English Language and Linguistics (2026), 1–6. © The Author(s), 2026. Published by Cambridge University Press
doi:10.1017/S1360674325100646

CAMBRIDGE UNIVERSITY PRESS

BOOK REVIEW

Giulia Rambelli, *Constructions and compositionality: Cognitive and computational explorations* (Cambridge Elements in Construction Grammar). Cambridge: Cambridge University Press, 2024. Pp. iv + 96. ISBN 9781009517393.

Remi van Trijp

Sony Computer Science Laboratories, Paris Lab, France
Email: remi.vantrijp@sony.com

(Received 25 November 2025)

Linguistic *productivity* is one of the greatest mysteries of human language: how is it possible that humans, given a *finite* inventory of linguistic constructions, are able to understand and produce an *infinite* number of novel utterances? Once brought to the forefront of modern linguistics by Chomsky's (1957) *Syntactic Structures*, this question has become less central in the current age of data-driven methods and usage-based linguistics, which concentrates on inductively extracting linguistic patterns from corpora. In her bite-sized Cambridge Element, Giulia Rambelli puts the spotlight on the question of productivity again, and why it still matters.

Rambelli's contribution is timely, as it fills a gap in the constructional literature that feels increasingly urgent. Even though the question of productivity (and, indeed, creativity) has been gaining traction in recent years, the majority of constructional analyses have remained mostly silent about exactly how constructions interact with each other and with cognition more generally. Fitting for the Cambridge Elements series, which fills the niche between concise research articles and extensive monographs, Rambelli's book gets readers up to speed with the topic's key approaches and challenges, all within the span of an evening's worth of reading.

The book opens with chapter 1, 'Introduction' (pp. 1–11), which does an excellent job of situating Construction Grammar's position on productivity and creativity within the field at large. Rambelli frames the discussion in terms of two broad perspectives on how language users understand utterances. In the traditional view, productivity is seen in terms of *compositionality*. Lexical items provide conceptual content, and syntax arranges this content in hierarchical structures. Meaning is then inferred from the linguistic structure through a separate component for semantic interpretation. This view has been criticized, especially in light of psycholinguistic and neurolinguistic evidence, suggesting that semantic expectations, among other factors, drive interpretation in a more *top-down* manner than the strict *bottom-up* composition posited by syntactic rules. Rambelli also notes that, in recent years, the debate has shifted towards viewing both perspectives as potentially complementary, though it remains unclear how they can be meaningfully integrated.

Researchers working in either perspective may find this dual characterization lacking in nuance, but given the scope of a Cambridge Element, I believe it offers a fair description that achieves its purpose: to situate the constructional approach within a broader landscape. Construction Grammar assumes that the primary units of language are constructions – mappings between form and function. Construction Grammar is compositional in that different constructions may combine with each other, yet it also has non-compositional

https://doi.org/10.1017/S1360674325100646 Published online by Cambridge University Press

aspects, as it assumes a *holistic* view in which not all meaning can be fully predicted from a construction's parts. As such, Rambelli presents Construction Grammar as a promising approach for reconciling the compositional and non-compositional perspectives. Rambelli also avoids the pitfall of presenting Construction Grammar as the final solution, instead highlighting open questions that the field continues to explore.

Rambelli then points out that the constructional literature distinguishes between *productivity* and *creativity*. Productivity refers to the creation and interpretation of novel utterances using the existing inventory of constructions – which Sampson (2016) calls *F-creativity* (fixed creativity). In the early years of generative linguistics, which assumed a strong autonomy of syntax and an idealized speaker-listener, this was the only type of creativity under investigation. The term *creativity*, then, is typically reserved for cases that extend the linguistic inventory to produce something truly novel, for instance by coining a new word – which Sampson (2016) defines as *E-creativity* (extending/enlarging creativity). The precise distinction between productivity and creativity is still a matter of debate, as well as the question of how a language user's existing grammar is employed in creative utterances (Hoffmann 2022: 280). Despite these disagreements, the distinction remains intuitively clear, offering readers a strong sense of what the remainder of the book sets out to explore.

Chapter 2, 'The problem of compositionality as a processing principle' (pp. 11–29), discusses the *Principle of Compositionality* in the tradition of Gottlob Frege (Partee 2004). In Fregean compositionality, the meaning of an utterance is determined by the meanings of its parts and how those parts are syntactically combined. This chapter is particularly valuable, as many cognitive linguists only have a strawman conception of compositionality as 'the meaning of the whole is the *sum* of its parts'. Rambelli's introduction convincingly shows that compositionality is a much more multifaceted concept with many arguments in its favor, and achieves a rare balance of accessibility, sprinkled with the occasional formal detail.

That said, Rambelli also subjects the principle of compositionality to criticism that, in my view, fails to distinguish the principle itself from how it has been incorporated in various theories. For instance, she writes that 'it is unclear how [formal representations] precisely capture the intricate nuances of the meanings of constituent expressions', and that 'Formal Semantics does not encode the full spectrum of human linguistic experiences' (p. 19). She thereby echoes a common misunderstanding of symbolic knowledge representations. Symbols are placeholders that stand for something – pointers, if you will – not the thing itself. To illustrate, the written word *cat* does not capture all the nuances of the feline world, but humans can *interpret* that symbol by mapping it to their knowledge about and experiences with cats.

This misunderstanding is common because many formal approaches – especially the pioneering ones between the 1950s and 1970s – have tried to use symbolic representations for capturing all semantically relevant distinctions. But, as Rambelli points out, such an approach makes it virtually impossible to distinguish between words like *man* and *dude*: both refer to male humans, but have quite different connotations and usage conditions. But ever since the philosopher and logician Richard Montague (posthumously) established formal semantics, a distinction is assumed between intensional and extensional semantics (Partee 2004). A logical term like *cat(x)* only represents the intensional meaning of 'cat', while its actual meaning in context (its extension) must still be computed through *semantic interpretation* – the process in which linguistic knowledge is mapped onto a non-linguistic domain. The same duality also exists for usage. In a neurosymbolic language model, for instance, the symbols 'man' and 'dude' can be translated into their respective vector representations (Verheyen *et al.* 2025). These vector representations, learned from distributional patterns and usage, allow more fine-grained semantic computation, while the symbolic representations excel at precision and composition. Indeed, the supposed opposition between symbolic

https://doi.org/10.1017/S1360674325100646 Published online by Cambridge University Press

and non-symbolic representations overlooks their complementary nature. In reality, both simply offer different views that afford different kinds of inferences.

The distinction between intensional and extensional meaning matters because one of the criticisms of compositionality is that the meaning of an expression's parts is (supposedly) fixed locally, and therefore unaffected by the larger context. Yet expressions such as *dark blue* show that this is not the case: our interpretation of *blue* is modulated by the modifier *dark*. This kind of modulation, however, poses no problem for a compositional approach. The modifier *dark* simply modulates the interpretation conditions of *blue* – that is, its intensional meaning – by shifting its value along the brightness dimension. When the actual meaning of the expression – its extension – is computed through semantic interpretation, this shift is reflected in the resulting color concept, demonstrating that context-sensitivity does not violate compositionality but operates at different levels of meaning computation (Partee 2004).

Given this background, the chapter is largely correct in identifying the limits of Fregean compositionality as it has been implemented in *most* theories, but it misses an opportunity to explore its potential within a constructional framework. More specifically, Rambelli discusses how Sign-Based Construction Grammar (SBCG) handles compositionality, yet SBCG is something of an outlier within the Construction Grammar community, as it remains fundamentally *lexicalist*. Lexicalist approaches encode most grammatical information in the lexicon and project it onto syntactic structures through general principles. In doing so, they forego much of the expressive power of Construction Grammar, which treats grammatical constructions as sources of top-down information that cannot be predicted from lexical items alone. Rambelli rightly notes that SBCG is a process-neutral model and therefore makes no predictions about actual processing. Instead, she briefly discusses two computational frameworks: Fluid Construction Grammar (FCG) and Embodied Construction Grammar (ECG). This inclusion is particularly welcome, as these frameworks represent the few constructional models with fully operational computational implementations. However, and I must acknowledge my bias here as one of the main developers of Fluid Construction Grammar, the description of FCG as relying on 'truth-conditional first-order predicate calculus' (p. 24) is somewhat misleading. While FCG's formalism indeed uses logic as a representational language for feature-value descriptions, its semantics is *embodied* and *context-grounded*: meaning construction involves mapping linguistic structures onto perceptual, conceptual and pragmatic representations that emerge dynamically during interpretation. In fact, one of FCG's central research threads over the past decades has been the study of context-sensitive meaning composition and situated interpretation (Beuls & Van Eecke 2024).

These comments aside, the chapter does an outstanding job of showing how compositionality has been implemented across different theories, which assumptions they make, and what benefits and drawbacks these entail. It concludes with a section on how distributional approaches handle compositionality, which will be of particular interest to readers concerned with how Large Language Models handle – or still struggle with – compositional meaning.

Chapter 3, 'Accessing meaning noncompositionally: Insights from experimental data' (pp. 29–45), examines evidence from psycho- and neurolinguistic experiments on how language users construct meaning without relying on syntactic composition. This chapter deserves praise not only because it offers experiment-driven insights, but also because it does so constructively: the evidence is presented not to undermine the Fregean tradition, but to reconcile different perspectives on sentence processing.

The chapter opens with a discussion on idiomatic expressions and multiunit sequences, which are widely regarded as problematic for compositionality and are typically represented as lexical units in grammatical theories. Recent evidence suggests that idiomatic

https://doi.org/10.1017/S1360674325100646 Published online by Cambridge University Press

expressions are indeed retrieved holistically – as single units – but that compositional processing can still play a role, depending on factors such as a language user's familiarity with the idiom. Rambelli then turns to multiword expressions, which are formulaic expressions that remain transparent in meaning, such as *in the meantime* or *bride and groom.* Again, there is ample evidence that such expressions are stored as chunks, though their retrieval and production depend on several factors, including schematicity and frequency. Beyond idioms and formulaic evidence, additional evidence against a purely compositional view comes from studies showing that language processing is often 'shallow': it is frequently good enough to yield correct interpretations, yet may also produce *semantic illusions* in which language users fail to recognize an inaccuracy or inconsistency in text. Complementary findings highlight the predictive nature of language comprehension, in which accurate parsing is facilitated by anticipating the upcoming input. Finally, a large body of psycholinguistic research supports the view that language users routinely draw on background knowledge and past experiences during processing – evidence that meaning construction is deeply interactive rather than purely compositional.

Rambelli argues that this body of evidence aligns naturally with constructional approaches to language. The observation that idioms and formulaic expressions form a continuum has long challenged traditional assumptions of a strict syntax–lexicon divide – an issue that poses difficulties for traditional theories, but is readily accommodated in Construction Grammar, which provides a uniform representation for all linguistic knowledge. Likewise, the finding that language users integrate all sources of information at their disposal – linguistic, conceptual and contextual – fits neatly with the constructional view of language as an interplay of interacting constraints. Rambelli cautions, however, that Construction Grammar has so far concentrated primarily on the representation of individual constructions, and has only recently begun to investigate the dynamic processes involved in their use.

Chapter 4, 'Explaining productivity through analogy' (pp. 45–59), marks the point where the Element turns from literature review into a constructive proposal. If syntactic compositionality is not the sole source of productivity, what other processes might play a role? According to Rambelli, the answer lies in domain-general cognitive processes – most notably *analogical reasoning.* Rambelli grounds her discussion in the analogy-as-structural-mapping tradition pioneered by Dedre Gentner (1983), where analogy is defined as mapping the relational structure of a source domain onto a target domain. A classic example is *the atom is like the solar system*, which makes sense of one domain (the atom) by invoking the relational structure of another (the solar system), highlighting shared roles – a central body (nucleus/sun) with orbiting bodies (electrons/planets) – while ignoring elements that do not correspond.

While the general exposition of analogy is admirably clear, the discussion becomes less coherent when Rambelli turns to its role in language use. Section 4.2, 'Analogy in language use', begins with a compelling premise: if not through syntactic composition, novel utterances can be constructed and interpreted through analogical reasoning based on our existing inventory of constructions. What follows is a comprehensive survey of the increasing body of work in usage-based linguistics and psychology that explores this idea. Where clarity falters, however, is in Rambelli's treatment of *exemplar-based* models – warning that a combination of exemplars with analogy 'can lead, in its most extreme version, to a complete repudiation of any form of abstraction' (p. 54), which she qualifies as 'not coherent with our cognitive architecture' (p. 55). Since Rambelli only discusses a very specific model – Ambridge's (2020) radical exemplar account – it is unclear whether the critique is limited to that model or whether she rejects exemplar-based approaches more generally, based on the assumption that linguistic knowledge requires explicit abstraction (that is, storage of *schematic* constructions).

https://doi.org/10.1017/S1360674325100646 Published online by Cambridge University Press

This part of her book would therefore have benefitted from a short explanation of how analogy- and exemplar-based models relate to each other. Both rely on similarity-based reasoning, but they compute different kinds of similarity. Exemplar-based models only extrapolate from *local* information – nearest neighbors – whereas analogy allows for non-neighbors to play a role as well, leading to important algorithmic and empirical differences. Analogical reasoning requires alignment of relational structures across domains – such as cause-effect, containment or hierarchy – where exemplar-based models require only local structural alignment. This locality is what allows exemplar-based models, such as Memory-Based Language Processing (Daelemans & Van den Bosch 2005), to make efficient generalizations not through abstraction, but by reusing the most similar stored instances when encountering something new. Both analogical reasoning and exemplar-based classification are, in my view, highly relevant and complementary; and the question of whether abstraction is needed only becomes pertinent for computing non-local, cross-domain structural mappings.

Fortunately, this remains a minor shortcoming in an otherwise strong chapter. What is particularly refreshing is Rambelli's survey of how analogy has been operationalized in computational models. As a computational linguist, I have long found it frustrating that linguistics and computer science often lead parallel lives with relatively little interaction, so it is gratifying to see a constructional account that actively bridges these domains. Rambelli highlights a range of approaches that make clear how different computational traditions formalize the mechanisms underlying analogical and similarity-based reasoning. By presenting these models side by side, she not only demonstrates the breadth of work on analogy, but also shows that the concept can be implemented in a variety of precise, testable ways.

Chapter 5, 'Rethinking compositionality: A constructivist perspective' (pp. 59–70), weaves together the book's various threads into a bold vision of productivity and creativity. Rather than treating compositionality as a property of how syntactic rules combine, Rambelli proposes that novel utterances are produced and comprehended through processes of efficient retrieval and adaptation through similarity and analogy.

Taken as a whole, *Constructions and Compositionality* succeeds in reopening a question that has too often been treated as settled: how language users manage to be endlessly productive with finite means. Rambelli offers a clear and engaging synthesis of the major strands of work on compositionality, analogy and usage, while pushing Construction Grammar toward a more explicit engagement with processing and computation. Even where certain distinctions might have been drawn more sharply, the book consistently invites the reader to think across theoretical boundaries and to consider how different levels of explanation – formal, cognitive and computational – might be integrated. For a field that has long wrestled with the relation between structure and use, Rambelli's Element provides both a lucid overview and a constructive nudge toward future work.

References

Ambridge, Ben. 2020. Against stored abstractions: A radical exemplar model of language acquisition. *First Language* 40(5–6), 509–59.

Beuls, Katrien & Paul Van Eecke. 2024. Humans learn language from situated communicative interactions. What about machines? *Computational Linguistic* 50(4), 1277–1311.

Chomsky, Noam. 1957. *Syntactic structures*. The Hague and Paris: Mouton de Gruyter.

Daelemans, Walter & Antal Van den Bosch. 2005. *Memory-based language processing*. Cambridge: Cambridge University Press.

Gentner, Dedre. 1983. Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7, 155–70.

Hoffmann, Thomas. 2022. *Construction Grammar*. Cambridge: Cambridge University Press.

Partee, Barbara H. 2004. *Compositionality in formal semantics: Selected papers*. Oxford: Blackwell.

https://doi.org/10.1017/S1360674325100646 Published online by Cambridge University Press

Sampson, Geoffrey. 2016. Two ideas of creativity. In Martin Hinton (ed.), *Evidence, experiment and argument in linguistics and the philosophy of language*, 15–26. Berlin: Peter Lang.

Verheyen, Lara, Jonas Doumen, Paul Van Eecke & Katrien Beuls. 2025. You shall know a construction by the company it keeps: Computational Construction Grammar with embeddings. In Claire Bonial, Melissa Torgbi, Leonie Weissweiler, Austin Blodgett, Katrien Beuls, Paul Van Eecke & Harish Tayyar Madabushi (eds.), *Proceedings of the Second International Workshop on Construction Grammars and NLP*, 75–83. Düsseldorf: Association for Computational Linguistics.

https://doi.org/10.1017/S1360674325100646 Published online by Cambridge University Press