

Exploring the Creative Potential of Computational Construction Grammar

Van Eecke, Paul; Beuls, Katrien

Published in:
Zeitschrift für Anglistik und Amerikanistik

DOI:
[10.1515/zaa-2018-0029](https://doi.org/10.1515/zaa-2018-0029)

Publication date:
2018

[Link to publication](#)

Citation for published version (APA):
Van Eecke, P., & Beuls, K. (2018). Exploring the Creative Potential of Computational Construction Grammar. *Zeitschrift für Anglistik und Amerikanistik*, 66(3), 341-355. <https://doi.org/10.1515/zaa-2018-0029>

Copyright

No part of this publication may be reproduced or transmitted in any form, without the prior written permission of the author(s) or other rights holders to whom publication rights have been transferred, unless permitted by a license attached to the publication (a Creative Commons license or other), or unless exceptions to copyright law apply.

Take down policy

If you believe that this document infringes your copyright or other rights, please contact openaccess@vub.be, with details of the nature of the infringement. We will investigate the claim and if justified, we will take the appropriate steps.

Exploring the creative potential of computational construction grammar

Paul Van Eecke and Katrien Beuls

*Artificial Intelligence Laboratory - Vrije Universiteit Brussel
Pleinlaan 2, 1050 Brussels, Belgium*

Computational construction grammar aims to provide concrete processing models that operationalise construction grammar accounts of the different aspects of language. This paper discusses the computational mechanisms that allow construction grammar models to exhibit, to a certain extent, the creativity and inventiveness that is observed in human language use. It addresses two main types of language-related creativity. The first type concerns the ‘free combination of constructions’, which gives rise to the open-endedness of language. The second type concerns the ‘appropriate violation of usual constraints’ that permits language users to go beyond what is possible when adhering to the usual constraints of the language, and be truly creative by relaxing these constraints and by introducing novel constructions. All mechanisms and examples discussed in this paper are fully operationalised and implemented in Fluid Construction Grammar.

1 Introduction

Computational construction grammar is a branch of linguistics that aims to operationalise the insights and analyses from construction grammar into concrete processing models. As these models need to run on computational platforms such as computers or robots, they are required to be very precise and explicit, both in terms of their representations and in terms of the processing mechanisms that are used. In this paper, we explore representations and mechanisms that allow computational construction grammar models to exhibit, to a certain extent, the creativity and inventiveness that is observed in human language use.

We address two quite different types of language-related creativity. The first type concerns the productivity of grammatical structures, or in other terms, the fact that a limited inventory of constructions (or words, rules and constraints in non-constructional approaches) can give rise to an open-ended set of sentences. While this is only creative

in a very specific sense of the word, the linguistic literature often refers to this type of creativity as ‘the creative potential of language’:

Constructional approaches share with mainstream generative grammar the goal of accounting for the creative potential of language (Chomsky 1957, 1965). That is, it is clear that language is not a set of sentences that can be fixed in advance. Allowing constructions to combine freely as long as there are no conflicts, allows for the infinitely creative potential of language. (Goldberg 2006, 22).

In this fragment, Goldberg explains that construction grammar deals with this type of creativity by allowing the constructions of the grammar to combine freely, as long as their combination causes no conflicts. A concrete example of how this free combination of constructions can be implemented in computational construction grammar is presented in section 3 of this paper, by means of the resultative ‘*Firefighters cut the man free*’ example, taken from Hoffmann (2018).

The second type of creativity that we address does not only require a free combination of the constructions of the grammar, but also requires either the introduction of novel constructions or the relaxation of certain constraints in existing constructions. An operationalisation of this type of creativity is exemplified in section 4 of this paper, by showing through which mechanisms the constraints on the idiomatic expression ‘*He’s not the sharpest tool in the box*’ can be violated, in order to allow novel, creative expressions, such as ‘*He’s not the brightest light in the harbour*’, ‘*He’s not the smartest suit in the wardrobe*’ or ‘*He’s not the quickest bunny in the forest*’.

All examples discussed in this paper are fully operationalised and implemented in Fluid Construction Grammar (FCG) (Steels 2011, 2017), a flexible and largely theory-neutral computational framework that provides the basic building blocks for implementing construction grammars. This paper is accompanied by an interactive web demonstration (<https://www.fcg-net.org/demos/creativity>), in which the examples can be explored in full detail.

2 Creativity in computational construction grammar

There exist multiple computational construction grammar frameworks, of which Embodied Construction Grammar (ECG) (Bergen and Chang 2005; Feldman et al. 2009) and Fluid Construction Grammar (FCG) (Steels 2011, 2017) are the most advanced projects¹. While there has been previous research on the modelling of metaphors in ECG (see e.g. Stickles et al. 2016), most research into creativity and innovation in language has happened in FCG. In fact, modelling creative and inventive language use has

¹Although Sign-Based Construction Grammar (SBCG) (Boas and Sag 2012) is an advanced project that certainly shares certain characteristics with ECG and FCG, it is normally not considered a computational construction grammar framework. This is mainly due to its focus on the mathematical formalisation of language rather than on the computational processing of language use.

been of central importance in FCG since its start. As the framework springs from research on simulating the emergence and evolution of language in multi-agent systems, it was always essential that individual agents in the simulation could add new lexical and grammatical constructions to their constructicon. Many examples of computational creativity and inventiveness in language emergence and evolution can be found in the collection of papers published in Steels (2012). For further reading about the link between emergence and creativity, see Wiggins et al. (2015).

The central focus of FCG on creativity is also clearly reflected in its meta-layer architecture (Steels and van Trijp 2011; Beuls et al. 2012; van Trijp 2012; Van Eecke and Beuls 2017). While the routine layer performs the standard activation of the constructions of the grammar, diagnostics constantly check whether everything goes well. When a diagnostic signals a problem, which means that the standard activation of constructions is not sufficient for processing the input utterance or meaning representation, FCG jumps to its meta-layer. At the meta-layer, the problem is repaired, often by introducing new constructions or by relaxing certain constraints in an existing construction, after which routine processing resumes. In multi-agent experiments on the emergence and evolution of language, creative behaviour is almost always driven by a need to express meanings that cannot be expressed using the constructicon of the speaker, or, in the case of the hearer, by a need to understand an utterance that cannot be comprehended using its constructicon. However, creative language can in FCG be triggered by any factor, for example a desire to be extravagant (Keller 1994; Haspelmath 1999).

The meta-layer architecture of FCG deals differently with the two types of creativity that are considered in this paper. The first type, namely the free combination of constructions, is handled completely by the routine layer. No additional mechanisms are required, as a free combination of constructions is at the heart of the FCG engine. Freely combining constructions as long as there are no conflicts is exactly how FCG processes its input utterances or meaning representations. The second type of creativity, in which new constructions need to be invented or in which certain constraints in existing constructions need to be relaxed are handled by the meta-layer. This kind of creativity closely resembles the kind that is common in evolutionary experiments and the same computational mechanisms, such as anti-unification, can be used.

Apart from the specific mechanisms involved, research into the computational simulation of language evolution has revealed three characteristics that are important when modelling creativity. First of all, creativity is a process that involves both the speaker and the hearer. The innovative utterances of the speaker will only be considered creative if the language processing system of the hearer is flexible enough to comprehend them, otherwise they will just be considered nonsensical. Second, it is not problematic when creative language use violates the usual constraints of a grammar, for example by relaxing constraints in existing constructions or by introducing novel constructions. These violations, and the mechanisms that cause and handle them, are necessary for a language to emerge and evolve. Lastly, and this is related to the first characteristic, the extent to which creative language use violates the usual constraints of the language needs to be limited, such that it is still understandable for the hearer.

3 The free combination of constructions

The first type of creativity that we consider, concerns the “infinitely creative potential of language” (Goldberg 2006, 22), which construction grammar accounts for by “allowing constructions to combine freely as long as there are no conflicts” (Goldberg 2006, 22). As we have mentioned earlier, the free combination of non-conflicting constructions is at the heart of how FCG processes language. Let us have a look at the concrete representations and processing mechanisms that are involved. We will use the analysis of the utterance ‘*Firefighters cut the man free*’, which was introduced as an example of creative language use in the introduction to this volume (Hoffmann 2018), and show how it can be computationally implemented in FCG. This sentence exhibits two interesting phenomena: (i) the word *firefighters* fills both the slot of the argument role ‘*Agent*₁’ of the resultative construction and the ‘*cutter*₁’ role of the verb-specific *cut*-construction, and (ii) the *cut.object* is left unexpressed “and will be contextually identifiable” (Hoffmann 2018, 5). In order to be able to see the entire FCG grammar and fully appreciate how the constructions become activated and combine together freely, we encourage the reader to visit the web demonstration that accompanies this paper. The interactive visualisations that are offered there are likely to be both clearer and more precise than those that can be printed on paper or explained in text.

In order to produce or comprehend the example utterance, seven constructions need to be activated and combined in the working memory: the lexical *firefighter-cxn*, the morphological *plural-N-cxn*, the lexical *cut-cxn*, the lexical *man-sg-cxn*, the definite-*np-cxn*, the lexical *free-cxn* and the *resultative-cxn*. Figure 1 shows the lexical construction for the verb *cut*, which has two participants: a ‘*cutter*’ and a ‘*cut.object*’. The construction has the typical lay-out of an FCG construction, with on the left-hand side the features that are contributed by the construction, and on the right-hand side the conditional features, i.e. the constraints that need to be fulfilled for the construction to become active. The conditional features are divided horizontally into two locks, the upper lock containing the constraints that need to be fulfilled for the construction to apply in production, and the lower lock containing the constraints that need to be fulfilled for the construction to apply in comprehension. In the *cut-cxn* in the figure, the conditional features concern the meaning of the verb (activating the construction in formulation) and its orthographic form (activating the construction in comprehension)². The contributing features include a number of syntactic categories such as *lex-class* (lexical class) and agreement information, syntactic valence (underspecified here since this will be determined by the argument structure constructions), and semantic valence (linking the agent to the cutter and the undergoer to the *cut.object*). Every construction also receives a score (in this case the default score of 0.50), which indicates the construction’s degree of entrenchment, and which will influence which construction will be applied in case of competition.

The constructicon forms a continuum of constructions that range from very concrete,

²The choice of features in FCG is completely up to the grammar designer. If, for example, the phonetic form is preferred over the orthographic form, it only requires a trivial change in the construction.

³For an elaborate description of the syntax and semantics of FCG, see <https://github.com/EvolutionaryLinguisticsAssociation/Babel2/wiki/Syntax-and-Semantics-of-FCG>.

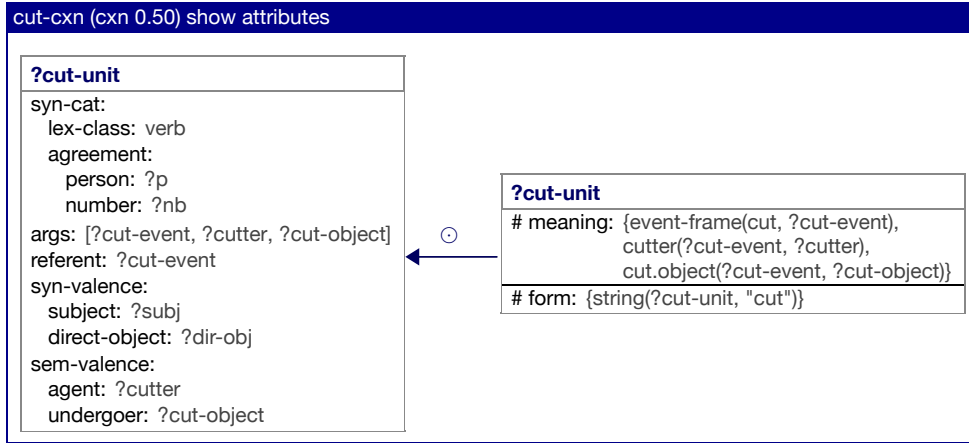


Figure 1: The lexical construction for the verb *cut*, with two event participant roles: a ‘cutter’ and a ‘cut.object’. The construction’s constraints are visualized on the right-hand side, with the production constraints on top and the comprehension constraints at the bottom. The features that are contributed by the construction are shown on the left-hand side. These can in turn be constraints for other constructions such as the resultative construction that can only become active if there is a verb. Symbols preceded by a question mark are variables.³

lexical constructions towards more abstract, grammatical constructions. For space reasons, we will skip over the other lexical constructions, as they look very much like the *cut-cxn*. All constructions can of course be inspected in detail in the web demonstration.

The definite NP construction, as described by Hoffmann (2018), matches on an instance of the determiner “the” followed by a noun. This construction is shown in Figure 2.

After the activation of the lexical constructions and the definite NP construction, the argument structure construction for the resultative becomes active. This construction maps between a subject NP (*firefighters*), a verb (*cut*), an object NP (*the man*) and an oblique unit (*free*), and its meaning “Agent causes Patient to become State by V-ing” (Hoffmann 2018). An FCG version of this construction is shown in Figure 3.

The FCG engine will try to freely combine these constructions. Constructions can be activated in any order, as soon as the constraints in their conditional part are satisfied. For each analysis, FCG automatically generates a visualisation that shows how the different constructions of the grammar have been combined to comprehend or formulate an utterance. This visualization for the utterance ‘*Firefighters cut the man free*’ is presented in Figure 4. It clearly shows that the *resultative-cxn* is conditioned on the activation of the *firefighter-cxn* as well as the *plural-n-cxn* for its subject unit. Its oblique unit is filled by the unit introduced by the *free-cxn* and the verb unit itself recruits material delivered by the *cut-cxn*. Finally, the object unit is linked to the NP that was created by the *definite-np-cxn*, which in turn depends on the *man-sg-cxn*.

The meaning that the grammar attributes to the utterance ‘*Firefighters cut the man*

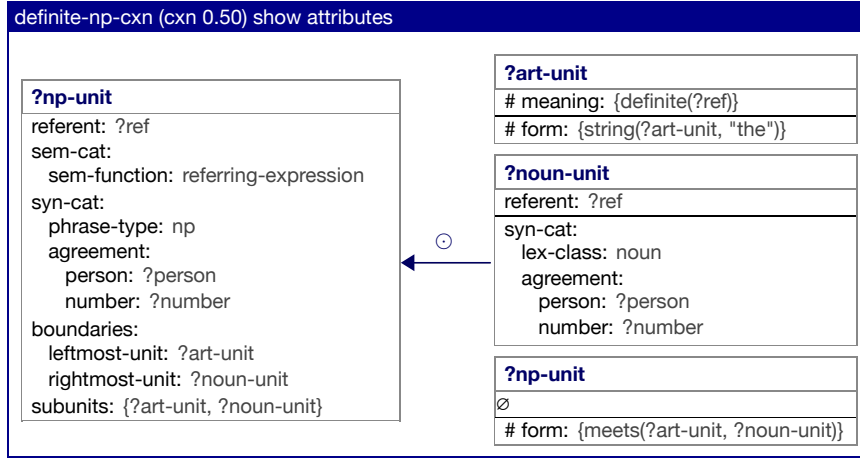


Figure 2: The construction for the definite NP combines “the” and an adjacent noun. In the conditional part (right-hand side), the meets predicate in the form feature of the np-unit imposes that the article is left-adjacent to the noun.

free’ is shown in Figure 5. Like in Hoffmann (2018)’s analysis, we can see that *firefighters* is both the cutter of the verb-specific cutting event and the agent of the resultative construction, and that the cut.object is left underspecified, as shown by the free variable ?cut-object-72.

4 The appropriate violation of usual constraints

In the previous section, creative language use consisted in combining the existing building blocks of language, i.e. the constructions in the constructicon, in novel ways. This can however only account for a very small part of the creativity that is observed in human language use. There, apart from combining existing building blocks in novel ways, new building blocks can be invented, often by altering one or more constraints in an existing building block. In computational construction grammar, this type of creativity cannot be handled by routine processing, as additional mechanisms are required to alter existing constructions, and to create new ones. This type of creativity is crucial in language, as it allows the language to emerge in the first place, but also to evolve and adapt to the environment.

An entire part of the architecture of the FCG framework, namely the *meta-layer*, is specifically designed to process language use that is not covered by the existing constructions of the grammar. When the FCG system detects at the routine layer that an input utterance or meaning representation cannot be handled in a satisfactory way by the existing constructions of the grammar, it jumps to the meta-layer. At the meta-layer, different strategies can be used to invent novel constructions (e.g. for novel words) or to alter existing constructions in such a way that they can become activated (e.g. by relaxing the constraints that blocked their activation). As the intention of a speaker

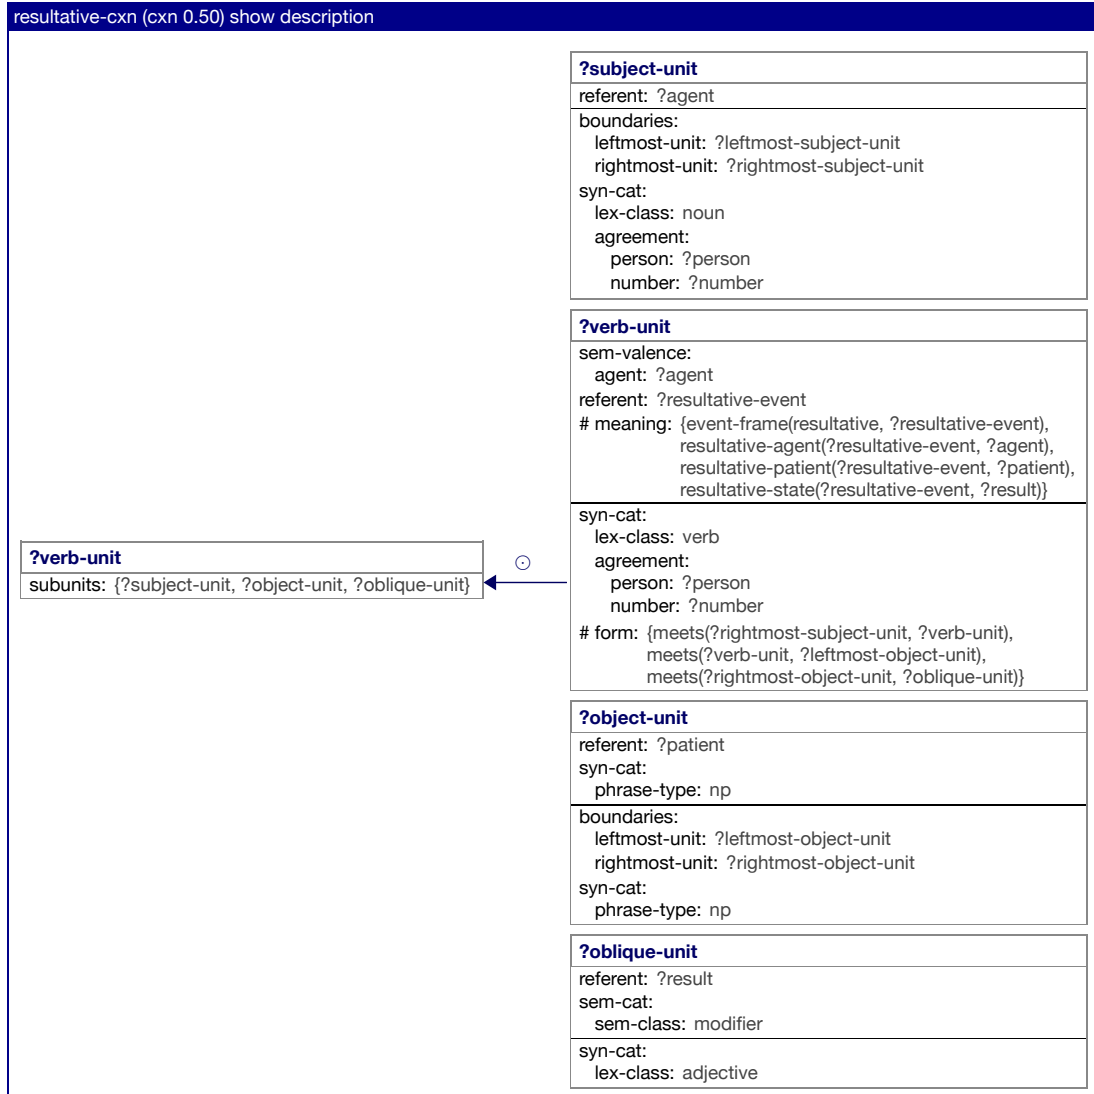


Figure 3: The resultative construction maps between a 'subject NP₁', a 'Verb₄', an 'object NP₂' and an 'oblique unit₃', and the meaning "Agent₁ causes Patient₂ to become State₃ by V₄-ing"

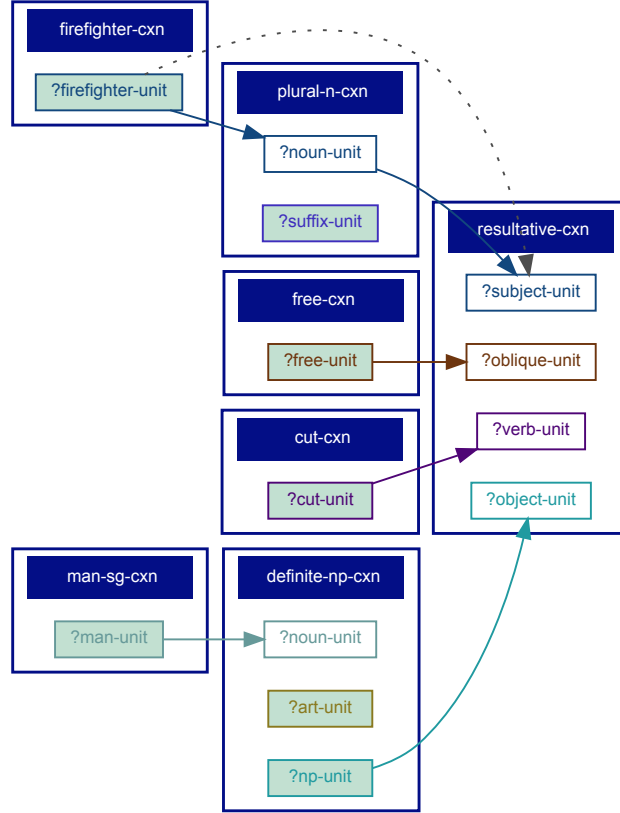


Figure 4: The free combination of seven constructions results in the analysis (comprehension or production) of the sentence ‘*Firefighters cut the man free*’. The arrows indicate the dependencies between the constructions. Dependencies that span over multiple constructions (e.g. the **resultative-cxn**, which depends on features added by the **firefighter-cxn** and **plural-n-cxn**) are visualised using dotted arrows.

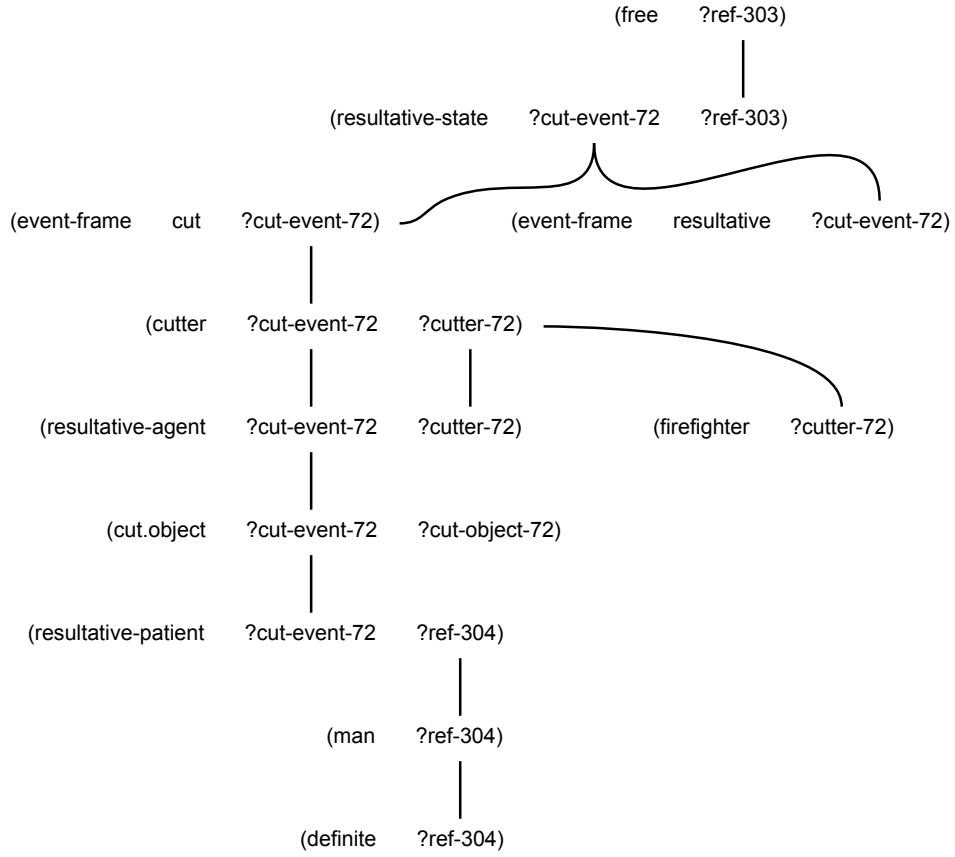


Figure 5: The meaning network that results from the comprehension process of the utterance ‘*Firefighters cut the man free*’. The network shows that the utterance evokes a **cut** frame represented by the variable **?cut-event-72**. The object slot of the frame is left open (**?cut-object-72**). The firefighters, the cutter of the event and the resultative agent all share the same referent (**?cutter-72**).

is almost always to be understood by the hearer, the new constructions that he introduces will rarely be completely novel and invented from scratch. New constructions will most often be variations on existing constructions, with one or more constraints altered. FCG’s meta-layer contains general mechanisms that allow this. *Anti-unification* is used for relaxing those constraints that block the activation of a construction, thus building a more general construction, and *pro-unification* is used to incorporate new, more specific constraints into an existing construction (Van Eecke and Beuls 2017).

We will now show a concrete example of creative language use that cannot be handled by the existing constructions of the grammar, and requires to create a new construction relaxing certain constraints of an existing one. The example that we use is the idiomatic expression ‘*he’s not the sharpest tool in the box*’, meaning that the subject of the utterance is not intelligent. Our FCG grammar contains an idiomatic construction that can handle this expression, namely the **not-the-sharpest-tool-in-the-box-cxn**, shown in Figure 6. This construction maps between the words *not*, *the*, *sharpest*, *tool*, *in*, *the* and *box*, and its meaning, namely that some person is not intelligent, that it is a metaphor, and that the semantic field of the metaphor has to do with hardware. The construction also contains a few additional constraints, namely that *tool* is an object, *box* is a container, that *sharp* is a positive property, and that all three words belong to the semantic field of hardware. Combined with the lexical constructions for the individual words, the **he-pronoun-cxn** and the **copula-cxn**, FCG can map between the utterance ‘*he’s not the sharpest tool in the box*’ and its meaning representation {**male-person(?x)**, **not-intelligent(?x)**, **metaphor(?metaphor, ?x)**, **semantic-field(?metaphor, hardware)**}. For more details about the specifics of these constructions, we refer the reader to the web demonstration.

One example of creativity would be to transpose the metaphor used in the idiomatic construction to a different semantic field than the field of hardware. This would require the grammar to generalise the **not-the-sharpest-tool-in-the-box-cxn** by relaxing the constraints on the semantic field of the metaphor and on the specific words that are used, while leaving intact all the other constraints in the construction. This generalised construction would then be able to handle utterances such as ‘*he’s not the brightest bulb in the box*’ or ‘*he’s not the quickest bunny in the forest*’. When the listener observes one of these utterances, FCG will detect that it cannot be processed using the existing constructions of the grammar. This triggers a jump to the meta-layer, where it will use anti-unification to find the existing grammatical construction that can process the utterance with the lowest number of generalisations needed. This process will generalise the **not-the-sharpest-tool-in-the-box-cxn** into a **not-the-x-est-y-in-the-z-cxn**, in which the references to the specific semantic field and the specific words *tool*, *sharp* and *box* have disappeared. All other constraints remain in the construction, including the constraints that *x*, *y* and *z* need to be of the same semantic field, that *x* needs to be an positive property, that *y* needs to be an object and that *z* needs to be some kind of container. The meaning of this generalised construction is the same as the meaning of the more specific construction, except for the semantic field of the metaphor. The generalised construction is shown in Figure 7.

Constructions obviously work in two directions, namely in comprehension and pro-

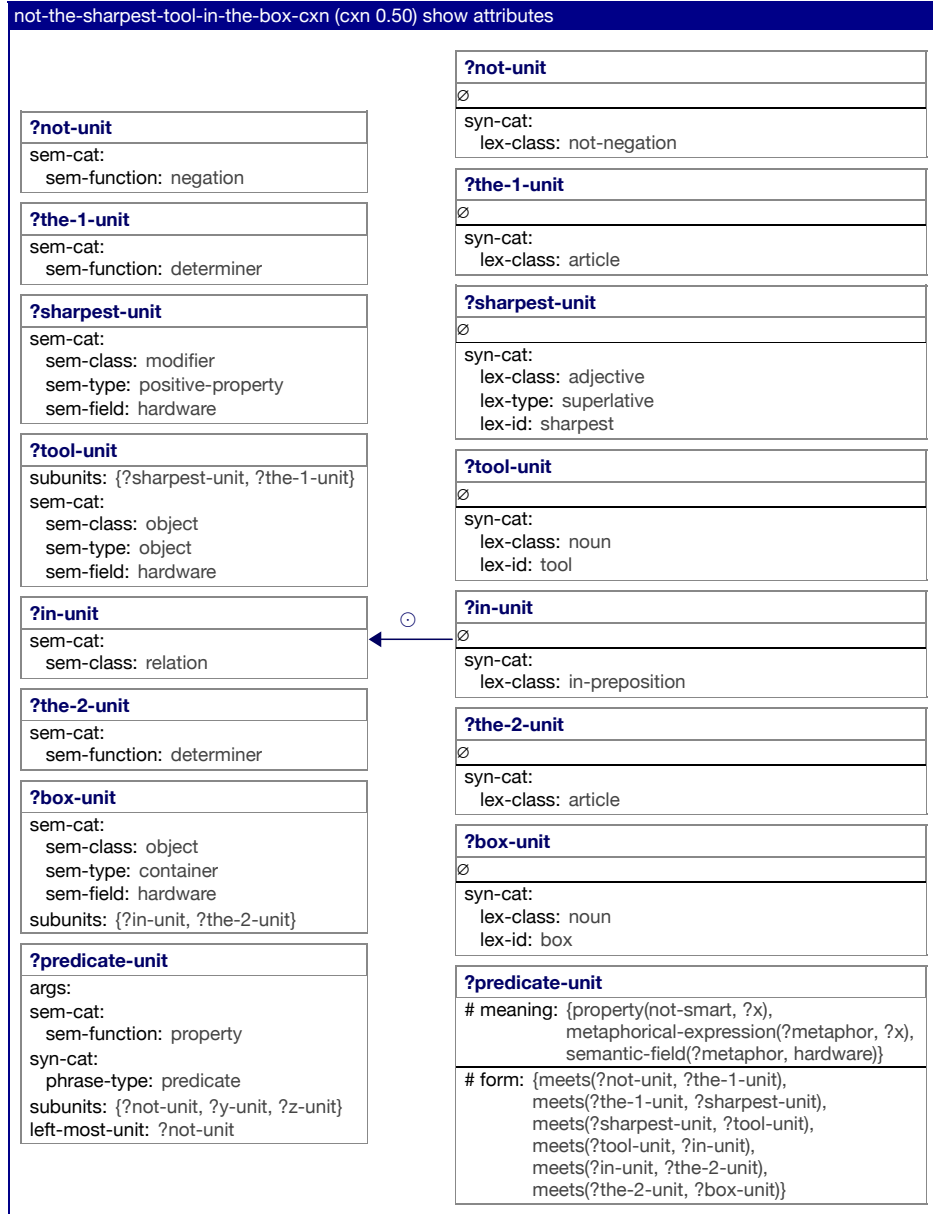


Figure 6: The idiomatic construction for ‘*not the sharpest tool in the box*’. Note that the construction includes the specific lex-ids for the words *sharp*, *tool* and *box* and specifies that the semantic field of the metaphor is ‘hardware’.

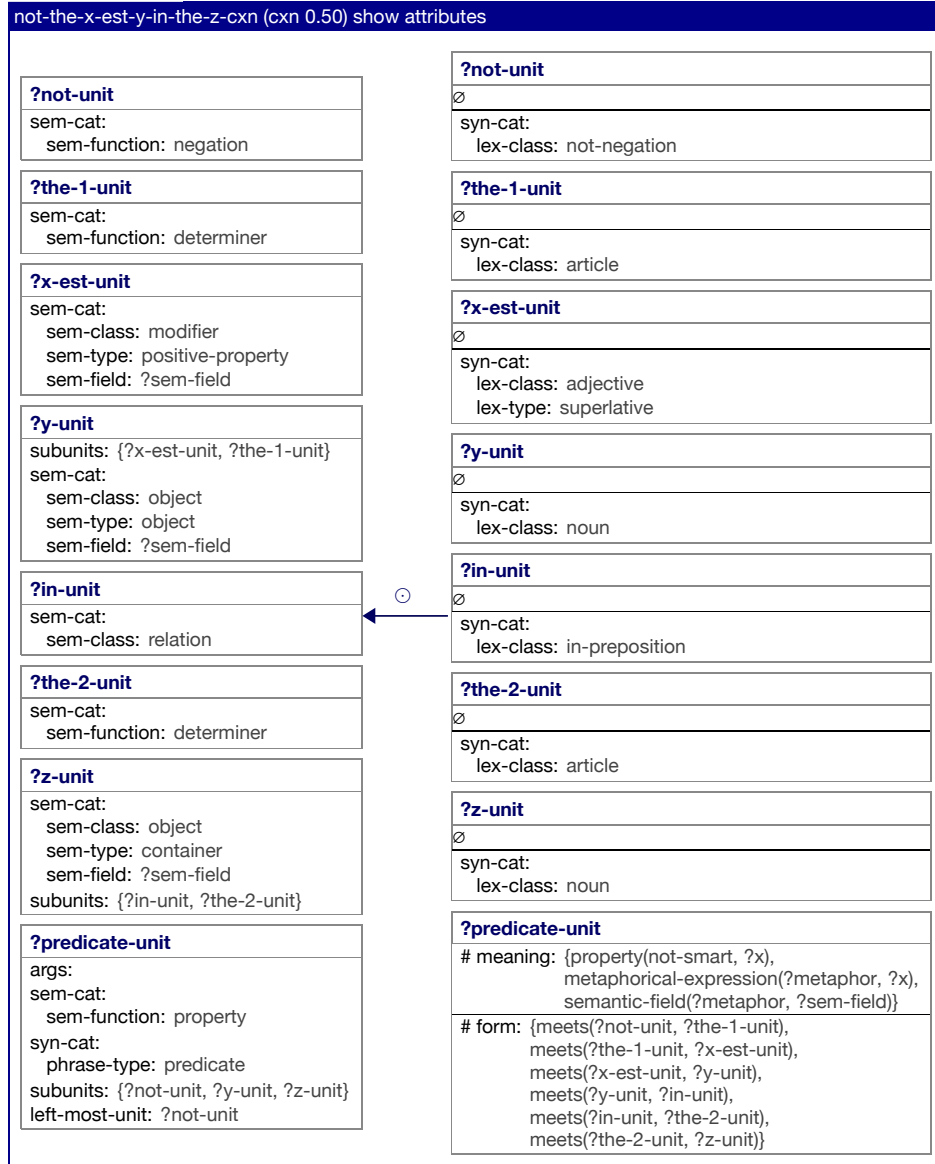


Figure 7: The generalised construction for 'not the x-est y in the z'. Note that (i) the semantic field is underspecified now, but is still required to be the same for *x*, *y* and *z*, and that (ii) the specific lex-id's of the words that will fill *x*, *y* and *z* have disappeared. All other constraints remain in place.

duction, and this is also the way in which they are implemented in FCG. In the case of the speaker, the generalised construction can be triggered by changing the argument of the **semantic-field** predicate in the meaning representation that it will express. If the semantic field is changed from ‘hardware’ to ‘clothing’ or ‘food’ respectively, the generalised construction might produce ‘*he’s not the smartest suit in the wardrobe*’, ‘*he’s not the crunchiest chip in the bag*’, or any other *the x-est y in the z* construct, as long as the other constraints on *x*, *y* and *z* are satisfied. More technical details and visualisations of these examples are shown in section 3 of the web demonstration.

When the generalisation and specialisation operators are used in formulation, there exists the challenge of defining which changes to the grammar are appropriate under which circumstances and which are not. In our system, this problem is handled in a very coarse way, by associating a cost to each kind of constraint violation (e.g. unit deletion, feature deletion and value relaxation) and limiting the overall cost of meta-layer operations. The design and implementation of a comprehensive theory of appropriate and inappropriate constraint violations falls outside the scope of this paper, but constitutes an interesting topic for further research.

5 Conclusion

In this paper, we have studied how computational construction grammar can account for creative language use. We have distinguished between two types of language-related creativity, and for each of these types, we have (i) discussed how they can be incorporated in computational construction grammar models, and (ii) provided concrete implementations of the required representations and processing mechanisms in Fluid Construction Grammar. The first type of creativity concerned the productivity of grammatical structures within the usual constraints of the language. This is handled by letting the constructions of grammar combine freely, as long as there are no conflicts. In FCG, this happens at the routine layer and is the standard way to process input utterances or meaning representations. The second type of creativity concerned utterances in which the usual constraints of the language are, to a certain extent, violated. Here, additional mechanisms for inventing novel constructions or for relaxing constraints within existing constructions are needed. In FCG, this happens at the meta-layer, where mechanisms such as anti-unification relax constraints in existing constructions and give rise to new, more general constructions. While the first type of creativity has drawn most attention in the history of linguistic theory, it is the second type that drives the innovations that allow language to emerge, evolve and adapt to new environments.

Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 732942. The authors would like to thank Thomas Hoffmann for the organization of an inspiring summer workshop on

construction grammar and creativity, as well as all researchers who participated in this workshop.

Works Cited

- Bergen, B. and Chang, N. (2005). Embodied construction grammar in simulation-based language understanding. In Fried, M. and Östman, J., editors, *Construction grammars: cognitive grounding and theoretical extensions*, pages 147–190. John Benjamins, Amsterdam.
- Beuls, K., van Trijp, R., and Wellens, P. (2012). Diagnostics and repairs in Fluid Construction Grammar. In Steels, L. and Hild, M., editors, *Language Grounding in Robots*, pages 215–234. Springer, Berlin.
- Boas, H. C. and Sag, I. A. (2012). *Sign-Based Construction Grammar*. CSLI Publications/Center for the Study of Language and Information, Stanford.
- Feldman, J., Dodge, E., and Bryant, J. (2009). Embodied construction grammar. In Heine, B. and Narrog, H., editors, *The Oxford Handbook of Linguistic Analysis*, pages 121–146. University Press, Oxford.
- Goldberg, A. (2006). *Constructions at Work: The Nature of Generalization in Language*. University Press, Oxford.
- Haspelmath, M. (1999). Why is grammaticalization irreversible? *Linguistics*, 37(6):1043–1068.
- Hoffmann, T. (2018). Grammar and creativity: Cognitive and psychological issues. *Zeitschrift für Anglistik und Amerikanistik*, 66(3).
- Keller, R. (1994). *On Language Change: The Invisible Hand in Language*. Routledge, London.
- Steels, L., editor (2011). *Design patterns in Fluid Construction Grammar*, volume 11 of *Constructional Approaches to Language*. John Benjamins, Amsterdam.
- Steels, L., editor (2012). *Experiments in Cultural Language Evolution*, volume 3 of *Advances in Interaction Studies*. John Benjamins, Amsterdam.
- Steels, L. (2017). Basics of Fluid Construction Grammar. *Constructions and Frames*, 9(2):178–225.
- Steels, L. and van Trijp, R. (2011). How to make construction grammars fluid and robust. In Steels, L., editor, *Design Patterns in Fluid Construction Grammar*, pages 301–330. John Benjamins, Amsterdam.
- Stickles, E., Oana, D., Dodge, E. K., and Hong, J. (2016). Formalizing contemporary conceptual metaphor theory. *Constructions and Frames*, 8(2):166–213.

- Van Eecke, P. and Beuls, K. (2017). Meta-layer problem solving for computational construction grammar. In *The AAAI 2017 Spring Symposium on Computational Construction Grammar and Natural Language Understanding Technical Report*”, number SS-17-02, pages 258–265, Stanford. Association for the Advancement of Artificial Intelligence.
- van Trijp, R. (2012). A reflective architecture for robust language processing and learning. In Steels, L., editor, *Computational Issues in Fluid Construction Grammar*, pages 51–74. Springer, Berlin.
- Wiggins, G., Tyack, P., Scharff, C., and Rohrmeier, M. (2015). The evolutionary roots of creativity: Mechanisms and motivations. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 370(1664).