DM & Neurolinguistics Chapter
Linnaea Stockall & Laura Gwilliams

Constructivist morphological theories in general, and Distributed Morphology (henceforth DM) in particular, are theories about how morphosyntactic elements are combined via the core syntax to create complex structures, and how those morphosyntactic structures are converted into phonological and semantic representations, themselves subject to further computation within those spellout domains. The theory is concerned with the architecture of the system that generates morphologically complex words in the mind of the speaker, before these representations are converted into instructions for an articulatory system or translated into representations that are interpretable by the conceptual system. In this sense, then, it is predominantly a theory about language production. Morphological processing research, by contrast, especially research using the tools of cognitive neuroscience, has overwhelmingly focused on language comprehension, which fundamentally alters the direction of the information flow. Figure 1 sketches how the core architecture of DM would need to be adapted to serve as a model of language comprehension. The goal of this chapter will be to review the existing research, engaging with each of the following processes and/or architectural distinctions that this model entails:

1. Form based morphological decomposition that provides the entry point into the grammatical system, and relies on 'unpacking' **P**honological **F**orm operations (such as readjustment rules, phonological rules and competitive vocabulary insertion)
2. Access to a stored mental lexicon of **V**ocabulary **I**tems, both roots and function morphemes, critically restricted to only grammatically relevant features, determined by the PF representations identified in (1)
3. Recombination of the constituent elements subject to syntactic structure building constraints operating over the features activated in (2)
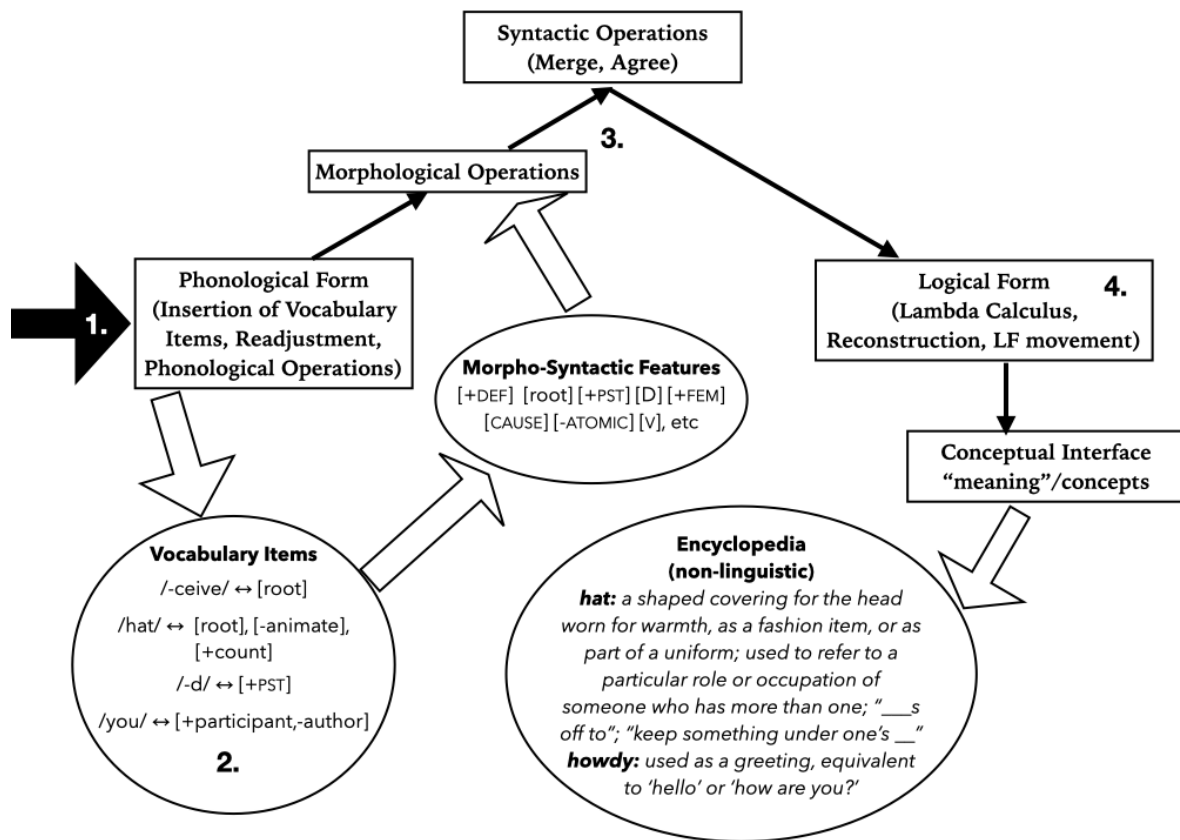4. Compositional semantic interpretation of the structures generated in (3)

Figure 1 A sketch of how the architecture and key components of DM could be adapted to provide a model of the parsing of morphologically complex words in comprehension. Black arrows trace the direction of the flow of information through the core Y model, from PF, through the syntax, to LF. Outlined white arrows represent the relationship between the various distributed 'lexicons' of DM and the grammar. Figure adapted from Harley & Noyer (1999).

We take the distributed architecture of DM as the organising principle of this chapter and discuss in turn neurolinguistic data as it relates to each of the key components of the system.

Before engaging with the literature relevant to testing the predictions of DM, however, we start with a focused introduction to the field of neurolinguistics, briefly outlining its goals and key methods.

1.1 The goal of neurolinguistics

To understand the potential relevance of neurobiological data to linguistic theory, we follow the lead of other recent papers and chapters addressing this issue (Sprouse & Lau 2013, Embick & Poeppel 2015, Gwilliams, 2019) in beginning the discussion with Marr's (1982) distinction between the computational, algorithmic/representational and implementational levels of analysis in cognitive science. While linguistic research in the cognitive domain over the past 60+ years (and here we include psycholinguistic and computational linguistic research as well as formal linguistic research) has made significant progress in developing and refining our understanding of the computational and the algorithmic/representation bases of the human linguistic capacity, the task of investigating the implementational level can only be addressed with the tools of neurobiology. The great promise of investigations at the

implementational level is that discovering *how* the brain implements linguistic computations and representations can ultimately lead to an understanding of *why* the human linguistic system has the computational and representation properties it has. Results from neurobiological research ought to provide significant explanatory power. This promise, however, is very far from being delivered upon. As David Poeppel has argued extensively over the past 15+ years, a major cause for the limited progress in this area is what he calls the Ontological Incommensurability Problem (OIP): computational/representational linguistic research and neuroscience research identify radically different primitive objects and operations, which bear no obvious relationship to each other (Embick & Poeppel 2015). We have no remotely reasonable linking hypotheses between linguistic objects like distinctive features and neurobiological units like neurons, or between linguistic operations like concatenate and neurobiological processes like oscillation. In the absence of any understanding of how the primitives in one domain map to the primitives in the other, we are in the position of 19th century chemists and physicists, awaiting the atomic revolution. The best we can currently do is discover and investigate *correlations* between linguistic stimuli and neural responses and use those correlations to test and refine predictions derived from our computational and algorithmic/representational theories.

By discovering and investigating those correlations, it is possible to gain brain-oriented insights (learning something about the brain using linguistically informed hypotheses) and theory-oriented insights (using neural recordings to refine and/or adjudicate between linguistic theories). In this chapter we focus on the second of these goals, and on results that are relevant to DM specifically, and to morphological processing in general. In this sense, then, we see brain measurements as just another type of empirical data that can be used to inform linguistic theories; on par with native-speaker judgements and corpus investigation. Readers interested in the ways in which linguistically informed research has enhanced our understanding of how the human brain works are directed to Poeppel et. al. 2012 (for a positive review of what we do know), and Embick and Poeppel 2015 for a discussion of the limitations of this research along the lines sketched above.

1.2 Approach

There are a number of different brain-recording techniques that can be used to conduct neuro-linguistic studies. Methods of measuring brain activity non-invasively (as opposed to placing electrodes directly on or in the brain) use different physiological measures as proxies for the underlying brain responses, which presents the trade-off between knowing *where* a neural process happens, versus *when* a neural process happens.

A technique that favours spatial precision over temporal precision is functional magnetic resonance imaging (fMRI). It works by measuring the level of oxygen in the blood (blood oxygenation level dependent imaging: BOLD imaging), under the assumption that neurons that are more active consume more oxygen. This provides an estimate of activity strength in a chunk of cortical tissue (a "voxel") about 1x1x1 millimeters in size (though, this is getting increasingly fine-grained as the precision of the machines improve). This comes with the caveat that activity in a voxel can only be recorded a few times per second -- quite a sparse sampling rate given how rapidly the underlying cognitive processes unfold.

Other techniques favour temporal precision, by taking measurements a few hundred times a second (so, on the scale of milliseconds). Commonly used techniques for language studies are

electroencephalography (EEG - measures the electrical current generated by synaptic potentials) and magnetoencephalography (MEG - measures the magnetic field generated around the electrical current). Here the assumption is that as more neurons fire together in synchrony, the measurable electro-magnetic signals correspondingly increase. With EEG, one can make claims about relative location of a response in the front/back of the brain, or the left/right hemisphere, but generally in not more detail than this. The relative benefit of MEG, therefore, is that one can localise responses to the cortical surface, providing a joint spatio-temporal measure of sensitivity to different linguistic features[1]. Most of this chapter will focus on results from EEG/MEG data, rather than fMRI data, favouring techniques with a higher temporal resolution.
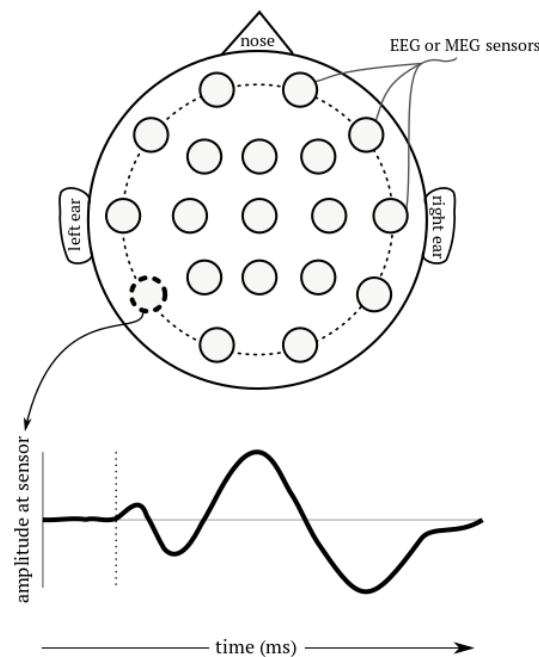

Figure 2 Diagram of how an EEG or MEG sensor in a particular location records continuous changes in electromagnetic activity. Above: schematic of an EEG cap / MEG helmet, where each circle represents a recording position. Below: An example recording from one sensor location over time.

These data make it possible to test whether responses at particular moments in time, and in particular cortical locations, are systematically modulated by certain aspects of the linguistic input. In Table 1, we provide a summary of some spatio-temporal response profiles that are consistently associated with particular morphological computations. By putting together this profile of neural sensitivity, one can infer the sequence of cognitive processes that the brain is performing on the visual or auditory sensory information. It is important that the spatial location and timing of responses are interpreted within a recording technique, as well as within the same cognitive tasks that are being asked of the participants, because the timing and sensitivity of responses shifts around depending on these factors.

Figure 3 summarises the left hemisphere brain areas reliably activated by visually presented morphologically complex words over time, as measured by MEG. Table 1 summarises the

---

[1] The reason MEG has not replaced EEG as the dominant tool to investigate evoked neural activity with high temporal resolution is its prohibitive cost (>$1 million, as compared to ~$20,000 for EEG). While MRI machines are equally expensive, they are essential medical equipment and thus much more widely available than MEG, which has comparatively underdeveloped clinical/medical uses.

distinct stages in processing morphologically complex words identified by MEG vs. EEG investigations.
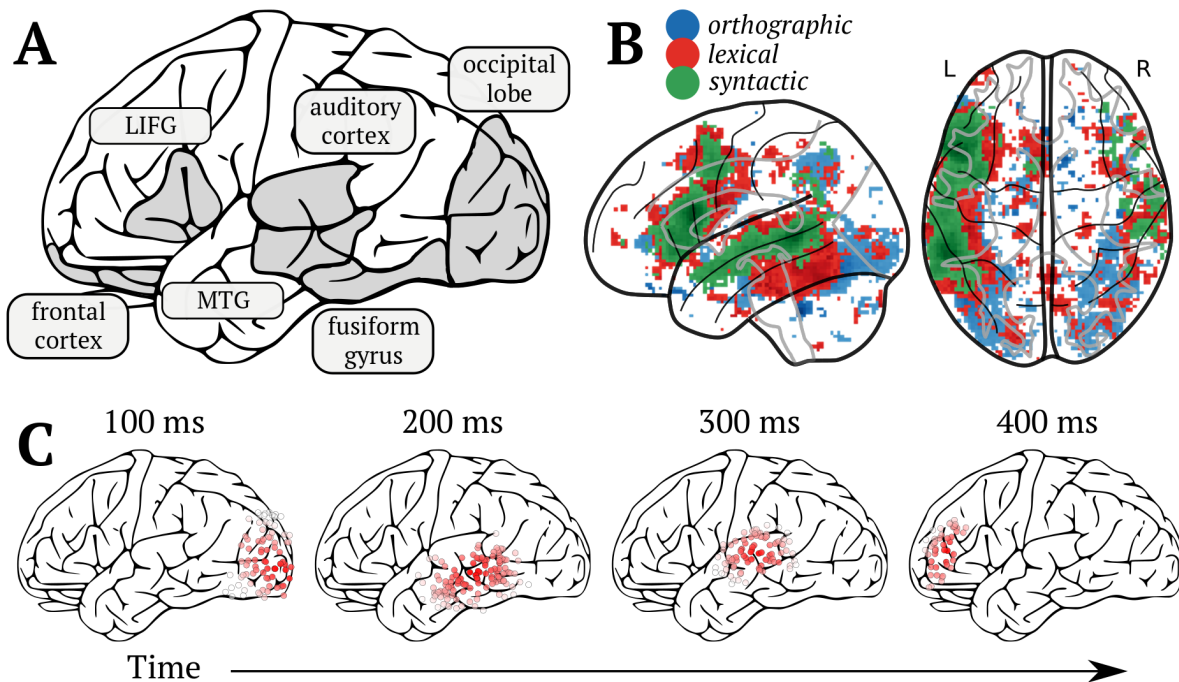

Figure 3 Schematic of how different brain areas are activated by morphologically complex written words at different times. A: Map of primary cortical regions involved in complex word processing. LIFG = Left Inferior Frontal Gyrus, MTG = Middle Temporal Gyrus. B: Voxels consistently activated for different language contrasts, when searching for the key words "orthographic", "lexical" or "syntactic". Maps were extracted from the open-source meta-analysis tool Neurosynth (http://neurosynth.org). C: Illustration of how activity moves across the cortex over time during the processing of a single written word.

Cognitive Process	EEG component	MEG component/response
Morpho-orthographic processing [1]	**N250** (110-300ms) *strongest over midline and anterior left hemisphere electrodes*	**M170** (150-200ms) originates in Left **F**usiform **G**yrus/**FG**
Morpho-phonological processing [1]	n/a	Bilateral **S**uperior **T**emporal **G**yrus/**STG** (120-170ms)
Morpho-syntactic processing [3]	**L**eft **A**nterior **N**egativity/**LAN** (300-450ms) *strongest over frontal electrodes*	Left **P**osterior **T**emporal **L**obe/**PTL** response (200-400ms)
Lexico-semantic processing [2]	**N400** (250-500ms) *strongest over centro-parietal sensors*	**M350** (250-400ms) Left **A**nterior **T**emporal **L**obe/**ATL** response
Morpho-semantic wellformedness processing [4]	n/a	**O**rbito-**F**rontal/**OF** response (300-600ms)
Morpho-syntactic reanalysis/structural-complexity processing	**P600** (500-900ms) *strongest over centro-parietal and posterior sensors*	n/a

The rest of this chapter now considers each of the four key commitments of DM that are outlined above, in detail.

1. Form based morphological decomposition

Undoubtedly the most significant and robust result in the neuro- and psycholinguistic morphological processing literature is that visually presented morphologically complex words are automatically, and very quickly, parsed into their constituent morphemes. Hammerly & Pfau (this volume) discuss the behavioural processing literature establishing this result (see also Rastle and Davis, 2008 for a review of much of the relevant work). The most convincing behavioural processing results, and some of the EEG and MEG research inspired by it (Lavric et. al. 2007, Morris et. al. 2007, 2008, Royle et. al. 2010, Lehtonen, Monahan and Poeppel 2011), rely on the masked morphological priming paradigm, in which a morphologically simplex target word (eg. teach) is preceded by either a morphologically related "prime" word (eg. teacher) or an unrelated word (eg. onion) presented for a very brief duration (typically less than 40ms). Responses to the target are significantly faster when preceded by the morphologically related word, as compared to the unrelated prime. A prime presented for fewer than 40ms is typically not even noticed, let alone consciously recognised (and such primes are often further 'masked' by being preceded by strings of symbols such as ######). Forster and Davis (1984). The fact that such brief exposure to a prime is sufficient to significantly speed lexical decision times to a morphologically related target provides evidence that morphological parsing occurs very rapidly, and without requiring conscious awareness. This inference is further strengthened by the EEG and MEG research establishing that masked morphological priming is associated with a significant decrease in evoked neural activity about 250ms after the onset of the target in the related condition as compared to the same target preceded by an unrelated prime (see Morris and Stockall 2012 for a review of this literature). These behavioural and neural results, however, still rely on an inference: in order to get priming effects (faster lexical decision time, decreased evoked amplitude) in response to the morphologically simplex target, it must have been the case that the masked morphologically complex prime was rapidly decomposed into its stem and affix, allowing the stored stem/lemma representation to be rapidly activated, and facilitating reactivation of the same stem when it is encountered as the target. By contrast, an approach that does not depend on inferring something about the processing of a prime, by measuring a response to a target, was innovated by Zweig and Pylkkänen 2009, who directly compared morphologically simplex and complex words in a single word reading paradigm. Although behavioural techniques such as lexical decision and naming experiments do not always detect a difference between complex and simplex words (see New et. al. 2004 for discussion), Zweig and Pylkkänen hypothesized that MEG should be sufficiently sensitive to detect a reliable difference in an evoked neural response.

Zweig and Pylkkänen (2009) aimed to test whether the morphological complexity of a word modulates responses in the putative "visual word form area" (VWFA; Cohen et. al., 2004). This is a region of cortex located in the left fusiform gyrus that has been shown to respond preferentially to word-like visual objects and is a candidate for processing high-level orthographic features (McCandliss et. al. 2003). Gwilliams, Lewis & Marantz (2016)

associated the VWFA with the "Type II" response as identified by Tarkiainen et. al. (1990) – a component that peaks ~150-200 ms after stimulus onset and responds preferentially to letters over symbols. The results of this study also link the VWFA and Type II response to the M170 component as described above. In the rest of this chapter, we'll refer to this response as the M170.

In Zweig & Pylkkänen's (2009) study, participants read bimorphemic words (e.g. *farm-er / re-fill*), mono-morphemic orthographic control words (e.g. *wint-er / re-ckon*) and mono-morphemic words (e.g. *switch / rotate*) while neural activity was recorded with MEG. In Experiment 1, the bimorphemic words were suffixed words; in Experiment 2, they were prefixed words. Activity in the fusiform gyrus, peaking at ~170 ms, elicited the strongest response to words that were morphologically complex, as compared to the mono-morphemic words. The results of this study therefore suggest that the VWFA, as reflected in Type II / M170 activity, is sensitive to the sublexical morphological structure of a word.

The Zweig & Pylkkänen 2009 results are extremely important, as the first study to confirm that morphological complexity is parsed within 200ms of the onset of a visually presented complex word, and that the brain response involved in this parsing has the same timing and spatial origin as the Visual Word Form Area response (McCandliss et. al. 2003). In a program of research launched in the immediate aftermath of this study, Marantz and colleagues have further investigated this M170 response and its sensitivity to morphological complexity across an increasingly wide range of languages and types of morphology.

Unlike the Zweig & Pylkkänen 2009 study, which employed a parametric design (morphologically simplex vs. complex), the key innovation of Marantz and his collaborators was to adopt a continuous measure of morphological complexity first proposed by Hay 2001: stem:whole word transition probability (also called lemma transition probability, or TPL), defined as the ratio between the frequency of the lexical stem to the frequency of the whole word. The term 'transition probability' reflects the fact that the closer the stem and whole word are to having the same frequencies, the more probable the whole word is as a continuation of (transition from) the stem. Words with a transition probability of 1 are words whose stem and whole word frequencies are identical. Hay 2001, and Hay and Baayen 2002 argue that the relative frequencies of the whole word and the stem determine how 'decomposable' a word is and argue for a 'parsing-line' threshold: only words whose stem frequency is sufficiently high relative to its whole word frequency (ie. words with a relatively low TPL) will be decomposed.

Without buying in to this interpretation, Solomyak and Marantz (2010) adopted the TPL measure as a tool to dissociate sensitivity to morphological structure from sensitivity to other, non-grammatical, factors that might also play a role in early visual word form processing. They measured the evoked brain responses to three categories of affixed words in English: free-stem suffixed words (eg. *tax-able, class-ic*), bound-stem suffixed words (eg. *toler-able, psych-ic*) and unique-stem suffixed words (eg. *vulner-able, galact-ic*), whose stems do not occur in any other context in the language, but which nonetheless exhibit the characteristics of a word composed with that suffix (eg. *vulnerable* is an adjective, and describes a characteristic property, just like *taxable*, and contrasts with a word like *vegetable*, which is a noun, and not plausibly composed of a stem *veget* plus the suffix *-able*). Solomyak and Marantz (2010) focussed their analysis on the M170 response and investigated which variables, for which categories of stimuli, significantly modulated this response. They found that affix frequency modulated M170 responses for all three stimuli types. Further, transition

probability modulated the M170 response for free-stem and bound-stem items, but it was not a significant correlate of M170 activation for the unique-stem items, all of which have a TPL of 1 (we return to a discussion of items of this type below). Notably, the non-morphological, purely orthographic variables such as orthographic affix frequency and orthographic bigram transition probability did not significantly correlate with M170 activation for any category of stimuli.

In a follow up study, Lewis, Solomyak and Marantz (2011) adopted the same correlational analysis techniques and single word reading paradigm to investigate the processing of pseudo-derived words like 'brother' and 'sandal' made so famous by the masked-morphological priming investigations discussed above (and in Hammerly & Pfau, this volume). Lewis et. al. find that pseudo-stem:whole word transition probability is just as robust a predictor of the M170 response to pseudo derived words as TPL was in the Solomyak and Marantz (2010) study for genuinely complex affixed words.

These two studies, then, not only confirm that form-based morphological decomposition is an early process that applies across the board to all potentially morphologically complex words, but also identify two critical sources of information that are used in that process: affix frequency, and stem:whole word transition probability. Though it may not seem earth shattering to confirm that the brain is, in fact, sensitive to morphological complexity, establishing that this sensitivity arises very early in the processing time course (prior to subsequent processing, as we shall see below), and appears to be accomplished on the basis of morphological visual form already speaks to debates in morphological theory. Across the board, purely form based (that is, meaning-blind) decomposition that precedes, and thus feeds, activation of grammatical and lexical representations is a fundamental commitment of Distributed Morphology, while it is not expected given approaches such as Word and Paradigm morphology, which privilege whole word representations and eschew the morpheme as a fundamental unit (Blevins 2003). As illustrated in Figure (1), the DM model requires that form-based detection of potential morphological units occurs prior to any subsequent processing, while activation of syntactic and semantic features, or whole word 'encyclopaedic' meaning can only occur later in the processing stream. This is in clear contrast to models in which meaning and form are not and cannot be dissociated.

In a number of subsequent studies, Marantz and colleagues have further investigated the sensitivity of the M170 decomposition response to a range of other kinds of morphological phenomena and sources of information, critically focusing on precisely the form-based processes that DM hypothesizes are implemented in PF: morphologically constrained phonological rule application and readjustment rules.

Fruchter, Stockall & Marantz 2013 used the masked morphological priming paradigm discussed above, combined with MEG recording, to investigate whether regularly and irregularly inflected English past tense verbs are processed by the same early form-based mechanisms. Participants saw regular past tense~stem prime-target pairs like *gained~gain* and irregular past tense~stem pairs like *sang~sing*. In both cases, with a 33ms prime duration, the evoked M170 response to the target (walk, sing) was significantly larger than it was when the same targets were preceded by unrelated primes (eg. *settle~gain, hand~sing*). Lexical decision responses were also significantly faster in the related conditions. A very similar result is reported by Morris & Stockall (2012) in an EEG experiment where they found equivalent masked priming effects for regular and irregular allomorphy in the N250 response (the related conditions evoked larger responses than the unrelated conditions). But having

found clear support for early, form-based decomposition of irregularly inflected forms, a question arises as to the mechanism. Unlike the cases of morphological complexity discussed above, there is no obvious root or stem in 'sang' or even 'kept' that could form the basis for a stem:whole word transition probability calculation which could then act as the trigger for decomposition. The probability that the whole word will be 'sang' if the stem is 'sing' is 0. However, Albright and Hayes 2003 offer a potential solution, which they term Islands of Reliability: while there are no categorical generalisations about which verb roots in English are irregular, or what form the irregular allomorph will take if they are, there are *stochastic* generalisations. For example, for 6/7 verbs in English which contain the final syllable /{r/l}id/ in the non-past form, the /i/ is an /ɛ/in the past tense form, and the past tense itself is realised with the Ø allomorph (eg. *bleed~bled, read~red*), so a correspondence rule such as i → ɛ/ [X {l, r} ___ d][$_{+past}$] is highly reliable. By contrast, a rule which maps a non-past verb ending in d or t to the same form with the Ø past tense allomorph only applies to 29/1234 possible targets[2] (eg. *hit, cost, spread*, but not *get, need*, etc), and thus is much less reliable. Albright and Hayes (2003) show that these stochastically weighted rules predict behaviour in a task requiring participants to rate past tense forms for novel verbs. Fruchter et. al. (2013) hypothesized that what they call a past tense allomorph's AlbrightScore (the reliability of the past-tense~stem correspondence rule) might be used by the form-based decomposition mechanism as a cue to decompose. And indeed they found that AlbrightScore was significantly correlated with the magnitude of the priming effect for irregular verbs, as predicted. The higher the AlbrightScore, the larger the priming difference. In fact, when they performed separate analyses on the past tense verbs with high AlbrightScores (>0.5) and low AlbrightScores (<0.5), they found that there was only a significant M170 priming effect for the high AlbrightScore verbs.

These results offer clear support for models of morphology like Distributed Morphology that account for stem allomorphy via rule based, phonologically constrained, readjustment rules, with the caveat that the early, form-based decomposition mechanism that identifies stems and affixes may require a rule to be sufficiently well supported by experience in order for it to serve as a trigger for decomposition[3]. It's worth noting, however, that, the AlbrightScore measure is calculated over phonological representations, while the experiment involved visually presented stimuli, so a similar score calculated over orthographic correspondences between past and non-past forms might significantly correlate with the M170 priming response for a larger range of items. It's also interesting to note that Fruchter et. al. (2013) included a pseudo-irregular priming manipulation in their study (eg. prime: *bell*, target: *ball*, on the pattern of *fell~fall*), in order to test whether early, form-based decomposition is over-applied to pseudo-irregulars the way it is for pseudo-derived words like *brother*. Lexical decision response times were about 15ms faster for both the pseudo-irregulars and real irregulars, as compared to the unrelated control condition. However, the difference was only statistically significant for the real irregulars, and there was no hint of any priming effect for the pseudo-irregulars in the M170 response. Research involving languages with regular stem

---

[2] Albright and Hayes derived their corpus from the English portion of the CELEX database (Baayen, Piepenbrock, & Gulikers, 1995), by selecting all the verbs that had a lemma frequency of 10 or greater. For verbs with more than one listed past tense (like dived/dove), they included both as separate entries (e.g. both dive-dived and dive-dove). The resulting corpus consisted of 4253 stem/past tense pairs, 4035 regular and 218 irregular.

[3] The possibility that a morphophonological readjustment rule of this type may need to meet a certain threshold of reliability to be learnable (and thus productively applied) has been extensively investigated by Charles Yang (Yang 2002, 2005; Schuler, Yang & Newport, 2021; inter alia).

allomorphy such as German ablaut is clearly required to understand better how readjustment rules are used as a cue for form-based decomposition.

The English irregular past tense is a case in which PF readjustment rules obscure the surface relationship between two morphologically related forms. While the regular past tense form clearly contains the same stem as the non-past form plus a suffix, the morphological relationship between *sold* and *sell* is not so transparent. However, PF morphophonological operations can also result in the preservation of a transparent morphological relationship, at the expense of phonological well-formedness.

Wray, Stockall and Marantz (2021) hypothesized that this kind of morphophonological information might also serve as a cue to trigger decomposition. To test this hypothesis, they investigated cases of reduplication and pseudo-reduplication in Tagalog, in a single word reading MEG study. Reduplication in Tagalog serves many purposes in both the inflectional and derivational systems of the language, and different reduplicating affixes target different sized portions of the base, with whole word and initial CV reduplication being particularly common (Carrier, 1979). In (1) we see a case of CV reduplication marking an infinitive verb form, and in (2) we see a case of whole word reduplication deriving a frequentive/iterative event nominal.

(1) súlat   ~ súsúlat
     √sulat   RED$_{INF}$+√sulat
     'write'   'to write'
(2) araw   ~ arawaraw
     √araw   RED$_{IT}$ +√araw
     'day'   'every day/daily'

Zuraw (2002) discusses pseudo-reduplicants in Tagalog, which she characterises as "roots of which one portion (the pseudo-reduplicant) is identical to another (the pseudo-base), but whose pseudo-base cannot stand alone, and which lack the morphosyntactic or semantic characteristics of a morphologically reduplicated Tagalog word.". Examples of such pseudo-reduplicants are in (3).

(3) a. tátal (N: 'wood chips, splinters, or shavings'
     b. laláʔ (A: 'acute')
     c. mismís (N: 'remnants of food left after a meal')
     d. sagság (A: 'split; blunt; sagging; at the peak of success')

As the examples in (3) show, pseudo-reduplication can appear to involve CV reduplication or whole word reduplication, but, critically, these words are not plausibly analysed as involving genuine reduplication because (a) their 'pseudo-base' is not only not an attested stem in Tagalog, but is also too small to be a legal phonological word, and (b) there is no plausible syntax or semantics associated with either the pseudo-base or the pseudo-RED morpheme.

However, Zuraw (2002) notes that pseudo-reduplicants fall into two separate classes, based on their well-formedness with respect to Tagalog phonological rules. Tagalog has phonological rules such as nasal coalescence, and intervocalic tapping, which systematically over or under apply to real reduplication in a way that preserves base-reduplicant faithfulness (McCarthy and Prince 1995). For example, if a prefix ending in a nasal consonant is attached

to a reduplicant prefix, the nasal coalescence triggered in the reduplicant also occurs in the base, even though the base is not adjacent to the trigger as in (4).

(4) paN+red$_{CV}$+pula → pa-mu-mulá

Most pseudo-reduplicants do not exhibit such non-transparent application of phonological rules, however, a small class do. Because the explanation for the over or under application of the phonological rules relies on a requirement that the reduplicant and base preserve their phonological similarity, the pseudo-reduplicants that also over and under apply these phonological rules appear to behave as if they have bound, morpho-phonological bases[4].

Wray et al's study therefore had multiple aims: first to see whether stem:wholeword transition probability (TPL) also served as a reliable cue to morphological complexity for a range of different Tagalog affixation processes (reduplication, infixation and circumfixation), and secondly to see whether pseudo-reduplicants which behave as if they have bases (+BR.FAITH pseudo-reduplicants) are parsed as morphologically complex during the form-based decomposition stage (M170). They found that the M170 reliably tracked transition probability in the true and +BR.FAITH pseudo-reduplication conditions (and in the infixed and circumfixed conditions), and that the timing and location of this effect corresponded to all the previous work on English discussed above. No such TPL correlation effects were found for the pseudo-reduplicants which transparently apply the expected phonological rules. This result suggests that morphological constraints on phonological rule application can serve as a cue to support a form-based decomposition analysis: if a phonological well-formedness rule is violated only in morphologically complex words, the detection of the rule violation may support decomposition.

Fruchter et. al. (2013) and Wray et. al. (2021) therefore both provide evidence that the process of form-based morphological decomposition indexed by the M170 response is sensitive to precisely the kinds of morpho-phonological operations that DM proposes occur in the PF component.

There are far fewer investigations using neuro-imaging to probe the mechanisms used to decompose spoken words. A number of studies using fMRI or other techniques which flatten the time-course do provide clear support that morphologically complex spoken words are processed differently than morphologically simple spoken words in English (Bozic et. al. 2014, Longe et. al. 2006 , Whiting et. al. 2013, Vannest et. al. 2005), Polish (Bozic et. al. 2013), and Finnish (Laine et al, 1999), with general, though not perfect agreement that complex words elicit increased activity in the left temporal and frontal lobe areas seen above in Fig. 3. For the processing of speech, because the signal unfolds over time, it presents an analytical challenge: with a written word, we can 'start the clock' at the instant the stimulus is presented on the screen. With a spoken word, it is much less clear when the critical information that might trigger morphological analysis is available (see, for example, Wurm et al 2006 for discussion). However, one very interesting study from the perspective of uncovering how spoken words might be parsed into morphological constituents on the basis of form properties is Gwilliams & Marantz, (2015). In this project, the authors hypothesized that comprehenders might leverage predictions on a phoneme-by-phoneme basis in order to detect morphological units. It is known that activity in superior temporal gyrus (STG) is

---

[4] Perhaps similar to bound latinate stems in English like -ceive and -sume, which have no meaning, but behave like vocabulary items with respect to phonological alternations, eg. receive~reception, perceive~perception

modulated by how expected a phoneme is to occur, given the previous phonemes in the word already encountered (e.g. the probability of /t/ given that /habita-/ has already occurred, formally: t | habita). Ettinger, Linzen & Marantz, 2014, using MEG, showed that the strength of the STG response to this probability is strengthened at morpheme boundaries (e.g. *bruis-**er*** vs. *bour**bon***). Gwilliams and Marantz (2015) investigated the processing of Arabic words consisting of a discontinuous tri-consonantal root (eg. K-T-B) and a CVCVCV template (eg. *kataba*). They predicted that if Arabic words were processed as morphologically complex, the STG phoneme-prediction response should be modulated by how predictable the final consonant is given the other root consonants (b | kt), not by how predictable it is given the whole, continuous string (eg. b | kata). They found that, indeed, about 150ms after the onset of the 3rd consonant, STG activity was reliably correlated with root~phoneme predictability, but not with whole-word~phoneme-predictability.

While more research into the processing of morphologically complex spoken words is clearly required, Gwilliams & Marantz (2015) suggests that in the auditory domain, morphological analysis is also rapid, and form-based, and is sensitive to statistical generalisations over relatively abstract morphological structure, consistent with the visual processing literature, and the DM model. Recent work in this direction (Gwilliams, King, Marantz & Poeppel, 2022) has shown that during continuous speech processing, phonetic properties of the input are maintained longer when the identity of morpheme units is less certain, and cease to be maintained at the point of morphological recognition. Furthermore, if a speech sound within a morpheme is more predictable, it will start to be processed earlier than if the speech sound is less predictable. This suggests that phonetic and morphological processing are tightly linked, and phonetic processing can be used as a window to explore higher order sub-lexical computations. Also see Gwilliams & Davis (2021) for a review of using predictability of speech inputs to investigate lexical and sub-lexical processing.

Before turning from decomposition to the subsequent processing of morphological constituents, one recent M170 study merits attention. This study, by Gwilliams and Marantz (2018), suggests that morpho-syntactic information may also be available during the form-based decomposition stage, in a way that is less obviously predicted by the DM architecture.

As described above, early stages of morphological processing appear to be blind to the "true" morphological relationship between stem and whole word. This leads to the decomposition of words like *corner* and *brother* which are not, in fact, morphologically complex. What is unclear, based on these results, is whether the morphological parser is purely using surface-based regularities in order to identify the morphemes, or whether it is instead initiating decomposition because the strings *broth* and *er* are both real morphemes in the language. Under the latter view, *corner* and *brother* are decomposed because *corn, broth* and *-er* all exist as entries in the lexicon, regardless of the lack of morphological relation to the whole word.

In order to address what cues the system uses in order to decompose written input, Gwilliams & Marantz (2018) conducted an MEG study contrasting three critical types of items: i) unarguably complex (e.g. *farmer*); ii) pseudocomplex, but containing an attested stem and an attested affix (e.g. *corner*) iii) and items containing an attested affix, but not an attested stem (e.g. *winter*). It was found that words that did not contain an existing stem (iii) were not decomposed, but the two other conditions were decomposed, suggesting that the presence of an affix alone is not sufficient to drive the morphological parser.

One possible explanation for this result is that the transition probability from *wint -> er* is 1, because the two pseudo-units always co-occur. An alternative explanation is that the system does not decompose *winter* because there is no potential for a grammatical parse: *winter* cannot be interpreted as *someone who wints* or *more wint than X* – the pseudo-affix 'er' does not make any predictable contribution to the grammar of the whole word. Based on these alternative hypotheses, and addressing the question of morpho-syntactic processing more specifically, the authors asked whether the decompositional process is initiated for words containing a root that only ever occurs with a single morpho-syntactic frame, just in case the affix *does* make the expected grammatical contribution (e.g. *excursion* is a valid *-ion* de-verbal nominalisation, similar to *donation* or *vibration,* that just happens to contain a stem that occurs in no other environments[5]). Critically, the transition probability for these items was also equal to 1.

M170 responses to these unique-root words were indistinguishable from the unarguably complex items, and significantly differed from the *winter*-type words, apparently suggesting that morpho-syntactic rules might be utilised by the early, morphological processing system to initiate morphological decomposition. Unlike the results pointing to the availability of morpho-phonological cues, which are precisely the kind of information that the architecture of DM predicts should be available to aid form-based grammatical analysis, this apparent early availability of syntactic information is surprising. Determining whether a potentially complex word is syntactically well-formed or not should be a later process.

However, more careful consideration of what this experiment is actually telling us helps to support the architecture of DM, not to undermine it. A plausible explanation for the Gwilliams & Marantz (2018) findings is that because words like *excursion* and *vulnerable* contain affixes that make predictable contributions to the whole word's grammatical properties, the learner analyses them as complex, and therefore stores the otherwise unattested stems *excurs-* and *vulner-* as root vocabulary items. No such analysis is ever viable for *winter* or *sausage* and therefore *wint-* and *saus-* are not added to the VI lexicon[6]. Participants in the Gwilliams & Marantz experiment weren't, therefore, actually doing morpho-syntactic analysis as part of the initial form-based parse, they were simply distinguishing between potential pieces that do correspond to a VI and those that don't. The parsing itself remains fully within the 'visual word form' domain.

---

[5] Note that *excursion* is more similar to *vibration* than *donation*, in that it does not seem to imply an external argument, and thus be grammatical as a transitive nominal. Compare (i)-(iii).

    (i)   John's donation of the money

    (ii)   *John's vibration of the strings

    (iii)   *John's excursion of the city

Marantz (1997) argues that the contrast between (i) and (ii) is an argument that nominalisations are built directly on a category neutral root, rather than on a verbal base. Under this analysis, $\sqrt{excur}$ is simply a category neutral root that does not imply external causation like $\sqrt{vibrat}$ and $\sqrt{grow}$.

[6] A learner encountering *winter* for the first time might attempt decomposition and posit a root *wint*, but such an analysis would require that they interpret *winter* as either an event nominal meaning 'one who wints' or a comparative, meaning 'more wint than x', or that they would have to posit yet a third *-er* morpheme, which occurs with just two stems (winter and summer) and assign it appropriate syntactic and semantic features. Note that this third *-er* would have to be consistent with both the nominal and 0-derived verbal uses of *winter*, and with the addition of further derivational affixes (e.g. *winterize*).

An extension of this result, therefore, is that even morphemes that *always* occur together – words for which the probability of transitioning from stem -> suffix is 1 -- are still processed and stored as separate pieces, just in case the grammar of the language supports their storage as such, as we have argued it does for *vulnerable* and *excursion* and other such unique-stem words. From this perspective then, the transition probability variable so consistently used to diagnose sensitivity to morphological structure by the early, form-based parsing system is not necessarily itself an actual cue used by this system, but rather a reflection of the frequency with which the units combine.

Overall, it is not resolved what the transition probability variable is quantifying. In most cases, we find that TPL is positively correlated with M170 amplitudes: the higher the TPL, the larger the evoked response. But in the Zweig & Pylkkänen (2009) study that kick-started the single-word reading study enterprise, the M170 response was larger for complex words (eg. teacher, refill) than pseudo-suffixed (eg. winter) and bound-stem prefixed (eg. resume) items[7]. And in Wray et. al. (2021), the authors find a significant correlation between TPL for reduplicated, infixed and circumfixed items, but only the reduplicated items show the positive correlation reliably seen in other languages – the infixed and circumfixed items show a negative correlation, such that M170 amplitudes are lower for the higher TPL words. This set of results highlights a more general issue for neurolinguistic study -- finding that the brain tracks a statistical property of language is easy: correctly interpreting what the tracking actually reflects is very difficult, underspecified and permits a degree of artistic license.

Critically for current purposes, however, the basic finding is not in dispute: within 200ms of encountering a word-form-based analysis that includes sensitivity to morphophonological rules and constraints detects all the elements within that word that plausibly correspond to a stored Vocabulary Item, consistent with the expectations of the DM model in Fig1. What neurolinguistic research can tell us about the storage and activation of these VIs is the topic of the next section.

2. Access to a stored mental lexicon of Vocabulary Items

The key innovation that gives Distributed Morphology its name is the proposal that there is no single Lexicon – the kinds of information that are stored in the lexicon in other theories are here distributed across three distinct lists or lexicons: (a) a set of purely syntactic features that serve as the input to the syntax; (b) the set of Vocabulary Items that store the relationship between a bit of phonology (and possibly orthography) and the syntactic features that bit of phonology can realize; and (c) the Encyclopedia, which relates vocabulary items to non-linguistic meanings. Thus a prediction of DM is that, in comprehension, there should be a stage of processing in which only the Vocabulary Item list is accessed. This list includes root morphemes, function morphemes, and arbitrarily long 'chunks' that have no internal syntactic structure (including, for example non-decomposable idioms/frozen phrases such as

---

[7] TPLs were not calculated for the Zweig & Pyllkänen items, and a complete item list is not available, but as we've just discussed, *winter* type items by definition have a TPL of 1, and so should be associated with a larger M170, but Zweig & Pyllkäanen find that these items, as well as their unambiguously monomorphemic items like *straight* evoke lower M170 peak amplitudes than their genuinely/transparently complex items such as *teacher*. Interestingly, the TPL for the bound-stem *resume* type items which were contrasted with the monomorphemic and complex prefixed items, varies considerably, *rebuke* and *refrain* have TPLs of 1, but *retain* has a TPL of only 0.027 (see, eg. contain, pertain, detain, etc). The items in the complex conditions likely varied considerably in their TPLs.

'by and large'), but critically it contains only phonological form information and morpho-syntactic information. Example VIs as proposed by Harley and Noyer (1999) are in (4).

(5) a. /i/ ←→ [___,+plural]
  A Russian affix (Halle 1997)
 b. /n/ ←→ [___,+participant+speaker,plural]
  A clitic in Barceloni Catalan (Harris 1997a)
 c. /y-/ ←→ elsewhere
  An affix in the Ugaritic prefix conjugation (Noyer 1997)
 d. Ø ←→ 2plu
  A subpart of a clitic in Iberian Spanish (Harris 1994)
 e. /kæt/ ←→ [DP D [LP ____]]
  Root inserted in a nominal environment (Harley & Noyer 1998a)

There is a substantial body of work pointing to the M350 response component as an index of 'lexical' access. The M350 is an MEG evoked response originating in left anterior medial temporal lobe and peaking approximately 350ms after the onset of visually presented single word stimuli. It has been shown to be sensitive to a whole host of variables associated with the storage and activation of lexical representations such as stem (but not whole word) frequency (Embick et. al. 2001), repetition priming (Pylkkänen et. al. 2000), phonotactic probability (Pylkkänen, Stringfellow and Marantz, 2002), morphological stem priming for both regular and irregular allomorphs (Stockall and Marantz 2006), morphological root and template priming in Hebrew (Kastner et. al. 2018), morphological stem priming for both transparent and opaque compound constituents (Fiorentino & Poeppel 2005), morphological family size (Pylkkänen et. al. 2004) and derivational family entropy (Fruchter & Marantz 2015). These variables all plausibly index the activation of the Vocabulary Item lexicon, as they could all track properties of the phonological chunks and their associated syntactic feature specifications. However, the M350 is also modulated by variables such as number of related senses (Pylkkänen, Llínas and Murphy 2006, Beretta, Fiorentio and Poeppel 2005), and has been argued to be an MEG correlate of the N400 ERP[8] response (Pylkkänen & Marantz 2003), which is robustly associated with semantic well-formedness and plausibility in a very large literature (see Lau, Phillips and Poeppel, 2008 for a review). This body of work, then, while supporting the general contention of DM that what we store are morphological pieces, not whole words, fails to clearly distinguish between the distributed lexicons model of DM and the single lexicon view of most other theories.

We are not aware of any neurolinguistic work that directly tests the question of whether we can dissociate 'lexical' lookup restricted to only the grammatically relevant information contained in the VI lexicon from 'lexical' lookup of encyclopedic information. However a partial answer comes from a body of work pioneered by Friederici and her colleagues in the late 90's that identified the Early Left Anterior Negativity (ELAN) ERP response. The ELAN was first identified by Hahne & Friederici (1998) in a study comparing the processing of auditorily presented grammatical German sentences with sentences containing an additional function morpheme that makes the sentences ungrammatical as in (6).

(6) a. Das Hemd wurde gebügelt.

---

[8] ERP = Event Related Potential – the conventional term for time-locked EEG signals evoked by a specific event or stimulus. ERF (Event Related Field) is the analogous term for time-locked MEG signal, but this is rarely used.

the shirt   was   ironed

   b. *Die Bluse wurde am gebügelt

    the blouse was   on ironed

Hahne & Friederici (1998) compared the ERPs to *'ironed'* in the two conditions and found that the ungrammaticality of *gebügelt* in (6b) evoked an increased negativity over left anterior channels 150-200ms after the onset of the word. In follow-up work, Hahne & Friederici (2002) compared sentences of the type in (6) with sentences like those in (7).

  (7) a. ?Das Gewitter   wurde gebügelt.

      the thunderstorm was   ironed

    b. *Das Gewitter   wurde im   gebügelt.

      The thunderstorm was   in-the ironed

They found that whereas sentences like the anomalous *'Das Gewitter wurde gebügelt.'* evoked the expected N400 response (an increased negativity strongest over centro-parietal electrodes peaking around 400ms), sentences like (7b), which are both semantically anomalous and syntactically ungrammatical, only evoke the ELAN response (and a later P600), but no N400. They interpret this finding to mean that syntactic well-formedness is determined more rapidly than, and independently from, semantic well-formedness. If a sentence 'crashes' syntactically, no issues around semantic interpretation arise, since there is no grammatical structure to interpret. Friederici et. al. (2000) replicated the initial ELAN ERP findings with MEG: syntactic violations such as (6b) evoked increased amplitudes between 130-160 compared to grammatical sentences like (6a), and dipole modelling located the source of the activation in left, frontal cortex. In these German studies, the element that makes the sentence ungrammatical is always a free-standing preposition, preceding a verb, so these studies are typically interpreted with respect to syntactic processing, not morphosyntax.

However, Dikker, Rabagliati and Pylkkänen (2009) adapted the ELAN paradigm to English with an important tweak. In addition to a preposition-triggered violation condition, they also included a condition in which an inflected verb occurred in a context that was strongly biased towards a noun, and a condition in which a (bare) noun occurred in a context in which a verb was expected as in (8). This modulation of the paradigm allowed them to test the hypothesis that "rapid word category identification relies on overt category-marking closed-class morphemes." (Dikker et. al. 2009, p294)[9].

  (8) a. *The boys heard Joe's **about** stories Africa. (compared to: The boys heard Joe's stories **about** Africa.)

    b. *The discovery was in the **reported**. (compared to: The discovery was **reported**.)

    c. *The discovery was **report**. (compared to: The discovery was in the **report.**)

Sentences were presented visually, word by word, so that when the critical word is presented on the screen, it is not actually ungrammatical, merely highly unlikely in context (in each

---

[9] This change in the design also allowed Dikker et. al. to avoid possible issues with the baseline in the Friederici lab studies. The debate about the significance of these issues is beyond the scope of this chapter but see Lau et. al. (2006) for an early discussion, and Steinhauer and Drury (2012) for a more substantial discussion of the problems with the initial ELAN design. Critically, the Dikker et. al. study avoids these issues by crossing violation type with context: the ungrammatical version of (6b) can be compared both with the grammatical version of (6b), to compare the same word in different contexts, and with the grammatical version of (6c) to compare the same context with minimally different words.

case, the sentences could have continued grammatically as, for example, *The discovery was in the reported story*.). The goal of this change was to focus on the generation and validation of syntactic expectations. Dikker et. al. found that when the unexpected word was either a free standing function morpheme (8a) or contained an overtly realised function morpheme (8b), a reliable increase in evoked amplitudes was elicited approximately 125ms after the onset of the critical word, and this response was localised to left posterior generators (plausibly located in peripheral visual areas in occipital cortex). The timing is consistent with the ELAN response, but the localisation is not, and instead points to a role for relatively low-level sensory areas being involved in evaluating form-based predictions. Critical for current purposes, Dikker et al's study points to very rapid processing associated with phonologically contentful functional elements, consistent with accessing representations that map phonological form to syntactic features, which is to say, DM Vocabulary Items. Neither Dikker et. al. (2009), nor any follow up project, has included a morphologically signaled syntactic vs. semantic manipulation along the lines of Hahne & Friederici (2002) that would dissociate access to the VI lexicon from access to the Encyclopedia.

The Dikker et. al. (2009) result suggests that form-based evaluation of whether a function morpheme is consistent with syntactic expectations can proceed extremely rapidly (within 130ms). Interestingly, Solomyak and Marantz (2010) found a significant affix frequency effect at around 130ms after the onset of single morphologically complex words presented in isolation, while TPL only significantly predicted the later M170 response. These results both suggest that activation of closed-class/function Vocabulary Items can be extremely rapid, and begin before whole word morphological parsing occurs.

Data from MEG single word reading studies lends additional support to the claim that activation of morpho-syntactic representations can be distinguished from activation of lexical semantic representations. Linzen, Pylkkänen and Marantz (2013) investigated the processing of morphologically simple verbs[10] in isolation. The verbs they chose varied both in what kinds of syntactic constituents they could be followed by, and the relative frequency of each of these possible continuations. For example, *sleep* is unergative and therefore rarely followed by a DP[11], while both *accept* and *prove* allow both DP and CP arguments but vary in the relative frequencies with which the two arguments occur (the probability of a nominal after *accept* is 98%, while for *prove*, a nominal is only probable 23% of the time). The variability in both the kind and frequency of the syntactic constituents that follow a particular verb can be captured in a variable called subcategorisation or continuation[12] entropy, and can be taken to reflect the degree of uncertainty about the syntactic category of the verb's complement (Moscoso del Prado Martín, Kostić, & Baayen, 2004). Linzen et. al. (2013) find that this continuation entropy variable reliably correlates with evoked MEG activity in the left Anterior Temporal Lobe (ATL) between 200 ms and 300 ms after the onset of visually presented single verbs, such that higher entropy verbs elicited less ATL activity. Critically,

---

[10] For the purposes of their study, Linzen et. al. define a verb as a word that is used as a verb at least twice as frequently as it is used as a noun, based on the CELEX corpus (Baayen & Piepenbrock, 1995).

[11] In fact in a sufficiently large corpus even *sleep* will occasionally be followed by a nominal, as in "…8 or more hours of sleep **a night**….", and thus the probability of a nominal following *sleep* will be > 0. *Sleep* can also occur followed by a nominal in in causative constructions such as 'We sleep five people in each room' (Levin 1993: 79, 82)

[12] Linzen et. al. (2013) use the term subcategorisation entropy, but Sharpe et. al. (2018) in a follow-up study, point out that because the variable is derived from corpus statistics, it does not discriminate between actual verbal complements, and constituents which happen to follow a verb but are not subcategorised for, and so they argue that 'continuation entropy' is a more accurate term. We adopt their usage here.

other variables associated with the lexical semantics of these same verbs (such as number of senses, and contextual distinctiveness[13]) had no such significant effect on ATL activation during processing. Number of senses had a marginal effect in Broca's area[14] between 300 ms and 400 ms, while contextual distinctiveness had no detectable effect in any of the ROIs investigated. Linzen and colleagues, then, find a spatial and temporal distinction between activation of a verb's syntactic subcategorisation possibilities (200-300ms, in the left ATL) and activation of a verb's semantic polysemy (300-400ms, in Broca's area), consistent with the distributed lexicons hypothesized by DM. Additionally, the direction of the continuation entropy correlation is also informative. Linzen et. al. find that higher entropy (that is, more possible continuations and/or more similar probabilities among possible continuations) is associated with a decreased evoked neural response. This is not what would be expected if what we were seeing was the simultaneous activation of a number of possible continuations and competition between them - increased competition should lead to an increase in activation, as in the case of competition between two unrelated meanings of an ambiguous word (Pylkkänen et al, 2006, Simon et. al., 2012). Rather, the decreased activation for the higher entropy words is in line with the finding by Fruchter and Marantz (2015) that increased morphological derivational family entropy[15] (Moscoso del Prado Martín, 2004) is associated with decreased evoked activity from the left medial temporal gyrus[16] between 241-387ms after the onset of visually presented suffixed words. The finding that syntactic constituent continuation entropy, and morpho-syntactic continuation entropy are both associated with reduced processing responses both supports the DM 'syntax all the way down' hypothesis, and supports the view that morphologically related forms are related syntagmatically (as variants of the same root) and not paradigmatically (as separate, similar items) (de Saussure 1916/1966, Baayen et. al. 2007). We have emphasised the syntactic continuation entropy results over the derivational family entropy results in this discussion since they point to a higher level of purely syntactic abstraction: syntactic category continuation entropy is calculated over DPs and CPs, not the individual VIs such as *-ate*, *-able* and *-ize* that determine derivational entropy[17].

---

[13] Contextual Distinctiveness measures the extent to which the collocational context of a word diverges from the average collocational context in the language. For example, *lane* occurs in several common collocations (fast lane, bike lane), and therefore is more distinct than other words from the average collocational context (McDonald & Shillcock, 2001)

[14] Here defined as left Inferior Frontal Gyrus and the pars triangularis (Brodmann area 45). Note that the term 'Broca's Area' is both very widely, and very variably used. See Tremblay & Dick, 2018 for a review of the many problems with both the term, and the overly simplistic and outdated models of language processing it often conjures up.

[15] A token frequency weighted measure of the number of different (non-∅) derivational affixes a root is attested with in the relevant corpus (in this case CELEX)

[16] Fruchter and Marantz (2015) do not further subdivide their ROIs into anterior and posterior regions, so it is not clear whether the brain area associated with their derivational family entropy variable is the same area as the ATL region Linzen et. al. (2013) find to be the locus of their syntactic continuation entropy variable.

[17] It would be extremely interesting to investigate an abstract version of derivational family entropy that counted over syntactic continuations such as n, v and a. King, Linzen and Marantz (2015) attempts something along these lines by investigating the processing of mono-morphemic, category ambiguous words in English such as kiss. The authors hypothesized that category ambiguity should evoke a similar processing response to affix ambiguity (Fruchter & Marantz 2015) and syntactic continuation ambiguity (Linzen et. al. 2013), namely increased ambiguity (increased noun/verb category entropy) should evoke decreased activation in left temporal lobe between 200-350ms. That is to say, syntactic category should behave like a syntactic continuation/affix on an acategorical root, not like a stored lexical property. In exploratory analyses, they find exactly this pattern, however the effect proved unstable when subjected to more systematic statistical analyses (Linzen, pc). Further research is clearly required.

3. Recombination of the constituent elements subject to syntactic structure building constraints operating over the features activated in (2)

We've considered the evidence that points to the activation of the morpho-syntactic feature representations stored in the VI lexicon. The obvious next step is to consider research investigating how those feature representations are used to build syntactic representations of morphologically complex words.

A pair of recent MEG studies addresses exactly this question, one investigating verbal suffixes in Greek (Neophytou et. al. 2018) and the other investigating verbal prefixes in English (Stockall et al, 2019). These two studies are concerned with testing a basic prediction of DM and all other constructivist or realisational syntactic theories, namely that syntactic structure building is a distinct, autonomous process that is necessarily prior to semantic interpretation (see, for example Marantz & Wood, 2017, and Alexiadou et. al. and McGinnis & Wood, this volume). These studies do this by taking affixes that have both clear syntactic restrictions on the category of the stem they can attach to (c-selectional restrictions) and clear restrictions on the type of verbal event/verbal argument structure they can attach to (s-selectional restrictions), and then systematically violating one or the other restriction and measuring the evoked neural response. In behavioural lexical decision experiments, c-selection violations are consistently associated with faster, and more robust rejections than s-selection violations across all languages tested (Greek: Manouilidou 2007, English and Greek: Manouilidou and Stockall, 2013, Slovenian: Manouilidou et. al. 2015).

Neophytou et. al. (2018) created their materials by combining three deverbal, adjective deriving suffixes (-simos, -tos, and -menos[18]) with either nominal stems (thus violating the suffix c-selections) or with unergative verbal stems (satisfying c-selection requirements, but violating the s-selection restrictions, see Anagnostopolou & Samioti 2014, Alexiadou 2018 for a discussion of the argument structure properties of these suffixes). These two categories of grammar-violating items were complemented by fully grammatical, familiar derived adjectives ending in the same three suffixes as in Table 2

Suffix	Grammatical	Ungrammatical	
		Category Violations	Argument Structure Violations
-simos	ananeo-simos	varel-imos	gela-simos
	(renewable)	(barrel-able)	(laugh-able)
-tos	sevas-tos	lahano-tos	tremi-tos
	(respectable)	(cabbage-able)	(tremble-able)

---

[18] Neophytou et al (2018) follow the common convention in Greek linguistics of citing these suffixes without any internal structure, however, the -os ending marks masc.sg.nom inflection, so that they are in fact morphologically complex

-menos	lavo-menos	ahino-menos	argi-menos
	(wounded)	(urchin-able)	(be late-able)

Table 2 Sample materials from Neophytou et. al. (2018). Materials were presented without hyphens.

Neophytou et. al. added additional grammatical derived adjectives ending in other suffixes (-ikos, -tikos, -inos, and -eros) to have equal numbers of grammatical and ungrammatical items. They investigated the evoked responses to the two types of violation item in two brain areas: the left Temporal Lobe, already discussed above as the area sensitive to a number of lexical and morpho-syntactic variables, and the Orbito-frontal (OF) Cortex, an area associated with compositional semantic well-formedness in a number of studies we'll discuss more fully in section (4). They found a significantly larger response to the category violation items than to the argument structure violation items in left posterior temporal lobe between 200-300ms after stimulus onset, across all three suffixes. Left posterior temporal lobe has been associated with processing syntactic structure (Flick and Pylkkänen 2020, Flick, Oseki, Kaczmarek, Al Kaabi, Marantz & Pylkkänen, 2018 and Matar & Marantz, 2021), supporting the interpretation of the violation type effect as an effect related to the syntactic features of the derived pseudowords. Critically, Neophtyou et.al. also found the opposite pattern, with more activity associated with the argument structure violations than the category violations, in the OF ROI, between 400-500ms, thus clearly dissociating the c-selection violation and s-selection violation responses in space and time.

Stockall et. al. (2019), ran essentially the same study, but with the English prefixes *re-*, *un-* and *out-* in lieu of adjectival suffixes. A key difference between the two studies besides the choice to use prefixes rather than suffixes is that whereas all 3 Greek suffixes require verbal predicates which include an affected internal argument, and thus the argument structure violation items for all 3 suffixes involved unergative verbs, the English prefixes are not so uniform. English *re-* is the most homogenous: it attaches only to verbal stems and requires that the vP it is part of is eventive, and contains a single, real (underlying) direct object (see Marantz 2007, 2009, Alexiadou, Anagnostopolou & Lechner 2014 and references therein for discussion)[19]. *Re-* is thus incompatible with unergatives, ditransitives and statives. English has two *un-* prefixes, the verb attaching *un-* we are interested in here, and an adjective attaching *un-*. The adjectival *un-* has fairly basic negation semantics (though see Horn 1995 for a more nuanced analysis), but the verbal *un-* has the more complicated meaning of restoring an affected object to its prior/natural state (eg. to *unbend a wire* is to restore a wire to a straight/unbent state) and therefore can only be attached to verbal predicates which denote a reversible change of state. Like *re-* it is incompatible with unergatives, ditransitives and statives. Finally, *out* is also ambiguous between a verbal prefix *out-* and the preposition *out*, which can occur in compounds (eg. *outlaw*, *outboard*, *outhouse*). Rather than requiring that the verbal constituent it attaches to already contains a DP direct object as *re-* does, *out-* adds a DP argument and establishes a comparison between this additional DP and the subject DP as in (9) (Ahn 2015).

(9) a. Zi ran (*the bus)
      b. Zi outran the bus = Zi ran faster/further than the bus.

---

[19] English also has a *re-* that attaches to a set of latinate bound stems as already discussed above (eg. *receive*, *resume*, *revive*, etc). But note that the *re-* that attaches to bound latinate roots lacks the precise restitutive semantics and the main-stress-bearing properties of the *re-* in *refill*, (to the extent we can say the bound, latinate *re-* has any semantics, it means something more like 'backwards' or 'towards the subject').

As (9) illustrates, *out-* can easily be attached to typically unergative verbs, precisely because they don't already have their internal argument position occupied. So violating the argument structure selection requirements of *out-* involves adding *out-* to obligatorily transitive predicates like *devour* (eg. *Zi outdevoured Ari) Sample materials for the English experiment are in Table 3.

| Prefix	Grammatical	Ungrammatical	
		Category Violations	Argument Structure Violations
re-	refill	recold	relaugh
un-	unbend[20]	unfaith	unthink
out-	outpace	outlarge	outkill

Table 3 Sample materials from Stockall et. al. (in press). Note that the selection of stems for both violation conditions was based on distributional corpus statistics, supplemented by native speaker intuitions about possible uses of the stems. Interested readers are directed to the paper for a fuller discussion of materials selection and norming.

The fact that the English prefixes represent a more heterogeneous set of items, with different category and argument structure restrictions, might lead one to expect more variation in the evoked neural responses, and indeed, that is what we see. However, across all the prefixes, the category violation items evoke a larger response from left posterior temporal sources than argument structure violation items, from 200-300ms, exactly as observed for Greek. The differences emerge in the later OF response, and we return to them below in section 4.

This pair of studies, then, successfully dissociates processing associated with evaluating the c-selectional requirements of a derivational affix from processing associated with its s-selectional requirements. This is expected in a model like DM in which syntactic computations are distinct from semantic computations. The results from these two studies also provide support for a distinction between grammatically contentful function morphology and grammatically underspecified/vacuous roots. Function morphemes like *-simos* or *re-* have a grammar: they are the exponents of a set of morpho-syntactic features that can only occur in a particular syntactic context. For Alexiadou et. al. 2014, for example, the VI *re-* is an aspectual head that serves as a lexical exponent of v°, and bears an accusative Case feature, which is checked via overt movement of the argument of a result predicate to SpecvP, where the root incorporates into *re-*, as in (10, adapted from Alexiadou et al, 2014, item 59), and *re-* has the semantic denotation in (11).

(10)

---

[20] Note that although *un-* only attaches to free stems, many of these stems are not unambiguously, or even predominantly used as verbs when they occur as free stems, so like with *re-*, it is more accurate to say that *un-* requires that the roots it attaches to be compatible with little v.
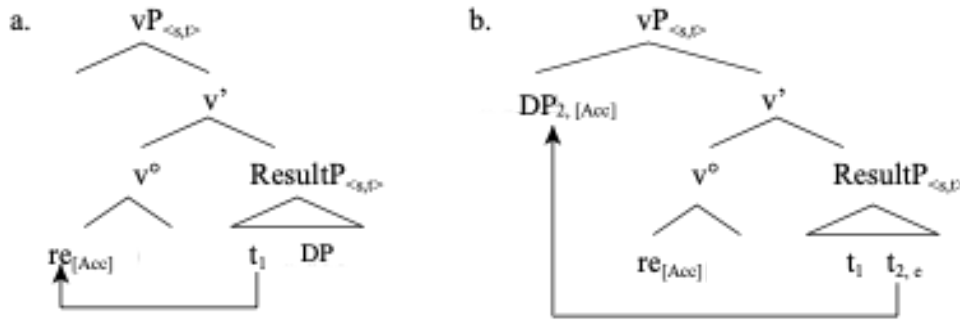
(11) [[re-]] = λP$_{<e,t>}$.λx$_e$.e$_s$Ǝs$_s$[P(x)(s) & CAUSE(s)(e)]: Ǝe'<e & P(x)(e')

From a processing perspective, then, activating the VI for *re-* is extremely informative, and allows the comprehender to generate an entire vP, which then creates an insertion context for the lexical root. If the root is *fill*, it has a syntactic distribution that is consistent with this context, allowing easy integration. If however, the root is *hat* or *idea*, its syntactic distribution is inconsistent with this context, and a processing difficulty ensues. The Neophytou et. al. (2018) and Stockall et. al. (2019) studies detect exactly this processing difficult in the Posterior Temporal Lobe response discussed above.

4. Compositional semantic interpretation of the structures generated in (3)

A small, but growing number of studies have used MEG to investigate processing associated with compositional semantic wellformedness in morphologically complex words. The first such study was Pylkkänen et. al. (2009) who used a sentence reading paradigm to compare the processing of fully grammatical *un-* participles such as *unclogged* and novel *un-* participles that violate *un-*'s **s-selection** restrictions as outlined above, such as *unflushed* (*flush* is not a reversible change of state predicate). They embedded such words in sentence contexts as in (12).

(12) a. ...the toilet was being unclogged…
　　 b. ...the toilet was being unflushed…

Pylkkänen et. al. (2009) found that the semantically ill-formed items such as (12b) evoked significantly more activity than the well-formed items from 290-310ms after the onset of the *un-* prefixed word, in ventro-medial Pre-Frontal Cortex (vmPFC). Critically, Pylkkänen et. al. compared the semantic violation condition with what they called a 'real world knowledge' violation as in (13).

(13) ...the towel was being unclogged…

Sentences such as (13), which are perfectly grammatically well-formed, but incompatible with our normal understanding of what a towel is and what being clogged is, did not evoke any vmPFC cost compared to the grammatical baseline (12a), but instead evoked a processing cost in left Inferior Prefrontal Cortex (LIPC, or Broca's area) between 300-330ms. The vmPFC effect is therefore plausibly an index of the cost of resolving the mismatch between the compositional semantic requirements of *un-* (a reversible change of state event

with an internal, affected, argument) and the semantic denotation typically associated with a verb like *flush* which describes a permanent change of state. Support for this interpretation comes from a series of other sentence reading studies from Pylkkänen and her colleagues which find similar vmPFC costs for complement coercion (*begin the book*) and aspectual coercion (*throughout the day the student sneezed*) (Pylkkänen & McElree, 2007, Brennan & Pylkkänen, 2008, Pylkkänen, Martin, McElree and Smart, 2009).

Fruchter and Marantz (2015) investigated the processing of semantic wellformedness in single word reading by creating a novel variable, **semantic coherence**, which is the ratio between the predicted surface frequency (a function of stem frequency, affix frequency, and biphone transition probability), and the actual corpus frequency of a complex word. Words whose actual surface frequency is as high, or higher than predicted, are considered to have a high semantic coherence, while words with a lower surface frequency than predicted have a low coherence[21]. Fruchter and Marantz found that this semantic coherence variable correlated reliably with evoked activity from the left Orbito-Frontal cortex (OF[22]) between 354-500ms after the onset of a morphologically complex word. Lending support to the association between the OF/vmPFC and compositional semantic well-formedness processing are the s-selection results from Neophytou et. al. (2018) and Stockall et. al. (2019) discussed above. Neophytou et. al. found that words like *gelasimos (laugh-able), which violate the s-selection requirements of their suffix, evoked increased OF activity compared to c-selection violation items between 400-500ms after stimulus onset. Stockall et. al. find the same pattern for their *re-* items: s-selection violations like *relaugh evoked significantly more OF activation than c-selection violations like *rehat from 325 - 440ms. However, they find that the effect goes in the opposite direction for the *out-* items (c-selection violation items like *outlarge evoked more activity than s-selection items like *outmurder), and that there was no significant difference for the *un-* items. Critically, the argument for expecting that items which violate the category selection requirements of their affix should be associated with less OF activation than items which match the c-selection requirements, but violate the s-selection requirements, is along the lines of the Hahne & Friederici (2002) ELAN/N400 results discussed in Section 2. In that study, if an ELAN violation response was observed, no N400 violation response was evoked - detecting the syntactic category error seemed to mean the comprehender just gave up on the sentence and did not attempt semantic interpretation. The temporal and spatial dissociation between the syntactic category violations and the semantic wellformedness violations in the Neophytou et. al. (2018) results, and Stockall et. al. (2019)'s results for *re-* are consistent with this account: only syntactically well-formed structures are semantically interpretable. Under this interpretation of the results, syntactically ill-formed items such as *reidea* or *varelimos* (barrel-able) evoke relatively little activity associated with semantic wellformedness evaluation, as they were not assigned a coherent syntactic representation, and thus there is no LF to interpret at stage (4), Fig1.

---

[21] Fruchter & Marantz's coherence variable is inspired by the work of Riddle (1985), who discusses pairs of words built on the same stem like *brutalness* and *brutality,* and argues that *-ness* carries semantic implications that make it "cohere" better to some stems than to others. Fruchter & Marantz hypothesise that this kind of coherence ought to be reflected in word usage. If, in their example, *acuteness* is less coherent than *awkwardness*, it ought to be used less often, and thus have a lower surface frequency than would be expected.

[22] The OF and vmPFC are adjacent subregions of the Ventral Prefrontal Cortex. Within the domain of language processing, the OF and vmPFC are considered anatomically synonymous (Pylkkänen et al., 2009; Phillips et al., 2002), and the spatial resolution of MEG does not permit strong claims about the functional separation between the two areas.

The fact that the *out-* category violation items evoked larger responses than *out-* semantic violation items in both the earlier temporal lobe response and the later OF response suggests comprehenders continued to try to interpret items such as **outlarge* and **outgrim* (while abandoning any attempt to interpret **rebald* and **redumb*). Any such attempts, however, were unsuccessful, as the category violation items were rejected more than 90% of the time across all 3 prefixes. A possible explanation for the difference is that once comprehenders had detected that *large* was a poor fit for the vP context associated with *out-*, they could reanalyse *out-* as the preposition *out,* and attempt a compound analysis. No such alternative is available for *re-*, or for the Greek *-simos*, *-menos* and *-tos*[23].

Putting aside this last complication, then, we see that across a range of different experimental paradigms (sentence reading, single word reading), and across different operalisations of the notion, compositional semantic wellformedness is associated with activation in the OF/vmPFC that is distinct from activation associated with real world (encyclopedic) knowledge (Pylkkänen et. al. 2009) and syntactic wellformedness (Neophytou et. al. 2018, Stockall et. al. in press), as predicted by the architecture of DM.

5. Wrapping Up

The goal of this chapter has been to try to take DM seriously as a comprehension model, and to review the evidence from neurolinguistic experiments that speaks to the specific predictions this model makes. We have not made any attempt to compare the predictions DM makes to those that a lexicalist morphological theory and/or a projectionist syntactic theory would make, and therefore we have not made the explicit case that these results adjudicate between theories. The assumption is that readers of this handbook are already persuaded that the distributed, constructivist approach of DM is at least broadly correct. Our focus, then, has been on taking a body of work that was written for a cognitive neuroscience audience, with minimal explicit links to formal linguistic theory, and demonstrating how this body of work speaks directly to fundamental questions about the nature of linguistic representations and the relationship of syntactic structure to form and meaning.

Works Cited:

Albright, A., and Hayes, B. (2003). Rules vs. analogy in English past tenses: a computational/experimental study. Cognition 90, 119–161.

Ahn, Byron. March 2015. Out-Sourcing Internal Arguments. Presented at WCCFL 33. Simon Fraser University.

Alexiadou, A. (2018). -Able adjectives and the syntax of psych verbs. Glossa: A Journal of General Linguistics, 3(1), 74.

---

[23] Although Greek does have a nominalising suffix -otos. Many of the stems in the Neophytou et. al. experiment ended in 'o', so that there may have been some parsing ambiguity here. However, the stem final 'o's were written with 'o', while the nominalizer 'otos' is written 'ωτός', so any such ambiguity was purely phonological, not orthographic.

Alexiadou, A., Anagnostopoulou, E., and Lechner, W. (2014). "Variation in repetitive morphemes: some implications for the clausal architecture," Presented at the Workshop on the State of the Art in Comparative Syntax, (York: University of York).

Anagnostopoulou, E., & Samioti, Y. (2014). Domains within words and their meanings: A case study. The syntax of roots and the roots of syntax, 81-111.,

Baayen, R. H., Piepenbrock, R., and Gulikers, L. (1995). The CELEX Lexical Database (Release 2) [CD ROM]. Philadelphia, PA : Linguistic Data Consortium, University of Pennsylvania.

Baayen, R. H., Wurm, L. H., and Aycock, J. (2007). Lexical dynamics for low-frequency complex words: a regression study across tasks and modalities. Mental Lexicon 2,

Beretta, A., Fiorentino, R., & Poeppel, D. (2005). The effects of homonymy and polysemy on lexical access: An MEG study. Cognitive Brain Research, 24(1), 57-65.

Blevins, J.P. (2003) Stems and paradigms. Language 79, 737–767

Bozic, M., Fonteneau, E., Su, L., & Marslen‑Wilson, W. D. (2015). Grammatical analysis as a distributed neurobiological function. Human brain mapping, 36(3), 1190-1201.

Brennan, J., & Pylkkänen, L. (2008). Processing events: Behavioral and neuromagnetic correlates of aspectual coercion. Brain and language, 106(2), 132-143.

Bozic, M., Szlachta, Z., & Marslen-Wilson, W. D. (2013). Cross-linguistic parallels in processing derivational morphology: evidence from Polish. Brain and language, 127(3), 533-538.

Carrier, J. L. (1979). The interaction of morphological and phonol- ogical rules in Tagalog: A study in the relationship between rule components in grammar [Unpublished Doctoral dissertation]. Massachusetts Institute of Technology.

De Saussure, F. (1916/1966). Course in general linguistics. New York: McGraw.

Dikker, S., Rabagliati, H., Pylkkänen, L., and Pylkkänen, L. (2009). Sensitivity to syntax in visual cortex. Cognition 110, 293–321.

Embick, D. and D. Poeppel (2015) 'Towards a computational(ist) neurobiology of language; Correlational, integrated, and explanatory neurolinguistics' Language, Cognition, and Neuroscience 30:4, 357-366.

Embick, D., Hackl, M., Schaeffer, J., Kelepir, M., & Marantz, A. (2001). A magnetoencephalographic component whose latency reflects lexical frequency. Cognitive brain research, 10(3), 345-348.

Ettinger, A., Linzen, T., & Marantz, A. (2014). The role of morphology in phoneme prediction: Evidence from MEG. Brain and Language, 129, 14-23.

Flick, G., Oseki, Y., Kaczmarek, A. R., Al Kaabi, M., Marantz, A., & Pylkkänen, L. (2018). Building words and phrases in the left temporal lobe. *Cortex, 106*, 213–236.

Flick, G. & Pylkkänen, L. (2020). Isolating syntax in natural language processing: MEG evidence for an early contribution of left posterior temporal cortex. Cortex, 127, 42-57.

Fiorentino, R., & Poeppel, D. (2007). Compound words and structure in the lexicon. Language and Cognitive processes, 22(7), 953-1000.

Friederici, A. D., Wang, Y., Herrmann, C. S., Maess, B., & Oertel, U. (2000). Localization of early syntactic processes in frontal and temporal cortical areas: a magnetoencephalographic study. Human brain mapping, 11(1), 1-11.

Forster, K.I., & Davis, C. (1984). Repetition priming and frequency attenuation in lexical access. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10, 680-698.

Fruchter, J., & Marantz, A. (2015). Decomposition, lookup, and recombination: MEG evidence for the full decomposition model of com- plex visual word recognition. Brain and Language, 143, 81–96.

Fruchter J, Stockall L and Marantz A (2013) MEG masked priming evidence for form-based decomposition of irregular verbs. Frontiers in Human Neuroscience. 7:798.

Gwilliams, L. (2019). How the brain composes morphemes into meaning. Philosophical Transactions of the Royal Society B.

Gwilliams, L., & Davis, M. H. (2022). Extracting language content from speech sounds: the information theoretic approach. In *Speech Perception* (pp. 113-139). Springer, Cham.

Gwilliams, L., King, JR, Marantz, A. & Poeppel, D. (2022) Neural dynamics of phoneme sequences reveal a position-invariant code for content and order. Nature Communications.

Gwilliams, L., Lewis, G. A., & Marantz, A. (2016). Functional characterisation of letter-specific responses in time, space and current polarity using magnetoencephalography. NeuroImage, 132, 320–333.

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 and Language.

Gwilliams, L., & Marantz, A. (2018). Morphological representations are extrapolated from morpho-syntactic rules. Neuropsychologia, 114, 77–87.

Hahne, A. and Friederici, A.D. (1998)_ERP evidence for autonomous first-pass parsing processes in auditory language comprehension. J. Cogn. Neurosci. (Suppl.) S125

Hahne, A. and Friederici, A.D. (2002) Differential task effects on semantic and syntactic processes as revealed by ERPs. Cogn. Brain Res.

Harley, H., & Noyer, R. (1999). Distributed morphology. Glot international, 4(4), 3-9.

Hay, J. (2001). Lexical frequency in morphology: is everything relative? Linguistics 39, 1041–1070.

Hay, J., & Baayen, H. (2002). Affix Productivity and Base Productivity. In Morphological Productivity Seminar, ESSE (Vol. 6).

Horn, L. (2005). An un-paper for the unsyntactician. in Polymorphous Linguistics: Jim McCawley's Legacy, eds S. Mufwene, E. Francis, and R. Wheeler (Cambridge, MA: MIT Press), 329–365.

King, J., Linzen, T., and Marantz, A. (2015). Syntactic categories as lexical features or syntactic heads: a MEG approach. lingbuzz/002477

Laine, M., Vainio, S., & Hyönä, J. (1999b). Lexical access routes to nouns in a morphologically rich language. Journal of Memory and Language, 40, 109–135.

Lau, E. F., Phillips, C., and Poeppel, D. (2008). A cortical network for semantics:(de) constructing the N400. Nat. Rev. Neuroscience. 9:920.

Lavric, A., Clapp, A., & Rastle, K. (2007). ERP evidence of morpho- logical analysis from orthography: A masked priming study. Journal of Cognitive Neuroscience, 19(5), 866–877.

Lehtonen, M., Monahan, P. J., and Poeppel, D. (2011). Evidence for early morphological decomposition: combining masked priming with magnetoencephalography. J. Cognitve Neuroscience. 23, 3366–3379.

Levin, B. (1993). English Verb Classes and Alternations: A Preliminary Investigation. Chicago, IL: University of Chicago Press.

Lewis, G., Solomyak, O., & Marantz, A. (2011). The neural basis of obligatory decomposition of suffixed words. Brain and Language, 118(3), 118–127.

Linzen, T., Marantz, A., and Pylkkänen, L. (2013). Syntactic context effects in visual word recognition: an MEG study. Mental Lexicon 8, 117–139.

Longe, O., Randall, B., Stamatakis, E. A., & Tyler, L. K. (2006). Grammatical categories in the brain: The role of morphological structure. Cerebral Cortex, 17(8), 1812-1820.

Manouilidou, C. (2007). "Thematic constraints in deverbal word formation: psycholinguistic evidence from pseudo-words," in Proceedings of the 7th International Conference on Greek Linguistics, (York: University of York).

Manouilidou, C., Dolenc, B., Marvin, T., and Pirtošek, Z. (2016). Processing complex pseudo-words in mild cognitive impairment: the interaction of preserved morphological rule knowledge with compromised cognitive ability. Clin. Linguist. Phonet. 30, 49–67.

Manouilidou, C., and Stockall, L. (2014). Teasing apart syntactic category vs. argument structure information in deverbal word formation: a comparative psycholinguistic study. Italian J. Linguist. 2, 71–98.

Marantz, A. (2007). Restitutive re and the First Phase Syntax/Semantics of the VP. Ms. NYU.

Marantz, A. (2009). Roots, re-, and affected agents: can roots pull the agent under little v? Paper Presented at Stuttgart Roots Workshop, Stuttgart.

Wood, J., & Marantz, A. (2017). The interpretation of external arguments. The verbal domain, 255-278.

Marr, D. (1982, 2010). Vision: A computational investigation into the human representation and processing of visual information. MIT press.

Matar, S., Dirani, J., Marantz. A., & Pylkkänen, L. (2021). Left posterior temporal cortex is sensitive to syntax within conceptually matched Arabic expressions. *Scientific Reports 11*, 7181.

McCandliss, B. D., Cohen, L., & Dehaene, S. (2003). The visual word form area: expertise for reading in the fusiform gyrus. Trends in cognitive sciences, 7(7), 293-299.

McCarthy, J. J., & Prince, A. (1995). Faithfulness and reduplicative identity. Papers in Optimality Theory, 10. Graduate Linguistic Student Association, Dept. of Linguistics, University of Massachusetts.

McDonald, S. A., & Shillcock, R. C. (2001a). Contextual distinctiveness: A new lexical property computed from large corpora. Informatics Research Report EDI-INF-RR-0042, University of Edinburgh.

Morris, J., Grainger, J., and Holcomb, P. J. (2008). An electrophys- iological investigation of early effects of masked morphological priming. Lang. Cogn. Process. 23, 1021–1056.

Morris, J., Grainger, J., & Holcomb, P. J. (2008). An electrophysio- logical investigation of early effects of masked morphological priming. Language and Cognitive Processes, 23(7–8), 1021–1056.

Morris, J., and Stockall, L. (2012). Early, equivalent ERP masked priming effects for regular and irregular morphology. Brain Lang. 123, 81–93.

Moscoso del Prado Martin, F., Bertram, R., Häikiö, T., Schreuder, R., & Baayen, R.H. (2004). Morphological family size in a morphologically rich language: The case of Finnish compared to Dutch and Hebrew. Journal of Experimental Psychology: Learning, Memory, and Cognition, 30, 1271-1278.

New, B., Brysbaert, M., Segui, J., Ferrand, L., & Rastle, K. (2004). The processing of singular and plural nouns in French and English. Journal of Memory and Language, 51(4), 568-585.

Neophytou, K., Manouilidou, C., Stockall, L., & Marantz, A. (2018). Syntactic and semantic restrictions on morphological ecompose- tion: MEG evidence from Greek. Brain and Language, 183, 11–20.

Poeppel, D., Emmorey, K., Hickok, G., & Pylkkänen, L. (2012). Towards a new neurobiology of language. Journal of Neuroscience, 32(41), 14125-14131.

Pylkkänen, L., Stringfellow, A., Flagg, E., & Marantz, A. (2000). A neural response sensitive to repetition and phonotactic probability: MEG investigations of lexical access. In Proceedings of biomag (pp. 363-367).

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(3), B35-B45.

Pylkkänen, L., Llinás, R., & Murphy, G. L. (2006). The representation of polysemy: MEG evidence. Journal of cognitive neuroscience, 18(1), 97-109.

Pylkkänen, L., and Marantz, A. (2003). Tracking the time course of word recognition with MEG. Trends in Cognitive Science. 7, 187–189.

Pylkkänen, L., & McElree, B. (2007). An MEG study of silent meaning. Journal of Cognitive Neuroscience, 19(11), 1905-1921.

Pylkkänen, L., Martin, A. E., McElree, B., & Smart, A. (2009). The anterior midline field: Coercion or decision making?.Brain and Language, 108(3), 184-190.

Pylkkänen, L., Oliveri, B., and Smart, A. (2009). Semantics vs. World knowledge in prefrontal cortex. Lang. Cogn. Process. 24, 1313–1334.

Pylkkänen, L. Stringfellow, A. and Marantz, A. (2002) Neuromagnetic evidence for the timing of lexical activation: an MEG component sensitive to phonotactic probability but not to neighborhood density. Brain Lang. 81, 666–678

Rastle, K., & Davis, M. H. (2008). Morphological decomposition based on the analysis of orthography. Language and cognitive processes, 23(7-8), 942-971.

Riddle, E. M. (1985). A historical perspective on the productivity of the suffixes - ness and -ity. In J. Fisiak (Ed.), Historical semantics, historical word-formation (pp. 435–462). Berlin: Mouton Publishers.

Royle, P., Drury, J. E., Bourguignon, N., and Steinhauer, K. (2012). The temporal dynamics of inflected word recognition: a masked ERP priming study of French verbs. Neuropsychologia 50, 3542–3553.

Schuler, K., Yang, C., & Newport, E. (2021). Testing the Tolerance Principle: Children form productive rules when it is more computationally efficient.

Sharpe, V., Reddigari, S., Pylkkänen, L., and Marantz, A. (2019). Automatic access to verb continuations on the lexical and categorical levels: evidence from MEG. Lang. Cogn. Neurosci. 34, 137–150.

Simon, D., Lewis, G., and Marantz, A. (2012). Disambiguating form and lexical frequency effects in MEG responses using homonyms. Lang. Cogn. Process. 27, 275–287.

Sprouse, Jon & Ellen Lau. 2013. Syntax and the brain. The Cambridge Handbook of Generative Syntax. Edited by Marcel den Dikken, Cambridge University Press.

Solomyak, O., and Marantz, A. (2010). Evidence for early morphological decom- position in visual word recognition. J. Cogn. Neurosci. 22, 2042–2057.

Steinhauer, K., & Drury, J. E. (2012). On the early left-anterior negativity (ELAN) in syntax studies. Brain and language, 120(2), 135-162.

Stockall, L., Manouilidou, C., Gwilliams, L., Neophytou, K., & Marantz, A. (2019). Prefix stripping re-re-revisited: MEG investi- gations of morphological decomposition and recomposition. Frontiers in Psychology, 10, Article 1964.

Tarkiainen, A., Helenius, P., Hansen, P. C., Cornelissen, P. L., and Salmelin, R. (1999). Dynamics of letter string perception in the human occipitotemporal cortex. Brain 122, 2119–2132.

Tremblay P, Dick AS. Broca and Wernicke are dead, or moving past the classic model of language neurobiology. Brain Lang. 2016 Nov;162:60-71.

Vannest, J., Polk, T. A., & Lewis, R. L. (2005). Dual-route processing of complex words: new fMRI evidence from derivational suffixation. Cognitive, Behavioral and Affective Neuroscience, 5, 67–76.

Whiting, C. M., Marslen-Wilson, W. D., & Shtyrov, Y. (2013). Neural dynamics of inflectional and derivational processing in spoken word comprehension: laterality and automaticity. Frontiers in human neuroscience, 7, 759.

Wurm, L. H., Ernestus, M. T., Schreuder, R., & Baayen, R. H. (2006). Dynamics of the auditory comprehension of prefixed words: Cohort entropies and conditional root uniqueness points. The Mental Lexicon, 1(1), 125-146.

Wray, S., Stockall, L., & Marantz, A. (2022). Early form-based morphological decomposition in Tagalog: MEG evidence from reduplication, infixation, and circumfixation. Neurobiology of Language, 3(2), 235–255.

Yang, C. D. (2002). Knowledge and Learning in Natural Language. Oxford, New York: Oxford University Press.

Yang (2005) On productivity. Yearbook of Language Variation, 5, 333-370.

Zuraw, K. (2002). Aggressive reduplication. Phonology, 19(3), 395–439.

Zweig, E., & Pylkkänen, L. (2009). A visual M170 effect of morpho- logical complexity. Language and Cognitive Processes, 24(3), 412–439.