21

Construction Grammar and Artificial Intelligence

Katrien Beuls and Paul Van Eecke

21.1 A Common Attitude towards Communication and Language

To many contemporary linguists, Construction Grammar (CxG) and Artificial Intelligence (AI) might not spring to mind as two scientific disciplines that are closely related. Yet, both fields share a long-standing interest in modeling human communication and language and adopt a similar attitude towards this area of research. This similar attitude most visibly encompasses the following aspects:

- *Language serves a communicative purpose.* The basic function of language is to serve the communicative needs of its users, facilitating the transfer of information from one language user to another. As such, language production corresponds to the process of expressing an idea in the form of a natural language utterance, while language comprehension corresponds to the process of reconstructing the communicative intention underlying an observed utterance.
- *Communication is a bidirectional process.* Adequate representations and processing mechanisms for linguistic knowledge need to support the bidirectional nature of human communication and language. This entails that language comprehension and production are performed using the same representations and processing mechanisms. It is crucial for both humans and artificial agents that they can use the linguistic knowledge they have acquired through language comprehension in the production direction and that they themselves can understand any utterances they produce.
- *Languages are acquired rather than innate.* An individual language user acquires the language of their community by actively taking part in situated, communicative interactions. Languages cannot be innate, as this would compromise their ability to dynamically adapt to changes in the environment or in the communicative needs of their users. As language processing is heavily intertwined with other cognitive processes, in particular reasoning and vision, it is preferably modeled through the same general cognitive mechanisms.

- *Languages emerge and evolve through communication.* Each individual language user has built up their own linguistic system based on the communicative interactions they have participated in. This linguistic system is unique to each language user, as it has been shaped by the history of their successes and failures in communication. The evolutionary processes of variation and selection that take place in each individual during communication ensure that the linguistic system of each individual is compatible on a communicative level with the linguistic systems of all other individuals in the population.
- *Languages are grounded in (knowledge of) the world.* As the basic function of language is to serve the communicative needs of its users, it is necessarily grounded in the world in which they live. Understanding and producing natural language expressions heavily relies on world knowledge and common-sense reasoning. Indeed, the intended meaning underlying a natural language expression crucially depends on the concrete situation in which it was uttered. This situation includes a variety of aspects, including objects and actions observed in the world, pragmatic and discursive factors, and interpersonal relations.

It is clear that the research fields of Construction Grammar and Artificial Intelligence adopt a similar attitude towards the study of human language and communication. Especially in the period from the late 1960s to the early 1990s, this could be witnessed by collaborations and close interactions between leading figures in both fields. Today, traces of these interactions are still visible through a close reading of contemporary articles and textbooks. For example, Charles Fillmore, who is often considered the founding father of the field of Construction Grammar, explicitly acknowledges in his seminal paper on the case of *let alone* (Fillmore et al. 1988) the advice of UC Berkeley AI professor Robert Wilensky and his student Peter Norvig, who later went on to become a key figure in AI education worldwide (Russell & Norvig 2021). The advice was bidirectional, as Fillmore had served as a member of the PhD committee of Peter Norvig in 1978. Fillmore's case grammar (Fillmore 1968) had a substantial influence on later natural language understanding systems and was even presented as a standard component of natural language understanding in the first edition of Patrick Winston's standard textbook on Artificial Intelligence (Winston 1977; Jurafsky 2014). Starting in the mid 1970s, the notion of a 'frame' as a situational representation emerged through an interdisciplinary dialogue between sociologists (Goffman 1974), AI researchers (Minsky 1974; Schank & Abelson 1977), linguists (Fillmore 1976), and psychologists (Rumelhart 1980). Fillmore's linguistic work on Frame Semantics has thereby been highly influential in the field of Artificial Intelligence, most notably through the eventual development of the FrameNet project (Baker et al. 1998; Fillmore & Baker 2001).[1] While these examples are anecdotes rather than evidence, they do reflect the fact that the idea that researchers in CxG and

[1] We refer the interested reader to Chapters 1 and 3 of this volume for more background on Frame Semantics and FrameNets, respectively.

Artificial Intelligence are working towards a common goal, namely, to understand and model human language use, was strongly present in the early days of CxG.

When taking a closer look at recent contributions to journals and conferences in CxG and Artificial Intelligence, one gets the impression that interactions between both fields are much scarcer today than they used to be in the past. At the same time, the knowledge that both fields used to be aware of their common ground seems to have vanished to a large extent from both communities. We can only speculate about the reasons for this divergence, and there is probably not a single cause for this effect. Perhaps it is a symptom of a broader tendency of research fields to specialize and isolate. Or it may be a consequence of the reaction of many cognitively inspired linguists against the dogmas of generative grammar, by which they have sometimes overreacted and thereby developed an aversion towards any form of formalization. In practice, informal theories of CxG are less attractive to Artificial Intelligence researchers, as these researchers often lack the extensive CxG expertise that is needed to formalize them. It could also be due to the progressive institutionalization of Artificial Intelligence research groups within computer science departments, by which fewer and fewer linguists are hired and by which research in AI focuses increasingly on statistics and data science at the expense of models involving domain knowledge.

It is an explicit goal of this chapter to draw renewed attention to the common goals and similar attitude towards language and communication that have motivated mutually beneficial collaborations between construction grammarians and AI scholars in the past, and to emphasize the great value that lies in further elaboration of this relationship. On the one hand, we focus on the influence of ideas and techniques from the field of Artificial Intelligence on the field of CxG, thereby discussing the importance of these techniques for operationalizing the basic CxG tenets, for validating the consistency and precision of CxG theories, for corroborating these theories with corpus data, and for scaling constructionist approaches to language. On the other hand, we zoom in on the importance and use of CxG insights and analyses in the field of Artificial Intelligence, thereby emphasizing the excellent fit between the foundational ideas underlying constructionist approaches to language and the needs of researchers aiming to build truly intelligent systems. We are convinced that a thorough understanding of the relationship between both fields is highly beneficial for the contemporary construction grammarian, and that further developments in this direction will play a key role in shaping the future of the CxG field.

This chapter focuses explicitly on approaches within the fields of CxG and Artificial Intelligence that explicitly model constructional language processing (the relationship between CxG and transformer-based large language models is covered in Chapter 22).

21.2 Artificial Intelligence for Operationalizing Construction Grammar

This section discusses how methods and techniques from the field of Artificial Intelligence have contributed to the formalization and computational operationalization of constructionist approaches to language. It first revisits the basic tenets of CxG and then continues with a stepwise explanation of how these basic tenets can be mapped to data structures and algorithms that are known from the field of Artificial Intelligence.

21.2.1 The Basic Tenets of Construction Grammar

CxG refers to a family of linguistic theories that share a number of foundational principles. These principles, as laid out by, among others, Fillmore (1988), Goldberg (1995), Kay and Fillmore (1999), Croft (2001), Fried and Östman (2004), and Hilpert (2014), are the following:

- *All linguistic knowledge is captured in constructions.* All linguistic knowledge that is needed for language comprehension and production can be represented in the form of form–meaning mappings, called constructions. These constructions can freely combine to comprehend and produce utterances, as long as no conflicts occur (Goldberg 2006; Van Eecke & Beuls 2018).
- *There exists a lexicon–grammar continuum.* Construction grammars do not distinguish between the traditional notions of 'words' and 'grammar rules'. Constructions can range from fully instantiated form–meaning mappings, as in the case of idioms, to abstract schemata, as in the case of argument structure or information structure constructions. Many constructions are partially instantiated and partially abstract, as exemplified by the famous *let alone* construction (Fillmore et al. 1988).
- *Constructions span all levels of linguistic analysis.* Constructions can include information from all levels of traditional linguistic analysis. The form side of a construction typically contains a combination of phonetic, phonological, lexical, morphosyntactic, and multimodal information, while its meaning side typically combines semantic and pragmatic information. Constructions do not need to contain information on each of these levels. For example, they can, but do not need to, include word order constraints.
- *Construction grammars are dynamic systems.* Constructions are not innate but constructed during communicative interactions. Based on the frequency of their success and failure in communication, constructions can become more or less entrenched. As a consequence, CxG always represents the linguistic knowledge of an individual language user, as opposed to modeling an imaginary ideal language user.
- *Construction grammars should account for all linguistic phenomena.* CxGs do not adhere to the generative core–periphery distinction, and all linguistic phenomena are considered to be of equal interest. The same machinery is used to handle all linguistic phenomena, whether they are traditionally seen as regular, semi-regular, irregular, or idiomatic.

Formalization was considered to be an important aspect of CxG research since the inception of the field, with initial formalizations being inspired by phrase structure grammars (e.g., Fillmore 1988). However, the focus on formalization faded into the background when the Lakovian/Goldbergian branch of CxG, called Cognitive Construction Grammar, became predominant. The focus was on the conceptual clarification of the refreshing ideas that laid the foundations of the field, rather than on precise formalizations or computational implementations. However, once the initial ideas had settled, a relatively small number of construction grammarians started to focus on how these ideas could be formalized, verified, implemented, and tested (Kay & Fillmore 1999; Steels 2004; Bergen & Chang 2005; Feldman et al. 2009; Sag 2012; Michaelis 2013). Traditional techniques that were commonly used to formalize and implement generative grammars, such as the unification of feature structures, had to be complemented with innovative machinery that could accommodate those aspects of constructions that substantially differ from phrase structure rules. These include among others the fact that constructions can be non-local, that they do not necessarily correspond to tree-building operations (van Trijp 2016), that they can, but do not need to, include word order constraints, and that they are acquired through communicative interactions.

The innovative machinery that was needed to formalize and implement the basic tenets of CxG was borrowed from the field of Artificial Intelligence. In particular, heuristic search strategies (Wellens & De Beule 2010; Bleys et al. 2011) and innovative unification algorithms (Steels & De Beule 2006; Sierra Santibáñez 2012) were used to operationalize the free combination of constructions. Multi-agent simulations (Steels 2005; van Trijp 2008; Beuls & Steels 2013; Nevens et al. 2022) were used to model the dynamic nature of CxGs, including on the one hand the constructivist emergence, evolution, and acquisition of constructions, and on the other hand the entrenchment processes that take place in the constructicon.

21.2.2 Towards Computational Construction Grammar

The basic function of language is to support communication, that is, the transfer of information from one language user to another. There are always two parties involved in a communicative interaction, namely, a party that produces a linguistic expression and a party that comprehends it. Language production amounts to expressing an idea or intention in the form of a natural language utterance, while language comprehension consists in reconstructing the idea or intention underlying an observed utterance. Language processing can therefore be seen as a bidirectional process of mapping between intentions or ideas, referred to as meaning, and natural language utterances that express them, referred to as form. In computational terms, this means that we need to (i) represent natural language utterances, (ii) represent semantic structures, and (iii) provide a model that maps between these representations both in the comprehension and the production direction.

In essence, computationally operationalizing CxG, or any linguistic theory for that matter, involves finding precise representations and processing mechanisms that correspond to all aspects of the underlying theory. In computational terms, representations take the form of data structures, while processing mechanisms take the form of algorithms that operate over these data structures. When designing and implementing data structures and algorithms that operationalize computational CxG, the basic tenets of CxG as laid out in the previous section serve as a logical starting point.

In the next sections, we will illustrate how the basic tenets of CxG can be captured and operationalized in the form of a computational CxG system, highlighting the important role of techniques and methods from the field of Artificial Intelligence in this endeavor. We will adopt the terminology and conceptual framework underlying Fluid Construction Grammar (FCG – www.fcg-net.org; Steels 2004; van Trijp et al. 2022; Beuls & Van Eecke 2023). FCG is a computational CxG implementation that takes the form of a special-purpose programming language for designing and computationally implementing CxG models. FCG has the explicit goal of providing computational counterparts to the basic tenets of CxG in the form of a library of ready-to-use building blocks, while remaining an open framework that provides the flexibility and customizability to explore novel CxG ideas. For a technical introduction to FCG, we refer the interested reader to Van Eecke (2018: chapter 3).

Representing Utterances and Meanings

There are many different ways in which natural language utterances and semantic representations can be computationally represented. For didactic reasons, we will adopt the representations that are most commonly used in the computational CxG literature. Utterances are represented as a combination of tokens and adjacency constraints between those tokens. A token is a sequence of characters, that is, a string which corresponds to the part of an utterance that is enclosed by white space or punctuation. In order to be able to unambiguously refer to a token, each token is assigned a unique identifier. Adjacency constraints use these unique identifiers to express that two tokens are adjacent to each other. The tokens and adjacency constraints can be expressed as predicates and an entire utterance can consequently be represented as a set of predicates. An example of such a representation for the utterance *The more you think about it, the less it makes sense* (example from Hilpert 2021) is shown in (1):

(1)
```
{string(the-1, "The"), string(more-1, "more"), string(you-1, "you"),
 string(think-1, "think"), string(about-1, "about"), string(it-1, "it"),
 string(,-1, ","), string(the-2, "the"), string(less-1, "less"),
 string(it-2, "it"), string(makes-1, "makes"), string(sense-1, "sense"),
 string(.-1, "."), adjacent(the-1, more-1), adjacent(more-1, you-1),
 adjacent(you-1, think-1), adjacent(think-1, about-1), adjacent(about-1, it-1),
 adjacent(it-1, -1), adjacent(,-1, the-2), adjacent(the-2, less-1),
 adjacent(less-1, it-2), adjacent(it-2, makes-1), adjacent(makes-1, sense-1),
 adjacent(sense-1, .-1)}
```

This representation consists of thirteen 'string' predicates that represent the tokens in the utterance along with their unique identifiers, and twelve 'adjacent' predicates that encode the order of the tokens within the utterance.

We represent semantic structures using the Abstract Meaning Representation (AMR) formalism (Banarescu et al. 2013). AMR is a meaning representation language that was developed for representing the meaning of utterances in a way that (i) abstracts away from syntactic idiosyncrasies, (ii) is easy to read for humans, and (iii) is easy to manipulate by computers (Banarescu et al. 2013). An example of the AMR representation for the utterance *The more you think about it, the less it makes sense* introduced above is shown in (2):

(2)
```
(c / correlate-91
   :ARG1 (m / more
            :ARG3-OF (h / have-degree-91
                          :ARG1 (t / think-01
                                    :ARG0 (y / you)
                                    :ARG1 (i / it))))
   :ARG2 (l / less
            :ARG3-OF (h2 / have-degree-91
                          :ARG1 (s / sense-02
                                    :ARG1 i))))
```

On the highest level, we can observe that the utterance evokes a correlation, introduced by the 'correlate-91' roleset (Bonial et al. 2018). By definition, this roleset describes the correlation between two degrees to which two things hold. The first degree, notated as ':arg1', is a relation, in this case *more*. This relation correlates with the second degree, notated as ':arg2', in this case the relation *less*. Thus, an increase in something leads to a decrease in something else. The first relation, *more*, corresponds to the degree (':arg3-of' of the 'have-degree-91' roleset) to which a thinking event of roleset 'think-01' holds. The agent/thinker (':arg0') of the thinking event is *you* while the undergoer/thought (':arg1') of the thinking event is *it*. Thus, the more-relation embodies the degree to which you think about it. The second relation, *less*, corresponds to the degree (':arg3-of' of the 'have-degree-91' roleset) to which a sense-making event of roleset 'sense-02' holds. The 'thing that makes sense' (':arg1') is the same entity (i) as the undergoer/thought of the thinking event, namely *it*. The *less*-relation thus embodies the degree to which it makes sense, with *it* being the thing you are thinking about. In sum, the AMR representation of the utterance *The more you think about it, the less it makes sense* expresses that there is a correlation between the increasing degree to which you think about a particular referent and the decreasing degree to which that referent makes sense.

The AMR representation shown in (2) is expressed using the Penman notation, which was designed to be maximally human-readable. For computational purposes, we use a different notation, which represents AMR structures in the form of sets of predicates. As a consequence, meaning representations can be represented using the same data structure as linguistic utterances. The translation from Penman notation to sets of predicates is loss-less and automatic. The corresponding set of predicates for the example above is in (3):

(3) {`correlate-91(c), more(m), have-degree-91(h), think-01(t), you(y), it(i), less(l), have-degree-91(h2), sense-02(s), :arg1(c, m), arg2(c, l), :arg3-of(m, h), :arg1(h, t), :arg0(t, y), :arg1(t, i), :arg3-of(l, h2), :arg1(h2, s), :arg1(s, i)`}

Language Comprehension and Production

Now that we have established representations for utterances and semantic structures, we can define language comprehension and production as processes that map between these representations. In computational CxG, the linguistic knowledge that drives these processes is captured in the form of constructions. Intuitively, the task of these constructions is to move from a representation of an utterance to a representation of its meaning and vice versa.

In order to operationalize constructional language processing, computational CxG frames language processing as a search problem. Search problems form a class of problems that have been extensively studied in Artificial Intelligence and whose foundations date back to the seminal work of Newell and Simon (1956). Search problems are characterized by three main components: (i) a representation of the state of the search problem, (ii) operators that can work on a problem state and create new problem states that are hopefully closer to a solution, and (iii) a 'goal test' that determines whether a problem state corresponds to a solution or not. In the case of constructional language processing, these three components are instantiated as follows (Van Eecke & Beuls 2017):

(1) *Transient structures* serve as the representation of the state of the search problem. A transient structure holds all information that is known about an utterance being comprehended or produced at a given point during processing.
(2) *Constructions* serve as the operators of the search problem. Given a transient structure, they can contribute new information and thereby give rise to a new transient structure.
(3) *Goal tests* verify whether a given transient structure qualifies as a solution to the search problem.

The search process starts from a representation of the problem to be solved. In the case of constructional language processing, this representation takes the form of an initial transient structure. By definition, the initial transient structure holds all information that is known before processing starts. In the

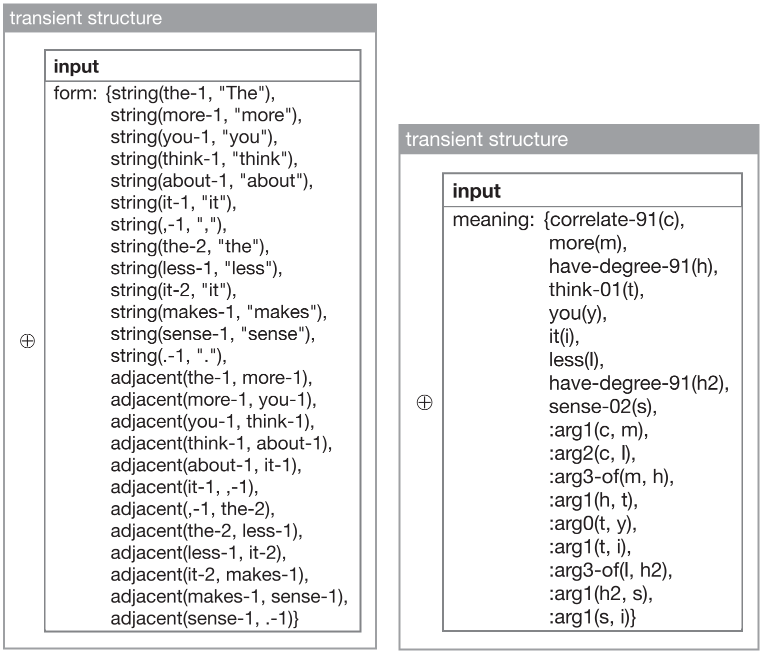

(a) Comprehension (b) Production

Figure 21.1 Initial transient structures in comprehension and production for the utterance *The more you think about it, the less it makes sense*

comprehension direction, the initial transient structure contains the form to be comprehended. In our example, this corresponds to the set of string and adjacency predicates introduced above. In the production direction, the initial transient structure contains a representation of the meaning to be expressed. In our example, this is the AMR representation shown above. The initial transient structures for comprehending and producing the example utterance, namely *The more you think about it, the less it makes sense*, are shown in Figure 21.1. The initial transient structures store their information in an 'input' unit under a 'form' and 'meaning' feature, respectively, denoting that it was the initial input to the problem-solving process.

Constructions capture linguistic information that can be used to advance the comprehension and production problem-solving processes. Given a transient structure, a construction can contribute new linguistic information and thereby create a new transient structure, which is hopefully closer to a solution. Constructions consist of two parts, a 'conditional pole' and a 'contributing pole'. The conditional pole contains the preconditions for a construction to apply and, thereby, create a new transient structure through its application. The contributing pole contains the postconditions of the construction, that is, information that will be added to the new transient structure during the application of the construction. As constructions support both the comprehension and production of utterances, they hold two sets of preconditions, one

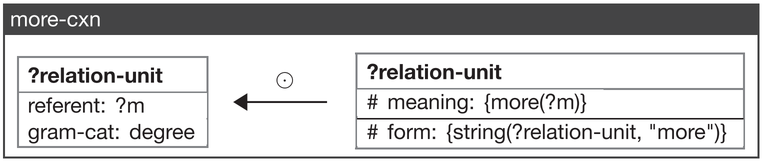

Figure 21.2 The *more-cxn*, which maps between the form *more* and the meaning predicate 'more(?m)', contributing the information that the result is of grammatical category 'degree' and that the referent of the unit is the argument of the 'more' predicate

for comprehension and the other for production. Preconditions in the comprehension direction serve as additional postconditions in the production direction and vice versa. During constructional language processing, constructions check whether their preconditions are compatible with a given transient structure in a given direction of processing, and if this is the case, they create a new transient structure that extends the current transient structure with the information contained in their contributing pole, combined with their preconditions of the other direction of processing.

An example of a construction is shown in Figure 21.2. The name of the construction, here *more-cxn*, is written in the dark box. The preconditions of the construction are written on the right-hand side of the horizontal arrow, while the postconditions are written on its left-hand side. The preconditions for comprehension and production are separated by a horizontal line, with the preconditions for production being written above the line and those for comprehension below it. On a conceptual level, the *more-cxn* maps between the form *more* and the AMR predicate 'more(?m)', contributing the information that the resulting unit is of grammatical category 'degree' and that its referent is the argument of the 'more' predicate. Technically, it does this through two features on its conditional pole and two features on its contributing pole. On its conditional pole, the construction contains a precondition for comprehension that a form predicate 'string(?relation-unit, "more")' should be part of the input, as well as a precondition for production that a meaning predicate 'more(?m)' should be part of the input.[2] If this is the case, the construction can apply and a new transient structure is created. This new transient structure starts as a copy of the current transient structure. In comprehension, the information is added that a meaning predicate 'more(?m)' is involved, along with the information that this unit is of grammatical category 'degree' and that its referent is the argument of the 'more' predicate. In production, the information is added that a string *more* is involved, along with the information that the unit is of grammatical category 'degree' and that its referent is the argument of the 'more' predicate. The result of the application of the *more-cxn* shown in Figure 21.2 on the transient structures shown in Figure 21.1 is provided in Figure 21.3.

[2] In fact, the symbol # that precedes these features explicitly indicates that they should be found in the 'input' unit rather than in any unit of the transient structure.

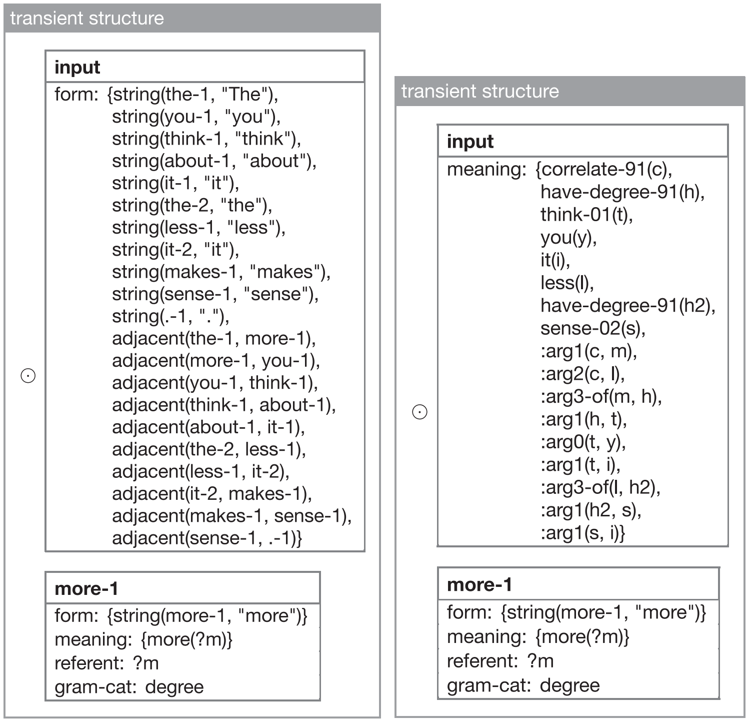

(a) Comprehension (b) Production

Figure 21.3 Transient structures in comprehension and production after applying the *more-cxn* from Figure 21.2 to the initial transient structures shown in Figure 21.1

Every time a new transient structure has been created as the result of a successful construction application, a number of goal tests are automatically run on this new transient structure to verify whether it qualifies as a solution state (Bleys et al. 2011). Typical goal tests for constructional language processing include (i) checking whether no more constructions can apply, (ii) verifying whether all 'string' or 'meaning' predicates have been processed, and (iii) checking whether the meaning comprehended so far consists of a fully connected network of predicates linked through their arguments. As soon as all goal tests succeed for a given transient structure, it is flagged as a solution state and the search process is halted. Depending on the direction of processing, all predicates under a 'meaning' (comprehension) or 'form' (production) feature are extracted from the solution transient structure. The result of the comprehension process is a set of meaning predicates, while the result of the production process is a set of string and adjacency predicates which can be automatically rendered as an utterance.

As it is typically the case that multiple constructions can apply to a given transient structure, the search space involved in the exploration of alternative construction applications quickly grows very large. To navigate the search

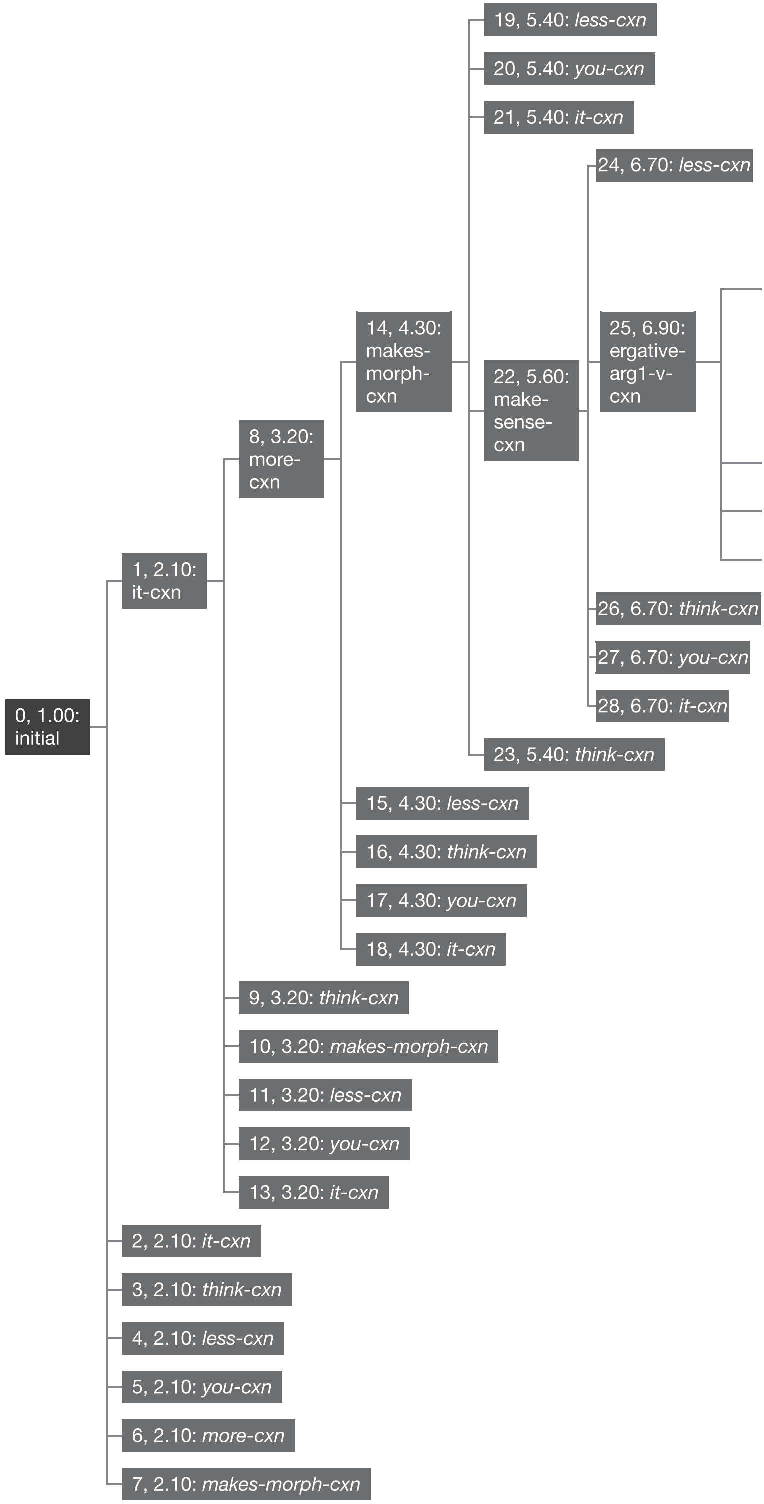

Figure 21.4 The search space involved in the comprehension of the utterance *The more you think about it, the less it makes sense*. A solution is found in node 41, after the application of eleven constructions.

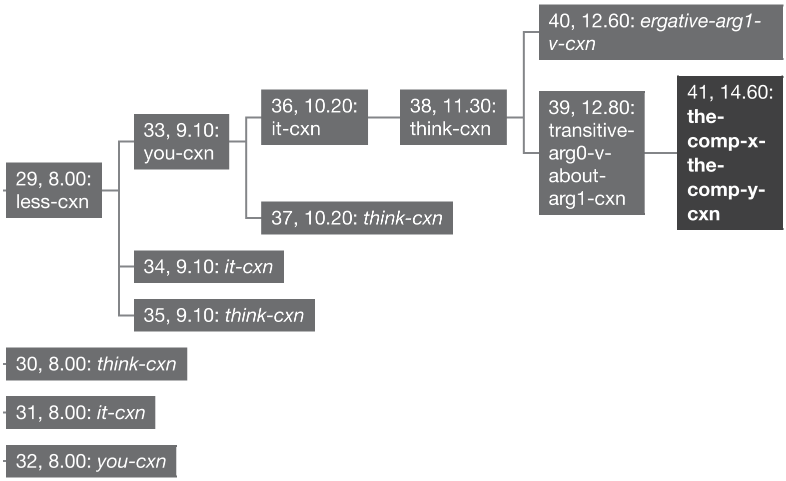

Figure 21.4 (cont.)

space in an informed way, researchers in AI have developed 'heuristic search' techniques. These techniques rank problem states according to their quality, estimating, for example, how close they are to a solution state (Pearl 1984; Russell & Norvig 2021). Common heuristics for steering the search space involved in constructional language processing include the number of constructions applied so far (favoring deeper solutions) and the number of units that were matched during construction application (favoring constructions that span larger patterns). More recently, it has been shown that neural sequence-to-sequence-based heuristics perform particularly well at ranking transient structures based on the sequence of constructions that have been applied in order to reach them (Van Eecke et al. 2022).

Figure 21.4 shows the search space involved in the comprehension process of the example utterance *The more you think about it, the less it makes sense*. The initial transient structure is represented by the leftmost box of the figure. The branching tree that is drawn to the right represents all construction applications that have taken place. The resulting transient structures are numbered according to when they were created (first number) and explored in the order obtained through their heuristic value (second number). In this example, all goal tests succeed for transient structure 41 shown with bold-faced title. This transient structure is the result of eleven construction applications that led from the initial transient structure to the solution transient structure. The last construction that applied was the *the-comp-x-the-comp-y-cxn*, a high-level construction that pairs the pattern ["the"–*degree-proposition*–","–"the"–*degree–proposition*] with its meaning representation that states that the extent to which the first degree holds for the first proposition is correlated with the extent to which the second degree holds for the second proposition. An implementation of this construction in FCG is shown in Figure 21.5. The construction includes adjacency constraints that capture the word order inherent to the pattern, as well as meaning predicates that integrate the referents of the different components of the pattern.

Acquisition, Evolution, and Entrenchment

As CxGs are dynamic systems that are 'constructed' during communicative interactions, the inventory of features and categories that is used in an individual grammar is open-ended. In FCG, this is reflected by the absence of an a priori specification of possible features and their values. The fact that new features and values can be dynamically added should the need arise is a necessary precondition for modeling the invention, adoption, and evolution of constructions in the context of both language acquisition and language emergence. In such experiments, inspiration is again drawn from the field of Artificial Intelligence, this time from research on learning in multi-agent systems. These experiments typically consist in a community of language users being modeled as a population of autonomous agents that participate in pairwise, goal-driven communicative interactions, referred to as 'language games' (Steels 1998, 2001).

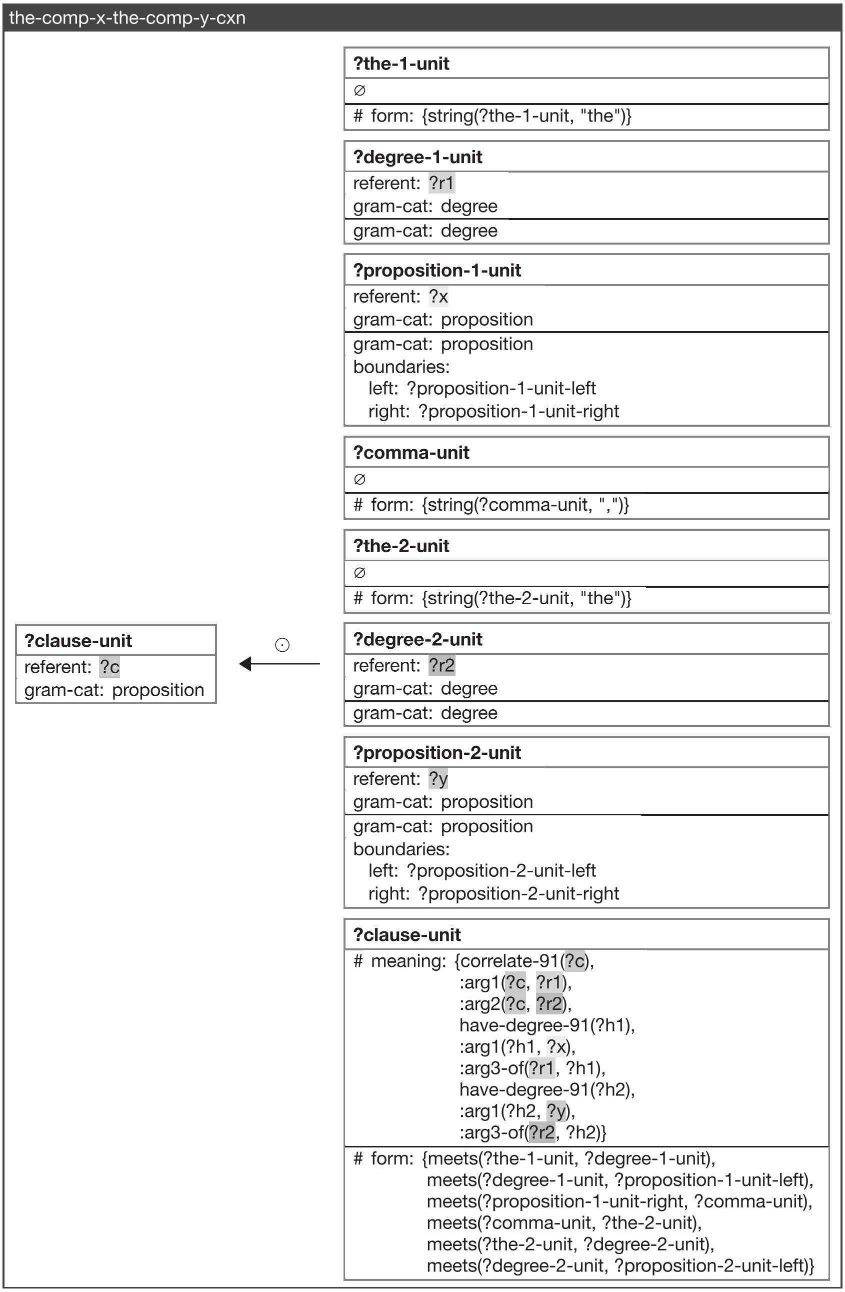

Figure 21.5 The *the-comp-X-the-comp-Y-cxn* pairs the pattern ["the"-*degree-proposition*-","-"the"-*degree-proposition*] with its meaning representation that states that the extent to which the first degree holds for the first proposition is correlated to the extent to which the second degree holds for the second proposition

Language games either adopt a tutor–learner scenario or a language emergence scenario. In a tutor–learner scenario, the goal is that one or more learner agents acquire the language of the community in a constructivist manner. In an emergence scenario, an entirely new language emerges that satisfies the communicative needs of its members. A typical language game in an emergence scenario proceeds as follows. At the beginning of each interaction, two agents are selected from the population and are assigned the role of either speaker or hearer. The agents are placed in a particular scene and need to successfully communicate to solve a given task, for example referring to objects or events that they observe in the scene. The agents are equipped with mechanisms for inventing and adopting linguistic means (i.e., constructions) that may be needed to achieve communicative success. After each interaction, the speaker provides feedback to the hearer about the outcome of the task. This allows the hearer to learn in the case that the agents did not reach communicative success. Additionally, both agents reward the constructions that were used in the case of a successful interaction and punish them in the case of a failed interaction. As more and more interactions take place, the agents in the population gradually converge on a shared language (De Vylder & Tuyls 2006). The language of each individual agent has been shaped by the communicative interactions it has participated in and is, therefore, well adapted to the task and the environment. As the scores of individual constructions reflect their frequency of successful application, they can be seen as a measure of their degree of entrenchment and are, as a consequence, often referred to by the term 'entrenchment scores'. When comprehending and producing linguistic utterances, the entrenchment scores are used to prioritize constructions where multiple constructions are in competition with each other. Notable applications of this language acquisition and emergence paradigm include experiments on the emergence and evolution of phonetic systems (de Boer 2000; Oudeyer 2006), vocabularies (Baronchelli et al. 2006; Steels 2015), domain-specific conceptual systems (Steels & Belpaeme 2005; Bleys & Steels 2009; Spranger 2016; Nevens et al. 2020), and grammatical structures (van Trijp & Steels 2012; Beuls & Steels 2013; Van Eecke 2018; Nevens et al. 2022; Doumen et al. 2023).

Typically, the grammatical categories that emerge during language game experiments are modeled in the form of a 'categorial network' with links between constructional slots and their (observed) fillers (Steels et al. 2022), very much in the spirit of Radical Construction Grammar (Croft 2001). Categories are thus construction-specific, emergent and ever-evolving as a result of language use. Figure 21.6 shows part of a categorial network that was acquired by an artificial agent in a question-answering game (Nevens et al. 2022). We see, for instance, that the 'ball' category is compatible with the 'how-big-is-the-?x(?x)' category, reflecting the fact that when a construction has contributed a 'ball' category to the transient structure, this category is compatible with the ?x slot of the *how-big-is-the-?x-cxn*. As a consequence, grammatical categories emerge as clusters within a graph that connects construction slots with their observed fillers. For example, as 'ball', 'cube', 'sphere', 'block', and 'cylinder' have often been observed as fillers of the

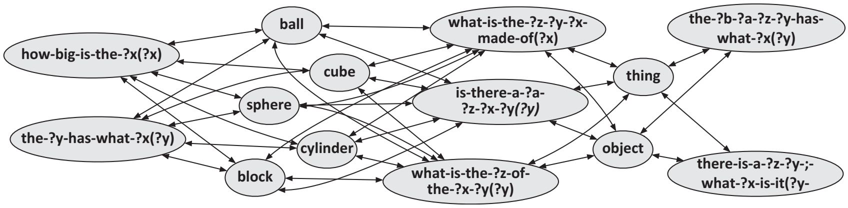

Figure 21.6 Snapshot of a small part of an agent's categorial network built up through a question-answering game

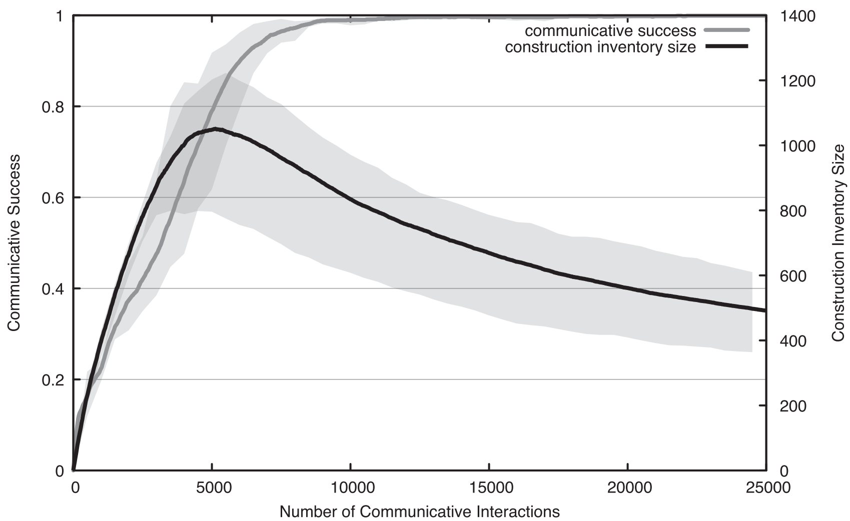

Figure 21.7 Typical learning dynamics of a language game experiment (graph plotted based on data from Nevens et al. 2022)

same construction slots, they are considered close in terms of grammatical category. A snapshot of the learning dynamics of the same language game experiment is captured in Figure 21.7, where the agent's communicative success and construction inventory size are plotted as a function of time. After an initial learning phase in which the number of constructions in the learner's grammar rises steeply, the size of the grammar starts to decrease steadily as a result of the entrenchment dynamics.

21.3 Construction Grammar for Operationalizing Artificial Intelligence

While the previous section discussed the influence of insights and techniques from the field of Artificial Intelligence on the field of CxG, the present section addresses the inverse direction of influence. We will focus in particular on how the foundational ideas underlying constructionist approaches to language form an excellent fit with the needs of AI researchers who aim to build truly intelligent agents that are capable of interacting through natural language. The section starts with an overview of the desirable properties of the linguistic capability of such agents. It then continues with two specific cases that illustrate the role and application of constructional research within the field of Artificial Intelligence.

21.3.1 Communicatively Capable Intelligent Agents

As introduced in Section 21.1, the fields of CxG and Artificial Intelligence adopt a similar attitude towards communication and language. Both fields acknowledge that language serves a communicative purpose and that language comprehension and production are equally important processes. Languages are acquired rather than innate, and they emerge and evolve as a result of verbal interactions between members of the linguistic community. Finally, languages are grounded in the world and are therefore strongly tied to the communicative needs of their users. A graphical representation of the processes involved in language processing viewed from this perspective is shown in Figure 21.8 in the form of the 'semiotic cycle'. The left-hand side of the semiotic cycle depicts the processes that involve the speaker and the right-hand side represents those that involve the hearer. The speaker and the hearer can both perceive the same world through their own sensors and act upon it through their own actuators. Through a conceptualization process, the speaker composes a conceptual structure based on his/her communicative intentions. In other words, the speaker decides what information s/he wishes to convey to the hearer and formalizes this information in the form of a semantic representation. The speaker then produces an utterance that expresses this semantic representation. The hearer observes this utterance and maps it to a semantic representation of his/her own. This semantic representation can be seen as a reconstruction of the conceptual structure underlying the speaker's utterance based on the hearer's knowledge of the world. The hearer then interprets this conceptual structure in relation to his/her view on the world and acts accordingly.

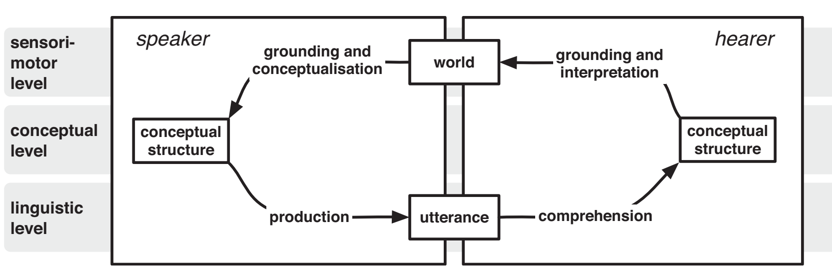

Figure 21.8 The semiotic cycle representing the processes involved in linguistic processing from a communicative perspective

Human languages are characterized by their remarkable robustness, flexibility, and adaptivity to changes in the environment and communicative needs of their users. These characteristics stem from the way in which these languages have emerged and continue to evolve. The language of a community corresponds in essence to a set of conventions on which its members have converged. This global convergence is a result of purely local interactions between community members, a phenomenon often referred to as self-organization. At the same time, such distributed systems are inherently robust against considerable perturbations, while maintaining the flexibility to adapt to environmental changes when they occur. CxG models, which start exactly from this idea and thereby incorporate the aforementioned desirable properties, are thus an important source of inspiration for the field of Artificial Intelligence, as the same properties can be seen as crucial properties of truly intelligent agents (Mikolov et al. 2016).

21.3.2 Application Case 1: Modeling Language Acquisition in Intelligent Agents

A first case that demonstrates the application of insights and analyses from CxG in the field of Artificial Intelligence concerns the constructionist acquisition of language by intelligent agents. Usage-based constructionist theories of language acquisition argue that the ability of children to learn language is based on two general cognitive capacities: intention reading and pattern finding (Tomasello 2003, 2009). Intention reading refers to the capacity of children to understand the communicative intentions of their interlocutors, while pattern finding refers to children's ability to recognize similarities and differences in sensorimotor experiences (Tomasello 2003: 3f.). In other words, intention reading allows a language learner to reconstruct the meaning of an utterance that they observe during a communicative interaction, while pattern finding provides mechanisms to learn constructions based on the combination of observed utterances and their reconstructed meanings. Computational models that implement these processes are of great interest to the field of Artificial Intelligence as the resulting grammars are learnable in a decentralized, data-efficient, and incremental manner.

This line of work has been embraced by a variety of researchers, pursuing goals that range from validating theories of language acquisition to finding practical solutions to problems faced by artificial agents. One class of models operationalizes pattern finding only, learning constructions from utterances paired with their meaning representation. These pairs are either provided in the form of an annotated corpus (Dominey & Boucher 2005; Chang 2008; Abend et al. 2017) or obtained through task-oriented communicative interactions in a tutor–learner scenario (see the section Acquisition, Evolution, and Entrenchment above; Beuls et al. 2010; Gerasymova & Spranger 2010; Spranger & Steels 2015). Another class of models, as introduced by Gaspers et al. (2011), is designed to learn form–meaning pairings under referential uncertainty. As such, the exact meaning representations of the input utterances are not provided to the learning algorithm, but constructions are learned based on the combination of input utterances and situational context snippets. A third class of models operationalizes both intention reading and pattern finding, whereby the results of the intention reading processes concern complex semantic structures and the pattern finding processes yield constructions that generalize over pairs of observed utterances and reconstructed meaning representations (Nevens et al. 2022; Doumen et al. 2023).

In general, computational models of intention reading and pattern finding operationalize task-based communicative interactions that follow the language game paradigm introduced in the section on Acquisition, Evolution, and Entrenchment above. They thereby implement the processes of grounding, conceptualization and interpretation, and language comprehension and production depicted in the semiotic cycle in Figure 21.8. Imagine that an agent needs to learn to answer questions about the world it visually observes. At the beginning of the experiment, the agent only knows how to perform a limited number of cognitive operations. These operations include, for example, segmenting an image, filtering a set according to a prototype, counting the number of items in a set, and querying properties of an object. The agent does not know any constructions or other linguistic entities such as grammatical categories or word boundaries at the start of the experiment. A tutor agent might ask the learner agent to name the color of the car that passes by, using the utterance *What is the color of the car?* At this point, the learner agent will signal that it does not understand the utterance and the tutor agent will provide the answer to the question as feedback (e.g., *yellow*). Based on this answer, the learner agent will make a hypothesis about the intended meaning of the observed utterance. In order to do this, it will compose a semantic network based on the primitive cognitive operations it knows, such that, upon evaluation, this network leads to the answer that was provided by the tutor, for example [segment image – filter car – query color]. If later the tutor asks *What is the color of the sheep?* and the learner hypothesizes after feedback that it means [segment image – filter sheep – query color], the learner can construct a generalized pattern that pairs *What is the color of the* ?X with the meaning representation [segment image – filter ?X – query color]. At the same time, the learner can learn two patterns which pair *sheep* and *car* with their respective meanings, as well as two links in its categorial network that express that *sheep* and *car* can fill the '?X' slot in the *what-is-the-color-of-the-?x-cxn*. A schematic representation of the intention reading and pattern finding processes involved in the processing of this example is shown in Figure 21.9.

The entrenchment dynamics of the game ensure that after many interactions, the linguistic model of the learner agent is compatible with the tutor's language use in their shared environment. The composition of semantic networks as hypotheses about the intended meaning of the other agent constitutes an operationalization of intention reading, while the syntactico-semantic generalization over pairs of utterances and semantic networks constitutes an operationalization of pattern finding.

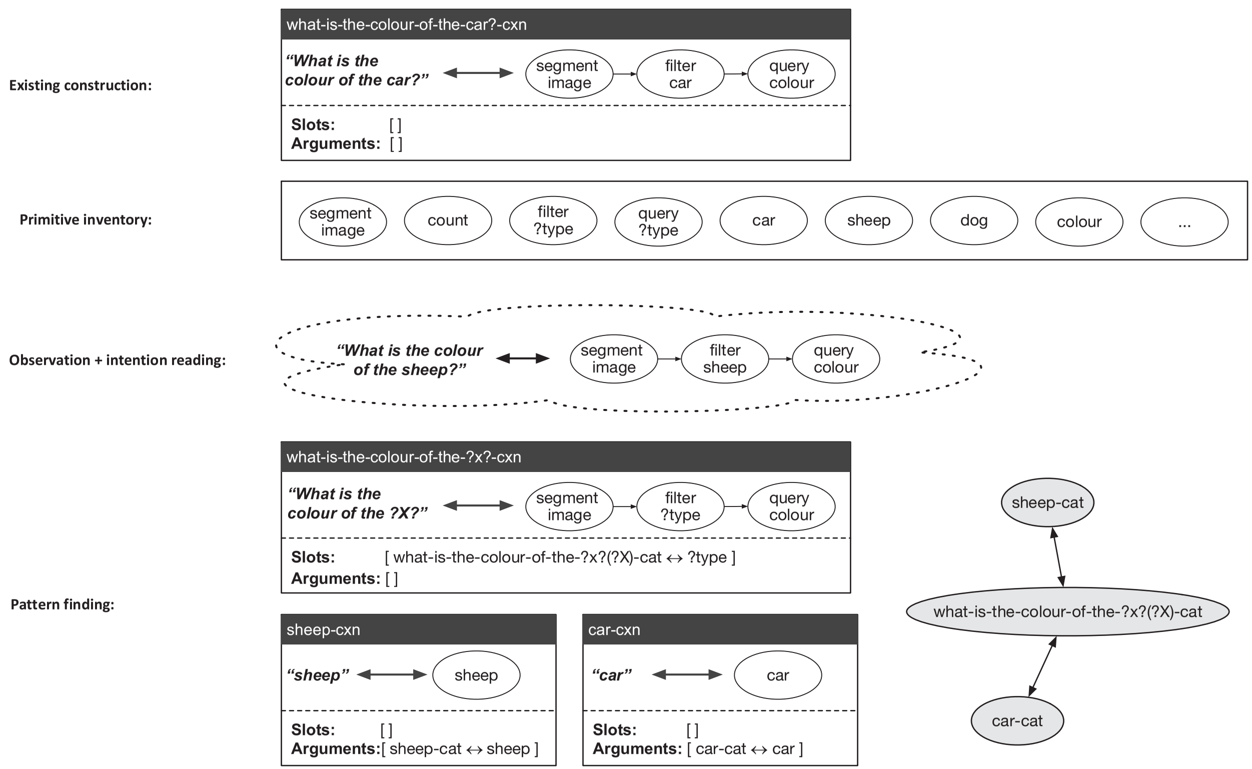

Figure 21.9 A schematic representation of the processes of intention reading and pattern finding taking place in an artificial agent. Three new constructions and two new categorial links were learned during a single interaction.

The constructionist acquisition of language through intention reading and pattern finding in task-based communicative interactions constitutes a paradigm that combines a number of features that are highly valued in the field of Artificial Intelligence. The paradigm assumes that agents are autonomous entities which sense, reason, and act independently. The global behavior that arises in the community stems from purely local interactions and is robust and adaptive as a consequence of its evolutionary nature. The paradigm focuses on the meaning and intentions underlying language as well as on their grounding in both the world and the knowledge of the agents. Semantic structures are composed by the agents themselves based on the environment, communicative feedback, and mental simulation. Learning is data-efficient and problem-driven, with one-shot learning of constructions being the norm. As the constructions that result from the learning process can generalize over the compositional aspects of the language (both in terms of form and meaning) and keep the non-compositional aspects within the constructions, the paradigm is compatible with any meaning representation. It is perhaps this insight from CxG, namely, that constructions can elegantly handle non-compositional forms and meanings, that has led to its appreciation in the AI community. The agents in the population do not even need to share the same primitive operations, morphology, or software architecture, making it possible to have communicatively adequate languages emerge in populations of heterogeneous agents. Finally, both the learning process and the resulting grammars are fully explainable and human-interpretable, which excellently fits the current focus on explainable and trustworthy AI.

21.3.3 Application Case 2: Modeling Opinion Dynamics for Understanding Society

The second case that showcases the application potential of CxG within the field of AI concerns the automatic analysis of opinions expressed on social media platforms. Today, such platforms play an important role in the formation of opinions and their propagation throughout society. The enormous amounts of data created every day make it impossible to grasp the dynamic landscape of opinions held by the members of a community. Automatic analysis tools therefore play an important role as research instruments for social scientists investigating this matter. Such tools need to be capable of analyzing social media posts and situate the opinions they express with respect to other opinions as well as real-world events. An important aspect of these tools is their ability to reason over the meaning of textual documents. Large-scale construction grammars can play an important role in the semantic analysis of these documents, as they are able to retrieve their underlying meaning through a transparent and interpretable model.

An illustrative example of an application that makes use of CxG for analyzing opinion dynamics is the Penelope opinion facilitator (Willaert et al. 2020, 2021). The opinion facilitator aims to help consumers of online news

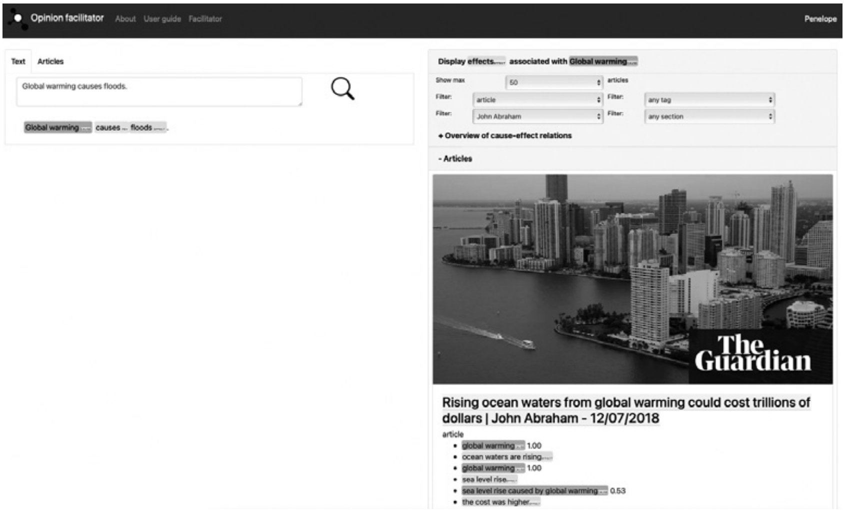

Figure 21.10 The Penelope opinion facilitator (Willaert et al. 2021)

media to investigate on the fly statements made in an article or newspaper comment by presenting other articles or comments that put forward concurring or diverging opinions. The news consumer is thereby offered a broad spectrum of opinions about a subject matter, reducing the risk of getting drawn into an echo chamber (see, e.g., Sunstein 2018).

An example of the use of the opinion facilitator is shown in Figure 21.10. On the left-hand side of the interface, the user can either enter a statement they wish to investigate or browse newspaper articles. A frame-semantic analysis of the statement or article is then visualized. In this visualization, the user can click on a participant role of a frame. Articles that contain the same frame with a semantically similar filler for the role that was clicked are then shown on the right-hand side of the interface. In these articles, the relevant frames are highlighted and a short summary is provided. In the example in Figure 21.10, the user has entered the statement *Global warming causes floods*. A causal frame was detected, with *global warming* filling the cause slot and *floods* filling the effect slot. The user has clicked on *global warming* and articles containing causal frames with global warming as cause are displayed on the right. The user has thereby found a broad spectrum of articles that mention the effects of global warming. The user can then build an informed opinion about the original statement based on the information conveyed through these articles. In this application, the frame-semantic analysis of the texts is performed by a computational CxG (Beuls et al. 2021). The main advantage of the use of such a grammar is that it is entirely transparent and human-interpretable. The detection of a frame and the assignment of participant roles is always the consequence of a construction application and is thereby linguistically motivated and explainable.

21.4 Discussion and Conclusion

Throughout this chapter, our aim has been to draw renewed attention to the common ground that is shared by the fields of CxG and Artificial Intelligence. We have done this on the one hand through a discussion of their historical ties and their common attitude towards communication and language, and on the other hand through an analysis of the way in which both fields have influenced and continue to influence each other.

When the field of CxG emerged in the 1980s, its common ground with the field of Artificial Intelligence was evident to its architects. Indeed, it was clear that both fields shared the objective of modeling human communication and language and that they held a similar attitude towards the nature of the subject matter. Both fields acknowledge that language serves as an instrument of communication between members of a community and that it has emerged and continues to evolve to serve its communicative purpose. As a logical consequence, both fields emphasize the importance of modeling language use, including the processes of language comprehension and production, rather than studying the competence of an ideal language user. Both fields acknowledge that languages are acquired rather than innate and that they emerge and evolve as a consequence of local communicative interactions between community members. The linguistic system of each community member is therefore unique as it has been shaped by their past successes and failures in communication. Finally, linguistic systems are grounded in the environment and world knowledge of community members and are adaptive to changes in the environment and communicative needs of their users. We are convinced that a renewed awareness of their shared objectives and attitude towards communication and language will benefit future research in both fields.

When it comes to the first direction of influence, that is, the influence of the field of Artificial Intelligence on the field of CxG, we have argued that ideas and techniques from AI have played a crucial role in the formalization and computational implementation of the basic tenets of CxG. A wide range of AI techniques has been deployed in this endeavor. Feature structures are used to formalize constructions and innovative unification algorithms have been developed to operationalize the processes of construction-based language comprehension and production. Constructional language processing is operationalized as problem solving through search, with heuristic search strategies making the free combination of constructions computationally tractable. Finally, multi-agent simulations are used to model the emergence, acquisition, and dynamic evolution of grounded constructicons within populations of language users. In sum, insights and techniques from the field of Artificial Intelligence have served as a cornerstone in the operationalization of computational CxG.

In regard to the second direction of influence, that is, the influence of the field of CxG on the field of Artificial Intelligence, we have highlighted the observation that the foundational ideas underlying CxG form an excellent fit with the Artificial Intelligence goal of building communicatively capable agents. First of all, the focus on the meaning of linguistic expressions, rather than on their form, and the grounding of this meaning in the world and knowledge of language users, supports the development of AI systems that can interact with their environment and each other through natural language. Second, the dynamic and usage-based nature of constructions, combined with the decentralized nature of their acquisition, facilitates the bootstrapping of communication systems that exhibit the robustness, flexibility, and adaptivity found in human languages. Finally, the inherent ability of constructions to generalize over the compositional aspects of language use and to capture any aspects of language use where the form and meaning are non-compositional with respect to each other is perhaps the most desirable property of CxG when it comes to building real-world AI systems. We have illustrated these aspects through two specific cases. One concerned the modeling of the acquisition of a constructicon that enables an autonomous agent to learn to answer questions about its environment. The other case presented an opinion facilitator tool in which frame-semantic analyses obtained through a human-interpretable computational CxG served as the basis for tracking opinions in online news media.

We strongly believe that a re-evaluation and further elaboration of the strong relationship between the research fields of CxG and Artificial Intelligence will play a key role in shaping the future of the CxG scholarship. Indeed, computational operationalizations of CxG bring important methodological advantages that carry the promise of leading to a number of substantial breakthroughs with respect to the state of the art. Most prominently, computational operationalizations are indispensable when it comes to scaling constructionist approaches to language. They facilitate the automatic validation of the precision and internal consistency of CxG theories and analyses, which is impossible to do by hand for grammars that consist of tens of thousands of constructions. Moreover, they allow us to corroborate constructionist analyses with large amounts of corpus data, unequivocally revealing what they can and cannot account for. The scalability advantages of computational CxG also support moving away from the study of individual constructions to the study of systemic relations between families of constructions, thereby directly contributing to theory formation. An additional benefit of computational operationalizations concerns the standardization of the way in which constructions are represented, thereby facilitating the exchange of ideas and results among researchers. Finally, computational operationalizations will play a crucial role in enhancing the application potential of CxG, both as a linguistic framework adopted in a variety of other scientific disciplines and as a central component of communicatively capable AI systems.

References

Abend, O., Kwiatkowski, T., Smith, N. J., Goldwater, S., & Steedman, M. (2017). Bootstrapping language acquisition. *Cognition*, 164, 116–143.

Baker, C. F., Fillmore, C. J., & Lowe, J. B. (1998). The Berkeley FrameNet project. In *Proceedings of the 17th International Conference on Computational Linguistics, Vol. 1*. Washington, DC: Association for Computational Linguistics, pp. 86–90.

Banarescu, L., Bonial, C., Cai, S., Georgescu, M., Griffitt, K., Hermjakob, U., Knight, K., Koehn, P., Palmer, M., & Schneider, N. (2013). Abstract Meaning Representation for Sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*. Washington, DC: Association for Computational Linguistics, pp. 178–186.

Baronchelli, A., Felici, M., Loreto, V., Caglioti, E., & Steels, L. (2006). Sharp transition towards shared vocabularies in multi-agent systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(6), P06014. https://doi.org/10.1088/1742-5468/2006/06/P06014.

Bergen, B. & Chang, N. (2005). Embodied Construction Grammar in simulation-based language understanding. In J.-O. Östman & M. Fried, eds., *Construction Grammars: Cognitive Grounding and Theoretical Extensions*. Amsterdam & Philadelphia: John Benjamins, pp. 147–190.

Beuls, K., Gerasymova, K., & van Trijp, R. (2010). Situated learning through the use of language games. In *Proceedings of the 19th Annual Machine Learning Conference of Belgium and The Netherlands (BeNeLearn)*, pp. 1–6.

Beuls, K. & Steels, L. (2013). Agent-based models of strategies for the emergence and evolution of grammatical agreement. *PLoS ONE*, 8(3), e58960. https://doi.org/10.1371/journal.pone.0058960.

Beuls, K. & Van Eecke, P. (2023). Fluid Construction Grammar: State of the art and future outlook. In *Proceedings of the First International Workshop on Construction Grammars and NLP (CxGs+NLP, GURT/SyntaxFest 2023)*. Washington, DC: Association for Computational Linguistics, pp. 41–50.

Beuls, K., Van Eecke, P., & Cangalovic, V. S. (2021). A computational construction grammar approach to semantic frame extraction. *Linguistics Vanguard*, 7(1), 20180015. https://doi.org/10.1515/lingvan-2018-0015.

Bleys, J., Stadler, K., & De Beule, J. (2011). Search in linguistic processing. In L. Steels, ed., *Design Patterns in Fluid Construction Grammar*. Amsterdam & Philadelphia: John Benjamins, pp. 149–179.

Bleys, J. & Steels, L. (2009). Linguistic selection of language strategies: A case study for color. In *Proceedings of the 10th European Conference on Artificial Life*, pp. 150–157.

Bonial, C., Badarau, B., Griffitt, K., Hermjakob, U., Knight, K., O'Gorman, T., Palmer, M., & Schneider, N. (2018). Abstract Meaning Representation of constructions: The more we include, the better the representation. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. Miyazaki: European Language Resources Association (ELRA), pp. 1677–1684.

Chang, N. (2008). *Constructing Grammar: A Computational Model of the Emergence of Early Constructions*. PhD thesis. University of California, Berkeley.

Croft, W. (2001). *Radical Construction Grammar: Syntactic Theory in Typological Perspective*. Oxford: Oxford University Press.

de Boer, B. (2000). Self-organization in vowel systems. *Journal of Phonetics*, 28(4), 441–465.

de Vylder, B. & Tuyls, K. (2006). How to reach linguistic consensus: A proof of convergence for the naming game. *Journal of Theoretical Biology*, 242(4), 818–831.

Dominey, P. F. & Boucher, J.-D. (2005). Learning to talk about events from narrated video in a construction grammar framework. *Artificial Intelligence*, 167(1), 31–61.

Doumen, J., Beuls, K., & Van Eecke, P. (2023). Modelling language acquisition through syntactico-semantic pattern finding. In *Findings of the Association for Computational Linguistics, EACL 2023*, pp. 1347–1357.

Feldman, J., Dodge, E., & Bryant, J. (2009). Embodied Construction Grammar. In B. Heine & H. Narrog, eds., *The Oxford Handbook of Linguistic Analysis*. Oxford: Oxford University Press, pp. 121–146.

Fillmore, C. J. (1968). The case for case. In E. W. Bach & R. T. Harms, eds., *Universals in Linguistic Theory*. New York: Holt, Rinehart and Winston, pp. 1–88.

Fillmore, C. J. (1976). Frame semantics and the nature of language. *Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech*, 280(1), 20–32.

Fillmore, C. J. (1988). The mechanisms of "Construction Grammar". *Annual Meeting of the Berkeley Linguistics Society* 14, 35–55.

Fillmore, C. J. & Baker, C. F. (2001). Frame semantics for text understanding. In *Proceedings of WordNet and Other Lexical Resources Workshop, NAACL*, 6.

Fillmore, C. J., Kay, P., & O'Connor, M. C. (1988). Regularity and idiomaticity in grammatical constructions: The case of *let alone*. *Language*, 64(3), 501–538.

Fried, M. & Östman, J.-O. (2004). Construction Grammar: A thumbnail sketch. In M. Fried & J.-O. Östman, eds., *Construction Grammar in a Cross-Language Perspective*. Amsterdam & Philadelphia: John Benjamins, pp. 1–86.

Gaspers, J., Cimiano, P., Griffiths, S. S., & Wrede, B. (2011). An unsupervised algorithm for the induction of constructions. *2011 IEEE International Conference on Development and Learning (ICDL)*, 2, 1–6.

Gerasymova, K. & Spranger, M. (2010). Acquisition of grammar in autonomous artificial systems. In M. Coelho, R. Studer & M. Wooldridge, eds., *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI-2010)*. Amsterdam: IOS Press, pp. 923–928.

Goffman, E. (1974). *Frame Analysis: An Essay on the Organization of Experience*. Cambridge, MA: Harvard University Press.

Goldberg, A. E. (1995). *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.

Goldberg, A. E. (2006). *Constructions at Work: The Nature of Generalization in Language*. Oxford: Oxford University Press.

Hilpert, M. (2014). *Construction Grammar and Its Application to English*. Edinburgh: Edinburgh University Press.

Hilpert, M. (2021). *Ten Lectures on Diachronic Construction Grammar*. Leiden: Brill.

Jurafsky, D. (2014). Charles J. Fillmore. *Computational Linguistics*, 40(3), 725–731.

Kay, P. & Fillmore, C. J. (1999). Grammatical constructions and linguistic generalizations: The *What's x Doing Y?* construction. *Language*, 75(1), 1–33.

Michaelis, L. A. (2013). Sign-Based Construction Grammar. In T. Hoffmann & G. Trousdale, eds., *The Oxford Handbook of Construction Grammar*. Oxford: Oxford University Press, pp. 133–152.

Mikolov, T., Joulin, A., & Baroni, M. (2016). A roadmap towards machine intelligence. In *Proceedings of the International Conference on Intelligent Text Processing and Computational Linguistics*, pp. 29–61.

Minsky, M. (1974). *A Framework for Representing Knowledge*. Cambridge, MA: MIT AI Laboratory.

Nevens, J., Doumen, J., Van Eecke, P., & Beuls, K. (2022). Language acquisition through intention reading and pattern finding. In *Proceedings of the 29th International Conference on Computational Linguistics*, pp. 15–25.

Nevens, J., Van Eecke, P., & Beuls, K. (2020). From continuous observations to symbolic concepts: A discrimination-based strategy for grounded concept learning. *Frontiers in Robotics and AI*, 7, 84.

Newell, A. & Simon, H. (1956). The logic theory machine – a complex information processing system. *IRE Transactions on Information Theory*, 2(3), 61–79.

Oudeyer, P.-Y. (2006). *Self-organization in the Evolution of Speech*. Oxford: Oxford University Press.

Pearl, J. (1984). *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Boston: Addison-Wesley Longman Publishing.

Rumelhart, D. E. (1980). Schemata: The building blocks of cognition. In R. Spiro, B. Bruce, & W. Brewer, eds., *Theoretical Issues in Reading Comprehension*. Hillsdale: Lawrence Erlbaum Associates, pp. 33–58.

Russell, S. & Norvig, P. (2021). *Artificial Intelligence: A Modern Approach*, 4th edition. Hoboken: Pearson.

Sag, I. A. (2012). Sign-Based Construction Grammar: An informal synopsis. In H. C. Boas & I. A. Sag, eds., *Sign-Based Construction Grammar*. Stanford: CSLI Publications, pp. 69–202.

Schank, R. C. & Abelson, R. P. (1977). *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Hillsdale: Lawrence Erlbaum Associates.

Sierra Santibáñez, J. (2012). A logic programming approach to parsing and production in Fluid Construction Grammar. In L. Steels, ed., *Computational Issues in Fluid Construction Grammar*. Berlin: Springer, pp. 239–255.

Spranger, M. (2016). *The Evolution of Grounded Spatial Language*. Berlin: Language Science Press.

Spranger, M. & Steels, L. (2015). Co-acquisition of syntax and semantics: An investigation in spatial language. In Q. Yang & M. Wooldridge, eds., *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*. Palo Alto: AAAI Press, pp. 1909–1915.

Steels, L. (1998). The origins of syntax in visually grounded robotic agents. *Artificial Intelligence*, 103(1-2), 133–156.

Steels, L. (2001). Language games for autonomous robots. *IEEE Intelligent Systems*, 16 (5), 16–22.

Steels, L. (2004). Constructivist development of grounded Construction Grammar. In W. Daelemans & M. Walker, eds., *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*. Barcelona: Association for Computational Linguistic Conference, pp. 9–19. https://doi.org/10.3115/1218955.1218957.

Steels, L. (2005). The emergence and evolution of linguistic structure: From lexical to grammatical communication systems. *Connection Science*, 17, 213–230.

Steels, L. (2015). *The Talking Heads Experiment: Origins of Words and Meanings*. Berlin: Language Science Press.

Steels, L. & Belpaeme, T. (2005). Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences*, 28(4), 469–488.

Steels, L. & De Beule, J. (2006). Unify and merge in Fluid Construction Grammar. In P. Vogt, Y. Sugita, E. Tuci, & C. Nehaniv, eds., *Symbol Grounding and Beyond*. Berlin & Heidelberg: Springer, pp. 197–223.

Steels, L., Van Eecke, P., & Beuls, K. (2022). Usage-based learning of grammatical categories. arXiv:2204.10201. https://doi.org/10.48550/arXiv.2204.10201.

Sunstein, C. R. (2018). *#Republic*. Princeton: Princeton University Press.

Tomasello, M. (2003). *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Cambridge, MA: Harvard University Press.

Tomasello, M. (2009). The usage-based theory of language acquisition. In E. L. Bavin, ed., *The Cambridge Handbook of Child Language*. Cambridge: Cambridge University Press, pp. 69–87.

Van Eecke, P. (2018). *Generalisation and Specialisation Operators for Computational Construction Grammar and Their Application in Evolutionary Linguistics Research*. PhD thesis. Vrije Universiteit Brussel.

Van Eecke, P. & Beuls, K. (2017). Meta-layer problem solving for computational Construction Grammar. In *The 2017 AAAI Spring Symposium Series*. Palo Alto: AAAI Press, pp. 258–265.

Van Eecke, P. & Beuls, K. (2018). Exploring the creative potential of computational Construction Grammar. *Zeitschrift für Anglistik und Amerikanistik*, 66(3), 341–355.

Van Eecke, P., Nevens, J., & Beuls, K. (2022). Neural heuristics for scaling constructional language processing. *Journal of Language Modelling*, 10(2), 287–314.

van Trijp, R. (2008). The emergence of semantic roles in Fluid Construction Grammar. In A. D. M. Smith, K. Smith, & R. Ferrer i Cancho, eds., *The Evolution Of Language: Proceedings of the 7th International Conference (EVOLANG7)*. Singapore: World Scientific, pp. 346–353.

van Trijp, R. (2016). Chopping down the syntax tree: What constructions can do instead. *Belgian Journal of Linguistics*, 30(1), 15–38.

van Trijp, R. & Steels, L. (2012). Multilevel alignment maintains language systematicity. *Advances in Complex Systems*, 15(3–4). https://doi.org/10.1142/S0219525912500397.

van Trijp, R., Beuls, K., & Van Eecke, P. (2022). The FCG editor: An innovative environment for engineering computational construction grammars. *PLoS ONE*, 17 (6). https://doi.org/10.1371/journal.pone.0269708.

Wellens, P. & De Beule, J. (2010). Priming through constructional dependencies: a case study in Fluid Construction Grammar. In A. D. M. Smith, M. Schouwstra, B. de Boer, & K. Smith, eds., *The Evolution of Language: Proceedings of the 8th International Conference (EVOLANG8)*. Singapore: World Scientific, pp. 344–351.

Willaert, T., Van Eecke, P., Beuls, K., & Steels, L. (2020). Building social media observatories for monitoring online opinion dynamics. *Social Media + Society*, 6 (2). https://doi.org/10.1177/205630511989.

Willaert, T., Van Eecke, P., Van Soest, J., & Beuls, K. (2021). An opinion facilitator for online news media. *Frontiers in Big Data*, 4, https://doi.org/10.3389/fdata.2021.695667.

Winston, P. H. (1977). *Artificial Intelligence*. Reading: Addison-Wesley.