Paradigms in Morphology

Petar Milin
University of Birmingham

James P. Blevins
University of Cambridge

To appear in *The Oxford Encylopedia of Morphology*, ed. by R. Lieber, S. Arndt-Lappe, A. Fàbregas, C. Gagné, and F. Masini. Oxford: Oxford University Press.

1. Morphological paradigms and grammatical organization

Inflectional paradigms have been central to models of morphological organization and change from the inception of the Western grammatical tradition. The oldest extant grammar in this tradition, the *Téchnē Grammatiké* (conventionally attributed to Dionysus Thrax — but see Law & Sluiter 1995), consists largely of a treatment of morphology under the title "An accurate account of analogies" (Davidson 1874: 3). Paradigmatic analogy in turn guided the 'inflection' or 'bending' of basic forms in the Priscianic tradition (Matthews 1994, 1991). Inflectional paradigms remained the primary locus of the analogical principles formulated within the Neogrammarian tradition (Morpurgo Davies 1998), and they preserve this role within contemporary models of morphological change (Fertig 2013). This consensus about the importance of paradigms stands in contrast to shifting assumptions about syntagmatic structure and units within this tradition.

The emergence of a paradigm-based perspective can be attributed in part to the flectional structure of classical languages and to the pedagogical goals of early grammatical descriptions. However, the subsequent development of this perspective owes more to the fact that inflectional paradigms exhibit properties that can be efficiently exploited by analogical principles. The high degree of interdependence exhibited by forms within inflectional paradigms is particularly favourable to the application of processes of analogical generalization. Analogical effects have been shown to influence a range of morphological phenomena, including the assignment of compound stress (Plag 2010) and the selection of linking elements (Krott et al. 2002), and have been claimed to condition patterns in wider linguistic – and cognitive – domains (Gentner et al. 2001; Itkonen 2005). Yet the descriptive potential of analogical principles is greatest within the relatively closed and uniform space of variation defined by inflectional paradigms and inflection classes.

In the modern period, a number of independent lines of research have converged on perspectives that make essential reference to paradigms in the analysis of morphological structure and use. One line of research grew out of attempts to measure the processing load associated with inflectional exponents (Kostić 1991, 1995). A largely complementary approach arose from studies

of system-level effects on the processing of individual words (Schreuder & Baayen 1997; Baayen et al. 2008). In the unification of these approaches (Moscoso del Prado Martín et al. 2004; Milin et al. 2009b), paradigms came to play a vital role in mediating between items and classes. The formal techniques developed to measure processing load inspired strategies for quantifying patterns of interpredictabilty within and across paradigms (Ackerman et al. 2009, Ackerman & Malouf 2103). The critical response to this literature in turn placed the study of paradigmatic structure within a larger investigation of notions of morphological complexity (Juola 1998, Bonami 2014, Pirrelli et al. 2015, Sagot 2018). These converged lines of inquiry fed back into the processing literature, as the goal of understanding paradigmatic effects on processing and the influence of implicational patterns on systems led in the direction of mechanisms of cognitively-plausible learning models (Baayen et al 2011; Ramscar et al. 2013).

The resurgence of interest in the paradigmatic structure of morphological systems has occurred against the backdrop of an emerging consensus about the probabilistic nature of language. Viewed from this perspective, paradigmatic structure provides a functional organization of the lexicon into patterns that permit speakers to extrapolate beyond the sparse and biased input that they encounter (Blevins et al. 2017). The informativity of these patterns helps to explain why speakers appear to show a high degree of sensitivity to paradigmatic structure, and, most strikingly, to systematic divergences between the distributional structure of paradigms and classes (Milin et al. 2009a).

The following sections trace the evolution of notions of morphological paradigms from their classical origins to contemporary formalizations. A longitudinal perspective on paradigms highlights the continuity within paradigm-based approaches, while identifying how this tradition has advanced by isolating idealizations and exploiting analytic techniques that allow it to address successively more fundamental questions.

2. The linguistic status of morphological paradigms

Much like the role of morphology itself, the place of paradigms in different frameworks of analysis tends to ebb and flow in ways that reflect larger assumptions about the structure and organization of linguistic systems. The exclusion of paradigms from descriptivist and generativist approaches is something of a historical anomaly, reflecting the Post-Bloomfieldian hypothesis that a system can be fully described in terms of minimal parts and combinatoric rules. This 'inventories and arrangements' perspective is implicit in the notion of "scientific compactness" in Bloomfield (1933: 238) and set out clearly by Hockett (1958):

> [E]very language has its own grammar. The grammar, or grammatical system, of a language is (1) the morphemes used in the language, and (2) the arrangements in which these morphemes occur relative to each other in utterances. (Hockett 1958: 129)

The omission of paradigms is not the outcome of a consideration of the status of paradigms *per se* but follows from the general goal of reducing all properties of a grammatical system to the syntagmatic features of individual expressions. Even expressions have no persistent status in a Post-Bloomfieldian model. Instead, they serve as 'inputs' to analytic procedures that extract recurrent units and patterns of arrangement, and as 'outputs' of constructive procedures that reverse this process, reconstituting 'surface' forms from inventories of abstract units and principles of arrangement. A radically reductionist model of this nature has no place for sub-systems such as paradigms or classes, let alone for system-level relations between their elements.

2.1. Classical paradigmatic approaches

In contrast, words and inflectional paradigms are both persistent components of approaches that conform to the classical 'word and paradigm' (WP) model (Hockett 1954). In these approaches, individual forms display patterns of arrangement, and sets of forms exhibit systematic patterns of cross-form similarity and contrast. The lexicon of a language is not conceived of as an inventory of atomic parts but as a 'word stock' of surface forms, organized into classes based on observable dimensions of variation, principally variation in shape, arrangement and distribution. The resulting conception treats a language as a complex system, containing parts whose meanings and functions depend essentially on contrasts and interactions involving other elements of the system.

Morphotactic structure can, at least in principle, be assigned to individual word forms on inventory- and system-based approaches. The key difference lies in the **interpretation** of this structure. In a classical WP model, morphotactic structure is not treated as an arrangement of independently meaningful parts but as a structural pattern that contributes to discriminating individual forms and to predicting systematic variation within sets of forms. Individual elements of a morphotactic structure have no status as persistent 'units' in this framework, and there is no requirement even that distinct analogical processes must abstract the same structure from sets of forms.

Significantly, it is the combination of **mutable** morphotactic structure and **stable** paradigmatic organization that accounts in large part for the descriptive success of the proportional analogies developed in the Neogrammarian period. In elucidating the representational agnosticism of Hermann Paul, Morpurgo Davies (1998) stresses that the 'structural analyses' assigned by proportions are intended to provide a task-specific notion of 'matching', one which carries no implications about the uniformity or persistence of matched segments:

> It is worth noticing that the proportion is a neutral form of notation, i.e. it may be interpreted in cognitive or structural terms and may or may not be reformulated in terms of morphological or morphophonemic rules. At the same time, it also offered an algorithm for a structurally based form of morphological segmentation, without making any claims about the segments in question. (Morpurgo Davies 1998: 258)

Establishing words and paradigms as persistent elements of a classical WP model in turn shifts the primary part-whole relations in the WP model to the affiliation of words with paradigms. The idea of 'paradigms as units' is stated succinctly by Matthews (1991):

> In the ancient model the primary insight is not that words can be split into roots and formatives, but that they can located in paradigms. They are not whole composed of simple parts, but are themselves the parts within a complex whole. In that way, we discover different kinds of relation, and, perhaps, a different kind of simplicity. (Matthews 1991: 204)

In addition, paradigms and paradigmatic relations underpin the creative use of language by defining stable domains over which processes of abstraction, discrimination and extrapolation apply. In codifications of the "ancient model", inflection classes provide the patterns that guide the formation of 'oblique' from 'basic' forms. As Hovdhaugen (1996) notes, inflectional 'rules' in this model express relations between pairs of attested forms, without treating the constant or variable portions of those forms as independent units:

> Priscian operates with a system of chains of rules where he starts with one basic form (e.g., nominative) and from that form derives another (e.g., genitive) which then forms the basis for the formation of the dative, and so on. (Hovdhaugen 1996: 117)

The role that paradigms play in the analogical deduction of forms is set out explicitly by Paul (1920) when he describes the acquisition of a foreign language as a process of 'learning a number of paradigms by heart' and 'memorizing only as many forms of individual words as is necessary to recognize their membership in this or that paradigm':

> Besonders klar sehen wir die Wirkungen der Analogie bei der grammatischen Aneignung der Flexionsformen einer fremden Sprache. Man lernt eine Anzahl von Paradigmen auswendig und prägt sich dann von den einzelnen Wörtern nur soviel Formen ein, als erforderlich sind, um die Zugehörigkeit zu diesem oder jenem Paradigma zu erkennen. (Paul 1920: 112)

For Paul, the use of paradigms to guide analogical generalization is not unique to foreign language acquisition but is a general process that is merely more perspicuous in this case.

A classical paradigmatic perspective survives into the modern era in analyses that describe morphological structure in terms of implicational rather than derivational relations. The 'Paradigm Structure Constraints' proposed by Wurzel (1984) in the context of Natural Morphology provide the most transparent bridge between classical and contemporary WP models. As in models of Priscian and Paul, it is inflectional paradigms that define the primary domain over which implicative patterns apply:

> [I]mplicative relations … exist throughout the whole paradigm: all paradigms (apart from suppletive cases) are structured on the basis of implicative patterns which go beyond the individual word, patterns of varying complexity. (Wurzel 1989: 208)

2.2. Realizational extensions

Strikingly, (surface) words, inflectional paradigms and paradigmatic relations play almost no role in the earliest efforts to reimplement the classical WP model in a modern setting. For the most part, the 'extensions' proposed within the family of 'extended word and paradigm' models followed the lead of Robins (1959) in restoring the grammatical word as the locus of meaning but attached no particular significance to word forms or sets of word forms. Instead, the central innovation in realizational tradition initiated by Matthews (1965) involved the separation of the feature bundles that expressed the 'grammatical meanning' associated to paradigm cells from the interpretive rules that specified the formal 'spell-out' of features. By encapsulating formatives in these rules, realizational approaches largely eliminated the inflectional 'morphemes' of Post-Bloomfieldian accounts. However, Matthews (1972, 1991) retains sub-word 'stems', which amount for all intents and purposes to lexical morphemes. Subsequent accounts in this tradition, notably Anderson (1992) and Aronoff (1994), make more essential use of a stem-based lexicon, and models of Paradigm Function Morphology (Stump 2001) go further in this direction by adopting a root-based perspective. In no variety of realizational approach are 'surface' word forms treated as persistent units. In most, morphotactic structure remains implicit in the 'derivational history' of a realized form; surface outputs exhibit no structure above the level of phonological elements.

Paradigms and paradigmatic relations are even more clearly deprecated in the realizational tradition. Without persistent word forms there can be no word-based lexicon, much less an organization of words into persistent paradigms. Instead, paradigms are reinterpreted either as abstract sets of feature bundles that correspond to paradigm cells or as virtual outputs obtained by applying a set of inflectional rules to a collection of abstract feature bundles. The second of these interpretations is articulated in Anderson (1992):

> We can now define the **paradigm** of a lexical item in terms of its lexical stem set: an item's paradigm is the complete set of surface word forms that can be projected from the members of its stem set by means of the inflectional Word Formation Rules of the language (Anderson 1992: 134)

This set of projected forms is virtual in the sense that it is the theoretical output of a set of rules and lexical bases. It is the rules and bases that are persistent in this model. The surface word forms defined by individual derivations have no persistent status, and, as in other realizational models, there is no provision for caching the output of independent derivations.

In part, the deprecation of paradigms reflects the influence of the syntagmatic focus of Post-Bloomfieldian approaches. Paradigmatic relations in a realizational model mainly regulate the

application of rules that 'compete' to define the spell-out of features at a given syntagmatic point. Although some realizational approaches follow Zwicky (1985) in expressing 'referral' relations between syncretic forms, a model without a lexicon of persistent word forms can only define referrals between rules or some other persistent element of a grammar.

Moreover, even the persistent status of rules and abstract stems is a symptom of the priority assigned to abstract grammars over surface expressions in Post-Bloomfieldian approaches. The status of grammars derives to some degree from the idealization of syntactic systems as infinite. Particularly in the generative tradition, this has been taken to support extensional rather than intensional characterizations of languages (i.e., via abstract enumerations rather than attestations or sample). Early generative accounts approached morphology in fundamentally the same way:

> Let us now consider various ways of describing the morphemic structure of languages. We ask what kind of grammar is necessary to generate all the sequences of morphemes (or words) that constitute grammatical English sentences, and only these. (Chomsky 1957: 18)

Whatever the merits of idealizing syntax as an infinite system (see, e.g., Pullum & Scholz 2010), justification for this idealization is entirely lacking in the domain of morphology. For all intents and purposes, inflectional systems are finite, with a finite number of form variants distributed over a closed, uniform space of variation. Although derivational patterns are less constrained and uniform than inflections, these features ultimately undermine rather than support the idealization of morphology as an infinite system. In principle, indefinitely many forms can be 'projected' from the elements of a derivational system. Nonetheless, the greater idiosyncrasy of derivational strategies entails that the meaning and function of potential neologisms cannot be reliably predicted until the form is actually coined and established. Hence the space of potential coinages has no direct relevance to an extant morphological system and no impact on a speaker. The studies of morphological family effects (see Mulder et al. 2014, and references cited therein) strongly reinforce this point. The morphological family literature demonstrates that speakers are acutely sensitive to family size, not solely to the wellformedness of individual members. Forms that are merely potential members fall outside the system and are as irrelevant as potential borrowings or other lexical coinages.

2.3. Modern paradigmatic perspectives

Contemporary attempts to rehabilitate a classical paradigmatic approach also trace their origins to the work of Hockett (1954), Robins (1959) and Matthews (1972, 1991). However, modern WP approaches develop the informal characterizations of classical models in these works. The most immediate obstacle to formalizing the insights of a classical model was the lack of applicable notions of 'variation' and 'implicational structure'. Classical WP approaches provided detailed case studies of systems that exhibited different types of variation and structural correspondences. Exceptionless patterns were sometimes characterized in terms of logical devices like material im-

plication (Wurzel 1989). But linguistic and philological traditions did not produce any general means of measuring variation or implicational structure, and efforts to borrow formal strategies from Post-Bloomfieldian approaches tended to subvert the goal of formalizing the classical model.

Ultimately, formalizations of classical conceptions of variation and structure arose as a by-product of psycholinguistic studies devoted to measuring system-level effects on processing. As the influence of paradigmatic organization on processing became clearer, techniques were developed to measure the influence of paradigms and classes on the comprehension and production of individual words. The same techniques offered a means of quantifying the interdependence of elements within a system and the general cohesion of the system.

3. Psychological and computational perspectives

Morphological paradigms had also attracted the attention of psychologists working on language from the dawn of the modern psycholinguistic era, in part to overcome the limitations imposed by focussing on the internal structure of individual words. Early psychological studies of language laid the groundwork by developing a wide range of language-specific experimental tasks and techniques for **presenting stimuli** such as words, phrases, or sentences (in isolation or in controlled contexts) and for **registering responses** of participants to these stimuli. Much of this methodological apparatus had been carried over from what was occasionally characterised as "the 'dark ages' or pre-Chomskyan investigations of verbal responses" (Greene 1972: 11). However, the development and deployment of these methodologies faced new challenges created by an intervening theoretical paradigm shift, as psycholinguistics struggled to integrate experimental results into linguistic models that placed them squarely on the performance side of a competence/performance divide (Chomsky 1965).

One of the central challenges that psycholinguistics addressed from the outset was a description of the organisational principles of what came to be called the 'mental lexicon'. Available linguistic descriptions of the mental lexicon tended to be unhelpfully vague, often explicitly acknowledging the metaphorical nature of the term. The characterizations offered in early generative accounts largely echoed the Bloomfieldian claim that "[t]he lexicon is really an appendix of the grammar, a list of basic irregularities" (Bloomfield 1933: 274), adding little more than a re-encoding of this perspective in terms of 'features' and 'rules':

> The lexicon consists of an unordered set of lexical entries and certain redundancy rules. Each lexical entry is a set of features … [T]he lexical entries constitute the full set of irregularities in a language. (Chomsky 1965: 153)

Imprecise linguistic notions of this kind could not provide a secure foundation for psychological approaches, which required constructs that were articulated clearly enough to permit operationalization in experimental studies. The challenges faced by experimental psycholinguists were further compounded by the steady expansion of types of 'mental representations' and the

proliferation of dimensions of representations, many of which were posited on the basis of a priori theoretical speculations.

Psycholinguistic studies were also impacted by the move from the corpus-based perspective of descriptivist models to the grammar-based orientation of generative accounts, as this reorientation was accompanied by a shift from the study of properties of attested expressions to the study of the 'wellformedness' of potential expressions. These changes devalued simple verbal responses (cf., Skinner 1957), along with the more sophisticated mediation processes involved in learning (Osgood & Sebeok 1954; Osgood et al. 1954; Osgood, 1966). Even more radically, the focus on wellformedness had the effect of excluding all distributional considerations. As a consequence, information theory (Shannon 1948), which had been embraced as an explanatory model for a range of psychological phenomena with almost uncritical enthusiasm (Shannon 1956), was abruptly disregarded as wholly irrelevant to the study of language. Information theory relied on the notion of probability, specifically the probability of a message being transmitted. This probabilistic orientation was irreconcilable with generative approaches, which endorsed the idea that "the notion 'probability of a sentence' is an entirely useless one, under any known interpretation of this term" (Chomsky 1969: 57).

The obstacles to psycholinguistic research provoked a re-evaluation of the theoretical assumptions that appeared to hamper progress. The syntagmatic orientation of Post-Bloomfieldian models was recognized as a particularly severe liability, since this focus excluded patterns of paradigmatic variation and competition that might be relevant to language processing. By jettisoning this restriction and integrating dimensions of paradigmatic structure, it became possible to use the clustering of word forms into classes as a window into the organizational principles of the mental lexicon. Experimental studies of processing effects soon established that paradigms and classes do induce measurable differences in behavioural responses, results subsequently replicated in neural studies. These robust effects provided evidence that the mental lexicon obeys specific organisational principles — those of a complex system. A conception of the lexical and morphological components as complex systems offered an explanation for the observation that speakers, when encountering an individual word, simultaneously activate its paradigmatic relatives.

Initially, the lines of psycholinguistic research that investigated paradigmatic effects evolved independently of prevailing theoretical approaches, which were in any case more concerned with syntax (and phonology) than with morphology or morphological paradigms. Independence from theoretical strictures created an opportunity to formulate and evaluate models that treated various frequency counts — essentially, probability estimates — as quantifiable properties of linguistic units. The influential studies of Taft (1979) and Colé et al. (1989) showed that a word's stem frequency (i.e., the sum of frequency counts of all of the inflected variants of that word) had significant power in predicting response latencies in a lexical decision task. By exploring the effects of a wider range of distributional factors on simple verbal responses, subsequent studies were able to probe successively more subtle and intricate properties of the mental lexicon.

The role of inflectional paradigms in language processing was, however, almost immediately contested by studies that approached language as a simple system whose properties were wholly determined by the properties of individual forms. These studies focussed primarily on languages with more intricate patterns of inflectional variation than those found in English (and investigated in early psycholinguistic accounts). One set of studies focussed on Italian (Burani et al. 1984; Caramazza et al. 1988), while another examined Serbian (Lukatela et al. 1980). During the 1990s and early 2000s this line of research was particularly active and extended to a number of additional languages, notably German (Clahsen 1999; Janssen & Penke 2002, among others), and Finnish (Hyönä et al. 1995; Bertram et al. 2000).

3.1. Classes and exponents

The study of paradigmatic effects on language processing provided the outline of a solution to the challenges faced by classical models of morphology. The components of this solution arose, largely independently, in a pair of psycholinguistic traditions. One tradition, associated with the Experimental Psychology Lab at the University of Belgrade (Kostić 1991, 1995; Kostić et al. 2003), sought to measure the 'information load' incurred in the processing of inflectional exponents. A second tradition, based largely at the Max Planck Institute in Nijmegen (Schreuder & Baayen 1997; Baayen et al. 2008), investigated system dependencies that arose in the processing of individual items. The focus of the two traditions reflected in part the languages that formed the initial basis of their investigations. The pioneering work on computational psycholinguistics conducted by the Belgrade School exploited the rich corpus resources available for Serbian. The prominent role that inflection class morphology plays in a Slavic language like Serbian naturally suggested analyses of the functional load attributable to the exponents that mark grammatical and class contrasts. A focus on inflectionally impoverished West Germanic languages likewise encouraged analysis of individual items within the Nijmegen School.

The salient typological differences between Slavic and West Germanic led the two traditions to approach broadly similar problems from complementary directions. The Belgrade School developed 'top-down' measures of informativity that compared distributional properties of classes with those of inflectional exponents. These measures were grounded in a notion of 'amount of information', I_e of an exponent, e, which stands in a simple inverse relationship to the probability of that exponent in its class π. The formula in Figure 1 expresses the relationship between the amount of information expressed by an element and its probability in a given set — here an inflectional exponent in an inflectional class — where inversion is expressed as negative log-transformed value of the probability. Put simply, the less frequently an element occurs, the more informative it is when the element does occur.

$$I_e = -log_2 \mathrm{Pr}_\pi(e)$$

Figure 1: The 'surprisal' of an inflectional exponent e in a class π

This measure of information is intrinsically paradigmatic; it uses the distribution of exponents in an inflectional paradigm to convert the raw frequencies of individual exponents to estimated

probabilities, given the summed frequency of all of the exponents of that paradigm (i.e., stem frequency). That is, information load is not a property of an exponent in isolation but is defined relative to its paradigmatic opponents. This measure is only meaningful if exponents show different distributions in the paradigms of items that belong to a common class. Interestingly, this quantified class informativity remains constant for all nouns belonging to the same class, as represented by the dashed lines in Figure 2 below (which are associated with representative feminine *a*-stem nouns in Serbian).

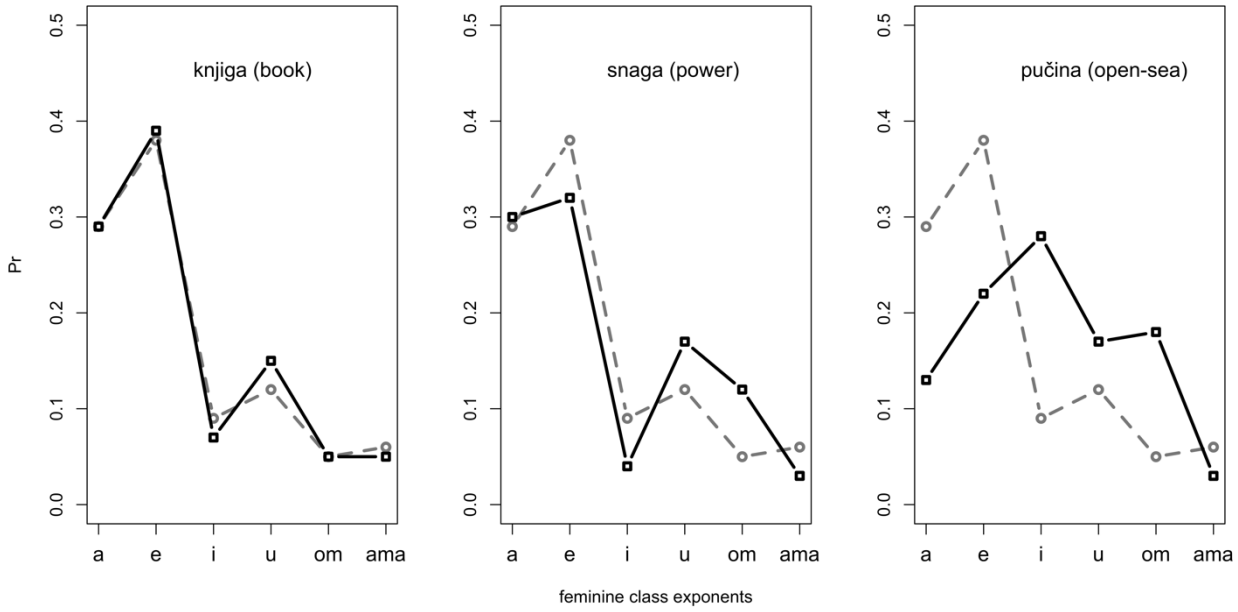

Figure 2: Variation in the distribution of inflectional exponents in Serbian nouns

Kostić and collaborators demonstrated that the class-level measures of 'amount of information' was a significant predictor of response time latencies in lexical decision tasks, despite the fact that bases and unattested inflected forms played no direct role in the experiments or in later analyses. The focus on classes and exponents in these studies can plausibly be attributed to the fact that exponents are the primary locus of inflectional variation in Slavic languages, while stems tend to remain relatively constant in a paradigm.

3.2. Words, paradigms and families

In contrast, the approaches developed by the Nijmegen School sought to measure processing effects in more of a 'bottom-up' fashion by considering relations between individual words and paradigms or word families. In the inflectional domain, one line of research probed system-level effects of the token frequencies of members of declensional paradigms. The West Germanic paradigms initially investigated were small and far less cohesive and uniform than their Serbian counterparts. Nonetheless, even these simple paradigms played a key role in the conversion of raw frequencies into estimated probabilities (i.e., proportions). A complementary line of research

in the derivational domain investigated type frequencies of members of morphological families, which are sometimes treated as an extended type of derivational paradigm (see, e.g., Robins 1959; Lieb 2005; Bonami & Strnadová 2018, among others).

Although the size of nominal paradigms in English and Dutch approach the lower limit of 2 forms, the investigation of paradigm-internal distributional contrasts revealed striking paradigmatic effects on processing. The point of departure for these studies was a contrast between **singular dominant nouns**, whose singular form is more frequent than the plural (e.g., *nose*), and **plural dominant nouns** whose plural form is more frequent than the singular (e.g., *hands*). As demonstrated in Baayen et al. (2003, 2008), the processing of each form in a 2-member noun paradigm is dependent on the other.

Even more interestingly, statistical dominance shows different effects on comprehension and production. In comprehension, a plural-dominant plural is understood more quickly than a singular-dominant plural that is matched for stem frequency. This contrast characterizes high-frequency and low-frequency nouns. In production, plurals and singulars are both slower when the plural is dominant. Baayen et al. (2008) attribute this delay to the greater information load of the inflectional paradigms of plural-dominant nouns, though it remains an open question how exactly information load induces the observed effects on production.

The paradigm-like effects on these inflected words correlated with relative frequencies of **tokens** – singular and plural forms. In contrast, the effects observed in the derivational domain correlated with the **type** frequency of items. Morphological family size is a count of the number of derivational formations containing a given word. As first observed for Dutch in Schreuder & Baayen (1997), the larger the family of derivational formations containing a word, the more quickly and accurately the word is recognized. This effect was subsequently replicated in a range of languages, including English, Hebrew, and Finnish (Moscoso del Prado Martín et al. 2004, 2005). As in the studies of statistical dominance, the explanation of family size appeals to system-level influence on the processing of individual words. Words with many paradigmatic or derivational connections are processed faster, while their relatives in turn constitute rich exemplar sets that facilitate analogical generalization.

3.3. Words, paradigms and classes

The contrast between the role of token frequency in the processing of inflection and the role of type frequency in the processing of derivation inspired attempts to derive a unified measure of processing load. Although ultimately unsuccessful, this search led to a standardisation of the measures that had been proposed to quantify processing load. As observed in Moscoso et al. (2004), the measures developed in 'bottom-up' approaches could be expressed in terms of the entropy measure proposed in Shannon (1948). This is repeated in Figure 4 below, where the entropy H of a paradigm π is defined as the frequency-weighted sum of the surprisal of its members w_e.

$$H = -\sum_{e} \Pr_{\pi}(w_e) \times \log_2 \Pr_{\pi}(w_e)$$

Figure 4: Shannon entropy as a measure of paradigmatic uncertainty

This measure preserves the bottom-up character of earlier studies by calculating the uncertainty in a paradigm from the distribution of its words. In addition, this measure underscored the close relationship between uncertainty and notions of 'processing load'.

However, measures of the entropy of individual paradigms and families did not take into account any influence attributable to inflection classes. In order to provide a unified measure of paradigm- and class-level effects, Milin et al. (2009a) proposed the relative entropy measure in Figure 5 below. This notion consolidates the empirical strategies developed in the two schools, modelling 'paradigmatic effects' by comparing the distribution of an inflectional paradigm with the distribution of a class that could be considered as a model.

$$D(P||Q;W) = \sum_{e} \Pr_{\pi}(w_e) \times \log_2 \frac{\Pr_{\pi}(w_e)}{\Pr_{\pi}(e)}$$

Figure 5: Relative entropy measure of divergence of a paradigm w_e from its class e

The integration of class- and paradigm-level effects revealed another remarkable dimension of paradigmatic structure. Using the relative entropy measure in Figure 5, Milin et al. (2009a) showed that speakers in a lexical decision task were sensitive to the **divergence** between the frequency distribution of an item's paradigm and that of its class. The patterns for three Serbian *a*-stem nouns are shown in Figure 2 above. Whereas the distribution of KNJIGA 'book' is prototypical for its class, those of SNAGA 'power' and PUČINA 'open sea' diverge progressively from the class average. For SNAGA 'power', the divergence from the class distribution lies precisely at the mid-point of all *a*-class nouns: exactly 50% diverge more than this word does from the class. These divergences are what are predictive of response latencies and captured by the relative entropy measure in Figure 5. The paradigmatic effects captured by the relative entropy measure provided striking confirmation that speakers are sensitive, simultaneously, to distributional factors at the level of the paradigm and class.

3.4. The information-theoretic turn

Although initially proposed to gauge processing costs, entropy-based measures immediately suggested strategies for quantifying classical conceptions of paradigmatic organisation in terms of 'morphological information'. Form 'variation' could be measured directly in terms of entropy, capturing the intuition that variation reflects a choice space, defined by an inventory of elements and their relationships within a set. Implicational 'structure' could likewise be modelled in terms of relational entropy measures. The simplest of these is **conditional entropy**, which, in this case, measures the amount of uncertainty about a paradigmatic choice that is eliminated by knowledge of a different choice or set of choices. These measures were introduced in the approach to the 'Paradigm Cell Filling Problem' outlined in Ackerman et al. (2009). Subsequent studies con-

firmed the usefulness of an 'information-theoretic perspective' in investigating a host of traditional issues: characterising inflectional variation and paradigm economy (Ackerman and Malouf 2013, 2015), distinguishing valid from spurious analogies (Blevins 2013), deriving notions of defectiveness (Sims 2015), among other problems.

Two general results have already emerged from this line of research. The first is the demonstration that the members of an inflectional (and, to a lesser degree, derivational) paradigm exhibit measurable patterns of interdependency, and that these patterns support effective strategies for deducing the shape and function of previously unencountered forms. Once idealisations about principal parts and exemplary choices are discarded as pedagogical artifacts, the existence of a strong implicational 'signal' in inflectional systems supports a classical WP conception of the organisation of such systems. The robustness of this paradigmatic structure is confirmed by the 'set-theoretic' approach of Stump & Finkel (2015), which models the same patterns of interdependency using a formally distinct approach.

The second general result, which is largely a corollary of the first, is that prevailing assessments of the relative 'complexity' of morphological systems based on global assessments of the amount of variation, are misleading, and arguably as artifactual as principal parts or exemplary paradigms. As argued at length in Ackerman and Malouf (2013), contrasts between the relative 'complexity' of morphological systems can be constrained if 'complexity' is determined by a cognitively-plausible measure that assumes partial familiarity with the system. For a speaker who is wholly unacquainted with any system, there will be a vast amount of variation in the space of alternatives across systems. However, partial knowledge of a system will allow speakers to exploit implicational patterns, and bring the variation within a more circumscribed range.

This literature has also sparked a robust critical discussion of the idealisations and limitations of initial information-theoretic approaches (see especially Sagot 2013, 2018; Bonami 2014). Many of these issues derived from idealisations or limiting assumptions incorporated in psycholinguistic operationalisations of morphological processing tasks. The entropy measures adapted from the psycholinguistic literature tended to carry over morphotactic idealisations, notably stem/exponent divisions in the characterisation of the 'random variables' used to calculate entropy values. Frequency counts likewise depended on particular treatments of allomorphy and other types of alternations. For studies of low-resource languages, the problems are even more acute. Descriptions of these languages tend not to provide frequency information, forcing analyses to fall back to the (counterfactual) assumption that many of unattested morphological alternatives could be treated as equiprobable, in order to obtain entropy 'ceilings'. These specific concerns about the use of corpora are accompanied by more general concerns about the integrity of corpora as aggregated data sources. It remains an open question whether corpora reflect factors such as socio-cultural conventions more than they represent a cognitively-valid sample of speakers' linguistic experience (cf. Divjak 2015).

At least some of these issues can be addressed through refinements of information-theoretic models. An implicit stem/affix segmentation is largely a holdover from the exponent-oriented

measures of the Belgrade School, again reflecting the morphological structure of Serbian. It is convenient (though not always unproblematic, as Bonami 2014 argues) to preserve this idealisation and define the entropy of a paradigm cell in terms of the number and distribution of exponents. However, the calculation of entropy is neutral with respect to the form variations. Paradigm size (as noted in Sagot 2018) could also be explicitly factored into the determination of entropy, to ensure commensurability. The challenges posed by low-resource languages and the concerns related to aggregation remain, as they do for any quantitative corpus-based approach.

Yet there is a sense in which these issues are symptomatic of deeper challenges. Cognitively-plausible analyses require descriptions of form variation that capture all of the contrasts to which speakers are sensitive. As demonstrated by studies of sub-phonemic variation in inflection (Baayen et al. 2003, Kemps et al 2005, Plag et al. 2015), speakers often perceive and produce phonetic properties that distinguish orthographically and phonemically identical elements. Hence eliminating the morphotactic idealisations inherited from psycholinguistic operationalizations still leaves the idealizations embedded in orthographic and phonemic representations.

Assessments of paradigm size raise a parallel set of issues. What notion of 'paradigm' should provide the basis for cognitively-plausible entropy measures? Given the Zipfian structure of a primary linguistic input (Blevins et al. 2017), speakers will encounter only partial paradigms of the open-class items of their language. Entropy calculations based on full paradigms thus represent an idealized extrapolation of a speaker's experience. Basing calculations on encountered forms (as estimated, e.g., from corpora) confronts the an acute problem of data sparsity, since approximately half of the inflected forms in a language will be hapax legomena, encountered only once, and another 10% are typically dis legomena, encountered twice.

A more fundamental issue concerns the cognitive status of global entropy measures. As noted above, entropy measures provide a means of quantifying the implicational structure that has long been attributed to morphological systems. But these global measures are akin to 'weather forecasting', providing no explanation for implicational structure, and making no contribution to an understanding of the cognitive mechanisms that might account for the uncertainty estimated by entropy measures. While information-theoretic measures are a useful diagnostic tool for probing implicational organisation, the explanation for this organisation must be grounded elsewhere, in cognitive mechanisms and structures.

4. Computational and learning perspectives

Most computational accounts of language processing in general, and paradigm processing in particular, share a common trait: they are all fundamentally learning-based. Even the most orthodox rule-driven model, strictly speaking, needs to address the questions that arise in the acquisition and/or learning of a system of rules and their interrelationships, in order to implement the formal algorithmic principles of the model. More often than not, models of paradigm processing are overtly based on specific, well-established learning processes. To date, the problem of paradigm processing has been addressed through the successful application of **memory-based**

learning and **error-driven** learning. A memory-based learning approach is implemented in the Tilburg Memory-Based Learner (TiMBL: Daelemans et al. 2007), and in Analogical Modeling (AM: Skousen 1989; Skousen et al. 2002). Prominent representatives of the error-driven learning paradigm include models such as connectionist parallel distributed processing (PDP: Rumelhart & McClelland 1986), naive discrimination learning (NDL: Ramscar et al. 2010; Baayen et al. 2011; Milin et al. 2017), and temporal self-organising maps (tSOM: Pirrelli et al. 2011, 2015).

The research traditions based on memory-based and error-driven learning principles have pursued the goal of broad coverage, modelling both the comprehension and production of morphological inflections. The empirical focus of these traditions also extends from full paradigms to specific paradigmatic subparts that exhibits allomorphic variation which appears to present noteworthy challenges for computational models.

4.1. Memory-based and error-driven learning of paradigms

The potential of memory-based models for modelling morphological patterns is illustrated by approaches that implement principles of analogical inference, using exemplars stored in memory to generate or classify novel input. Keuleers and Daelemans (2007) used TiMBL to produce correct Dutch noun plurals from singular pairs and a larger set of known plural forms. The same principles of analogical inference were applied with considerable success in studies of allomorphic variations in more intricate morphological systems such as South-Slavic nominal paradigms. Milin et al. (2011) used TiMBL to model allomorphic variation in Serbian masculine instrumental singular nouns. Lečić (2016) compared TiMBL and AM on modelling Croatian allomorphs (i.e., 'overabundance' or 'doubletism') in masculine instrumental singulars and feminine genitive plurals.

Error-driven learning strategies are implemented in pioneering connectionist models (e.g. perceptrons). These models simulated the task of producing an inflected output (e.g., *went*) from a base input (*go*), starting with a fully connected network of input-to-output mappings or links. Links were either direct or mediated by a number of units in hidden-layer(s), all were weighted to reflect the strength of connections. Weights were gradually adjusted by error-correction, typically involving back-propagation (cf. Werbos 1977). The success of early models lead researchers to address more challenging and intricate problems. In the domain of inflectional morphology, Mirković et al (2011) applied a four-layer PDP network (i.e., one with input and output layers and two hidden layers) to the task of producing phonologically correct inflected forms for a sample of randomly-selected Serbian nouns. These results built on the earlier study of Miković et al. (2005), which had modelled paradigm-related patterns of nominal gender, again in Serbian.

Temporal self-organising maps provide a network-based model that implements a two-stage learning strategy. In the first stage, these maps perform a spatial (or topological) coding of input units onto a map. In the second stage, they perform a discriminative coding of the temporal contingencies of the map locations, i.e., the consecutive activations of the map nodes. Pirrelli and collaborators have obtained an impressive range of results in which tSOM systems have learned

morphological paradigms in a diverse range of languages, including Italian, German, Greek, and Spanish (Pirrelli et al. 2011, 2015; Marzi et al. 2014, 2016; Pirrelli 2018). A distinctive aspect of this class of models is that self-organising maps express an explicit commitment not only at the functional (behavioural or cognitive) level, but also at the implementational (neural) level.

Although naive discrimination learning models share some implementational traits with other error-driven models, notably PDP and tSOM models, computational NDL approaches provide a more straightforward representation of error-driven learning principles. Unlike PDP models, NDL approaches do not use layers of hidden units. Instead, an NDL model learns to calibrate direct links between input and output units, which then proximate an (untransformed) input cue's potency to predict (discriminate) an outcome. Unlike tSOM models, NDL approaches do not use intermittent topological representations, nor do they commit to emulating neural architecture in any strict sense. Instead, the NDL model provides the simplest approach to modelling input cue competition. Formally, it is an implementation of the delta rule of Widrow and Hoff (1960), which is familiar in computer engineering, and in machine learning in particular. Rescorla and Wagner (1972) independently arrived to the same formal specification of error-driven learning to define what is arguably the most successful model of human and animal learning (see, e.g., Milin, Nenadić, & Ramscar, submitted).

The differences between memory-based learning and naive error-driven learning are reflected in the divergent predictions that they make. As discussed in Milin et al. (2016), TiMBL and NDL produce predictions that conflict in intriguing ways when applied to a large sample of Serbian inflected forms. TiMBL finds it 'useful' to rely on lemma letter sequences that are present in many other lemmata. These highly diverse letter sequences provide the foundation for predictions based on the dense exemplar space in TiMBL's memory. NDL follows exactly the opposite 'reasoning', as it eventually un-learns these cues and weights them lightly, as they are divergent and, thus, not particularly discriminative of any specific outcome. The opposing reasoning strategies can be attributed, at least in part, to the tasks to which these models have primarily been applied. TiMBL's probabilities are tailored to language **production**. In contrast, the weights in NDL have been developed mainly for **comprehension** tasks such as lexical decision or reading, in which the diversity of letter (or phones) sequences can be harmful.

4.2. Information load and learning

Learning, and learning theory in particular, can offer an explanation for the predictive value of the information-theoretic measures discussed in Section 3 above. As observed by Milin, Nenadić, & Ramscar (submitted), learning is driven by "the systematicity of contextualised experience", which amounts to the requirement that any relationship that can be learned cannot exhibit completely random behaviour. In mathematical terms, there must be a "divergence between the observed entropy of a potentially informative event and the maximum entropy" of the event (Gallistel 2003: 93).

The information-theoretic tradition has securely established that the distributional properties of morphological paradigms are at some remove from theoretically maximal uncertainty (e.g., maximum entropy). This entails that paradigms are learnable, at least in principle; a point that is confirmed by the fact that they are learned by speakers. It is a matter of increasing interest and importance to understand how precisely paradigms are learned, and what exactly can be discovered about paradigms from computational models. Methodologies and techniques for addressing these questions form some of the most active areas of research into morphological paradigms (cf. Ferro, Marzi, & Pirrelli 2018), complementing work on the systematicity of language (cf. Ramscar et al. 2010, 2013; Dye et al. 2018) and the complex dynamics that guide linguistic organization and evolution (Beckner et al. 2009).

Different learning models offer alternative perspectives on the mechanisms that underlie the organizational principles captured by global uncertainty measures. One the one hand, as noted above, exemplar-based models such TiMBL and AM provide the most direct representation of the analogical principles proposed in classical WP models. On the other hand, NDL models suggest the most transparent characterization of classical conceptions of 'variation' and 'structure'. From an NDL perspective, questions about the variation in a set of forms amount to questions about the contrasts that speakers discriminate in learning the forms. Questions about the size of the paradigms that speakers internalize can likewise be recast as questions about learning inputs and effects on the state of a learning model. Questions about the mechanisms that underlie global entropy measures are in turn recast as questions about the paradigmatic "connections to other words" in a learning model that represents the state of a speaker's knowledge of a morphological system.

Baayen at al. (2011) appear to have been the first to propose a learning-based explanation for the divergence between the distributions of inflectional paradigms and inflection class captured by the relative entropy measure of Milin et al. (2009a). This study used an idealized example of a small lexicon containing a number of lemmata and their (deliberately ambiguous) inflected variants. The NDL model was trained on this data to find support for three central claims. The first was that there existed a degree of 'attraction' (or 'repulsion') between a lemma and the inflectional exponents of its class. The second was that the process of discrimination learning determined association weights between those elements, reflecting positive connections (attraction) as well as negative connections (repulsion). The third claim was that relative entropy is an unsigned measure of this association, one that shows only the strength of a connection but not its directionality (e.g., attraction or repulsion).

As acknowledged in Milin et al. (2017), this initial NDL study was implicitly decompositional, as it discriminated morphological constituents (not only lemmata but also sub-word stems and exponents). To demonstrate that this commitment was excisable and indeed disadvantageous, Milin et al. (2017) showed that an NDL model that dispensed with sub-word units enjoyed advantages over the earlier 'decompositional' model in predicting lexical decision latencies in a large experimental study. The revised NDL model also shared closer ties with classical and con-

temporary WP models (cf. Blevins 2016), which recognize no morphological unit smaller than the word.

The fruitfulness of this approach is reflected in the way that it has not only clarified traditional notions but also opened up entirely new vistas. For example, the study of paradigmatic entropy has a natural counterpart in surprisal-based accounts of syntagmatic surprisal (Hale 2003, 2006). In what appears to be the first investigation of the relationship between paradigmatic and syntagmatic uncertainty, Filipović Đurđević and Milin (2018) conducted an experimental study of the entropies of Serbian inflected adjectives. The study revealed a systematic trade-off between sources of uncertainty: adjectives could exhibit high paradigmatic uncertainty or high syntagmatic uncertainty, but not both. This result suggests a dynamic relation between uncertainty along paradigmatic and syntagmatic axes. As in earlier studies in this tradition, information theory provides a cognitively-relevant measure of structure. Learning theory, and the NDL model in particular, offers an explanatory framework for grounding measures of information load in basic principles of discrimination learning.

5. Conclusions

Until comparatively recently, virtually all treatments of morphological paradigms, whether theoretical, psycholinguistic or computational, were based on descriptions provided by grammars, dictionaries and other traditional lexical resources. Description-based approaches enjoy a number of familiar advantages: they can be applied to low-resource languages, they can be comprehensive in scope, and they can abstract away from contingent factors, such as the quantity and nature of the input that speakers may encounter. Yet these advantages all come paired with disadvantages for models that aspire to cognitive relevance. The cognitive value of description-based analyses correlates with their success in capturing contrasts to which speakers are sensitive; the same idealizations that enhance the suitability of materials for pedagogical or other practical purposes impair their usefulness as models of speakers' knowledge of a language.

In the domain of morphology, it is clear that standard orthographic and phonemic descriptions fail, systematically, to express morphologically relevant sub-phonemic variation. The heavy hand of history also exerts a largely unacknowledged influence. The morphotactic analyses assigned in philologically-sophisticated descriptive traditions often mirror the processes of grammaticalization that produced the modern forms. But the synchronic relevance of these historical patterns stands in need of validation in any account that aims to model a speaker's morphological knowledge. In general, there is no evidence that analyses based on idealized descriptions can be successfully retrofitted to incorporate cognitive motivation for patterns and structures obtained using linguistic or philological methods.

Yet, as the language sciences adapt to the creative disruption unleashed by the big data revolution, the approaches developed to investigate the structural organization and psychological effects of morphological paradigms find themselves well positioned to exploit the opportunities created by this revolution. Many questions and challenges remain to be addressed. These in-

clude: (i) a latent ambiguity about the object of study in formal analyses of paradigmatic systems, (ii) concerns about the ecological validity of experimental psycholinguistic evidence, (iii) the need for critical re-evaluations of idealizations and simplifications incorporated in computational models, and other fundamental issues of this nature. There are, nonetheless, grounds for cautious optimism. Tendencies "toward a wedding of insufficiencies" (Osgood 1968) within initially independent research traditions have coalesced into a broadly compatible theoretical and methodological perspective (cf. Milin et al. 2016) that appears capable of producing answers to long-established questions that lead in turn to new and deeper questions.

References

Ackerman, Farrell, James P. Blevins, and Robert Malouf. 2009. Parts and wholes: Implicative patterns in inflectional paradigms. In *Analogy in Grammar: Form and Acquisition*, eds. James P. Blevins and Juliette Blevins, 54–81. Oxford University Press.

Ackerman, Farrell, and Robert Malouf. 2013. Morphological organization: The Low Conditional Entropy Conjecture. *Language* 89, 429–464.

Anderson, Stephen R. 1992. *A-Morphous Morphology*. Cambridge: Cambridge University Press.

Aronoff, Mark. 1994. *Morphology by Itself: Stems and Inflectional Classes*. Cambridge, MA: MIT Press.

Baayen, R. Harald, Rochelle Lieber, and Robert Schreuder. 1997. The morphological complexity of simple nouns. *Linguistics* 35, 861–877.

Baayen, R. Harald, James M. McQueen, Teun Dijkstra, and Robert Schreuder. 2003. Frequency effects in regular inflectional morphology: Revisiting Dutch plurals. In *Morphological Structure in Language Processing*, eds. R. Harald Baayen, and Robert Schreuder, 355–370. Berlin: Mouton de Gruyter.

Baayen, R. Harald, Lee H. Wurm, and Joanna Aycock. 2008. Lexical dynamics for low-frequency complex words: A regression study across tasks and modalities. *The Mental Lexicon* 2, 419–463.

Beckner, Clay, Richard Blythe, Joan Bybee, Morten H. Christiansen, William Croft, Nick C. Ellis et al. 2009. Language is a complex adaptive system: Position paper. Language learning, 59, 1–26.

Bertram, Raymond, Matti Laine, R. Harald Baayen, Robert Schreuder, and Jukka Hyönä. 2000. Affixal homonymy triggers full-form storage, even with inflected words, even in a morphologically rich language. *Cognition* 74, B13—B25.

Blevins, James P. (2013). The information-theoretic turn. *Psihologija* 46(3), 355–375.

Blevins, James P. 2016. *Word and Paradigm Morphology*. Oxford: Oxford University Press.

Blevins, James P., Petar Milin, and Michael Ramscar. 2017. The Zipfian Paradigm Cell Filling Problem. In *Perspectives on Morphological Structure: Data and Analyses*, eds. Ferenc Kiefer, James P. Blevins, and Huba Bartos, 141–158. Leiden: Brill.

Bloomfield, Leonard. 1933. *Language*. Chicago: University of Chicago Press.

Bonami, Olivier. 2014. *La structure fine des paradigmes de flexion: Études de morphologie descriptive, théorique et formelle*. Habilitation, Université Paris 7 – Denis Diderot.

Burani, Cristina, Dario Salmaso, and Alfonso Caramazza. 1984. Morphological structure and lexical access. *Visible Language* 18, 342–352.

Caramazza, Alfonso, Alessandro Laudanna, and Cristina Romani. 1988. Lexical access and inflectional morphology. *Cognition* 28, 297–332.

Chomsky, Noam. 1957. *Syntactic Structures*. The Hague: Mouton.

Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Chomsky, Noam. 1969. Quine's empirical assumptions. In *Words and objections: Essays on the work of W. V. Quine*, eds. D. Davidson and J. Hintikka, 53–68. Dordrecht: Reidel.

Clahsen, Harald. 1999. Lexical entries and rules of language: A multidisciplinary study of German inflection. *Behavioral and Brain Sciences* 22(6), 991–1060.

Colé, Pascale, Cécile Beauvillain, and Juan Segui. 1989. On the representation and processing of prefixed and suffixed derived words: A differential frequency effect. *Journal of Memory and Language* 28, 1–13.

Daelemans, Walter, Jakub Zavrel, Ko Van der Sloot, and Antal Van den Bosch. 2010. *TiMBL: Tilburg memory-based learner, version 6.3, Reference Guide*. ILK Technical Report – ILK 10–01. http://ilk.uvt.nl/downloads/pub/papers/ilk.1001.pdf

Davidson, Thomas. 1874. *The Grammar of Dionysus Thrax*. St Louis: R. P. Studley & Co.

Divjak, Dagmar. 2015. Four challenges for usage-based linguistics. In *Change of paradigms – New paradoxes; Recontextualizing language and linguistics*, eds. J. Daems, E. Zenner, K. Heylen, D. Speelman, and H. Cuyckens. Vol. 31, 297–310. Berlin: Walter de Gruyter.

Dye, Melody, Petar Milin, Richard Futrell, and Michael Ramscar. 2018. Alternative Solutions to a Language Design Problem: The Role of Adjectives and Gender Marking in Efficient Communication. Topics in cognitive science, 10(1), 209–224.

Fertig, David. 2013. *Analogy and Morphological Change*. Edinburgh: Endinburgh University Press.

Ferro, Marcello, Claudia Marzi, and Vito Pirrelli. 2018. Discriminative word learning is sensitive to inflectional entropy. *Lingue e Linguaggio* 17(2), 307–327.

Gallistel, C. Randy. 2003. Conditioning from an information processing perspective. *Behavioural Processes* 62(1–3), 89–101.

Gentner, Dedre, Keith James Holyoak, and Boicho N. Kokinov. 2001. *The Analogical Mind: Perspectives from Cognitive Science*. Cambridge, Mass.: MIT Press.

Hale, John (2003). The information conveyed by words in sentences. *Journal of Psychological Research* 32(2), 101–123.

Hale, John (2006). Uncertainty about the rest of the sentence. *Cognitive Science* 30(4), 643–72.

Hockett, Charles F. 1954. Two models of grammatical description. *Word* 10, 210–231.

Hockett, Charles F. 1958. *A Course in Modern Linguistics*. New York: MacMillan.

Hovdhaugen, Even. 1996. Roman Ars Grammatica, including Priscian. In *Concise History of the Language Sciences: From the Sumerians to the Cognitivists*, eds. E. K. Koerner, and R. E. Asher, 115–115. Oxford: Pergamon.

Hyönä, Jukka, Matti Laine, and Jussi Niemi. 1995. Effects of a word's morphological complexity on readers' eye fixation patterns. In *Eye movement research: Mechanisms, processes and applications*, eds. J. M. Findlay, R. W. Kentridge, & R. Walker, 445–452. Amsterdam: Elsevier.

Itkonen, Esa. 2005. *Analogy As Structure And Process: Approaches in Linguistics, Cognitive Psychology, and Philosophy Of Science*. Amsterdam: Benjamins.

Janßen, Ulrike, and Martina Penke. 2002. are inflectional affixes organized in the mental lexicon? Evidence from the investigation of agreement errors in agrammatic aphasics. *Brain and Language* 81(1–3), 180–191.

Juola, Patrick (1998). Measuring linguistic complexity: The morphological tier. *Journal of Quantitative Linguistics* 5(3), 206–213.

Kostić, Aleksandar. 1991. Informational approach to processing inflectional morphology: Standard data reconsidered. *Psychological research* 53, 62–70.

Kostić, Aleksandar. 1995. Informational load constraints on processing inflectional morphology. In *Morphological Aspects of Language Processing*, ed. Laurie B. Feldman, 317–344. Hillsdate, NJ: Lawrence Erlbaum.

Krott, Andrea, Robert Schreuder, and R. Harald Baayen. 2002. Analogical hierarchy: Exemplar-based modeling of linkers in Dutch noun-noun compounds. In *Analogical Modeling: An Exemplar-Based Approach to Language*, eds. Royal Skousen, Deryle Londsdale, and Dilworth B Parkinson, 181–181. Amsterdam: Benjamins.

Lečić, Dario. 2016. *Morphological Doublets in Croatian: A Multi-methodological Analysis*. PhD dissertation, University of Sheffield.

Lieb, Hans-Heinrich (2005), Notions of paradigm in grammar, in *Lexikologie / Lexicology: Ein internationales Handbuch zur Natur und Struktur von Wörtern und Wortschätzen / An international handbook on the nature and structure of words and vocabularies*, eds. D. Alan Cruse et al., 1613–1646. Berlin: de Gruyter.

Lukatela, Georgije, Branislaw Gligorijević, Aleksandar Kostić, and Michael T. Turvey. 1980. Representation of inflected nouns in the internal lexicon. *Memory & Cognition* 8, 415–423.

Marzi, Claudia, Marcello Ferro, and Vito Pirrelli. 2014. Morphological structure through lexical parsability. Lingue e Linguaggio, 13(2), 263–290.

Marzi, Claudia, Marcello Ferro, Franco Alberto Cardillo, and Vito Pirrelli. 2016. Effects of frequency and regularity in an integrative model of word storage and processing. Italian journal of linguistics, 28(1), 79–114.

Matthews, Peter H. 1994. Greek and Latin linguistics. In *History of Linguistics II: Classical and Medieval Linguistics*, ed. Lepschy Guilio, 1–133. London: Longman.

Milin, Petar, Dušica Filipović Đurđević, and Fermín Moscoso del Prado Martín. 2009a. The simultaneous effects of inflectional paradigms and classes on lexical recognition: Evidence from Serbian. *Journal of Memory and Language* 60, 50–64.

Milin, Petar, Victor Kuperman, Aleksandar Kostić, and R. Harald Baayen. 2009b. Words and paradigms bit by bit: An information-theoretic approach to the processing of inflection and derivation. In *Analogy in Grammar: Form and Acquisition*, eds. James P. Blevins, and Juliette Blevins, 214–253. Oxford University Press.

Milin, Petar, Emmanuel Keuleers, and Dušica Đurđević. 2011. Allomorphic responses in Serbian pseudo-nouns as a result of analogical learning. *Acta Linguistica Hungarica* 58(1), 65–84.

Milin, Petar, Dagmar Divjak, Strahinja Dimitrijević, and R. Harald Baayen. 2016. Towards cognitively plausible data science in language research. *Cognitive Linguistics* 27(4), 507–526.

Milin, Petar, Filip Nenadić, and Michael Ramscar. How contextualized experience shapes task-specific performance. Submitted to *Cognitive Science*.

Mirković, Jelena, Mark S. Seidenberg, and Marc F. Joanisse. 2011. Rules versus statistics: Insights from a highly inflected language. *Cognitive science* 35(4), 638–681.

Mirkovic, Jelena, Maryellen C. MacDonald, and Mark S. Seidenberg. 2005. Where does gender come from? Evidence from a complex inflectional system. *Language and Cognitive Processes* 20(1–2), 139–167.

Morpurgo Davies, Anna. 1998. *Nineteenth-century Linguistics*. London: Longman.

Moscoso del Prado Martín, Fermín, Aleksandar Kostić, and R. Harald Baayen. 2004. Putting the bits together: An information-theoretical perspective on morphological processing. *Cognition* 94, 1–18.

Mulder, Komberley, Teun Dijkstra, Robert Schreuder, and R. Harald Baayen. 2014. Effects of primary and secondary morphological family size in monolingual and bilingual word processing. *Journal of Memory and Language* 72, 59–84.

Osgood, Charles E., and Thomas A. Sebeok. 1954. Psycholinguistics: A survey of theory and research problems. A supplement to *Journal of Abnormal and Social Psychology* 49, xi—203.

Osgood, Charles E., George J. Suci, and Percy H. Tannenbaum. 1957. *The measurement of meaning*. Urbana: University of Illinois Press.

Osgood, Charles E. 1966. Meaning cannot be r_m? *Journal of Verbal Learning and Verbal Behavior* 5(4), 402–407.

Osgood, C. E. (1968). Toward a wedding of insufficiencies. In Tomas R. Dixon and Deryck L. Horton (Eds.), Verbal behavior and general behavior theory (pp. 495–519). Englewood Cliffs: Prentice-Hall.

Paul, Hermann. 1920. *Prinzipien der Sprachgeschichte*. Tübingen: Max Niemayer Verlag.

Pirrelli, Vito, Marcello Ferro, and Basilio Calderone. 2011. Learning paradigms in time and space. Computational evidence from Romance languages. In M. Maiden, J. C. Smith, M. Goldbach, and M. O. Hinzelin (Eds.) Morphological Autonomy: Perspectives from Romance Inflectional Morphology. Oxford: Oxford University Press.

Pirrelli, Vito, Marcello Ferro, and Claudia Marzi. 2015. Computational complexity of abstractive morphology. In Matthew Baerman, Dunstan Brown, and Greville G. Corbett (Eds.), Understanding and Measuring Morphological Complexity (pp. 141–166). Oxford: Oxford University Press.

Pirrelli, Vito. 2018. Morphological theory and computational linguistics. In Jenny Audring and Francesca Masini (Eds.), The Oxford Handbook of Morphological Theory, (pp. 573–593). Oxford: The Oxford University Press.

Plag, Ingo. (2010). Compound stress assignment by analogy: The constituent family bias. *Zeitschrift für Sprachwissenschaft* 29, 243–282.

Pullum, Geoffrey K., and Barbara C. Scholz. 2010. Recursion and the infinitude claim. In *Recursion and Human Language*, eds. Harry van der Hulst, Jan Koster, and Henk van Riemsdijk, 113–138. De Gruyter Mouton.

Ramscar, Michael, Daniel Yarlett, Melody Dye, Katie Denny, and Kirsten Thorpe. 2010. The effects of feature-label-order and their implications for symbolic learning. *Cognitive Science* 34, 909–957.

Ramscar, Michael, Melody Dye, and Stewart M. McCauley. 2013. Error and expectation in language learning: The curious absence of mouses in adult speech. *Language* 89(4), 760–793.

Robins, Robert H. 1959. In defense of WP. *Transactions of the Philological Society* 58, 116–144.

Rumelhart, David E. and James McClelland. 1986. On learning the past tense of English verbs. In *Parallel distributed processing: Explanations in the microstructure of cognition*, eds. David E. Rumelhart and James McClelland. MIT Press. 216–271.

Sagot, Benoît. 2018. *Informatiser le lexique — Modélisation, développement et exploitation de lexiques morphologiques, syntaxiques et sémantiques*. Habilitation, Sorbonne Université.

Schreuder, Robert, and R. Harald Baayen. 1997. How complex simplex words can be. *Journal of Memory and Language* 37, 118–139.

Shannon, Claude E. 1948. A mathematical theory of communication. *Bell System Technical Journal* 27(3), 379–423.

Shannon, Claude E. 1956. The bandwagon. *IRE Transactions on Information Theory* 2(1), 3.

Skinner, Burrhus F. 1957. *Verbal behavior*. New York: Appleton-Century-Crofts.

Skousen, Royal. 1989. *Analogical modeling of language*. Berlin: Springer.

Skousen, Royal, Deryle Lonsdale, and Dilworth B. Parkinson. 2002. *Analogical modeling: An Exemplar-based Approach to language*. Amsterdam: John Benjamins.

Stump, Gregory T. 2001. *Inflectional Morphology: A Theory of Paradigm Structure*. Cambridge: Cambridge University Press.

Taft, Marcus. 1979. Recognition of affixed words and the word frequency effect. *Memory & Cognition* 7, 263–272.

Werbos, Paul. 1974. *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. PhD dissertation, Harvard University.

Wurzel, W. U. 1989. *Inflectional Morphology and Naturalness*. Dordrecht: Kluwer.

Zwicky, Arnold M. 1985. How to describe inflection. In *Proceedings of the Eleventh Annual Meeting of the Berkeley Linguistics Society*. Berkeley: Berkeley Linguistics Society, 372–386.