PHILOSOPHICAL TRANSACTIONS B

royalsocietypublishing.org/journal/rstb

Research

Check for updates

**Cite this article:** Gwilliams L. 2019 How the brain composes morphemes into meaning. *Phil. Trans. R. Soc. B* **375**: 20190311.
http://dx.doi.org/10.1098/rstb.2019.0311

Accepted: 7 August 2019

One contribution of 16 to a theme issue 'Towards mechanistic models of meaning composition'.

**Subject Areas:**
cognition, neuroscience

**Keywords:**
morpho-syntax, semantic composition, neurolinguistics, natural language processing

**Author for correspondence:**
Laura Gwilliams
e-mail: leg5@nyu.edu

THE ROYAL SOCIETY PUBLISHING

# How the brain composes morphemes into meaning

Laura Gwilliams

Psychology Department, New York University, New York, NY 10003, USA

LG, 0000-0002-9213-588X

Morphemes (e.g. [tune], [-ful], [-ly]) are the basic blocks with which complex meaning is built. Here, I explore the critical role that morpho-syntactic rules play in forming the meaning of morphologically complex words, from two primary standpoints: (i) how semantically rich stem morphemes (e.g. explode, bake, post) combine with syntactic operators (e.g. -ion, -er, -age) to output a semantically predictable result; (ii) how this process can be understood in terms of mathematical operations, easily allowing the brain to generate representations of novel morphemes and comprehend novel words. With these ideas in mind, I offer a model of morphological processing that incorporates semantic and morpho-syntactic operations in service to meaning composition, and discuss how such a model could be implemented in the human brain.

This article is part of the theme issue 'Towards mechanistic models of meaning composition'.

## 1. Introduction

That you are understanding the words on this page; that you can understand me still, when I talk to you in a noisy pub or over the telephone; that you and I are able to use language to communicate at all, usually effortlessly, exemplifies one of the most critical cognitive faculties belonging to human beings. Here, I focus on a specific aspect of this process, namely how the brain derives the meaning of a word from a sequence of morphemes (e.g. [dis][appear][ed]).[1]

A morpheme is defined as the smallest linguistic unit that can bear meaning. The kind of meaning that it encodes depends on what type of morpheme it is. For instance, *lexical morphemes* primarily encode semantic information (e.g. [house], [dog], [appear]); *functional morphemes* primarily encode grammatical or morpho-syntactic information (e.g. [-s], [-ion], [dis-]), such as tense, number and word class. In English, these usually map to root and affix units, respectively, though this differs considerably cross-linguistically. Each morpheme is an atomic element, groups of which are combined in order to form morphologically complex words. For example, to express the process *appear* in the past, one can combine the stem morpheme with the inflectional suffix *-ed* to create *appeared*; to convey the opposite, add a negating prefix: *disappeared*.

In what follows, I will first overview what I consider to be the main neural processing stages recruited during morphological processing. In many senses, the selection and description of these stages build from previous models of behavioural data, such as [10], updated to incorporate results from neurolinguistics and natural language processing (NLP). For each processing stage, I will review the relevant literature regarding both language comprehension and production.

As should become clear from this review, there are a number of aspects of morphological processing that remain highly under-specified. The goal of the second part of this paper, therefore, is to put forward a composite model of morphological processing, with the aim of offering directions to guide future research. I am as explicit as possible regarding the semantic and morpho-syntactic features at play, the transformations applied at each stage, and where in the brain these processes may be happening. In this sense, then, the discussion will focus on the representational and algorithmic level of analysis, as defined by David Marr [11].

royalsocietypublishing.org/journal/rstb *Phil. Trans. R. Soc. B* **375**: 20190311

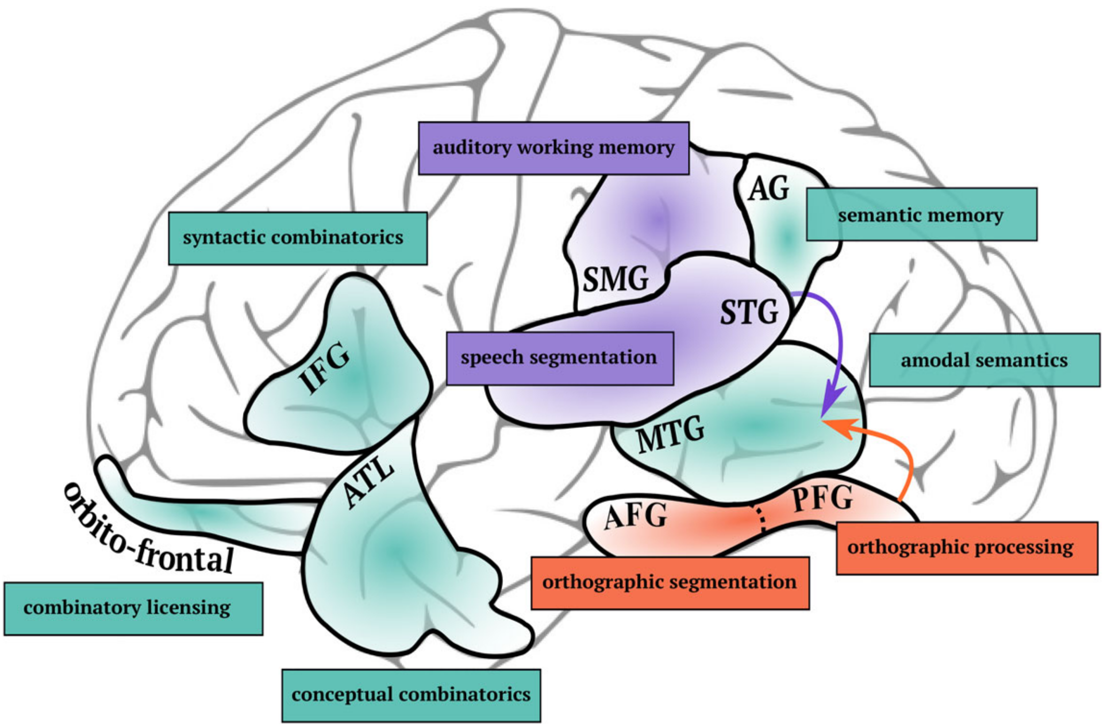


**Figure 1.** Putative brain regions associated with different stages of morphological processing. Orange colour refers to modality-specific written word processes. Purple colour to modality-specific spoken word processes. Turquoise refers to a-modal processes. PFG, posterior fusiform gyrus; AFG, anterior fusiform gyrus; STG, superior temporal gyrus; SMG, supramarginal gyrus; AG, angular gyrus; MTG, middle temporal gyrus; ATL, anterior temporal lobe; IFG, inferior frontal gyrus.

## 2. Overview of processing stages

In the case of language comprehension, the job of the language listener is to undo the work of the speaker, and reconstruct the intended meaning from the produced expression—to understand the concept *disappeared* rather than to articulate it [12–14]. I propose that, to achieve this, the following processing stages are involved:

- (i) **Segmentation**. Identify which morphemic units are present.
- (ii) **Look-up**. Connect each of those units with a set of semantic and/or syntactic features.
- (iii) **Composition**. Following the morpho-syntactic rules of the language, combine those features to form a complex representation.
- (iv) **Update**. Based on the success of the *sentence* structure, adjust the morphemic representations and combinatory rules for next time.

Note that these operations need not unfold under a strictly serial sequence, such that the previous stage completes before the next is initiated. Rather, based on previous work, it is likely that operations unfold under a more cascaded architecture, such that many computations occur in parallel (e.g. see [15]). The rest of this section will review the neural evidence in favour of these stages.

### (a) Morphological segmentation: identifying the building blocks

One of the earliest neural processes is morphological segmentation. The goal is to locate the morphological constituents (roots, derivational and inflectional affixes—defined fully in §§2b(i–iii), respectively, below) within the written or spoken input, and link them to a modality-specific representation (sometimes referred to as a form-based 'lexeme'[2] [16,17]).

Evidence for morphological segmentation comes from both written and spoken language processing. Putative anatomical locations for these processes are presented in figure 1.

#### (i) Written word processing

During reading, it appears that the brain segments written words into morphemes based on an automatic morpho-orthographic parser ([18–21], among others). Whenever both a valid root (either free or bound) and a valid suffix are present in the input, the parser is recruited (e.g. *farm-er*, *post-age*, *explode-ion*) [22]. This is true even for words like *vulner-able*, *excurs-ion*, whose stems never occur in any other stem–affix combination [23]. At this stage, the system is not yet sensitive to the semantic relatedness between the stem and the complex form. This has been shown to lead to false parses of mono-morphemic words like *corn-er* and *broth-er*; the parser is not fooled, however, when a stem is present without a valid affix (e.g. *broth-el* [21]). Overall, this suggests that the system segments input based on the semantically blind identification of morphemes that contain an entry in the lexicon.

Visual morpheme decomposition has been associated with activity in the fusiform gyrus using fMRI [24], overlapping with the putative visual word form area [25,26]. Corroborating evidence from magneto-encephalography (MEG) has also associated this area with morphological segmentation: responses in posterior fusiform gyrus (PFG) around 130 ms after visual presentation are modulated by bi-gram and orthographic affix frequency [27–29]. This is consistent with research focused on orthographic processing, which associates this area with the identification of recurring substrings [30–32]. Slightly more anterior along the fusiform,

**Table 1.** Summary of experimental variables that have been used to tap into different stages of morphological processing. Here, ($a$, $b$), letters of the input; ($X$), stem; ($Y$), affix; ($Z$), root; ($W$), morphologically complex word; ($C$), cohort of morphologically related words; ($E$), expected value.

| feature | formula | process, timing | example study |
|---|---|---|---|
| bigram frequency | $\log \sum (a, b)$ | orthographic, 130 ms | Simon *et al.* [29] |
| stem frequency | $\log \sum (X)$ | segmentation, 170 ms | Gwilliams *et al.* [27] |
| affix frequency | $\log \sum (Y)$ | segmentation, 170 ms | Solomyak and Marantz [35] |
| transition probability | $P(Y\|X)$ | segmentation, 170 ms | Gwilliams and Marantz [23] |
| root frequency | $\log \sum (Z)$ | lexical access, 350 ms | Lewis *et al.* [34] |
| family entropy | $-\sum (P(W\|C) \log_2 P(W\|C)$ | lexical access, 350 ms | del Prado Martin *et al.* [36] |
| semantic coherence | $E(\text{freq}(W)) - \text{freq}(X + Y)$ | composition, 400–500 ms | Fruchter & Marantz [33] |

responses around 170 ms are modulated by morpheme-specific properties such as stem frequency, affix frequency and the frequency with which those units combine (i.e. *transition probability*) [23,33–35]. Anterior fusiform gyrus is therefore associated with decomposing written input into morphemes.

Table 1 summarizes these different metrics of orthographic and morphological structure, and specifies the cognitive process that each putatively taps into.

#### (ii) Spoken word processing

Less research has been conducted on morphological segmentation of spoken language. Work on speech segmentation, in general, suggests that words and syllables are identified using statistical phonotactic regularities [37], acoustic cues such as co-articulation and word stress [38,39] and lexical information [40]. Similar kinds of statistical cues appear to be used for morpheme boundaries as well. In particular, there is sensitivity to the transition probability between spoken morphemes [41,42] in superior temporal gyrus (STG) at around 200 ms after phoneme onset. These moments of low transition probability may be used as boundary cues to bind phonological sequences into morphological constituents.

For both auditory and visual input, then, the system applies a morphological parser on the input, which segments the signal based on sensory (pitch, intensity, bigrams), form-based (identification of affix or stem string) and statistical information (transition probability).

### (b) Lexical access: figuring out what the blocks mean

Identifying the meaning of the segmented morphemes is often referred to as 'lexical access'. This stage involves linking the form-based morpheme to the relevant bundle of semantic and syntactic features. Each word consists of at least three pieces: the root, (any number of) derivation(s) and inflection, even if one of the pieces is not spelled out in the written or spoken language [43]. Depending on the type of morpheme being processed, the features are different.

#### (i) Root access

Root morphemes are the smallest atomic elements that carry semantic properties; units like: *house*, *dog*, *walk*, *love*. Based on previous theoretical work (e.g. [44]), the root is assumed not yet to be specified for its word class; so, the use of *love* as a verb (*to love*) and a noun (*the love*) contains the same root morpheme.

The middle temporal gyrus (MTG) has been implicated in semantic lexical access in a number of processing models [45–47], along with the superior temporal sulcus [48]. The angular gyrus has also been implicated in semantic memory more broadly (see [49–51] for reviews of anatomical locations associated with semantic processing).

A particular response component found in MEG, whose neural source originates from MTG at around 350 ms after word onset, has been associated specifically with access to the decomposed *root* of the whole word [52–54]. This has been corroborated by the finding that neural activity in this area, at this latency, is modulated by lemma frequency [35], polysemy [54] and morphological family entropy [33,53]—perhaps reflecting competition between the recognition of different roots.

#### (ii) Derivation access

Derivational morphology refers to a constituent (e.g. *-ion*, *-ness*, *-ly*) that creates a new lexeme from that to which it attaches. It typically does this by changing part of speech (e.g. employ → employment), by adding substantial non-grammatical meaning (e.g. child → childhood; friend → friendship), or both [55].

There are data from cross-modal priming studies indicating that derivational suffixes can be primed from one word to another (darkNESS—happiNESS) [56]. This suggests that (i) there is an amodal representation (i.e. not bound to the visual or auditory sensory modalities) that can be accessed and therefore primed during comprehension; (ii) the representation is somewhat stable in order to generalize across lexical contexts. Furthermore, findings from fMRI link the processing of derived forms with activity in the left inferior frontal gyrus (LIFG) [57]—i.e. Broca's area—which is traditionally associated with syntactic processing, broadly construed. This region has also been associated with the processing of verbal argument structure [58], further implicating the LIFG in derivational morphology, though, overall, the processing of derivation is not as clear as it is for roots [59].

#### (iii) Inflection access

Inflectional morphology invokes no change in word class or semantic features of the stem. Instead, it specifies the grammatical characteristics that are obliged by the given syntactic category [55]. In English, inflectional morphemes would be units like *-s*, *-ed*, *-ing*.

Similar to derivational morphology, and with more empirical support, inflectional morpho-syntactic properties appear to be processed in the LIFG [60,61]. This area is recruited for both overt and covert morphology (i.e. inflections that are realized with a suffix (ten lamb + s) and those that are silent (ten sheep + ∅) [62]), which suggests that the same processing mechanisms are recruited even when the morphology is not realized phonetically or orthographically.

### (c) Morphological combination: putting the blocks together

In order to comprehend the meaning of the complex item, the system needs to put together the semantic and syntactic content of the constituents. This is referred to as the *composition* stage of processing.

The majority of work on morphological composition has employed quite coarse categorical distinctions between semantically valid and invalid combinations. For example, morphologically valid combinations like *farm* + *er* elicit a stronger EEG response at around 400–500 ms when compared to invalid combinations like *corn* + *er* [19,63–66]. MEG work has associated this with activity in the orbito-frontal cortex [33,67–69].

In a more fine-grained comparison, Fruchter & Marantz [33] tested just morphologically valid combinations, and found that the extent to which the meaning of the whole word is expected given its parts (termed 'semantic coherence') also drives orbito-frontal activity at around 400 ms. This has been interpreted as reflecting a stage that assesses the compatibility between the composed complex representation and the predicted representation of the whole word given the parts of the word. Broadly though, the mechanism by which morphemes are combined is extremely under-specified, and is a rich avenue for future study.

## 3. Composite model of morphological processing

The discussion so far has reviewed the literature regarding three stages of morphological processing. While neurobiological research provides quite a comprehensive explanation of how the brain segments sensory input into morphological constituents, our understanding remains poorly defined in terms of (i) what linguistic features make up the representation of each morphological unit; (ii) what operations are applied to those features at each stage of composition. Therefore, the rest of this article is dedicated to putting forward a framework that is explicit on both of these points, in order to guide future studies. It is informed by the neural data reported above, as well as research from linguistics and NLP.

### (a) Morpheme segmentation

It is interesting to note that the neurophsyiological findings regarding morphological decomposition are echoed in the engineering solutions developed in NLP. Some tools employ a storage 'dictionary' of morphological constituents that are compared to the input to derive units from the speech or text stream [70]. This is similar to the stem + affix look-up approach of Taft [22]. Other morphological segmentation tools such as *Morfessor* work by picking up on statistical regularities in the input and maximizing the likelihood of the parse in a unsupervised manner [71]. This is related to sensitivity to grapheme and phoneme transition probability as attested in the PFG and STG, respectively. Overall, both types of NLP segmentation—dictionary lookup and statistical regularity—are attested in the cognitive neuroscience literature as methods the brain uses for segmentation.

### (b) Morpheme representations

The framework presented here treats morphemes (and words) as a set of semantic and syntactic features. Figure 2 visualizes a collection of some example features for four words.

Computationally, this means that each morpheme is represented as a list of numbers. Each slot in the list corresponds to a particular feature (e.g. *cute*, *loud*, *plural*) and the number associated with that slot reflects how relevant a feature is for that morpheme. The list of numbers that represent each morpheme, or each word, will be referred to as a *vector* from now on.

Precisely how many semantic dimensions define a word is likely impossible to answer. But, I propose that all words are defined relative to the same set of semantic features, and the feature-slot correspondence is the same across all words. This systematicity between the index of the vector, and the meaning of that dimension, is what allows the brain to keep track of what elements contain what information, and apply the appropriate computation.

In terms of neural implementation, each dimension of the vector could be realized by a neuron or neuronal population that codes for a particular feature. The vector format has been chosen for modelling purposes because it supports basic mathematical operations, but the vector format itself is not critical for how these processes work in the brain.

#### (i) Root features

The root morpheme is proposed to consist only of semantic properties. The root in *dogs*, for example, is *dog*, which may have a feature for *brown*, *fluffy*, *cute*, *small*, *mammal* and *barks*. The root itself is assumed to not be associated with any particular word class—this is only determined once it is combined with derivational morphology, as explained below. So, in its acategorical form, I hypothesize that the weights for each feature correspond to the average weight given all the contexts of use of the word.

In NLP, the meaning of words are routinely represented as *word embeddings* (e.g. [5–8]). These are vectors created from the statistical co-occurrence of words, capitalizing on the fact that words that mean something similar often occur in similar contexts.[3] They provide a powerful way of representing meaning and obey simple geometric transformations (for example, subtracting $\vec{\textit{\textbf{man}}}$ from $\vec{\textit{\textbf{king}}}$ and adding $\vec{\textit{\textbf{woman}}}$ results in a vector that closely approximates $\vec{\textit{\textbf{queen}}}$ [74]). These vectors are typically created to contain around 50–300 features, where each feature is extracted using unsupervised techniques such as principal component analysis. As a consequence of this data-driven approach, it is often not possible to interpret what each dimension of the vector actually means.

I want to highlight that even though this is a highly successful method for solving language engineering problems, I am not proposing that this is how the brain acquires dimensions of understanding. Rather, I believe that the dimensions of

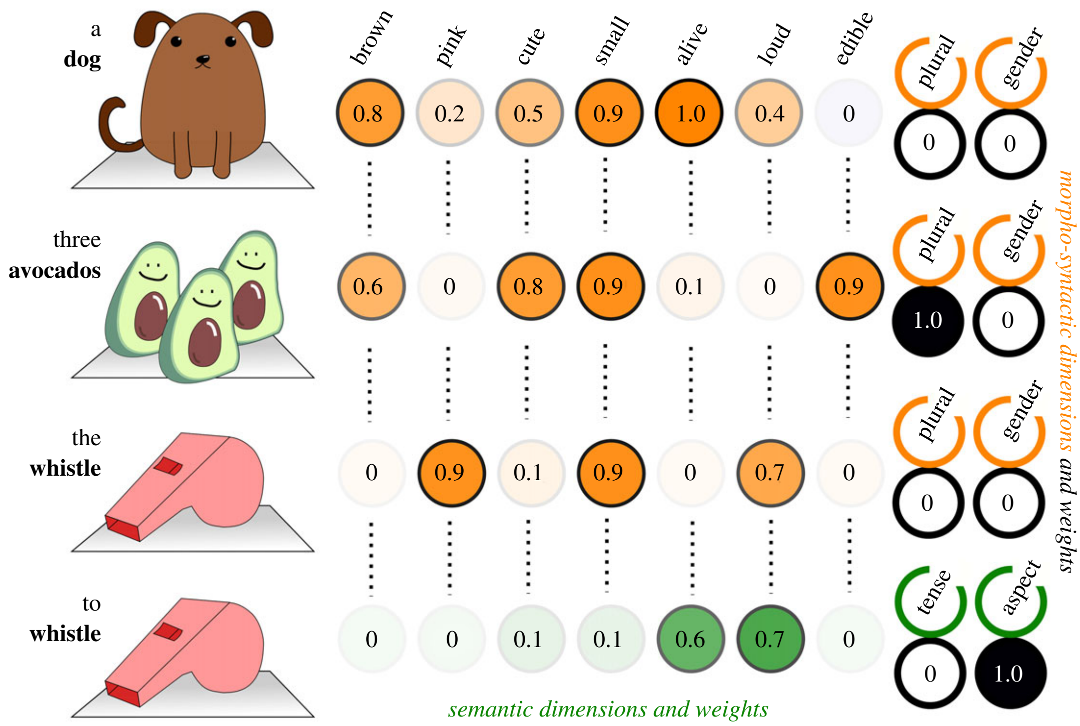


**Figure 2.** Representing words as vectors. For the three morphologically composed words *dog, avocados, whistle*, the semantic and syntactic features are displayed. The first seven features refer to the semantics of the item; the last two features refer to syntactic properties. The labelled disks correspond to the morpho-syntactic dimensions that are relevant to the word, as primarily provided by the derivation; the weights of those dimensions are specified by the inflection, as shown in greyscale. Each word is expressed under the same set of semantic dimensions, but the syntactic dimensions differ depending on the word class of the item. The orange colour denotes that this item is a noun; the green colour, a verb.

word embeddings, as learnt through corpus statistics, are conveniently correlated with the 'true' dimensions of meaning used by the brain, thus leading to their correlation with neural activity [75]. Determining what those 'true' features are, though, will likely require more manual inspection and human intuition (e.g. along the lines of [51]).

### (ii) Derivation features

Unlike the potentially infinite number of semantic features, syntactic features are of a closed finite set. Given the significant stability of word classes and morpho-syntactic features cross-linguistically, I propose that the derivation contains a place-holder for all possible morpho-syntactic properties of the particular syntactic category [76]. This is very much in line with *categorical structure* as proposed by Lieber [77]. For example, the morpheme *-ful*, which derives an adjective from a noun, would contain adjectival morpho-syntactic features such as *function*, *complementation*, *gradation*. The suffix *-ion* contains nominal features such as *case*, *number*, *gender*. The suffix *-ify* contains verbal features such as *tense*, *aspect*, *mood*. So, even though English does not mark grammatical gender, a native English speaker would still have a slot for this morpho-syntactic feature in their representation of nouns. The index location of these morpho-syntactic features would also always be stable within the representation of the whole word (so, which entry in the vector, for the computational implementation; which feature-sensitive neuron(s), for the neural implementation). This way, the system knows from where to extract the derivational dimensions, and on what features to apply the relevant compositional functions.

Critically, the derivation only serves to specify which morpho-syntactic features are of potential relevance. It does not actually contain the weights for each dimension—that is the job of the inflection, as expanded upon below. In this way, then, the derivation acts as a kind of coordinate frame: it determines the word class of the whole word by specifying the relevant syntactic dimensions within which it should be expressed. In figure 3, the coordinate frame is depicted by the colour of the vector.

I propose that the derivation is also specified in terms of the same semantic dimensions that define the root morpheme. For example, this semantic feature would be shared between child*hood*, man*hood* and woman*hood*; and between charm*er*, kiss*er* and walk*er*. Furthermore, there may also be semantic similarities within word classes more generally, either expressed as an explicit feature, or as an emergent property of occupying similar syntactic roles. For example, the semantic noun-y-ness associated with *mountain* may be shared with the noun-y-ness of *petrol*.

### (iii) Inflection features

I propose that the inflectional morpheme serves to specify the value for each of the morpho-syntactic dimensions identified by the derivation. For example, if the derivation recognizes *number* as a relevant dimension for the stem, it is the inflection that specifies whether the word is singular or plural. If the derivation specifies a feature that is not applicable to the word being processed, such as *gender* for the English word *table*, then the inflection simply allocates a zero weight. This is consistent with Levelt *et al.*'s [81] lexical access model of speech production.

## (c) Compositional computations

Now I move to a critical aspect of morphological processing, which is how the morphemes are combined in order to form a complex meaning. The basic proposal here is that the three

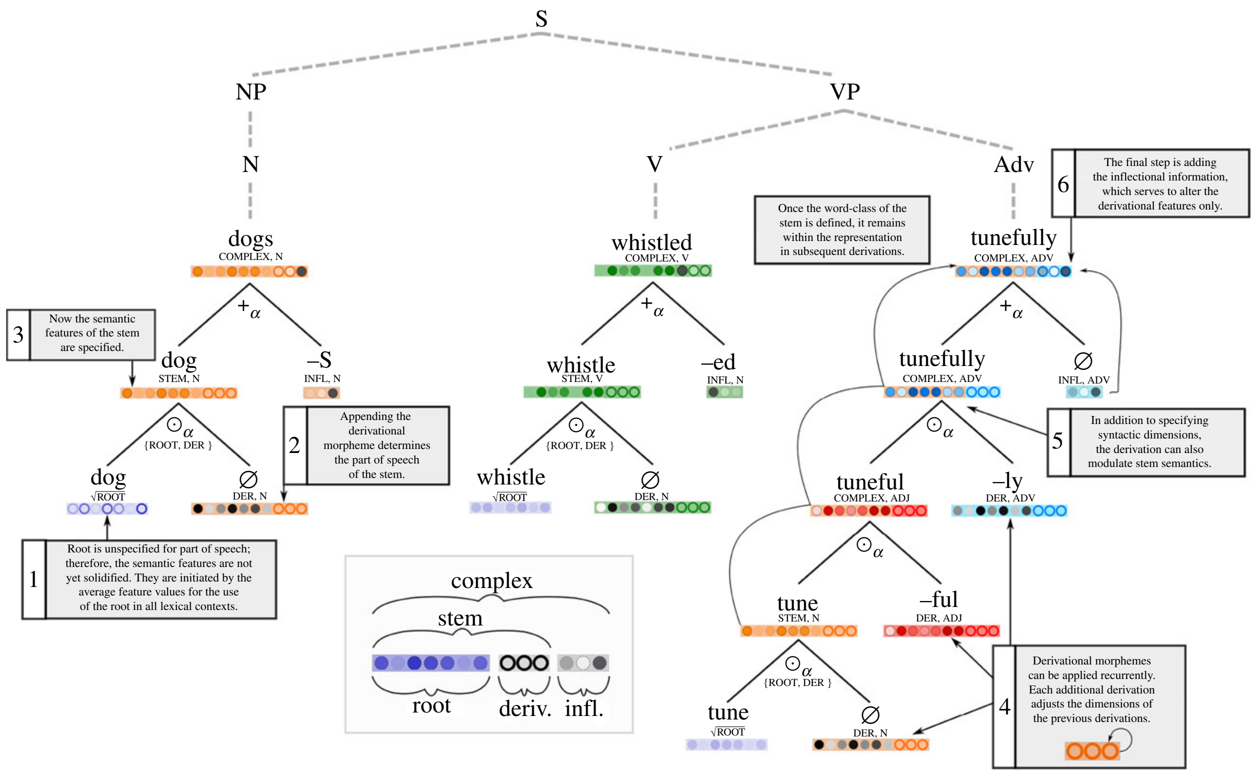


**Figure 3.** Morpho-syntactic composition, detailing the morphological composition recruited during the processing of the sentence *dogs whistled tunefully.* The colour specifies the word class of the unit (purple, acategorical; orange, nominal; green, verbal; red, adjectival; blue, adverbial). The index of the feature in the vector is what determines its semantic/syntactic function. $\varnothing$, a morpheme with no phonological or orthographic realization; {ROOT, DER}, concatenation of vectors; $\odot$, element-wise multiplication; +, addition; $\alpha$, learnt combinatorial weight. This allows the modulation of the semantic features to be adjusted with experience with a particular language. I am assuming, in line with linguistic theories [78–80], that the root is unspecified for word class (e.g. adjective, noun, verb). Adj, adjective; Adv, adverb; Der, derivation; Infl, inflection; N, noun; NP, noun phrase; S, sentence; V, verb; VP, verb phrase.

**Table 2.** Mathematical example of the derivational processes leading to the formation of the complex word *whistler*. $\odot$ refers to element-wise multiplication between vectors. + refers to vector addition. Top part of the table corresponds to semantic dimensions; bottom part to syntactic dimensions. The semantic dimension labels were taken from [51].

| | whistle (root) | | Ø (verbal) | | whistle (stem) | | -er (nominal) | | whistler (complex noun) | | -s (inflection) | | whistlers (inflected) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| sound | 0.9 | | 1.0 | | 0.9 | | 1.0 | | 0.9 | | Ø | | 0.9 |
| social | 0.3 | | 0.8 | | 0.24 | | 5.0 | | 1.2 | | Ø | | 1.2 |
| shape | 0.5 | $\odot$ | 0.2 | = | 0.1 | $\odot$ | 0.5 | = | 0.05 | + | Ø | = | 0.05 |
| happy | 0.5 | | 2.0 | | 1.0 | | 1.0 | | 1.0 | | Ø | | 1.0 |
| human | 0.2 | | 0.3 | | 0.06 | | 10.0 | | 0.6 | | Ø | | 0.6 |
| | Ø | | tense | | tense | | plural | | plural | | 1 | | 1 |
| | Ø | | aspect | | aspect | | gender | | gender | | 0 | | 0 |

types of morphemes obey three types of combinatoric operations, which unfold in a particular order, and with predictable consequences for the semantic and syntactic outcome of the word. Below, each combinatorial stage of processing is explained in turn.

### (i) Concatenation and multiplication to create the stem morpheme

The first operation involves combining the semantics of the root morpheme (purple vector in figure 3, step 1) with morpho-syntactic dimensions of the initial derivation (step 2). This forms the stem morpheme, which is specified for its word class (step 3) [44].

Computationally, I propose that this involves two steps. One is appending the syntactic dimensions, which are of potential relevance given the word-class of the derivation, to the representation of the root morpheme. Second is modulating the semantic features of the root, relative to the word-class of the derivation. This second procedure can be achieved through element-wise multiplication, such that the derivation systematically increases the weight of certain dimensions and decreases the weight of others. This forms a stem morpheme, whose semantic and syntactic properties are defined relative to its established word class. An explicit example of this process is shown in table 2.

### (ii) Multiplication at each derivation

After this initial stem-formation stage, any number of additional derivational operations can then be applied (step 4 in figure 3). Each additional transformation involves adjusting which morpho-syntactic dimensions are relevant, and at the same time, modulating the semantic dimensions in line with the syntactic category. For example, if transforming a noun into a verb, the relevant syntactic dimensions change from *number* and *gender* to *tense* and *aspect*. Furthermore, the semantic dimensions change from highlighting visual aspects of the concept to kinetic aspects of the concept.

In line with this idea, behavioural studies have shown that listeners are sensitive to the word class of the stem, even when the word as a whole ends up being a different syntactic category (e.g. the verb *explode* in the nominalization *explosion*) [82]. This suggests that the history of syntactic dimensions is accessible during comprehension.

Critically, I propose that the semantic consequences for the properties of the word are predictable given knowledge about the semantic input and the syntactic operators being applied. In theory, this means that a derivation will always modulate the features of the root in the same way, and this will generalize across lexical contexts. As expanded upon below, this connection between syntax and meaning is precisely what allows for the comprehension of novel words.

### (iii) Addition at the inflection

The final stage involves combining the lexical structure with the inflectional morpheme (step 6 in figure 3). Here, I have denoted the combinatorial operation as simple element-wise vector addition (+) between the morpho-syntactic features of the derivation (all of which are zero) and those of the inflection (non-zero). In this way, the inflection works to specify the weights of the derivational suffix, making explicit which morpho-syntactic properties are relevant and to what extent. This is also depicted on the right side of figure 2.

This idea of vector manipulation, in service to meaning composition, has enjoyed success in previous NLP research. Studies have used morpheme vector representations in a broad sense, though not strictly coding for semantic versus morpho-syntactic properties of the units [5–8,83]. Even when using a simple composition rule such as addition or element-wise multiplication between morpheme vectors [84], these composed vectors reasonably approximate semantic representations of morphologically complex whole words [85,86]. This suggests that implementing even a very basic composition function could serve to generate complex lexical meaning.

### (iv) Cross-linguistic coverage

As is true for any model of language processing, it is important that it is equally applicable across different languages. Although in English the usual order of morphological units is: root, derivation, inflection, this is not the case for languages with different typologies. For example, some languages make use of infixation: the embedding of an affix *within* a root, rather than before (prefixation) or after (suffixation). Furthermore, Semitic languages, such as Arabic, have discontinuous morphemes. In this case, morphemes are received by the listener in an interleaved fashion, with no neat boundary between one morpheme and the next. For example, in the word *kataba*, the root is expressed as the consonants *k-t-b* and the pattern *-a-a-a* expresses derivational information.

The current proposal predicts that the same basic set of operations (stem formation; derivation; inflection) are always applied, and always in the same order, regardless of the order in which the information is received. From this perspective, knowledge of the language-specific grammar is used to organize the input into an appropriate syntactic structure. The input is then processed under the language-general architecture described here. This leads to the strong prediction that, regardless of the language being processed, the neural signatures of these processes should always occur in the same order.

## (d) Feedback from the sentence structure

Once the full sentence structure has been created, the sentence can similarly be represented in terms of a set of semantic and syntactic features. This process—of generating a semantic–syntactic representation of the sentence—is not discussed here, but I point the reader to [87] for a related proposal of how this may be achieved. In abstract, the idea is that the sentence representation is built using similar compositional rules to those currently described, as well as incorporating additional non-structural meaning from the broader situational and pragmatic context.

The final part of the current framework describes how the system can use this sentence structure to inform morphological structure in two ways: (i) *update* or strengthen the constituent representations; (ii) *create* representations of novel morphemes.

### (i) Update

Once the sentence is built, its semantic content and syntactic structure can be used to compute the most likely representations of morphological constituents. One possible mathematical implementation of this would be Bayesian inference, where the most likely representation of morpheme$_j$ is inferred based on the representation of the sentence as computed *without* morpheme$_j$:

$$P(\boldsymbol{morpheme}_j|\boldsymbol{sentence}_{-j}) \propto P(\boldsymbol{sentence}_{-j}|\boldsymbol{morpheme}_j) \times P(\boldsymbol{morpheme}_j). \tag{3.1}$$

The discrepancy between the inferred morpheme (the posterior) and the constituent representation (the prior), can be thought of as the representational error:

$$\mathbf{error} = P(\boldsymbol{morpheme}_j) - P(\boldsymbol{morpheme}_j|\boldsymbol{sentence}_{-j}). \tag{3.2}$$

If the sentence structure has low semantic or syntactic validity, the error will be high, which can be used as a feedback signal in order to update the constituent representation. This will improve validity for next time by making the morpheme representation more similar to the inferred posterior representation given the sentence structure, whereas if error is low, the signal will simply strengthen the prior representations that are already in place.

One obvious prediction borne out of this model is that the *update* process would have most influence during language acquisition, because lexical representations are still in the process of being formed. An interesting avenue for further study would be to test how neural correlates

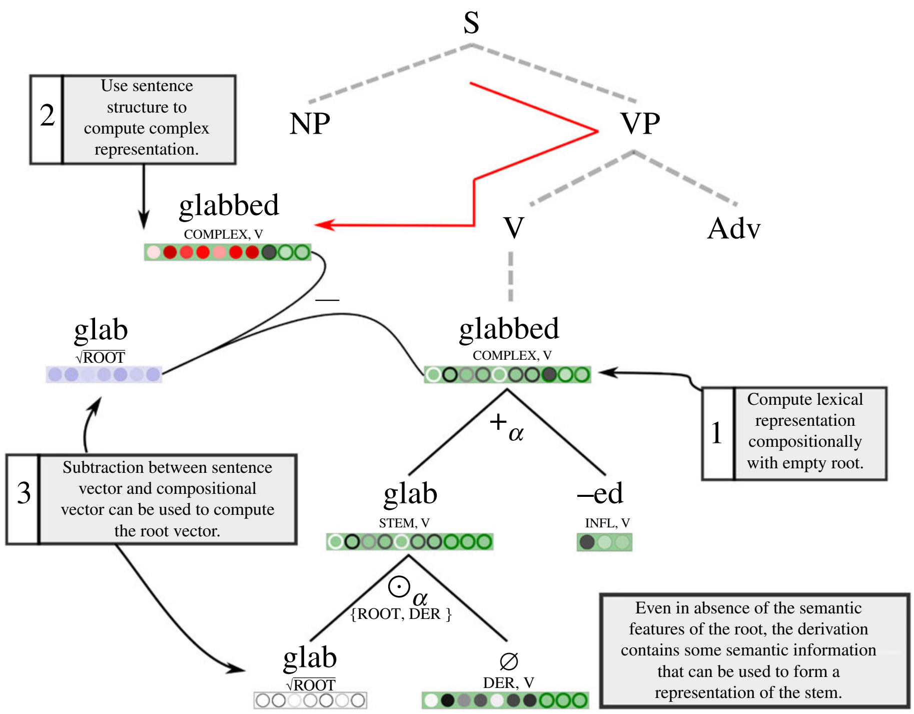


**Figure 4.** Using sentence structure and morphological rules to generate representations. The meaning of the novel word *glabbed* can be estimated using the sentence structure. From there, the meanings of all the atomic morphemes are also generated, by performing simple mathematical operations on the known representations.

of representational updates correspond to behavioural improvements in comprehension and production.

### (ii) Create

If the brain primarily conducts language processing via the atomic morphological units, it is critical that it ensures full coverage over those units. When encountering a morpheme for the first time, the sentence structure can be used to generate the required constituent representations. This idea is quite similar to *syntactic bootstrapping*—using structural syntactic knowledge to understand the semantic content of novel words [88–90]. Here, I propose this can be achieved by applying the following sequence of operations.

First, the compositional representation of the whole word is computed using the steps described in §3c (also depicted in step 1 of figure 4). The derivational and inflectional morphemes are defined in line with the user's knowledge of the language, but because the word has not been encountered before, the root contains null entries. As shown in equation (3.3), the composed meaning of the novel word *glabbed* can be computed as:

$$\boldsymbol{glabbed}_{\text{composition}} = \emptyset \odot \boldsymbol{DER}_V + \boldsymbol{INFL}_V. \tag{3.3}$$

Because the composition is applied using an empty root, the resulting representation only reflects the morpho-syntactic and semantic properties of the affixes.

Second, a contextual representation of the whole word can be estimated from the sentence structure, using the same Bayesian method as described above (and step 2 in figure 4):

$$\boldsymbol{glabbed}_{\text{sentence}} = P(\text{glabbed}|\text{sentence}). \tag{3.4}$$

Third, a subtraction between the compositional meaning of the word as computed with a null root and the interpreted meaning of the word given the sentence can be used to interpolate an atomic representation of the root (step 3 in figure 4):

$$\boldsymbol{glab}_{\text{new}} = \boldsymbol{glabbed}_{\text{sentence}} - \boldsymbol{glabbed}_{\text{composition}}. \tag{3.5}$$

This interpolation process may serve to 'initialize' the representation of a root vector. Then, at each subsequent use, the update function described above (and shown in equations (3.1)–(3.2)) can be used to stabilize the weights of the morphological representation.

Again, while this process is most obviously recruited during language learning, the same mechanism is hypothesized to still be in effect in proficient speakers. Every time a listener is faced with a morphologically complex word, all of the atomic constituents of that word can be computed through an iterative subtraction of affixes. This makes the prediction that the system will hold representations of constituents that are never encountered in isolation: for example, of the root *excurse* from *excursion*. Recent evidence from MEG suggests that this is indeed the case [23].

Filling gaps in the lexicon through either interpolation of constituents or combination into complex forms is precisely the main advantage owed to morphological over word-based representations in NLP. The use of vector representations of morphemes provides better predictive power, especially for out-of-vocabulary words that do not exist in corpora [5–8].

Overall, this highlights that the systematicity between structure and meaning provides a powerful framework for generating missing semantic representations based on the syntactic (both lexical and sentential) situation alone. In this way, it is plausible that the language faculty *in general* is primed to associate syntactic information with particular semantic information. As discussed, this has clear advantages for language acquisition, along the lines of syntactic bootstrapping, as the process of meaning generation would be employed every time a new word is encountered.

Taking this idea a step further, one can imagine that not only the representation is computed, but also a relative activation strength of that representation. Typically, the activation level of a word is thought to be a consequence of how often that word is retrieved from the lexicon: frequent words are accessed more often and therefore have a higher activation level. An alternative explanation is that the system *wants* to make frequently accessed words easier to recognize, and so keeps them activated. From this perspective, activation level would be based on the statistical properties of the word—its orthographic, phonological, morphological structure—rather than (or independently from) frequency of exposure and age of acquisition [91]. If this is true, then the activation level of a novel word could be computed based upon those regularities, and how active a word is may be reflected in the *magnitude* of the corresponding vector representation.

To my knowledge, this has not yet been tested, but it would be easy to do. Although it is a simple distinction, whether word frequency effects are a *consequence* of lexical access or a *engineered processing advantage* has sizeable consequences for the structure of language process and of the mental lexicon more generally.

## 4. Discussion

The goal of this paper is to offer a model of morphological composition that makes explicit (i) what linguistic features make up the representation of a morphological unit; (ii) what operations are applied to those features. While the proposal is based as much as possible on extant literature, it also includes some untested, but testable, educated guesses. There are a number of aspects of morphological processing in the brain that remain highly under-specified; therefore, a fruitful avenue for future research will be to explicitly test the predictions of this model with neural data.

For example, each stage of processing is associated with a particular compositional rule: concatenation and multiplication at the first derivation, multiplication at subsequent derivations and addition at the inflection. Furthermore, the proposal suggests these operations are always performed in the same order, regardless of the language being processed. Whether these are indeed the type and order of neural operations needs to be investigated, perhaps by testing whether a sequence of intermediate representations are encoded in neural activity before arriving at the final complex representation. It would also be informative to correlate the features of both the simple morpheme vectors and the complex word vectors with neural activity to test whether the input/output sequence as tracked by the brain indeed obeys the mathematical operations outlined here.

Although this article has focused on morphological processing, it is possible that these basic principles hold true across multiple units of language. The most obvious analogy is between the syntactic operations used to generate phrasal structures and those used to generate word structures. In line with linguistic theory [78–80], the current proposal makes no meaningful distinction between the two. This is an intuitive idea. For instance, there is very little difference between the composed meaning of *sort of blue* and *blueish*, even though one is made of a phrasal structure and the other a lexical structure. Furthermore, some languages may choose to encode information using multiple words (e.g. *in the house*, English) whereas others may use multiple morphemes (e.g. *extean*, Basque). That one contains orthographic spaces and the other does not is quite arbitrary, and it is not clear whether there are any meaningful processing differences [92], above and beyond things like differences in unit size [93]. Indeed, morphologically rich languages, such as Ojibwe, allow a very large number of morphemes to be combined to form a sentence structure that essentially only consists of a single word. In such cases, it is hard to draw a meaningful distinction between the structure of morphemes to make words, and the structure of words to make sentences.

Overall, this work shows that combining insight from NLP, linguistics and cognitive neuroscience to develop hypotheses for biological systems is potentially very powerful, as each field is essentially tackling the same problem from differing standpoints. For example, in NLP, the goal is to engineer a system to achieve language comprehension with (at least) human-level ability; for neuroscience, the goal is to understand the system that has already done that: the human brain. Consequently, each field has developed tools and insights that (perhaps with a bit of tweaking in implementation or terminology) are mutually beneficial.

## 5. Conclusion

Composition of morphological units provides insight into the infinite potential of meaning expression and the critical systematicity between syntactic structure and semantic consequence. Here, I have briefly reviewed research across cognitive neuroscience, linguistics and NLP in order to put forward a model of morphological processing in the human brain. I hope that this serves as a useful overview and highlights fruitful avenues for further discovery.

Data accessibility. This article has no additional data.

Competing interests. I declare I have no competing interests.

Funding. This work was conducted through support from the Abu Dhabi Institute (grant no. G1001).

Acknowledgements. I especially want to thank Alec Marantz for invaluable discussion and comments on the manuscript. I am also very grateful for extensive and thoughtful feedback from Arianna Zuanazzi, and thank Linnaea Stockall, Samantha Wray and Ryan Law for their very helpful suggestions. Finally, I thank Kiel Christianson and one anonymous reviewer for very insightful comments on earlier versions of this work.

## Endnotes

[1]There is still some contention as to whether morphemes are neurally represented. An alternative possibility is that lexical information gets stored as whole words: the units that would orthographically be flanked by white space [1,2]. However, given (i) the substantial behavioural and neurophysiological evidence that morphemes are in fact represented (for reviews see [3,4]); (ii) the advantage morphological representations provide to speech and text recognition systems [5–8]; and (iii) the need to move the discussion forward, I take for granted that in representing lexical information, the brain does indeed encode morphological units, likely in combination with, but possibly instead of, morphologically complex wholes [9].

[2]A lexeme relates to all inflected forms of a morpheme: play, plays, played, playing would all be grouped under the lexeme *play*.

[3]This is highly consistent with the *Distributional Hypothesis* of semantics—a prevalent usage-based theory of word meaning [72,73].

# References




1. Giraudo H, Grainger J. 2000 Effects of prime word frequency and cumulative root frequency in masked morphological priming. *Lang. Cogn. Process.* **15**, 421–444. (doi:10.1080/01690960050119652)
2. Devlin JT, Jamison HL, Matthews PM, Gonnerman LM. 2004 Morphology and the internal structure of words. *Proc. Natl Acad. Sci. USA* **101**, 14984–14988. (doi:10.1073/pnas.0403766101)
3. Rastle K, Davis MH. 2008 Morphological decomposition based on the analysis of orthography. *Lang. Cogn. Process.* **23**, 942–971. (doi:10.1080/01690960802069730)
4. Amenta S, Crepaldi D. 2012 Morphological processing as we know it: an analytical review of morphological effects in visual word identification. *Front. Psychol.* **3**, 232. (doi:10.3389/fpsyg.2012.00232)
5. Bojanowski P, Grave E, Joulin A, Mikolov T. 2017 Enriching word vectors with subword information. *Trans. Assoc. Comput. Linguist.* **5**, 135–146. (doi:10.1162/tacl_a_00051)
6. Creutz M, Lagus K. 2005 Unsupervised morpheme segmentation and morphology induction from text corpora using Morfessor 1.0. Technical Report A81. Helsinki University of Technology, Finland: Publications in Computer and Information Science.
7. Luong T, Socher R, Manning C. 2013 Better word representations with recursive neural networks for morphology. In *Proc. of the Seventeenth Conf. on Computational Natural Language Learning*, pp. 104–113.
8. Snyder B, Barzilay R. 2008 Unsupervised multilingual learning for morphological segmentation. *Proc. of 46th Annual Meeting of the Association for Computational Linguistics, Columbus, OH, USA, 15–20 June 2008*, pp. 737–745.
9. Marantz A. 2013 No escape from morphemes in morphological processing. *Lang. Cogn. Process.* **28**, 905–916. (doi:10.1080/01690965.2013.779385)
10. Schreuder R, Baayen RH. 1995 Modeling morphological processing. *Morphol. Aspects Lang. Process.* **2**, 257–294.
11. Marr D. 1982 *Vision: a computational approach*. San Francisco, CA: Freeman.
12. Bever TG, Poeppel D. 2010 Analysis by synthesis: a (re-) emerging program of research for language and vision. *Biolinguistics* **4**, 174–200.
13. Halle M, Stevens K. 1962 Speech recognition: a model and a program for research. *IRE Trans. Inf. Theory* **8**, 155–159. (doi:10.1109/TIT.1962.1057686)
14. Poeppel D, Monahan PJ. 2011 Feedforward and feedback in speech perception: revisiting analysis by synthesis. *Lang. Cogn. Process.* **26**, 935–951. (doi:10.1080/01690965.2010.493301)
15. Gwilliams L, Linzen T, Poeppel D, Marantz A. 2018 In spoken word recognition, the future predicts the past. *J. Neurosci.* **38**, 7585–7599. (doi:10.1523/JNEUROSCI.0065-18.2018)
16. Caramazza A. 1997 How many levels of processing are there in lexical access? *Cogn. Neuropsychol.* **14**, 177–208. (doi:10.1080/026432997381664)
17. Laudanna A, Badecker W, Caramazza A. 1992 Processing inflectional and derivational morphology. *J. Memory Lang.* **31**, 333–348. (doi:10.1016/0749-596X(92)90017-R)
18. Crepaldi D, Rastle K, Coltheart M, Nickels L. 2010 'Fell' primes 'fall', but does 'bell' prime 'ball'? Masked priming with irregularly-inflected primes. *J. Memory Lang.* **63**, 83–99. (doi:10.1016/j.jml.2010.03.002)
19. Lavric A, Elchlepp H, Rastle K. 2012 Tracking hierarchical processing in morphological decomposition with brain potentials. *J. Exp. Psychol.: Human Perception Perform.* **38**, 811–816. (doi:10.1037/a0028960)
20. McCormick SF, Rastle K, Davis MH. 2008 Is there a 'fete' in 'fetish'? effects of orthographic opacity on morpho-orthographic segmentation in visual word recognition. *J. Memory Lang.* **58**, 307–326. (doi:10.1016/j.jml.2007.05.006)
21. Rastle K, Davis MH, New B. 2004 The broth in my brother's brothel: morpho-orthographic segmentation in visual word recognition. *Psychon. Bull. Rev.* **11**, 1090–1098. (doi:10.3758/BF03196742)
22. Taft M. 1979 Recognition of affixed words and the word frequency effect. *Memory Cogn.* **7**, 263–272. (doi:10.3758/BF03197599)
23. Gwilliams L, Marantz A. 2018 Morphological representations are extrapolated from morpho-syntactic rules. *Neuropsychologia* **114**, 77–87. (doi:10.1016/j.neuropsychologia.2018.04.015)
24. Gold BT, Rastle K. 2007 Neural correlates of morphological decomposition during visual word recognition. *J. Cogn. Neurosci.* **19**, 1983–1993. (doi:10.1162/jocn.2007.19.12.1983)
25. Cohen L, Lehéricy S, Chochon F, Lemer C, Rivaud S, Dehaene S. 2002 Language-specific tuning of visual cortex? Functional properties of the visual word form area. *Brain* **125**, 1054–1069. (doi:10.1093/brain/awf094)
26. McCandliss BD, Cohen L, Dehaene S. 2003 The visual word form area: expertise for reading in the fusiform gyrus. *Trends Cogn. Sci.* **7**, 293–299. (doi:10.1016/S1364-6613(03)00134-7)
27. Gwilliams L, Lewis G, Marantz A. 2016 Functional characterisation of letter-specific responses in time, space and current polarity using magnetoencephalography. *NeuroImage* **132**, 320–333. (doi:10.1016/j.neuroimage.2016.02.057)
28. Pammer K, Hansen PC, Kringelbach ML, Holliday I, Barnes G, Hillebrand A, Singh KD, Cornelissen PL. 2004 Visual word recognition: the first half second. *Neuroimage* **22**, 1819–1825. (doi:10.1016/j.neuroimage.2004.05.004)
29. Simon DA, Lewis G, Marantz A. 2012 Disambiguating form and lexical frequency effects in meg responses using homonyms. *Lang. Cogn. Process.* **27**, 275–287. (doi:10.1080/01690965.2011.607712)
30. Binder JR, Medler DA, Westbury CF, Liebenthal E, Buchanan L. 2006 Tuning of the human left fusiform gyrus to sublexical orthographic structure. *Neuroimage* **33**, 739–748. (doi:10.1016/j.neuroimage.2006.06.053)
31. Dehaene S, Cohen L, Sigman M, Vinckier F. 2005 The neural code for written words: a proposal. *Trends Cogn. Sci.* **9**, 335–341. (doi:10.1016/j.tics.2005.05.004)
32. Vinckier F, Dehaene S, Jobert A, Dubus JP, Sigman M, Cohen L. 2007 Hierarchical coding of letter strings in the ventral stream: dissecting the inner organization of the visual word-form system. *Neuron* **55**, 143–156. (doi:10.1016/j.neuron.2007.05.031)
33. Fruchter J, Marantz A. 2015 Decomposition, lookup, and recombination: Meg evidence for the full decomposition model of complex visual word recognition. *Brain Lang.* **143**, 81–96. (doi:10.1016/j.bandl.2015.03.001)
34. Lewis G, Solomyak O, Marantz A. 2011 The neural basis of obligatory decomposition of suffixed words. *Brain Lang.* **118**, 118–127. (doi:10.1016/j.bandl.2011.04.004)
35. Solomyak O, Marantz A. 2010 Evidence for early morphological decomposition in visual word recognition. *J. Cogn. Neurosci.* **22**, 2042–2057. (doi:10.1162/jocn.2009.21296)
36. del Prado Martın FM, Kostić A, Baayen RH. 2004 Putting the bits together: an information theoretical perspective on morphological processing. *Cognition* **94**, 1–18. (doi:10.1016/j.cognition.2003.10.015)
37. Saffran JR, Aslin RN, Newport EL. 1996 Statistical learning by 8-month-old infants. *Science* **274**, 1926–1928. (doi:10.1126/science.274.5294.1926)
38. Cutler A, Butterfield S. 1992 Rhythmic cues to speech segmentation: evidence from juncture misperception. *J. Memory Lang.* **31**, 218–236. (doi:10.1016/0749-596X(92)90012-M)
39. Johnson EK, Jusczyk PW. 2001 Word segmentation by 8-month-olds: when speech cues count more than statistics. *J. Memory Lang.* **44**, 548–567. (doi:10.1006/jmla.2000.2755)
40. Mattys SL, White L, Melhorn JF. 2005 Integration of multiple speech segmentation cues: a hierarchical framework. *J. Exp. Psychol.: Gen.* **134**, 477. (doi:10.1037/0096-3445.134.4.477)
41. Ettinger A, Linzen T, Marantz A. 2014 The role of morphology in phoneme prediction: evidence from MEG. *Brain Lang.* **129**, 14–23. (doi:10.1016/j.bandl.2013.11.004)
42. Gwilliams L, Marantz A. 2015 Non-linear processing of a linear speech stream: the influence of morphological structure on the recognition of spoken Arabic words. *Brain Lang.* **147**, 1–13. (doi:10.1016/j.bandl.2015.04.006)
43. Pesetsky DM 1996 *Zero syntax: experiencers and cascades*, Number 27. Cambridge, MA: MIT Press.
44. Embick D, Marantz A. 2008 Architecture and blocking. *Linguist. Inq.* **39**, 1–53. (doi:10.1162/ling.2008.39.1.1)
45. Friederici AD. 2012 The cortical language circuit: from auditory perception to sentence

comprehension. *Trends Cogn. Sci.* **16**, 262–268. (doi:10.1016/j.tics.2012.04.001)

46. Hickok G, Poeppel D. 2007 The cortical organization of speech processing. *Nat. Rev. Neurosci.* **8**, 393. (doi:10.1038/nrn2113)
47. Indefrey P, Levelt WJ. 2004 The spatial and temporal signatures of word production components. *Cognition* **92**, 101–144. (doi:10.1016/j.cognition.2002.06.001)
48. Binder JR, Frost JA, Hammeke TA, Bellgowan PS, Springer JA, Kaufman JN, Possing ET. 2000 Human temporal lobe activation by speech and nonspeech sounds. *Cereb. Cortex* **10**, 512–528. (doi:10.1093/cercor/10.5.512)
49. Binder JR, Desai RH, Graves WW, Conant LL. 2009 Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cereb. Cortex* **19**, 2767–2796. (doi:10.1093/cercor/bhp055)
50. Binder JR, Desai RH. 2011 The neurobiology of semantic memory. *Trends Cogn. Sci.* **15**, 527–536. (doi:10.1016/j.tics.2011.10.001)
51. Binder JR, Conant LL, Humphries CJ, Fernandino L, Simons SB, Aguilar M, Desai RH. 2016 Toward a brain-based componential semantic representation. *Cogn. Neuropsychol.* **33**, 130–174. (doi:10.1080/02643294.2016.1147426)
52. Fiorentino R, Poeppel D. 2007 Compound words and structure in the lexicon. *Lang. Cogn. Process.* **22**, 953–1000. (doi:10.1080/01690960701190215)
53. Pylkkänen L, Feintuch S, Hopkins E, Marantz A. 2004 Neural correlates of the effects of morphological family frequency and family size: an MEG study. *Cognition* **91**, B35–B45. (doi:10.1016/j.cognition.2003.09.008)
54. Pylkkänen L, Llinás R, Murphy GL. 2006 The representation of polysemy: MEG evidence. *J. Cogn. Neurosci.* **18**, 97–109. (doi:10.1162/089892906775250003)
55. Anderson SR. 1985 Typological distinctions in word formation. *Lang. Typology Syntactic Descr.* **3**, 3–56.
56. Marslen-Wilson WD, Ford M, Older L, Zhou X. 1996 The combinatorial lexicon: priming derivational affixes. In *Proc. of the Eighteenth Annual Conf. of the Cognitive Science Society: July 12–15, 1996, University of California, San Diego* (ed. GW Cottrell), vol. 18, pp. 223–227. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
57. Carota F, Bozic M, Marslen-Wilson W. 2016 Decompositional representation of morphological complexity: multivariate fMRI evidence from Italian. *J. Cogn. Neurosci.* **28**, 1878–1896. (doi:10.1162/jocn_a_01009)
58. Thompson CK, Bonakdarpour B, Fix SC, Blumenfeld HK, Parrish TB, Gitelman DR, Mesulam M-M. 2007 Neural correlates of verb argument structure processing. *J. Cogn. Neurosci.* **19**, 1753–1767. (doi:10.1162/jocn.2007.19.11.1753)
59. Leminen A, Smolka E, Duñabeitia JA, Pliatsikas C. 2018 Morphological processing in the brain: the good (inflection), the bad (derivation) and the ugly (compounding). *Cortex* **116**, 4–44. (doi:10.1016/j.cortex.2018.08.016)
60. Marslen-Wilson WD, Tyler LK. 2007 Morphology, language and the brain: the decompositional substrate for language comprehension. *Phil. Trans. R. Soc. B* **362**, 823–836. (doi:10.1098/rstb.2007.2091)
61. Whiting C, Shtyrov Y, Marslen-Wilson W. 2014 Real-time functional architecture of visual word recognition. *J. Cogn. Neurosci.* **27**, 246–265. (doi:10.1162/jocn_a_00699)
62. Sahin NT, Pinker S, Cash SS, Schomer D, Halgren E. 2009 Sequential processing of lexical, grammatical, and phonological information within Broca's area. *Science* **326**, 445–449. (doi:10.1126/science.1174481)
63. Dominguez A, De Vega M, Barber H. 2004 Event-related brain potentials elicited by morphological, homographic, orthographic, and semantic priming. *J. Cogn. Neurosci.* **16**, 598–608. (doi:10.1162/089892904323057326)
64. Lavric A, Rastle K, Clapp A. 2011 What do fully visible primes and brain potentials reveal about morphological decomposition? *Psychophysiology* **48**, 676–686. (doi:10.1111/j.1469-8986.2010.01125.x)
65. Morris J, Frank T, Grainger J, Holcomb PJ. 2007 Semantic transparency and masked morphological priming: an ERP investigation. *Psychophysiology* **44**, 506–521. (doi:10.1111/j.1469-8986.2007.00538.x)
66. Morris J, Grainger J, Holcomb PJ. 2008 An electrophysiological investigation of early effects of masked morphological priming. *Lang. Cogn. Process.* **23**, 1021–1056. (doi:10.1080/01690960802299386)
67. Neophytou K, Manouilidou C, Stockall L, Marantz A. 2018 Syntactic and semantic restrictions on morphological recomposition: MEG evidence from Greek. *Brain Lang.* **183**, 11–20. (doi:10.1016/j.bandl.2018.05.003)
68. Pylkkänen L, McElree B. 2007 An MEG study of silent meaning. *J. Cogn. Neurosci.* **19**, 1905–1921. (doi:10.1162/jocn.2007.19.11.1905)
69. Pylkkänen L, Oliveri B, Smart AJ. 2009 Semantics vs. world knowledge in prefrontal cortex. *Lang. Cogn. Process.* **24**, 1313–1334. (doi:10.1080/01690960903120176)
70. Qiu S, Cui Q, Bian J, Gao B, Liu T-Y. 2014 Co-learning of word representations and morpheme representations. In *Proc. of COLING 2014, the 25th Int. Conf. on Computational Linguistics: Technical Papers, Dublin, Ireland, 23–29 August 2014*, pp. 141–150. Dublin, Ireland: Dublin City University and Association for Computational Linguistics.
71. Creutz M, Hirsimäki T, Kurimo M, Puurula A, Pylkkönen J, Siivola V, Varjokallio M, Arisoy E, Saraçlar M, Stolcke A. 2007 Morph-based speech recognition and modeling of out-of-vocabulary words across languages. *ACM Transactions on Speech and Language Processing (TSLP)ACM Trans. Speech Lang. Proces. (TSLP)* **5**, 3.
72. Firth JR. 1957 A synopsis of linguistic theory, 1930–1955. *Stud. Linguist. Anal.* **1952–59**, 1–32.
73. Harris ZS. 1954 Distributional structure. *Word* **10**, 146–162. (doi:10.1080/00437956.1954.11659520)
74. Pennington J, Socher R, Manning C. 2014 Glove: Global vectors for word representation. In *Proc. of the 2014 Conf. on empirical methods in natural language processing (EMNLP), Doha, Qatar, 25–29 October 2014*, pp. 1532–1543. Stroudsburg, PA: Association for Computational Linguistics.
75. Huth AG, de Heer WA, Griffiths TL, Theunissen FE, Gallant JL. 2016 Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature* **532**, 453. (doi:10.1038/nature17637)
76. Cinque G. 2006 *Restructuring and functional heads: the cartography of syntactic structures volume 4*. New York, NY: Oxford University Press.
77. Lieber R 1992 *Deconstructing morphology: word formation in syntactic theory*. Chicago, IL: University of Chicago Press.
78. Halle M, Marantz A. 1994 Some key features of distributed morphology. *MIT Work. Pap. Linguist.* **21**, 88.
79. Halle M, Marantz A. 1993 Distributed morphology and the pieces of inflection. In *The view from building 20* (eds K Hale, J Keyser), pp. 111–176. Cambridge, MA: MIT.
80. Harley H, Noyer R. 1999 Distributed morphology. *Glot Int.* **4**, 3–9.
81. Levelt WJ, Roelofs A, Meyer AS. 1999 A theory of lexical access in speech production. *Behav. Brain Sci.* **22**, 1–38. (doi:10.1017/S0140525X99001776)
82. Gwilliams LE, Monahan PJ, Samuel AG. 2015 Sensitivity to morphological composition in spoken word recognition: evidence from grammatical and lexical identification tasks. *J. Exp. Psychol.: Learn., Memory, Cogn.* **41**, 1663–1674. (doi:10.1037/xlm0000130)
83. Soricut R, Och F. 2015 Unsupervised morphology induction using word embeddings. In *Proc. of the 2015 Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, CO, 31 May–5 June 2015*, pp. 1627–1637. Stroudsburg, PA: Association for Computational Linguistics.
84. Mitchell J, Lapata M. 2010 Composition in distributional models of semantics. *Cogn. Sci.* **34**, 1388–1429. (doi:10.1111/j.1551-6709.2010.01106.x)
85. Cotterell R, Schütze H. 2018 Joint semantic synthesis and morphological analysis of the derived word. *Trans. Assoc. Comput. Linguist.* **6**, 33–48. (doi:10.1162/tacl_a_00003)
86. Lazaridou A, Marelli M, Zamparelli R, Baroni M. 2013 Compositional-ly derived representations of morphologically complex words in distributional semantics. In *Proc. of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Sofia, Bulgaria, 4–9 August 2013*, vol. 1, pp. 1517–1526. Stroudsburg, PA: Association for Computational Linguistics.
87. Baroni M, Bernardi R, Zamparelli R. 2014 Frege in space: a program of compositional distributional semantics. *LiLT* **9**, 5–110.
88. Gleitman L. 1990 The structural sources of verb meanings. *Lang. Acquis.* **1**, 3–55. (doi:10.1207/s15327817la0101_2)
89. Landau B, Gleitman LR, Landau B 2009 *Language and experience: evidence from the blind child*, vol. 8. Cambridge, MA: Harvard University Press.

90. Lidz J, Gleitman H, Gleitman L. 2003 Understanding how input matters: verb learning and the footprint of universal grammar. *Cognition* **87**, 151–178. (doi:10.1016/S0010-0277(02)00230-5)

91. Gerhand S, Barry C. 1998 Word frequency effects in oral reading are not merely age-of-acquisition effects in disguise. *J. Exp. Psychol.: Learn., Memory, Cogn.* **24**, 267–283. (doi:10.1037/0278-7393.24.2.267)

92. Moro A, Tettamanti M, Perani D, Donati C, Cappa SF, Fazio F. 2001 Syntax and the brain: disentangling grammar by selective anomalies. *Neuroimage* **13**, 110–118. (doi:10.1006/nimg.2000.0668)

93. Ullman MT. 2004 Contributions of memory circuits to language: the declarative/procedural model. *Cognition* **92**, 231–270. (doi:10.1016/j.cognition.2003.10.008)