**Simulating the syntax and semantics of linguistic constructions about time**

Joachim De Beule
Artificial Intelligence Lab
Vrije Universiteit Brussel
Brussels - Belgium
joachim@arti.vub.ac.be

Joachim De Beule
Artificial Intelligence Lab
Vrije Universiteit Brussel
Pleinlaan 2
1050 Brussels
Belgium

In this paper we motivate and report on the implementation of a computer experiment to investigate the syntax and semantics of linguistic constructions about time. It is argued that the way in which a domain like time is conceptualized is not universal and evolves over time. To investigate this we want to simulate a population of agents evolving their proper language and ontology of time in order to succeed in communicating temporal information. Such simulations can be done using a formalism proposed by Steels (2004). Some advances in applying the formalism to the domain of time are reported and examples of actual simulations are presented.

Keywords: Language, Evolution, Meaning, Time

## 1. Introduction

In this paper we motivate and report on the implementation of a computer experiment to investigate the syntax and semantics of linguistic constructions about time. It is argued that the way in which a conceptual domain, like time, is structured and conceptualized by humans is not fixed but differs from culture to culture and evolves over time. This is especially true for the semantics of linguistic constructions since language is a conventional system, and the meaning of for example a grammatical tense category is conventionally determined (Lapolla n.d.)

A language is not an isolated system of word to meaning mappings and grammar rules that can be used to express things about the world in some universal way. A language cannot exist outside the context of an environment inhabited and shaped by a population of language users. When a member of a language community wants to engage in a communicative interaction he has to conform to the consensus in that community, both in the way he conceptualizes the topic that he wants to express as in the way he verbalizes the conceptualization. A child hearing a new word or grammatical construction has to hypothesize about what is intended and how the word or construction should be used. She will have succeeded when she is herself able to use the construction successfully. In other words: when she finds that the other members of her language community seem to correctly understand the hypothesized meaning or communicative intent expressed by the new construction as she uses it. It is possible however that she does not use the construction as was intended, that she uses it in a novel but successful way, or even that she uses a construction that wasn't even there (grammaticalization), but still is successful in using it. At this point the new conceptualisation or rule she introduced might become part of the language and she is herself constructing the language that she is trying to learn.

Hence, to study language we need to look at it as an evolutionary and self organizing system in which language users try to figure out their environment (to *survive* in it) by learning the language but in the mean time constantly shaping and creating that environment themselves. Moreover, both the meaning and form of a language are subject to this evolution. There is no universal way to conceptualize the world for language. Rather, conceptualization schemes should be valued by their usefulness in communication. As Nathalie Gontier (this volume) explained, looking at language in this way and studying it from this perspective means implementing evolutionary epistemology in the field.

## 2. The Artificial Intelligence approach

In artificial intelligence (AI) we try to build computational or robotic systems that act and interact in some intelligent way. One way to study cultural evolution and phenomena like language is to put together a population of such systems and see what it takes for them to evolve a language. In this case it is common to call an individual system or instance thereof an *agent* and a population or group of interacting agents a *multi-agent system*. Thus an agent is a robot or program representing an individual and in a multi-agent simulation a population of such agents interacts with each other and with the environment.

To study language as an evolutionary system we simulate a population of interacting agents with the goal of letting them evolve their proper language. This way we can gain insight in both what capacities and mechanisms an individual agent needs to accomplish this goal and in the structure and evolution of language as a cultural phenomenon. It should be clear however that these two levels cannot exist separately: the evolution of an individual agent both depends on and determines the evolution of the language and vice versa.

The ultimate goal is to find out how a population of agents can establish a set of conventions about the expression of temporal information. For this to be possible, the agents have to agree upon an ontology in the domain of time. By an ontology we mean a set of categories or distinctions with which an agent can conceptualize the world.

As is explained in Steels (1998), the emergence of a *shared* set of categories requires the interaction among the members of the population and a co-evolutionary coupling between the categories that make up the ontology and the forms to express them. Thus, an agent should be capable of doing at least three things. First, he should be capable of creating new categories. Second, he should be capable of incorporating these categories into his language repertoire. Third, he should be capable of adapting his language and ontology in order to conform to the (emerging) consensus in the population.

In another paper (De Beule 2004) a mechanism is proposed to accomplish the first prerequisite: the creation of new temporal categories by playing a generalized discrimination game. In short, when an agent has to conceptualize a topic event he builds a description for it. A description could for example state that the topic is a walk event. Such a description is considered adequate if it uniquely describes the topic with respect to the other events in the context. If the agent is unable to build such a description the creation of a new category is triggered. For example if the topic event is a walk event in the past, but also another walk event is taking place in the present, a new distinction or category distinguishing present from past events is created. The description of the topic can now be extended by specifying that it is the walk event in the past. The creation of a new category and the subsequent adoption of the category into the agent's ontology (category repertoire) is called a discrimination game.

The current paper focuses on how the resulting temporal categories can be incorporated into an agent's language repertoire using a formalism proposed by Steels (2004).

The third step, the interaction between the different agents and the resulting co-evolution of meaning and form is not discussed in depth (although see section 5.2 for some first results.) Still, the underlying assumption is that every act of communication attempts to guide the interpretation process of a hearer to arrive at the topic. Consequently a new conceptual category or syntactic extension to the language should only be introduced when it will help the interpretation process. This allowed Steels to postulate a mechanism by which a speaker can decide when to extend the syntax of the emerging language: he first listens to himself and interprets his utterance as if he were a hearer. As such he can detect and solve ambiguities before actually pronouncing an utterance.

This mechanism and the language formalism have been successfully used to simulate the emergence of a case grammar (Steels 2004). In this paper we will apply it to the domain of time.

## 3. Investigating language as an evolving conventional System

### *3.1. Motivation*

One of the primal aspects of language is that it is used for communication. If the communication is to be successful the users of the language need to agree upon both what to express and how to express it. Both of these are partly predetermined by the world and the physiology of the agents. But they are also partly conventional and optimized to make them more efficient for the purpose of communication.

For example there exist only a finite set of vowel systems used in human language. As was shown by de Boer (2001), each such system is a tradeoff between its expressive power and its understandability given the constraints of human physiology, of the physics of air as transmission medium etc.

The set of categories and relations that are made explicit in a language is partly conventional as well. For example, the distribution of colors in the world and the properties of the human visual system only partly explain the way in which color is categorized for language. Each particular color categorization scheme is also the emergent result of a negotiation between the members of a language community (Belpaeme 2002).

Similar observations can be made in the domain of time. There are many commonalities between the ways in which time is conceptualized across different cultures. Most if not all cultures conceptualize time as progressively and irreversibly moving forward from the past to the future on a linear time line (Comrie 1985). Time is a deictic system so a reference is needed to specify a location in time. In probably all languages the current moment or moment of speech is used for this. Most cultures also have notions for today, yesterday, etc. (although Hopi might be an exception, see Whorf (1956) or Yee (n.d.) for a discussion on Whorf's claims about the Hopi concept of time.) Because of the nature of time and its importance to daily life it is no surprise to find these commonalities.

Still, there are indications that the categorization of time is not universal but also in part cultural or conventional. For example, if we move from daily life to a more philosophic or religious domain, some cultures clearly have concepts of time that are cyclic. As another example, the more technically advanced a culture is the more elaborate and precise the temporal categorization becomes.

In addition there is great variety in the way in which temporal information is encoded in different languages. If it is true that form and meaning co-evolve to adapt to the need for efficient communication there should also be some variety in the way time is conceptualized. In the following we will show that there *is* great diversity in expressing temporal information across different languages and even within one language.

### *3.2. Diversity in the Expression of Temporal Information*

First, not all languages express similar information in the same way. For example, almost all cultures make metaphorical associations between the future and the front of the body. This convention is used extensively while make gestures accompanying an utterance. It can also be observed in words like "be-fore" where the "fore" part means in front. There are however cultures that associate the *past* with the front of the body and accordingly express temporal information with opposite gestures and different metaphors (Nunez 1999).

Second, not all languages require the same temporal information to be made explicit. For example in Chinese there is no tense system and temporal information is specified lexically or with aspectual categories. Chinese has a relative aspectual system, e.g. in the sentence:

*kui sik joh faan siin jau*
he eat ASP dinner then leave

the aspectual marker marks completion of dinner relative to leaving. But it is not specified whether this happens in the past, present or future so there are still different interpretations possible:

(1) He ate dinner, then left.
(2) He always eats dinner before leaving.
(3) He will eat dinner, then leave.

These have to be disambiguated by the conversational context or by providing extra information (Yee n.d.)

In contrast with this, in Bantu languages it is common to have *several* tense categories for different degrees of remoteness in time, e.g. making a distinction between *a few days ago* and *more then a few days ago* (Comrie 1985).

Even among the members of the same language community there can be differences in the meaning associated with a temporal category. An example from Dutch is the meaning of the future tense form. The future tense is constructed by combing the verb *gaan* (to go) with the infinitive. In standard Dutch "*ik ga eten*" means "*I will eat*", placing the eat event in the future without any additional connotation. In many Belgian dialects however this construction means "*I will eat* here and now". This is in opposition with the construction *"ik ga gaan eten"*. This construction is not grammatical according to standard Dutch but is very widely used. The meaning of *"Ik ga gaan eten"* is more or less captured by the sentence "*I will go and eat/with the purpose of eating"*. There seems to be an additional meaning of displacement putting the eat event in the more distant future. This construction is not limited to eating but can be made with practically all verbs, i.e. *"Ik ga gaan zwemmen"* (I will go and swim). An exception might be the verb gaan (to go) itself: it is normally not acceptable to say "*ik ga gaan gaan"* although some people consider even this grammatical.

This diversity can be explained by the fact that language is a complex dynamic phenomenon and the emergent global result of local peer to peer negotiations. New members entering a language community might trigger changes in the meanings and forms of the language. For example, the origin of the English future tense form (will+infinitive) is closely related with that of the still separately existing verb *"willen"* (to want) in Dutch. Wanting something is closely related with the future and in English *"will"* became used to indicate the future tense.

Hence, one way to gain insight into language is to build systems evolving their proper language. This has been done in a variety of ways (e.g. De Jong 2000; Kirby 2001; Steels&Kaplan 2002;…). In the rest of the paper we will report on some advances in building such an experimental setup to investigate the conceptualization and expression of temporal information.

## 4. Formalism and Mechanism

As explained, the aim is to simulate a population of agents evolving a shared ontology and language. The force that drives the evolution is the *desire* of the agents to be successful in communication. While trying to accomplish this goal the agents are forced to invent new categories and extend the language. In Steels (2004) a formalism for construction grammars and accompanying learning strategies are presented that are specifically designed to do such simulations. In the following two sections we will summarize this formalism and illustrate the mechanism by which new rules are introduced by an agent. We refer to Steels (2004) for details.

### *4.1. Formalism*

The agents in our experiments need a way to represent semantic and syntactic information. This is done with feature/value structures called units. A unit also has a name. A collection of units is called a unit structure or simply structure. An example of a semantic structure is shown in figure 1[1].

unit-1

meaning : {*fall(X)*}

Figure 1: A semantic structure containing one unit named unit-1 and containing one feature named meaning. Meanings are represented in standard predicate calculus-style, i.e. the value of the meaning feature is *fall(X).*

Semantic structures specify constraints on the meaning to be expressed or interpreted (interpreting a semantic structure means finding bindings for all the variables in the meaning

[1]Predicates and variables are written in *italics*, variables also with a capital first letter. The value of a unit's meaning feature is a predicate calculus like logical expression possibly containing variables.

to things in the world.) Every semantic structure is mirrored by a syntactic structure containing units with the same names as those in the semantic structure. Syntactic structures contain constraints on the form to be expressed or interpreted

Agents also need a way to represent lexicon entries and other rules of language. This is done by specifying a transformation between unit structures. A rule normally contains a semantic pole and a syntactic pole. The application of a rule in production is done by first matching the semantic pole against a semantic structure to get a set of bindings and then (if a match is found) unifying the syntactic pole with the corresponding syntactic structure. In interpretation the matching is done against the syntactic structure and the semantic pole is

```
rule fall-past
  Unit
    meaning :+ {fall(X) & past(X)}
<-->
  Unit
    form :+ {stem("fallpast")}
```

Figure 2: A lexical rule associating the word "fallpast" with the meaning *fall(X) &past(X)*. A rule is printed as follows: first the left pole (a unit structure containing variables) followed by an arrow (<-->) followed by the right pole. The ':+' means that the feature value 'should at least contain' the specified elements.

unified with the corresponding semantic structure. An example of a lexical rule is given in figure 2. It specifies that the part of the meaning that is equal to *fall(X)} & past(X)* is expressed by the word "fallpast" and vice versa. The formalism is thus closely related with construction grammar and cognitive grammar formalisms.

*4.2. Mechanism*

Imagine an agent placed before a screen where movies of falling and rolling objects are shown. Assume he gets to see a fall event followed by a roll event followed by another fall event. The second fall event is still taking place at the moment that he wants to draw another agents attention to the past fall event (see figure 3).

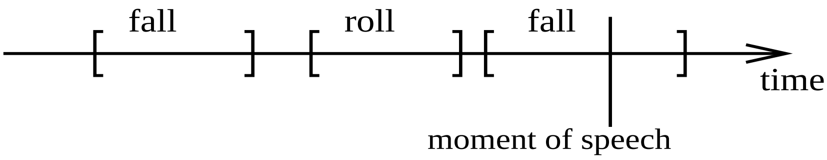


Figure 3: A fall event followed by a roll event followed by a fall event.

First he has to build a semantic structure describing the event, for example stating that it is a fall event as is shown in figure 1. However, this description is not specific enough to discriminate the topic from the other fall event in the context. The addition of a temporal category could solve this problem. Assuming that the agent's set of temporal categories is still empty the agent can decide to create a category *past*. Such a discriminating category can be found by playing a (generalized) discrimination game as is explained in De Beule

(2004). Adding the category *past* to the semantic structure in figure 1 results in the semantic structure shown in figure 4.

```
unit-1
    meaning : {fall(X) & past(X)}
```

Figure 4: The semantic structure of figure 1 extended with a temporal category *past*.

Next, the agent has to identify applicable lexical rules by matching their semantic pole against this semantic structure. When a matching rule is found, the unification of its syntactic pole with the syntactic structure will result in the addition of syntactic (lexical) information specifying the form. For clarity, in this section the semantic and syntactic structures will be shown as one structure containing both meaning and form features. It is possible that (part of) the meaning is uncovered by the lexicon. In this case new lexical rules have to be introduced. There are several possibilities. If the entire meaning is uncovered one word is associated with it. For example associating the word "fallpast" with the structure shown in figure 4 results in the rule shown in figure 2.

```
unit-1
    subunints : {fall-unit, past-unit}
fall-unit
    meaning : {fall(X)}
    form : {stem("fall")}
past-unit
    meaning : {past(X)}
    form : {stem("did")}
```

Figure 5: The result of applying lexical rules for "fall" and "past" to the semantic structure of figure 4.

If however the agent's lexicon already contains an entry for *fall* he only has to introduce a rule associating a new word with the *past* part of the meaning. New units are introduced if the meaning of a unit is broken down in the lexicon into more than one word. Thus, the application of two separate lexical entries for *fall* and *past* transform the structure of figure 4 into the one shown in figure 5. The syntactic information in this structure can be used to assemble the utterance. The result is either "fall did" or "did fall" (there is no word order defined.)

Now assume a hearer knowing the word "fall" receives this utterance. He has to figure out two things: (1) the meaning of the word "did" and (2) the fact that both words "fall" and "did" are about the same thing.

To solve problem (1) the hearer should in some way be able to deduce that the topic is the past fall event, e.g. by joint attention or by the speaker pointing to it. But even if this succeeds the hearer might still make a different conceptualization for the topic, for example stating that it is the fall event *before the roll event*, thereby mistakingly associating the meaning *before(X,Y) & roll(Y)* with the word "did". Such an error will become apparent in future interactions and cause the adaptation of the agent's ontology and/or lexicon.

To solve problem (2) some additional syntactic constraints on the utterance should be added reflecting the constraints at the semantic side. The fact that the variables in the fall-unit and the past-unit are the same (the variables to the *fall* and the *past* predicates in figure 5 are both *X*) can be considered as a constraint at the semantic level. This constraint should also be specified in some way at the syntactic level. This is comparable with the lexical rule in figure 2 which connects a constraint on the form feature of a unit to contain the word “fallpast” (syntactic level) with a constraint on the meaning feature to contain the meaning *fall(X) & past(X)* (semantic level).

To express the equality of the variables in both units another syntactic constraint is needed, for example word order. The agents could postulate a rule stating that if the word “did” directly precedes the word “fall” it can be concluded that the corresponding meanings are about the same thing. This is accomplished with the rule shown in figure 6. This rule introduces word order as can be seen by the *precedes* constraint in the syntactic pole of the rule. Hence, this is a first step toward a language with non-trivial syntax. The mechanism responsible for this step is the introduction of syntactic constraints on the form (e.g. word order) reflecting semantic constraints on the meaning (e.g. the equality of variables in different units.)

```
rule fall+past
   Parent-unit
      subunits :+ {Subunit-1, Subunit-2}
   Subunit-1
      meaning :+ {fall(X)}
   Subunit-2
      meaning :+ {past(X)}
<-->
   Parent-unit
      subunits :+ {Subunit-1, Subunit-2}
      form :+ {precedes(Subunit-1, Subunit-2)}
   Subunit-1
      form :+ {stem(“fall”)}
   Subunit-2
      form :+ {stem(“did”)}
```

Figure 6: Example grammatical rule introducing word order.

```
rule cat1+cat2
    Parent-unit
        subunits :+ {Subunit-1, Subunit-2}
    Subunit-1
        semantic-category :+ {semcat1(X)}
    Subunit-2
        semantic-category :+ {semcat2(X)}
<-->
    Parent-unit
        subunits :+ {Subunit-1, Subunit-2}
        form :+ {precedes(Subunit-1, Subunit-2)}
    Subunit-1
        syntactic-category :+ {syncat1}
    Subunit-2
        syntactic-category :+ {syncat2}
```

Figure 7: A more general version of the rule shown in figure 6 using semantic and syntactic categories.

The rule in figure 6 is not general but only applies to the specific combination of *fall* and *past*. A way to arrive at a more general rule is to introduce semantic and syntactic categories[2]. For example, if *fall(X)* would be of category *semcat1(X)*, "fall" of *syncat1*, *past(X)* of *semcat2(X)* and "did" of category *syncat2*, the rule in figure 7 would achieve the same result as the fall+past rule of figure 6. But this rule is more general because it can be applied to any combination of units having the correct semantic and syntactic categories. It is still not very clear however how to determine the members of a particular semantic or syntactic category. Some very promising results are reported in Steels (2004) where the re-use of categories is driven by analogy.

To summarize, new lexical rules are introduced by a speaker when the lexicon does not cover the entire meaning that is to be expressed. To detect opportunities to introduce new grammatical rules the speaker can interpret his utterance as if he were a hearer. Some of the constraints present in the semantic structure of the production phase will not be present in the semantic structure of the interpretation phase. For example, an equality of two variables might be missing. He can then add grammatical rules to solve this problem.

## 5. Some simulations and results

### *5.1. Simulating one agent*

In this section the details of some actual simulations are shown in which an agent acts as a speaker and looks for opportunities to add new rules. The world presented to the agent is

[2]These categories are not to be confused with ontological categories like *past*. A semantic category (like object, event, agent,...) is a categorization of a conceptual relation used to constrain the semantic pole of grammatical rules. A syntactic category (like verb, noun, nominative,...) is a categorization of a word or group of words used to constrain the syntactic pole of a grammatical rule.

shown in figure 3, with the past fall event called fall-1 and the present fall event fall-2. In other words, the context that is presented to the agent is given by the conjunction *fall(*fall-1*)* *& roll (*roll*) & fall(*fall-2*).* The agent is also aware of every event's begin and end time and of the current moment. His goal is to produce an utterance describing the past fall event.

5.1.1. Conceptualization

In all examples the agent starts with conceptualizing the topic by building the semantic structure shown in figure 1, but with the variable *X* instantiated as (replaced by) fall-1. Because this semantic structure is not specific enough to uniquely describe the topic (also the present fall event is described by the semantic structure) the agent adds a temporal category. In figure 8 three possible solutions are shown that were observed during different simulations.Three different solutions were found in different simulations as shown in figure 8.

unit-1
Referent : fall-1
Meaning :
{ *end*(fall-1)
< *end*(fall-2)
*& fall*(fall-1)
*& fall* (fall-2) }

(a)

unit-1
Referent : fall-1
Meaning :
{ beg*in*(fall-1)
< *end*(roll)
*& fall*(fall-1)
*& roll* (roll) }

(b)

unit-1
Referent : fall-1
Meaning :
{*end*(fall-1)
< *begin*(now)
*& fall* (fall-1)
*& now* (now) }

(c)

Figure 8: Three different conceptualizations observed in simulations for conceptualizing the past fall event in the context of figure 3.

The first conceptualization (figure 8a) describes the topic as the fall event which ends before the end of the fall-2 event. The second conceptualization (figure 8b) describes the topic as the fall event that begins before the end of the roll event. The third (figure 8c) describes it as the fall event that ends before the current moment (*now*). All conceptualizations uniquely describe the fall-1 event within the context of figure 3. In the following we will continue with the semantic structure given by figure 8c.

5.1.2. Application of lexical rules

After successfully conceptualizing the topic and constructing a semantic structure the agent looks for matching lexical rules. As explained in section 4.2, the particular combination of lexical rules applied also determines how the semantic (and syntactic) structure is decomposed. If part of the meaning is uncovered by the lexicon new lexical rules are invented. On the left the result is shown for an agent with an empty lexicon. The agent created a new rule (figure 9a) covering the entire meaning of the structure in figure 8c. Applying this rule to the structure results in the structure of figure 9b. On the right the result is shown for an agent who already had a rule associating the *fall*(fall-1) part of the meaning with the word "zapaxo". He created a new rule associating the remaining part with the word "fovuxi" (figure 9c). These rules transform the structure of figure 8c into the one of figure 9d.

5.1.3. Creation and application of pre-grammatical rule

We will continue the example with the unit structure shown in figure 9d. Before uttering the phrase specified by this structure the agent listens to himself to determine ambiguities. The phrase specified by the unit of figure 9d is "zapaxo fovuxi" or "fovuxi zapaxo" (there is no word order specified.)

```
Rule zivatu
    Unit
       referent : X
       meaning :+ {fall(X) & now(Y)
                       & end(X)<begin(Y) }
<-->
    Unit
       form :+ {stem("zivatu")}
```

(a)

```
rule zapaxo
    Unit
       referent : X
       meaning :+ {fall(X)}
<-->
    Unit
       form :+ {stem("zapaxo")}
rule fovuxi
    Unit
       referent : X
       meaning :+ {now(Y)
                       & end(X)<begin(Y) }
<-->
    Unit
       form :+ {stem("fovuxi")}
```

(c)

```
unit-1
   referent : fall-1
   meaning : { end(fall-1)<begin(now)
                & fall(fall-1)
                & now(now) }
   form :{ stem("zivatu") }
```

(b)

```
unit-1
     subunints : {unit-2, unit-3}
unit-2
    referent : fall-1
    meaning : {fall(fall-1)}
    form : { stem("zapaxo") }
unit-3
    referent : fall-1
    meaning : { end(fall-1)<begin(now)
                & now(now) }
    form : { stem("fovuxi") }
```

(d)

Figure 9: Two different lexical rule sets and unit structures resulting from applying the rules to the structure shown in figure 8c. In the simulation the rules are created by the agents as they need them to cover the meaning in the structure.

```
unit-4
     subunints : {unit-5, unit-6}
     form : precedes(unit-5,unit-6)
unit-5
    referent :  X-1
    meaning : {fall(X-1)}
    form : { stem("zapaxo") }
unit-6
    referent : X-2
    meaning : { end(X-2)<begin(Y)
                & now(Y) }
    form : { stem("fovuxi") }
```

Figure 10: The unit structure resulting from interpreting the utterance "zapaxo fovuxi" using the lexical rules shown in figure 9c.

Translating the utterance “zapaxo fovuxi” back into a unit structure using the lexical rules of figure 9c results in the structure shown in figure 10. As can be seen, the main difference with the structure of figure 9d is that now variables are introduced. In addition, the variable for the *fall* predicate in unit-5 is *X-1*, while the variable that should be bound to fall-1 in unit-6 is *X-2*. Interpreting this structure within the context of figure 3 means finding bindings for the variables *X-1, X*-2 and *Y* such that all meaning parts in the structure become true. The *now*(*Y*) part of the meaning feature in unit-6 specifies that the variable *Y* should be bound to the current moment. The other variables *X-1* and *X-2* should be bound to one of the events fall-1, fall-2 or roll. Since it is not specified that *X-1* and *X-2* should be equal there are several solutions: *X*-1 can be bound to both fall-1 and fall-2 and *X*-2 can be bound to both fall-1 and to roll. To solve this, the agent should add an extra constraint on the form of the utterance, for example specifying word order or introducing case. In the specific simulation being discussed the agent introduced the 5 rules shown in figure 11a to 11e. These specify that

(a) the word “zapaxo” is of syntactic category syncat-1 (rule syncat-1),
(b) the word “fovuxi” is of syntactic category syncat-2 (rule syncat-2),
(c) the meaning *fall(X)* is of semantic category semcat-1 (rule semcat-1),
(d) the meaning *end(X)<begin(Y) & now(Y)* is of semantic category semcat-2 (rule semcat-2) and finally that
(e) if these categories are combined then the word of syncat-1 gets a suffix “-lo” and the word of syncat-2 gets a suffix “-ri”.

If the agent now re-verbalizes the semantic structure of figure 8c by applying these rules to the semantic structure of figure 9d the resulting utterance would be “zapaxo-lo fovuxi-ri” or “fovuxi-ri zapaxo-lo” (there is still no word order specified.) And in interpretation the rule of figure 11d will specify that the variables *X-1* and *X-2* in figure 10 have to be equal because they will both get bound to the same variable *X* of the rule (the variables in the values of the referent features of units *Subunit-1* and *Subunit-2* are both *X*.)

### *5.2. Simulating a population of interacting agents*

In the previous section we looked at the internals of a single agent. In this section we discuss a simulation of a population of 10 interacting agents. These agents are simplified with respect to the agents discussed in the previous section and results are therefore preliminary. The agents in this section are only allowed to use holistic lexical rules, i.e. lexical rules that cover the entire conceptualization as does for example the rule in figure 9a. Consequently, the application of the lexicon does not fragment the unit structure. This means that no grammar rules like the rules of figure 11 are used as they are not needed. Still some interesting results can be shown, for example that such a population is indeed capable to converge to a shared ontology and language to communicate temporal information. It will also be shown that, in this simulation, both the conceptualization and the syntax of language are, at least partly, conventional.

#### 5.2.1. Setup of the Experiment

The experiment consists of a number of successive two-agent interactions. Each interaction, two agents are randomly selected from the population. One of them is the speaker, the other the hearer. Both agents are presented with the same context, on average consisting of five events. The speaker randomly selects one of these events as the topic. Next he tries to conceptualize this event as was the case in the previous section, thereby using a temporal category from his category repertoire. If there is no category in his category repertoire that can be used to uniquely describe the topic, the agent creates a new temporal category for it,

**rule** syncat-1

*Unit*
form :+ {stem("zapaxo")}
-->
*Unit*
syncat : syncat-1

(a)

**rule** syncat-2

*Unit*
form :+ {stem("fuvuxi")}
-->
*Unit*
syncat : syncat-2

(b)

**rule** semcat-1

*Unit*
referent : *(X)*
meaning :+ { *fall(X)* }
-->
*Unit*
semcat : semcat-1

(c)

**rule** syncat-2

*Unit*
referent : *(X)*
meaning :+ {*end(X)<begin(Y) & now(Y)*}
-->
*Unit*
semcat : semcat-2

(d)

**rule** cat-1+cat2

*Parent-unit*
subunits :+ { *Subunit-1, Subunit-2* }
*Subunit-1*
referent : (*X*)
semcat : semcat-1
*Subunit-2*
referent : (*X*)
semcat : semcat-2
<-->
*Parent-unit*
subunits :+ { *Subunit-1, Subunit-2* }
*Subunit-1*
syncat : syncat-1
form :+ {suffix("-lo")}
*Subunit-2*
syncat : syncat-2
form :+ {suffix("-ri")}

(e)

Figure 11: (pre-)grammatical rules introduced by the agent to make the variables *X*-1 and *X*-2 of the structure in figure 10 equal, this way solving the multiple interpretation problem.

thereby extending his category repertoire (see also De Beule 2004). It is also possible that several categories in the repertoire uniquely conceptualize the topic. In that case he chooses the category that was most successful in previous interactions. Next, the speaker searches his lexicon for entries that express the conceptualization. If no entry can be found, his lexicon is extended with a new entry associating the temporal category with a new word. If several entries are found the one that was most successful in previous interactions is chosen. Finally, the speaker utters the word to the hearer who has to interpret it in the context, i.e. the hearer has to determine the event that the speaker chose as topic.

To do so, the hearer first determines the lexical entries that match the uttered word. It is possible that several entries match. Only those with an associated temporal category that uniquely determines one event in the current context are considered. In other words, only the *relevant* entries are considered. Of these, the entry that has been most successful in the past is chosen to determine the topic. If this event indeed is the speaker's topic, the interaction ends successfully. If not, the interaction ends unsuccessfully. It is also possible that there was no matching lexical entry. In this case also the interaction ends unsuccessfully. In the case of an unsuccessful interaction the speaker points to the intended topic so that the agents can adjust their language repertoires (categories and lexicon.)

The precise details of the adjustment procedure are not discussed in this paper. The philosophy behind it is that an agent should try to mimic the majority of the other agents in the population. Every time an agent interacts with another agent he samples the population and gains information about the language repertoires of the other agents. This information can be used by the agent to adjust his own language repertoire. Information is gained by both the speaker and the hearer and in both successful and unsuccessful interactions. For example, note that even when an interaction ends successfully, this does not mean that both the speaker and the hearer used the same temporal category. Nor does it mean that the hearer would have used the same word to describe the topic since he could for example have conceptualized it differently. In this case, the adjustment procedure will encourage the hearer to use the speaker's word in future interactions and discourages the use of the word that he himself would have used. In addition, when a word was unsuccessful too many times or when its use was discouraged too many times, it is removed from the agent's lexicon. The same holds for temporal categories.

#### 5.2.2. Results

In figure 12 the evolution of the communicative success among the agents in the population is shown. The communicative success is the probability of having a successful interaction between two randomly picked agents when presenting them with a random context[3]. One can see that communicative success increases fast in the beginning with the increase rate slowing down as time proceeds. After 10000 interactions the communiucative success reached 99.9 percent (not shown in figure 12.)

As was already explained, a successful interaction does not per se mean that both agents have the same category repertoire or lexicon. Similarly, a high communicative success does not mean that the population has converged to a small but shared set of temporal categories or words to express them.

[3] This context should of course be similar to the contexts shown to the agents during the experiment, in the current experiment for example containing on average 5 events.

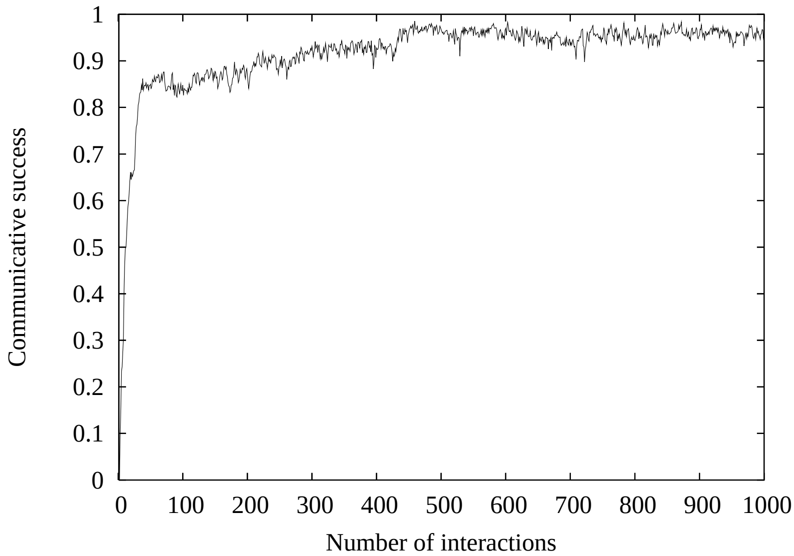


Figure 12: Evolution of the communicative success in time (number of interactions) in a population of 10 interacting agents (see text for details.)

It is also possible that although agents start to understand the words used by other agents, they would have used another word themselves. There is a high degree of synonymy of which most agents are aware (hence a high communicative success), but no consensus about which word to prefer has been reached yet. This is indeed illustrated by figure 13 which shows the average number of words associated with a single category in an agent (averaged over both agents and temporal categories).

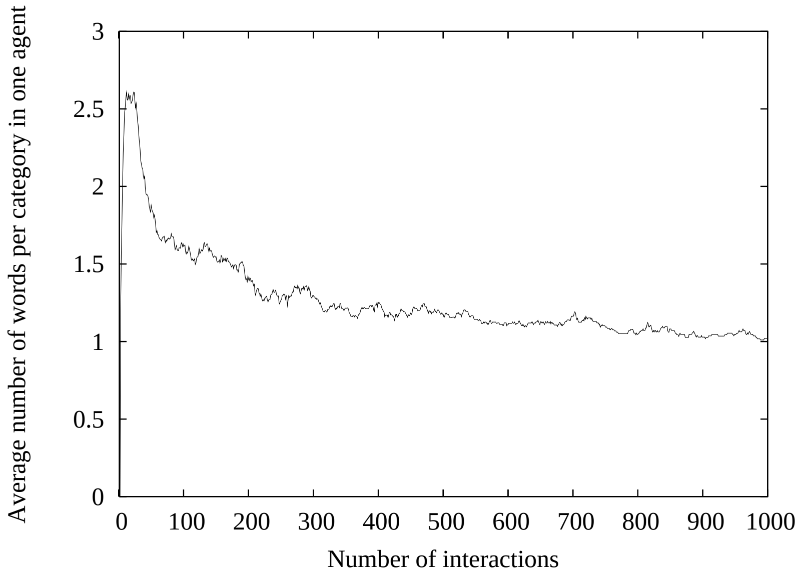


Figure 13: Evolution of the number of words associated with a temporal category averaged over all categories and all agents in the population.

As explained, the agents not only need to agree upon a lexicon, they also have to agree upon an ontology for time with which they can conceptualize the world. What is a useful set of temporal distinctions or categories to conceptualize the temporal aspects of the world is partly determined by the world itself. But there could be many useful sets of temporal categories compatible with the world. And in the case that the world or environment includes the population in which temporal information is to be expressed, usefulness is also partly conventional. Thus, we should also expect an increase in the similarity of category repertoires between different agents and in the average number of categories associated with a word as the agents are converging to a consensus. This is shown in figures 14 and 15. Note that the category similarity does not reach 100%. Apparently, to reach almost complete communicative success it is not necessary for all category repertoires of all agents to be equal. This of course depends on the structure of the world and the kind of contexts the agents are presented with.

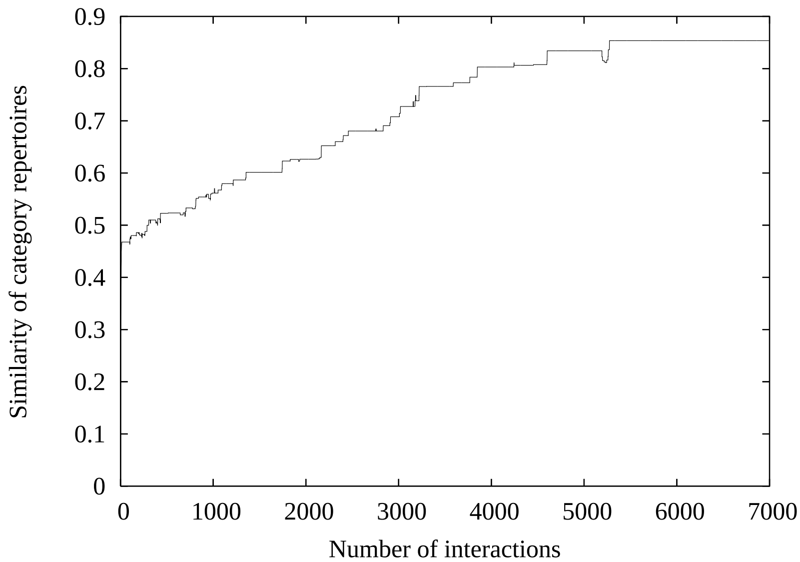


Figure 14: Evolution of the similarity between the repertoires of temporal categories of the different agents in the population, averaged over all pairs of agents. The similarity of two repertoires is 1 iff both repertoires are of equal size (same number of categories) and contain the same categories. The similarity is zero if the two repertoires have no category in common.

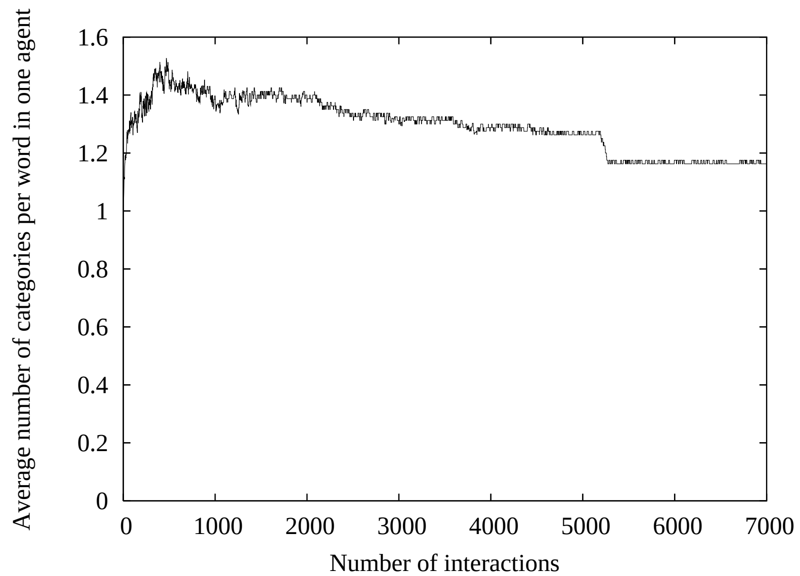


Figure 15: Evolution of the number of temporal categories associated with a word averaged over all words and all agents in the population.

Finally, figures 16 and 17 show that the average size of an agent's language repertoire (lexicon plus categories) does not simply increase (which would explain communicative success as it does in the first 50 interactions) but converges to a minimal but sufficient and (partly) conventional set of words and categories.

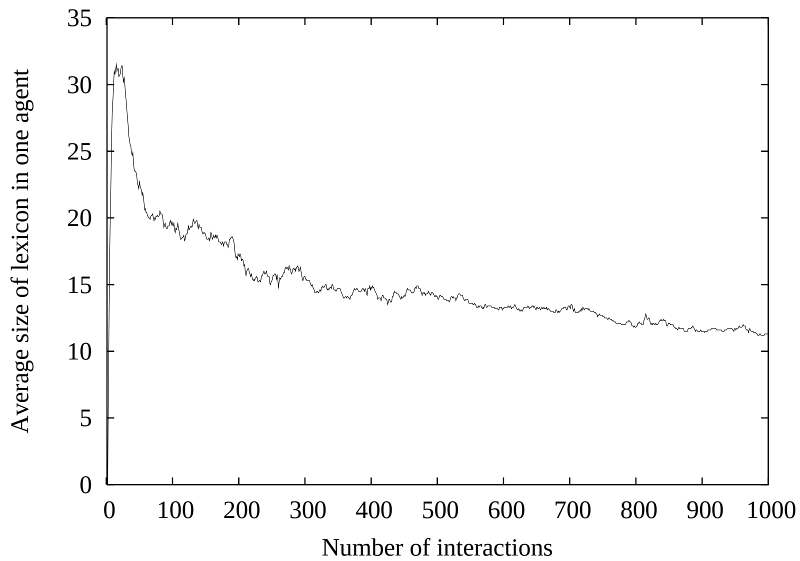


Figure 16: Evolution of the average size of an agent's lexicon (category/word associations.).

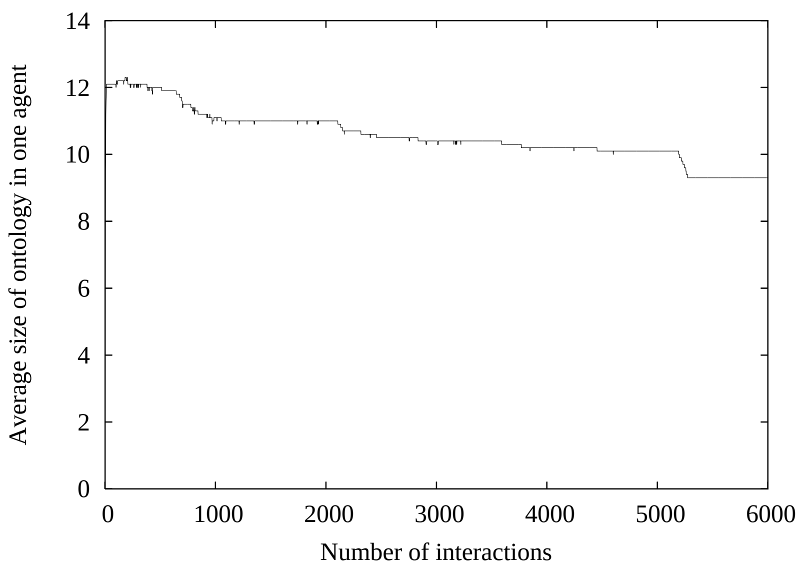


Figure 17: Evolution of the size of the temporal category repertoire of an agent averaged over all agents.

The (primarily quantitative) differences between words and categories (compare for example figures 13 and 15 or figures 16 and 17) in these experiments can be explained by the fact that the temporal category repertoire an agent adopts (by creating a new category when needed or removing one from its repertoire when unsuccessful etc.) is biased by the world (all agents experience the same world) and by the algorithm that creates new temporal categories. In contrast, the choice of  which word to use for some conceptualization purely conventional (new words are created in a random fashion.) Still, part of the resulting temporal category repertoires is clearly conventional as well since the reperoires of different agents become smaller and more similar as time passes and agents are adjusting to conform to the emerging consensus in the population.

In conclusion the experiment shows that it is possible to simulate a population of artificial agents inventing their proper language to communicate temporal information. Both the temporal category repertoire that is used to conceptualize temporal information and the lexicon that is used to express these conceptualizations are subject to evolution and (at least partly) conventional.

## 6. Conclusion

We have argued that language is a complex dynamic phenomenon and the emergent solution of a community to the problem of efficient communication. In order to be successful in communication the members of a language community need to agree on both what to express and how to express it. As such, both the syntax and semantics of a language are in part conventional. In addition, every language user is constantly shaping his own environment and creating himself the language he is learning. Investigating language from this perspective means implementing evolutionary epistemology in the field of language research.

One way to investigate these claims is to look at natural languages and try to find proof for them. Several examples illustrating the diversity in which temporal information is expressed in different languages were given. Another way to approach these claims is to build actual systems that explain the observed diversity in a computational and plausible manner. This can be done by using techniques from artificial intelligence and computer science, which also illustrates the fact that language is a complex cognitive and cultural phenomenon and that language research requires the integration of different fields as evolutionary epistemology predicts.

First steps in building such a system for the domain of time were presented. It was shown how an agent, driven by the desire to be as clear as possible, can propose both semantic and syntactic extensions to the language. New temporal categories are introduced into the agent's ontology of time and syntactic extensions are added to his language repertoire. This was illustrated with actual detailed simulations of one agent.

It was also shown that a population of simplified agents is, through repeated peer to peer negotiations, capable of evolving a shared set of temporal categories to conceptualize temporal aspects of the world and a shared associative lexicon to express these conceptualizations.

There are still many issues left unsolved. For example the emergence of a more complex language with non trivial syntax will require the introduction of more complex learning mechanisms. The emergence of tense as a grammatical system can only occur when a mechanism is found to re-use semantic and syntactic categories. It will also require a convention to *always* specify the time frame of certain semantic or syntactic components to be established.

## 7. Acknowledgements

This research is partly funded by the European Science Foundation within the Origins of Man, Language and Languages program (OMML, CRP 01- JA02: The cultural self-organisation of cognitive grammar.)

## References

Belpaeme, Tony. 2001. *Factors influencing the origins of colour categories*. PhD thesis, Vrije Universiteit Brussel, ARTI-lab.

Comrie, Bernard. 1985. *Tense*. Cambridge: Cambridge University Press.

De Beule, Joachim. 2004. Forthcoming. "Creating temporal categories for an ontology of time". To appear in *Proceedings of the 16th Belgian-Dutch Conference on Artificial Intelligence (BNAIC'04)*, Groningen.

De Boer, Bart. 2001. *The origins of vowel systems*. Oxford: Oxford University Press.

De Jong, Edwin. 2000. *Autonomous formation of concepts and communication*. Brussel: PhD thesis, Vrije Universiteit Brussel.

Kirby, Simon; and Hurford, Jim. 2001. "The emergence of linguistic structure: an overview of the iterated learning model". In: Cangelosi, Angelo; Parisi, Domenico (eds.) *Simulating the evolution of language,* 121-148. London: Springer Verlag.

LaPolla, Randy J. (n.d.) "Language as culture: the conventionalization of constraints on infernece". http://citeseer.ist.psu.edu/215204.html

Nunez, Rafael E. 1999. "Could the future taste purple?". In: Nunez, Rafael; Freeman Walter (eds.) Reclaiming Cognition: The primacy of action, intention and emotion 41-60. Thorverton: Imprint Academic.

Steels, Luc. 1998. "The origins of ontologies and communication conventions in multi-agent systems". Autonomous Agents and Multi-Agent Systems 1: 169-194.

Steels, Luc; and Kaplan, F. 2002. "AIBO's first words: The social learning of language and meaning". Evolution of Communication 4(1): 3-32.

Steels, Luc. 2004. Forthcoming. "Constructivist development of grounded construction grammars". To appear in Proceedings of the 42nd Assoc. For Comp. Linguistics Conference, Barcelona.

Whorf, Benjamin Lee. 1956. Langauge, Thought and Reality. Cambridge: The MIT Press.

Yee, Nick. (n.d.) "What Whorf really said". http://www.nickyee.com/ponder/whorf.html.

Yee, Nick. (n.d.) "Implications of language: Tone, tense and time". http://www.nickyee.com/ponder/chinese.html.