions of different error types produced for six-letter words.

The model data in Figure 6 are restricted to errors on six-letter words and to responses including only a single apparent error to reduce the possible ambiguity in classification. The closest

Table 1. Sample errors from the primary Type A and B models

Error type	Examples from models
Insertion	"KEPT" → KEPDT (A)
	"ENVY" → ENVNY (A)
Deletion	"POWER" → POER (A)
	"VANITY" → VANTY (B)
Exchange	"MAIN" → MIAN (A)
	"PLAYING" → PYALING (A)
Substitution	"PERIOD" → PEDIOD (A)
	"STORY" → SCORY (B)
Fragment-correct	"SPIRIT" → SPIR (B)
	"CORD" → C (B)
Fragment-similar	"SCORN" → SO (B)
	"RADIO" → AI (B)
Fragment-unrelated	"ADORN" → I (B)
	"FLOAT" → SO (B)

Note: The model that produced the error is indicated as (A) or (B).

equivalent data for Type A patients LB and AS and Type B patients DA and BA are also provided. Some caution must be exercised in comparing these results for two reasons. Different studies have used slightly different criteria to classify error types (for example, they have included or excluded multiple-error responses). The distribution of error types can also be affected by the distribution of word lengths included in the study, as proportions of different error types can vary with word length. However, we can make some general observations.

Type A dysgraphic patients vary somewhat in the relative proportions of different error types. ML and DH (Hillis & Caramazza, 1989) and SE (Posteraro et al., 1988) show a predominance of deletions, while FV (Miceli et al., 1985), for example, produces relatively few deletions (10%). For LB (Caramazza & Miceli, 1990; Caramazza et al., 1987), CW (Cubelli, 1991), and JH (Kay & Hanley, 1994) substitutions are the most common error (64% of FV's errors are substitutions). HE (McCloskey et al., 1994) and AS (Jonsdottir et al., 1996) produce similar numbers of substitutions and deletions. Most patients also produce a fair number of insertion and transposition errors. For example, LB's errors include 6% insertions and 17% transpositions, while for AS the proportions are 22% and 14%, respectively. For those patients where the distinction has been made (LB, JH, and AS), transposition errors are predominantly exchanges rather than shifts, and indeed shift errors were very infrequent for these patients. Mixed errors, containing more than one of the error types listed above, also occur, with different incidences in different patients although such errors are generally at least as frequent as any of the individual error types.

The Type A models produce all four error types at this word length (deletion, exchange, insertion, and transposition errors). The Type B models produce deletion and substitution errors at this word length, though they also produce relatively low rates of insertion and transposition errors at other word lengths. Examining the qualitative error patterns of Figure 6 it is immediately clear that the two model types have different qualitative error profiles.

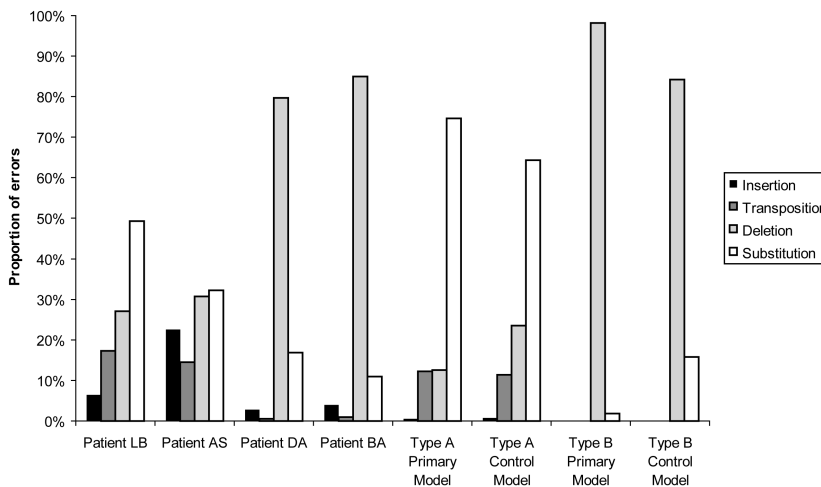


Figure 6. Distribution of different error types produced by the models (combining results for single-error responses and fragments over all word lengths) compared with similar analyses for Type A patients LB and AS and Type B patient DA.

The Type A models produce a pattern with insertions as the least common error type, transpositions the next common, deletions next, and substitutions most common. The same qualitative pattern is shown by Type A patient LB. Type A patient AS has the qualitative ranking of insertion and transposition errors reversed and has a more similar incidence of deletions and substitutions. The proportion of substitutions produced by the Type A models is somewhat higher than is usual among Type A patients, although it is comparable with that of FV, for example.

The Type B patients show a much higher incidence of deletion errors and few transpositions. The Type B models both also show very high incidence of deletion errors and low rates of transpositions. For both models the second ranking error type is substitutions, although the incidence is considerably lower than that for deletions. This also accords with patients DA and BA.

Although shift and exchange errors have been combined as “transpositions” in the figure, all of the models produce very low numbers of letter shift errors. This is an apparently robust feature of CQ models, and it is a feature of all Type A patients for whom the distinction has been made.

Ward and Romani (1998) introduced a new type of error, the fragment, in their analysis of the Type B dysgraphic patient BA. Fragments are errors that are two or more letters shorter than the target word. Type B patients BA and DA produce large numbers of fragment errors, whereas they are not a prominent feature of Type A dysgraphics. Table 2 shows the errors produced by each model that were classified as fragments using the same criteria as those for the analysis of DA’s errors. Fragments are classed as “correct” if there are only letters missing from the end of the word. Otherwise fragments are “structurally similar” if more than 50% of the letters in the response are present in the target word. All other fragments are classed as “unrelated”. (Examples of each fragment type from the model are given in Table 1.) Table 2 also includes comparable figures for Type B patients BA and DA.

There is a striking effect of locus of damage on the overall proportion of fragments. The Type A models produce few fragments; the Type B models produce many. Virtually all of the fragments produced by the Type A models are of the “similar” subtype, whereas the Type B models also produce many “correct” fragments.

Table 2. Errors produced by each model that were classified as fragments

Fragment type	Type A				Type B				Patient			
	Primary		Control		Primary		Control		BA		DA	
	No.	Percent.	No.	Percent.	No.	Percent.	No.	Percent.	No.	Percent.	No.	Percent.
Total	1,602	2	12,555	7	129,644	73	30,221	39	194	13	242	20
Correct	19	1	604	5	32,668	25	8,542	28	126	65	89	37
Similar	1,581	99	11,504	92	90,695	70	21,236	70	58	30	78	32
Unrelated	2	0	447	4	6,281	5	443	1	10	5	75	31
First letter correct	1,288	80	6,520	52	73,501	57	21,344	71	(Not reported)		148	61

Note: For comparison, fragment errors are also listed from Type B patients BA (Ward & Romani, 1998, Table 2) and DA (Cipolotti et al., 2004, Table 11). Single and multiple letter fragments are combined. Percent. = percentage of all errors.

The patients themselves do not appear to show a consistent pattern in the relative proportions of different fragment subtypes. The Type B models produce a low rate of “unrelated” fragments, which is similar to patient BA but is lower than DA. The rate of “correct” fragments is lower than either patient, though it is more comparable to DA than BA. Finally the model produces considerably more of the “similar” subtype than either patient. However the model does concur with the high rate of preservation of the first letter in fragment errors, and in fact the primary Type B model closely matches this figure for patient DA.

The explanation for fragment errors in the model is very simple. A pair of letter activation thresholds must be exceeded for spelling of a word to continue and for a particular letter to be articulated. Degraded input to the GOB leads to subnormal activation levels within the model, which in turn leads to more frequent omission of letters and the early ending of the spelling attempt. The large increase in deletion and fragment errors in the Type B model than in the Type A models follows straightforwardly. The large number of Type B model fragment errors that preserve the first letter of the word bears this explanation out, suggesting that stopping spelling too early is a common error. To check this we have confirmed in the Type B models that the number of fragments in which sequence generation appears to have stopped early falls off

monotonically, as one would expect, with decreasing fragment length (i.e., fragments in which the final letter only is deleted are more frequent than those in which the final two letters are deleted, which in turn are more common than those with the final three letters deleted).

Consonant/vowel status

The GOB subsystem of the model is forced to represent the consonant or vowel status of letters during training. This was done to investigate the effect on the error pattern after damage of shared features in the internal representations created in such a model. Preservation of domain-dependent constraints in sequencing errors is common in many types of serial behaviour, and we see the relative preservation of consonant/vowel (C/V) status in the errors of some Type A patients as a manifestation of this general principle. In any CQ system where some feature of the items being sequenced is shared across items within particular categories, errors will tend to preserve these categories. In the case of movement errors this is equivalent to a higher probability of confusion errors occurring between items of the same category (Glasspool & Houghton, 1997). It makes most sense to speak of C/V status being preserved in substitution and exchange errors, and Table 3 shows the proportion of these types of error that preserve consonant/vowel status in

Table 3. Percentages of substitution, exchange, and shift errors produced by the models in which consonant/vowel status is preserved

Type A		Type B	
Primary	Control	Primary	Control
93	91	100	91

each model. The sharing of C/V status information in the internal representations of letters in each model leads to a relative preservation of C/V status in errors, as expected.

This feature is included only to confirm that shared features can lead to preservation of domain-specific serial constraints in sequencing errors. We make no theoretical claim for the particular method that we use, and we have not attempted a quantitative fit to a particular patient.

GENERAL DISCUSSION

The disorder that we have termed “Type B graphemic buffer disorder” (see Cipolotti et al., 2004; Ward & Romani, 1998) has presented theorists with a paradox. Typical patients of this type—for example, BA and DA—have a pattern of errors that is, in most respects, qualitatively equivalent to that found in what was assumed to be the sole form of graphemic buffer disorder (Type A), of which prototypic patients are LB (Caramazza & Miceli, 1990; Caramazza et al., 1987) and AS (Jonsdottir et al., 1996). These patients are widely assumed to have a disorder at the level at which representations of letters are held in a buffer before being produced in writing or spelling aloud (Caramazza & Miceli, 1990; Caramazza et al., 1987) or at the roughly equivalent level in a connectionist architecture—that of letter form units (Houghton et al., 1994; Shallice et al., 1995). However, at the same time graphemic buffer Type B patients generally exhibit characteristics not normally associated with the GOB, such as the presence of semantic errors (as in deep dysgraphia) or more commonly an effect of word concreteness or part of speech on writing, or

both. Conversely, Cipolotti et al. (2004) argue that typical deep dysgraphic patients show GOB Type B characteristics. Moreover, the Type B GOB patients described in the literature also show as a frequent error type the fragment error (Ward & Romani, 1998), in which the patient produced a response two or more letters shorter than the word itself. Often these fragments begin with the same first letter as that of the target word.

In this paper we have begun the process of situating the CQ_model of the GOB within the context of the wider spelling system. Modelling the relationship between the semantic system and the GOB in this way results in a model which predicts the existence of the two forms of GOB disorder, A and B. Moreover, it also explains their most salient characteristics and similarities and differences between them.

The model also has a number of limitations, however, in particular relating to more fine-grained details of patient error characteristics. We begin, therefore, by discussing the major successes and limitations of the model and considering what the implications of the limitations might be for our theoretical position. Despite the complexity of the model we consider this work to be essentially preliminary, laying the foundations for more comprehensive modelling of the spelling system. It is especially important, therefore, to identify the theoretically significant features of the model with respect to the data against which it has been compared. We finish this section by identifying the main features that are critical to our account of GOB disorders Types A and B.

Successes of the model

The GOB element of the model has a form of competitive queuing architecture (Houghton, 1990) initially applied to spelling by Houghton et al. (1994) and Shallice et al. (1995). The major development in the model is that it uses distributed rather than localist internal representations. When noise is added to the activation levels of the competitive queuing mechanism the model reproduces qualitatively similar

characteristics to the localist form of the CQ model when noise is added at a similar point (Houghton et al., 1994). In particular the model reproduces the four characteristic error types of the empirical disorder Type A (GOB disorder), it shows a word length effect, and it shows preservation of consonant/vowel status in the errors.

A primary aim of the simulations that we have described was to test our prediction that degraded input to an intact CQ sequence generation mechanism would result in errors with the same general character as those produced by a damaged CQ mechanism. This is supported by the production of deletion, substitution, and exchange errors (and insertion errors in some situations, see below) by the Type B models. The effects of word length on error rate and preservation of C/V status are also present in the Type B models.

On this model not only can the GOB Type A pattern be explained by a damaged CQ network but also those features that it shares with the GOB Type B error pattern (and hence that it shares with the typical deep dysgraphic pattern) can be explained in terms of degraded input to the GOB. This is compatible with our hypothesis that deep dysgraphia generally leads to weak output from central semantic representations so that only a weak input reaches the graphemic buffer in writing. A further aim was to test the assumption that the fundamental differences between GOB Type A and Type B disorders can also be explained by the same basic account.

At a gross level these include the effect on error rate of word frequency and of semantic factors such as concreteness in Type B but not Type A disorder. On the model Type B disorder is the consequence of damage to the semantic-to-orthographic system, whereas Type A disorder does not involve damage to systems carrying semantic representations. This difference is thus naturally explained as the result of the semantic pathway locus of damage in Type B but not Type A. Indeed the simulation shows strong effects of concreteness and frequency when the input to graphemic output systems from semantic representations is degraded, but not when the graphemic buffer itself is damaged.

At a finer level three major differences were noted in the pattern of errors produced by Type A and Type B patients. Type B patients produce monotonically increasing rather than bowed serial error incidence curves, they produce a majority of deletion errors, and they produce a new type of error in addition to those characteristic of GOB Type A disorder—the fragment error. All three of these differences have been identified in the simulations presented above, and they have been explained as straightforward consequences of the operation of a disrupted CQ sequencing system versus an intact CQ system operating on degraded input. Most critically the effect of reduced input to the CQ mechanism leads to many fragment errors, which occur only rarely in the Type A models. On the model, fragments occur as a direct consequence of the reduced activation of letter level units. It should be noted that the assumption of Schiller et al. (2001) that in such patients fragment errors occur due to an excessively rapid decay in the graphemic buffer could be incorporated in a different form in the model. However, it would be redundant. The qualitative similarity between mild and severe versions of each type of simulated disorder in the model confirms that the observed differences are due to damage locus rather than severity.

The model therefore produces the following advances:

1. It allows us to explain why the pattern of GOB disorder errors—substitution, deletion, insertions, and transpositions—occurs in the context of at least two qualitatively different patterns of overall impairment.
2. It explains why one of these patterns should contain as a critical subcomponent the error pattern associated with “deep dysgraphia” and normally thought of as resulting from a more central locus of impairment than classical GOB disorder.
3. It explains the observed differences in the pattern of errors between the two types of disorder—specifically, the difference in serial error incidence, the difference in incidence of deletion errors, and the occurrence within the

GOB Type B pattern of an additional type of error—the fragment error.

In all but the most advanced and narrow domains of science a model is valuable if it captures in a succinct way important aspects of the empirical data set. Particularly when it is produced fairly early in the theoretical investigation of a domain it is unrealistic to expect a model to explain the details of all findings that are potentially relevant. However, to be useful it must be potentially extendable to deal with aspects of the data domain where there is a possible discrepancy. The current model has the strong advantage of succinctly capturing the areas of overlap and difference between GOB disorder and deep dysgraphia, which no previous model has done in such detail. There are, however, a number of limitations in the model's account for patient characteristics. How serious are its failures to capture empirical features, and where are the areas where its account of the data needs to be extended?

Limitations of the model

We can identify four main areas where the model displays significant limitations in its account for patient characteristics:

1. An inaccurate quantitative, and in some cases qualitative, match to the error patterns of patients.
2. An incomplete account for semantic errors.
3. No account for nonword spelling.
4. No account for errors involving double letters.

The first issue listed is the most important from the point of view of the theoretical adequacy of our approach, as it is the issue for which it is least clear whether modelling simplifications or problems with the underlying theory are to blame.

1. An inaccurate match to patient error patterns

As our aims for this model, outlined above, were rather broad we have not at this stage paid great attention to anything beyond the most general fit to empirical data. The model does in fact show a number of quantitative and in some cases

qualitative deviations from the data. We have in most cases made no attempt to quantitatively model patient performance, so it is the qualitative mismatches that are of most concern. There are four main problems with the simulation results in this regard. First, the models aim to capture “prototypic” Type A and Type B patients, but as noted earlier a number of patients exist that appear to combine features of both types, or are otherwise nontypical. The models do not capture the error patterns of these patients. Second, the qualitative fit of the proportions of errors of different types in both typical A and typical B patients is imperfect. Third, the qualitative fit of the proportions of different types of fragment errors to those of Type B patients is imperfect. Finally the Type A models produce only weakly bowed serial error curves when all incorrect letters are taken into account.

A number of patients reported in the literature do not fall cleanly into either Type A or Type B as we have characterized them. We have proposed that, since we relate Type A and B disorders to orthogonal types of damage, they represent end points of a continuum of possibilities corresponding to cases in which both types of damage may be present to some degree. We have not yet attempted to simulate such cases so we cannot categorically say what the predictions of the model would be. However, it is important to consider the ways in which these atypical patients diverge from the model's predictions. We begin by summarizing the model's predictions for the “prototypical” end-point A and B cases. With internal disruption to the competitive queuing process in the GOB the model predicts effects of word length, bowed serial error curves, letter-sequencing errors with substitutions, deletions, transpositions, and insertions, and a tendency to preserve consonant–vowel status in errors. Spelling accuracy is not sensitive to word frequency or concreteness. This we consider the model's prototypic Type A profile. With degraded input to an intact GOB the model again predicts word length effects, consonant–vowel preservation, and letter-sequencing errors of the same general types as the Type A case. However, the

proportion of letter deletion errors is much greater, and a new error, the fragment, occurs in substantial numbers. Additionally the serial error plot changes to a monotonically increasing curve, and spelling errors are sensitive to word frequency and semantic factors such as concreteness. This is the model's predicted Type B profile.

Turning to patients apparently intermediate between these profiles, the majority of these fall into three classes. The most common are those essentially of Type A but which show a significant effect of word frequency or other lexical features (e.g., patients AS, JH, DH, JES, and HE). There appear to be two ways in which one might account for this within the present model. Most straightforwardly one could propose minor disruption to the semantic representation reaching the GOB (mild Type B disorder) combined with more serious disruption to the serial output stage (strong Type A disorder), in line with our "continuum" proposal.

Alternatively, if our assumptions that representations at the level of "word identity" are entirely free from semantic content, and that the GOB portion of the model is trained entirely to ceiling so that word frequency does not affect the strength of representations in this subsystem, do not hold completely for all patients, semantic content or word frequency would affect the error rate of the pure Type A model. In the present model we arrange for consonant–vowel status to be represented in the connection weights activating letter words, and this leads to a bias in the competition for output resulting in relative preservation of CV status. Features such as word frequency, age of acquisition, and concreteness, if they are represented at all in the system that activates letter representations in the GOB, may in the same way lead to biases in the output competition and effects on error patterns. In the present model we have assumed a "classical" interpretation of Type A GBD, with no lexical influences on activations at this level of the model. Sage and Ellis (2004) argue that such features, as well as the number of near "orthographic neighbours" that a word has, do often influence spelling in Type A GBD, and the present model has the

advantage that such effects could be readily accounted for by relaxing the assumption of freedom from lexical influence in GOB representations. We would expect that relaxing the requirement that the GOB model be trained to ceiling would allow these influences to be felt in the Type A simulations.

Less common are patients who show anomalous serial position curves, either patients apparently of Type B with bowed serial error curves (for example FM and AM) or apparently of Type A but with a flat or monotonic serial position curve (e.g., FV and GSI, Miceli et al., 2004). The most obvious way in which the model might account for these patterns is through a combination of Type A and B damage. As this has not been simulated we cannot easily predict its effect on the model.

A third set of anomalous patients are apparently of Type B but produce no semantic errors (HR, AZO, and GSI). The present model simulates Type B disorder through degraded input to GOB. However, this degraded input might have a number of causes, some of which (for example, damage to projections from semantics to GOB) would not lead to semantic errors. Such patients might then correspond to damaged transfer of representations from an intact semantic system to an intact GOB system. In the absence of a full simulation of the semantic system (see below) the present model does not distinguish this case from a damaged semantic system.

None of these more complex situations has yet been simulated, so intermediate patients of these types remain an important issue for the model.

Two further types of patient are relevant to the theoretical position that we have taken. First, we claim that degraded input to the GOB from the semantic system necessarily results in Type B GOB symptoms. However, many types of semantic disorder involve no such dysgraphic symptoms. This would be critical to our account if it were assumed that the presence of semantic processing problems necessarily led to degraded input to the GOB. However, other types of input to the GOB may compensate (the phonological spelling route may be operating correctly, for example),

and additionally there may be ways in which semantic processing could be disrupted while presenting a superficially normal output to the GOB (a fully activated pattern of activation representing an incorrect word, for example). A failure mode of this type is in fact a prediction of the Plaut and Shallice (1993) model of lexical semantic representation. A central aspect of that model is the emergence of “attractors”, which, under some types of damage, clean up errors in semantic to word-identity mapping so that a known word is always produced rather than, for example, a blend of words. Damage to the model can result in degradation of the output signal by comparison with the intact model; however, some types of damage can allow the attractor system to “clean up” a disrupted internal representation to produce a normal output, albeit possibly representing an incorrect word. This type of output from the semantic processing system would allow the GOB as we have modelled it to operate normally resulting in deep dysgraphic symptoms in the absence of GOB Type B errors. While the review of Cipolotti et al. (2004) suggests that this is not the norm in deep dysgraphia, patient RCM (Hillis, Rapp, & Caramazza, 1999) appears to be close to this pattern. This patient produced many (56%) whole-word substitution errors, but relatively few (21%) spelling errors that resulted in nonwords. We would expect this type of error pattern under some types of damage to a model of the present type combined with a Plaut and Shallice type of lexical semantic system. However, a considerably more complex simulation will be necessary to test this aspect of the model's predictions. Similarly, our position that the Type B error profile is associated with degraded input to the GOB does not necessarily imply a locus for the corresponding damage. It would seem that a lexical-semantic locus is common in such cases, but some patients with aspects of the Type B profile show no indication of lexical-semantic damage (e.g., GSI). Any damage affecting the transmission of information into the GOB would affect performance in this way, for example.

Second, we claim that certain key features of the error pattern in Types A and B GBD follow

directly from the compromise of the GOB, whether through damage or degraded input. These include the production of letter sequence errors including insertions, omissions, transpositions, and substitutions, and the effect of word length on accuracy of spelling. Patients in which the former pattern appeared in the absence of the latter would challenge this position.

In fact a few patients have been reported who apparently show letter-sequencing errors in the absence of word length effects. However, we are aware of no patient of this type where there is no indication that a substantial proportion of errors arise in systems other than GOB (such as phonological spelling or allographic conversion). Thus Hillis, Chang, Breese, and Heidler (2004) report three dysgraphic patients who produce spelling errors resulting in nonwords in the absence of significant effects of word length. However, these patients all show a large proportion of phonologically plausible errors (62–90%). This suggests that they have preserved phonological spelling systems, which can be assumed to support input to the GOB. The two that were tested also showed many stroke-related errors, evidence of damage to allographic conversion processes subsequent to the GOB. It appears possible that these patients' spelling errors do not involve abnormal operation of the GOB, which would be consistent with the absence of word length effects. RCM (Hillis et al., 1999) also shows no effect of word length on spelling accuracy. However, as discussed above this patient appears to fit the pattern of disrupted lexical semantic representations that are largely “cleaned up” prior to the GOB, and again the majority of the patient's spelling errors may not involve abnormal operation of the GOB.

Turning now to the distribution of different error types (substitution, insertion, deletion, and transposition), this is the least stable feature of the model and is easily affected by changes in free parameters. Viewed positively this accords with the relative variability of the error distribution across different patients. The variability of patients makes it less easy to directly compare the models with data. However, certain common features appear to be present across a majority of

patients. One of these, the universally low rate of shift errors, is robustly shown by the model (and by previous CQ spelling models). The most obvious quantitative discrepancy is the low rate of insertion errors in the simulations that we have reported (Figure 6). However, other types of manipulation to the existing model affect this feature. For example, reducing the value of the response threshold T_r from 0.8 to 0.7, during recall only, leads to higher rates of insertion errors on shorter words. It is, however, possible that an additional error mechanism may be at work in patients, which is not present in the model.

As yet the different mechanisms by which errors can occur within a simple CQ sequence generator are not well characterized (see Glasspool, 1998, 2005, for a preliminary discussion). This makes it difficult to predict in detail the effects of small changes to a model on the proportions of different error types produced. A number of factors influence the balance between error types. For example, the relative rate of insertions, transpositions, and substitutions that is produced on the model depends on specific details of the implementation of the impairment as well as the core assumptions of the model. It is likely to depend on:

1. How noise is added—whether it changes by the letter or by the word.
2. The distribution of noise—exponential, normal, or rectangular distributions, for example.
3. The number of thresholds used. Two activation thresholds are used in the simulation. There are two mechanisms for deletion errors in the model, one involving failure of any letter nodes to exceed one or both of the activation thresholds (for letter articulation or continuation of spelling) and another that is more basic to the CQ approach in which initially an anticipation error occurs, and then spelling continues without the letter that was omitted in the initial error ever being produced. Using only a single activation threshold would affect the first of these mechanisms, producing a reduction in the rate of deletion errors (the

manipulation of the response threshold mentioned above has a similar effect and increases the rate of insertions relative to deletions). Various factors such as the slope of the letter activation gradient can affect the incidence of the second type of deletion (Glasspool, 1998).

The discrepancies between model and patients in the proportions of different types of error are a concern for the CQ approach. However, we do not believe that it would be appropriate to argue from this alone that the model is flawed, for two reasons. First, the model does not tightly constrain performance on this feature. Performance is affected by changes to free parameters of the model, whereas other features of the approach are much more stable (see Glasspool, 1998, and Glasspool & Houghton, 2005, for discussion). Second, factors outside the model can be expected to have a relatively large impact on the finer details of error proportions. For example, this aspect of the data is likely to be affected most by strategic factors such as the patients' tendency to guess.

A relatively clear general account exists on the model for fragment errors. However, the relative incidence of different subtypes of fragment errors is not predicted by that account. Indeed the rate of correct and similar fragments is reversed between the model and patient BA. Again, though, patients are variable on this measure, and we would expect the same range of factors as that mentioned above to influence this feature. However the mechanism for fragment errors in the model does provide a straightforward explanation for two effects seen in the patients: the qualitative difference in incidence of fragment errors between Type A and Type B disorder, and the tendency for letters nearer the start of the word to be preserved in these errors (observed in both DA and BA). These appear to be somewhat more gross features than the relative proportions of correct and similar fragments, which is what we would expect to see if the discrepancy were due to the crude nature of our simulation of this type of error.

Finally, the Type A models produce only weakly bowed serial error curves when all incorrect

letters are taken into account. Previous CQ models have demonstrated strongly bowed serial error curves, both in spelling and other domains (including models of Type A GOB, e.g., Houghton et al., 1994), so this does not appear to be a limitation of the CQ approach per se. It is possible that the move to a multilayer network architecture has in some way affected this aspect of the model's operation. However, we do not yet understand the cause of this difference.

2. Incomplete account for semantic errors

From the perspective of modelling Type B dysgraphia the most serious omission in the present model is the lack of a full model for the semantic spelling route. Our general assumption has been that a system of the type advanced by Plaut and Shallice (1993) would provide this function.

While it is beyond the scope of this paper to combine two complex models, we would expect that full simulation of the semantic system would produce similar results to the present model with the exception that some errors would be incorrect but intact words semantically related to the target. However, the manner in which output from a damaged Plaut/Shallice type network is degraded may differ in detail from that assumed in the present model, and further simulations will be required to confirm that this does not materially affect our overall account.

3. No account for nonword spelling

The present model does not attempt to cover the entire spelling system. The main omission is the lack of a route via phonology, which could be used to spell nonwords. However, we do not believe that incorporating a phonological spelling route would impact on the general account provided by the present model.

The extension of the CQ approach to a multilayer network architecture allows us to propose a simpler account for the integration of semantic and phonological routes at the level of the GOB than was possible in our previous localist spelling models (Shallice et al., 1995). The multilayer framework has the potential to learn a generalized

mapping from a phonemic representation to a spelling output, allowing the model to incorporate a phoneme-to-grapheme conversion system as well as the graphemic output lexicon. Two options are possible—the phonological system might project to Hidden Layer 1, in which case the word identity field would represent a combined phonological and semantic orthographic lexicon, or the word representation field could be split into semantic and phonological sections, each receiving input via their own hidden layers from semantic or phonological representations. The phonological representation might represent the entire word and be held constant throughout spelling, or it could be a dynamic pattern such as the shifting window of models such as NETspell (Olson & Caramazza, 1994). We leave exploration of these possibilities for future work.

4. No account for spelling with geminate (double) letters

Our previous modelling work (e.g., Glasspool, 1998; Houghton et al., 1994; Shallice et al., 1995) as well as other CQ modelling work (e.g., Rumelhart & Norman, 1982) has demonstrated the general effectiveness of a separate geminate representation in modelling the phenomena associated with doublings in serial behaviour, including spelling. The arguments supporting this approach hold good for the current model, and we do not believe it will be problematic in principle to add a separate geminate representation to the current GOB model.

The issue of repeated letters

Ward and Romani (1998) consider a CQ sequencing explanation for the GOB-type errors of the Type B patient BA. They conclude that the approach provides a promising model for this aspect of the deficit, with one reservation concerning spelling of words with repeated letters. A *prima facie* prediction of the CQ approach is that repeated items in sequences will cause difficulties, as one might expect that the temporary inhibition of items following their production should make the second occurrence of a repeated

item more error prone than a nonrepeated letter in the same sequence position. Ward and Romani conclude that words with repeated letters (e.g., *fence*) should be more error prone than words with no repeats (e.g., *lance*) on any model of the CQ type. Patient BA, though, shows no difference in spelling performance on this variable.

However, the prediction from any moderately complex CQ model is unlikely to be so straightforward. The most basic form of CQ mechanism with a static activation gradient is unable to repeat items within a sequence, so most implemented models allow the gradient of activations over the “queue” of responses to vary over time. When an item is to be repeated its representation may then receive additional excitement later in the sequence in order to overcome its temporary inhibition (e.g., Houghton, 1990). This is the basic form of the model presented here. When a letter is repeated temporary inhibition is balanced by an increase in excitation, and it is not clear that repeated letters will necessarily be disadvantaged (although *doubled* letters—that is, immediate repeats—do seem to be difficult for such models to represent unaided). The effect of repeated letters depends on the balance between these two influences.

The present model has the additional feature that it learns to spell during repeated attempts. It thus has the opportunity to learn to compensate for any additional fragility in repeated letters during the training process. Based on these considerations we did not expect a strong effect of repeated letters on the present model. To test this expectation the models' responses to high-frequency, high-concreteness words were sorted according to the presence or absence of repeated letters in the target word. The percentage of correct spellings at each word length was averaged to correct for word length effects. The mean results were 60% correct (with repeated letters) versus 55% correct (no repeated letters) for the Type A primary model and 21% correct (repeats) versus 24% correct (no repeats) for the Type B primary model. We conclude that similar performance on both word types in a patient does not present a difficulty for the model.

Critical features of the model

The computational model that we have advanced has several components and is moderately complex. Some complexity is necessary in order to produce a self-contained process model in a nontrivial domain. However, the essence of our approach is straightforward, so it is useful to consider which features of the model are critical to our theoretical position and which are present simply in order to provide a concrete realization of that position in a workable model.

To summarize our theoretical position, we consider Type A disorder to be due to disruption of the operation of a sequence generation component working according to the principles of competitive queuing (the disruption being equivalent to increased stochastic variability in the selection of the “winning” letter at each serial position), so that intact word identity information reaches a damaged serial output system. Type B disorder involves damage to an “upstream” activating system so that the word identity information is degraded, while the serial output mechanism itself is undamaged. The similarities between the two disorders arise because in both cases the operation of a CQ output system is disrupted. Some of the differences between the two disorders are due to the different locus of damage to the system (the differential effects of frequency and semantic class arise because damaged areas in Type B, but not Type A, are sensitive to these features), while others are due to differences in the type of errors typical of a CQ system under noise disruption compared with degraded input (the different serial error incidence curves, the occurrence of fragment errors, and the different profile of error types).

We can identify seven features of the model that are critical with respect to this account:

1. *The output of the semantic spelling route is a representation of word identity that is free from semantic content.* This is implied by the fact that Type A errors are insensitive to semantic class. Presumably then a level exists in the spelling system where such information is not present, and the proposed location for

this representation is at the last point that it could reasonably be expected: at the input to the final sequence generation stage in the shared output path of the system. The model's explanation for the lack of semantic class sensitivity in Type A disorder depends on this feature.

2. *The semantic spelling route, when damaged, produces degraded output sensitive to word frequency and to semantic class.* This is necessary to the explanation for the frequency and semantic class sensitivity of Type B errors.
3. *The final shared output from the spelling system is a sequence generation system operating according to the principles of competitive queuing.*
4. *Damage to the CQ output system has the effect of introducing uncertainty in letter selection.* These two assumptions are critical to our explanation for the types of error common to the two disorders, for the effects of word length and serial position, and for the differences in these features between the two disorders.
5. *The CQ output system is not sensitive to word frequency when damaged.* This feature is required to explain lack of frequency sensitivity typical of Type A disorder. It is implied empirically by the existence of patients producing GBD type errors but with no sensitivity to frequency. This feature follows from the type of learning required in a CQ system.
6. *The CQ system represents consonant/vowel status of letters (and any other linguistic information subject to observed regularities in the errors of both Type A and Type B dysgraphics, see Sage & Ellis, 2004).* The model implies that any regularities in the errors of both types of dysgraphic must arise from information represented in the gradient of activations of the CQ system. The effect of such regularities has been demonstrated in the model by forcing the representation of consonant/vowel information, but any information leading to systematic biasing of letter activations in this part of the model will have similar effects (notably the tendency for letters of the same "class" to participate jointly in errors). Any such information must be present in the word ID field or in

representations or connections downstream of that point.

7. *The CQ output system stops sequence generation when no letter is activated above a threshold.* This mechanism for stopping sequence production has been proposed on independent grounds as the most elegant means of stopping sequence production in a CQ model (e.g., Glasspool, 1998; Houghton, 1990,). This is critical to the explanation for fragment errors in Type B but not Type A disorder. The present model also incorporates a response threshold for letter production, which leads to omission rather than stopping on subthreshold letters. This assumption appears less critical to the approach but contributes to the finer level performance, in particular the proportion of deletion errors.

The extension of CQ to a multilayer network offers a number of methodological advantages, although it also suffers from disadvantages (e.g., the long training time required for backpropagation learning algorithms). It also has a number of effects on the fine level performance of the model. However, we do not consider that it is critical to the theoretical position advanced here. There is no fundamental reason why a localist model incorporating the critical features outlined above could not offer essentially the same explanation for Type A and B disorders. The multilayer approach allows the separation of the timing signal from the word identity representation. This is interesting from the point of view of processing requirements and may have implications for physiological plausibility. The approach may also offer a more tractable solution to the problem of integrating semantic and assembled spelling routes at the level of the GOB. However these features are not critical to the explanation of Type A and B syndromes.

CONCLUSION

Previously deep dysgraphia and graphemic buffer disorder have been seen as two distinct syndromes.

The present model supports the conclusion that there is an overlap: Certain critical systems are involved in both. This is consistent with the observation (Cipolotti et al., 2004) of similarities in the error patterns in the two syndromes. It should be noted, however, that the idea of examining in that paper whether deep dysgraphic patients also make graphemic buffer type errors only followed the analysis of the properties of the model.

On the model, what differs between the two types of patient is whether an intact output mechanism operates on degraded input, or a damaged output mechanism operates on intact input. On this view the Type A-like properties that Cipolotti et al. (2004) find in Type B disorder are a consequence of the presence of a CQ sequencing mechanism, which receives input from a damaged system.

There is perhaps a wider implication for computer simulation in neuropsychology. The internal operation of a particular element in a cognitive model may have a profound effect on the predictions following from damage to the model. We would argue that process models of the operation of modules and their interactions are required in addition to abstract reasoning about high-level "box-and-arrow" models of cognitive architecture, or about the type and structure of information that is represented within them.

We have identified a number of limitations of the model in its present form. In particular it does not provide a close quantitative fit to patient data in all areas, and it does not give a full qualitative account for some patients described in the literature. Further simulation work is required to confirm that these cases can be accommodated within the general theoretical framework that we propose. This will involve the development of a more accurate model of the spelling system, which will in turn require a careful investigation of the influence of lexical parameters such as word concreteness and frequency at different points in the system, and it will no doubt also require a better understanding of the mechanisms of error in CQ models. If these issues can be addressed then we believe that the benefits of widening the modelling enterprise to encompass

multiple subsystems, and their interaction, will be substantial.

Manuscript received 20 November 2002

Revised manuscript received 19 July 2005

Revised manuscript accepted 20 July 2005

PrEview proof published online December 2005

REFERENCES

- Aliminos, D., McCloskey, M., Goodman-Schulman, R., & Sokol, S. M. (1993). Remediation of acquired dysgraphia as a technique for testing interpretations of deficits. *Aphasiology*, *7*, 55–69.
- Bub, D., & Kertesz, A. (1982). Deep agraphia. *Brain and Language*, *17*, 146–165.
- Bullinaria, J. A., & Chater, N. (1995). Connectionist modelling: Implications for cognitive neuropsychology. *Language and Cognitive Processes*, *10*, 227–264.
- Burgess, N., & Hitch, G. J. (1992). Towards a network model of the articulatory loop. *Journal of Memory and Language*, *31*, 429–460.
- Burgess, N., & Hitch, G. J. (1996). A connectionist model of STM for serial order. In S. E. Gathercole (Ed.), *Models of short-term memory* (pp. 51–72). Hove, UK: Psychology Press.
- Caramazza, A., & McCloskey, M. (1991). The poverty of methodology. *Behavioural and Brain Sciences*, *14*, 444–445.
- Caramazza, A., & Miceli, G. (1990). The structure of graphemic representations. *Cognition*, *37*, 243–297.
- Caramazza, A., Miceli, G., Villa, G., & Romani, C. (1987). The role of the graphemic buffer in spelling: Evidence from a case of acquired dysgraphia. *Cognition*, *26*, 59–85.
- Christiansen, M. H. (1997). Improving learning and generalisation in neural networks through the acquisition of multiple related functions. In J. Bullinaria, D. Glasspool, & G. Houghton (Eds.), *Connectionist representations. Proceedings of the 4th Neural Computation and Psychology Workshop* (pp. 58–70). London, Springer-Verlag.
- Cipolotti, L., Bird, C., Glasspool, D., & Shallice, T. S. (2004). The impact of deep dysgraphia on graphemic output buffer disorders. *Neurocase*, *10*(6), 405–419.
- Coltheart, M. (1981). The MRC Psycholinguistic Database. *Quarterly Journal of Experimental Psychology*, *33A*, 497–505.

- Cooper, R., & Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive Neuropsychology*, *17*, 297–338.
- Cubelli, R. (1991). A selective deficit for writing vowels in acquired dysgraphia. *Nature*, *353*, 258–260.
- De Partz, M.-P. (1995). Deficit of the graphemic buffer: Effects of a written lexical segmentation strategy. *Neuropsychological Rehabilitation*, *5*, 129–147.
- Ellis, A. W. (1982). Spelling and writing (and reading and speaking). In A. W. Ellis (Ed.), *Normality and pathology in cognitive function* (pp. 251–275). London: Academic Press.
- Ellis, A. W. (1984). *Reading writing and dyslexia: A cognitive analysis*. Hove, UK: Lawrence Erlbaum Associates Ltd.
- Glasspool, D. W. (1998). *Modelling serial order in behaviour: Studies of spelling*. Unpublished Ph.D. thesis, University College London, London.
- Glasspool, D. W. (2005). Modelling serial order in behaviour: Evidence from performance slips. In G. Houghton (Ed.), *Connectionist models in psychology* (pp. 241–270). Hove, UK: Psychology Press.
- Glasspool, D. W., & Houghton, G. (1997). Dynamic representation of structural constraints in models of serial behaviour. In J. Bullinaria, D. Glasspool, & G. Houghton (Eds.), *Connectionist representations. Proceedings of the 4th Neural Computation and Psychology Workshop* (pp. 269–282). London, Springer-Verlag.
- Glasspool, D. W., & Houghton, G. (in press). Serial order and consonant-vowel structure in a model of disordered spelling. *Brain and Language*.
- Glasspool, D. W., Houghton, G., & Shallice, T. (1995). Interactions between knowledge sources in a dual-route connectionist model of spelling. In L. S. Smith & P. J. B. Hancock (Eds.), *Neural computation and psychology* (pp. 209–226). London: Springer-Verlag.
- Glasspool, D. W., Shallice, T., & Cipolotti, L. (1999). Neuropsychologically plausible sequence generation. In D. Heinke, G. W. Humphreys, & A. Olson (Eds.), *Connectionist models in cognitive neuroscience* (pp. 40–51). London: Springer-Verlag.
- Hartley, T., & Houghton, G. (1996). A linguistically constrained model of short-term memory for nonwords. *Journal of Memory and Language*, *35*, 1–31.
- Hillis, A. E., & Caramazza, A. (1989). The graphemic buffer and attentional mechanisms. *Brain & Language*, *36*, 208–235.
- Hillis, A. E., Chang, S., Breese, E., & Heidler, J. (2004). The crucial role of posterior frontal regions in modality specific components of the spelling process. *Neurocase*, *10*, 175–187.
- Hillis, A. E., Rapp, B. C., & Caramazza, A. (1999). When a rose is a rose in speech but a tulip in writing. *Cortex*, *35*, 337–356.
- Hinton, G. E. (1989). Connectionist learning procedures. *Artificial Intelligence*, *40*, 185–234.
- Houghton, G. (1990). The problem of serial order: A neural network model of sequence learning and recall. In R. Dale, C. Mellish, & M. Zock (Eds.), *Current research in natural language generation* (pp. 287–319). London: Academic Press.
- Houghton, G. (1994). *Some formal variations on the theme of competitive queueing* (Internal Technical Report, UCL-PSY-CQ1). University College London, Department of Psychology.
- Houghton, G., Glasspool, D., & Shallice, T. (1994). Spelling and serial recall: Insights from a competitive queueing model. In G. D. A. Brown & N. C. Ellis (Eds.), *Handbook of spelling: Theory, process and intervention* (pp. 365–404). Chichester, UK: John Wiley and Sons.
- Houghton, G., & Hartley, T. (1996). Parallel models of serial behaviour: Lashley revisited. *PSYCHE*, *2*(25). Retrieved from <http://psyche.cs.monash.edu.au/v2/psyche-2-25-houghton.html>
- Jonsdottir, M., Shallice, T., & Wise, R. (1996). Language-specific differences in graphemic buffer disorder. *Cognition*, *59*, 169–197.
- Jordan, M., (1986). Attractor dynamics and parallelism in a connectionist sequential machine. *Proceedings of the 8th Annual Conference of the Cognitive Science Society* (pp. 10–17). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Katz, R. (1991). Limited retention of information in the graphemic buffer. *Cortex*, *27*, 111–119.
- Kay, J., & Hanley, R. (1994). Peripheral disorders of spelling: The role of the graphemic buffer. In G. D. A. Brown & N. C. Ellis (Eds.), *Handbook of spelling: Theory, process and intervention* (pp. 295–315). Chichester, UK: John Wiley and Sons.
- Lashley, K. S. (1951). The problem of serial order in behaviour. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior* (pp. 341–392). New York: Wiley.
- Lennox, C., & Siegel, L. S. (1994). The role of phonological and orthographic processes in learning to spell. In G. D. A. Brown & N. C. Ellis (Eds.), *Handbook of spelling: Theory, process and intervention* (pp. 93–109). Chichester, UK: John Wiley and Sons.
- Margolin, D. I. (1984). The neuropsychology of writing and spelling. Semantic, phonological, motor and processes. *Quarterly Journal of Experimental Psychology*, *36A*, 459–489.

- McCloskey, M., Badecker, W., Goodman-Shulman, R., & Aliminosa, D. (1994). The structure of graphemic representations in spelling: Evidence from a case of acquired dysgraphia. *Cognitive Neuropsychology*, *11*, 341–392.
- McCloskey, M., & Caramazza, A. (1991). On crude data and impoverished theory. *Behavioural and Brain Sciences*, *14*, 453–454.
- Miceli, G., Benvegnu, B., Capasso, R., & Caramazza, A. (1995). Selective deficit in processing double letters. *Cortex*, *31*, 161–171.
- Miceli, G., Capasso, R., Benvegnu, B., & Caramazza, A. (2004). The categorical distinction of vowel and consonant representations: Evidence from dysgraphia. *Neurocase*, *10*, 109–121.
- Miceli, G., Capasso, R., Ivella, A., & Caramazza, A. (1997). Acquired dysgraphia in alphabetic and stenographic handwriting. *Cortex*, *33*, 355–367.
- Miceli, G., Silveri, M. C., & Caramazza, A. (1985). Cognitive analysis of a case of pure dysgraphia. *Brain and Language*, *25*, 187–196.
- Olson, A., & Caramazza, A. (1994). Representation and connectionist models: The NETspell experience. In G. D. A. Brown & N. C. Ellis (Eds.), *Handbook of spelling: Theory, process and intervention* (pp. 337–363). Chichester, UK: John Wiley and Sons.
- Plaut, D. (1995). Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology*, *17*, 291–321.
- Plaut, D. C., & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, *10*, 377–500.
- Posteraro, L., Zinelli, P., & Mazzucchi, A. (1988). Selective impairment of the graphemic buffer in acquired dysgraphia: A case study. *Brain and Language*, *35*, 274–286.
- Quinlan, P. T. (1993). *The Oxford Psycholinguistic Database*. Oxford, UK: Oxford University Press.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by back-propagating errors. *Nature*, *323*, 533–536.
- Rumelhart, D. E., & Norman, D. A. (1982). Simulating a skilled typist: A study of skilled cognitive-motor performance. *Cognitive Science*, *6*, 1–36.
- Sage, K., & Ellis, A. W. (2004). Lexical influences in graphemic buffer disorder. *Cognitive Neuropsychology*, *21*, 381–400.
- Schiller, N. O., Greenhall, J. A., Shelton, J. R., & Caramazza, A. (2001). Serial order effects in spelling errors: Evidence from two dysgraphic patients. *Neurocase*, *7*, 1–14.
- Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge, UK: Cambridge University Press.
- Shallice, T., Glasspool, D. W., & Houghton, G. (1995). Can neuropsychological evidence inform connectionist modelling? Analyses of spelling. *Language and Cognitive Processes*, *10*, 195–225.
- Tainturier, M.-J., & Caramazza, A. (1996). The status of double letters in graphemic representations. *Journal of Memory and Language*, *35*, 53–73.
- Ward, J., & Romani, C. (1998). Serial position effects and lexical activation in spelling: Evidence from a single case study. *Neurocase*, *4*, 189–206.
- Wing, A. M., & Baddeley, A. D. (1980). Spelling errors in handwriting: A corpus and a distributional analysis. In U. Frith (Ed.), *Cognitive processes in spelling* (pp. 251–285). London: Academic Press.

APPENDIX

Formal description of the model

The operation of the network can be separated into two passes. During the forward pass, an activation pattern is applied to the input layer, and activation propagates forward to the output layer. The backward pass operates only during training, when an error signal is propagated back from the output layer to adjust weights in proportion to their contribution to the error. The two subnetworks of the model are trained separately. During training of the semantic network, the semantic input layer is treated as the input layer and the word identity layer is treated as the output. During training of the GOB network, the word identity field and the position field are treated as the input layer and the letter layer is treated as the output. During

recall the model is treated as a single network, with the word identity field operating as a hidden layer. The following applies equally to both subnetworks except where otherwise indicated.

Forward pass

Nodes in input fields have their activation levels set to 1.0 or 0.0 according to the current input pattern for that field. The net input $net_i(t)$ to node i in a hidden or output layer at time step t is given by:

$$net_i(t) = \sum_{j=1}^n A_j(t)W_{ji} \quad (1)$$

where A_j is the activation of node j in the previous layer, and W_{ji} is the weight from node j to node i . The function f is the logistic function standardly used in backpropagation networks (Rumelhart, Hinton, & Williams, 1986):

$$f(x) = \frac{1}{e^{-x} + 1} \quad (2)$$

As well as receiving input from nodes in the previous layer, each node in hidden and output layers also receives a bias input, which may be thought of as an additional weight from a unit that is permanently set to an activation of 1.0.

The activity $A_i(t)$ of node i in a hidden layer at time step t is given by:

$$A_i(t) = f(\text{net}_i(t)) \quad (3)$$

The C and V nodes also obey Equation 3.

Letter nodes recover slowly from inhibition. The activation $A_i(t)$ of letter node i at time t is given by:

$$A_i(t) = \begin{cases} \text{net}_i(t) + \eta & \text{if } A_i(t-1) \geq 0 \\ \text{net}_i(t) + rA_i(t-1) + \eta & \text{otherwise} \end{cases} \quad (4)$$

where η is a noise value drawn from a uniform distribution around 0 (actual values used in simulations are given in the text) and r is a parameter that governs the rate of recovery from inhibition.

During recall the most active letter node is determined by a simulated competitive filter, and the activation level of this letter node is then set to a standard negative (inhibited) activation level, *Inb*. During learning the letter that should be produced in the current position is inhibited, regardless of which letter actually has the highest activation level (this immediate correction of errors during learning prevents errors from disrupting subsequent correctly learned letters: see Houghton, Glasspool, & Shallice, 1994).

Backward pass

Hinton (1989) shows that in a network where binary-valued output vectors are desired, and real-valued output vectors may be interpreted as probability distributions over binary vectors (the CQ noisy selection procedure is straightforwardly interpretable in this way), the appropriate error measure to use in a backpropagation training procedure is the cross-entropy, C , between the desired and actual probability distributions. The cross-entropy between an actual probability vector A , with elements a_i , and a desired probability vector D , with elements d_i , is given by:

$$C = - \sum_i d_i \log_2(a_i) + (1 - d_i) \log_2(1 - a_i) \quad (5)$$

When used in a backpropagation algorithm the derivative of C is multiplied by the derivative of the logistic function, and this product reduces to the difference between the desired and actual outputs (Hinton, 1989). An error value δ_u is thus generated for each output layer unit:

$$\delta_u = (d_u - a_u) \quad (6)$$

where d_u is the desired activation value for output unit u and a_u is the actual value. The desired values d_u are generated according to a “lazy” learning rule as follows. In determining the desired value for a letter node during spelling we distinguish five cases:

1. The most active letter node corresponds to the target letter for the current position within the word, and its activation exceeds the response threshold T_r .
2. The most active letter node corresponds to the target letter, but its activation does not exceed T_r .
3. The most active letter node is not the target letter for the current position within the word.
4. The current letter position is past the end of the word, and the most active letter node does not exceed the stopping threshold T_s .
5. The current letter position is past the end of the word but the most active letter node exceeds T_s .

In Cases 1 and 4 we score the letter position as “correct” for the purposes of monitoring learning. However, weight changes are only made in the “error” cases, Cases 2, 3, and 5. In Cases 2 and 3 the target letter is insufficiently active and must be reinforced: In these cases the desired activation for the target node is set to 1.0. In Cases 3 and 5 the incorrect winning letter is too active; in these cases the desired activation for the winning node is set to 0.0. For all other letter nodes the desired activation is set to the actual activation (i.e., no weight changes are made). A small training margin m_T is added to threshold T_r and subtracted from T_s during training. This means that the model is trained against slightly more stringent thresholds than those that are used in recall, which prevents undue fragility of recall due simply to the model operating very close to its thresholds. An equivalent effect could be achieved over a longer training period if a small amount of random noise were present in activation levels or thresholds during training.

The C and V nodes are treated more conventionally since they do not participate in the CQ process: The desired activation for these nodes is simply set to 1.0 if the current target letter is a consonant or vowel, respectively, and 0.0 otherwise.

A weight change is calculated for the weights from a hidden layer to an output layer according to:

$$\Delta W_{ub} = \varepsilon \delta_u A_b \quad (7)$$

where ΔW_{ub} is the required change in the weight from hidden unit b to output unit u , A_b is the activation of hidden unit b , and ε is a small constant, the “learning rate”.

The error value δ_b for each hidden unit b is given by:

$$\delta_b = A_b(1 - A_b) \sum_u \delta_u W_{ub} \quad (8)$$

The weight change for the input-to-hidden layer weights is then:

$$\Delta W_{bi} = \varepsilon \delta_b A_i \quad (9)$$

where ΔW_{bi} is the required change in the weight from input unit i to hidden unit b , and A_i is the activation of input unit i . For each weight W in the network, the weight changes are now applied using:

$$W(t) = W(t-1) + \Delta W(t) + m\Delta W(t-1) \quad (10)$$

where $W(t)$ is the new weight value for time step t , $W(t-1)$ is the value of the weight at the previous time step, $\Delta W(t)$ is the required weight change for the current time step, $\Delta W(t-1)$ is

Table A1. Parameter values used in simulations unless otherwise stated

Parameter	Symbol	Value
Learning rate	ε	0.001
Momentum	m	0.9
Letter node recovery rate	r	0.8
Letter node inhibition level	Inb	-1.0
Response threshold	T_r	0.8
Stopping threshold	T_s	0.6
Training margin	m_T	0.025

the required weight change calculated on the previous time step, and m ($0 < m < 1$) is the momentum value. At the start of learning, all weights $W(t)$ are set to random values between ± 0.5 and all $\Delta W(t-1)$ are assumed to be 0.

Table A1 gives the parameter values used in the simulations.