1 Introduction

Understanding and generating novel sentences is often considered the hallmark of human language. It is striking how, given a finite inventory of means, speakers can always process and produce new expressions through which they can convey a new trend, realize a poem, express emotions, or produce scientific theories. In that sense, language is a remarkable example of human creativity. Theories of language refer to this phenomenon as the *productivity* (sometimes named *creativity*) of language, and they usually explain this by the fact that language processing is driven by a mechanism that composes the meaning of words into larger semantic units to create novel combinations.

Nevertheless, the influence of context on the interpretation or generation of a sentence is equally characteristic of language understanding. People have a great deal of previous linguistic and extralinguistic experience they use to build rich, elaborated representations of texts and conversations. In each utterance or discourse, each word acts as a cue to "activate" and retrieve related background knowledge, creating anticipations (or expectations) about sentence completion. In a general sense, this trait could be designated as the *context-sensitivity* of language: Sentence comprehension results from how preexisting linguistic knowledge and contextual constraints combine to make a particular linguistic form.

The following examples illustrate these observations:

(1) a. *The child spilled the milk all over the floor.*
b. *The child spilled the rice all over the floor.*

Both sentences are syntactically well-formed and generate meaningful semantic representation; that is, they are both semantically plausible. The first utterance includes nonnovel combinations of words: *Milk* is something we experienced to *spill* very often, and specifically, *spill the milk* is a chunk of words recurrent in texts. Usually, interpreting this kind of sentence eases the processing determined by several rationales (the expression is stored in long-term memory, the words match semantic expectations, etc.). Conversely, the situation reflected in the second sentence is unexpected. A comprehender could have never directly experienced that scene or heard this sequence of words (which is quite rare[1]). Therefore, interpreting this unusual utterance should rely on a specific process. While the classic mechanism proposed in the literature relies on a building-block strategy, where the final meaning is somehow built

[1] Observing the bigram frequencies in a large corpus of English – the enTenTen18 corpus (Jakubíček et al., 2013) – the word *milk* occurs as the object of the verb *spill* 2,392 times, while *spill* * *rice* occurs just 37 times.

from single utterance components, an alternative hypothesis is to characterize this process as a generalization based on similar previous experiences driven by other inference mechanisms (co-activated network of representations, analogical inferences, and so forth). Despite the vast amount of research in sentence processing offered by linguistic and psycholinguistic literature, it is still debatable how language users deal with the challenge of interpreting sentences presented in real time, incrementally, word by word. Accordingly, efforts are focused on formalizing a linguistic theory that can provide an adequate description and computational model of language processing.

This Element (1) proposes a review of what compositionality has been defined in general and its concrete transposition in different formalisms, (2) summarizes several experimental works that demolish (or largely downsize) the role of compositionality in language processing, (3) introduces analogy as a mechanism that could be used to build meaning, and (4) defines the role of compositionality in a usage-based constructionist perspective. In all cases, the theoretical perspectives, experimental data supporting these observations, and the potential applications to a computational model of language processing are presented.

From a broad perspective, the Element provides an overview of mechanisms proposed to explain language comprehension and their potential integration into a comprehensive theory of language processing. Before delving into specific details, the next section will introduce the two concurrent general mechanisms that linguistic theories have posited as the driving force of comprehension.

1.1 The Dual-Route Access to Meaning

The traditional view of language understanding relies on *compositionality*: The meaning of a sentence (or a discourse) is a function of the meaning of its constituents and the way they are syntactically combined (Partee, 1995). More precisely, the associated theoretical strategy sustained in the generative tradition (Chomsky, 1965; Pinker, 1999) is that syntax provides the primary mode of meaning combination in language: Syntactic rules do no more than determine which symbol sequence works as units for syntactic purposes, while meaning derives from the lexical conceptual structure. In other words, reading a sentence consists of linearly accessing the meaning of the words stored in the lexicon and then integrating them within the abstract hierarchical structure. This stance encourages a *bottom-up*, or building block, model of meaning where the interpretation mechanism is incremental: The meanings of the words are composed into syntactic units and aggregated until reaching a complete interpretation.

Nevertheless, there is extensive experimental evidence from psycholinguistic and neurolinguistic research against traditional compositionality (Baggio, 2021; Baggio, Van Lambalgen, & Hagoort, 2012; Ferreira, Bailey, & Ferraro, 2002; Ferreira & Patson, 2007; Mollica et al., 2020, among others). These findings can be synthesized into two fundamental observations (Baggio, 2018, p. 19):

- A comprehender generates interpretations based on the semantic relations between words, not necessarily encoded or reflected by the grammar, and
- A comprehender tends to generate semantic representations that could bypass or collide with syntactic analyses, resulting in superficial and even inaccurate interpretations.

These observations argue for a more *top-down* model, in which our semantic expectations drive comprehension: The linguistic elements function as cues to activate our extensive linguistic knowledge stored in the long-term memory. This knowledge encompasses various facets, such as the frequency of use of certain expressions, common associations, and event schemas. Consequently, linguistic knowledge transcends mere lexical components and transforms into the "cognitive organization of one's experience with language" (Bybee, 2006, p. 771). Moreover, experimental studies have revealed a consistent pattern: When the comprehender accesses information congruent with their preexisting or pre-activated knowledge, facilitation effects in processing occur (faster reading times, low N400, etc.). The linguistic theories rooted in these assumptions endorse a *noncompositional* mechanism of meaning interpretation. They advocate a model of language where linguistic knowledge comes from direct linguistic experience and sentence processing is constraint-based, probabilistic, and reliant on expectations.

Although the ongoing debate has opened the possibility that these two strategies (compositional and noncompositional access to meaning) are not mutually incompatible, few linguistic theories have thus far bothered to elucidate how the two mechanisms could be integrated within a unified framework of sentence processing. Notably, theories within the fields of psycholinguistics, theoretical linguistics, and computational linguistics have offered different perspectives on uncovering the cognitive systems behind language, describing their characteristics, and modeling them. On one end, experimental results from different research areas have shed light on the mechanisms that could underlie human language comprehension. However, our knowledge about the language system remains scattered: Psycholinguistic studies usually focus on language processing subtasks (e.g., lexical access) or modules (e.g., morphology, syntax) without being aggregated into a unified framework. Conversely, linguistic

theories provide a rigorous formalism for language description, but they still have difficulty integrating the variability of language productions observed in behavioral experiments. To this day, it remains challenging to find in the existing literature a comprehensive model that unifies the different observations on language understanding into a unique architecture (Blache, 2017). However, there is a family of linguistic theories whose fundamental assumption could make it possible to integrate findings from multiple research fields.

1.2 A Constructionist View of Language

The behavioral evidence surrounding noncompositionality points to a model of linguistic representation that is in line with the assumptions of the usage-based models of language (Bybee, 2010; Croft, 1991, 2001; Langacker, 1987; Tomasello, 2009) and the Construction Grammar (CxG) paradigm (Hilpert, 2019; Hoffmann, 2022b; Hoffmann & Trousdale, 2013; Ungerer & Hartmann, 2023). CxG refers to a family of models based on the assumption that grammar is more than simply a formal system consisting of stable but arbitrary rules for defining well-formed sequences. Besides their specificities, all constructionist theories agree on a fundamental claim: Grammar consists of meaningful and symbolic form–meaning mappings, called *constructions* (Goldberg, 1995, 2006, 2019). The definition and operationalization of "construction" are still under debate, and each formalism proposes a slightly different criterion (see Ungerer & Hartmann, 2023 for a recent discussion on the definition of constructions). In the more general sense, constructions are processing units or chunks, from morphemes or words, to partially and fully lexicalized expressions, to schematic and productive patterns of language – such as Passive or Ditransitive constructions (Goldberg, 2003), up to even genres and text types (Hoffman & Bergs, 2018). For instance, the comparative correlative construction (Hoffmann, Brunner, & Horsch, 2020; also "Covariational Conditional"; cf. Culicover & Jackendoff, 1999) "The Xer, the Yer," such as

(2) *The more you know, the less you understand.*

has specific syntactic and semantic properties. First, both clauses are introduced by an element that "resembles" the English definite article *the*, which instead is followed by comparative phrases. Semantically, English speakers identify the cause–effect relationship between the two clauses, which is not marked at any syntactic level. Thus, both the syntax and the meaning of the construction are not entirely predictable by any abstract rule.

At the same time, a syntactic pattern like the Double Object construction has a meaning independently of the words that compose the construction. People reading a sentence like

(3) *She mooped him something.*

interpret the made-up word *moop* as "to give" because the abstract pattern V Obj_1 Obj_2 communicates itself the concept of transfer between two persons (Goldberg, 2019, p. 29).

Constructions form a structured inventory of a speaker's knowledge of the conventions of their language, called the *constructicon* (Diessel, 2023): Each construction constitutes a node in the taxonomic network of constructions, and taxonomic relations allow us to distinguish different types of grammatical knowledge. However, there is no complete agreement about how such taxonomies emerge. Formal models such as Sign-Based CxG (Sag, 2012) assume that only "idiosyncratic morphological, syntactic, lexical, semantic, pragmatic or discourse-functional properties must be represented as an independent node in the constructional network in order to capture a speaker's knowledge of their language" (Croft & Cruse, 2004, p. 263). Conversely, usage-based approaches – that is, Radical CxG (Croft, 2001) or Cognitive CxG (Goldberg, 2003) – advocate that constructions can be of any linguistic pattern used enough to be memorized (or *entrenched*; cf. Blumenthal-Dramé, 2012) in the long-term memory (Goldberg, 2006). Specifically, the assumption is that linguistic units that are more frequently encountered become more accessible and are preferred. According to this thesis, the most entrenched linguistic units tend to shape the language system in terms of patterns of use, at the expense of less frequent and thus less well-entrenched words or phrases. This account opens to a more *redundant* view of the lexicon: Although we do not technically need to memorize the word form *cats* because, in principle, it can be formed with a productive rule *cat+s*, it could be memorized in the lexicon because we have encountered it thousands of times in everyday language (Hilpert, 2021, p. 21). The same observation could be done for phrases. For instance, *read a book* is a semantic transparent chunk; that is, we can identify the meanings of its parts and how they were combined to generate the final interpretation. Even though we could build the meaning of the expression "on the fly" using a compositional, incremental mechanism, this sequence was heard and used so many times to be stored as a whole; thus, interpretation becomes the act of retrieving the stored meaning. Using Goldberg's words, "memory is cheap. There is a good deal of evidence that we retain an enormous amount of information about the language(s) we witness" (Goldberg, 2019, p. 54).

However, that does not mean that people retain frequently observed word combinations as atomic units (as this would quickly result in a combinatorial explosion), but memory traces have an internal structure. Thus, representations of related memories overlap neurally, mitigating the concern about a

combinatorial explosion (Goldberg, 2019). Moreover, the cognitive capacity of pattern detection and schematization (Bybee, 2010) allows the storage of more abstract construction from specific instances. Going back to the previous example, speakers redundantly store frequent plural forms in addition to a general plural construction, or an entire expression together with the abstract transitive pattern.

Placing constructions as the fundamental unit of language has the consequence of blurring the distinction between words of the lexicon and the rules of grammar. Contrary to generative theories, CxG argues that the architecture of the grammar is not layered in distinct modules, but different properties (morphological, prosodic, syntactic, semantic) together constitute the form that allows the construction to be identified, and when it is recognized, it is possible to access the associated meaning directly. This holistic view emphasizes the importance of surface structure, that is, the concrete utterances that a hearer is exposed to, as opposed to mainstream generative grammar, which primarily focuses on hidden syntactic processes not directly observable in the final output (Goldberg, 2013). As a joint representation of syntax and semantics, constructions provide a powerful mechanism for investigating many different linguistic phenomena (Diessel, 2019).

Despite the vast possibilities this framework offers for linguistic description and language modeling, some issues are still to be addressed. For instance, while we agree with the assumption that the lexicon is a repository of constructions, it is unclear which factors drive the memorization of specific chunks. The question about which constructions are stored in long-term memory and which aspects can be constructed online in working memory is yet to be fully answered and has consequences on the mechanisms governing sentence processing. In that regard, one more issue has to be figured out: What is the most appropriate and acceptable representation of constructional *meaning*? Toward a complete model of language comprehension, a further challenge is to give a semantic representation that could be coherent with the usage-based perspective and could account for the evidence that lexical knowledge is quite detailed, often idiosyncratic and verb specific, and often accessible at the earliest possible stage in sentence processing. Finally, while the majority of approaches have focused on grammatical description, only a few efforts have been carried on to operationalize CxG as computational modeling, with the exception of Fluid CxG (Steels, 2011) and Embodied CxG (Bergen & Chang, 2013).

To conclude, the core of this Element is investigating the relationship between constructions and compositionality, as the title of this work suggests. Indeed, constructionist approaches do not refute compositionally per se, but

reformulate this paradigm in terms of the combination of constructions ("weak compositionality"; cf. Michel, 2023, p. 566):

> By recognizing the existence of contentful constructions we can save the compositionality in a weakened form: the meaning of an expression is the result of integrating the meaning of the lexical items into the meanings of constructions. (Goldberg, 1995, p. 16)

Consequently, the CxG framework turns out to be the best way to unify the different compositional and noncompositional mechanisms observed in language due to its key assumptions. Nevertheless, there is no consensus regarding the manner in which constructions interact with each other to license a specific utterance (Boas, 2021, p. 64). Overall, this Element is focused on reviewing the different insights about language processing, which aspects are already depicted in CxG, and which ones still need to be addressed by future research.

Theoretical Claims of Construction Grammar

Adapted from Goldberg (2013, p. 15–16)

1. **Language as language use**. Our linguistic knowledge comes from linguistic experience: our lexicon and grammar are shaped by repeated exposure to specific utterances.
2. **Construction are the fundamental units of language**. Constructions are conventionalized associations of a form and a function, which apply not only to words but also to syntactic structures, thus guaranteeing a certain uniformity of representation of linguistic facts.
3. **The importance of the surface structure** Meaning is directly associated with surface structure, without derivations or transformations.
4. **The construct-i-con**. Grammar is a network of constructions, hierarchically organized through inheritance relations.
5. **There is a continuum from what is stored to what is processed**. There is no dichotomy between interpreting stored linguistic units and assembling expressions "on the fly"; there is just a *continuum* from stored items, highly predictable sequences, and completely compositional ones.
6. **Meaning emerges through context**. The meaning of a construction is inherently rooted in its contexts of use.

1.3 Compositionality, Productivity, Creativity

As introduced in the first paragraph of this Element, the surprising fact about language is that people can constantly generate (and understand)

never-ever-produced utterances. Chomsky defined this property as the "creative aspect" of language: "[A]n essential property of language is that it provides the means for expressing indefinitely many thoughts and for reaching appropriately in an indefinite range of new situations" (Chomsky, 1965, p. 6). More recently, Adger affirmed:

> The fact that sentences hardly reoccur shows us that we use our language in an incredibly rich, flexible, and creative way . . . . Virtually every sentence we utter is novel. New to ourselves, and quite often new to humanity. We come up with phrases and sentences as we need to, and we make them express what we need to express. We do this with incredible ease. We don't think about it, we just do it. We create language throughout our lives, and respond creatively to the language of others. (Adger, 2019, p. 2)

Chomsky and the generative tradition thus seem to suggest that linguistic creativity is "combinatorial" and "productive" (Bergs, 2018, p. 278): It involves creating something entirely new using existing rules in almost infinite ways. The success of the principle of compositionality thus relies on its ability to explain the most attractive property of language: *creativity*. However, before introducing compositionality in the next section, it is essential to step back and understand how linguistic creativity is defined (especially in the realm of CxG). Let us start with some creative expressions:

(4) a. *She smiled him in the door* (Goldberg, 2019, p. 61).
b. *The mother of all battles* (Hartmann & Ungerer, 2023, p. 5).
c. *Messi is the Mozart of football* (Hoffmann, 2019, p. 5).
d. *Weapons of mass distraction* (Giora et al., 2004).

The following expressions are likely to be unfamiliar to most readers and, therefore, by the previous definition can be considered creative utterances. However, according to Sampson (2016), there are two distinct types of creativity: F-creativity (fixed creativity), which produces examples drawn from a predetermined and established inventory, and E-creativity (extending/enlarging creativity), which goes beyond the system rules. According to this dichotomy, many linguistic phenomena traditionally assumed as "creative" are, in fact, examples of F-creativity, as new sentences are the result of grammatical rules (Hoffmann, 2018).

The term *productivity* is used in linguistics to refer to the "original use of established possibilities of the language" (Leech, 2014, p. 24). For instance, syntactic productivity concerns "the range of lexical items that may fill the slots of constructions" (Perek, 2016, p. 66). In accordance with CxG's assumptions, one uses and extends preexisting constructions to generate novel utterances. This is exemplified in (4a), where the use of a typical intransitive verb

(e.g., *smiled*) as transitive in a caused-motion construction forces a creative new meaning, such as "She caused him to move the door by smiling." Mismatches between the typical environments in which a verb is used and its occurrence in a new and creative way are widely discussed as *valency coercion* (Goldberg, 1995). Several studies in CxG have investigated this construction productivity, and in particular Goldberg (2019) offers an extensive review focusing on explaining "the partial productivity of grammatical constructions."

Even though many new expressions arise from productivity (of F-creativity), the question of "how do speakers use their grammar to create E-creative utterances" remains a topic of debate (Hoffmann, 2022a, p. 280). According to Bergs (2018), a source of E-creativity relies on the "intentional manipulation of linguistic structure" (p. 281), usually exemplified by linguistic *extravagance*, that is, to talk in such a way that you are noticed (Haspelmath, 1999; Ungerer & Hartmann, 2020). The use of formulaic patterns drawn from a fixed template, like in (4b) (namely, *snowclones*; cf. Hartmann & Ungerer, 2023) can be considered creative. They represent an interesting case because, even if they transmit a hyperbolic meaning fulfilling a specific pragmatic function, they still derive from a partially fixed construction. As such, these expressions illustrate the complex interplay between creativity and productivity (Ungerer & Hartmann, 2023).

Other examples of proper creative constructs are metaphorical expressions, like the one in (4c), which are governed by the general cognitive process of Conceptual Blending (Fauconnier & Turner, 2002; Turner, 2018). This mental operation constructs a partial match between two input mental spaces (FOOTBALL and CLASSIC MUSIC, in this example) and selectively projects from those inputs into a novel "blended" mental space, resulting in a new meaning (Messi is a genius on the football pitch, just as Mozart was a musical genius; cf. Hoffmann, 2019). However, even apparently, rule-breaking phenomena like the production of a novel metaphor rely on established patterns (i.e., entrenched construction X-is-the-Y-of-Z; cf. Fauconnier & Turner, 2002) and on established mechanisms. As Bergs and Kompa (2020, p. 14) observes: "Still, even the most creative metaphor has to use established means (analogy) and comply with most of the rules governing language use and linguistic interaction. Thus, metaphors are actually also examples of F-creativity in the widest sense; they do not expand the rules of language as such."

Other research domains propose alternative models for linguistic creativity. One such theory is the Optimal Innovation hypothesis, which posits that the aesthetics of creative productions are best explained by variations of familiar material (Giora et al., 2004). According to this theory, specific minimal modifications of familiar expressions can be more pleasurable than entirely novel

creations. For instance, the neologism in (4d) is optimally innovative because it induces a novel response while enabling the retrieval of a salient stimulus, the familiar expression "weapons of mass destruction." The question is: Are utterances of this type examples of F-creativity (as they relate to the familiar and use a specific mechanism) or of pure E-creativity?

Despite the considerable research on linguistic creativity, the examples above reveal that a consensus on what constitutes a creative expression has yet to be reached. As Maybin (2015, p. 34) stated: "While everyday language creativity is now an established area of ongoing linguistic research, there is a continuing lack of clear agreement about the precise definition and scope of creativity itself." Generally, the complex relation between productivity and creativity is far from being defined: Given that language is a complex system, it is challenging to define expressions entirely unconstrained by any rules ("All use of natural human language ultimately is F-creative"; cf. Bergs & Kompa, 2020, p. 18). In a broader sense, creativity can be viewed as a gradient phenomenon ranging from systematic productivity to extravagant stimuli that generalize from existing schemata. This Element focuses more on the F-creativity aspect of language, a "constrained" form of creativity (Goldberg, 2019), focusing on how the generation and comprehension of new expressions relate to the familiar and the mechanisms we can exploit to generate novel (but not necessarily creative) utterances apart from compositionality.

1.4 Roadmap

What is compositionality's role in today's models of language and sentence processing? How do we process both familiar and novel expressions? How can observations from experimental data be transposed into a formal theory of language representation and processing? This Element connects various linguistic theories and behavioral observations about processes governing semantic interpretation, arguing that CxG provides a more suitable linguistic formalism to explain language comprehension.

This Element is organized as follows. First, Section 2 introduces one of the two protagonists of the title: compositionality. Specifically, it discusses the notion of Fregean compositionality, traditionally believed to be the sole explanation for our ability to understand and create new sentences, and illustrates how this principle was used for describing the mechanism of meaning composition in traditional formalisms, constructionist approaches, and distributional models of meaning. Complementary, Section 3 examines studies in psycholinguistics and neurolinguistics that challenge this traditional view. The behavioral outcomes suggest a model of linguistic representation consistent