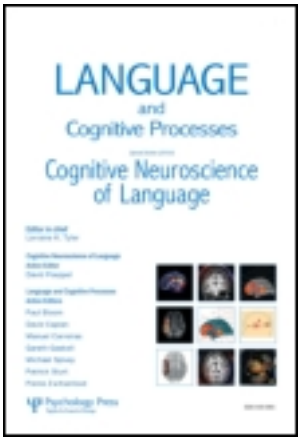

Language and Cognitive Processes
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/plcp20

No escape from morphemes in morphological processing

Alec Marantz [a b]

[a] Departments of Linguistics and Psychology , New York University , New York , NY , USA

[b] NYUAD Institute, New York University Abu Dhabi , Abu Dhabi , UAE

Published online: 29 Apr 2013.

To cite this article: Alec Marantz (2013) No escape from morphemes in morphological processing, Language and Cognitive Processes, 28:7, 905-916, DOI: 10.1080/01690965.2013.779385

To link to this article: http://dx.doi.org/10.1080/01690965.2013.779385

LANGUAGE AND COGNITIVE PROCESSES, 2013
Vol. 28, No. 7, 905–916, http://dx.doi.org/10.1080/01690965.2013.779385

No escape from morphemes in morphological processing

Alec Marantz[1,2]

[1]Departments of Linguistics and Psychology, New York University, New York, NY, USA
[2]NYUAD Institute, New York University Abu Dhabi, Abu Dhabi, UAE

Any approach to understanding morphological processing must begin with assumptions about the role of morphemes in linguistic representations. Contemporary linguistic theory proposes that such representations are centred on a syntactic organisation of morphemes, where a morpheme is an abstract syntactic unit that finds an interpretation in form and in meaning. From the linguist's perspective, then, morphological processing falls together with syntactic processing, and both involve the central exploitation of a grammar of morphemes. Recent computational work has challenged this perspective, proposing instead that morphological structure emerges from the acquisition of form/meaning connections. Here we show that these challenges to the role of morphemes in word and sentence processing rest on a misunderstanding of the morpheme within linguistic theory. Once we understand the notion of a morpheme and its explanatory role within linguistic theory, we can see that apparent challenges to the morpheme, rather than dispensing with the morpheme in the architecture of linguistic representations, always in fact presuppose a syntactic structure of morphemes as the unlearned input to language learning models. Since computational modelling is indispensable for work connecting theories of linguistic representation to theories of (morphological) processing, it is crucial that those constructing computational models appreciate why there is no escape from morphemes in morphological processing.

Keywords: Morphology; Linguistics; Paradigmatic dimension; Syntagmatic dimension; Generative grammar.

The "morpheme" is a slippery notion within linguistic theory. On one tradition, morphemes are sound/meaning connections that serve as the minimal building blocks of language (e.g., in "lexicalist" approaches; see Lieber (1992)). However, following lines of research always in play within linguistics, most morphologists over the last 20 years have converged on a view that e.g., Robert Beard (1995) and Mark Aronoff (1994) have labelled "separationist". On this view, morphemes are abstract units assigned phonological values by the grammar. Each morpheme might be realised by a variety of phonological "exponents" (for example, the English plural morpheme is realised generally by *–(e)s* but by *–en* after "ox") and each phonological exponent

Correspondence should be addressed to Alec Marantz, Department of Linguistics, New York University, 10 Washington Place, New York, NY 10003, USA. Email: marantz@nyu.edu

I would like to thank the editors of this volume and the anonymous referees for their helpful comments as well as Allyson Ettinger for her editorial support. This material is based upon work supported by the National Science Foundation under Grant No. BCS-0843969 and by the NYU Abu Dhabi Research Council under Grant No. G1001 from the NYUAD Institute, New York University Abu Dhabi.

might realise a number of different morphemes (e.g., *–ing* realises progressive aspect—"He is leaving"—and is also the realisation of a nominalising morpheme on some verbs—"the running of a business"). But the morpheme itself has an identity independent of its phonological realisation, an identity that determines its "syntactic" behaviour, i.e., its distribution within words and phrases relative to other morphemes.

The notion of "abstract" when applied to morphemes characterises not only the independence of a particular morpheme from a specific realisation in form but also its independence from a specific semantic value. For example, the "plural" morpheme in English is the same morpheme, appearing in the same syntactic position with respect to the noun stem, whether it conveys a meaning associated with "more than one" (as in "cats" and "oxen") or whether it does not (as in "(eye-)glasses" or "pants"). More generally, the features associated with morphemes by linguists, while connected to their syntactic and semantic properties, find their grammatical import (their role in the syntax in particular) independent of their possible semantic (or phonological) interpretations.

As described in some detail within the theoretical framework of Distributed Morphology (Halle & Marantz, 1993, 1994; Harley & Noyer, 1999), the separationist understanding of morphemes allows for a complete integration of morphology with the syntax (the "grammar" in everyday language) such that the internal structure of words finds analysis within the same syntactic architecture and subject to the same syntactic principles as the internal structure of phrases and sentences. With this understanding of morphemes and morphology, the psycho- and neurolinguistic investigation of morphological processing should fall together with the investigation of the processing of phrases and sentences. On this view, the goal of language acquisition is not primarily to find the mapping between form and meaning for a language but to discover the morphemes and grammar of morphemes that accounts for observed connections between form and meaning.

In opposition to this point of view, one finds claims in the computational and psycholinguistic literatures that morphemes and the internal syntax of words are emergent properties of learning the connections between sound and meaning (or the connections among sound, orthography and meaning) and are therefore not fundamental units in the knowledge/internal representation of language. Some work in this anti-morpheme tradition can be taken as making a strong cut between words, whose potential internal structure would only be emergent, and phrases and sentences, whose analysis would involve learning a grammar that governs the distribution of words. However, in a recent paper, Baayen, Milin, Durdevic, Filpovic, and Hendrix (2011) have fully embraced a position, also associated with Distributed Morphology that dissolves a word-internal vs. word-external distinction. In contrast to the assumptions of Distributed Morphology, on their view, there are no morphemes at all—neither at the phrase and sentence level nor internal to words—and all apparent syntactic structure would be emergent from the computation of sound(orthography)/meaning connections.

In the first section of this paper, "The Morpheme in Linguistic Theory", I will explain, at the most general level, why linguists believe that morphemes are fundamental to our knowledge of language and what kinds of data are explained by analyses in which words and sentences are treated as structures of morphemes. In the next section, "No Escape from Morphemes", I will turn to two examples of claims that morphemes are emergent properties of systems learning sound/meaning connections. Straightforwardly, the instantiation of these claims in the literature involve more levels of abstract units than the linguistic systems described in the first section of this

paper, and always build (abstract) morpheme structure into their "semantics". That is, the existing computational work does not demonstrate that morphemes might be an emergent property of systems learning language. Although this work is extremely important for linguists and other cognitive scientists, it does not present a challenge to the linguistic conception of knowledge of language at the moment, but does point to the possibility of such a challenge. The final section of the paper will briefly explore how one might try to investigate the role of morphemes and syntax (since, on the view presented in the first section, syntax is entirely about morphemes, and each requires the other) in language processing and really test whether a theory with morphemes does better in explaining or accounting for something than a theory without morphemes.

THE MORPHEME IN LINGUISTIC THEORY

Abstracting away from most differences among theoretical approaches in linguistics, linguistics can be said to have discovered that the syntactic structure of phrases and words is a structured list of morphemes, where morphemes are the minimal structural units that anchor the connection between form (sound, orthography, signs in signed languages) and meaning. Within linguistics, we contrast the analysis of language along the syntagmatic (co-occurrence) and paradigmatic (complementary distribution) dimensions. On the syntagmatic dimension, the grammar describes the distribution and co-occurrences of morphemes within the structured list of a syntactic representation, with morphemes sharing a distribution belonging to the same syntactic category. For example, determiners like "the" generally co-occur with nouns like "cat". Also along the syntagmatic dimension, we can place the usually obligatory presence of a tense morpheme on a verb in the main clause of an English sentence. That is, the co-occurrence of a tense morpheme with a verb is a property of the syntax of English, similar in important respects to the co-occurrence of determiners and nouns. In this respect, word-internal syntax (verb plus tense marking) and word-external syntax (noun plus determiner) are clearly parallel.

The paradigmatic dimension includes two aspects. First, there are choices of morphemes within a category of morphemes whose distribution along the syntagmatic dimension is constrained by the grammar. To oversimplify a bit, in a position for morphemes of the category noun, one might choose "cat" or "dog", or "lion"—these morphemes belong, in a sense, to the paradigm of the category of "noun" and one, and only one, such morpheme from the paradigm will fill a syntactic slot for nouns. Also along the paradigmatic dimension of "choice" we find the analysis of allomorphy—the conditioned choice of phonological form for a morpheme. For example, the plural morpheme is realised by *–s* in the environment of "cat" but by *–en* in the environment of "ox".

Much theoretical debate in the linguistics morphology literature has centred around the apportioning of data to analyses within the paradigmatic vs. the syntagmatic dimensions. For example, Distributed Morphology takes inflectional morphology such as tense, number, aspect and agreement to involve morphemes whose distribution is governed by syntax, i.e., determined along the syntagmatic dimension. Morphologists in the Word and Paradigm framework (e.g., Matthews, 1965), on the other hand, tend to treat inflected forms of a word, say a verb in all tenses, aspects and varieties of agreement, as filling a single slot along the syntagmatic dimension, being distributed in syntax as verbs, with the choice of inflected forms

being negotiated along the paradigmatic dimension. That is, "eat, ate, eaten, eats and eating" would belong to the paradigm of the verb "eat" and only one member of the paradigm would be chosen to fill any syntactic position for verbs. On the Distributed Morphology view, the internal structure of inflectional morphology receives a syntactic treatment, while under the Word and Paradigm view, inflectional morphology does not decompose into morphemes governed by the syntax.

Note, however, that this contrast between a syntagmatic and paradigmatic approach to inflectional morphology masks a broader agreement among morphologists about the morphological essence of "morphology" in the crucial sense of "morphemes as abstract units seeking realization in form and meaning". For Word and Paradigm morphologists as well as so-called "item and arrangement" morphologists (those that treat inflection along the syntagmatic dimension, including Distributed Morphologists), an inflected word includes a representation of the stem as well as a set of features (e.g., for verbs, tense, aspect and agreement features) that govern the phonological realisation of the word, its syntactic behaviour and its semantic interpretation. On every view, these features are abstract units which by themselves or bundled together would constitute morphemes in the crucial grammatical sense. Placing a set of these features with a verb stem for phonological realisation (organising paradigms of verbs along the dimensions defined by the features) constitutes reliance on a syntax of morphemes—why are we putting these features on the verb and not others (say case features), why are some features in complementary distribution (e.g., past tense precluding present tense), etc. Within linguistics, it is the job of morphology and syntax to account for the organisation of these (abstract) features within constituents that are realised phonologically. Again, the debates in morphological literature relevant to the current discussion concern the treatment of morphological features along the syntagmatic and paradigmatic dimensions (thus, the contrast between "item and arrangement" and "word and paradigm"), not the existence of (abstract) morphological features nor the reliance on the syntax to account for the distribution of these features.

On most views of morphology (putting aside now extreme versions of Word and Paradigm theories in which all morphology is analysed along the paradigmatic dimension), derivational morphology (characterised as changing the category of a word) at least involves a syntax identical for present purposes to the syntax of words—an analysis involving combinations of morphemes. For example, the *–er* that appears to attach to verbs to create a noun describing a person or thing that habitually does what the verb describes ("walk-er") is generally assumed to be the form of a morpheme that combines with a verb stem in more or less the same way, syntactically, as a determiner combines with a noun. Putting aside issues about the possible similarities and differences between word-internal and word-external syntax, actual debate within Linguistics about the extent to which words display an internal syntax usually centres around "fusional" inflectional morphology in which multiple linguistic features are associated with affixes that do not seem to decompose into phonological pieces that can be associated with individual features or particular sets of features. Take for example the English present tense ending *–s*, which signals not only tense but also third person singular subject agreement in, "John walks". Are there two (or more) morphemes organised along the syntagmatic dimension in the structure spelled out orthographically as *–s*? As a morphologist, I can present arguments that there is a right answer to this question, but the validity of the general morphological approach to word structure does not depend on settling such questions.

Another live debate within linguistic morphology concerns the nature of the grammar appropriate for describing the distribution of morphemes within words. Within Distributed Morphology and related theories, the grammar internal to words is the same as the grammar external to words, involving hierarchical organisation of morphemes distributed according to their categories (such a "noun" or "verb" or "tense"). A different tradition proposes that word-structure involves templates with "slots" designated for sets of morphemes that do not necessarily share syntactic or semantic features (see the discussion in Rice, 2000). Baayen et al. (2011) suggest that Anderson's A-morphous morphology (1992) supplies an alternative to the view that the internal structure of morphemes is governed by grammars or templates, but this is a misunderstanding of Anderson's system. First, Anderson supposes that derivational morphology does involve a syntactic structure of morphemes—A-morphous principles apply only to a narrowly defined set of inflectional material. Second, Anderson supposes that inflectional morphology involves features assigned to, e.g., nouns and verbs by the syntax and these features themselves have an internal organisation that is reflected in their phonological realisation (e.g., there is a hierarchical relation among sets of agreement features on a verb). As suggested above, these feature bundles should be considered as morphemes on the characterisation of morphemes being explained here, and their placement on a verb or noun, as well as their organisation, involves a syntax in the standard sense. A-morphism for Anderson involves the way in which the syntactic organisation of morphemes is realised phonologically, not a lack of commitment to morphemes or their syntactic arrangement in the analysis of language. Finally, Anderson's organisation of the morphophonology into disjunctive rule blocks along with his blocking principles actually imposes a templatic structure on words, with each rule block defining a templatic slot. In short, almost every morphologist attempting to explain complementary distribution of morphological material along the paradigmatic dimension and co-occurrence of morphological material along the syntagmatic dimension has concluded that words have syntactic structure of some sort, involving bundles of abstract features which we call morphemes. The disagreements among morphologists are highlighted in the literature; the enormous agreement among morphologists based on a long history of empirical work on numerous languages goes much deeper.

On what grounds could we conclude that morphologists over the years have been misguided in their belief that words derive from a grammatical combination of morphemes? Clearly, alternatives to morpheme-centric accounts of language would need to confront the evidence from co-occurrence and complementary distribution that motivate existing grammatical analyses and provide alternative accounts of the data. The next section explains that recent computational critiques of the linguistic morpheme, rather than confronting the linguist's evidence, presuppose the results of the morphemic analysis of words in their models. Morphemes, then, do not emerge from computational learning but serve as givens in the learning models.

NO ESCAPE FROM MORPHEMES

For present purposes, I will restrict attention here to two representative examples of the claim that morphemes are an emergent property of a system that learns the connections between form and meaning: the "triangle" model of Seidenberg (Seidenberg & McClelland, 1989), Plaut (Plaut, McClelland, Seidenberg, & Patterson, 1996) and others, and the recent A-morphous model of Baayen et al. (2011). Any

implemented computational model must make certain somewhat arbitrary decisions to allow for coding and testing, so we must be careful to understand the general claims of the model independent of potentially irrelevant features of the actual test cases of the model presented in the literature. Here we want to give the models every benefit of any doubt. So, for example, since computational models of morphological structure usually involve inputs to learning mechanisms in the form of structured phonological and semantic representations, we should suppose that a child learning language has extracted much from the regularities in his/her environment that might serve as the basis for establishing form/meaning connections for language. In particular, we might suppose a relatively sophisticated knowledge of the sound system of the language to be learned, such that incoming linguistic information is organised to some degree—idealised as a string of phonetic units, perhaps. Similarly, we should allow that the child has acquired concepts that make sense of his/her experience, grouping say dogs together in a category, and perhaps understanding aspects of cause/effect, etc., that might underlie verb meanings.

Looking at the architecture of linguistic theory with an eye to a comparison with the computational models, the linguistic morphologist is assuming one "layer" or level of abstract units: the level of syntax/morphology that assigns each utterance a representation consisting of a structured list of morphemes, themselves abstract units. The morphemes connect to form and to meaning, but the levels of phonological and semantic representations are not necessarily abstract. That is, as described by phonologists, phonological units are grounded in acoustics and articulation; similarly semantic units are grounded in "meaning", variously theorised by different groups of semanticists. While contemporary phonologists would not necessarily consider a "phoneme" (unit of phonological analysis) to be a set of phones (acoustically and/or articulatorily defined units), the properties of phonological units are nonetheless contaminated by their phonetic grounding, with some phonological generalisations partially understood in terms of their phonetic motivations. Similar considerations hold for units or features of semantic analyses. Baayen et al. (2011), however, in doing without a level of abstract morphemes, correctly identify their phonological and semantic representations as involving abstract units—they are acknowledging both the lack of any specified relation between their phonological (actually orthographic) and semantic categories and a phonetic or real-world ground and the lack of contamination from grounding in the model (nothing about the "meanings" of the orthographic or semantic units play a role in the model). So Baayen et al. suppose more sets of abstract units than the standard linguist.[1] For the triangle model, the issue is less clear. In principle, this model is close to the standard linguistic model; one might suppose the "input" levels of meaning, sound, and orthography to be grounded (and "contaminated") in some way, leaving only the hidden unit layer that mediates among the other layers as the locus of abstraction. However, in practice the semantic layer in particular is treated as abstract, in the sense of Baayen et al., and we will see

[1] A referee points out that the Baayen model could follow linguistic theory and incorporate semantic and phonological/orthographic units that are grounded or contaminated by their connection to interpretation; the abstract nature of these units in the model is not a matter of principle or architecture, in the way that the abstract nature of the linguists' morphemes is crucial to their identity. Below we will see that the abstractness of the semantic level in the Baayen model is tied to the actual functioning of the model; the semantic level builds in a syntactic structure of abstract morphemes without which the model would not perform as it does. The comparison of numbers of abstract levels between approaches to language here is loose; the real point is that adequate approaches to language will find one level of abstract morphemes inescapable, and for the Baayen model under consideration, this level is identified as the semantic one.

why this is important shortly. In sum, though, we see that a comparison among the linguist's proposal about the nature of language and the two representative computational approaches that lack morphemes would not rest on some notion that the linguist is proposing greater abstraction than the computationalists.[2]

Nor, at this general level of discussion, is there a difference among the approaches on the postulation of specific built-in (or universal) knowledge. For the morphologist, what is "innate" (brought to language learning by the child) is the assumption that there is a grammar, i.e., that the connection between form and meaning is mediated by a structured list of (abstract) morphemes and that learning a language involves learning the grammar of this structured list. For the linguist, the typology of languages—the range of ways in which they go about doing their form/meaning connecting business—displays their reliance on grammatical structure. When we as linguists imagine what a language might be like if, for example, the connection between form and meaning were "amorphous" (without the mediation of morpheme units), we can describe types of languages that do not exist (see the discussion in Embick & Marantz, 2008; Embick, 2010). Many of these restrictions involve "locality": relations between form and meaning are mediated by the syntax such that form (sound) realisations of meanings are constrained to syntactically local parts of sentences. For example, although tense marking in a language might "spread" to various constituents within a verb phrase in a manner that follows syntactic structure (see Richards, 2007), a language would not signal past tense by raising all the vowels in a sentence, for example, although that would involve a direct and transparent connection between form and meaning. Of course it would be quite interesting if computational models could show that the "look" of morphism (being governed by structure of morphemes) follows from constraints on language learning derivable from the most general mechanisms for learning form/meaning connections and so the "impossible" languages that linguists identify are impossible because they are effectively unlearnable (not unlearnable in principle, but unlearnable in practice). For example, explanations for some apparent locality constraints on grammars might be approached through modelling. However, at the moment linguists have been enormously successful at explaining what occurs and does not occur across languages through reference to the grammar of morphemes that mediates form/meaning connections.

Note again that there are no obvious simplicity arguments that favour amorphous approaches to language (the ones under discussion at least) to approaches with a syntax internal to words—the approaches assume similar degrees of "abstractness" and complexity. However, morphologists should believe that the assumption of a grammar will help acquisition of form/meaning connections, given the way the world is, rather than putting an extra acquisitional burden on the child learning language. Some evidence for this comes from the fact that the amorphous approaches sneak morpheme structure into their systems, as we show below, presumably because this helps make the models work.

[2] A referee points out that many linguistic models of morphology assume, in addition to abstract morphemic units, a set of morphological rules (e.g., copying, movement, insertion) that could be conceptualised as adding additional levels of abstract units or structure to the theory. The necessity of these rules, and thus the need for a syntax of morphemes more powerful than a context-free phrase structure rule grammar, is of course a matter for debate among linguists and computational scientists, not between them. The points made in this paper concern the necessity of morphemes and SOME syntactic structure that arranges them; the necessity for various types of grammatical rules can be argued and disputed separately.

The undeniable fact is that the computational approaches under discussion, although claiming to be amorphous (without morphemes) in an essential sense, build morphological structure into the inputs to their learning models, most obviously in the semantic units and their organisation. Phonology perhaps goes a long way to identifying and categorising units based on distribution of sounds independent of their connection to meaning or morphemes. One can imagine learning much about the phonological units and the phonological "grammar" of a language simply through distributional analysis of the utterances of speakers. Similarly, some semantic concepts and categories seem available to a pre-linguistic child ready to be correlated with the sounds produced by the people in his/her environment. It is beyond the realm of plausibility, however, to suppose that the meaning of "past tense" in English is something uniform and available as some simple environmental input to the child in the context of hearing the past tense form of a verb. Linguistic semanticists have a difficult time characterising the meaning of past tense, and in any case the uses of past tense verbs are quite varied in any language.

Implementations of the triangle model present to the model "learning" the past tense some representation of the meaning of the stem of the verb, which is invariant each time the verb is presented, plus a representation of the meaning of the tense morpheme, either past or non-past in (see e.g., Joanisse & Seidenberg, 1999, p. 7593). If we present to a learning model as "semantic input" a vector of binary values associated with the meaning of the stem of a verb and a one or a zero indicating past vs. present tense, we are in fact presenting the syntactic structure of the verb to the learning model, not anything that could reasonably be called a semantic input. In the terms introduced earlier, we are giving the model the syntactic representation of the word, analysed along the syntactic and paradigmatic dimensions. Along the syntactic dimension, we are indicating the co-occurrence of a verb with a tense morpheme—actually, the obligatory co-occurrence of these categories. We are also excluding any other aspects of meaning from consideration in the mapping to the form of a verb (e.g., semantic characteristics of the direct object are not part of the semantic input, although we know that some languages spell out object agreement on the verb). Along the paradigmatic dimension, we are indicating that only one verb meaning is relevant to the form of an English verb (we don't in English, for example, encode conjunction of verbs ("John ran and played") by realizing some blend of forms associated with the individual verbs), and one and only one tense meaning (so a "recent past" meaning is not coded as some weighted combination of past and present). Similar considerations apply to Baayen et al.'s (2011, p. 452) approach to Serbian case inflection. Their model assumes that nouns come with "grammatical meanings" for number (singular vs. plural) and case (nominative, genitive, dative, accusative, locative and instrumental). Every noun comes with one meaning from the number set and one from the case set, but no other "grammatical" meanings. This model clearly builds the morphological structure of the Serbian noun into the "semantics", demanding number and case marking along the syntagmatic dimension and requiring only a single instantiation of number and one of case along the paradigmatic dimensions.

We can unpack the issues here into three subcomponents. First, it is unrealistic enough with the meaning of stems like "walk" or "cat" to present a uniform "semantic" input for the stem on each occurrence of the phonological form. For present purposes, we are more than willing to overlook this issue, since it is reasonable to suppose that some concepts are acquired prior in some sense to learning the language and that attention or focus to one concept at a time might be an extra-linguistic cognitive

strategy brought to the language learning task. However, the notion of a uniform semantic input associated with a tense marking does not correspond to any reality nor to anyone's idea about meaning (compare the current time and habitual use of the "present tense" in "John knows the answer" and "John runs every day" or the past time and counterfactual uses of the "past tense" in "John knew the answer" and "If John knew the answer, he would have spoken up"). The types of features associated with inflectional morphology are not the type that one would suggest are distilled into features as generalisations about the environment independent of learning a language, and the use and distribution of such features is governed by the syntax, not primarily by the meanings conveyed by sentences.

Second, the kind of representation that presents one semantic value for tense (say either past or present, or either past or non-past) or case on nouns presupposes complementary distribution for tense marking and case marking, while complementary distribution for tense/case marking is a syntactic, not a semantic fact about languages (like English or Serbian). Third, and related to the second point, the input to the learning system involves a co-occurrence of the stem "meaning" and a "tense" meaning (or a noun meaning plus a number and a case meaning) and nothing else. That is, learning is based on a structured, minimal list of the units to be realised phonologically; a syntax of morphemes. In general, the learning models I have seen that claim to show that morphemes are an emergent property of the learning process provide to the system the end product of morphological learning—a structured list of morphemes—in the semantic input layer, choosing a subset of the semantic properties associated with a word as input and imposing co-occurrence restrictions along the syntagmatic dimension and complementary distribution along the paradigmatic dimension. Baayen et al. acknowledge this in their upfront identification of their semantic layer as abstract, but the abstractness—syntacticness—of the "semantic" layer is not acknowledged openly within the triangle model tradition.

Linguists may sense in the computational and psychological literature denying the existence of morphemes a misunderstanding of what a morpheme might be within linguistic theories. It is quite clear from the discussion in Baayen et al. (2011) that the authors take criticism of the lexicalist notion of a morpheme (where morpheme units are like Sausarian signs, linking sound and meaning independent of the grammar) to be criticism of morpheme-centred theories of grammar in general. They seem unaware that contemporary theories of morphology such as Distributed Morphology also reject the lexicalist position on the role of morphemes. Similarly, discussions of the nature of morphemes within the literature on the triangle model seem disconnected from linguistic morphological theory. Readers familiar with morphological theories will recognise immediately that the "tests" of the triangle model described within the literature build the morphological structure of words into the input to the learning models. The models then learn how to realise the morphological structure phonologically, learning about allomorphy (alternative phonological realisations of the same morpheme, dependent on the context, i.e., the particular other morphemes with which the morpheme is in construction). As long as the semantics is pre-morphologised and pre-syntactised in these models, they are not testing the feasibility of doing without a morpheme/syntax layer.

One might attribute the differences between linguists' morpheme-centric approaches to language and computational amorphous approaches to a difference in the data to be explained by the approaches. For example, the computational models addressed here take reaction times and errors in psycholinguistic experiments as (among) primary data for explanation, while linguists tend to emphasise the

distribution of linguistic forms in a language and speakers' intuitions about the well-formedness of words and sentences. Elsewhere I have explained that the generative approach to language in the Chomskyan tradition has always aimed to characterise a speaker's knowledge of language that must be part of the explanation of anything the speaker does with language, including his/her reaction times in experiments involving linguistic stimuli (Marantz, 2005). That is, the linguist is responsible for the same types of data that drive the computational theories explored here. The emphasis of the present paper has been aimed directly at the computational approaches and the data they present as relevant to the issue of the existence of morphemes. It has been argued that models discussed have incorporated the linguist's syntax of morphemes into their semantic "layers" and this adoption of morphemic structure is in fact necessary to achieve the goals of the models. That is, any difference in emphasis between linguistic morphologists and computationalists here as to the data to be explained by a science of language is irrelevant to the inescapability of morphemes within any approach to language.

THE FUTURE OF MODELLING, WITH MORPHEMES AND SYNTAX

Computational modelling of the acquisition and use of morphological knowledge should prove essential and exciting to linguists and cognitive neuroscientists investigating the structure of language and its use by speakers and hearers. There are many open questions in linguistic morphology that might be clarified via modelling and the necessary "theories of the task" for cognitive neuroscience experiments (e.g., what is a participant doing when making a lexical decision) could usefully depend on modelling of the participants' knowledge of language. Modelling is required as we move from making qualitative to quantitative predictions about experimental data. From the standpoint of contemporary morphological theory, any approach to modelling word structure would involve modelling the acquisition of grammatical knowledge, a point also emphasised by Baayen et al. (2011) and supporters of the triangle model.

Baayen et al. (2011) explain the value of taking very seriously the observation that words are not learned in isolation; the linguistic context of words/morphemes in learning is crucial to their representation and processing. For the linguist, the "structured list" of morphemes that constitutes the syntactic representation of an utterance is the proper locus of context in learning and processing; that is, the relevant context is (local) syntactic context. Syntax-based processing models describe the ways in which already processed input builds expectations for following morphemes and structures and the ways in which newly encountered input forces updating of representations and expectations (see Roark, Bachrach, Cardenas, & Pallier, 2009). Baayen et al. (2011) as well as my own work on morphological processing (e.g., Solomyak & Marantz, 2010) suggest that this syntax-centred approach to language processing needs to be extended equally to the within-word (between morpheme) and between-word contexts.

Standard computational approaches explicitly treat time (for auditory processing; left–right or right–left spatial dimension for reading) as a crucial dimension for the syntax, with the notion of predicting upcoming morphemes not reduced to the prediction of upcoming form or upcoming meaning. The models discussed in "No escape from morphemes" reduce the temporal (or orthographic spatial) dimensions to the form level of representations, with any temporal organisation of other levels parasitic on the connection to the form level. This reduction is radical indeed, and, as far as I can tell, completely without support from the modelling or experimental literature.

I believe that the linguistic world would be interested in experiments (in both the modelling and human participant senses) that explicitly test amorphous and a-syntactic representations of language against the standard syntactic models, where the notions like entropy over possible representations and surprisal at input based on the probability distribution over expected input are strictly based on the form dimension (i.e., temporal sequence for auditory processing), and the organisation of semantic and other information in sentences is tied only to this form dimension. So, what linguists attribute to the syntagmatic and paradigmatic dimensions of linguistic knowledge would need to be explained via learned connections between the temporally ordered form dimension and everything else (the "semantics"). Other than the linguist's assumption—that people learned grammars of morphemes—what architectural or other features could one add to the amorphous, a-syntactic models to approach accounts of what the linguist explains with a grammar? As I have indicated in the section "No escape from morphemes", no computational approach actually asks this question, with current approaches building syntactic knowledge in the linguist's sense into the semantic representations, which are not learned.

From the linguist's standpoint, the most promising way to proceed computationally would be to model the acquisition of morphemes and syntax—i.e., the acquisition of the grammar of a language. The approaches mentioned above—the Baayen et al. (2011) and the triangle models—that presuppose that acquisition of words begins with the syntactic morphological structure as input (within the semantic input) is not likely to shed light on how the child acquires knowledge of morphology. Existing computational approaches to learning grammar could usefully be extended to learning the grammar internal to words (see Perfors, Tenenbaum, & Regier, 2011 on Bayesian grammar induction).

It might be possible to pit a model in which there is truly no morphological layer and no syntax of morphemes against a model in which learning morphology is learning this syntax. As explained above, neither of these models, at this level of abstraction, is a priori simpler; neither is cut from the other via Occam's razor. The long history of linguistic research suggests there should be utility in learning the grammar, since we have every reason to believe that language is characterised by a word-internal syntax and every reason from experiments to believe that this syntax is internalised by speakers and used in language comprehension and production. Morphemes are units with interpretations in form and in meaning and whose distribution, both co-occurrence and complementarity, is described by a syntax. The understanding that language involves a structure of morphemes is so obvious that it is perhaps easy to overlook that one is presupposing this structure to language at the same time that one is denying a central role for abstract morphemes in linguistic knowledge.

REFERENCES

Anderson, S. R. (1992). *A-morphous morphology*. Cambridge, UK: Cambridge University Press.

Aronoff, M. (1994). *Morphology by itself: Stems and inflectional classes*. Cambridge, MA: MIT Press.

Baayen, R. H., Milin, P., Durdevic, D. F., Hendrix, P., & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review*, *118*(3), 438–481. doi:10.1037/a0023851

Beard, R. (1995). *Lexeme-morpheme base morphology: A general theory of inflection and word formation*. Albany: State University of New York Press.

Embick, D. (2010). *Localism versus globalism in morphology and phonology*. Cambridge, MA: MIT Press.

Embick, D., & Marantz, A. (2008). Architecture and blocking. *Linguistic Inquiry*, *39*(1), 1–53. doi:10.1162/ling.2008.39.1.1

Halle, M., & Marantz, A. (1993). Distributed morphology and the pieces of inflection. In K. Hale & S. J. Keyser (Eds.), *The view from building 20* (pp. 111–176). Cambridge, MA: MIT Press.

Halle, M., & Marantz, A. (1994). Some key features of distributed morphology. *MIT Working Papers in Linguistics*, *21*, 275–288.

Harley, H., & Noyer, R. (1999). Distributed morphology. *Glot International*, *4*(4), 3–9.

Joanisse, M. F., & Seidenberg, M. S. (1999). Impairments in verb morphology after brain injury: A connectionist model. *Proceedings of the National Academy of Sciences*, *96*(13), 7592–7597. doi:10.1073/pnas.96.13.7592

Lieber, R. (1992). *Deconstructing morphology: Word formation in syntactic theory*. Chicago, IL: University of Chicago Press.

Marantz, A. (2005). Generative linguistics within the cognitive neuroscience of language. *The Linguistic Review*, *22*(2–4), 429–445. doi:10.1515/tlir.2005.22.2-4.429

Matthews, P. H. (1965). Some concepts in word-and-paradigm morphology. *Foundations of Language*, *1*(4), 268–289. http://www.jstor.org/stable/25000186

Perfors, A., Tenenbaum, J. B., & Regier, T. (2011). The learnability of abstract syntactic principles. *Cognition*, *118*(3), 306–338. doi:10.1016/j.cognition.2010.11.001

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*(1), 56–115. doi:10.1037/0033-295X.103.1.56

Rice, K. (2000). *Morpheme order and semantic scope: Word formation in the Athapaskan verb*. Cambridge, UK: Cambridge University Press.

Richards, N. (2007). Lardil 'case stacking' and the structural/inherent case distinction. Retrieved from http://ling.auf.net/lingBuzz/000405

Roark, B., Bachrach, A., Cardenas, C., & Pallier, C. (2009). Deriving lexical and syntactic expectation-based measures for psycholinguistic modeling via incremental top-down parsing. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing* (pp. 324–333). Association for Computational Linguistics.

Seidenberg, M.S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*(4), 523–568. doi:10.1037/0033-295X.96.4.523

Solomyak, O., & Marantz, A. (2010). Evidence for early morphological decomposition in visual word recognition. *Journal of Cognitive Neuroscience*, *22*(9), 2042–2057. doi:10.1162/jocn.2009.21296