



# Integrative views of representations and processes in morphology: an introduction

Claudia Marzi<sup>1</sup> · Vito Pirrelli<sup>1</sup>

Received: 14 July 2023 / Accepted: 27 July 2023 / Published online: 14 August 2023  
© The Author(s) 2023

## Abstract

One of the most enduring conceptualisations of the language architecture rests on a modular subdivision of work between lexical representations of stored items on the one hand, and dynamic processes, modelled as procedural rules working on such items, on the other hand. In morphology, network-based approaches have suggested an alternative “integrative” view of word representations and processes, where lexical representations consist of partially overlapping activation patterns spreading over several processing units. From this integrative perspective, the resulting network is *both* a lexicon *and* a word processor. We argue that the network-based view provides a stimulating research framework for several complementary levels of language inquiry (including theoretical, computational and neuro-psychological approaches) to be fruitfully integrated into a novel, comprehensive understanding of morphology. We discuss some implications of this view and delineate prospects of progress in this area.

**Keywords** Morphology · Mental lexicon · Connectionism · Network science · Discriminative learning

## 1 The *network* metaphor

### 1.1 Two models of grammatical description

The dualism between representations and processes has deep roots in morphological theory. This is hardly surprising, given the hybrid ontological status of words, halfway between stable holistic units, committed to a speaker’s long-term memory,

---

✉ V. Pirrelli  
[vito.pirrelli@ilc.cnr.it](mailto:vito.pirrelli@ilc.cnr.it)

C. Marzi  
[claudia.marzi@ilc.cnr.it](mailto:claudia.marzi@ilc.cnr.it)

<sup>1</sup> Institute for Computational Linguistics, Italian National Research Council (ILC-CNR), Via Moruzzi 1, 56124 Pisa, Italy

and changeable, context-sensitive units, which are themselves parts of a large network of paradigmatic and syntagmatic relations with other words. Hockett's (1954) *Two models of grammatical description* laid the theoretical foundations of this dualism from a structuralist perspective. According to Hockett, a process-based approach to morphology (which he dubbed "item-and-process") maintains that the morphological relation between two words is modelled "as a process which yields one form out of the other". Conversely, the essence of an arrangement-based model (or "item-and-arrangement" in Hockett's terminology) is "to talk simply of things and the arrangement in which those things occur".<sup>1</sup> Although he conceded that neither model was completely satisfactory, Hockett's preference went to the arrangement-based approach, as he was not happy with the process-based assumption that one form could be *derived* from another form.

## 1.2 The cognitive revolution

In fact, Hockett's dualism did not survive the advent of what, years later, will come to be known as the "cognitive revolution" (Miller, 2003), when the structuralist emphasis on meticulous data analysis and system-based factors was replaced by the bold, anti-behaviourist metaphor that the mind is a *computer*, i.e. an information processing machine. According to this metaphor, linguistic information can be represented symbolically, and cognitive processes can be described, with no loss of generality or scientific accuracy, in terms of algorithmic rules that operate on these symbols.

One of the most enduring consequences of the computer metaphor on views of language architecture was a modular subdivision of work between representations of atomic items stored in the (mental) lexicon, and procedural rules combining such items (Baayen, 2007). From this perspective, stored items are taken to be stable units, which are recalled from long term memory in the same form as they were originally memorised. Their representation is fundamentally independent of processing principles and is not affected by rule application. Complex structures (e.g. morphologically complex words, phrases and sentences) are the outcome of rule-based processes of online symbolic manipulation. As such, they are not stored, but computed on the fly. Under the spell of the computer metaphor, Hockett's original dualism was bound to vanish. The two models appeared to differ mainly in the types of rules they allow for: namely, combinatorial rules for arrangement-based approaches, and fusional rules for process-based ones. In cognition, the computationalist view, most radically advocated within Pinker and Ullman's *Declarative/Procedural* model (Pinker & Ullman, 2002; Ullman, 2001, 2004), went on to maintain that speakers' knowledge of word inflection is subserved by two distinct, functionally segregated human brain systems: the declarative memory, where irregularly-inflected forms are stored as atomic units, and the procedural memory, where regularly-inflected forms are produced by assembling their (stored) sublexical units (see also Marzi & Pirrelli, this issue). Accordingly, only arrangement-based morphology survived in Pinker's view of English inflection.

---

<sup>1</sup>In passing, Hockett also mentioned a third view, which he dubbed "Word and paradigm". He considered this third model as deserving the same consideration given to the other two, but he probably found it too "loose" to describe in structuralist terms.

### 1.3 Connectionism

The advent of connectionism (Rumelhart et al., 1986) brought to the fore a very different metaphor. The mind is the *brain*, i.e. a system of highly interconnected *units of information* (akin to single neurons or neuron clusters) that are also *processing units*, i.e. which “fire” (change their state of activation) in response to a stimulus or a class of stimuli. The consequences of this metaphor on views of language architecture were far reaching, and went well beyond the limitations of early connectionist simulations of child’s language learning (Pinker & Prince, 1994). The idea that linguistic representations are distributed over very many processing units (nodes) that activate and compete in parallel, mimicking the connectivity structure among neurons in the brain, proved to be instrumental in simulating many aspects of human cognition.

This move radically challenged the computationalist view of representations. Neural network representations are no longer “things” (to use Hockett’s term): rather, i) they consist in real-valued *activation patterns* spreading over a great number of nodes; ii) they are not enumeratively identifiable and exhaustively listed, since partially overlapping activation patterns can be associated with new representations the network was never trained on, iii) they can change dynamically both at short time scales (in response to the current input stimulus) and at long time scales (as a result of small, incremental changes of connection weights after repeated exposure to many input stimuli), iv) they are strongly context-sensitive and probabilistic, as they depend on where and when an input stimulus activates a node or a node cluster.

What about processes? In principle, in a neural network, representations exist because some cognitive processes apply to input signals, and because the human brain tends to memorise its most successful processing responses. However, early connectionist models did not deal with temporal input representations in a satisfactory way. They used bigram or trigram nodes (so-called “Wickelphones”) to encode lexical forms as static activation patterns. Wickelphones looked like a computationalist left-over in a parallel processing architecture. The problem was solved few years later through *recurrent* neural networks (Elman, 1990; Jordan, 1986), which use *recurrent* connections to have access to their own activation history. What initially appeared to be a simple technical solution to a technical problem, turned out to be a fundamental principle of the processing brain, where representations and processes are in fact *mutually implied* (Marzi & Pirrelli, 2015). On the one hand, representations consist in (stored) successful processing patterns. Processes, in turn, consist of the transient, task-related activation of long-term memory patterns (Wilson, 2001; D’Esposito, 2007; D’Esposito & Postle, 2015; Ma et al., 2014; Sreenivasan et al., 2014).

### 1.4 Network science

In the late 70s and early 80s, connectionism appeared to resonate well with growing empirical evidence in cognitive psychology that the application of the mathematical tools and graph-theoretical concepts of *network science* could make substantial contributions to modelling the structure of lexical nodes in the mental lexicon in terms of pairwise relationships (e.g. semantic similarity) between stored entities (Anderson & Bower, 1972; Collins & Quillian, 1969). Network science aimed at shedding light

on how the topology and distribution of multilevel relations holding between lexical nodes may affect high-level lexical processes such as priming, lexical retrieval, lexical association, lexical competition and cognitive search (see Castro & Siew, 2020, for a comprehensive overview). Connectionism and network science had different aims, the former focusing on low-level processing issues, the latter on the structural topology of underlying lexical representations. Nonetheless, they both model the mental lexicon as a network, while sharing a few fundamental algorithmic notions such as *spreading activation* (the idea that the activation of one node in memory can activate other connected nodes), *random walk* (the idea that observation of the information flow in a network can be used to understand the network's structure of connectivity) and *network growth* (the observation that the number of connections and the distribution of their weights change with increasing input and decreasing levels of network plasticity). In addition, networks were conducive to information-theoretical analyses of their internal structure, based on measuring a network's probabilistic expectation for a particular node to be activated given a history of activated nodes.

## 2 Morphological networks

Such a sweeping range of innovative principles took some time to find its way in language studies. In morphological theory, however, the network metaphor did not go unnoticed, and started influencing theoretical modelling in direct and indirect ways since the 90s (Bybee, 1995; Corbett & Fraser, 1993). The idea that linguistic structure can emerge from the self-organisation of unstructured input is nowadays key to understanding language acquisition (Hopper & Bybee, 2001; Ellis & Larsen-Freeman, 2006; MacWhinney, 1999; MacWhinney & O'Grady, 2015). Nonetheless, it had to await the challenging test of successful computer simulations before it was given wide currency in the psycholinguistic (Baayen et al., 2011) and theoretical literature (Blevins, 2016) on word structure. A recent conceptualisation of morphological generalisation known as the "Cell Filling Problem" (Ackerman et al., 2009; Ackerman & Malouf, 2013) hinges on modelling the implicative structure of morphological paradigms as a word network, capitalising insights and mathematical tools (e.g. conditional entropy) from network science. Exemplar-based machine learning models successfully operationalised analogy-based relations among fully stored lexical items, and questioned the need to resort to multiple levels of abstraction (Daelemans & Van den Bosch, 2005; Keuleers et al., 2007; Pirrelli & Yvon, 1999). Recent advances in distributional semantics (Baroni & Lenci, 2010; Padó & Lapata, 2007; Mitchell & Lapata, 2010) have thrown into sharp relief the role of lexical semantics in morphological processing, particularly for compounding and derivation (Marelli et al., 2017; Marelli & Baroni, 2015; Günther & Marelli, 2019), while helping draw a measurably graded distinction between derivation and inflection (Bonami & Paperno, 2018).

Most of these studies lie at the cross-road of neighbouring linguistic and cognitive disciplines, bearing witness to the fruitful prospects of an interdisciplinary integration of different approaches to language inquiry: from linguistic theories and cognitive models of human language processing, to computational and neuropsychological

language architectures. Of these disciplines, some are undergoing technological developments at a faster rate than others, and this makes it more difficult to take stock of the ways in which field-specific advances can sharpen our understanding of word knowledge in general. Technical details tend to obscure underlying assumptions, and theoretical implications are not always easy to spell out. The present issue was intended to report on recent progress in the interdisciplinary allegiance between neighbouring language disciplines.

Four different realms of morphological competence are explored here: lexicosemantics (with a specific focus on compounds), inflection, suprasegmental phonology and child language acquisition. Our goal was twofold. First, we aimed to understand the theoretical implications of data-driven, quantitative approaches for the individual morphological domain which they deal with. Secondly, we were wondering whether all these approaches, taken together, could delineate a new view of the role of the lexicon in language architecture, and, possibly, of morphology in language theory. It should be appreciated that each study in this issue takes a different approach, and makes use of a specific computational framework. Nonetheless, there are a few assumptions that all present contributions share. In this section, we anticipate what we consider the most important such assumptions using word semantics by way of exemplification of a general, integrative approach. Implications that are more specific of the four single realms are considered in detail in the following section.

The (mental) lexicon is *redundant*, *token-based* and *context-sensitive*. Focusing on the semantic lexicon, it is often forgotten that what we call the meaning of a word is in fact a convenient abstraction, a loosely defined concept associated with the use of the word in a variety of contexts (which, for a highly polysemous word such as *table*, can in fact be very different from one another). So-called “word embeddings” (e.g., Mikolov et al., 2013) operationalise this pre-theoretical definition by computing a word meaning as a vector representation, whose values are averaged across the contexts where the word was encountered. Although distributional vectors can be enumerated and independently stored, their computation is crucially token-based, and depends on local input conditions (e.g. time and context of input stimuli) in ways that require a rather profligate usage of memory resources. This makes the lexicon’s content redundant, but exquisitely sensitive to context and token frequency effects.

Lexical relations are *graded* and *analogy-based*. Due to their gradient, real-valued nature, semantic vectors can be used to process and understand the meaning of a novel word, by measuring the similarity of its context to stored representations. More generally, the fact that vector representations collectively define a continuous, multidimensional space leads naturally to the view that an individual word meaning is understood only in terms of its *relative* position to other word meanings in the semantic space (whose relevance is an inverse function of their distance from the target word). This resonates well with paradigm-based approaches to inflection, and with child’s acquisition of morphology as a continually expanding network of interconnected lexical nodes.

Lexical structure is *emergent* and *relational*. Structure emerges from discriminable patterns of variation exhibited by sets of related representations that are encoded by exemplars. For example, it is the systematic (spatial) relation between the semantic vector of a word occurring freely in a syntactic context and the semantic vector of

the same word occurring as a compound constituent that allows a speaker to perceive the interdependence between word meaning and compound structure. The ability of a speaker to infer the meaning of that word in a novel compound is grounded in the perception of this systematic relation.

Lexical representations are *optimised for processing*. Since vector distance is measured across very many dimensions, meaning perception is the result of *maximising* the advantages offered by both token-based representations (i.e. representations of individual contexts) and their averaged centroids (e.g. word embeddings), in keeping with Gary Libben's principle of *maximisation of opportunities* (Libben, 2006, 2010). In fact, while individual word usages are associated with specific contexts and feature clusters, more abstract representations are maximally discriminative and optimised for processing. Technically, this can be implemented in different ways, depending on the specific computational framework and the task at hand: through recurrent predictive connections in artificial neural networks, cue-outcome weights in Naive Discriminative Learning, information gain in exemplar-based learning, levels of connectivity in network science etc. Such a predictive bias appears to be one of the strongest drive in language processing and learning, and complies with one of the most deeply-rooted objectives driving the way the human brain responds to the environment: anticipating an incoming stimulus to facilitate adaptive functioning (Heilbron et al., 2022; Tanovic & Joormann, 2019).

### 3 The papers

The contribution of complex morphology to word meaning is the main focus of the first paper in this collection: "CAOSS and transcendence: Modeling role-dependent constituent meanings in compounds", by Fritz Günther and Marco Marelli. Here, the authors explore Gary Libben's notion of *morphological transcendence* by applying regression models of semantic compositionality to *distributional word vectors* (Mikolov et al., 2013).

According to the *morphological transcendence hypothesis* (Libben, 2014), morphological structure affects meaning, to the extent that words can take on position-specific meanings depending on how frequently they appear in either the first position (the modifier's slot) or the second position (the head's slot) of an English compound. For example, the specific bird-related meaning of *bill* when used as a head in compounds like *shoebill* and *hornbill* is only weakly connected with the meaning representation of *bill* as a free word. An implication of this hypothesis is that lexical representations in the mental lexicon should be sensitive to the specific structural context where a lexical entry occurs. Günther and Marelli lend considerable computational support to Libben's hypothesis, showing that transcendent meaning representations are predicted by a compositional model of semantic vectors (the CAOSS model, Marelli et al., 2017) where the word's vector representation of *bill* as-a-head in *hornbill* results from the product of *bill*'s semantic vector as a free-word with a matrix of position-dependent weights. In the end, the authors claim that their model dispenses with the need to store transcendent representations in the lexicon, as the latter can be derived by general, linear manipulations of free-word vectors. Although

this is a strong argument for having abstract vectors as well as algebraic operations applying to them, it is not clear how a lexicon containing these vectors/operations can get rid of token representations in an acquisitional perspective, while remaining compatible with Libben's principle of maximisation of opportunities.

In "Stratification effects without morphological strata, syllable counting effects without counts – modelling English stress assignment with Naive Discriminative Learning", Sabine Arndt-Lappe, Robin Schrecklinger and Fabian Tomaschek use Naive Discriminative Learning (NDL) to model stress assignment in the English orthographic lexicon. In NDL, learning is shaped by prediction and prediction error. Association weights between surface forms (represented as letter bigrams and trigrams) and stress placement are increased every time the predicted stress outcome co-occurs with the current cue, and are decreased whenever the predicted stress outcome does not occur. This gives rise to cue competition, through which cognitively plausible representations emerge.

Predicting English stress assignment is a notoriously difficult task, particularly in connection with the distinction between stress-preserving (e.g. *happiness*) and stress-shifting (e.g. *popularity*) derivational suffixes. According to stratal accounts (e.g. Kiparsky, 1982), stress-shifting suffixes are attached before phonological stress rules apply, stress-preserving suffixes are attached after stress rule application; hence the need for abstract lexical representations being stored in the mental lexicon. The paper provides clear evidence that English stress position can be learned successfully with no information of either morphological strata or syllable counting, based on the observed competition between stem-as-a-cue weights and suffix-as-a-cue weights. In stress-preserving derivatives, the relative balance of cue weights between the word's stem and its suffix is tilted towards the stem, while the opposite obtains for stress-shifting derivatives. This evidence supports a view of the morphonological lexicon which makes abstract lexical representations and syllable-rich information dispensable, while making testable predictions about the development of stress-related morphological categories in child language acquisition. Here, complex outcomes result from the strong interaction between parts of orthographic words and their stress patterns, providing a nice example of how word structure effects can emerge from densely interconnected surface representations.

A strongly related question is addressed by our paper "A discriminative information-theoretical analysis of the regularity gradient in inflectional morphology", which shows how orthographic forms are represented in a recurrent self-organising lexical map learning the conjugation system of a language. Here, following NDL principles, lexical access is modelled as consisting in *discriminating* between time-bound cues (e.g. a time-series of letters in dynamic competition for their predictive value) for a target lexical unit to be accessed. Between-cue competition proceeds through a continuous, incremental update of each cue's predictive bias, based on the number of times the cue is seen (or is not seen) be associated with the outcome.

Cue competition at learning time is shown to shape inter-node connections and proves to have far-reaching consequences on the processing behaviour of the lexical map. This mechanism appears to provide an explanatory link between the amount of competition in the input and the structural entropy of the forward connections emanating from a chains of activated nodes after training. Evenly distributed con-

nections create a balanced competition that maximises processing uncertainty. Conversely, when one forward connection is much stronger than other connections from the same node, one member of the family will be pre-activated more strongly than other members. We show that this fundamentally predictive bias can account for a variety of effects in the speakers' word processing behaviour, including their sensitivity to word frequency, paradigm entropy and perception of the inflectional (ir)regularity gradient. Having simulated these effects with superpositional patterns of node activation makes it hard to define the resulting network as either a lexicon or a word processor.

In "Explaining dynamic morphological patterns in acquisition using Network Analysis", Elitzur Dattner, Orit Ashkenazi, Dorit Ravid and Ronit Levie offer a *network analysis* of the development of morphological patterns across stages of acquisition of the Hebrew verb system in different contexts. In Hebrew, a verb token is a triple link between a *root* (a semantic concept), a *binyan* (a schematic event structure), and a *temporal pattern* (a specific reference to time and/or modality). Bipartite networks consisting of dyadic non-inflected combinations of a root consonantal skeleton (e.g. *k-t-b* 'write') and a bynian-specific temporal pattern (e.g. *Qal*.present *kotev* 'writes/writing', *Qal*.past *kataav* 'wrote'), are built based on corpus evidence sampled from different conversational settings and text sources. Unlike artificial neural networks, which are typically intended to model lexical processing/learning, dyadic morphological networks of this kind are used to model the emergence of the Hebrew verb system's structure across different ages, types of interaction (e.g. parent-child or child-child) and communication modes (oral vs. written language).

This is done by counting the number of nodes (both roots and patterns) in the network and their level of connectivity (or degree centrality), measured as the number of links that a node has with other nodes in the network. Accordingly, a morphological network with highly-connected nodes that are not interconnected will result in a repetitive lexicon, where roots tend to be associated with one pattern only. Conversely, a morphological network with a high number of interconnected important nodes makes the final state of a speech event less predictable, increasing the entropy and the productivity of the system. The paper shows that these changes characterise the developmental path of the morphological system of the Hebrew verb lexicon through time. Acquiring new forms is in fact necessary for (young) speakers to communicate in an ever changing social environment, as confirmed by evidence that language acquisition processes in young children are sensitive to the specific structure of the (language) environment they are exposed to. This is an important insight offered by Dattner and colleagues' paper, showing the limits of approaches to language learning that focus on the internal structure of a child's existing vocabulary only. The parental language input as well as the variety of sensory stimuli coming from their environment play a fundamental role in shaping the structure of a child's lexicon.

#### 4 Future prospects

Many open challenges and outstanding issues remain to be addressed. Here, we limit ourselves to mentioning two of them only. It has recently been argued (Jamieson et



al., 2022) that humans store individual experiences in episodic memory, and that abstractions such as conceptual categories and word meanings emerge solely during retrieval, e.g. as a by-product of the competitive activation of multiple lexical nodes. This view has also been forcefully advocated for language learning by Ambridge (2020), whose recent manifesto is a programmatic vindication of exemplar-based machine learning algorithms, which are argued to be consistent with empirical findings from the deep learning and neuroimaging literature. Network models, while being in principle compatible with the same evidence, do not subscribe to a radical exemplar view. Nonetheless, such a rekindling of interest in exemplar-based models as a domain-general framework for cognitive psychology looks like a promising arena for progress in this area.

This brings us to our second and final point. All present papers focus on one level of morphological analysis only. A real lexicon is expected to contain hundreds of thousands of nodes that are mutually related in myriad ways, at multiple levels of linguistic analysis. In this connection, the issue of scale is thus important not only for developing practical applications, but also to understand how principles of network structure address issues of optimally efficient lexical self-organisation, whereby lexical nodes can be accessed rapidly for effective communication to occur in real time. We look forward to seeing multilevel scale issues be addressed from such an integrative network perspective. The rationale for this convergence was presciently epitomised, back in the early 80s, by David Marr's (1982) hierarchy of levels of understanding of a complex processing system:

1. the *computational level* answers the "semantic" question of what a system does, by providing a precise characterisation of what types of functions and operations are to be computed for a specific cognitive process to occur;
2. the *algorithmic level* answers the "syntactic" question of how a system does what it does, by specifying how computation takes place in terms of detailed algorithmic steps and programming instructions;
3. the *implementation level* finally states how representations and processes are actually realised at the physical level, e.g. as electronic circuits or as patterns of neurobiological connectivity.

Recent advances in computational and neurocognitive approaches to language sciences have provided the level of material continuity between linguistic functions (level 1), algorithmic operations (level 2) and neuro-functional correlates (level 3) that is a necessary pre-condition to successful integration of language sciences along Marr's hierarchy (Alvargonzález, 2011). Interdisciplinary progress in this direction is well underway, and is likely to lead to a different understanding of traditional linguistic issues. This will probably require an effort to depart from the hypothesis of a *direct correspondence* (Clahsen, 2006) between modular components of the language architecture (lexicon vs. rules), processing correlates (memory vs. computation) and their neuro-anatomical localisation (prefrontal vs. temporo-parietal perisylvian areas of the left hemisphere). In the end, it may turn out that some ontological units of linguistic theory (e.g. stems, words or phrases) cannot be readily matched to the fundamental processing nodes that are central to a neural network architecture. Even a cognitive pillar such as the mental lexicon may call for a radical reappraisal (Elman,

2009). We concur with Poeppel and Embick (2005) that one promising solution to the lack of correspondence between Marr's levels may require shifting our focus away from language-specific units and representations to primitive functional and biological processes (e.g., storage, co-activation, competition, prediction and retrieval), with a view to investigating the role these processes play in both language and cognition across the mind and brain. In the end, one can argue that it is precisely the complex overlapping of these dynamic processes that is responsible for sophisticated effects on language processing. Simple processing/learning principles, operating across different time scales, can eventually yield complex outcomes.

**Acknowledgements** The authors express their deep gratitude to Ingo Plag for his constant support and guidance throughout the various phases of this editorial work.

**Funding** Open access funding provided by ILC – PISA within the CRUI-CARE Agreement.

## Declarations

**Competing Interests** The authors hereby declare that they have no conflict of interests.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Ackerman, F., & Malouf, R. (2013). Morphological organization: The low conditional entropy conjecture. *Language*, 89(3), 429–464.
- Ackerman, F., Blevins, J. P., & Malouf, R. (2009). Parts and wholes: Implicative patterns in inflectional paradigms. In J. P. Blevins & J. Blevins (Eds.), *Analogy in grammar: Form and acquisition* (pp. 54–82). London: Oxford University Press.
- Alvargonzález, D. (2011). Multidisciplinarity, interdisciplinarity, transdisciplinarity, and the sciences. *International Studies in the Philosophy of Science*, 25(4), 387–403. <https://doi.org/10.1080/02698595.2011.623366>.
- Ambridge, B. (2020). Against stored abstractions: A radical exemplar model of language acquisition. *First Language*, 40(5–6), 509–559.
- Anderson, J. R., & Bower, G. H. (1972). Recognition and retrieval processes in free recall. *Psychological Review*, 79(2), 97.
- Baayen, R. H. (2007). Storage and computation in the mental lexicon. In G. Jarema & G. Libben (Eds.), *The mental lexicon: Core perspectives* (pp. 81–104). Amsterdam: Elsevier.
- Baayen, R. H., Milin, P., Đurđević, D. F., Hendrix, P., & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review*, 118(3), 438–481.
- Baroni, M., & Lenci, A. (2010). Distributional memory: A general framework for corpus-based semantics. *Computational Linguistics*, 36(4), 673–721.
- Blevins, J. P. (2016). *Word and paradigm morphology*. Oxford: Oxford University Press.
- Bonami, O., & Paperno, D. (2018). Inflection vs. derivation in a distributional vector space. *Lingue e Linguaggio*, 17(2), 173–196.

- Bybee, J. (1995). Regular morphology and the lexicon. *Language and Cognitive Processes, 10*(5), 425–455.
- Castro, N., & Siew, C. S. (2020). Contributions of modern network science to the cognitive sciences: Revisiting research spirals of representation and process. *Proceedings of the Royal Society A, 476*(2238), 20190825.
- Clahsen, H. (2006). Linguistic perspectives on morphological processing. In D. Wunderlich (Ed.), *Advances in the theory of the lexicon* (Vol. 13, pp. 355–388). Berlin: Mouton de Gruyter.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8*(2), 240–247.
- Corbett, G. G., & Fraser, N. M. (1993). Network morphology: A datr account of Russian nominal inflection. *Journal of Linguistics, 29*(1), 113–142.
- Daelemans, W., & Van den Bosch, A. (2005). *Memory-based language processing*. Cambridge: Cambridge University Press.
- D'Esposito, M. (2007). From cognitive to neural models of working memory. *Philosophical Transactions of the Royal Society B: Biological Sciences, 362*(1481), 761–772.
- D'Esposito, M., & Postle, B. R. (2015). The cognitive neuroscience of working memory. *Annual Review of Psychology, 66*, 115–142.
- Ellis, N. C., & Larsen-Freeman, D. (2006). Language emergence: Implications for applied linguistics—introduction to the special issue. *Applied Linguistics, 27*(4), 558–589.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science, 14*(2), 179–211.
- Elman, J. L. (2009). On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive Science, 33*(4), 547–582.
- Günther, F., & Marelli, M. (2019). Enter sandman: Compound processing and semantic transparency in a compositional perspective. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 45*(10), 1872.
- Heilbron, M., Armeni, K., Schoffelen, J. M., Hagoort, P., & De Lange, F. P. (2022). A hierarchy of linguistic predictions during natural language comprehension. *Proceedings of the National Academy of Sciences, 119*(32), e2201968119.
- Hockett, C. F. (1954). Two models of grammatical description. *Word, 10*(2–3), 210–234.
- Hopper, P. J. & Bybee, J. L. (Eds.) (2001). *Frequency and the emergence of linguistic structure*. Amsterdam: Benjamins.
- Jamieson, R. K., Johns, B. T., Vokey, J. R., & Jones, M. N. (2022). Instance theory as a domain-general framework for cognitive psychology. *Nature Reviews Psychology, 1*(3), 174–183.
- Jordan, M. (1986). *Serial order: A parallel distributed processing approach*. Tech. Rep. 8604, University of California.
- Keuleers, E., Sandra, D., Daelemans, W., Gillis, S., Durieux, G., & Martens, E. (2007). Dutch plural inflection: The exception that proves the analogy. *Cognitive Psychology, 54*(4), 283–318.
- Kiparsky, P. (1982). *Lexical phonology and morphology. Linguistics in the morning calm*.
- Libben, G. (2006). Why study compound processing? An overview of the issues. In *The representation and processing of compound words* (pp. 1–22). London: Oxford University Press.
- Libben, G. (2010). Compound words, semantic transparency, and morphological transcendence. In S. Olsen (Ed.), *New impulses in word-formation* (pp. 317–330). Hamburg: Buske.
- Libben, G. (2014). The nature of compounds: A psychocentric perspective. *Cognitive Neuropsychology, 31*(1–2), 8–25.
- Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience, 17*(3), 347.
- MacWhinney, B. (Ed.) (1999). *The emergence of language*. London: Taylor & Francis.
- MacWhinney, B. & O'Grady, W. (Eds.) (2015). *The handbook of language emergence*. New York: Wiley.
- Marelli, M., & Baroni, M. (2015). Affixation in semantic space: Modeling morpheme meanings with compositional distributional semantics. *Psychological Review, 122*(3), 485.
- Marelli, M., Gagné, C. L., & Spalding, T. L. (2017). Compounding as abstract operation in semantic space: A data-driven, large-scale model for relational effects in the processing of novel compounds. *Cognition, 166*, 207–224.
- Marr, D. (1982). *Vision. A computational investigation into the human representation and processing of visual information*. New York: Freeman.
- Marzi, C., & Pirrelli, V. (2015). A neuro-computational approach to understanding the mental lexicon. *Journal of Cognitive Science, 16*(4), 493–535.

- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 26, pp. 3136–3144).
- Miller, G. A. (2003). The cognitive revolution: A historical perspective. *Trends in Cognitive Sciences*, 7(3), 141–144.
- Mitchell, J., & Lapata, M. (2010). Composition in distributional models of semantics. *Cognitive Science*, 34(8), 1388–1429.
- Padó, S., & Lapata, M. (2007). Dependency-based construction of semantic space models. *Computational Linguistics*, 33(2), 161–199.
- Pinker, S., & Prince, A. (1994). Regular, and irregular morphology, and the psychological status of rules of grammar. In S. D. Lima, R. Corrigan, & G. K. Iverson (Eds.), *The reality of linguistic rules* (pp. 321–351). Amsterdam: Benjamins.
- Pinker, S., & Ullman, M. T. (2002). The past and future of the past tense. *Trends in Cognitive Sciences*, 6(11), 456–463.
- Pirrelli, V., & Yvon, F. (1999). The hidden dimension: A paradigmatic view of data-driven nlp. *Journal of Experimental & Theoretical Artificial Intelligence*, 11(3), 391–408.
- Poeppel, D., & Embick, D. (2005). Defining the relation between linguistics and neuroscience. In A. Cutler (Ed.), *Twenty-first century psycholinguistics: Four cornerstones* (Vol. 1, pp. 103–120). Hillsdale: Erlbaum.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533–536.
- Sreenivasan, K. K., Curtis, C. E., & D’Esposito, M. (2014). Revisiting the role of persistent neural activity during working memory. *Trends in Cognitive Sciences*, 18(2), 82–89.
- Tanovic, E., & Joormann, J. (2019). Anticipating the unknown: The stimulus-preceding negativity is enhanced by uncertain threat. *International Journal of Psychophysiology*, 139, 68–73.
- Ullman, M. T. (2001). The declarative/procedural model of lexicon and grammar. *Journal of Psycholinguistic Research*, 30, 37–69.
- Ullman, M. T. (2004). Contributions of memory circuits to language: The declarative/procedural model. *Cognition*, 92(1–2), 231–270.
- Wilson, M. (2001). The case for sensorimotor coding in working memory. *Psychonomic Bulletin & Review*, 8(1), 44–57.

**Publisher’s Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.