Language and Cognitive Processes

Morphological decomposition based on the analysis of orthography

Kathleen Rastle [a]; Matthew H. Davis [b]

[a] Royal Holloway University of London, London, UK [b] MRC Cognition and Brain Sciences Unit, Cambridge,
UK

First Published:November2008

PLEASE SCROLL DOWN FOR ARTICLE

Psychology Press
Taylor & Francis Group

Morphological decomposition based on the analysis of orthography

Kathleen Rastle

Royal Holloway University of London, London, UK

Matthew H. Davis

MRC Cognition and Brain Sciences Unit, Cambridge, UK

Recent theories of morphological processing have been dominated by the notion that morphologically complex words are decomposed into their constituents on the basis of their semantic properties. In this article we argue that the weight of evidence now suggests that the recognition of morphologically complex words begins with a rapid morphemic segmentation based solely on the analysis of orthography. Following a review of this evidence, we discuss the characteristics of this form of decomposition, speculate on what its purpose might be, consider how it might be learned in the developing reader, and describe what is known of its neural bases. Our discussion ends by reflecting on how evidence for semantically based decomposition might be (re)interpreted in the context of the orthographically based form of decomposition that we have described.

One of the key topics in research on visual word processing over the past 30 years has concerned the recognition of words comprising more than one morpheme (e.g., trusty, untrusting, distrust). Though there is wide agreement that such words are 'decomposed' into their constituent morphemes during visual word perception (e.g., 'distrust' is segmented into {dis-} + {trust}), there is less consensus on precisely how or when this decomposition is

Correspondence should be addressed to Kathleen Rastle, Department of Psychology, Royal Holloway University of London, Egham, Surrey TW20 0EX, UK. E-mail: Kathy. Rastle@rhul.ac.uk or Matthew H. Davis, MRC Cognition & Brain Sciences Unit, 15 Chaucer Road, Cambridge, CB2 7EF, UK. E-mail: matt.davis@mrc-cbu.cam.ac.uk

achieved. The earliest theoretical account (Taft, 1981; Taft & Forster, 1975) considered morphological decomposition to be achieved through the analysis of sublexical orthographic information, such that it would be applied indiscriminately to affixed (e.g., repaint) and pseudoaffixed (e.g., restore) words alike. Originally formulated in the context of a search theory of visual word recognition, this account was later reformulated (Taft, 1994) so that it could be expressed in terms of the influential interactive-activation model (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982). In spite of this advance, however, the tide soon turned toward an understanding of morphological decomposition as a higher-level phenomenon guided by semantic knowledge (see Marslen-Wilson, Tyler, Waksler, & Older, 1994). Bolstered by theoretical insights from distributed-connectionist modelling (Davis, van Casteren, & Marslen-Wilson, 2003; Plaut & Gonnerman, 2000; Rueckl & Raveh, 1999), this conceptualisation of morphological decomposition subsequently dominated the next decade.

Our aim in writing this article is to argue that it is time to return to a theory in which the recognition of morphologically complex words begins with a rapid morphemic segmentation based purely on the analysis of orthography. In building this case, we begin by reviewing a recent yet substantial body of literature demonstrating that the recognition system rapidly decomposes *any* printed stimulus that has the appearance of morphological complexity, irrespective of whether or not that stimulus is semantically related to its stem. Our discussion then turns to the characteristics of this form of decomposition (hereafter, referred to as 'morpho-orthographic' decomposition), to some hypotheses about what purpose it might serve in visual word recognition, and to an examination of how morpho-orthographic decomposition might be learned by the developing reader. Following a discussion of what is known of the neural bases of morpho-orthographic decomposition, we close by considering how evidence for semantically based decomposition might be (re)interpreted in the light of the evidence for the orthographically based form of decomposition that we describe.

BEHAVIOURAL EVIDENCE FOR MORPHO-ORTHOGRAPHIC DECOMPOSITION

Morphemes are defined as 'minimal meaning-bearing units'. They allow us to express a vast range of concepts with a much smaller range of orthographic or phonological units, and provide to us our most productive means of creating new words (e.g., George Bush's recent claim 'I'm the *decider* and I decide what's best'). It is not surprising, therefore, that theories of morphological processing introduced over the past 10 years or so have

generally seen decomposition in the context of meaning (e.g., Davis et al., 2003; Giraudo & Grainger, 2000; Gonnerman, Seidenberg, & Andersen, 2007; Marslen-Wilson et al., 1994; Plaut & Gonnerman, 2000; Rueckl & Raveh, 1999).

These theories (two of which are described visually in Figure 1) claim that decomposition is applied only to *some* morphologically structured letter strings – namely, those that are semantically transparent (e.g., unbeatable). Distributed-connectionist theories, for example, propose that complex words are represented componentially in the learned internal representations mediating orthography and semantics (e.g., the distributed representation of 'darkness' overlaps that of 'dark'; see Rueckl & Raveh, 1999). However, these componential representations develop only to the extent that the complex word is related in meaning to its stem. Morphologically structured words that have no relationship to their stems (i.e., pseudomorphological constructions like 'corner') or that have a historical relationship to their stems that is no longer apparent (i.e., opaque constructions like 'witness') have representations that are unlike those of their stems in these models (Plaut & Gonnerman, 2000). Similarly, the supralexical theory of Giraudo and Grainger (2000) posits that local morphemic representations act as an interface between orthographic representations of whole words and representations of their meanings. Thus, this theory also claims that morphologically complex words are decomposed only if they are related in meaning to their stems, with morphologically structured words that have no semantic relationship with their stems being represented as full forms in the
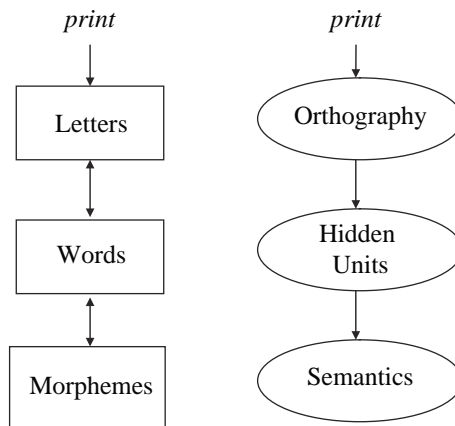

Figure 1. Examples of semantically constrained theories of morphological decomposition. The left panel is based on the supralexical theory of Giraudo and Grainger (2000) while the right panel is based on various distributed-connectionist theories (e.g., Davis et al., 2003; Plaut & Gonnerman, 2000; Rueckl & Raveh, 1999).

morphemic layer. For the rest of this article, we refer to this latter type of morphologically structured stimulus as 'opaque', irrespective of whether it has a historical relationship with its stem or not.

These semantically based theories of morphological decomposition have derived support from a variety of tasks including cross-modal priming (Gonnerman et al., 2007; Longtin, Segui, & Hallé, 2003; Marslen-Wilson et al., 1994; Meunier & Longtin 2007), visual priming with fully-visible primes (Rastle, Davis, Marslen-Wilson, & Tyler, 2000), long-lag priming (e.g., Drews & Zwitserlood, 1995; Marslen-Wilson & Zhou, 1999; Rueckl, Aicher, & Yovanovich, 2008 this issue), and unprimed lexical decision (Ford, Marslen-Wilson, & Davis, 2003; Schreuder & Baayen, 1997). For example, Marslen-Wilson et al. (1994; see also Longtin et al., 2003) demonstrated that robust cross-modal priming effects are observed for semantically related morphological relatives (e.g., hunter-hunt) but not for opaque morphological relatives (e.g., gingerly-ginger). Similar effects are apparent in visual priming with fully visible primes: robust priming for semantically related morphological relatives but no priming for opaque morphological relatives (Rastle et al., 2000).

The problem with these theories is that they fail to explain the pattern of morphological priming effects observed in the context of masked priming. Key studies on this topic were reported by Longtin et al. (2003) and by Rastle, Davis, and New (2004). Critical to both studies was the comparison of masked priming effects for semantically related morphological relatives (e.g., darkness-DARK), for prime-target pairs that had an opaque morphological relationship (e.g., corner-CORN), and for prime-target pairs that had a non-morphological form relationship (e.g., brothel-BROTH; –el never functions as a suffix in English). Results showed robust and equivalent masked priming effects against an unrelated baseline for both of the conditions in which primes were morphologically structured, and critically, that these effects were significantly larger than those obtained from orthographic overlap alone. These results suggest that the 'darkness' and 'corner' primes were being analysed in terms of their apparent morphemic constituents, thus enabling savings in the recognition of their respective targets. Semantically based theories of morphological processing have no explanation for these results since these theories claim that opaque words are never decomposed into their constituent morphemes. On these theories, the opaque primes should have produced effects of a similar magnitude to those produced by the non-morphological form primes.

It is always possible that the materials in these studies were unsatisfactory in some respect (e.g., that some confound existed across the manipulation of priming condition), so it is fortunate that the weight of evidence demonstrating non-semantic morphological effects in masked priming is now much greater than a couple of experiments. Table 1 summarises the results of every

TABLE 1

Studies investigating masked priming of opaque morphological relatives against masked priming of semantically-transparent morphological relatives and/or non-morphological masked form priming.

Article	Language	Prime duration	Transp. related (darkness-DARK)	Transp. unrelated (freedom-DARK)	Opaque related (corner-CORN)	Opaque unrelated (banker-CORN)	Form related (brothel-BROTH)	Form unrelated (warfare-BROTH)	Transp. priming	Opaque priming	Form priming
Kazanina et al. (2008)	Russian	59 ms	619	663	638	689	684	672	44	51	−12
Marslen-Wilson et al. (2008)	English	36 ms	495	513	507	528	525	539	18	21	14
Marslen-Wilson et al. (2008)	English	48 ms	493	529	513	536	539	547	36	23	9
Gold & Rastle (2007)	English	30 ms			571	589	582	589		18	7
McCormick et al. (2008, Exp 4)	English	42 ms	597	617	618	636	620	623	20	18	3
Lavric et al. (2007)	English	42 ms	650	682	675	700	714	723	32	25	9
Morris et al. (2007)*	English	50 ms	626	669	655	682	675	676	43	27	1
Diependaele et al. (2005)*	Dutch	53 ms	602	628	625	623	640	623	26	−2	−17
Diependaele et al. (2005)	French	40 ms	560	581	574	566	579	583	21	−8	4
Devlin et al. (2004)	English	33 ms	605	631	606	631			26	25	
Feldman et al. (2004)	English	48 ms	733	747	727	747			14	20	
Rastle et al. (2004)	English	42 ms	570	597	598	620	635	639	27	22	4
Longtin et al. (2003, Exp 1)**	French	46 ms	612	650	611	646	698	672	38	35	−26
Rastle & Davis (2003, Exp 1a)	English	52 ms	574	614	573	614			40	41	
Rastle & Davis (2003, Exp 1b)	English	52 ms			641	659	652	659		18	7
Rastle & Davis (2003, Exp 2a)	English	52 ms	563	593	571	593			30	22	
Rastle & Davis (2003, Exp 2b)	English	52 ms			601	623	619	623		22	4
Rastle et al. (2000, Exp 1)	English	43 ms	561	607	582	617	594	613	46	35	19
Feldman & Soltano (1999)	English	48 ms	unreported	unreported	unreported	unreported			19	23	
								Average	**30**	**23**	**2**

Transp. = Transparent. * These studies included a further backward mask that separated prime and target

** Opaque RTs are averaged across the opaque and pseudosuffixed conditions in this study

study of an Indo-European language that has examined masked priming of opaque morphological relatives against masked priming of semantically related morphological relatives and/or against non-morphological masked form priming. It includes only those studies in which primes were displayed for less than 60 ms, as there is mounting evidence that a semantically based form of decomposition becomes evident when primes become partially or fully visible (see Rastle et al., 2000, and the next section). Morphologically structured primes in these studies all comprised a stem plus a suffix, except in the case of Kazanina, Dukova-Zheleva, Geber, Kharlamov, & Tonciulescu (2008), in which these primes comprised a stem plus multiple suffixes. Further, the stem–suffix combinations used in the opaque primes in most of these studies were not constrained by syntactic legality (i.e., they contained both syntactically legal morphemic combinations like {whisk} + {-er} and syntactically illegal morphemic combinations like {quest} + {ion}). Longtin et al. (2003) reported that these two types of prime yield masked priming effects of the same magnitude.

Overall, the pattern of data closely follows the results of Rastle et al. (2004), with Diependaele, Sandra, and Grainger (2005) being the only outlier.[1] Priming effects yielded by morphologically structured words that have no semantic relation to their stems (e.g., corner-CORN) are of approximately the same magnitude as priming effects yielded by morphologically structured words that are semantically related to their stems (e.g., darkness-DARK). No priming effects are observed when primes comprise the target plus some non-morphological ending (e.g., brothel-BROTH), rendering it highly unlikely that the effects observed with morphologically structured pairs are due to simple orthographic overlap between prime and target. It seems from the data in Table 1 that masked morphological priming effects emerge whenever a morphologically structured prime *appears* to have a morphological relationship with its target. This evidence suggests strongly that there is a form of morphological decomposition that is based on orthographic rather than semantic information.

One potential limitation of the data in Table 1 is that they deal only with prime stimuli that have a {stem} + {suffix} structure, thus leaving open the possibility that morpho-orthographic decomposition is a process specific to suffixed items. Though we cannot conclusively rule out this possibility, there

---

[1] In their second experiment, Diependaele et al. (2005) attempted to replicate the findings of Longtin et al. (2003). Their study used substantially similar stimuli and a comparable prime duration to that of Longtin et al. (2003) but found priming effects only for morphologically related pairs that were also semantically related. It is unclear why Diependaele et al. (2005) failed to replicate Longtin et al. (2003). However, one potentially important difference between these studies was that primes and targets were repeated several times each in the Diependaele et al. study whereas they appeared only once in the Longtin et al. study.

is increasing evidence that morpho-orthographic decomposition is a more general process. One important line of evidence for this claim comes from studies of the recognition of semantically transparent (e.g., carwash) and semantically opaque (e.g., mayhem) compound words. Eye-movement studies in English (Frisson, Niswander-Klement, & Pollatsek, 2008) and in Finnish (Pollatsek & Hyönä, 2005) have consistently shown no effect of semantic transparency on the processing of such items. Further, related research has shown that while compounds with letter transpositions within morphemes (e.g., sunhsine) serve as effective primes for the recognition of non-transposed targets (e.g., sunshine), compounds with letter transpositions across morpheme boundaries (e.g., susnhine) do not. Critically, this holds for semantically transparent and semantically opaque compounds alike. Overall, though further research is needed to draw a definitive conclusion, these data are suggestive that morpho-orthographic decomposition is a general process that applies to any stimulus that has a morphological structure.[2]

CHARACTERISTICS OF MORPHO-ORTHOGRAPHIC DECOMPOSITION

Though this evidence has been obtained fairly recently, some of the key functional characteristics of morpho-orthographic decomposition have emerged already. One is that it appears to be a sublexical phenomenon (i.e., it applies to stimuli irrespective of their lexical status). The evidence for this locus is based on masked priming studies conducted by Longtin and Meunier (2005) investigating the decomposition of morphologically structured French pseudowords. Longtin and Meunier (2005) reported that the masked priming effects yielded by morphologically structured pseudowords (e.g., darkism-DARK) were of the same magnitude as those yielded by semantically transparent derived words (e.g., darkly-DARK). This was the case irrespective of whether the pseudoword primes formed syntactically legal (e.g., quickify-QUICK) or syntactically illegal (e.g., sportation-SPORT) combinations. Similar effects did not arise when primes were pseudowords comprising a stem plus a non-morphological ending (e.g., canalast-CANAL), suggesting that the facilitation observed for morphologically structured pseudowords was not the result of simple orthographic overlap.

---

[2] Further suggestive evidence for this conclusion comes from research conducted by Forster and Azuma (2000), who reported masked priming effects for prefixed items that shared a bound stem (e.g., debate-REBATE). These priming effects were significantly greater than those obtained for simple orthographic overlap (e.g., shallow-FOLLOW), potentially implicating a process of morpho-orthographic decomposition. Unfortunately, some of the prime-target pairs in the experiment were semantically related to one another (e.g., survive-REVIVE; demote-PROMOTE), making it difficult to rule out a (partial) semantic locus for the effects observed.

These data support the view that morphological decomposition is a process that is applied to *all* morphologically structured stimuli, irrespective of their lexical, semantic, or syntactic characteristics.

It also appears that morpho-orthographic decomposition is a phenomenon that arises early in visual word recognition. This claim is based on evidence that decomposition of opaque words seems to be restricted to masked priming situations in which derived words are presented so briefly that they are unavailable for conscious report. Priming from opaque derivations is generally not apparent in situations in which primes are fully perceptible such as cross-modal priming (e.g., Gonnerman et al., 2007; Longtin et al., 2003; Marslen-Wilson et al., 1994), visual priming with fully visible primes (Rastle et al., 2000), or long-lag priming (Drews & Zwitserlood, 1995; Rueckl et al., 2008 this issue; but see Bozic, Marslen-Wilson, Stamatakis, Davis, & Tyler, 2007). Similarly, while the syntactic legality of morphologically structured pseudowords has no influence on masked priming effects (Longtin & Meunier, 2005), it does have an impact on the magnitude of cross-modal priming effects (Meunier & Longtin, 2007). Though it may be impossible to map prime duration onto a precise description of the time course of recognition (a prime presentation duration of 40 ms need not imply that decomposition occurs within 40 ms of stimulus presentation), the fact that evidence for morpho-orthographic decomposition falls away with increasing prime duration (see Rastle et al., 2000) makes us comfortable in concluding that we are dealing with a process that occurs relatively rapidly in visual word perception.

The influence of prime perceptibility on the pattern of morphological priming effects is important for at least two further reasons. The first is that it allows us to be reasonably confident that the robust masked priming effects produced by opaque morphological constructions (e.g., corner-CORN) are not the result of strategic processes. Indeed, if one wishes to make an argument that these priming effects are the result of some strategy (e.g., repetition of suffixes, characteristics of nonwords, etc.), then one also has to explain why that strategy is not at work under the very conditions (i.e., long exposures) that strategic processes are most likely to arise. The second important point is that it differentiates the pattern of data observed in Indo-European languages from that observed in Semitic languages. Like in the Indo-European languages, robust masked priming effects are observed for morphologically related words with no semantic relationship in Hebrew (e.g., Frost, Forster, & Deutsch, 1997) and in Arabic (e.g., Boudelaa & Marslen-Wilson, 2001). However, unlike in the Indo-European languages, effects of semantic transparency on priming do *not* emerge with increasing prime perceptibility in the Semitic languages (Boudelaa & Marslen-Wilson, 2005; Frost, Deutsch, Gilboa, Tannenbaum, & Marslen-Wilson, 2000). It is as yet unknown whether this difference from Indo-European languages reflects

morphological richness (Plaut & Gonnerman, 2000), non-concatenative morphology (Boudelaa & Marslen-Wilson, 2001) or some other distinctive property of the Semitic languages.

Though morpho-orthographic decomposition appears to be a sublexical phenomenon that arises early in the time course of recognition, it also seems to be a 'smart' process that survives the regular orthographic alterations found in complex words. This claim is based on studies conducted by McCormick, Rastle, & Davis (2008) that investigated masked morphological priming effects using primes that could not be parsed straightforwardly into their morphemic constituents because of a missing 'e' (e.g., adorable-ADORE), a shared 'e' (e.g., writer-WRITE), or a duplicated consonant (e.g., metallic-METAL) at the morpheme boundary. Results of their experiments showed that masked priming effects observed under these conditions were of the same magnitude as those observed when primes could be parsed perfectly into their constituents. Results of their fourth experiment demonstrated that this robustness to orthographic alteration also applies to the decomposition of opaque words. Opaque prime-target pairs such as fetish-FETE that consist of a regular orthographic alteration yield robust priming effects that are significantly greater than those produced by simple orthographic overlap (e.g., blister-BLISS). Once again, this result provides support for a form of morphological decomposition that is insensitive to the semantic characteristics of complex stimuli.

WHAT IS MORPHO-ORTHOGRAPHIC DECOMPOSITION FOR?

Consistent with the earliest models of morphological processing (Taft & Forster, 1975; Taft, 1981), the results from studies described in Table 1 suggest that morphological decomposition is applied indiscriminately to any stimulus that has the appearance of morphological complexity. Thus, stimuli like 'corner' are decomposed into their constituents (e.g., {corn} + {-er}) in visual word perception, despite the fact that {corn} + {-er} is not a syntactically legal morphemic combination (nouns cannot take the suffix –er), and indeed that segmenting this stimulus leads to an incorrect semantic interpretation (i.e., a corner is not someone who corns) that must yield a processing cost. Even though stimuli like 'corner' are relatively few in number, it is somewhat difficult to understand why the recognition system would develop in a manner that would allow these kinds of 'processing mistakes' to occur. How, then, might we characterise the function of morpho-orthographic decomposition?

The simple answer, of course, is that morpho-orthographic segmentation constitutes an efficient computational process that allows rapid access to the meanings of morphologically structured stimuli *most of the time*. Perhaps a

more detailed means of expressing this is provided through insights from distributed-connectionist modelling (e.g., Davis et al., 2003; Plaut & Gonnerman, 2000; Rueckl & Raveh, 1999). Morphology in these models consists of learned correlations across the largely arbitrary mapping between orthography and meaning. Morphemes form 'islands of regularity' (Rastle et al., 2000) in this mapping because (a) the meanings of groups of letters corresponding to stems are usually preserved in their derivations (e.g., the meaning of 'design' is preserved in 'designer', 'redesign', etc.); and (b) groups of letters corresponding to affixes alter the meanings of stems in consistent ways (e.g., -less denotes 'without' when applied to a stem as in 'ageless', 'fearless', and 'passionless'). Distributed-connectionist networks are sensitive to these regularities and thus develop componential representations for morphologically complex words in the internal units mediating orthographic and semantic representations (Davis et al., 2003; Rueckl & Raveh, 1999).

However, in order to discover form-meaning regularities, these networks must have an input representation that allows them to recognise ortho-graphic similarity across particular sets of semantically related words. For example, the network must be able to discover that the semantically related words 'trusty', 'distrust', and 'untrustworthy' share significant orthographic overlap in the form of the stem 'trust'. This is a non-trivial task on current theories of orthographic input coding. On left-aligned slot-based coding (e.g., Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001; Grainger & Jacobs, 1996), for example, these derived words share no orthographic overlap whatsoever. Further, while they are more similar to one another on relative coding schemes such as Wickelcoding (e.g., Seidenberg & McClelland, 1989), open-bigram coding (e.g., Schoonbaert & Grainger, 2004), and spatial coding (e.g., Davis, 1999), the orthographic codes representing 'trust' are nevertheless non-identical. To the extent that the input representation for the stem 'trust' is non-identical in different contexts, then, the network will learn the form-meaning mapping for this stem independently in each different context and will hence fail to activate the appropriate meaning representa-tion when presented with a novel complex word that includes the stem 'trust'.[3]

This alignment problem has thus far been dealt with in simulations of the orthography-to-semantics mapping (e.g., Davis et al., 2003; Plaut & Gonner-man, 2000; Rueckl & Raveh, 1999) by providing the network with an input representation that is *already segmented* into its morphemic constituents. Essentially, these modellers assumed that a morpho-orthographic segmenta-tion had already taken place prior to the processing stages simulated in the

---

[3] This is one version of the translation invariance problem that is widely acknowledged in the computational literature on visual object recognition. For an introduction to this problem and description of neural network solutions see Plunkett and Elman (1997, Chapter 7).

model, thus ensuring that the input representations for stems across morphologically related words (e.g., trusty, distrust) were identical. By implication from this modelling work, we suggest that one function of morpho-orthographic decomposition may be to allow the developing reader to discover the morphological regularities that characterise the mapping between orthography and meaning (see also Rastle & Davis, 2003). Of course, the fact that a morphemically structured orthographic representation may be required to discover the morphological regularities across the orthography-to-semantics mapping naturally begs the question of how orthographic representations become morphemically structured in the first place (see e.g., Plaut & McClelland, 2000). Thus, we now turn to some hypotheses about the acquisition of morpho-orthographic knowledge in beginning readers.

THE ACQUISITION OF MORPHO-ORTHOGRAPHIC KNOWLEDGE

The data summarised in Table 1 and the modelling work described in the previous section both suggest that the initial visual processing of printed words requires segmentation into morphemic constituents. This segmentation could be modelled in a localist interactive-activation framework in a manner similar to that proposed by Taft (1994) in which sublexical morphemic units are contacted prior to the activation of whole-word orthographic units. Similarly, it might be modelled in a distributed-connectionist framework in terms of componential representations in the orthographic layer (e.g., in which the distributed orthographic representation of 'corn' overlaps that of 'corner'). These potential models of morpho-orthographic decomposition are depicted visually in Figure 2.

However one chooses to model morpho-orthographic decomposition, though, a complete theory of this phenomenon will require an account of how readers come to acquire morphemically structured orthographic representations. This is in itself a substantial computational problem because visual presentations of written words do not come pre-marked with orthographic cues to identify morpheme boundaries. How, then, are readers to acquire knowledge of morphemes without knowing the locations of morpheme boundaries?

Fortunately, morphologists are not alone in being faced with this computational problem. The literature on the segmentation and identification of words in connected speech has long grappled with an analogous problem both for adult performance (i.e., how do listeners recognise words in connected speech given the paucity of bottom-up segmentation cues in speech?; e.g., Davis, Marslen-Wilson, & Gaskell, 2002) and during
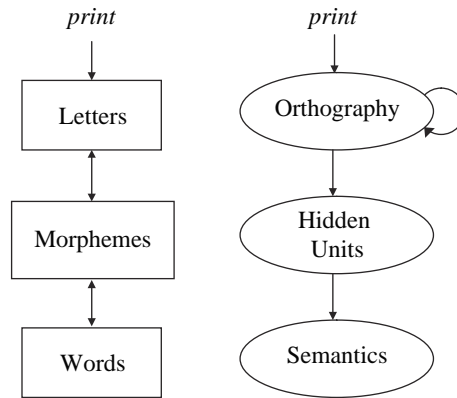
Figure 2. Potential theories of morpho-orthographic decomposition. The left panel is based on the sublexical theory of Taft (1994). The right panel is based on a potential distributed-connectionist theory in which morpho-orthographic representations are learned through recurrent connections in the orthographic layer. These connections are trained using a learning algorithm and thus become sensitive to the sequential dependencies of letter strings (see Moscoso del Prado Martín et al. 2004b, for simulations using a similar method).

development (i.e., how do infants learn words without knowing the location of word boundaries?; e.g., Jusczyk, 1997). Our focus here is on the second of these problems: How is it that the developing reader learns what letter sequences form morphological units in written text? Various learning strategies have been proposed in the domain of speech segmentation with recent evidence favouring accounts that include a combination of these (see Brent, 1999; Christianson, Allen, & Seidenberg, 1998; Davis, 2003, for reviews). We review three of these strategies that we believe might also permit the reading system to discover appropriately segmented orthographic representations for complex words. These strategies include (a) marking low probability sequences as containing boundaries; (b) grouping high-probability sequences into single units; or (c) discovering units that provide regularities in the orthography-to-semantics mapping.

MARKING LOW-PROBABILITY SEQUENCES AS CONTAINING BOUNDARIES

One method by which readers could acquire morpho-orthographic knowledge is through the analysis of sequential probabilities of letter combinations in printed text (e.g., bigram or trigram troughs, Seidenberg, 1987; see also Rastle et al., 2004). Statistical and connectionist implementations of these n-gram methods (cf. Elman, 1990) are highly effective at finding word boundaries in child-directed speech (Cairns, Shillcock, Chater, & Levy,

1997; Christiansen et al., 1998). Further, both adults and 8-month-old infants use these sequential probabilities in segmenting words from connected speech sequences in artificial language studies (Saffran, Newport, & Aslin, 1996a; Saffran, Aslin, & Newport, 1996b). By analogy, then, we propose that readers may use bigram and trigram probabilities to discover which letter sequences cohere as morphemic units in print.

Preliminary corpus analyses suggest that placing morpheme boundaries within low-frequency transitions can segment many, though not all, polymorphemic words (Rastle et al., 2004). Though words like 'helpful' would be correctly segmented (because the low-frequency bigram 'pf' straddles the morpheme boundary) other words like 'hopeful' might not be (because the bigram 'ef' is of higher frequency and occurs within monomorphemic words). Existing data, however, suggest that morphemic effects arise for both kinds of stimuli (Rapp, 1992), suggesting that n-grams may not provide a sufficient account of *online* segmentation in skilled readers.

However, such data need not contradict the suggestion that n-gram information could be used to *acquire* orthographic representations. The speech segmentation literature proposes a similar distinction between acquisition and online use: phonotactic probabilities are a valuable prelexical cue for *acquisition*, but are overruled by higher-level lexical information during *online* processing in adults (Mattys, White, & Melhorn, 2005). On this account, then, one can view the acquisition of morphologically structured orthographic representations as a separate computational problem that can be solved prior to learning the form–meaning mapping for morphemically structured words.

One example of how this kind of account could be implemented in a computational system is directly inspired by the simple recurrent neural networks (SRNs) that have been used in simulations of infant speech segmentation (e.g., Cairns et al., 1997; Christiansen et al., 1998). Moscoso del Prado Martín, Schreuder, and Baayen (2004b) show that training an SRN on a letter prediction task generates internal representations that encode the sequential structure of English or Dutch orthography. Critically, when an 'accumulation of expertise' method is used to generate orthographic representations from these networks, the structure of these representations encodes the shared, morphemic units found in English words that end in –ity or –ness (Moscoso del Prado Martín et al., 2004b). Further simulations show that this method can provide appropriate orthographic representations for a large-scale model of Dutch past-tense formation (Moscoso del Prado Martín, Ernestus, & Baayen, 2004a). Thus the statistical structure of letter sequences provides sufficient cues to morphological segmentation to assist in the construction of connectionist models of language. Further application of

these methods to the form-meaning mapping for English derivational morphology would be of interest.

Note, however, that although this account specifies processing mechanisms that function during initial acquisition, we might still expect to see downstream consequences of bigram- and trigram-based segmentation in adult processing. Evidence in support of this account of orthographic segmentation could therefore be obtained if the degree of morpho-orthographic segmentation (e.g., as reflected by priming data) were predicted by the distribution of bigram or trigram profiles for words containing a particular affix over the lexicon as a whole. On this account we would predict that those affixes that consistently surface in words with reliable low-level segmentation cues (e.g., a robust bigram trough separating the affix from its stem) would be more readily segmented by readers and hence produce more reliable masked priming effects than would those affixes that do not surface in the context of such segmentation cues.

GROUPING HIGH-PROBABILITY SEQUENCES INTO SINGLE UNITS

Rather than dividing words on the basis of the low-frequency sequences that they may contain, a second approach to the acquisition of morphemically structured orthographic representations involves grouping high-frequency letter sequences into single units. This approach is at the heart of an account of speech segmentation based on the detection of sequential regularities in phoneme sequences that are assumed to be single lexical units (Brent & Cartwright, 1996; Wolff, 1977). Such accounts of segmentation have already been proposed for the acquisition of morphemic units (Brent, 1993) and provide a ready explanation for differences in the segmentation of opaque ('corner') and non-morphological ('brothel') items: the increased frequency of the letter sequence -er compared to -el leads the former but not the latter to be learned as an orthographic affix. However, differences in letter frequencies may not be sufficient as the sole explanation for why certain letter sequences function as orthographic affixes. For instance, the affix -able occurs in 484 lemmas in the CELEX database. This type frequency is not markedly different from that of the non-affix ending -el (242 items) which experimental evidence suggests does not support segmentation.

Perhaps a more critical difference between these endings is that affixes occur in combination with other linguistic units (e.g., stems). This characteristic provides for highly efficient chunking and therefore segmentation (Brent & Cartwright, 1996; Brent, 1997; see also Davis, 1999, for an analogous process in the SOLAR model of visual word recognition). Such a strategy would therefore favour detection of orthographic units (like -able)

for which the majority of occurrences are in combination with other units (rather than in simple items like 'stable'). By this account, acquisition of affix units might also be assisted by the existence of pseudo-affixed forms (such as 'tenable'), which despite being semantically opaque, would nonetheless support orthographic segmentation since they consist of a stem ('ten') plus an existing affix (-able).

A number of models in the speech segmentation literature have suggested computational mechanisms that group together frequently occurring sequences into single units or chunks. For instance, the PARSER account of word segmentation can develop a lexicon from exposure to continuous sequences of spoken syllables that are composed of trisyllabic 'words' (Perruchet & Vintner, 1998). A similar statistical approach has been proposed by Brent and colleagues in word and morpheme discovery (Brent, 1993; Brent & Cartwright, 1996), though these implementations require a perhaps implausibly large memory for unanalysed sequences. The discovery of orthographic chunks also forms an important part of a recent account of visual word recognition and word learning (the SOLAR model; Davis, 1999). In this model, new lexical nodes are assigned to frequently occurring letter sequences in a self-organising fashion inspired by the SONNET model of sequence learning (Nigrin, 1990). However, since large-scale simulations of morpheme learning in SOLAR have not been presented, it is difficult to know whether this model can account for existing evidence on morpho-orthographic segmentation. For instance, would morpheme recognition in SOLAR be disrupted by orthographic changes in {stem} + {suffix} combinations like 'metallic', 'writer', or 'adorable' even though these do not appear to disrupt human participants (cf. McCormick et al., 2008)?

USING FORM-MEANING REGULARITIES TO DRIVE ORTHOGRAPHIC LEARNING

Our final account of the acquisition of morphemically structured orthographic representations suggests that higher-level regularities learned across the form-meaning mapping drive lower-level orthographic learning. Though this style of account has been less favoured in the literature on speech segmentation (understandably given the sparseness of conceptual representations in pre-linguistic infants), neural network simulations have nonetheless shown that if conceptual representations can be assumed *a priori*, then consistencies in the form-meaning mapping do provide for the acquisition of form-based lexical segmentation (Davis, Gaskell, & Marslen-Wilson, 1997). Experimental investigations have similarly shown that form-meaning consistencies can be exploited in word learning by adult listeners (Yu & Smith, 2007). Finally, recent computational simulations of

segmentation and word learning (Davis, 2003), along with empirical investigations in infants (Graf-Estes, Evans, Alibali, & Saffran, 2007) converge in showing that form-meaning correspondences are learnt most effectively in conjunction with form-based segmentation processes. This might suggest that higher-level learning mechanisms operate in conjunction with lower-level orthographic segmentation processes.

In applying this theory to the acquisition of morpho-orthographic segmentation we should point out that the beginning reader has a head-start in using form-meaning regularities to segment written words into morphemes. New readers already have a well formed spoken vocabulary, including lexical and semantic representations for many of the more common stems and affixes in their language. Form-meaning correspondences therefore have a much greater opportunity to inform morpho-orthographic segmentation than they do for speech segmentation. The learning process involves readers detecting that certain letter sequences are consistently associated with morphemic elements already learnt from spoken language. In this way, readers have higher-level interpretations available that they can use to detect consistencies in the spelling of multiple different words that share stems and inflectional or derivational affixes.

Though a number of computational models have been proposed that learn the form-meaning mapping for morphologically complex words (Davis et al., 2003; Plaut & Gonnerman, 2000; Rueckl & Raveh, 1999), these models have so far all assumed that morphemically structured representations are provided as the input during training. This pre-segmented input is what allows these models to recognise orthographic similarity across sets of semantically related words (e.g., distrust, trust, untrustworthy). Thus, some mechanism is required that explains how it might be that form-meaning correspondences drive morphemic segmentation at the orthographic level. One potential mechanism can be derived from a 'NetTalk' inspired (Sejnowski & Rosenberg, 1987) model of reading aloud developed by Bullinaria (1995, 1997). Successful generalisation in this model depends on orthographic input and phonological output representations being aligned so as to emphasise consistent orthography-to-phonology correspondences. Rather than specifying these correspondences manually (as in Sejnowski & Rosenberg, 1987), Bullinaria (1995, 1997) demonstrated that output error during training can be used to select the correct input-to-output alignments from an exhaustive set of possible representations. We propose that the same method might be used to discover appropriately segmented orthographic input representations for complex words. This process is illustrated for a simple slot-based coding scheme in Figure 3. From a large set of possible input representations for a complex form like 'untrustworthy', a measure of output error at the semantic level for the stem 'trust' would suggest a single preferred input representation. Over the course of training, the coding
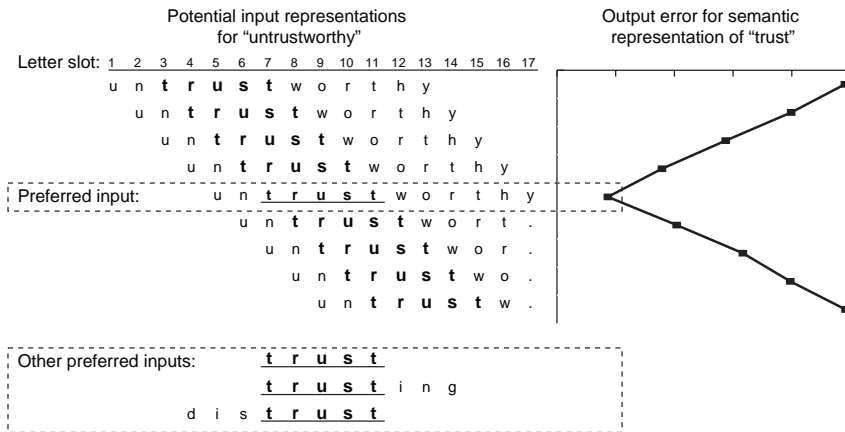
Figure 3. Illustration of how feedback from semantic representations could lead to successful morpho-orthographic segmentation using a method adapted from models of reading aloud (Bullinaria, 1995, 1997). An exhaustive set of orthographic input representations are generated and presented to a distributed connectionist network similar to the right panel of Figure 1. This exhaustive set of orthographic inputs is illustrated for the word 'untrustworthy' in a simple slot-based orthographic coding scheme. The input that produces the minimum output error for the stem 'trust' (as shown in the right hand graph) is marked as preferred and used in training the network. Across many iterations of testing different input representations and training the preferred input, the network will converge on aligned representations in which the same letter units code for the stem 'trust' in a variety of morphological contexts (e.g., trust, trusting, distrust, as shown in the bottom panel). This method of using learning and generalisation of meaning to select appropriately aligned input representations provides a mechanism by which feedback from semantics can assist in the discovery of morpho-orthographic segmentation.

scheme that consistently minimises output error at the semantic level should be the one in which the stem 'trust' is represented over the same set of input units in related forms like 'trust', 'trusting', and 'distrust'. By employing the same process for other morphemes (e.g., 'un-', and '-worthy', in 'untrustworthy'), the network will discover consistent morpho-orthographic units in a manner that is informed by feedback from semantics.

In proposing that form-meaning regularities contribute to the acquisition of morpho-orthographic segmentation we must make clear that (as for the bigram trough account) we can distinguish between mechanisms that support the *acquisition* of morpho-orthographic segmentation and those involved *online* in orthographic segmentation. Though on a form-meaning account, the acquisition of orthographic representations would be informed by shared affixes in semantically transparent complex words (e.g., 'darker', 'taller', 'smarter', etc.), the resulting orthographic representations could also be used to segment complex words with opaque meanings such as 'corner'. Such a situation might be expected if (as is the case for the affix -er),

non-compositional, semantically opaque items are relatively rare exceptions to a family of largely-consistent affix interpretations. The majority of semantically transparent forms drive learning, but the orthographic representation that is generated on the basis of these consistencies applies to multiple items, irrespective of semantic transparency (for similar arguments see Plaut & Gonnerman, 2000). Nonetheless, by this account we would expect to see an influence of higher-level factors such as the proportion of semantically transparent forms, affix consistency, and productivity on the effectiveness of morpho-orthographic segmentation for specific affixes (see also Chateau, Knudsen, & Jared, 2002).

Overall, then, experimental evidence of differences between various stems and affixes on the degree of morpho-orthographic decomposition that they support would be informative in evaluating all three of these accounts. By relating these empirical data to orthographic, morphological, and semantic properties of the family of lexical items that use each morpheme we could obtain evidence from adults to support one or more of these theories of acquisition. However, as is often the case, it is likely that the three sets of predictions will be hard to distinguish using the limited set of naturally occurring stems and affixes. Thus, it may be that investigations of artificial languages and laboratory analogues of morphemic acquisition will prove as valuable in investigations of morpho-orthographic segmentation as they have in studies of speech segmentation (Dahan & Brent, 1999; Saffran et al., 1996a,b). Such studies could also provide evidence concerning the relative effectiveness of each of these multiple cues either singly or in combination.

THE NEURAL BASES OF MORPHO-ORTHOGRAPHIC DECOMPOSITION

We now turn to a review of what is known of the neural instantiation of morpho-orthographic decomposition. This is of particular interest since (as described by acquisition theories) it provides an example of abstract, language-specific knowledge that is employed early on in the reading process.

Given the recent historical advent of reading in general, and mass-literacy in particular, it is implausible that the neural organisation of visual recognition of written words reflects anything other than a learnt specialisation, based on pre-existing cortical circuitry for the identification of visual objects and the translation of object representations into spoken language. Thus, we should be unsurprised to learn that the initial stages of visual word recognition build on the hierarchical cortical anatomy for visual feature identification (Hubel & Wiesel, 2005) and for the recognition of complex objects as established from cell recordings in non-human primates (Riesenhuber & Poggio, 2002; Tanaka, 1993) and from functional imaging studies of

object perception in humans (Malach, Levy, & Hasson, 2002). One recent review (Dehaene, Cohen, Sigman, & Vinckier, 2005) presents a precisely characterised hierarchical account of early visual processing of written words. This account starts from the recognition of simple letter elements (lines and curves) in primary visual cortex and proceeds in ascending levels of visual complexity along the ventral visual pathway. In this account the identification of letters and letter sequences occurs at later stages on the undersurface of the occipital and temporal lobe including portions of the fusiform gyrus. This hierarchical account would predict that the earliest form of morphological knowledge that is activated during visual word recognition corresponds to letter combination detectors in regions of the fusiform gyrus that are sensitive to the orthographic form of commonly occurring morphemic units (at an approximate coordinate of $y = -60$ in the MNI standard brain, cf. Dehaene et al., 2005, Figure 1).

Our review includes functional imaging and electrophysiological data that pertain to this account, and focuses in particular on studies that have the potential to inform our understanding of morpho-orthographic decomposition. In trying to identify the neural correlates of this form of decomposition, we decided to include in our review only those studies that use repetition priming paradigms. This decision rules out most functional imaging studies which focus on explicit morphological processing operations such as generating the past tense of regular and irregular verbs from their stems (e.g., Jaeger et al., 1996), detecting morphological violations in tense/agreement marking (Penke, Weyerts, Gross, Zander, Munte, & Clahsen, 1997), or performing other forms of explicit judgement that might be specifically sensitive to morphological variables (e.g., phonological same/different judgements; Tyler, Stamatakis, Post, Randall, & Marslen-Wilson, 2005). This omission is not intended to suggest that these studies are without value, only that they index neural correlates of explicit morphological processes that are later and more dependent on task manipulations than the early, obligatory morpho-orthographic decomposition that is the focus of the present paper. One other method that is frequently employed in functional imaging studies is to compare responses to complex and simple words using simple word recognition tasks (such as lexical decision or semantic judgements; e.g., Davis, Meunier & Marslen-Wilson, 2004; Laine, Rinne, Krause, Teras, & Sipila, 1999; Zweig & Pylkkänen, in press). Should response differences be observed for well-matched complex and simple words, then these differences may provide information about the neural correlates of morphemic processing. However, because these studies reflect both late semantically constrained decomposition as well as early morpho-orthographic decomposition they are also excluded from this review.

In considering functional imaging evidence we will focus on two techniques that provide complementary information concerning the neural

processes underlying reading: functional Magnetic Resonance Imaging (fMRI) and Electro/Magnetoencephalography (E/MEG). fMRI provides a slow, haemodynamic measure (BOLD) that, although only indirectly associated with spiking activity (Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001), provides good spatial precision in localising neural activation. E/MEG offers greater temporal precision by directly measuring electrical activity in the brain (or magnetic fields that are concomitant with electrical activity); however, responses can only be approximately localised to underlying neural generators (see Johnsrude & Hauk, 2005, for a more detailed review). In the recent literature there are both fMRI and E/MEG studies that use variants of the priming methods used in traditional behavioural investigations of morphological processing (specifically, repetition priming studies exploring the effect of morphological and/or orthographic overlap on word recognition).

Assessing the behavioural impact of morpheme repetition has long provided critical data on which to construct psychological theories of morphological processing (e.g., Marslen-Wilson et al., 1994; Stanners, Neiser, Hernon, & Hall, 1979). However, unlike behavioural experiments that provide only a simple measure of total facilitation or inhibition, using this method in the context of neuroimaging allows us to observe multiple differential priming effects that are localised to specific brain regions. The clearest example of the value of measures of neural priming comes from MRI studies in which region-specific priming effects provide a means of establishing the nature of representations found in specific brain areas and for inferring how different stages of neural processing contribute to an overall priming effect (for discussion see Grill-Spector, Henson, & Martin, 2006; Nacacche & Dehaene, 2001). For instance, fMRI studies have used masked priming to characterise a sequence of visual areas in the fusiform gyrus that generate an abstract representation of printed words independent of the case and retinal position of the constituent letters (Dehaene et al., 2004). Thus, neural priming can reveal spatially separate, and functionally dissociable processing stages within the hierarchy of regions involved in visual word recognition.

One early and influential neuroimaging study that applied this repetition priming method to morphological processing was conducted by Devlin, Jamison, Matthews, and Gonnerman (2004). They assessed neural repetition priming during masked presentation of orthographically (corner-CORN) and semantically (imitate-COPY) related word pairs, as well as morphologically related pairs that had both orthographic and semantic overlap (hunter–HUNT). Both sets of orthographically related word pairs produced repetition-related reductions (i.e., neural priming) in the fusiform gyrus consistent with form based processes hypothesised by Dehaene et al. (2004, 2005). However, as is apparent from the example stimulus pairs listed above,

all primes in these conditions included affix endings, and thus would be expected to elicit behavioural (and hence perhaps also neural) priming. For this reason one possible interpretation of the results reported by Devlin et al. (though not the one that they favoured) is that the priming effects observed in the fusiform gyrus reflect morpho-orthographic decomposition (Davis, 2004).

However, a follow-up study conducted by Gold and Rastle (2007) confirmed that response reductions in regions of the fusiform and posterior middle-occipital gyri were also equivalent for non-morphological form pairs (e.g., brothel-BROTH). These findings suggest that the fusiform and middle occipital gyri make an equivalent functional contribution to encoding letter sequences in both morphological and non-morphological contexts. In contrast to the Devlin et al. (2004) study, however, Gold and Rastle (2007) observed an additional region of the anterior middle occipital gyrus that showed neural priming specific to those stimulus pairs in which form overlap occurred in the context of a morphological affix (i.e., priming for corner-CORN, but not for brothel-BROTH). Gold and Rastle (2007) argued that the posterior-to-anterior orthographic-to-morphological gradient of neural priming effects observed in their study reflects the fact that the processing stream proceeds in the anterior direction as linguistic operations become more abstract. Because morphemes are letter clusters that play a functional role within words they can be regarded in a hierarchical model as having greater abstraction than letters themselves (Gold & Rastle, 2007).

One further recent study of morphological priming in fMRI that is of note was conducted by Bozic and colleagues (2007). In contrast to the masked fMRI priming studies, Bozic et al. employed a long-lag repetition priming paradigm in which neural correlates of morpheme repetition with multiple intervening items were assessed. Previous results obtained from this paradigm (Drews & Zwitserlood, 1995; Marslen-Wilson & Zhou, 1999; Rueckl et al., 2008 this issue) have yielded greater behavioural priming for semantically transparent pairs than for semantically opaque pairs. However, this study reported the intriguing finding of equivalent behavioural and neural priming (in this case in left inferior frontal regions) for transparent and opaque pairs (i.e., both hunter-HUNT and corner-CORN showed priming). Such results suggest that regions of prefrontal cortex that are distant from visual analysis of written words may contribute to morphemic analysis under conditions in which complex words are fully visible. One speculative interpretation of this finding is that these prefrontal regions provide top-down support for morpho-orthographic analysis.

Two recent studies have employed masked priming and a similar morphological versus non-morphological repetition design with time-locked EEG measures of neural activity (Lavric, Clapp, & Rastle, 2007; Morris, Frank, Grainger, & Holcomb, 2007). Though electrophysiological measures

are difficult to localise to critical sources of neural activity, their exquisite temporal resolution offers the potential to dissociate different time points during the processing of transparent, opaque, and simple words. However, despite using very similar methods and priming conditions (semantically transparent, semantically opaque, and non-morphological orthographic), there are some salient differences in the results obtained from these two studies. Both studies observed significant priming of electrophysiological responses approximately 400 ms after target onset in the transparent condition. This priming effect on the N400 component mirrors that observed in previous masked priming studies (Brown & Hagoort, 1993; Holcomb, Reder, Misra, & Grainger, 2005; Kiefer, 2002). Interestingly, however, the two studies differ in whether this N400 component is described as common to transparent and opaque items and significantly diminished for orthographic pairs (Lavric et al., 2007) or as showing a graded effect with progressively reduced N400 responses for both opaque and orthographic pairs (Morris et al., 2007).

Further disagreements between the two studies arise in considering earlier response components that also reflect masked repetition priming (approximately 200 ms after target onset). Both studies observed a significant ERP effect on the transparent items (Morris et al. labels this an N250 effect, while Lavric et al. analysed a longer time range between 140 and 260 ms after target onset). However, Morris et al. once more reported a graded effect with weaker neural priming for opaque items (primarily in posterior electrodes) and no effect for orthographic pairs, while Lavric et al. presented a more complex picture with priming effects for all three conditions, a reliable difference in topography between transparent and orthographic pairs, and an intermediate (or perhaps combined) topography in the opaque condition. Further differences are also observed in the behavioural measures of priming, with Lavric et al. reporting equivalent priming for transparent and opaque pairs and Morris et al. reporting a graded pattern (with intermediate and non-significant priming for the opaque condition).

One potential explanation for the differing outcomes of these studies can be traced to their SOAs. While Lavric used an SOA of 42 ms, Morris et al. used an SOA of 70 ms comprising a 50 ms prime and a 20 ms backward mask. The backward mask was used to reduce prime visibility, although data from a prime visibility test was not reported. Previous masked priming studies that have directly contrasted 42 and 70 ms SOAs (Rastle et al., 2000) report a reduction in the magnitude of behavioural priming for opaque items at the longer SOA which might explain apparent differences in neural and behavioural priming between these two studies. Follow-up experiments that examine the neural consequences of changes to prime presentation duration will be required, however, if we are to assess the significance of this methodological change.

The results of these EEG studies combine with fMRI data in suggesting that neural priming (like behavioural priming) can contribute to accounts of the recognition of complex words. However, both of these neurophysiological measures provide an amalgam measure of the processing of a prime-target pair. Even for EEG measures with high temporal resolution, critical differences between conditions do not emerge until around 200 ms after the onset of the target item – a time point at which processing of the target would be well underway. It is therefore unclear whether neural priming methods can provide an unambiguous measure of the initial processing of affixed words. Such data might be obtained from studies in which early responses to single written words (rather than pairs of written words) are assessed. However, E/MEG studies have not so far distinguished between decomposition processes that result from processing of semantically transparent complex words like 'hunter', and early orthographic decomposition that is also observed for opaque words like 'corner' (see Zweig & Pylkkänen, in press). We hope that this review will galvanise researchers in the cognitive neurosciences to conduct further psycholinguistically informed investigations of the neural basis of the identification of morphologically complex and simple words.

RECONCILING EVIDENCE FOR SEMANTICALLY BASED DECOMPOSITION WITH MORPHO-ORTHOGRAPHIC DECOMPOSITION

This article has provided evidence for a form of morphological decomposition based on the analysis of orthography, and has considered hypotheses about how the representations underlying this form of decomposition may be acquired. However, at the outset of this article we cited several pieces of evidence that would seem to be inconsistent with a form of decomposition based purely on orthography, and would instead support a form of decomposition constrained by semantic knowledge (e.g., Gonnerman et al., 2007; Longtin et al., 2003; Marslen-Wilson et al., 1994; Meunier & Longtin, 2007; Rastle et al., 2000). The critical finding in this respect is that transparent stimuli like 'darkness' but not opaque stimuli like 'corner' prime their stems in paradigms in which primes are of sufficient duration that they can be perceived consciously. How might these findings be reconciled with the form of decomposition that we have described?

Before considering this issue, we need to evaluate just how compelling these data are. One problem with using data from long-SOA priming paradigms to argue for semantically constrained decomposition is that it remains possible that these effects arise not because of a morphological relationship between prime and target but because prime and target are

related on both semantic and form dimensions (e.g., 'darkness-dark' are related morphologically, semantically, and orthographically). In their study of long SOA visual priming, for example, Rastle et al. (2000) were unable to distinguish priming of 'darkness-dark' items either from pure semantic priming (e.g., violin-cello) or from priming between pairs that had a semantic and orthographic relationship (e.g., screech-scream; brunch-lunch). Indeed, we are not aware of any priming study using a long SOA that has been able to distinguish morphological priming from that yielded by these kinds of pairs. Similarly, though Rueckl et al. (2008 this issue) demonstrates that priming for 'darkness-dark' items survives multiple intervening items while pure semantic priming does not, their work does not exclude the possibility that the 'darkness-dark' priming is due to the semantic *and* orthographic relationship between these primes and targets. The way to demonstrate this would be to include items like 'brunch-lunch' in a long-lag study like the one that they reported, an experiment that (to our knowledge) has not been done.

However, there are some priming studies that offer evidence for semantically constrained decomposition that are more difficult to reduce to semantic and/or form overlap. Marslen-Wilson et al. (1994) argued that the cross-modal priming effects that they observed for 'darkness-dark' pairs could not have been due to a combination of semantic and phonological overlap because they observed *inhibition* between pairs of suffixed items (e.g., darkness-darkly). Though this inhibitory effect does not hold up in visual priming (Rastle et al., 2000), it does suggest that the darkness-dark priming effects that they observed were due (at least in part) to shared morphology. Similarly, the finding that syntactically legal derived pseudowords (e.g., rapidify) facilitate recognition of their stems in cross-modal priming but that syntactically illegal derived pseudowords (e.g., sportation) do not (Meunier & Longtin, 2007) seems hard to reduce to a semantic effect for the simple reason that derived pseudowords do not have pre-existing semantic representations. These data also implicate a form of decomposition that is semantically informed. It thus seems that our account of morphological processing does require some explanation for why decomposition that appears morpho-orthographic in nature gives way at later periods in the time course of recognition to a form of decomposition that appears to be semantic in nature.

One possibility is that the two forms of decomposition observed behaviourally (orthographically based and semantically based) reflect decomposed representations at two separate levels of processing in visual word recognition. Specifically, the recognition system may contain two hierarchically organised processing stages: (a) a level of morpho-ortho-graphic decomposition that characterises the earliest stages of visual word perception; and (b) a level of 'morpho-semantic' decomposition that characterises a later stage of processing. This possibility is exemplified by

the distributed-connectionist theory pictured in Figure 2. Distributed representations for 'darkness' and 'dark' and for 'corner' and 'corn' overlap at the orthographic level in this theory. However, in the hidden units mediating orthographic and semantic representations, only those distributed representations for 'darkness' and 'dark' are overlapping. This theory is consistent with the observation of functionally distinct forms of decomposition, as long as it can be assumed that masked priming effects reflect orthographic levels of processing while priming effects from long-SOA paradigms reflect higher levels of processing.

The two forms of decomposition observed behaviourally might also be consistent with decomposed representations at just a single level of processing in the recognition system. The idea is that a single processing stage would produce an initial morpho-orthographic segmentation of the input, with inappropriate decompositions (e.g., interpreting 'corner' as 'corn' + '-er') being ruled out at later periods in the time course of recognition through a process of semantic integration. Schreuder and Baayen (1995; see also Meunier & Longtin, 2007) proposed a 'licensing' procedure along these lines that assesses the appropriateness of morphemic combinations (i.e., whether morphemes can legally be combined). Only if this licensing process succeeds is the meaning of the stimulus computed from its morphemic constituents. One interesting problem with this theory concerns words like 'whisker'. Though the licensing process would fail for stimuli like 'corner' (because nouns cannot take the suffix –er), it would succeed for stimuli like 'whisker' (because verbs can take the suffix –er). The problem here is that the usual meaning for the word 'whisker' is not that derived from its morphemic elements (i.e., 'someone who whisks'). One would have to propose a parallel non-decompositional process in order to explain how the meaning of this word is accessed, and even if such a process were proposed, a cost would still be predicted in the recognition of such words. To our knowledge this prediction has not been investigated.

CONCLUSIONS

We have reviewed a range of behavioural and neural data consistent with the proposal that a form of morphological decomposition based purely on orthographic analysis arises in the early stages of visual word processing. Though the functional properties of this morpho-orthographic decomposition are beginning to be established, many questions of scientific and applied importance remain unanswered. For example, though we have outlined three theories concerning the acquisition of morpho-orthographic information, there are virtually no relevant empirical or computational data to adjudicate between these. However, a full understanding of how children develop these

segmentation processes will likely be of considerable importance in considering different methods of reading instruction. Similarly, it is likely that morpho-orthographic segmentation makes an important contribution to the efficiency of speeded reading, particularly for users of morphologically rich languages. Thus, teaching methods that enhance morpho-orthographic segmentation should be favoured in school classrooms. Finally, a further important role for morpho-orthographic segmentation is in the context of understanding the neural basis of visual word recognition. Though relatively detailed accounts of the early stages of visual analysis of written words are available, questions concerning the functional and neural bases of morphological analysis and semantic interpretation remain largely unanswered. The stage is set for the cognitive accounts generated by psycholinguistics to be mapped onto neural circuitry. We predict that long-standing theoretical questions concerning the functional organisation of morphological processing will be advanced by parallel investigations of mind and brain.

REFERENCES

Boudelaa, S., & Marslen-Wilson, W. D. (2001). Morphological units in the Arabic mental lexicon. *Cognition*, *81*, 65–92.

Boudelaa S., & Marslen-Wilson W. D. (2005). Discontinuous morphology in time: Incremental masked priming in Arabic. *Language and Cognitive Processes*, *20*, 207–260.

Bozic, M., Marslen-Wilson, W. D., Stamatakis, E. A., Davis, M. H., & Tyler, L. K. (2007). Differentiating morphology, form, and meaning: neural correlates of morphological complexity. *Journal of Cognitive Neuroscience*, *19*, 1464–1475.

Brent, M. R. (1993). *Minimal generative explanations: A middle ground between neurons and triggers.* Paper presented at the Fifteenth Annual Conference of the Cognitive Science Society.

Brent, M. R. (1997). Towards a unified model of lexical acquisition and lexical access. *Journal of Psycholinguistic Research*, *26*, 363–375.

Brent, M. R. (1999). Speech segmentation and word discovery: A computational perspective. *Trends in Cognitive Sciences*, *3*, 294–301.

Brent, M. R., & Cartwright, T. A. (1996). Distributional regularity and phonotactic constraints are useful for segmentation. *Cognition*, *61*, 93–125.

Brown, C., & Hagoort, P. (1993). The processing nature of the N400: Evidence from masked priming. *Journal of Cognitive Neuroscience*, *5*, 34–44.

Bullinaria, J. A. (1995). Neural network learning from ambiguous training data. *Connection Science*, *7*, 99–122.

Bullinaria, J. A. (1997). Modelling reading, spelling and past tense learning with artificial neural networks. *Brain and Language*, *59*, 236–266.

Cairns, P., Shillcock, R., Chater, N., & Levy, J. (1997). Bootstrapping word boundaries: a bottom-up corpus based approach to speech segmentation. *Cognitive Psychology*, *33*, 111–153.

Chateau, D., Knudsen, E. V., & Jared, D. (2002). Masked priming of prefixes and the influence of spelling-meaning consistency. *Brain and Language*, *81*, 587–600.

Christiansen, M. H., Allen, J., & Seidenberg, M. S. (1998). Learning to segment speech using multiple cues: a connectionist model. *Language and Cognitive Processes*, *13*, 221–268.

Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). The DRC model: A model of visual word recognition and reading aloud. *Psychological Review*, *108*, 204–256.

Dahan, D., & Brent, M. R. (1999). On the discovery of novel word-like units from utterances: An artificial-language study with implications for native-language acquisition. *Journal of Experimental Psychology: General*, *128*, 165–185.

Davis, C. J. (1999). *The Self-Organising Lexical Acquisition and Recognition (SOLAR) model of visual word recognition*. Unpublished doctoral dissertation.

Davis, M. H. (2003). Connectionist modelling of lexical segmentation and vocabulary acquisition. In P. Quinlan (Ed). *Connectionist models of development: Developmental processes in real and artificial neural networks*. Hove, UK: Psychology Press.

Davis, M. H. (2004). Units of representation in visual word recognition. *Proceedings of the National Academy of Sciences USA*, *101*, 14687–14688.

Davis, M. H., Gaskell, M. G., & Marslen-Wilson, W. D. (1997). Recognising embedded words in connected speech: Context and competition. In J. A. Bullinaria, D. W. Glasspool, & G. Houghton (Eds.), *Proceedings of the Fourth Neural Computation and Psychology Workshop: Connectionist Representations*. London: Springer-Verlag.

Davis, M. H., Marslen-Wilson, W. D., & Gaskell, M. G. (2002). Leading up the lexical garden-path: Segmentation and ambiguity in spoken word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 218–244.

Davis, M. H., Meunier, F., & Marslen-Wilson, W. D. (2004). Neural responses to morphological, syntactic, and semantic properties of single words: an fMRI study. *Brain and Language*, *89*, 439–449.

Davis, M. H., van Casteren, M., & Marslen-Wilson, W. D. (2003). Frequency effects in processing inflected Dutch nouns: A distributed connectionist account. In R. H. Baayen & R. Schreuder (Eds.), *Morphological structure in language processing* (pp. 427–462). Berlin: Mouton de Gruyter.

Dehaene, S., Cohen, L., Sigman, M., & Vinckier, F. (2005). The neural code for written words: a proposal. *Trends in Cognitive Sciences*, *9*, 335–341.

Dehaene, S., Jobert, A., Naccache, L., Ciuciu, P., Poline, J. B., Le Bihan, D. et al. (2004). Letter binding and invariant recognition of masked words: Behavioral and neuroimaging evidence. *Psychological Science*, *15*, 307–313.

Devlin, J. T., Jamison, H. L, Matthews, P. M., & Gonnerman, L. M. (2004). Morphology and the internal structure of words. *Proceedings of the National Academy of Sciences USA*, *101*, 14984–14988.

Diependaele, K., Sandra, D., & Grainger, J. (2005). Masked cross-modal morphological priming: Unravelling morpho-orthographic and morpho-semantic influences in early word recognition. *Language and Cognitive Processes*, *20*, 75–114.

Drews, E., & Zwitserlood, P. (1995). Morphological and orthographic similarity in visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 1098–1116.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*, 179–211.

Feldman, L. B., & Soltano, E. G. (1999). Morphological priming: The role of prime duration, semantic transparency, and affix position. *Brain and Language*, *68*, 33–39.

Feldman, L. B., Soltano, E. G., Pastizzo, M. J., & Francis, S. E. (2004). What do graded effects of semantic transparency reveal about morphological processing? *Brain and Language*, *90*, 17–30.

Ford, M., Marslen-Wilson, W. D., & Davis, M. H. (2003). Morphology and frequency: Contrasting methodologies. In R. H. Baayen & R. Schreuder (Eds.), *Morphological structure in language processing* (pp. 89–124). Berlin: Mouton de Gruyter.

Forster, K. I., & Azuma, T. (2000). Masked priming for prefixed words with bound stems: Does revive prime survive? *Language and Cognitive Processes*, *15*, 539–561.

Frisson, S., Niswander-Klement, E., & Pollatsek, A. (2008). The role of semantic transparency in the processing of English compound words. *British Journal of Psychology*, *99*, 87–107.

Frost, R., Forster, K. I., & Deutsch, A. (1997). What can we learn from the morphology of Hebrew? A masked-priming investigation of morphological representation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 829–856.

Frost, R., Deutsch, A., Gilboa, O., Tannenbaum, M., & Marslen-Wilson, W. (2000). Morphological priming: Dissociation of phonological, semantic and morphological facts. *Memory and Cognition*, *28*, 1277–1288.

Giraudo, H., & Grainger, J. (2000). Effects of prime word frequency and cumulative root frequency in masked morphological priming. *Language and Cognitive Processes*, *15*, 421–444.

Gold, B., & Rastle, K. (2007). Neural correlates of morphological decomposition during visual word recognition. *Journal of Cognitive Neuroscience*, *19*, 1983–1993.

Gonnerman, L. M., Seidenberg, M. S., & Andersen, E. S. (2007). Graded semantic and phonological similarity effects in priming: Evidence for a distributed connectionist approach to morphology. *Journal of Experimental Psychology: General*, *136*, 323–345.

Graf Estes, K., Evans, J. L., Alibali, M. W., & Saffran, J. R. (2007). Can infants map meaning to newly segmented words? Statistical segmentation and word learning. *Psychological Science*, *18*, 254–260.

Grainger, J., & Jacobs, A. (1996). Orthographic processing in visual word recognition: A multiple read-out model. *Psychological Review*, *103*, 518–565.

Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: neural models of stimulus-specific effects. *Trends in Cognitive Sciences*, *10*, 14–23.

Holcomb, P. J., Reder, L., Misra, M., & Grainger, J. (2005). The effects of prime visibility on ERP measures of masked priming. *Cognitive Brain Research*, *24*, 155–172.

Hubel, D. N., & Wiesel, T. N. (2005). *Brain and visual perception: The story of a 25-year collaboration.* Oxford: Oxford University Press.

Jaeger, J. J., Lockwood, A. H., Kemmerer, D. L., Van Valin, J., Robert, D., Murphy, B. W. et al. (1996). A PET study of English verb morphology. *Language*, *2*, 451–497.

Johnsrude, I., & Hauk, O. (2005). Neuroimaging. In N. Braisby (Ed.), *Cognitive psychology: A methods companion.* Oxford, UK: Oxford University Press.

Jusczyk, P. W. (1997). *The discovery of spoken language.* Cambridge, MA: MIT Press.

Kazanina, N., Dukova-Zheleva, G., Geber, D., Kharlamov, V., & Tonciulescu, K. (2008). Decomposition into multiple morphemes during lexical access: A masked priming study of Russian nouns. *Language and Cognitive Processes, 23*, 800–823.

Kiefer, M. (2002). The N400 is modulated by unconsciously perceived masked words: further evidence for an automatic spreading activation account of N400 priming effects. *Cognitive Brain Research*, *13*, 27–39.

Laine, M., Rinne, J. O., Krause, B. J., Teras, M., & Sipila, H. (1999). Left hemisphere activation during processing of morphologically complex word forms in adults. *Neuroscience Letters*, *271*, 85–88.

Lavric, A., Clapp, A., & Rastle, K. (2007). ERP evidence of morphological analysis from orthography: a masked priming study. *Journal of Cognitive Neuroscience*, *19*, 866–877.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, *412*, 150–157.

Longtin, C. M., & Meunier, F. (2005). Morphological decomposition in early visual word processing. *Journal of Memory and Language*, *53*, 26–41.

Longtin, C. M., Segui, J., & Hallé, P. A. (2003). Morphological priming without morphological relationship. *Language and Cognitive Processes*, *18*, 313–334.

Malach, R., Levy, I., & Hasson, U. (2002). The topography of high-order human object areas. *Trends in Cognitive Sciences*, *6*, 176–184.

Marslen-Wilson, W. D., Bozic, M., & Randall, B. (2008). Early decomposition in visual word recognition: Dissociating morphology, form, and meaning. *Language and Cognitive Processes, 23*, 394–421.

Marslen Wilson, W., Tyler, L. K., Waksler, R., & Older, L. (1994). Morphology and meaning in the English mental lexicon. *Psychological Review, 101*, 3–33.

Marslen-Wilson, W., & Zhou, X. (1999). Abstractness, allomorphy, and lexical architecture. *Language and Cognitive Processes, 14*, 321–352.

Mattys, S. L., White, L., & Melhorn, J. F. (2005). Integration of multiple speech segmentation cues: a hierarchical framework. *Journal of Experimental Psychology: General, 134*, 477–500.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review, 88*, 375–407.

McCormick, S. F., Rastle, K., & Davis, M. H. (2008). Is there a 'fete' in 'fetish'? Effects of orthographic opacity on morpho-orthographic segmentation. *Journal of Memory and Language, 58*, 307–326.

Meunier, F., & Longtin, C. M. (2007). Morphological decomposition and semantic integration in word processing. *Journal of Memory and Language, 56*, 457–471.

Morris, J., Frank, T., Grainger, J., & Holcomb, P. J. (2007). Semantic transparency and masked morphological priming: an ERP investigation. *Psychophysiology, 44*, 506–521.

Moscoso del Prado Martín, F., Ernestus, M., & Baayen, R. H. (2004a). Do type and token effects reflect different mechanisms? Connectionist modeling of Dutch past-tense formation and final devoicing. *Brain and Language, 90*, 287–298.

Moscoso del Prado Martín, F., Schreuder, R., & Baayen, R. H. (2004b). Using the structure found in time: Building real-scale orthographic and phonetic representations by accumulation of expectations. In H. Bowman & C. Labiouse (Eds.), *Connectionist models of cognition, perception and emotion: Proceedings of the Eighth Neural Computation and Psychology Workshop* (pp. 263–272). Singapore: World Scientific.

Naccache, L., & Dehaene, S. (2001). The priming method: imaging unconscious repetition priming reveals an abstract representation of number in the parietal lobe. *Cerebral Cortex, 11*, 966–974.

Nigrin, A. L. (1990). SONNET: A self-organizing neural network that classifies multiple patterns simultaneously. *Proceedings of the International Joint Conference on Neural Networks, 2*, 313–318.

Penke, M., Weyerts, H., Gross, M., Zander, E., Munte, T. F., & Clahsen, H. (1997). How the brain processes complex words: An event-related potential study of German verb inflections. *Cognitive Brain Research, 6*, 37–52.

Perruchet, P., & Vintner, A. (1998). PARSER: A Model for Word Segmentation. *Journal of Memory and Language, 39*, 246–263.

Plaut, D. C., & Gonnerman, L. M. (2000). Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? *Language and Cognitive Processes, 15*, 445–485.

Plaut, D. C., & McClelland, J. L. (2000). Stipulating versus discovering representations [Commentary on M. Page, Connectionist modeling in psychology: A localist manifesto]. *Behavioral and Brain Sciences, 23*, 489–491.

Plunkett, K., & Elman, J. L. (1997). *Exercises in Rethinking Innateness: A Handbook for Connectionist Simulations.* Cambridge, MA: MIT Press.

Pollatsek, A., & Hyönä, J. (2005). The role of semantic transparency in the processing of Finnish compound words. *Language and Cognitive Processes, 20*, 261–290.

Rapp, B. (1992). The nature of sublexical orthographic organization: The bigram trough hypothesis examined. *Journal of Memory and Language, 31,* 33–53.

Rastle, K., & Davis, M. H. (2003). Reading morphologically complex words: Some thoughts from masked priming. In S. Kinoshita & S. J. Lupker (Eds.), *Masked priming: State of the art* (pp. 279–305). New York: Psychology Press.

Rastle, K., Davis, M. H., Marslen-Wilson, W. D., & Tyler, L. K. (2000). Morphological and semantic effects in visual word recognition: A time-course study. *Language and Cognitive Processes*, *15*, 507–537.

Rastle, K., Davis, M. H., & New, B. (2004). The broth in my brother's brothel: Morpho-orthographic segmentation in visual word recognition. *Psychonomic Bulletin and Review, 11*, 1090–1098.

Riesenhuber, M., & Poggio, T. (2002). Neural mechanisms of object recognition. *Current Opinion in Neurobiology*, *12*, 162–168.

Rueckl, J., & Aicher, K. (2008). Are corner and brother morphologically complex? Not in the long term. *Language and Cognitive Processes, 23*, 972–1001.

Rueckl, J. G., & Raveh, M. (1999). The influence of morphological regularities on the dynamics of a connectionist network. *Brain and Language*, *68*, 110–117.

Rumelhart, D. E., & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception: Part 2. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, *89*, 60–94.

Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996a). Word segmentation: the role of distributional cues. *Journal of Memory and Language*, *35*, 606–621.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996b). Statistical learning by 8-month-old infants. *Science*, *274*, 1926–1928.

Schoonbaert, S., & Grainger, J. (2004). Letter position coding in printed word perception: Effects of repeated and transposed letters. *Language and Cognitive Processe*s, *19*, 333–367.

Schreuder, R., & Baayen, R. H. (1995). Modelling morphological processing. In L. B. Feldman (Ed.), *Morphological aspects of language processing* (pp. 131–154). Hillsdale, NJ: Erlbaum.

Schreuder, R. & Baayen, R.H. (1997). How complex simplex words can be. *Journal of Memory and Language*, *37*, 118–139.

Seidenberg, M. S. (1987). Sublexical structures in visual word recognition: Access units or orthographic redundancy. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 245–263). Hove, UK: Lawrence Erlbaum Associates.

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*, 523–568.

Sejnowski, T. J., & Rosenberg, C. R. (1987). Parallel networks that learn to pronounce English text. *Complex Systems*, *1*, 145–168.

Stanners, R. F., Neiser, J. J., Hernon, W. P., & Hall, R. (1979). Memory representation for morphologically related words. *Journal of Verbal Learning and Verbal Behavior*, *18*, 399–412.

Taft, M. (1981). Prefix stripping revisited. *Journal of Verbal Learning and Verbal Behavior*, *20*, 289–297.

Taft, M. (1994). Interactive-activation as a framework for understanding morphological processing. *Language and Cognitive Processes*, *9*, 271–294.

Taft, M., & Forster, K. I. (1975). Lexical storage and retrieval of prefixed words. *Journal of Verbal Learning and Verbal Behavior*, *14*, 638–647.

Tanaka, K. (1993). Neuronal mechanisms of object recognition. *Science*, *262*, 685–688.

Tyler, L. K., Stamatakis, E. A., Post, B., Randall, B., & Marslen-Wilson, W. (2005). Temporal and frontal systems in speech comprehension: an fMRI study of past tense processing. *Neuropsychologia*, *43*, 1963–1974.

Wolff, J. G. (1977). The discovery of segmentation in natural language. *British Journal of Psychology*, *68*, 97–106.

Yu, C., & Smith, L. B. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science*, *18*, 414–420.

Zweig, E., & Pylkkänen, L. (in press). A visual M170 effect of morphological complexity. *Language and Cognitive Processes*. E-pub ahead of print, DOI: 10.1080/01690960802180420