

Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing?

David C. Plaut

Departments of Psychology and Computer Science and the Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, USA

Laura M. Gonnerman

Department of Psychology and the Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, USA

On a distributed connectionist approach, morphology reflects a learned sensitivity to the systematic relationships among the surface forms of words and their meanings. Performance on lexical tasks should thus exhibit graded effects of both semantic and formal similarity. Although there is evidence for such effects, there are also demonstrations of morphological effects in the absence of semantic similarity (when formal similarity is controlled) in morphologically rich languages like Hebrew. Such findings are typically interpreted as being problematic for the connectionist account. To evaluate whether this interpretation is valid, we carried out simulations in which a set of morphologically related words varying in semantic transparency were embedded in either a morphologically rich or impoverished artificial language. We found that morphological priming increased with degree of semantic transparency in both languages. Critically, priming extended to semantically opaque items in the morphologically rich language (consistent with findings in Hebrew) but not in the impoverished language (consistent

Requests for reprints should be addressed either to David Plaut (plaut@cmu.edu) or to Laura Gonnerman (laurag@cncb.cmu.edu), Mellon Institute 115-CNBC, Carnegie Mellon University, 4400 Fifth Avenue, Pittsburgh PA 15213-2683.

The research was supported by an NIMH FIRST award (MH55628) to the first author, by an NIMH NRSA Grant (DC00374) to the second author, and by NIMH Program Project Grant MH47566 (J. McClelland, PD). The computational simulations were run using customised software written within the Xerion simulator (version 3.1) developed by Drew van Camp, Tony Plate, and Geoff Hinton at the University of Toronto. We thank Marlene Behrmann, Steve Joordens, Jay McClelland, Brian MacWhinney, Jay Rueckl, and the CMU PDP research group for helpful comments and discussion.

with findings in English). Such priming arises because the processing of all items, including opaque forms, is influenced by the degree of morphological organisation of the entire system. These findings suggest that, rather than being challenged by the occurrence of non-semantic morphological effects in morphologically rich languages, the connectionist approach may provide an explanation for the cross-linguistic differences in the occurrence of these effects.

INTRODUCTION

Language is among the most fundamental of human cognitive abilities, and one of its most remarkable properties is its productivity: a relatively small number of familiar elements can be recombined in highly flexible ways to express a virtually limitless range of ideas. This productivity stems both from *morphology* (i.e., how words are formed) and from *syntax* (i.e., how they are ordered into sentences). Languages differ widely in the extent to which they rely on morphology versus syntax for productivity, and how morphology is expressed in the written and spoken forms of words. Indeed, it is generally agreed that morphology is the area of language that exhibits the greatest degree of cross-linguistic variation and, consequently, that learning must play a particularly important role in its organisation.

In these respects, English and Hebrew constitutes a very informative contrast. English makes relatively limited use of morphology, and the process of word formation mostly involves simple concatenation (e.g., UNBREAKABLE = UN- + BREAK + -ABLE; see Marchand, 1969). In Hebrew, by contrast, almost every word is structurally complex and word formation involves both concatenative and non-concatenative processes, including interdigitating a largely vowel-based word pattern into a three-consonant root (e.g., ZMR [related to music] + _ I _ _ A = “zimra” [singing]; see Berman, 1978). Thus, it is a considerable challenge to theories of language to explain how a child can learn either English or Hebrew (or both) depending only on the language environment in which he or she grows up.

The standard view of morphology is that words are built out of discrete units called *morphemes* that contribute systematically to the meanings of words containing them (e.g., RE- typically means to do again, as in REWASH, RESEAL, etc.). Morphological systematicity is important in part because it underlies the formation of novel words. For example, given the relatively new English verb FAX, we can automatically say and understand words like REFAXED and UNFAXABLE. In fact, the same holds for novel verbs: regardless of what the verb WUG might mean, something is UNWUGGABLE if it cannot be WUGGED.

Critically, as in other linguistic domains, systematicity in morphology is only partial. For example, although most English verbs form the past tense by adding -ED, there are also about 150 *irregular* verbs that form clusters

which undergo similar changes (e.g., SING) “sang”, DRINK) “drank”) along with a few very high-frequency arbitrary forms (e.g., GO) “went”). A similar state of affairs holds for other aspects of morphology (see Marchand, 1969). For instance, given an English verb, the most productive way of deriving the word for one who does the action is to add -ER; thus, someone who RUNS is a RUNNER, someone who THINKS is a THINKER, and so on. Such words are *transparent* because their meaning can be derived directly from the meanings of their morphemes. However, there are also *opaque* items: someone who - IES a plane is not a - IER but a PILOT; a CORNER is not someone who CORNS.

Many theories of morphological processing (e.g., Burani & Caramazza, 1987; Marslen-Wilson, Tyler, Waksler, & Older, 1994; Pinker, 1991; Prasada & Pinker, 1993; Schreuder & Baayen, 1995) respond to the fact that morphology is only partially systematic by handling the systematic and idiosyncratic aspects separately. Idiosyncratic words are stored as unanalysed wholes whereas systematic words are represented and processed compositionally in terms of stems and affixes. Recognition of the latter thus involves decomposing a complex word into its identifiable stem and affix and then accessing the stored stem (Marslen-Wilson et al., 1994). Such *dual-mechanism* theories run into problems, however, because morphological structure is not just partial but graded—that is to say, there are many intermediate cases that are neither completely transparent (regular) nor completely opaque (irregular). Rather, there are sub-regularities among the irregulars and between them and regulars. For example, the spoken forms of irregular verbs that do not change between present and past tense (e.g., HIT, SPREAD) all end in either /t/ or /d/—the same phonemes that occur at the end of regular past-tense forms. Thus, these items are not arbitrary exceptions; the fact that they already share the phonological structure of a past-tense form contributes to their patterning. Similarly, a word like DRESSER is neither fully semantically transparent nor fully opaque—although a DRESSER is not someone who dresses but a piece of furniture that holds clothes, it is clearly related to the activity of dressing. More generally, the isolation of idiosyncratic from systematic knowledge, which is at the core of dual-mechanism theories and which seems necessary to prevent the exceptions from compromising generalisation to novel forms, ends up being a serious handicap in the face of the rich, graded structure of language.

These difficulties have prompted many researchers to explore an alternative way of expressing language knowledge, based on connectionist or neural network modelling, that can handle graded degrees of systematicity more naturally. Connectionist models attempt to capture the essential properties of the neural mechanisms that give rise to behaviour by implementing cognitive processes in terms of cooperative

and competitive interactions among large groups of simple, neuron-like units (see McClelland, Rumelhart, & PDP Research Group, 1986; McLeod, Plunkett, & Rolls, 1998; Quinlan, 1991; Rumelhart, McClelland, & PDP Research Group, 1986b).

From a connectionist perspective, morphology is a characterisation of the learned mapping between the surface forms of words (orthography, phonology) and their meanings (semantics) (see also Cottrell & Plunkett, 1995; Gonnerman, Devlin, Andersen, & Seidenberg, 2000b; Hoeffner & McClelland, 1993; Joanisse & Seidenberg, 1999; Rueckl, Mikolinski, Raveh, Miner, & Mars, 1997; Rueckl & Raveh, 1999). To the extent that a particular surface pattern occurs in many words and maps consistently to certain aspects of meaning, the internal representations will come to reflect this structure and treat the pattern *componentially*—that is, represent and process it relatively independently of the other parts of the word (Plaut & McClelland, 1993; Plaut, McClelland, Seidenberg, & Patterson, 1996). As surface patterns behave less systematically, their internal representations will be correspondingly less componential. At the extreme end, if an item violates what is otherwise a systematic aspect of the domain (e.g., CORNER), its internal representation must be very non-componential in order to avoid being given a transparent interpretation (e.g., someone who CORNS).

To make the connectionist approach to morphology more concrete, consider the framework for lexical processing depicted in Figure 1 (see also Gonnerman et al., 2000b; Joanisse & Seidenberg, 1999). The orthography, phonology and semantics of words are represented as

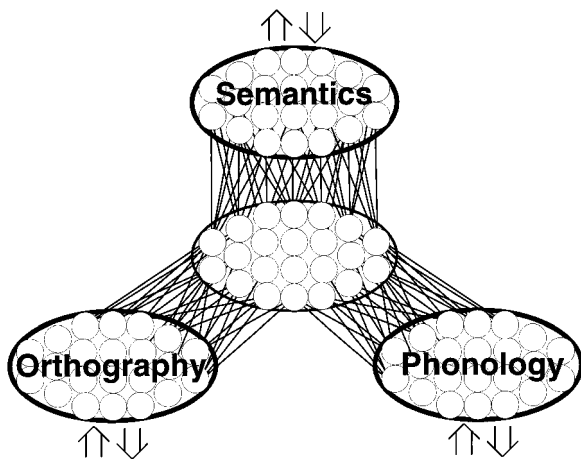


Figure 1. A connectionist framework for lexical processing. The large arrows depict inputs and outputs of the system.

distributed patterns of activity over separate groups of processing units. Within each of these domains, similar words are represented by similar (overlapping) patterns of activity. For example, CAT and DOG would be encoded by similar patterns in semantics but by dissimilar patterns in orthography and phonology, whereas CAT and CAP would be similar in orthography and phonology but dissimilar in semantics. Lexical tasks involve transformations between these patterns—for example, written word comprehension requires that the orthographic pattern for a word generate the appropriate semantic pattern. Such transformations are accomplished via interactions among the units, including additional intermediate or *hidden* units that mediate between the other groups of units.¹ Unit interactions are governed by weighted excitatory and inhibitory connections between them. When the network is presented with input (e.g., the orthography of CAT), units repeatedly update their activations as a function of the total input they receive via connections from other units, until the network as a whole settles into a stable pattern of activity constituting its response to the input (e.g., the meaning and pronunciation of CAT). The behaviour of the network in response to a given input thus depends on the specific values of the connection weights, which collectively encode the knowledge of the system. The network learns these values by gradually adjusting the weights to improve performance based on exposure to written words, spoken words, and their meanings. This learning has the effect of inducing new, internal representations over the intermediate units that are effective in accomplishing the relevant tasks.

At a general level, the reliance within the connectionist approach on *learning* internal representations is in sharp contrast to more traditional approaches to language which typically make very strong assumptions about the structure of internal representations and the processes that manipulate them. For example, it is often assumed that underlying linguistic knowledge takes the form of explicit rules which operate over discrete, symbolic representations (Chomsky, 1957; Chomsky & Halle, 1968; Fodor & Pylyshyn, 1988; Pinker, 1991; Prasada & Pinker, 1993) and, moreover, that this knowledge is, in large part, innately specified (Chomsky, 1965; Crain, 1991; Pinker, 1994). By contrast, rather than stipulating the specific form and content of the knowledge required for performance in a domain, the connectionist approach instead specifies the

¹ Readers may note that the current framework differs from the “triangle” model of word reading (e.g., Harm & Seidenberg, 1999; Plaut et al., 1996; Seidenberg & McClelland, 1989) in having one rather than three groups of hidden units mediating among orthography, phonology, and semantics. Although this distinction is a matter of ongoing research, Kello and Plaut (in press) report strategic effects in word reading that seem problematic for the triangle framework but can be simulated effectively by a model with a common group of hidden units (Kello & Plaut, 1998, in preparation).

tasks the system must perform but then leaves it up to the learning process to develop the necessary internal representations and processes (McClelland, St. John, & Taraban, 1989). The result is often that learning develops representations and processes which are radically different from those proposed by traditional theories, resulting in novel hypotheses and testable predictions concerning human cognitive behaviour.

In particular, the connectionist approach to morphology makes the strong prediction that the magnitudes of behavioural effects that reflect morphological processing should vary continuously as a function of the degree of semantic and phonological transparency of words. This prediction was recently confirmed in English by Gonnerman, Andersen, and Seidenberg (2000a), who found that prime-target pairs with intermediate degrees of semantic or phonological transparency produced intermediate degrees of priming.

Other empirical findings, however, would seem to challenge the connectionist claim that morphology is derived from learning among orthography, phonology, and semantics. In particular, it seems difficult to reconcile this perspective with demonstrations of morphological effects that are independent of both semantic and surface similarity. Many of these demonstrations are in languages, such as Hebrew, that have much richer morphological systems than English. For example, morphological priming in Hebrew can occur in the absence of semantic priming or can be insensitive to variations in semantic relatedness when surface similarity is controlled (Bentin & Feldman, 1990; Frost, Deutsch, & Forster, in press-a; Frost, Deutsch, Gilboa, Tannenbaum, & Marslen-Wilson, in press-b; Frost, Forster, & Deutsch, 1997). It thus appears that the existence of priming under these conditions is governed by morphological relatedness *per se*, and not by semantic or surface similarity. If, according to the connectionist approach, morphological structure ultimately derives from graded systematicity among surface forms of words and their meanings, there would seem to be no basis for morphological effects that are independent of semantic and surface similarity.

The current article reports on two simulations designed to evaluate whether the existence of morphological priming in the absence of semantic similarity in morphologically rich languages is inconsistent with a distributed connectionist approach to morphology. A common set of primes and targets, ranging in transparency, were embedded among a larger set of words that were either transparent or opaque, forming two artificial languages: a morphologically rich language analogous to Hebrew, and a morphologically impoverished language analogous to English. A network was then trained to map orthography to semantics for either language, and then tested for priming. Both languages yielded morphological priming for the more transparent targets, but priming extended to

semantically opaque targets only in the rich language. The results suggest that non-semantic morphological priming is, in fact, compatible with the connectionist approach but that its occurrence is predicted to depend on the degree of morphological richness of the language.

The next section reviews the relevant empirical findings that support or challenge a connectionist approach to morphology. We then discuss the computational principles that are essential for understanding both the general approach and the basis for the results of the simulations which follow. We conclude with a more general discussion of the strengths and limitations of the current work as well as remaining challenges to the connectionist approach.

OVERVIEW OF RELEVANT EMPIRICAL FINDINGS

One of the earliest studies on the role of morphology in lexical processing was conducted by Taft and Forster (1975). Using a lexical decision task, they found that subjects were able to reject nonwords with pseudo-stems (e.g., *DEPERTOIRE*) faster than nonwords with real stems (e.g., *DEJUVENATE*). Based on these findings, Taft and Forster argued for a processing strategy that first strips affixes from complex words and then searches for a stored stem representation. Many subsequent studies, using a variety of experimental techniques, have examined the extent to which morphological information is used to decompose complex words in lexical processing (see Feldman, 1995). These studies have produced conflicting results, leading to disagreement about the precise nature and locus of morphological decomposition. For example, Andrews (1986) argued that compound words are decomposed in lexical access, suffixed words are optionally decomposed in access, and all types of complex words are represented as whole forms. Alternatively, Stanners, Neiser, Hernon, and Hall (1979) claimed that derived forms are processed as wholes, but inflected forms are decomposed. Another account proposes that prefixed words are processed as whole forms but that suffixed words are decomposed (Colé, Beauvillain, & Segui, 1989). Yet another approach advocates two types of lexical access procedures—whole word and decompositional—which are activated in parallel (e.g., Burani, Salmaso, & Caramazza, 1984; Laudanna, Badecker, & Caramazza, 1989, 1992; Laudanna, Cermele, & Caramazza, 1997). While there has been no consensus as to exactly when and where morphological processing occurs, the majority of researchers agree that some form of discrete, all-or-none decomposition based on morphological structure operates somewhere within the lexical system.

The connectionist approach, by contrast, holds that decomposition is not an all-or-none phenomenon and that behavioural effects should be graded, reflecting the degree of convergence among semantic, phonological, and orthographic codes. One source of evidence for continuous, graded similarity structure in the semantic relatedness of morphologically complex words and their stems comes from subjects' similarity judgements (Gonnermann, 1999; Gonnerman, et al., 2000a). Gonnerman and colleagues found that subjects produced highly consistent ratings that fell on a continuum, judging some word pairs to be unrelated (e.g., CORNER-CORN), some moderately related (e.g., DRESSER-DRESS) and some highly related (e.g., TEACHER-TEACH). With regard to formal similarity, Rueckl et al. (1997) present findings from both masked and standard fragment completion tasks in which priming effects for morphologically related word pairs varied in magnitude as a function of orthographic similarity.

Further empirical support for the connectionist perspective comes from a series of cross-modal priming studies in which subjects made lexical decisions to written targets preceded by auditorily presented primes (Gonnerman, 1999; Gonnerman et al., 2000a; see also Marslen-Wilson et al., 1994). Gonnerman and colleagues compared the magnitude of priming as a function of the degree of semantic transparency between prime and target (e.g., high: BREAKAGE-BREAK; intermediate: SHORTAGE-SHORT; low: MESSAGE-MESS). Crucially, all of the items were phonologically transparent, such that the stems were contained unchanged in the complex words. Gonnerman and colleagues found that prime-target pairs with intermediate degrees of transparency also produced intermediate degrees of priming. Thus, highly semantically-related word pairs (e.g., BAKER-BAKE) primed twice as much (40 vs. 19 ms) as moderately related pairs (e.g., DRESSER-DRESS), and unrelated pairs (e.g., CORNER-CORN) showed no priming effects. Strikingly similar results were also obtained for pairs of prefixed words and their stems; highly related pairs (e.g., PREHEAT-HEAT) primed twice as much (42 vs. 20 ms) as moderately related pairs (e.g., MIDSTREAM-STREAM), whereas unrelated pairs (e.g., REHEARSE-HEARSE) showed no priming. Again, as with the suffixed-stem pairs, phonological relatedness was held constant. A third experiment showed that these graded effects hold for derived-derived word pairs. Highly related pairs (e.g., SAINTLY-SAINTHOOD) produced significant facilitation (34 ms), whereas less related pairs (e.g., OBSERVATION-OBSERVANT) produced only moderate facilitatory effects (14 ms) that failed to reach significance.

Gonnerman et al. (2000a) also found evidence for graded effects of phonological relatedness among word pairs that were equated on measures of semantic similarity. Thus, priming was greater among phonologically related pairs involving a consonant change (e.g., DELETION-DELETE; 65 ms)

than among pairs involving a vowel change (e.g., VANITY–VAIN; 48 ms), which, in turn, was greater than priming among pairs involving both vowel and consonant changes (e.g., INTRODUCTION–INTRODUCE; 35 ms).

In two final experiments, Gonnerman et al. (2000a) found similar priming effects among word pairs that are both semantically and phonologically similar but morphologically unrelated. In one experiment, significant priming (24 ms) was demonstrated for ‘pseudo-suffixed’ words that do not share morphological stems (e.g., TRIVIAL–TRI–E), but crucially, only if the words overlapped sufficiently in both meaning and sound. Data from a second experiment indicated priming (35 ms) for words taken from semantic and phonological clusters where there are no morphological relationships (e.g., SNARL–SNEER). Again, priming occurred only if there was sufficient overlap in both meaning and sound; SWITCH–SWIM pairs did not prime. These findings confirm that priming effects similar to those found for morphologically related items hold for word pairs that are correlated in meaning and sound, but have no morphological relationships (see also Rueckl & Dror, 1994, for a related demonstration that subjects learn an artificial vocabulary better if formal and semantic similarity are correlated).

Thus, Gonnerman et al. (2000a) provided clear evidence for graded effects of similarity in semantics, phonology, and orthography on cross-modal priming. Furthermore, the graded semantic effects were shown to generalise across morphological types, holding for morphologically related suffixed-stem, prefixed-stem, and suffixed-suffixed pairs. The effects of graded semantic and phonological similarity were demonstrated even for morphologically unrelated word pairs. Taken together, these results support a distributed connectionist approach where morphology characterises the regularities in the learned mapping between semantics and phonology.

There are, of course, findings that seem to challenge the notion that morphological effects can be reduced entirely to the joint influence of semantic and formal similarity. For example, with regard to formal similarity, Murrell and Morton (1974) found priming among formally related word pairs only if they were morphologically related (e.g., CARS–CAR but not CARD–CAR) and concluded that priming is determined by morphological structure rather than by formal relatedness per se. With regard to semantics, Kempley and Morton (1982) found that regularly inflected word pairs (e.g., RE–ECTED–RE–ECTING) produced significant facilitation, whereas irregularly inflected pairs (e.g., HELD–HOLDING) did not. They argued on the basis of these results that semantic properties of words do not influence processing, whereas morphology does. Many additional studies have come to a similar conclusion: that morphological structure, and not semantic or formal properties, underlies lexical

processing (e.g., Fowler, Napps, & Feldman, 1985; Grainger, Colé, & Segui, 1991; Marslen-Wilson et al., 1994; Napps, 1989; Napps & Fowler, 1987; Stolz & Besner, 1998).

It is important to note, however, that none of these studies examined or manipulated semantic and formal similarity *jointly*. For example, Napps and colleagues failed to find effects of semantic factors on morphological priming in one study (Fowler et al., 1985), and failed to find effects of phonological factors in another (Napps & Fowler, 1987), and concluded that “morphemic priming is not the result of the convergence of semantic, orthographic, and phonological relationships but rather that morphemic relationships are represented explicitly in the lexicon” (Napps, 1989, p. 729). Similarly, Marslen-Wilson et al. (1994) failed to find an effect of degree of phonological change between stem and derived forms, but they (unlike Gonnerman et al. 2000a) did not equate semantic similarity across conditions. In a subsequent experiment, they found evidence that priming was influenced by morphological type—suffixed-stem word pairs exhibited priming but suffixed-suffixed pairs did not—even though the conditions were matched on semantic similarity. Gonnerman et al. (2000b) showed, however, that the word pairs used by Marslen-Wilson et al. (1994) in the suffixed-suffixed condition were systematically less phonologically similar than the pairs in the suffixed-stem condition; in fact, suffixed-suffixed pairs do yield priming if they are sufficiently related (Gonnerman et al., 2000a). Gonnerman and colleagues also demonstrated that a distributed connectionist model trained on the experimental stimuli closely replicated the empirically observed pattern of priming.

More recently, Stolz and Besner (1998) explicitly question the ability of the connectionist approach to account for their morphological effects. They conducted a series of experiments where subjects made lexical decisions after performing a letter search task. Two experiments showed priming for morphologically related items (e.g., MARKED—MARK) compared to either an unrelated or an orthographic baseline (e.g., MARKET—MARK). In a third experiment, morphologically related items produced significant priming, whereas items that were only semantically related (e.g., GOLD—SILVER) did not. The authors also collected semantic relatedness ratings and demonstrated that these ratings could not account for the differences in priming between morphologically and semantically related items. Based on the combined results from their separate experiments, Stolz and Besner argue that morphological effects are independent of either orthographic or semantic similarity and thus challenge a connectionist approach. Recall, however, that form and meaning interact on our connectionist account, so we would in fact predict greater priming for items that are similar in *both* meaning and sound (e.g., morphologically related pairs) than for those related only on one of the two dimensions. Indeed, the Stolz and Besner

results, far from presenting a challenge, are entirely consistent with a distributed connectionist approach.

Overall, the evidence for morphological effects in English that are not reducible to semantic or formal similarity is, we believe, questionable at best. More convincing evidence of pure morphological effects comes from experiments in languages, such as Hebrew, with much richer morphological structure than English. As mentioned in the Introduction, virtually every word in Hebrew is morphologically complex, being formed by interdigitating a mostly vowel-based *word pattern* into a three-consonant *root* (see Bentin & Frost, 1995; Berman, 1978). Words derived from the same root exhibit a range of semantic relatedness. Bentin and Feldman (1990) took advantage of this fact to compare priming of a target word (e.g., MIGDAL [tower], root GDL) by morphologically related primes that either were semantically related to the target (e.g., GADOL [big]) or were not (e.g., GYDUL [tumor]), as well as by semantically related but morphologically unrelated primes (e.g., TZERIAH [turret]). For long-lag repetition priming, they found no reliable non-morphological semantic priming and equivalent levels of morphological priming regardless of whether the primes were semantically related or not (28 and 24 ms, respectively). Analogous results have also been obtained using brief masked priming (Frost et al., 1997). Note that non-semantic morphological priming in Hebrew is not restricted to within-modality (e.g., visual-visual) presentation as Frost et al. (in press-b) also found reliable priming for morphologically related but semantically opaque items under cross-modal presentation (26 ms). In this case the magnitude of priming did increase if the items were also semantically related (42 ms), although this is not surprising as the paradigm also yields semantic priming among morphologically unrelated items (Gonnerman, et al., 2000a; Marslen-Wilson et al., 1994).

There are analogous items in English that are morphologically related on a linguistic analysis but have little if any semantic similarity. For example, Aronoff (1976) has argued that -MIT operates like a morpheme in that it undergoes systematic morphological transformation across a broad class of words (e.g., TRANSMIT-TRANSMISSION, REMIT-REMISSION, COMMIT-COMMISSION) even though the semantic contribution of -MIT in these words is largely opaque. Note, however, that such items do not seem to produce reliable behavioural priming effects, under either cross-modal presentation (Marslen-Wilson et al., 1994) or visual-visual presentation with 250 ms prime duration (Feldman & Soltano, 1999). Thus, English appears to contrast with Hebrew in this respect.

An even more direct challenge to the connectionist approach comes from recent demonstrations that 'structural' manipulations of surface form that retain surface and semantic similarity can nonetheless have a profound impact on morphological priming (Frost et al., in press-a). Frost

and colleagues found that, among verbs, prime-target pairs in Hebrew with different roots but a common word pattern produce reliable priming if roots have three consonants, but that this priming is eliminated when the prime is a 'weak' form containing only two consonants. The word-pattern priming is reinstated, however, if the missing consonant is replaced with another random consonant to form a 'pseudo-root' (resulting in a nonword prime). The authors were careful to ensure that the manipulation of the consonantal roots did not alter the orthographic overlap or introduce semantic similarity among the resulting primes and targets. Thus, it appears as if the existence of priming in this context is governed, not by semantic or surface similarity, but by the presence of a structurally intact, morphologically defined three-consonant root.


Taken at face value, the evidence of pure morphological effects in Hebrew would seem to conflict with the connectionist proposal that morphology derives from learned relationships among the surface forms of words and their meanings. Indeed, Frost et al. (in press-a) claim that their results are incompatible with distributed connectionist systems, of which 'all that can be said is that a level of hidden units picks up the correlations between phonology and semantics, or orthography and semantics, and these underlie the morphological effects'. To more fully understand what the connectionist approach actually predicts in this context, however, we must consider in detail some of the computational principles that underlie the approach. It will turn out that the implications of these principles for morphological processing are far more subtle than is often assumed, and that the degree to which morphologically related items without semantic similarity should exhibit priming on a connectionist account depends on the degree of morphological structure of the entire language. These claims are supported by simulations described in this article.

PRINCIPLES UNDERLYING THE CONNECTIONIST APPROACH


The connectionist approach instantiates a number of computational principles that are relevant to morphological processing (see Figure 2). We discuss five central ones in some detail because they are important for understanding the conditions under which the approach predicts morphological effects in the absence of semantic and/or phonological similarity (for additional background on principles of connectionist modelling, see Chauvin & Rumelhart, 1995; Hertz, Krogh, & Palmer, 1991; McClelland et al., 1986; Rumelhart, Hinton, & Williams, 1986a; Smolensky, Mozer, & Rumelhart, 1996).

a) Distributed representations

Items are represented by patterns of activity such that similar items are assigned similar patterns

**b) Systematicity**


Similar inputs tend to produce similar outputs; mappings that maintain similarity structure are easier to learn

**c) Learned internal representations**

"Hidden" representations develop similarity structure midway between input structure and output structure

**d) Componentiality**

Parts of the output may depend only on parts of the input; supports combinatorial generalization

**e) One System**

All items are processed by the same set of weights; systematic and unsystematic aspects coexist but interact




Figure 2. Principles of connectionist modelling. In the diagrams, the circles represent units and the ovals represent groups of units. The rectangles represent similarity spaces in which the points correspond to different entities; the distance between points indicates the relative similarity (overlap) of the corresponding representations. Arrows indicate mappings between representations.

Principle 1: Distributed representations

Although many influential connectionist models employ localist representations, in which individual units correspond to familiar entities like words and concepts (see Page, in press, for discussion), a richer set of theoretical principles is available through the use of distributed representations (Hinton, McClelland, & Rumelhart, 1986). In a distributed representation, familiar entities are encoded by alternative patterns of activity over the same set of units, where each unit participates in representing many entities. For any pair of entities, their similarity in a given domain is reflected by the degree of overlap (i.e., number of shared active units) in their representations over the units for that domain. The full set of pairwise similarities among the representations for a given domain constitutes the *similarity structure* of that domain. Note that this similarity structure is intrinsic to a set of distributed representations, whereas, with localist representations, it must be imposed externally by appropriate associative weights between units.

Principle 2: Systematicity

Whereas similarity structure concerns the relationships among representations within a domain, *systematicity* concerns the relationship between the similarity structures of two domains. In order to perform a task, a network must map from the representations in one domain to the representations in another. For example, oral reading involves mapping the orthographic representations of words to their phonological representations, whereas spoken word comprehension involves mapping phonological representations to semantic representations. In general, the ease of performing a mapping will depend on its systematicity—that is, the extent to which patterns that are similar in one domain map to patterns that are similar in the other domain. Thus, in English, oral reading is highly systematic because words that are spelled similarly (e.g., CAT and CAP) tend to be pronounced similarly (see the left side of Figure 2b). By contrast, spoken comprehension of monomorphemic words is highly unsystematic because phonological similarity is unrelated to semantic similarity (see the right side of Figure 2b).

Systematic mappings are easier for connectionist networks to perform because networks have an inherent tendency to produce similar responses to similar inputs. This is because unit activations are determined by first computing a linear weighted sum of inputs from other units, and then applying a smooth, non-linear (typically sigmoidal) function to this summed input. A similar pattern of activity over the sending units will generally produce a similar summed input to the receiving unit and, hence, a similar activation value. The tendency to give a similar output to similar inputs is violated only when the differences between the sending patterns happen to be over units with very large connection weights to the receiving units. It is possible for learning to develop large weights when it is necessary to perform a task, but it takes a long time to do so (Hinton & Sejnowski, 1986).

Principle 3: Learned internal representations

Mappings that are highly systematic can often be carried out by a network with only direct connections from input units to output units (see, e.g., Zorzi, Houghton, & Butterworth, 1998, for such a model in the domain of oral reading). Such two-layer networks are, however, severely restricted in the mappings they can perform (Minsky & Papert, 1969). In order to carry out more complicated mappings, a network must have the ability to re-represent the input patterns such that they can be mapped effectively to the output patterns. More specifically, the network requires additional *hidden* units between the inputs and outputs, and it must develop representations over these units such that the similarity structure among

hidden representations is sufficiently similar to that of both the input and output representations. Each separate transformation (i.e., input-to-hidden and hidden-to-output) is then more systematic than the full input-to-output mapping.

In effect, one can think of the hidden representations as “splitting the difference” between the structure of the input and the structure of the output (this is depicted in Figure 2c by the fact that the arrows map in straight lines through the intermediate similarity space). For a relatively systematic mapping, the hidden representations will preserve much of the common similarity structure between the inputs and outputs, deviating from this only for the more idiosyncratic aspects of the mapping. By contrast, for an unsystematic mapping in which the similarity structures of the inputs and outputs are very different, the structure of the hidden representations may need to be rather different from each of them.

The tendency for hidden representations to adopt intermediate similarity structure was illustrated by Plaut (1991) for an attractor network that mapped from orthography to semantics via a single hidden layer. This task is a particularly useful one to investigate as orthographic similarity is largely unrelated to semantic similarity. Figure 3 shows, over the course of eight iterations of settling, the degree to which the activations over the hidden and semantic layers exhibit semantic structure versus orthographic structure. This was measured by computing the correlation between the pairwise similarities of representations at each layer with the pairwise similarities among the orthographic or semantic representations, respectively. The figure shows that the similarity structure of the network’s hidden representations (open symbols) falls between those of orthography and semantics (closed symbols). Specifically, over the last three iterations, the semantic activations are perfectly correlated with semantic similarity structure because the network is successful at generating the correct semantic patterns. The low degree of orthographic structure of these final semantic activations indicates that semantic structure is largely uncorrelated with orthographic structure. Most interesting, however, is that the final hidden activations are only moderately (and about equally) correlated with both orthographic structure and with semantic structure. Thus, in learning to map orthography to semantics, the network has developed hidden representations that compromise between orthographic and semantic similarity structure.

Principle 4: Componentiality

A form of systematicity that is particularly relevant to morphology is *componentiality*—the degree to which parts of the input can be mapped independently from the rest of the input. According to standard theories of

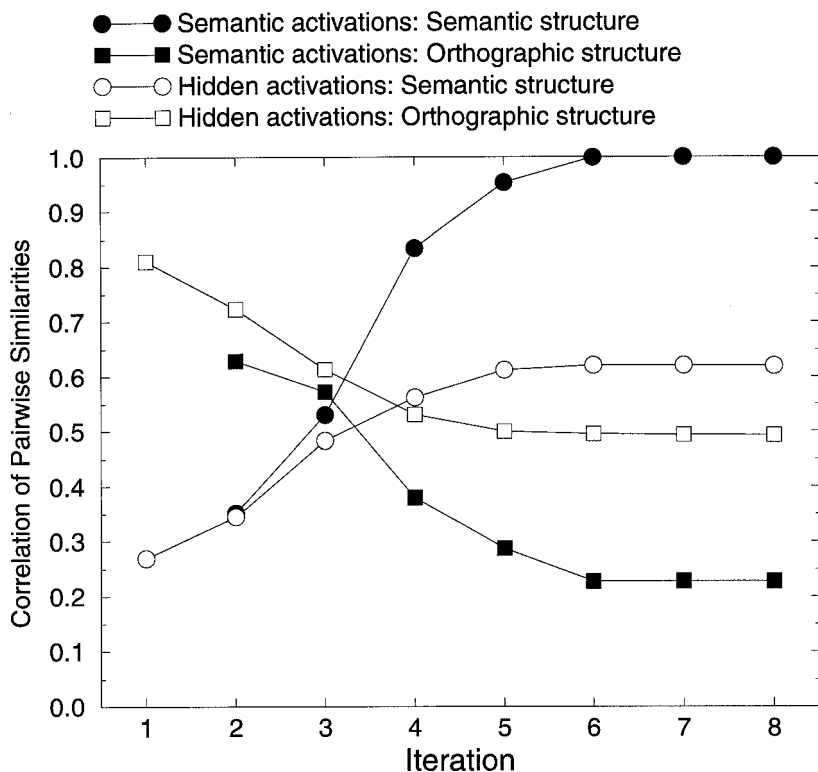


Figure 3. Changes in the similarity structure among hidden representations and among semantic representations during settling in a network that mapped orthographic representations to semantic representations (Plaut, 1991; Plaut & Shallice, 1993). 'Semantic structure' and 'orthographic structure' refer to the pairwise similarities among specified semantic or orthographic representations, respectively. (Adapted from Plaut, 1991.)

morphology, morphemes are discrete units that either are or are not contained in a given word. From a connectionist perspective, by contrast, the notion of 'morpheme' is an inherently graded concept because the extent to which a particular part of the orthographic or phonological input behaves independently of the rest of the input is always a matter of degree (Bybee, 1985). Also note that the relevant parts of the input need not be contiguous, as in prefixes and suffixes in concatenative systems like English. Even non-contiguous subsets of the input, such as roots and word patterns in Hebrew, can function morphologically if they behave systematically with respect to meaning.

A network comes to exhibit degrees of componentiality in its behaviour because, on the basis of exposure to examples of inputs and outputs from a

task, it must determine not only what aspects of each input are important for generating the correct output, but also what aspects are uninformative and should be ignored. This knowledge can then apply across large classes of items, only within small subclasses, or even be restricted to individual items. In this way, the network learns to map parts of the input to parts of the output in a way that is as independent as possible from how the remaining parts of the input are mapped. This provides a type of combinatorial generalisation by allowing novel recombinations of familiar parts to be processed effectively. In short, a network can develop componential representations that handle the systematic aspects of the task and that generalise to novel forms, while simultaneously developing less componential representations for handling the idiosyncratic aspects of the task.

The ability of connectionist networks to learn internal representations with varying degrees of componentiality has been illustrated by Plaut et al. (1996) in the domain of English word reading. After training an attractor network to map from orthography to phonology for about 3000 monosyllabic words, Plaut and colleagues determined the minimum level of activation over each orthographic input cluster (onset, vowel, coda) required to activate each phonological output cluster correctly. They contrasted *regular consistent* words that obey the standard spelling-sound correspondences of English (i.e., GAVE, MINT, analogous to morphologically transparent words) with *exception* words that violated these correspondences (i.e., HAVE, PINT, analogous to opaque words). They also considered the intermediate case of so-called *ambiguous* words which contain spelling patterns with multiple common pronunciations (e.g., -OWN in DOWN, TOWN, BROWN, etc. vs. KNOWN, SHOWN, GROWN, etc.).

The results are shown in Figure 4. The network developed highly componential representations for regular consistent words, in that each phonological cluster depended only on the corresponding orthographic cluster (i.e., onset) onset, vowel) vowel, coda) coda). This componentiality supports effective generalisation to pronounceable non-words (e.g., MAVÉ, RINT) because each cluster in the input is mapped more-or-less independently of the others. Onsets and codas in the other word classes are also treated componentially because they typically obey the standard spelling-sound correspondences. By contrast, the vowel cluster of exception words—which is the primary locus of exceptionality in English—is treated non-componentially. Correct activation of the vowel in the phonological output was sensitive to activation across the entire orthographic input. Moreover, the representations of the ambiguous words, with intermediate levels of spelling-sound consistency, showed intermediate degrees of componentiality. Thus, the network solved the task by learning internal representations with the appropriate degree of componentiality given the specific nature of the statistical structure of the task.




Figure 4. Componential representations in the word reading model of Plaut et al. (1996, Simulation 3), for Taraban and McClelland's (1987) regular consistent words (top) and exception words (bottom), and for orthographically matched 'ambiguous' words with an intermediate degree of spelling-sound consistency (middle). The height of each bar indicates the minimum activity level of the indicated orthographic cluster to correctly activate a particular phonological cluster. Words lacking either an onset or coda consonant cluster in orthography were excluded from the analysis. (Adapted from Plaut et al., 1996.)

By the same token, a network trained on inputs with graded morphological structure would be expected to develop representations with the corresponding degrees of componentiality (see, e.g., Rueckl & Raveh, 1999). That is, items with a relatively high degree of morphological transparency (e.g., *TEACHER*) will be represented and processed componentially (*TEACH* + *-ER*), whereas relatively opaque items (e.g., *CORNER*) will be treated more holistically. Critically, the connectionist approach offers a continuum of compositionality, such that items with an intermediate degree of transparency (e.g., *DRESSER*), analogous to the ambiguous words, will participate in the more general componential structure to some degree but will also deviate from this structure somewhat.

Principle 5: One system

The final principle is perhaps the one most directly relevant to the current simulation work. It concerns the simple fact that all items within a task are processed by the same system, so that the knowledge of how to handle all aspects of the task must be superimposed on the same set of weights. Thus, although a single network can handle both the systematic, regular, transparent aspects of the task and the idiosyncratic, irregular, opaque aspects (and everything in between), these aspects are not treated independently of one another. Rather, every item is processed in the context of knowledge about every other item, and the general organisation of the hidden representations is determined by the overall nature of the task. If a task contains a large proportion of opaque items, it will be difficult for a network to discover and take advantage of the systematicities underlying the few transparent items. On the other hand, if a task is largely systematic, the hidden representations will be dominated by this structure, making it more difficult to learn the opaque items; indeed, the opaque items will, to some extent, be carried along by this structure and their representations will have to compensate.

This property may have important implications for understanding the general pattern of empirical results reviewed in the previous section, suggesting that morphological effects that are independent of semantic and surface similarity are found only in languages with rich morphological structure (e.g., Hebrew) but not in languages with relatively impoverished morphology (e.g., English). On the connectionist account, a given set of opaque items will be represented and processed differently depending on the degree of systematicity of the rest of the task. In an unsystematic task, the opaque items will be free to behave idiosyncratically because the other items do not combine to exert a coherent influence. In an otherwise systematic task, the same opaque items will be subject to the strong influence of the overriding transparent structure of the task and, thus, will

necessarily be represented in a way that coexists effectively with this structure. This is depicted in Figure 2e by the fact that mapping for the lone opaque item is distorted to be more consistent with the mappings for transparent items than it would have been in isolation (straight line through open circle). The influence of overall task structure on the way in which opaque items are represented and processed within connectionist networks forms the basis of our account of empirical findings, within morphologically rich languages like Hebrew, of morphological effects that are not reducible to semantic and/or phonological similarity.

SIMULATION 1

The preceding discussion of the computational principles underlying a distributed connectionist approach to morphology suggests that, although morphology ultimately derives from learned systematic relationships between word forms and meanings, the representational structure that emerges from this learning may, at least under some circumstances, extend to the processing of morphologically simple, opaque forms. In this way, the approach may be consistent, not only with graded influences of semantic and formal similarity on morphological priming (e.g., Gonnerman et al., 2000a), but also with findings of effects of morphology that are independent of semantic and formal similarity (e.g., Bentin & Feldman, 1990; Frost et al., 1997, in press-a). Although, it may seem at first glance as if the approach is underconstrained if it can account for such inconsistent patterns of performance, it actually provides a possible explanation for the diversity of empirical observations. In particular, the connectionist approach may provide insight into why the specific pattern of morphological priming observed depends on the overall degree of morphological richness of the language.

To further substantiate this account, we carried out a simulation in which a network was trained on an abstract version of the task of comprehending morphologically complex written words in either of two languages.² One of the languages had a rich degree of morphological structure, analogous to Hebrew, whereas the other had far less morphological structure, analogous to English. For present purposes, the morphological 'richness' of a language was operationalised in terms of the frequency and extent to which the structure of words within the

² Although our general framework for lexical processing includes interactions of both orthography and phonology with semantics (see Figure 1), we included only orthography in the current simulation because the current work is directed at understanding findings from languages, such as English and Hebrew, in which orthographic and phonological structure are highly correlated. Thus, for present purposes, the inclusion of phonology would have been largely redundant with orthography.

language—both in terms of their surface forms and their meanings—could be systematically decomposed into components that participate in the construction of many other words. More specifically, each language employed a small number of ‘affixes’; languages differed both in the proportion of words containing an affix and in the degree to which the meanings of such words could be predicted from the ‘transparent’ meanings of their stems and affixes. The instantiation of morphological structure within the simulation was, however, very abstract and not intended to approximate the actual morphology of any real language in any detail. Rather, it was designed to capture the most general characteristics of morphological structure and to permit tightly controlled manipulations of this structure in a way analogous to variations of morphological richness across languages. The goal was to demonstrate that, on a connectionist account, the occurrence of non-semantic morphological effects depends on the degree of morphological structure of the entire language. We consider the implications of the limitations in our approximation of morphological structure in the General Discussion.

Method

Network architecture. The architecture of the network is shown in Figure 5. Thirty orthographic units were fully connected to 300 hidden units which, in turn, were fully connected to 50 semantic units. The hidden and semantic units also had trainable *biases* that are equivalent to connections from a unit whose activation is always equal to 1.0. Including the biases, the network had a total of 24350 connections.

The activation a_j of each unit j was computed using the standard sigmoid function $\sigma(\cdot)$,




Figure 5. The architecture of the network.

$$a_j = \sigma \left(\sum_i a_i w_{ij} \right) = \frac{1}{1 + \exp(-\sum_i a_i w_{ij})} \quad (1)$$

where w_{ij} are the weights on connections from sending units i .

Each weight was initialised to a random number drawn from a uniform distribution between ± 0.1 , except for the biases. These were initialised to $\sigma^{-1}(0.1) = -2.197$ for hidden units, and $\sigma^{-1}(0.2) = -1.368$ for semantic units. Using relatively large negative initial biases improves stability at the onset of learning.

Representations. The orthographic representations of all words contained two syllables chosen from 100 possible first syllables and 100 possible second syllables. Each syllable was represented in terms of a random binary pattern over 15 units with approximately one-third of the units active for any syllable. Thus, no attempt was made to introduce any particular similarity structure among the syllables.

The abstract morphology was modelled after a system with stems and affixes. We will refer to the first syllables as *stems* for ease of exposition, although they ranged in transparency from those behaving like fully independent morphemes to those with no independent semantic status whatsoever. Ten of the 100 possible second syllables were selected at random and designated as *affixes*; these syllables behaved in a systematic fashion with regard to their occurrence in the language and their implications for meaning. In particular, the transparency of a stem was reflected in the frequency with which it combined with the affixes (as opposed to other second syllables) and in the degree to which the meanings of these stem-affix combinations preserved the canonical meanings of the stem and affix.

The construction of semantic representations was designed to permit a graded degree of transparency between the orthographic forms of words and their meanings. This was done by first assigning each of the 200 possible syllables a canonical meaning over 50 semantic features, containing 10 randomly selected semantic features for stems, and 5 for affixes and other second syllables. For any pair of syllables, a transparent meaning was constructed by including any semantic feature that was in the meaning of either syllable (i.e., bitwise OR). Note that the resulting pattern has a high degree of overlap with each of the component patterns. This transparent meaning was then distorted to create an actual word meaning, by randomly turning off a specified proportion of its active features and randomly turning on the same proportion of features elsewhere.

Each of the 100 stems was associated with one of four levels of distortion, resulting in four word classes: (a) *transparent* words retained all

of the features of their transparent meaning; (b) *intermediate* words retained two-thirds of the features; (c) *distant* words retained one-third of the features; and (d) *opaque* words retained none of the features of their transparent meaning. Thus, words sharing a stem had varying degrees of semantic similarity as a function of the levels of distortion used in creating their meanings.

Using these procedures for generating representations, two languages were created containing 1200 words each. In the morphologically *rich* language, the first 60 stems, forming 720 words, were all transparent; in the *impoverished* language, they were all opaque. This contrast formed the main experimental manipulation of the simulation. Critically, the remaining 480 words were identical across the two languages, both orthographically and semantically. These were formed from 10 transparent stems, 10 intermediate stems, 10 distant stems, and 10 opaque stems. The simulation was designed to evaluate the degree of morphological priming among this shared set of words as a function of the nature of the remaining words in each of the two languages.

Opaque stems differed from the three classes of non-opaque stems (i.e., transparent, intermediate, distant) not only in their degree of transparency but also in their distributional properties with respect to the 10 designated affixes. In particular, each opaque stem was combined with 2 randomly chosen affixes and 10 randomly chosen non-affix second syllables to form 12 words, all of which were assigned opaque meanings. The combinations with affixes are analogous to pseudo-affixed words like *CORNER*. By contrast, each non-opaque stem was combined with all 10 affixes and with 2 randomly chosen non-affixes. The affixed items were assigned semantic representations with the specified degree of transparency; the non-affixed items were assigned opaque meanings and are analogous to words like *SPINACH* which are not morphologically complex but happen to start with syllables that can be stems in other contexts.

Training procedure. Starting with the identical set of initial random weights, the network was trained on either the morphologically rich or impoverished language. For each word presentation, the activations of orthographic units were set to the assigned orthographic representation of the word, and hidden and semantic activations were computed in turn. Performance error was computed in terms of the cross-entropy (Hinton, 1989) between the semantic activations generated by the network and the correct meaning of the word, using targets of 0.8 or greater for present features and 0.05 or less for absent features. The partial derivative of this error with respect to each weight and bias in the network was then computed using back-propagation (Rumelhart et al., 1986a; Werbos, 1974) and accumulated over word presentations. Weights were updated only

after each sweep through the full 1200 words, using a learning rate of 0.0005, momentum of 0.9 (applied only after the first 10 sweeps), and weight decay of 0.0001.

The two versions of the network were trained until they were fully accurate at activating only the correct semantic features for each word and had reached equivalent (and extremely low) levels of quantitative error. This required 300 training sweeps for the morphological rich language and 600 sweeps for the impoverished language, yielding means of cross-entropy per unit per pattern of 0.0034 and 0.0037, respectively.

Testing procedure. The network was evaluated for its ability to produce facilitation or inhibition in reaction time (RT) in a paired priming paradigm for each of the four levels of morphological transparency in each of the two languages. Given that the standard computations in a feedforward network do not produce a time-course of processing over the output units, we altered the computation in the network to *cascade* unit activations (McClelland, 1979) according to the equation

$$a_j^{[t+\tau]} = \tau \sigma \left(\sum_i a_i^{[t]} w_{ij} \right) + (1 - \tau) a_j^{[t]} \quad (2)$$

with a time constant $\tau = 0.01$. We then defined RT to be the amount of time required to reach a stability criterion such that the mean change in semantic activations fell below 0.01. Note that Equation 2 is equivalent to Equation 1 for $\tau = 1.0$ and reaches stability when $a_j^{[t+\tau]} = a_j^{[t]}$, which occurs when $a_i^{[t]} = \sigma \left(\sum_i a_i^{[t]} w_{ij} \right)$, thereby satisfying Equation 1. Thus, the cascaded network asymptotes to the same activations as the standard network (McClelland, 1979). The technique of cascading unit activations within a feedforward network has been applied effectively to modelling detailed RT data in other contexts (e.g., Cohen, Dunbar, & McClelland, 1990).

On a given trial, a prime word was presented to the network by setting orthographic activations to its orthographic representation and initialising hidden and semantic activations to 0.1. The network then processed the prime for 1.0 unit of time (100 unit updates according to Equation 2 with $\tau = 0.01$). The orthography of a target word then replaced that of the prime without altering any other activations in the network, and processing continued until the semantic activations satisfied the stability criterion (mean semantic activation change < 0.01) in settling to the meaning of the target word. The time from the onset of the target to when this criterion was reached was taken as the network's RT to the target following the prime.

As the primary goal of the simulation was to evaluate the influence of language structure on morphological processing, all primes and targets were drawn only from the 480 words that were common across the two languages. In this way, all potential confounds due to differences in surface and semantic structure across languages were eliminated. Targets consisted of each affixed word generated from the transparent, intermediate, distant, and opaque stems. There were 100 targets in each of the first three classes and 20 in the last class (as these are the pseudo-affixed items).

Many empirical studies employ primes that are individually matched to each target word. To derive a more reliable estimate of priming effects, however, we computed RTs for each target word preceded by every prime word having the appropriate relationship to the target word. *Identical* primes were, of course, the targets themselves. *Morphological* primes consisted of all non-identical affixed words with the same stem (and, hence, the same level of morphological transparency). There were nine morphological primes for each non-opaque target (900 total prime-target pairs) and one for each opaque target (20 prime-target pairs). *Control* primes consisted of non-affixed opaque words with no syllables in common with the target. There were 90 of these for each opaque target (9 non-identical stems \times 10 non-affixed second syllables; 1800 total prime-target pairs). As there were far more non-opaque targets, we randomly sampled among possible control primes for these targets to produce approximately 2000 prime-target pairs for each level of transparency. Table 1 gives examples from English that are analogous to each of the experimental conditions used in testing the network.

It should be pointed out that the identity condition for the network is not directly analogous to that for human subjects in that no change in the stimulus occurs between prime and target—it is exactly equivalent to measuring the network's RT to the target and then subtracting the prime duration of 1.0. By contrast, identity conditions in empirical studies employ some physical change in the stimulus, such as increasing its size or changing from lower to upper case. Insofar as these alterations slow RTs to the target, the network's RTs will underestimate those of human subjects.

TABLE 1
English examples of experimental conditions

<i>Transparency</i>	<i>Target</i>	<i>Priming condition</i>		
		<i>Identical</i>	<i>Morphological</i>	<i>Control</i>
Transparent	RUNNER	RUNNER	RUNNING	SPINACH
Intermediate	DRESSER	DRESSER	DRESSING	SPINACH
Distant	TENDER	TENDER	TENDING	SPINACH
Opaque	CORNER	CORNER	CORNISH	SPINACH

Results and discussion

Figure 6 shows the total error produced by the network over the course of training on either the morphologically rich or impoverished languages. With the same architecture and starting from the same initial random weights, the network clearly acquires the morphologically rich language more quickly than the impoverished one, requiring only half as many sweeps through the 1200-word corpus to achieve the same final level of performance. This difference arises directly from the different degrees of systematicity across the two languages, as discussed under Principle 2 earlier. It is also broadly consistent with cross-linguistic differences in rates of morphological acquisition (see, e.g., Andersen, 1992).

Each version of the network was then run on the identical set of prime-target pairs, as described in the Method section, and its RTs were measured. Table 2 presents the mean RTs of the network to target words of each level of systematicity as a function of the type of the preceding prime word and the language on which it was trained. The RTs were then entered into a $2 \times 3 \times 4$ ANOVA with items as the random variable, *language* (rich, impoverished) and *priming condition* (identical, morpho-

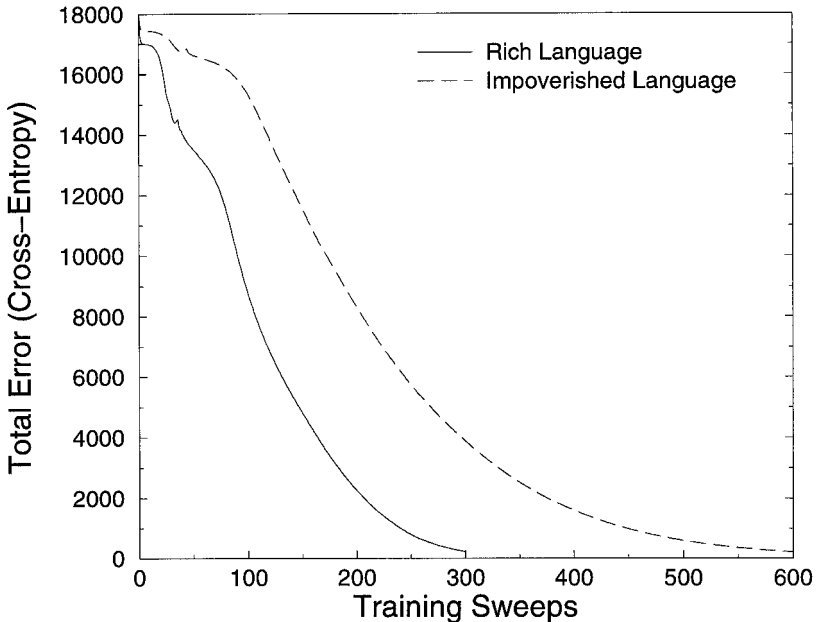


Figure 6. Total cross-entropy error for the entire training corpus produced by the network as a function of sweeps through the corpus for either the morphologically rich or impoverished language.

TABLE 2

Mean RTs (and standard errors) to target words as a function of language and priming condition, and magnitudes of priming effects

<i>Transparency</i>	<i>Priming condition</i>			<i>Priming</i>	
	<i>Identical</i>	<i>Morphological</i>	<i>Control</i>	<i>Identity</i>	<i>Morphological</i>
<i>Rich language</i>					
Transparent	3.28 (0.02)	3.95 (0.02)	4.41 (0.01)	1.11	0.45
Intermediate	3.58 (0.02)	4.47 (0.02)	4.70 (0.02)	1.13	0.26
Distant	3.69 (0.02)	4.67 (0.02)	4.82 (0.02)	1.13	0.14
Opaque	3.63 (0.05)	4.67 (0.05)	4.75 (0.04)	1.13	0.08
<i>Impoverished language</i>					
Transparent	3.28 (0.02)	4.07 (0.02)	4.47 (0.02)	1.17	0.37
Intermediate	3.60 (0.03)	4.62 (0.03)	4.77 (0.02)	1.18	0.17
Distant	3.72 (0.03)	4.84 (0.02)	4.89 (0.02)	1.17	0.05
Opaque	3.71 (0.06)	4.91 (0.05)	4.89 (0.05)	1.18	-0.02

logical, control) as within-item factors, and morphological *transparency* (transparent, intermediate, distant, opaque) as a between-item factor. There was no need to perform any data trimming as the network operated deterministically and made no errors.

The ANOVA revealed a main effect of language, with targets producing faster RTs in the morphologically rich language compared with the impoverished language (means of 4.19 vs. 4.27, respectively; $F_{1,316} = 34.42$, $p < .001$); a main effect of priming condition, with identity primes yielding the fastest RTs, followed by morphological primes, and then control primes (means of 3.53, 4.46, 4.69, respectively; $F_{2,316} = 34526$, $p < .001$); and a main effect of morphological transparency (means of 3.91, 4.29, 4.44, and 4.43, for transparent, intermediate, distant, and opaque targets, respectively; $F_{3,316} = 252.2$, $p < .001$). However, priming condition interacted with both language and transparency, such that the effect of priming condition was greater in the impoverished compared with the rich language ($F_{2,632} = 144.9$, $p < .001$), and with decreasing morphological transparency ($F_{6,632} = 206.6$, $p < .001$).

Given the concerns about the equivalence of the identity condition for the network compared with that in empirical studies, a second, $2 \times 2 \times 4$ ANOVA was carried out excluding this priming condition. This ANOVA produced exactly the same pattern of results as the first.

The effects of priming condition and morphological transparency are straightforward, and follow from the principles underlying the behaviour of connectionist networks. For both factors, RTs were faster when the prime had greater orthographic and semantic overlap with the target. The effects of language are more interesting, as these reflect not the

primes and targets themselves—which were held constant across the languages—but the different contexts in which these words were learned (see Principle 5 earlier). RTs were faster in the morphologically rich language than in the impoverished language because the overall organisation of the internal representations within the network is more morphologically structured in the former case. This organisation supports more effective processing of primes and targets that, at least to some degree, share this morphological structure. The effect of priming condition was stronger in the impoverished language essentially because there is more room for improvement in performance on the targets when base RTs are slower (see also Plaut, 1995; Plaut & Booth, in press; Plaut et al., 1996).

Critically, the effect of language context on performance is not restricted to morphologically structured words. Indeed, in a separate planned comparison restricted only to opaque targets and the morphological and control priming conditions, there was a reliable interaction of language and priming condition ($F_{1,19} = 7.269, p = .0143$), such that priming was greater in the rich language. In fact, separate *t*-tests showed that morphological priming (i.e., RTs following control primes minus RTs following morphological primes) was reliable in the rich language ($t_{19} = 3.780, p = .001$) but not in the impoverished language ($t_{19} = 0.536, p = .598$). Figure 7 shows the magnitudes of morphological priming at each level of transparency for the two languages.

Thus, as suggested in the discussion of Principle 5, the way in which a set of semantically opaque pseudo-affixed items (e.g., CORNER) is represented and processed by a connectionist network depends on the nature of the entire task on which it is trained. If the overall language has relatively little morphological structure, as in English, then opaque items are free to develop fairly idiosyncratic representations, even though their surface forms happen to contain segments (e.g., CORN, -ER) that operate as morphemes in other contexts. Hence, processing one such item as a prime will have little impact on processing another as target, relative to non-affixed control primes. By contrast, in a language that is highly morphologically structured, the processing of the stem and pseudo-affix in opaque items is more strongly influenced by how these elements are handled throughout the rest of the language. Their representation in opaque items necessarily has some similarity with their representation in more semantically transparent items, and hence with each other. Thus, the processing of a pseudo-affixed item will be facilitated by processing another such item with the same stem or affix. Note that, in this context, it isn't strictly appropriate to call such items *pseudo*-affixed as, although they are semantically opaque, they nonetheless participate in the morphological structure of the entire language.

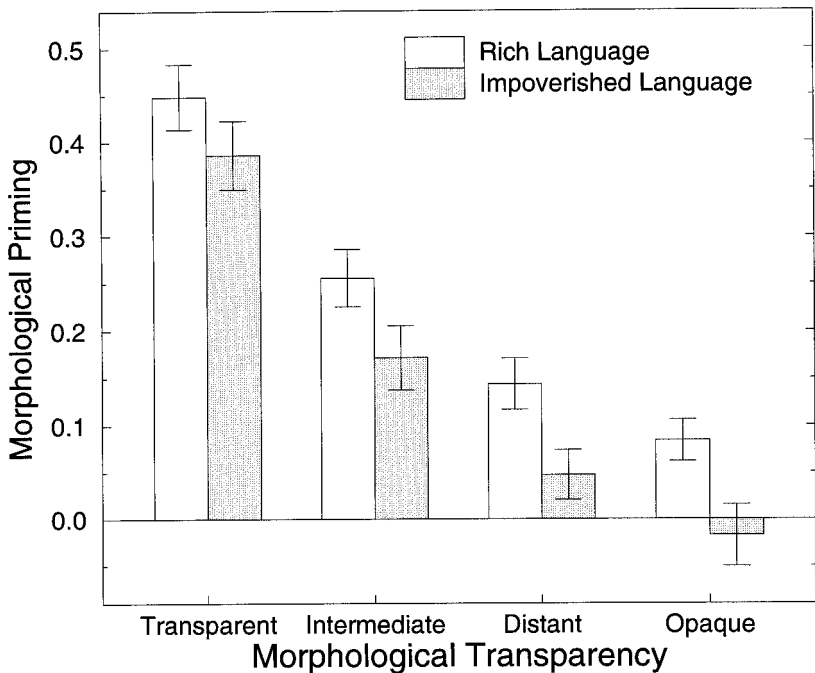


Figure 7. Mean priming, measured as mean RT following control primes minus mean RT following morphological primes, as a function of the transparency of the target word and the language on which the network was trained.

To demonstrate the effect of training context on the internal representations developed by the network, we determined the hidden representation in each version of the network of each word shared across the two languages. We then compared their pairwise similarities, measured in terms of Pearson correlation coefficients, as a function of the level of transparency of the two words and the nature of their orthographic overlap, if any.

Overall, the representations of all non-identical pairs of these common items were more similar in the morphologically rich language (mean $r = 0.14$) than in the impoverished language (mean $r = 0.06$; $t_{102078} = 75.65$, $p < .001$). Not surprisingly, this effect was carried largely by pairs with some degree of orthographic overlap. Across languages, the representations of pairs of items sharing a stem were more similar to each other (mean $r = 0.52$) than were pairs sharing second syllables (mean $r = 0.33$, $t_{99358} = 200.23$, $p < .001$) which, in turn, were more similar than pairs with no common syllables (mean $r = 0.06$, $t_{12646} = 64.78$, $p < .001$). The latter effect can be attributed to orthographic factors alone, but the difference

between overlapping in stem versus second syllable reflects the relative number of features in the canonical semantic representations of these items.

Table 3 lists the mean correlation coefficients for word pairs that overlap in either stem or affix, as a function of morphological transparency and language (and *t*-tests for effects of language). Pairs of items sharing stems were more similar in the rich compared with the impoverished language at each level of morphological transparency. These differences give rise to the effects of language on morphological priming that are apparent in Figure 7. Note, in particular, that the representations of opaque items sharing a stem are substantially more similar when learned within the morphologically rich language than when the same items are learned within the impoverished language.

Interestingly, none of the comparisons for affix overlap were reliable, although there may be a weak trend (given so few items) in favour of the rich language for opaque items. This suggests that the affixes themselves are treated in a systematic fashion regardless of the overall level of transparency of the language; this is presumably because there are only 10 of them and they are very common, amounting to 25% of all second syllables in the impoverished language and 85% in the rich language. These levels of overlap are, of course, much higher than for non-overlapping word pairs (mean $r = 0.06$). These findings suggest that the system would exhibit affix priming (i.e., DARKNESS priming TOUGHNESS) in both morphologically rich and impoverished languages. It is important to note, though, that the affixes in the simulation, even in the impoverished language, are all highly transparent, productive, and have unreduced (e.g.,

TABLE 3
Mean Pearson correlation coefficients for word pairs as a function of language and morphological transparency, and *t*-tests for effects of language

<i>Transparency</i>	<i>Language</i>		<i>t</i> -test
	<i>Rich</i>	<i>Impoverished</i>	
Stem overlap			
Transparent	0.58	0.42	$t_{898} = 23.8, p < .001$
Intermediate	0.62	0.46	$t_{898} = 23.7, p < .001$
Distant	0.63	0.43	$t_{898} = 28.6, p < .001$
Opaque	0.60	0.39	$t_{19} = 4.22, p = .001$
Affix overlap			
Transparent	0.34	0.34	$t_{898} = 0.58, p = .57$
Intermediate	0.34	0.33	$t_{898} = 1.03, p = .30$
Distant	0.33	0.33	$t_{898} = 0.47, p = .64$
Opaque	0.37	0.29	$t_{26} = 1.27, p = .21$

syllabic) phonological forms. In fact, empirical studies employing affixes of this sort have found reliable priming, both in Hebrew (e.g., Deutsch, Frost, & Forster, 1998) and in English (e.g., Marslen-Wilson, Ford, Older, & Zhou, 1996).

SIMULATION 2

The first simulation employed an abstraction of morphological structure in which all of the words formed from a given stem had the same degree of semantic transparency. Although this organisation may be sufficient to demonstrate graded influences of transparency, it is in many ways a poor approximation of the morphology of real languages, in which there is generally a range of transparency among words generated from a common stem. For example, the English stem *SHORT* is found in the transparent forms like *SHORTNESS* and *SHORTEST*, intermediate forms like *SHORTAGE*, *SHORTLY* and *SHORTS*, and opaque forms like *SHORTENING*. Thus, it seems important to confirm that the cross-linguistic differences in non-semantic morphological priming observed in Simulation 1 were not due to idiosyncratic aspects of the form of morphological structure that was employed. Accordingly, we carried out a second simulation that was as closely analogous to the first simulation as possible except that semantic transparency varied in a graded manner within the set of words derived from each stem.

Method

Network architecture. The network architecture was identical to that used in Simulation 1 (see Figure 5).

Representations. The orthographic representations and “transparent” semantic representations were the same as in Simulation 1; the only difference was in how these semantics were distorted to form word meanings. As in Simulation 1, 12 words were formed from each of 100 stems, with 40 “experimental” stems forming words that were common across the two languages. Ten of the words derived from each experimental stem were created by combining with each of the 10 affixes. Of these 10 words, two were *transparent* (no distortion to transparent semantics), two were *intermediate* (one-third of features regenerated randomly), two were *distant* (two-thirds of features regenerated), two were *opaque* (all features regenerated), and the remaining two were assigned to one of these transparency levels randomly. The other two words derived from each experimental stem were formed using randomly selected non-affix second syllables and were assigned opaque semantics. Thus, as in Simulation 1, the 40 experimental

stems created (approximately) 100 words at each of four levels of semantic transparency—although now every stem contributed to every transparency level—as well as 80 additional opaque words containing non-affix second syllables.

The remaining 60 stems (720 words) were treated exactly as in Simulation 1. In the morphologically *rich* language, each of these stems formed 10 semantically transparent words using each of the 10 affixes, and two opaque forms using other second syllables. In the *impoverished* language, they formed 12 semantically opaque words (two with randomly selected affixes and 10 with other second syllables).

Training procedure. The training procedure was the same as in Simulation 1. The network had to be trained slightly longer than in the first simulation in order to reach equivalent levels of cross-entropy error. For the rich language, the network required 360 sweeps through the 1200-word training corpus to achieve a mean cross-entropy error per semantic unit per pattern of 0.0033; for the impoverished language, 650 sweeps were required to achieve a mean of 0.0029.

Testing procedure. The testing procedure was the same as in Simulation 1.

Results and discussion

Table 4 presents the mean RTs of the network to target words of each level of systematicity as a function of priming context and language. The RTs were entered into a $2 \times 3 \times 4$ ANOVA with stem as the random variable and *language* (rich, impoverished), *priming condition* (identical, morphological, control) and morphological *transparency* (transparent, intermediate, distant, opaque) as within-item factors. The ANOVA revealed main effects of language ($F_{1,39} = 100.2, p < .001$), of priming condition ($F_{2,78} = 14727, p < .001$), and of transparency ($F_{3,117} = 64.11, p < .001$), and two-way interactions of transparency with language ($F_{2,78} = 185.6, p < .001$) and with priming context ($F_{6,234} = 186.7, p < .001$). In addition, the three-way interaction of transparency, context, and language was also reliable ($F_{6,234} = 3.64, p < .01$), such that the decrease in priming as a function of transparency was greater in the impoverished compared with rich language. The directions of these effects were all in accordance with the results of Simulation 1.

A second $2 \times 2 \times 4$ ANOVA that compared only the morphological and control priming conditions produced the same pattern of effects. Figure 8 shows the magnitudes of morphological priming (i.e., control condition minus morphological condition) at each level of transparency for the two

TABLE 4

Mean RTs (and standard errors) to target words as a function of language and priming condition, and magnitudes of priming effects

<i>Transparency</i>	<i>Priming condition</i>			<i>Priming</i>	
	<i>Identical</i>	<i>Morphological</i>	<i>Control</i>	<i>Identity</i>	<i>Morphological</i>
Rich language					
Transparent	3.43 (0.02)	4.06 (0.03)	4.71 (0.02)	1.28	0.65
Intermediate	3.54 (0.03)	4.54 (0.03)	4.84 (0.02)	1.30	0.30
Distant	3.54 (0.03)	4.66 (0.02)	4.86 (0.02)	1.30	0.20
Opaque	3.64 (0.02)	4.76 (0.02)	4.96 (0.02)	1.31	0.18
Impoverished language					
Transparent	3.51 (0.02)	4.19 (0.03)	4.72 (0.02)	1.22	0.54
Intermediate	3.60 (0.03)	4.68 (0.03)	4.80 (0.03)	1.20	0.12
Distant	3.64 (0.03)	4.82 (0.03)	4.85 (0.03)	1.20	0.02
Opaque	3.74 (0.03)	4.93 (0.02)	4.94 (0.02)	1.20	0.01

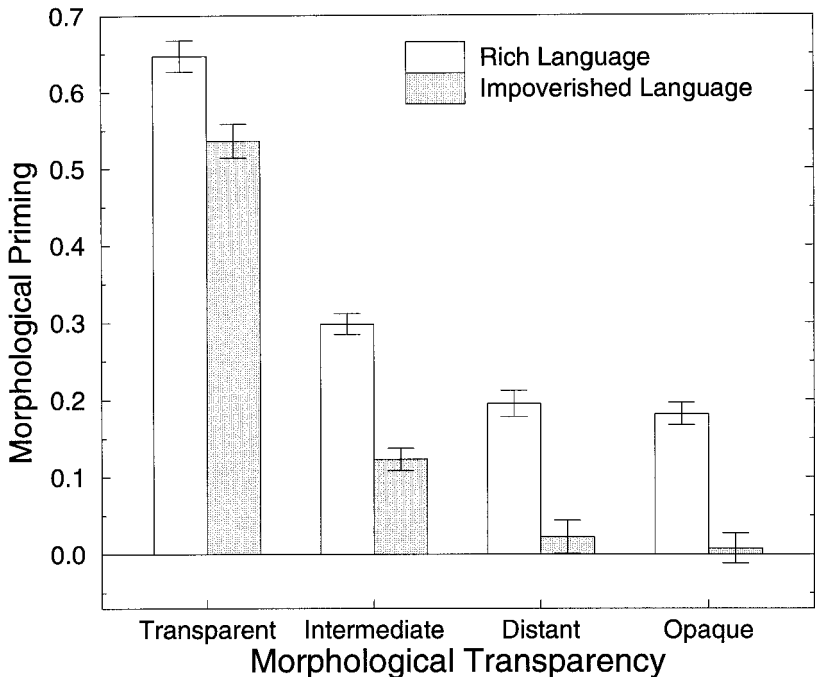


Figure 8. Mean priming, measured as mean RT following control primes minus mean RT following morphological primes, as a function of the transparency of the target word and the language on which the network was trained.

languages. As in Simulation 1, a planned comparison restricted only to opaque targets found a reliable interaction of language and priming condition ($F_{1,39} = 122.3, p < .001$), such that priming was greater in the rich language. Separate t -tests showed that morphological priming was reliable in the rich language (paired $t_{39} = 12.56, p < 0.001$) but not in the impoverished language (paired $t_{39} = 0.40, p = 0.69$).

Thus, the current simulation replicated the main findings of Simulation 1: (a) morphological priming effects varied as a function of graded degrees of semantic transparency of primes and targets; and (b) semantically opaque forms yielded priming only in the morphologically rich language. Indeed, the use of more natural morphological structure, in which semantic transparency varied both within and between stems, yielded an even stronger cross-linguistic difference in priming for opaque (and even distant) forms.

GENERAL DISCUSSION

Traditional theories of lexical processing assume that morphology is reflected explicitly in the structure of the lexicon: words are built out of discrete units called morphemes, and a given word either is or isn't decomposed into constituent morphemes at each stage of processing. In fact, given a reliance on symbolic, rule-based computation as the basis for language knowledge, it is difficult to imagine an alternative. What would it mean to say that one symbol (a word) can be decomposed into component symbols (morphemes) *partially*?

Distributed connectionist modelling offers an alternative. Words are represented as alternative patterns of activity over multiple groups of units, and subpatterns of those patterns can exhibit varying degrees of stability and independence in a natural way. Connectionist networks are sometimes described as simple 'associationist' (Pinker, 1991) or 'correlational' (Frost et al., in press-a) systems, but such descriptions misrepresent the computational sophistication of the approach. Networks can learn internal representations that reflect the underlying functional structure of the domains in which they are trained, and these representations can support effective handling of exceptional knowledge without compromising generalisation to novel forms (Plaut et al., 1996).

On a connectionist account, morphological processing reflects a learned sensitivity to the systematic relationships among the surface forms of words and their meanings. Given that these relationships exhibit graded degrees of systematicity, the approach predicts that the magnitudes of behavioural measures of morphological processing, such as priming, should vary continuously as a function of the degree of semantic and

formal transparency of stimuli. These predictions were recently confirmed by Gonnerman et al. (2000a).

The connectionist account would also seem to be supported by failures to find priming among morphologically related words in English that are not semantically related (Feldman & Soltano, 1999; Marslen-Wilson et al., 1994). By the same token, however, evidence for non-semantic morphological priming in Hebrew (Bentin & Feldman, 1990; Frost et al., 1997, in press-a, in press-b) has been assumed to present a challenge for the account. It turns out, however, that the predictions of the connectionist account in this context depend on the morphological structure of the language as a whole, and Hebrew is morphologically much richer than English.

We reported on simulations in which a network was trained on abstract versions of either a morphologically rich or impoverished language and then tested for morphological priming among a set of words in both languages that varied in semantic transparency. The network exhibited graded morphological priming as a function of semantic transparency in both languages (consistent with Gonnerman, et al., 2000a). More importantly, there was reliable priming for morphologically related but semantically opaque items in the rich language (consistent with Bentin & Feldman, 1990; Frost et al., in press-a, in press-b) but not in the impoverished language (consistent with Feldman & Soltano, 1999; Marslen-Wilson et al., 1994). Morphologically related but semantically opaque items in the rich language exhibit priming because the organisation of the internal representations in the network are dominated by the pervasive morphological structure of the language to such an extent that even opaque items participate in it. By contrast, in the impoverished language, the same items are free to behave idiosyncratically and, thus, exert little influence on each other. These findings suggest that, rather than being challenged by the occurrence in Hebrew of morphological effects in the absence of semantic similarity, the connectionist approach may be able to explain the cross-linguistic differences in the occurrence of these effects.

To be clear, we do not consider the current work, in and of itself, to constitute an account of the relevant empirical findings in English or in Hebrew. Indeed, we employed tasks that bear only the most abstract relationship to actual morphological systems. To be more realistic, the system would need to be modified both with regard to the tasks it would perform, and the input it would receive. For example, children learn the relationships between meaning and sound before they learn anything about orthography. While effects of literacy on interpretation of morphologically complex words have yet to be spelled out for English or Hebrew, knowledge of orthographic patterns will undoubtedly impact ultimate representations, since some morphological relationships in

English are maintained in orthography but obscured in speech (e.g., SIGN-SIGNAL). An adequate account of the empirical results, in our view, would be based on a simulation in which orthographic, phonological, and semantic representations all interact in the process of interpreting written and spoken words (see Figure 1). Moreover, the orthographic and phonological representations would need to approximate the actual systems of the two languages, and the semantic representations would need to capture salient aspects of the relationships among word meanings. In particular, constructing an adequate approximation of Hebrew orthography presents a number of challenges in that some aspects of word patterns are represented with diacritics and others are typically omitted in material for adults. Moreover, the richness of Hebrew morphology manifests not only in graded semantic transparency but also (and perhaps primarily) in the extensive number of inflections and derivations which common phonological structure across words. By contrast, our abstract rich and impoverished languages differed only in the token frequencies of affixes but not in their type frequencies. Additionally, while we examined degree of semantic transparency both across word forms (Simulation 1) and within them (Simulation 2), we did not manipulate other distributional characteristics that have been shown to be empirically relevant. For example, differential effects for morphologically complex words based on cumulative root frequency compared to surface-form frequency have been demonstrated for English (Taft, 1979) and in other languages (Italian: Burani & Caramazza, 1987; French: Colé et al., 1989). A more adequate model would thus also incorporate realistic type and token frequencies of both complex words and related simple words. We leave the development of such simulation to future work.

It is also important to be clear that the connectionist approach does not entail the elimination of morphology as an important level of analysis of linguistic structure and psycholinguistic behaviour. To the contrary, on our view, the morphological structure of a language has a fundamental impact on the organisation of lexical representation and processing. The critical difference between our perspective and more traditional ones is that, rather than stipulate morphological principles as facts that are independent of other aspects of the lexicon, we seek to derive them from independently motivated assumptions about the nature of orthographic, phonological and semantic representations in a language and the nature of learning and processing in the underlying (connectionist) mechanism. Note that morphology is no different than other levels of representation in this respect. For example, phonological representations can be understood as arising from learned relationships among acoustic, semantic, and articulatory information (Plaut & Kello, 1999) and semantic representations can be understood as arising from learned relationships among multiple input

and output modalities (McGuire & Plaut, 1997; Plaut, in press, in preparation; Rogers & Plaut, in press). By adopting a multi-pronged approach in which different simulations make linked approximations in the service of addressing a range of specific issues in lexical processing, we seek to develop a computationally explicit and empirically adequate theoretical account of language behaviour that spans perception to action.

REFERENCES

- Andersen, E.S. (1992). Complexity and language acquisition: Influences on the development of morphological systems in children. In J.A. Hawkins & M. Gell-Mann (Eds.), *The evolution of human language*, pp. 241–272. Reading, MA: Addison-Wesley.
- Andrews, S. (1986). Morphological influences on lexical access: Lexical or nonlexical effects. *Journal of Memory and Language*, 25, 726–740.
- Aronoff, M. (1976). *Word formation in generative grammar*. Cambridge, MA: MIT Press.
- Bentin, S., & Feldman, L.B. (1990). The contribution of morphological and semantic relatedness to the repetition effect at long and short lags: Evidence from Hebrew. *Quarterly Journal of Experimental Psychology*, 42A, 693–711.
- Bentin, S., & Frost, R. (1995). Morphological factors in visual word identification in Hebrew. In L.B. Feldman (Ed.), *Morphological aspects of language processing*, pp. 271–292. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Berman, R.A. (1978). *Modern Hebrew structure*. Tel Aviv, Israel: University Publishing Projects.
- Burani, C., & Caramazza, A. (1987). Representation and processing of derived words. *Language and Cognitive Processes*, 2, 217–227.
- Burani, C., Salmaso, D., & Caramazza, A. (1984). Morphological structure and lexical access. *Visible Language*, 18, 419–427.
- Bybee, J.L. (1985). *Morphology: A study of the relation between meaning and form*. Philadelphia, PA: John Benjamins.
- Chauvin, Y., & Rumelhart, D.E. (Eds.) (1995). *Back-propagation: Theory, architectures, and applications*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Chomsky, N. (1957). *Syntactic structures*. The Hague: Mouton.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, N., & Halle, M. (1968). *The sound pattern of English*. New York: Harper & Row.
- Cohen, J.D., Dunbar, K., & McClelland, J.L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, 97(3), 332–361.
- Colé, P., Beauvillain, C., & Segui, J. (1989). On the representation and processing of prefixed and suffixed words: A differential frequency effect. *Journal of Memory and Language*, 28, 1–13.
- Cottrell, G.W., & Plunkett, K. (1995). Acquiring the mapping from meanings to sounds. *Connection Science*, 6, 379–412.
- Crain, S. (1991). Language acquisition in the absence of experience. *Behavioral and Brain Sciences*, 14, 597–650.
- Deutsch, A., Frost, R., & Forster, K.I. (1998). Verbs and nouns are organized and accessed differently in the mental lexicon: Evidence from Hebrew. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(5), 1238–1255.
- Feldman, L.B. (Ed.) (1995). *Morphological aspects of language processing*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

- 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.
- Fodor, J.A., & Pylyshyn, Z.W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3–71.
- Fowler, C., Napps, S., & Feldman, L.B. (1985). Relations among regular and irregular morphologically related words in the lexicon as revealed by repetition priming. *Memory and Cognition*, 13, 241–255.
- Frost, R., Deutsch, A., & Forster, K.I. (in press). Decomposing morphologically complex words in a nonlinear morphology. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Frost, R., Deutsch, A., Gilboa, O., Tannenbaum, M., & Marslen-Wilson, W. (in press). Morphological priming: Dissociation of phonological, semantic, and morphological factors. *Memory and Cognition*.
- 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(4), 829–856.
- Gonnerman, L.M. (1999). *Morphology and the lexicon: Exploring the semantics-phonology interface*. Unpublished doctoral dissertation, Department of Linguistics, University of Southern California, Los Angeles, CA.
- Gonnerman, L.M., Andersen, E.S., & Seidenberg, M.S. (2000a). *Graded semantic and phonological effects in priming: Evidence for a distributed connectionist approach to morphology*. Manuscript submitted for publication.
- Gonnerman, L.M., Devlin, J.T., Andersen, E.S., & Seidenberg, M.S. (2000b). *Derivational morphology as an emergent, interlevel representation*. Manuscript submitted for publication.
- Grainger, J., Colé, P., & Segui, J. (1991). Masked morphological priming in visual word recognition. *Journal of Memory and Language*, 30, 370–384.
- Harm, M.W., & Seidenberg, M.S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*, 106(3), 491–528.
- Hertz, J., Krogh, A., & Palmer, R.G. (1991). *Introduction to the theory of neural computation*. Reading, MA: Addison-Wesley.
- Hinton, G.E. (1989). Connectionist learning procedures. *Artificial Intelligence*, 40, 185–234.
- Hinton, G.E., McClelland, J.L., & Rumelhart, D.E. (1986). Distributed representations. In D.E. Rumelhart, J.L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations*, pp. 77–109. Cambridge, MA: MIT Press.
- Hinton, G.E., & Sejnowski, T.J. (1986). Learning and relearning in Boltzmann Machines. In D.E. Rumelhart, J.L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations*, pp. 282–317. Cambridge, MA: MIT Press.
- Hoeffner, J.H., & McClelland, J.L. (1993). Can a perceptual processing deficit explain the impairment of inflectional morphology in developmental dysphasia? A computational investigation. In E.V. Clark (Ed.), *Proceedings of the 25th Annual Child Language Research Forum*, pp. 38–49. Stanford, CA: Center for the Study of Language and Information.
- Joanisse, M.F., & Seidenberg, M.S. (1999). Impairments in verb morphology after brain injury: A connectionist model. *Proceedings of the National Academy of Science, USA*, 96, 7592–7597.

- Kello, C.T., & Plaut, D.C. (1998). Modulation of lexical effects in word reading without strategic control of pathways: Evidence from tempo naming [Abstract 73]. *Abstracts of the Psychonomic Society*, 3, 7. (Proceedings of the 39th Annual Meeting, Dallas, TX.)
- Kello, C.T., & Plaut, D.C. (in preparation). *Strategic control by gain manipulation in a computational model of word reading in the tempo naming task.*
- Kello, C.T., & Plaut, D.C. (in press). Strategic control in word reading: Evidence from speeded responding in the tempo naming task. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*
- Kempey, S.T., & Morton, J. (1982). The effects of priming with regularly and irregularly related words in auditory word recognition. *British Journal of Psychology*, 73, 441–454.
- Laudanna, A., Badecker, W., & Caramazza, A. (1989). Priming homographic stems. *Journal of Memory and Language*, 28, 531–546.
- Laudanna, A., Badecker, W., & Caramazza, A. (1992). Processing inflectional and derivational morphology. *Journal of Memory and Language*, 31(3), 333–348.
- Laudanna, A., Cermele, A., & Caramazza, A. (1997). Morpho-lexical representations in naming. *Language and Cognitive Processes*, 12(1), 49–66.
- Marchand, H. (Ed.). (1969). *The categories and types of present-day English word-formation.* Munich: C.H. Beck'sche Verlagsbuchhandlung.
- Marslen-Wilson, W., Tyler, L.K., Waksler, R., & Older, L. (1994). Morphology and meaning in the English mental lexicon. *Psychological Review*, 101(1), 3–33.
- Marslen-Wilson, W.D., Ford, M., Older, L., & Zhou, X. (1996). The combinatorial lexicon: Priming derivational affixes. In G. W. Cottrell (Ed.), *Proceedings of the 18th Annual Conference of the Cognitive Science Society*, pp. 223–227. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- McClelland, J.L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, 86, 287–330.
- McClelland, J.L., Rumelhart, D.E., & the PDP Research Group (Eds.) (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 2: Psychological and biological models.* Cambridge, MA: MIT Press.
- McClelland, J.L., St. John, M., & Taraban, R. (1989). Sentence comprehension: A parallel distributed processing approach. *Language and Cognitive Processes*, 4, 287–335.
- McGuire, S., & Plaut, D.C. (1997). Systematicity and specialization in semantics: A computational account of optic aphasia. In *Proceedings of the 19th Annual Conference of the Cognitive Science Society*, pp. 502–507. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- McLeod, P., Plunkett, K., & Rolls, E.T. (1998). *Introduction to connectionist modelling of cognitive processes.* Oxford, UK: Oxford University Press.
- Minsky, M., & Papert, S. (1969). *Perceptrons: An introduction to computational geometry.* Cambridge, MA: MIT Press.
- Murrell, G.A., & Morton, J. (1974). Word recognition and morphemic structure. *Journal of Experimental Psychology*, 102, 963–968.
- Napps, S.E. (1989). Morphemic relationships in the lexicon: Are they distinct from semantic and formal relationships? *Memory and Cognition*, 17, 729–739.
- Napps, S.E., & Fowler, C.A. (1987). Formal relationships among words and the organization of the mental lexicon. *Journal of Psycholinguistic Research*, 16, 257–272.
- Page, M. (in press). Connectionist modelling in psychology: A localist manifesto. *Behavioral and Brain Sciences.*
- Pinker, S. (1991). Rules of language. *Science*, 253, 530–535.
- Pinker, S. (1994). *The language instinct.* New York: Morrow.

- Plaut, D.C. (1991). *Connectionist neuropsychology: The breakdown and recovery of behavior in lesioned attractor networks*. Unpublished doctoral dissertation, School of Computer Science, Carnegie Mellon University. (Available as Technical Report CMU-CS-91-185.)
- Plaut, D.C. (1995). Semantic and associative priming in a distributed attractor network. In *Proceedings of the 17th Annual Conference of the Cognitive Science Society*, pp. 37–42. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Plaut, D.C. (in preparation). *Systematicity, topography, and graded modality-specific specialization in semantics: An account of optic aphasia*.
- Plaut, D.C. (in press). Systematicity and specialization in semantics. In G.W. Humphreys, D. Heinke, & A. Olson (Eds.), *Proceedings of the Fifth Neural Computation and Psychology Workshop: Connectionist models in cognitive neuroscience*. New York: Springer-Verlag.
- Plaut, D.C., & Booth, J.R. (in press). Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. *Psychological Review*.
- Plaut, D.C., & Kello, C.T. (1999). The interplay of speech comprehension and production in phonological development: A forward modeling approach. In B. MacWhinney (Ed.), *The emergence of language* (pp. 381–415). Mahwah, NJ: Lawrence Erlbaum Associates Inc.
- Plaut, D.C., & McClelland, J.L. (1993). Generalization with componential attractors: Word and nonword reading in an attractor network. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society*, pp. 824–829. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Plaut, D.C., McClelland, J.L., Seidenberg, M.S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*, 56–115.
- Plaut, D.C., & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, *10*(5), 377–500.
- Prasada, S., & Pinker, S. (1993). Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes*, *8*, 1–56.
- Quinlan, P. (1991). *Connectionism and psychology: A psychological perspective on new connectionist research*. Chicago: University of Chicago Press.
- Rogers, T.T., & Plaut, D.C. (in press). Connectionist perspectives on category-specific deficits. In E. Forde & G.W. Humphreys (Eds.), *Category-specificity in brain and mind*. New York: Psychology Press.
- Rueckl, J.G., & Dror, I.E. (1994). The effect of orthographic-semantic systematicity on the acquisition of new words. In C. Umiltà & M. Moscovitch (Eds.), *Attention and Performance XV: Conscious and nonconscious aspects of cognitive processing*, pp. 571–588. Cambridge MA: MIT Press.
- Rueckl, J.G., Mikolinski, M., Raveh, M., Miner, C.S., & Mars, F. (1997). Morphological priming, fragment completion, and connectionist networks. *Journal of Memory and Language*, *36*(3), 382–405.
- 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., Hinton, G.E., & Williams, R.J. (1986a). Learning internal representations by error propagation. In D.E. Rumelhart, J.L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume I: Foundations*, pp. 318–362. Cambridge, MA: MIT Press.
- Rumelhart, D.E., McClelland, J.L., & the PDP Research Group (Eds.) (1986b). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations*. Cambridge, MA: MIT Press.

- Schreuder, R., & Baayen, H. (1995). Modeling morphological processing. In L.B. Feldman (Ed.), *Morphological aspects of language processing*, pp. 131–154. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Seidenberg, M.S., & McClelland, J.L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*, 523–568.
- Smolensky, P., Mozer, M.C., & Rumelhart, D.E. (Eds.), (1996). *Mathematical perspectives on neural networks*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Stanners, R.R., Neiser, J.J., Herson, W.P., & Hall, R. (1979). Memory representation for morphologically related words. *Journal of Verbal Learning and Verbal Behavior*, *18*, 399–412.
- Stolz, J.A., & Besner, D. (1998). Levels of representation in visual word recognition: A dissociation between morphological and semantic processing. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(6), 1642–1655.
- Taft, M. (1979). Lexical access via an orthographic code: The Basic Orthographic Syllable Structure (BOSS). *Journal of Verbal Learning and Verbal Behaviour*, *18*, 21–39.
- Taft, M., & Forster, K. (1975). Lexical storage and retrieval of prefixed words. *Journal of Verbal Learning and Verbal Behaviour*, *14*, 638–647.
- Taraban, R., & McClelland, J.L. (1987). Conspiracy effects in word recognition. *Journal of Memory and Language*, *26*, 608–631.
- Werbos, P.J. (1974). *Beyond regression: New tools for prediction and analysis in the behavioral sciences*. Unpublished doctoral dissertation, Harvard University.
- Zorzi, M., Houghton, G., & Butterworth, B. (1998). Two routes or one in reading aloud? A connectionist “dual-process” model. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1131–1161.