

The Emergence of Internal Agreement Systems

Katrien Beuls, Luc Steels and Sebastian Höfer

This paper is the author's draft and has now been officially published as:

Beuls, Katrien, Steels, Luc and Sebastian Höfer (2012). The Emergence of Internal Agreement Systems. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*, 233 – 256. Amsterdam: John Benjamins.

Abstract

Grammatical agreement means that two linguistic units share certain syntactic or semantic features such as gender, number or person. Agreement has a variety of grammatical functions. One of them, called internal agreement, is to signal which words are grouped together as part of the same phrase. This chapter explores how a population might self-organize such an agreement system. We argue that this happens when speakers attempt to reduce processing effort and avoid ambiguities.

1. Introduction

Grammatical agreement takes place when two linguistic units (for example two words) share certain features with each other, such as gender, number, person, or case. For example, in the Spanish utterance (1), the definite article “la”, the noun “chica” and the adjective “estudiosa” all agree with respect to gender (feminine) and number (singular).

- (1) *la* *chic-a* *estudios-a*
DEF.F.SG girl(F)-SG diligent-F.SG
'the diligent girl'

One word normally acts as the *source* of features (in this example the noun “chica”) and other words, called *targets*, obtain features from the source. Agreement can be explicitly marked by the form of the word (“el” versus “la”) or by morphemes (“-o” versus “-a”). Agreement is used in many, although certainly not all, of the world’s languages and there is a huge literature describing the many variations that have been observed in the roles, features and forms used (Corbett, 2006; Boeckx, 2006).

1.1. Variation in Roles of Agreement

Here are some examples of common roles for agreement:

- *To support pronominal reference.* For example, in the utterance “the boy gave his mother her book back”, we know that most probably “his” refers to the referent of “the boy” and “her” to the mother because of the agreement in number and gender between the possessive pronouns and their referents. Notice that in French, this information would not be expressible this way because possessive pronouns have to agree with their head nouns, as in Example 2.

(2) *le garçon a donné son livre a sa mère*
 DEF.M.SG boy(M)-SG AUX give-3sg.PST POSS.M book(M) to
 POSS.F mother(F)
 (literally) ‘The boy returned ‘his’ book to ‘her’ mother’.
 The boy gave his mother her book back.

- *To express subject-predicate-object structure.* For example, in the utterance “the boy gives a book”, subject and verb agree with respect to person and number. Hungarian requires under certain circumstances that there is not only agreement between subject and verb but also between verb and object (Beuls, 2011) as shown in Example 3, where the verb has second person singular marking “-l-” due to the direct object (DO) “you” and first person singular marking “-ek” from the subject (S) “I”.

(3) *Szeret-l-ek té-ged.*
 love-(2sg.DO-1sg.S) you-ACC
I love you.

- *To express which constituents of a nominal phrase belong together.* For example, in the German utterance in Example 4, we know that the article “der” (the), the adjective “aufmerksame” (attentive) and the noun “Jäger” (hunter) all belong to the same phrase, because they all agree with respect to case (nominative), number (singular) and gender (masculine).

- (4) *Der* *aufmerksame* *Jäger*
 the.NOM.M.SG attentive-NOM.M.SG hunter(M).NOM.SG
schießt *das* *schnelle*
 shoot-3SG.PRES the.ACC.N.SG fast-ACC.N.SG
Reh
 deer(N).ACC.SG
 ‘The attentive hunter shoots the fast deer.’

This use of agreement is usually called *internal agreement* (Lehmann, 1988) and is an alternative to marking constituent structure with word order or intonation. Consequently many languages with a strong agreement system support significantly freer word order, as in the Latin example in 5 (from Horace, Carminum I) where agreement between the noun “spe-m” (hope) and the adjective “long-am” (long) is enough to signal that they belong to the same noun phrase even if they are far apart in the sentence. Their ordering could be reversed or these words could be located elsewhere in the sentence without leading to incomprehension.

- (5) *spe-m* *nos vet-at* *inchoare long-am*
 hope-acc us inhibit-pres.3sg start-on long-acc.
 ‘[...] stops us from cherishing long hope.’

1.2. Variation in Features used in Agreement

Substantial variation is also found across the world’s languages with respect to which features are used in agreement systems. We usually see that these features must have been motivated, at least in an early stage, by semantic criteria. For example, the distinction between masculine and feminine animate nouns goes back to the natural gender distinction between male and female. Other typical semantic features include definiteness or animacy (Corbett, 2006). Occasionally we find more fine-grained distinctions which reflect

a particular way of structuring the world. For example, Bantu languages have an agreement system based on classes that reflect distinctions between persons, plants, fruits, things, animals, locations, and items associated with these (Katamba, 2003). Learning a language thus requires learning a complex categorization of reality which hangs together with the belief systems, cosmologies and world knowledge of the language community.

Although semantics provides usually the original foundation for an agreement system, two grammaticalization paths are often observed (Heine & Kuteva, 2007):

1. Semantic distinctions become obscured. A famous example is the category containing women, fire and dangerous things which is one of the classes playing a role in the agreement system of the Australian aboriginal language Dyirbal (Lakoff, 1990).
2. Agreement features become conventional. This is a further step away from real world semantics and then requires that the language learner memorizes feature assignments in the lexicon. For example, if a language has only a masculine/feminine distinction (as in French or Spanish) an issue arises to which category inanimate objects are to be assigned and this choice is necessarily conventional. Thus, table is feminine (“la table”) in French but masculine (“der Tisch”) in German. “Tram” is masculine in Polish (“tramwaj”), feminine in Czech (“tramvaj”) and neuter in Romanian (“tramvai”).

1.3. Variation in Marking of Agreement

Finally there is also substantial variation in terms of the ways in which agreement shows up in the utterance itself. We get the following possibilities:

- No overt marking. The features involved in agreement could be implicit in the meaning of the word and therefore do not need to be marked. For example, “fille” (French. girl) is feminine because it refers to a female person.
- A marker is used and it is the same for all units which share the feature. This is the case in Example 6 from Swahili (Welmers, 1973). The same prefix “ki-” is simply attached to every word which shares feature Class 7 (with core meaning utensils and hand tools).

- (6) *ki-kapu* *ki-kubwa ki-moja ki-li-anguka*
 SG-basket(7/8) 7-large 7-one 7-PAST-fall
 ‘One large basket fell.’

We see a similar situation in Latin, as in “*illorum saevorum virorum*” (of those savage man) where the “-orum” suffix marks in all cases plural genitive.

- A marker is used but its form varies according to the type of units that share the feature. This is the case in the Example 7 from Polish where the accusative forms in the noun phrase “*duż-ego człowiek-a*” vary between the adjective and noun inflections.

- (7) *Widzę duż-ego* *człowiek-a.*
 (I) see big-ACC.M.SG man(M)-ACC.SG.
 ‘I see a big man.’

- Different markers are used and moreover the same marker may express many different functions. For example, the German article “*der*” can express nominative masculine singular (“*der Mann*” - the man), genitive or dative feminine singular (“*der Frau*” - (of) the woman) or genitive masculine singular or plural (“*der Männer*” - of the men). This phenomenon is known as *syncretism* and implies much more sophisticated language processing to handle multiple hypotheses van Trijp (2011).

The deep variation in agreement systems shows that there is a strong cultural component in their origins and evolution. In this chapter, we are interested in the question how agreement systems may culturally emerge, become shared, and get more complex. Agreement systems that become too complex eventually collapse, particularly when the population of speakers is large. This is presumably because such systems have become too hard to learn to still sustain proper cultural transmission or because the morphological endings have eroded so that there is not enough evidence for practicing or learning the basis of an agreement system.

We follow the theoretical framework of linguistic selectionism as outlined in the first chapter of this book (Steels, 2012), which argues that the cultural origins and evolution of language systems (and therefore also agreement) is driven by the need to optimize communicative success. Speakers and hearers strive

for enough expressive adequacy while minimizing cognitive effort. As language users streamline or complexify their language, they recruit different cognitive capacities to configure effective language strategies. But only those strategies that lead to consistently successful language interactions in the population are retained.

To substantiate these hypotheses implies going through the following steps:

1. Identify language strategies for the chosen domain with increasing degrees of complexity.
2. Operationalize these language strategies and show how they allow a population to self-organize a language system.
3. Demonstrate the selective advantage of each strategy. Selection is not taken here in the biological sense but in terms of improvement of communicative success while minimizing cognitive effort.
4. Show how populations could move progressively from one strategy to another.

We apply this methodology in this paper towards understanding the origins of agreement systems, focusing exclusively on internal agreement systems. Moreover we only show the first initial stages in the application of this methodology.

2. The Sticker Theory

We start from the hypothesis that internal agreement systems arise because speakers want to signal to hearers how some of the words in a sentence are to be grouped together into word-groups or phrases. The simplest way to do so is to add a ‘sticker’ to each of the words that forms part of the same group. This is a common sense principle that is used by everyone, for example, to organize clothes or utensils in different boxes, or to signal that certain people belong to a particular group. In these cases stickers may be circles of different colors or figures of different shapes, markers on uniforms, styles of hats, badges, etc. It does not matter which form is used for a sticker, as long as all members of the same group use the same. In the case of language, a sticker could just be an arbitrary string, like “bo” which is attached as a prefix or suffix to a word.

Although the sticker principle is a good way to start thinking about the emergence of agreement systems, human natural languages seldom (maybe

never) use formal stickers as their starting point (Corbett, 2006). Instead, the stickers appear to be meaningful. Sometimes they are words that exist already in the language and then undergo a grammaticalization process, losing their original meaning and some of the complexity of their form. Meaningful agreement markers are like iconic stickers, such as a picture of a vegetable which is tagged on all boxes that contain vegetables, or a small airplane icon which is used as an insignie by pilots.

One can see several advantages for using meaningful stickers as a basis of an agreement system:

1. If full fledged words are used as initial stickers, their meaning can be immediately inferred.
2. No additional stickers are needed for words acting as sources (such as words referring to female persons when feminine is chosen as agreement feature).
3. Human memory works best when items to be memorized have existing semantic associations. Meaningful stickers are therefore easier to remember compared to formal stickers.

In the case of formal stickers, language users need only as many different markers as there can potentially be words in a clause, which is always a small finite number. When stickers are meaningful, there is the risk of an exploding number of stickers, because if existing semantic categories are not applicable to a specific case. For example, gender distinctions are not applicable to inanimate objects, and therefore a new semantic distinction must be introduced. But this runs counter to the general principle of ‘economy of means’, which is to use as few forms as possible in a single utterance in order to convey as much meaning or structure as possible. The more stickers, the more memory is needed to remember them and the more time learners need to acquire them.

The historical development of languages shows that often markers begin to lose their original semantic clarity as the features they express begin to reflect wider and wider classes, and hence they become entirely conventional. For example the masculine/feminine distinction started out as natural gender and then became applied to inanimate objects in arbitrary ways. Or, class 7 in the agreement system of Swahili started out as being the class of artifacts and then became extended to objects which resemble artifacts or actions which involve artifacts.

Typological studies have shown that there are two more developments that have happened or are happening in the languages of the world:

1. Often a particular marker is no longer used for a single purpose (namely marking agreement) but becomes used for other purposes as well (for example marking the syntactic category of a word or its case), i.e. they are multi-functional. This reflects again the ‘economy of means’ principle.
2. The economy of means principle comes even further into play when the total inventory of markers is reduced as much as possible, which means that markers become ambiguous and their meaning or function has to be determined on the basis of context. This is what we get in German for example, where the article “der” marks for many different combinations of case, gender and number.

Given these observations it is useful to research agreement systems in steps:

```
no stickers > formal stickers >  
  meaningful stickers > conventional stickers >  
    multi-functional markers > syncretism
```

There are hints for evidence of similar stages in the historical evolution of language, as discussed for example by Greenberg (1978); Givón (1976); Corbett (2006), but that is not the main purpose here. We want first of all clear models for each stage before tackling their empirical relevance.

In the rest of this chapter, we present experiments concerning only the first three steps, from no stickers to formal stickers and then to meaningful stickers. Research results have been achieved concerning the remaining steps as well. For example, van Trijp (2011) has proposed processing models for syncretism and shown that the German syncretic agreement system leads to greater optimality in the marker inventory.

3. Experimental set-up

In order to simulate the emergence of an agreement system, we set up a language game in which an agreement system might potentially be useful. We employ a Description Game similar to the one used in the case grammar experiments reported in the chapter by van Trijp (2012). Both agents that participate in a language game are presented with a situation that contains a

number of objects and the game is a success if the description given by the speaker fits with the situation as perceived by the hearer.

More concretely, assume a population P of agents, and a world W consisting of a set of individual objects. Two agents are randomly selected from the population and take on the roles of speaker and hearer respectively. The situation S contains a subset of the world W , represented as a set of predicates and arguments in prefix notation. For example, the following situation involves two objects (labeled `indiv-1` and `indiv-2`) which each have a series of properties, such as `big`, `ball`, `chair`, etc.

```
(big indiv-1) (ball indiv-1)
(chair indiv-2) (red indiv-2)
```

These situations are automatically generated by a script but they could be derived from real world scenes as shown in experiments reported in other chapters, such as in the Grounded Naming Game experiment (Steels & Loetzsch, 2012). For the present experiment, there is no uncertainty in perception, i.e. both agents are assumed to have perceived the situation in exactly the same way. Moreover the situation always contains at least two objects.

The speaker chooses at least two objects and chooses at least two predicates for each object. The speaker then describes these objects by expressing these predicates, making use of the initial lexicon and grammar. For example, assume the speaker has the following lexical associations:

```
(big ?x) ⇔ "big"
(ball ?y) ⇔ "ball"
(chair ?z) ⇔ "chair"
(red ?u) ⇔ "red"
```

then he might produce the utterance "big red chair ball" to describe the two objects in the situation given earlier. Agents never use word order as a syntactic device, as we want to focus on agreement only. Agreement features will be expressed using morphology, more specifically suffixes. Thus, when the speaker has already morphological markers, he may produce an utterance such as "big-dug chair-bo ball-dug red-bo", where `-dug` and `-bo` are agreement suffixes.

The hearer parses the utterance using his own lexicon and the agreement markers learned so far and signals success when the reconstructed meaning applies to his own situation model. If this is not the case, the game is a

failure. If he detects markers he does not know, the hearer tries to deduce the function of the markers by inspecting his internal situation model and then stores new hypotheses in his memory.

Because we want to focus on exploring agreement, the lexicons of the agents have been scaffolded, which is justified as many experiments now exist on how lexicons may arise, some of them shown in earlier chapters of this book. Agents are initialized with an inventory of lexical constructions implemented in Fluid Construction Grammar (FCG) Steels (2011b) (see the appendix for an example). To keep the experiment simple and easier to follow, we use a one-to-one correspondence between words and predicates and we use English-like words. Agents therefore start out with a kind of Pidgin lexicon without grammar and then a creolization process starts introducing grammar in the form of agreement.

When an agent parses the utterance “big red chair ball”, he reconstructs a possible meaning which is equal to:

(big ?x) (ball ?y)
(chair ?z) (red ?u)

Each word introduces a different variable and the hearer cannot know, from linguistic information contained in the utterance itself, which of these are co-referential. There are three possible interpretations - before inspecting the situation model:

1. “red ball” and “big chair” so that ?x and ?y and ?z and ?u are co-referential, which means that ideally the meaning expressed by language should have been

(big ?x) (chair ?x) (red ?y) (ball ?y)

2. “big red” and “ball chair” so that the meaning should have been

(big ?x) (red ?x) (chair ?y) (ball ?y)

3. “red chair” and “big ball” whose meaning is:

(big ?x) (ball ?x) (chair ?y) (red ?y)

Even with the restriction of only two words per object, the space of possible combinations grows rapidly. For an utterance of length 10 describing 5 objects, we have already 945 possible solutions.¹ If an object can be described with a variable number of words, which is the case in the present experiment, then the combinatorics becomes even more daunting, which demonstrates why marking phrase structure is useful.

We are now going to study different language strategies that signal phrase structure and test them in populations of agents playing language games. The general framework for playing the game is always the same. We discuss first an experiment where there is no agreement system nor any other marking of phrase structure, then an experiment using formal stickers, and next an experiment based on meaningful stickers.

4. Absence of an agreement system

Let us experiment first with a strategy where there is no marking of constituent structure through agreement or any other means, in other words where agents use a lexical language without syntax. However agents already use a grammatical construction to group together the words whose meanings are about the same object. The principles behind this construction are now explained as they form the basis of the agreement system to be discussed in the next sections.

Let us look at production first. The lexical items cover parts of the meaning of the utterance and introduce new units for the relevant word strings. The lexical item also has to specify what this meaning is about, i.e. the object in the situation over which the predicates holds. For example, the word “big” is about ?x in the association:

$$(\text{big } ?x) \iff \text{"big"}$$

To express the meaning

$$\begin{array}{l} (\text{big indiv-1}) (\text{ball indiv-1}) (\text{black indiv-1}) \\ (\text{chair indiv-2}) (\text{red indiv-2}) \end{array}$$

1. The set of possible combinations for two words per object is given by formula 8 (De Beule, personal communication), where r stands for number of objects in a situation that are described in the utterance.

$$(2r - 1)!! = \frac{(2r)!}{r!2^r} \tag{8}$$

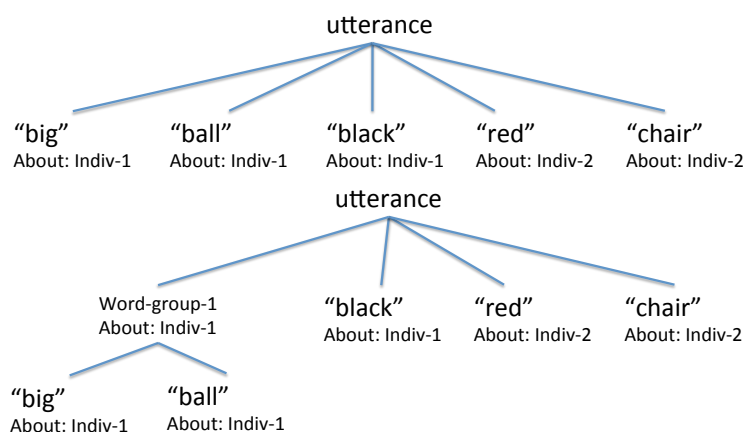


Figure 1. *The top picture shows the structure after lexicon lookup. Units have been created for each word. The bottom picture shows the structure after the first application of the `create-word-group` construction.*

we then get after lexicon lookup the initial structure shown in Figure 1, top. The precise definition of the scaffolded lexical constructions in Fluid Construction Grammar is given in the appendix.

The phrasal constructions operate over the units for the words. The first construction, called `create-word-group`, is on the look-out for two words that are about the same object, in this example “big” and “ball”, and combines them into a word-group, here labeled `word-group-1`, as shown in Figure 1, bottom. The word-group as a whole is about the same object as the individual words. The same construction can also create a second word-group for “chair” and “red”.

Because there can be more than two words about the same object, an additional construction is needed, called `expand-word-group`, which adds a word to an existing word group if both of them are about the same object. And this is how “black” gets added to the first word-group, giving the final result in Figure 2. Precise definitions of the `create-word-group` and `expand-word-group` constructions are given in the appendix.

All this is relatively straightforward. But for parsing the situation is more complex. The hearer performs lexicon look-up and ends up with a list of units

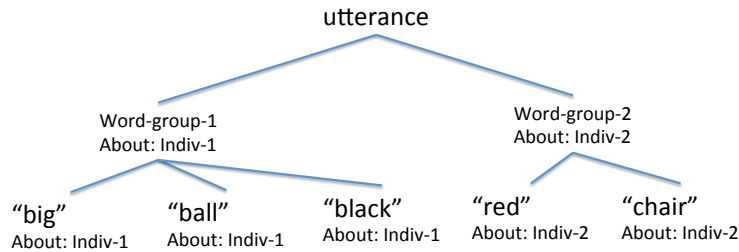


Figure 2. *This figure shows the resulting structure after expand-word-group has added the unit for “black” to word-group-1 and create-word-group has created a second phrase labeled word-group-2 for “red chair”.*

for words, but now, because there is no syntax, he has no clue about which words should be put together into word-groups and so a search space arises. As we have seen earlier, an exhaustive search of all possible combinations would lead to a combinatorial explosion - indeed otherwise we would not need syntax. However agents can still parse within reasonable bounds if they make use of the situation model at every step in the search process to test whether a search node should be expanded further or not. Clearly if the meaning obtained so far does not match with the situation model, further search can be abandoned.

We have implemented this strategy and it leads to a search process as shown in Figure 3. The speaker applies the grammatical construction to combine two words, but if the meaning of these words is not about the same object according to the situation model, this path is no longer pursued. Figure 3 shows the process of parsing the utterance “big chair black ball red”. The light nodes indicate that a search path has been abandoned due to a failed match with the situation model.

Figure 4 shows typical experimental runs. We plot the cognitive effort which is the number of nodes in the search tree whose meanings do not match with the current situation. Because agents do not perform exhaustive search this is the only way to gauge cognitive effort. As shown in Figure 4 the effort stays more or less constant, because in this experiment the utterance length is on average 8. Of course if the situation model would not be used or available to reduce the search, cognitive effort would be much higher.

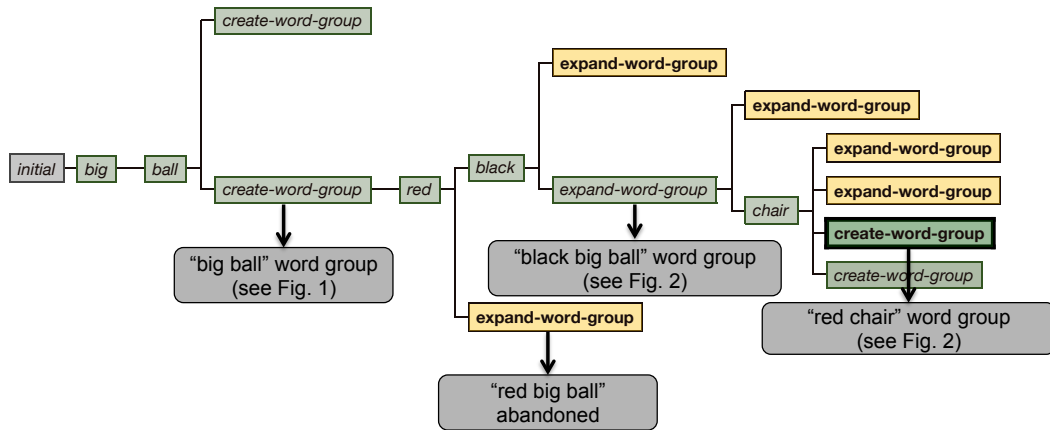


Figure 3. *The search tree when parsing “big chair black ball red”. Each node represents the application of a construction. The agent inspects the situation model after every step to find out whether a particular path in the tree is still viable and immediately cuts away nodes (in light color) which do not make sense from a semantic point of view.*

It should now be clear that without a syntactic marking of phrase structure, agents need a lot of cognitive effort, either because they need to search a very large space or because they have to check all the time their situation model to cut away parses that do not fit. The latter strategy is only possible when a situation model is available and when there is no or little risk of semantic ambiguity.

5. Agreement based on formal stickers

Our next step is to investigate how a system of formal stickers can arise. As explained earlier, formal stickers signal phrase structure by adding a sticker to every word within the same word-group. In production, agents create word-groups based on whether their meanings are about the same object and they add the required stickers in the form of suffixes. In parsing, agents use these suffixes to decide which words belong together. The remainder of this section provides technical details of the constructions involved, the language strategy

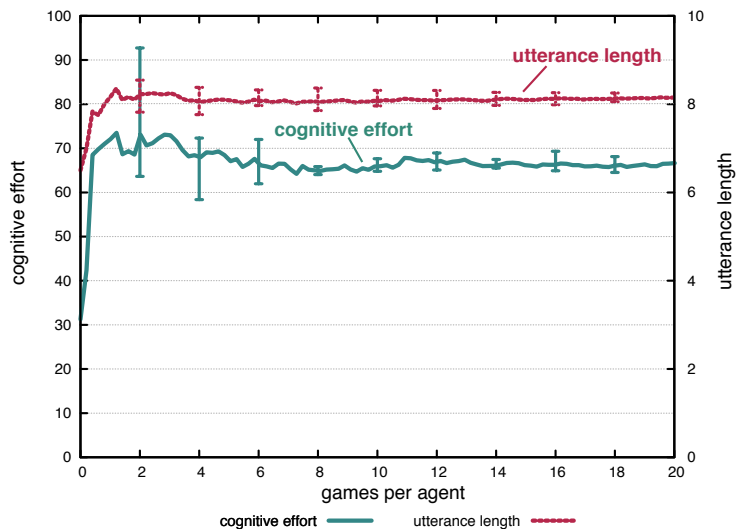


Figure 4. The graph shows cognitive effort and utterance length for four series of 100 games played by a population of 10 agents. Error bars: standard deviation.

that gives rise to a formal agreement system and results of experiments with this strategy.

5.1. Grammatical Constructions

The addition of formal stickers can be implemented in a straightforward way by extending the `create-word-group` construction discussed earlier with the additional information that the same suffix must be present for the two words (see the appendix for technical details). We call such constructions `create-word-group-sticker` where *sticker* is the sticker itself, as in `create-word-group-cui` for the sticker “-cui”. One construction is needed per sticker. The `expand-word-group-sticker` construction also needs to be extended to take the sticker into account, but we only need a single construction which looks into the word-group to find which sticker is already used and then ensures that the word has the same one.

Speakers create new formal stickers when they need them (as discussed in the next subsection). They also learn the stickers used by others. So after a while the issue arises which sticker they should prefer.

The meaning of all the stickers is the same, namely to express that the words to which they are attached to are about the same object in the situation model. But this does not mean that the stickers are synonyms in the traditional sense of competing forms for the same meaning that need to be damped. They cannot be considered to be homonyms either. So a lateral inhibition dynamics as normally used in lexicon formation cannot be employed here. Instead we use a frequency-based approach.

If an agent observes that a sticker is used more, he should prefer this sticker himself as well. This can be achieved when agents track the use of a particular sticker, which is done here with a simple counter. Agents increase the usage counter when, as a hearer, they encounter the sticker. When they need to decide which sticker to use, they should use the sticker which has the highest usage counter so far. This leads to a self-enforcing effect because the more a sticker is used in the population the more agents start preferring it, and the faster the counters will increase.

5.2. Formal Sticker Strategy

At the start of the experiment, the agents are scaffolded with the lexicon and the two word-group constructions discussed in the previous section. There are no constructions yet for stickers. The language strategy used by the agents to self-organize a system of formal stickers works as follows:

Invention (by speaker) The speaker first produces an utterance using the available lexicon and grammar. But before transmitting the resulting utterance, the speaker re-enters the utterance in his own parsing system and computes the number of possible parses, which is taken as an indication of the cognitive effort needed. If there is more than one possible parse, the speaker introduces a new sticker using two steps:

1. The speaker creates a new form by assembling a random combination of consonants (C) and vowels (V) according to three possible patterns: CV, CVC or CVCV.
2. The speaker builds a new *create-word-group-sticker* construction. The newly created form is used as the morphological suffix. If no

`expand-word-group-sticker` construction exists yet, such a construction is made as well, but only one is ever needed.

A new sticker construction is made for every phrase that was not explicitly marked. Once the necessary constructions have been built, the speaker reproduces the utterance and transmits it to the hearer.

Adoption (by hearer) The hearer first parses an utterance. But when unknown stickers are present in the input, parsing will be incomplete. The hearer can then first construct new sticker constructions for each of the unknown suffixes. If no `expand-word-group-sticker` construction exists yet, such a construction is made as well, but only one is ever needed. Once all these constructions have been made, the hearer can re-parse and interpret the sentence.

Alignment Every sticker construction has a usage counter, which is initialized at 1 when the construction is created. Each time the sticker is encountered by the hearer, he increases the counter of the relevant construction.

5.3. Results

We will track the cognitive effort of the agents, similar to the previous experiment, by measuring the size of the search space. Figure 5 shows the outcome of an experiment with 10 agents. Cognitive effort is equal to $\frac{T_s}{M}$ where T_s is the total number of stickers in use in the population and M the maximum number of objects. The cognitive effort indeed drops dramatically as soon as agents start to use the formal sticker strategy. We are also interested to see whether agents reach a shared inventory, i.e. whether their preferences for the usage of stickers is converging so that the inventory is easier to learn for new incoming agents, less memory is needed to store them, and lookup is faster. We therefore track the variation in the population by measuring the average number of markers per feature. If this is equal to 1, then there is only one marker per feature and hence agents have reached total convergence. Figure 5 shows that this is also happening. So the proposed language strategy entirely achieves its objective.

A snapshot of the top five highest frequency sticker constructions of two different agents in one experiment is shown in Table 1. The markers that belong to the top five are the same for both agents, but the frequency is different for constructions on rank three and four. In practice, this means that

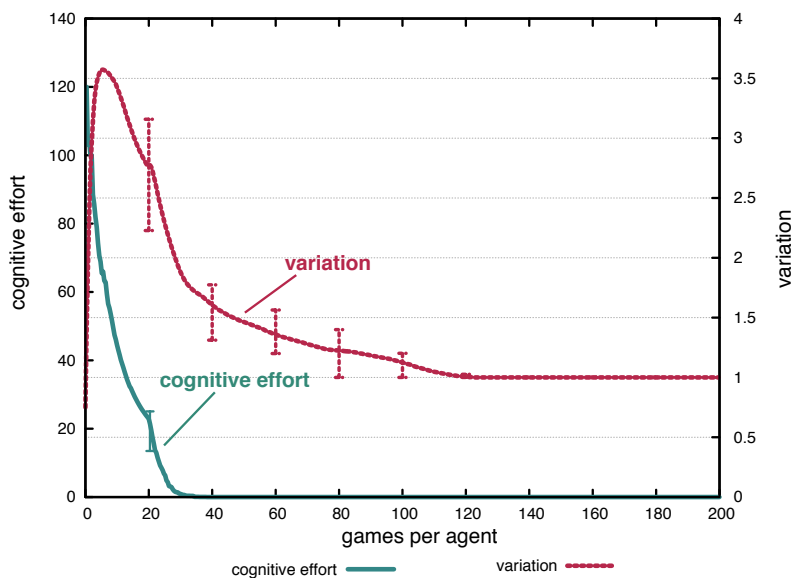


Figure 5. The cognitive effort (hearer) and the variation in sticker usage is plotted in function of the number of games played per agent. Settings: 10 agents, 1000 language games, 4 game series, maximum number of objects = 5. Error bars: standard deviation.

Table 1. Top five highest frequency sticker constructions for two agents after a number of games.

rank	agent-1		agent-2	
	cxn name	freq.	cxn name	freq.
1	-cui	25	-cui	18
2	-xu	25	-xu	17
3	-xufi	15	-buc	14
4	-buc	13	-xufi	12
5	-saev	5	-saev	8

when agent 1 is describing a situation with three individuals he will use for the third individual -xufi, whereas agent 2 would prefer -buc.

When one examines the marker frequencies of individual agents, we see a clear tendency where the markers that are used more will rapidly increase even further. Thus Figure 6 shows the development of a single agent's markers in terms of their frequency scores. There are five markers that distinguish themselves from the rest. The difference in frequency between these five can be explained by looking at the size of the meaning spaces that are generated. These contain at least two and at most five objects, and thus require always two and occasionally five markers.

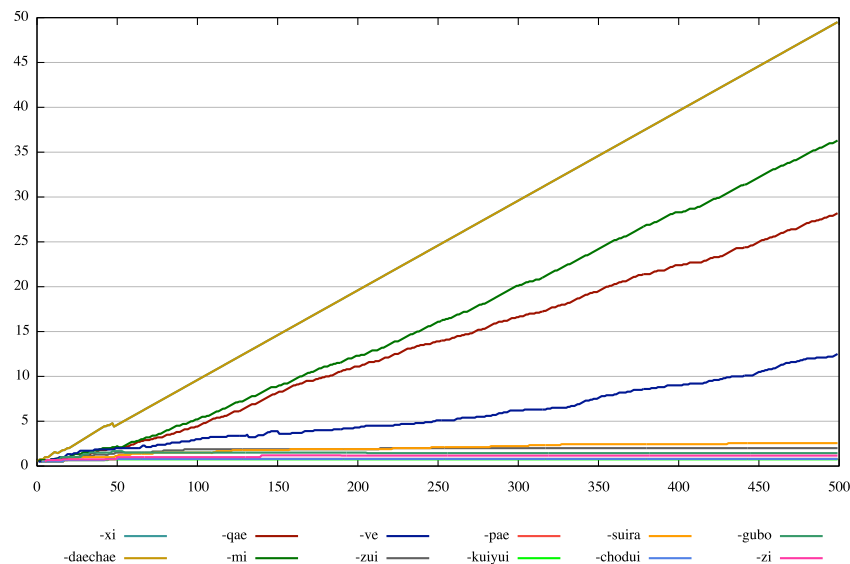


Figure 6. *The development of markers in the inventory of a specific agent. The agent uses five markers most frequently: -xi and -daechae (same rank), -mi, -qae, and -ve. The lower ranked markers only occur when enough objects appear in the utterance.*

6. Agreement based on meaningful stickers

Now we look at the processes which could explain how agents can self-organize an inventory of meaningful stickers. We take inspiration from the agreement system of Swahili - although the experiment should not be seen as a serious reconstruction effort. The next subsections discuss again first the nature of the needed constructions, then the language strategy and then the experimental results obtained.

6.1. Grammatical Constructions

As explained in the introduction to this chapter, agreement systems assume a particular categorization of the world which is based on cultural and world knowledge. For example, Swahili uses the following classes for single objects: person (e.g. child, man, teacher, mother, doctor), nature (e.g. tree, moon, mountain, river, fire), groups (e.g. glasses, eyes, scissors, feathers, corn), animals (e.g. bird, fish, goat, snake, cow) and artefacts (e.g. chair, book, knife, house, pen) Katamba (2003). It would lead too far from our current objective in this paper to model the underlying ontology. Instead we extend lexical constructions so that they have already information about this categorization, and we use these Swahili categories. Thus, the lexical construction for "child" now not only specifies the meaning, the form, and the referent of the meaning, but also that this word introduces an object that belongs to the class person.

To clearly understand the dynamics of the experiment, we assume a one-to-one mapping between these classes and meaningful stickers, although obviously the situation is much more dynamic in human languages and is reminiscent of the multi-dimensional lexical meaning dynamics discussed in an earlier chapter Wellens et al. (2008). In fact, a follow up experiment, reported in Beuls & Höfer (2011), shows already how out of a multitude of possibilities, agents can select one feature bundle as the meaning of a sticker.

The definition of semantic features for agreement in the lexicon is only relevant for words which can act as the source of agreement, which are words that have a referring semantic function, such as "box" or "woman". A word such as "big" can be applied to any kind of object, whether it is a person, a natural object, an animal, etc. We say that such words have a modifying function, because they cannot be used on their own to refer to objects. This difference in possible semantic function is also declared in the lexicon.

The grammatical constructions that are required are similar to the ones used in the formal sticker experiment. There is a *create-word-group-feature-sticker*, where *feature* is the meaning of the sticker and *sticker* is the form, as in *create-word-group-person-cui*. In this kind of construction the constraints on semantic categories have become explicit and specific to each sticker. For example, if the "-cui" suffix is for objects belonging to the class person, then the grammatical construction contains information about this category as well (see the appendix for details). Words with a referring function can act as a source and are therefore possible

bases for creating a new word-group. Words with a modifying function will obtain the semantic features from their source.

There is also a single generic **expand-word-group-feature-sticker** which expands an existing word-group when they are about the same referent. There is agreement with respect to the feature for the word and the word-group and the word gets the same sticker as used for the other words in the word-group. The appendix provides the precise definition of this construction.

6.2. Semantic Sticker Strategy

At the start of the experiment, the agents are scaffolded again with the lexicon and the two word-group constructions discussed in section . There are no constructions for stickers yet. The language strategy used by the agents to self-organize a system of meaningful stickers works as follows:

Invention (by speaker) The speaker first produces an utterance using the available lexicon and grammar and re-enters the utterance in his own parsing system to compute the cognitive effort. If there is more than one possible parse, the speaker introduces a new sticker by creating a new form and then builds a new **create-word-group-feature-sticker** construction with the newly created form. The difference with the formal sticker strategy is that the speaker also adds semantic constraints to the construction, namely those that are associated through the lexicon with a word that has a referring function. If no **expand-word-group-feature-sticker** construction exists yet, such a construction is made as well, but only one is ever needed. Once the necessary constructions have been built, the speaker reproduces the utterance and transmits it to the hearer.

Adoption (by hearer) The hearer first parses an utterance. When unknown stickers are present in the input, parsing will be incomplete. Rather than simply constructing a new formal sticker construction, the hearer first needs to find out which semantic features can be used to constrain the unknown sticker. Once this has been found a new **create-word-group** construction can be built for the sticker. If no **expand-word-group-feature-sticker** construction exists yet, such a construction is made as well, but only one is ever needed.

Alignment Given that stickers have now a particular meaning, a lateral inhibition update scheme (which is much more efficient than pure frequency-based alignment) can be employed. We use a learning rule that does not

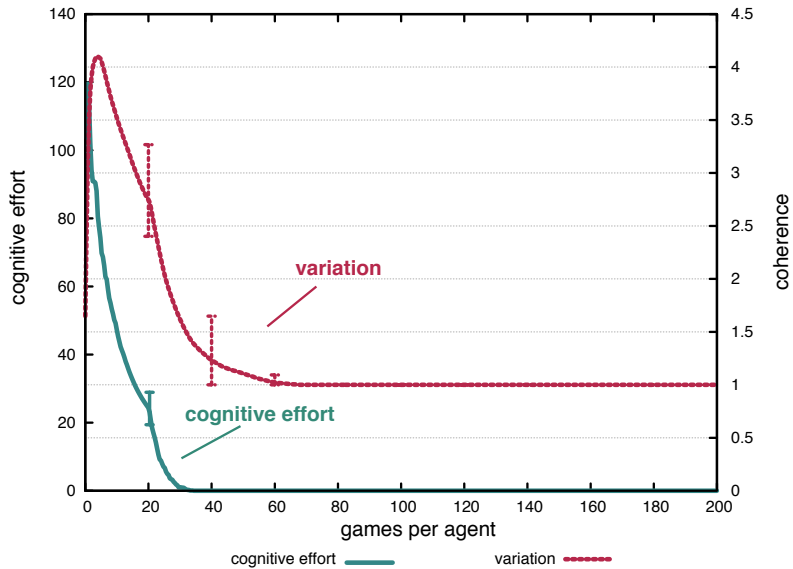


Figure 7. The cognitive effort (hearer) and the variation in sticker usage in function of the number of games played per agent. Settings: 10 agents, 1000 language games, 4 game series, maximum number of objects per situation = 5.

simply increase or decrease the score with a fixed δ but uses a more progressive update rule: In case of success, the score s_i of the used name is updated by $s'_i = s_i(1 - \delta) + \delta$ and in case of failure by $s'_i = s_i(1 - \delta)$, where $\delta = 0.2$. To make a comparison possible with the formal sticker strategy, only the hearer uses this update rule.

6.3. Results

We track again the cognitive effort of the agents by measuring the size of the search space and the degree of convergence towards a common inventory of meaningful stickers. Figure 7 shows the outcome of an experiment with 10 agents. The cognitive effort again drops dramatically as soon as agents start to use the meaningful sticker strategy. So this language strategy also achieves its objectives. The rate of cognitive effort optimization is the same as for formal

stickers, but the convergence towards a shared norm is faster because agents can use a lateral inhibition strategy.

7. Conclusions

This chapter discussed the question how a group of agents could self-organize an internal agreement system that allows agents to spend less cognitive effort to find which word groups should be organized into phrases and to avoid possible semantic ambiguities. Natural agreement systems can be quite complex, involving a mixture of formal and meaningful suffixes, multiple functions for these suffixes and syncretism. Here we studied only the first two steps towards the emergence of such complex agreement systems, namely a step where agents develop purely formal stickers and then a step where they use meaningful stickers based on semantic features declared in the lexicon. We have shown language strategies that allow agents to establish such agreement systems through situated language games.

Based on the results in this paper, we can already see what are possible next steps for further experiments. Firstly, we can now investigate what happens when the sticker itself is not provided as part of the phrase building construction but becomes an independent morphological construction, which would open the door for integrating more features of the word to decide about how to express an agreement feature. Secondly we can now study the process whereby a sticker becomes conventional because the semantic features do not apply and hence the meaning of the sticker has to be stretched. Beyond these two steps, many more experiments can be envisioned before we can simulate the complete grammaticalization paths from a language without an agreement system to one where we find an agreement similar in complexity to human languages.

Acknowledgements

This research was conducted at the Vrije Universiteit Brussel, financed by a strategic basic research grant (IWT-489) from the agency for Innovation by Science and Technology (IWT), and at the Sony Computer Science Laboratory in Paris. Additional funding came from the European research project ALEAR (FP7, ICT-214856). We especially want to thank Pieter Wellens and Remi van Trijp for the many discussions on the actual implementation of the experiments reported in this paper.

References

- Beuls, Katrien (2011). Construction sets and unmarked forms: A case study for Hungarian verbal agreement. In Luc Steels (Ed.), *Design Patterns in Fluid Construction Grammar*. Amsterdam: John Benjamins.
- Beuls, Katrien, Sebastian Höfer (2011). Simulating the emergence of grammatical agreement in multi-agent language games. In AAAI (Ed.), *Proceedings of the Twenty-second International Joint Conference on Artificial Intelligence*.
- Boeckx, Cedric (2006). *Agreement systems*. Amsterdam: John Benjamins.
- Corbett, Greville (2006). *Agreement*. Cambridge Textbooks in Linguistics: Cambridge.
- Givón, Talmy (1976). Topic, pronoun and grammatical agreement. In Charles N. Li (Ed.), *Subject and topic*, 149–188. New York: Academic Press.
- Greenberg, Joseph (1978). How does language acquire gender markers? In Joseph H. Greenberg (Ed.), *Universals of Human Language*, vol. 4, 47–82. Stanford University Press.
- Heine, Bernd, Tania Kuteva (2007). *The Genesis of Grammar*. New York: Oxford University Press.
- Katamba, Francis (2003). Bantu nominal morphology. In Derek Nurse, Gérard Philippson (Eds.), *The Bantu languages*, chap. 7. London: Routledge.
- Lakoff, George (1990). *Women, Fire, and Dangerous Things: What categories reveal about the mind*. Chicago: University Of Chicago Press.
- Lehmann, Christian (1988). On the function of agreement. In Michael Barlow, Charles Albert Ferguson (Eds.), *Agreement in Natural Language: approaches, theories, descriptions*, chap. 3, 55 – 65. Stanford University: Center for the Study of Language and Information.
- Steels, Luc (2011a). A design pattern for phrasal constructions. In Luc Steels (Ed.), *Design Patterns in Fluid Construction Grammar*. Amsterdam: John Benjamins.
- Steels, Luc (Ed.) (2011b). *Design Patterns in Fluid Construction Grammar*. Amsterdam: John Benjamins.

- Steels, Luc (2012). Self-organization and selection in cultural language evolution. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins.
- Steels, Luc, Martin Loetzsch (2012). The grounded naming game. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins.
- van Trijp, Remi (2011). Feature matrices and agreement: A case study for German case. In Luc Steels (Ed.), *Design Patterns in Fluid Construction Grammar*. Amsterdam: John Benjamins.
- van Trijp, Remi (2012). The evolution of case systems for marking event structure. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins.
- Wellens, Pieter, Martin Loetzsch, Luc Steels (2008). Flexible word meaning in embodied agents. *Connection Science*, 20(2), 173–191.
- Welmers, William (1973). *African Language Structures*. Berkeley: University of California Press.

Appendix

The language systems discussed in this paper are all implemented in Fluid Construction Grammar (FCG) Steels (2011b). This appendix defines in more technical detail how the constructions involved in each experimental step have been defined.

Absence of agreement

The agents' grammar is scaffolded with a lexicon and two phrasal constructions that link multiple words into a word-group.

Lexical constructions A lexical construction is a mapping between meaning and form. The meaning is a set of predicates (although in this experiment we use only one for simplicity) and the form is a string. The lexical construction also has to define what the meaning is about. Lexical constructions are defined using templates (see (Steels, 2011b)), as in the following example:

```
(def-lex-cxn girl-cxn
  (def-lex-skeleton girl-cxn
    :meaning (== (girl ?x))
    :about (?x)
    :form "girl")
  (def-lex-cat girl-cxn
    :syn-cat (== word)))
```

This builds a construction called `girl-cxn`. It has a lexical skeleton which defines the meaning, the form, and what the meaning is about, and adds a syntactic category (which is very general, namely being a word) and a semantic-category.

Phrasal constructions The initial grammar contains two grammatical constructions. The first one, called `create-word-group` combines two words into a word-group. Which words can be combined is not constrained by their syntactic categorization (they just have to be words) nor their semantic categorization. However the two words have to be about the same referent. The construction is specified using the `def-phrasal-linking` template (discussed in (Steels, 2011a)). Note that the word-group as a whole inherits the `:about` feature from the words.

```
(def-phrasal-cxn create-word-group
```

```

(def-phrasal-skeleton create-word-group
  :phrase
  (?word-group
   :phrase-type word-group)
  :constituents
  ((?word-1
   :syn-cat (==1 word))
   (?word-2
   :syn-cat (==1 word))))
(def-phrasal-linking create-word-group
  (?word-1
   :about (?referent))
  (?word-2
   :about (?referent))
  (?word-group
   :about (?referent)))

```

The second construction, called `expand-word-group` is entirely similar, except that a word is added to an existing word-group, if they are about the same object:

```

(def-phrasal-cxn expand-word-group
  (def-expand-phrasal expand-word-group
   :existing-phrase
   (?word-group
    :syn-cat (==1 word-group))
   :constituent
   (?word
    :syn-cat (==1 word)))
  (def-phrasal-linking expand-word-group
   (?word-group
    :about (?referent))
   (?word
    :about (?referent))))

```

Formal stickers

Formal stickers are handled by a straightforward expansion of the two phrasal constructions used earlier. The suffixes are added by a new template

`def-phrasal-morph` which attaches a suffix to a stem. Here is an example for a sticker “-bla”. The sticker is stored in the unit for the word-group so that it can easily be reused when the word-group is expanded later with more words.

```
(def-phrasal-cxn create-word-group-bla
  (def-phrasal-skeleton create-word-group-bla
    :phrase
    (?word-group
      :syn-cat (==1 word-group)
      :sticker "-bla")
    :constituents
    ((?word-1
      :syn-cat (==1 word))
     (?word-2
      :syn-cat (==1 word))))
  (def-phrasal-linking create-word-group-bla
    (?word-group
      :about (?referent))
    (?word-1
      :about (?referent))
    (?word-2
      :about (?referent)))
  (def-phrasal-morph create-word-group-bla
    (?suffix-word-1
      :stem ?word-1
      :string "-bla")
    (?suffix-word-2
      :stem ?word-2
      :string "-bla")))
```

The expansion from `expand-word-group` to `expand-word-group-sticker` is straightforward because the `create-word-group` construction stores the sticker of the word-group at the level of the sticker:

```
(def-phrasal-cxn expand-word-group-sticker
  (def-expand-phrasal expand-word-group-sticker
    :existing-phrase
    (?word-group
```

```

      :syn-cat (==1 word-group)
:sticker ?sticker)
  :constituent
    (?word
      :syn-cat (==1 word)))
(def-phrasal-linking expand-word-group-sticker
  (?word-group
    :about (?referent))
  (?word
    :about (?referent)))
(def-phrasal-morph create-word-group-bla
  (?suffix-word
    :stem ?word
    :string ?sticker)))

```

Meaningful stickers

The lexical constructions used in the meaningful sticker experiment are expanded with the semantic categories that are going to play a role in the agreement system. We have seen that the lexical constructions need to specify whether the predicates involved can be used as a means to find the referent (**referring**) or only as a way to further constrain or modify

For example, for the girl construction we get the semantic category **person** rather than the generic **entity**.

```

(def-lex-cxn girl-cxn
  (def-lex-skeleton girl-cxn
    :meaning (== (girl ?x))
    :about (?x)
    :form "girl")
  (def-lex-cat girl-cxn
    :syn-cat (== word)
    :sem-cat (== (sem-function referring)
              (class person))))

```

Constructions for meaningful stickers are the same as for formal stickers, except that now the semantic criteria are tighter. For example, if the "-bla" suffix is linked to the semantic category of **person**, we get the following construction. Note that one of the words needs to have a referring function.

```

(def-phrasal-cxn create-word-group-bla
  (def-phrasal-skeleton create-word-group-bla
    :phrase
    (?word-group
      :syn-cat (==1 word-group)
      :sticker "-bla"
      :sem-cat (==1 (class person)))
    :constituents
    ((?word-1
      :syn-cat (==1 word)
      :sem-cat (==1 (class person)
                (sem-function referring)))
      (?word-2
        :syn-cat (==1 word)
        :sem-cat (==1 (class person))))))
  (def-phrasal-linking create-word-group-bla
    (?word-group
      :about (?referent))
    (?word-1
      :about (?referent))
    (?word-2
      :about (?referent)))
  (def-phrasal-morph create-word-group-bla
    (?suffix-word-1
      :stem ?word-1
      :string "-bla")
    (?suffix-word-2
      :stem ?word-2
      :string "-bla")))

```

Notice how the semantic category is also stored in the word-group so that the `expand-word-group-sticker` construction can easily find out what feature is to be used for linking further words into the word-group.

```

(def-phrasal-cxn expand-word-group-sticker
  (def-expand-phrasal expand-word-group-sticker
    :existing-phrase
    (?word-group

```

```
      :syn-cat (==1 word-group)
      :sticker ?sticker
:sem-cat (==1 (class ?sem-cat)))
  :constituent
    (?word
      :syn-cat (==1 word)
      :sem-cat (==1 (class ?sem-cat))))
(def-phrasal-linking expand-word-group-sticker
  (?word-group
    :about (?referent))
  (?word
    :about (?referent)))
(def-phrasal-morph create-word-group-bla
  (?suffix-word
    :stem ?word
    :string ?sticker))
```