Emergent Adaptive Lexicons

Luc Steels
Artificial Intelligence Laboratory
Vrije Universiteit Brussel
Pleinlaan 2, B-1050 Brussels, Belgium
E-mail: `steels@arti.vub.ac.be`

Abstract

The paper reports experiments to test the hypothesis that language is an autonomous evolving adaptive system maintained by a group of distributed agents without central control. The experiments show how a coherent lexicon may spontaneously emerge in a group of agents engaged in language games and how a lexicon may adapt to cope with new meanings that arise or new agents that enter the group. The lexicon has several characteristics of natural language lexicons, such as polysemy, synonymy and ambiguity.

Keywords: origins of language, lexicon acquisition, self-organization.

1 Introduction

The origins and evolution of language is still clouded in mystery, despite an extensive literature within linguistics, psychology, anthropology and neurobiology (see a recent overview in [15]). The most common hypothesis being explored in American linguistics is that language is based on a species-specific innate ability (a kind of language organ) and on the refinement of innate knowledge (universal grammar) by a parameter setting process [3]. This hypothesis suggests in turn that the language faculty and universal grammar came into existence due to a series of genetic mutations each giving an adaptive advantage [8] or alternatively that there has been a single 'catastrophic' mutation giving rise to syntax and thus full language [2]. Such proposals are coherent and in principle ammenable to computational experimentation. For example, Batali [1] has investigated whether recurrent neural networks, with prior weights resulting from an evolutionary process, might explain the rapid learning and critical periods found in human language acquisition. Much work remains to be done however to make precise proposals for universal grammar and to show how a simple process of parameter setting may make it possible to acquire all the languages currently found in the world.

Current research on the origins of complexity in general [6], suggests an alternative hypothesis: Language could also be an emergent phenomenon, and this in the two senses of emergence. Language is a 'mass phenomenon' actualised by the different agents interacting with each other. In this sense, language is like a cloud of birds which attains and keeps its coherence based on individual rules enacted by each bird. No single individual has a complete view of a language nor does anyone control a language. The processes underlying language use and language formation become invisible once one focuses on a single 'idealised' speaker-hearer, just like a path of an ant society disappears or becomes incomprehensible when only a single ant is investigated. Second, language might be emergent in the sense that (1) it could spontaneously form itself once the appropriate physiological, psychological and social conditions are satisfied and (2) it could autonomously become more complex, based on the same mechanisms that cause the growth of complexity in other fields [16]: evolution, co-evolution, self-organisation and level formation.

I am exploring this hypothesis in a series of experiments on robotic and software agents that span all aspects of language: grounded meaning creation, lexicon formation, syntax, and emergent phonology. An overview of these experiments is given in [13]. This paper only focuses on lexicon formation. It describes a system that gives rise to adaptive lexicons in a group of distributed agents. The system is discussed here in an abstract fashion, but it has been coupled with meaning creation processes [12], implemented on physical robots, and tested in applications such as spatial vocabulary formation [11]. A more detailed description of the lexicon formation process is found in [10].

Other researchers have been exploring the 'emergent language' hypothesis as well. For example, McLennan [7] and Werner and Dyer [14] have conducted experiments in the origins of communication from an 'alife' perspective. These experiments show that communication arises as a side effect of cooperation if it is beneficial for communication. However these emergent communication systems do not constitute a language in the normal definition of the word. The number of agents is small and fixed. The repertoire of symbols is small (8) and fixed. None of the properties of a natural language such as hierarchical structure, synonymy, ambiguity, etc. are observed. Most importantly, a defining characteristic of natural language is missing, namely that it is open: At all times a natural language keeps growing and changing to express new and different meanings with a finite but open set of building blocks. These growth and adaptation processes must be sufficiently fast to explain for example that a typical adult has acquired a vocabulary of 100.000 words. Moreover the McLennan and Werner-Dyer experiments are based on genetic evolution as the driving force, which means that the language stays fixed within a single individual, whereas languages clearly evolve within the lifetime of individuals.

Another experiment is discussed by Yanco and Stein [18]. They describe how a communication protocol could emerge in a small group of robots using reinforcement learning. Again, the size of the group involved (2 or 3) and the size of the language (between 2 and 20 words) is small and fixed. A scale up on each of these two dimensions quickly leads to an explosion (12,105,480 iterations for a 20 element language with 3 agents, i.e. 24 hours of processing time in simulation). These combinatorial explosions have also been experienced by other researchers and indicate that the problem of emergent languages is a very difficult one.

The present paper reports a significant advance in the problem how a shared vocabulary could emerge in a group of distributed agents. The advance is partly in speed to reach coherence. But also in terms of the presuppositions that are made, and the complexity of the language generated. More concretely, the results are significant in the following way:

1. The system is open. Agents may enter or leave at all times the population (of course within certain stability bounds). New agents quickly pick up the existing lexicon.

2. New meanings may enter the pool of expressable meanings. The agents will rapidly develop new words for the new meanings (after a few hundreds of interactions).

3. Only those meanings are lexicalised that are relevant from the viewpoint of the environment and the communications that arise.

4. Occasionally incoherences arise, i.e. two agents associate a different meaning with the same word. These incoherences cannot be detected if they do not impact communicative success, but they resolve when it is required to make more fine-grained distinctions.

5. It is not assumed that meanings are uniquely identifiable from the context, as is not the case in normal lexicon acquisition either. Pointing to a brown table and saying “wa” can not only mean ‘table’ but also ‘brown’ or ‘place to sit’. This uncertainty causes the introduction of ambiguity in the language which gets resolved in further interactions.

6. Context plays a role in disambiguation of a given sentence and may provide information to progressively disambiguate a word (a phenomenon also found in natural lexicon acquisition [4]).

7. Multiple word sentences emerge in order to disambiguate single words.

The rest of the paper is in two parts. Part one presents the proposed mechanism formally and gives some examples. Part two discusses results.

2 The basic mechanism

2.1 *Features*

We assume a set of agents A where each agent $a \in A$ has a set of features $F_a = \{f_0, ..., f_n\}$. A feature f_i consists of a pair $(a\ v)$ where a is called an attribute and v a value. For example, the agents could have attributes like *weight*, *size*, and *shape*, with respective values $\{oval, round, square\}$, $\{tall, small, medium\}$, or $\{weight, heavy, light, average\}$.

A *distinctive feature set* for identifying an agent a as different from the agents in a group B is a set $D_{a,B} \subset F_a$ such that $\forall b \in B, D_{a,B} \not\subset F_b$. There can be several distinctive feature sets for the same a and B. There can also be none if there is an agent $a_1 \in B$ such that $F_a = F_{a_1}$. For example, assume the following agents and associated feature sets: a_5 : $\{(weight\ heavy)(size\ medium)(shape\ oval)\}$,
$a_4 : \{(weight\ heavy)(size\ small)$
$(shape\ oval)\}$,
$a_3 : \{(weight\ heavy)(size\ medium)(shape\ oval)\}$,
$a_2 : \{(weight\ light)$
$(size\ tall)(shape\ oval)\}$,
$a_1 : \{(weight\ light)(size\ tall)(shape\ oval)\}$.
Then $\{(size\ small)\}$ is a distinctive feature set for distinguishing a_4 from $\{a_5, a_1, a_3\}$. $\{(size\ tall)\}$ and $\{(weight\ light)\}$ are two possible distinctive feature sets for distinguishing a_2 from $\{a_3, a_5, a_4\}$. But there is no way to distinguish a_5 from $\{a_3$.

A set of features K can be used to filter a set of agents M resulting in a subset $C_{K,M} = \{a | K \subset F_a\}$. For example, given the set $M = \{a_5, a_1, a_3, a_4\}$ and $K = \{(size\ small)\}$ then $C_{K,M} = \{a_4\}$.

2.2 *The lexicon*

A word is a sequence of letters drawn from a finite shared alphabet. An utterance is a set of words. In the experiments reported here, word order does not play a role. A lexicon L is a relation between feature sets and words. A single word can have several associated feature sets and a given feature set can have several associated words. Each agent $a \in A$ is assumed to have a single lexicon L_a which is initially empty. A feature set of a word in L is denoted as $F_{w,L}$. We can then define the following functions:

- $cover(F, L)$ defines a set of utterances U such that $\forall u \in U, \{f \mid f \in F_{w,L}\ and\ w \in u\}$
- $uncover(u, L)$ defines a feature set F such that $F = \{f \mid f \in F_{w,L}\ and\ w \in u\}$.

2.3 Coherence through self-organisation

The lexicon formation process proposed in this paper, assumes that agents have the capability to create new words (by random combinations of letters from the alphabet) and associate it with a feature set. They also have the capability to form a new association between a word and a feature set. This association can be detected in conversations with another agent when both the word is known and a possible meaning (in the sense of feature set) can be derived from the context. However due to the generative capacity of each agent this adoption of words does not guarantee coherence because different subgroups of agents (in the extreme case each agent) may create their own words and associations instead.

Self-organisation in the sense of the spontaneous formation of dissipative structures by the amplification and damping of random fluctuations [9], [5] is a well known mechanism for achieving coherence without central control. It is here applied as follows: Agents randomly couple words to meanings, and engage in communication. They record the success of a particular word-meaning pair and preferably use that in future communications. This establishes a positive feedback mechanism: A word that is used a lot will have a high communicative success and will hence be used even more.

2.4 Language games

A language game involves a dialog between two agents, a speaker and a hearer, within a particular contextual setting which includes other agents. The language game succeeds when the agents manage to identify a particular object (which is in the present case also an agent), further called the topic of the dialog. There are two types of language games. In the first type the topic is first introduced using extra-linguistic means, for example by pointing. Then the speaker identifies the topic again using linguistic means. The hearer can use such a language game to learn part of the language or check whether the right meaning is associated with the right words. Alternatively a language game could contain only verbal communication which is only possible if the language is already sufficiently developed. In this paper, only language games where speaker and hearer already share the topic are considered.

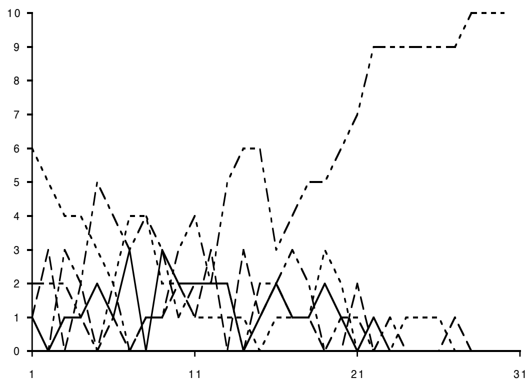

Figure 1: This figure plots the results of an experiment in lexicon formation with 10 agents, 5 possible words, and 1 meaning. It plots the communicative success of each word (y-axis) over time (x-axis). We see a search period in which different words compete to express the same meaning until one gains complete dominance.

The scenario for a typical game is as follows:

1. A speaker and hearer is randomly identified against a background of other agents.
2. The speaker selects another agent which will be the topic and points to the agent so that the hearer shares the topic.
3. Both speaker and hearer identify possible distinctive feature sets that set the topic apart from the background.
4. The speaker selects one set and translates it to words using the cover function.
5. The hearer interprets the utterance using the uncover function and compares it with his expectations.

2.5 Language formation steps

As a side effect of such a language game, various language formation steps now take place. Each of these is discussed in turn.

1. No differentiation possible(step 3 fails)

This situation occurs when the available features are not enough to make a distinction. This stimulates a feature creation process because the available distinctions are not enough to support effective communication. This situation is not further considered in this paper.

2. The speaker does not have a word.(step 4 fails)

In the second case, at least one distinctive feature set S is detected but the speaker s has no word(s) yet to express it. The language game obviously fails. However the speaker may create a new word (with a probability $p_w = 0.05$) and associate it in his lexicon with S. This is for example the case in a dialog printed out as follows:

```
Dialog 102 between a-2 and a-1 about a-4.
```

```
    Context: {a-4 a-1 a-5 a-3 a-2}
    Distinctive: ((size small))
 a-2: ((size small)) -> ? -> a-1: nil
   New word:  a-2: ((size small)) -> (p o)
```

What happened here is the following. a_2 is the speaker, a_1 is the hearer and they try to differentiate a_4 from $\{a_1, a_5, a_3, a_2\}$. There is only one distinctive feature set: $\{(\textit{size small})\}$. a_2 does not have a word for this feature set and therefore a-1 cannot extract its meaning. a_2 creates a new word. The new word '(p o)' is associated with the set $\{(\textit{size small})\}$. Newly created words are not directly used but will be in the next conversation.

3. The hearer does not have a word.

In the next case, at least one distinctive feature set S is detected and the speaker s can construct an utterance to express it, i.e. $\textit{cover}(S, L_s) = W$. However, the hearer does not know the word. Because the hearer has a hypothesis about possible feature sets that might be used, he is able to extend his lexicon to create associations between the word used and each possible feature set. Note that if there is more than one possibility, the hearer cannot disambiguate the word and retains the ambiguity in the lexicon.

4. The speaker and the hearer know the word.

4.1. The meanings are compatible with the situation.

The next situation is one where the speaker and the hearer both manage to formulate distinctive feature sets and know a word covering one or more of them. The speaker produces the word and the feature set uncovered by the hearer is one of the feature sets that are distinctive for the topic. In this case the dialog is a success and both speaker and hearer achieve communicative success. This is for example the case in the following dialog:

```
Dialog 125 between a-2 and a-1 about a-4.
    Context: {a-1 a-3 a-2 a-4 }
    Distictive: ((size small))
 a-2: ((size small)) -> ((p o))
              -> a-1: (((size small)))
```

Note that it is possible that the speaker and the hearer use different feature sets, but because the communication is a success there is no way to know this. Semantic incoherences persist until new distinctions become important and disambiguate them.

4.2. The meanings are not compatible with the situation.

The same situation as before may arise, except that the feature set uncovered by the hearer is not one of the feature sets expected to be distinctive. In this case, there is no communicative success, neither for the speaker or the hearer.

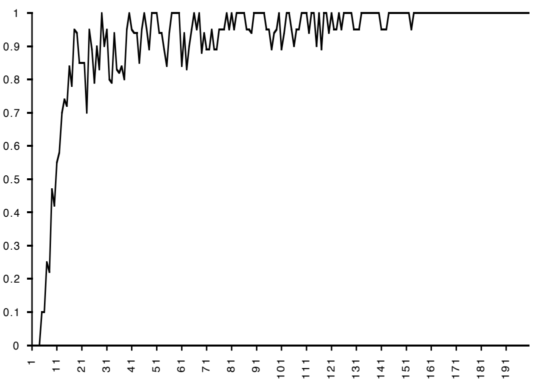

Figure 2: This figure plots the formation of a language from scratch. 4000 language games are shown, involving 5 agents and 10 meanings. The x-axis plots the number of language games (scale 1/20). The y-axis shows the average communicative success.

3 Experimental Results

3.1 *One-word utterances*

Here are the results of a typical experiment (see fig.2). It starts with the 5 agents given earlier and a dozen features (size, weight, shape) that can be used to differentiate one agent from the others. After a dozen conversations, the first word "(p o)" has been created and is used consistently by all agents to mean *(size medium)*. A dozen conversations later, the word "(f o)" appears to have propagated in the population to mean *(size small)*. Further conversations lead to further increases of the lexicon but also to the first occurrences of ambiguity: the word "(j u)" is used both for *(weight light)* and for *(size tall)*. The reason is that both can be used to distinguish a-2 from the others and an agent who does not know the word yet will retain the ambiguity as seen in the next conversation.

```
Dialog 423 between a-3 and a-1 about a-2.
    Context: {a-3 a-2 a-5}
    Distinctive:
       (((SIZE TALL)) ((WEIGHT LIGHT)))
 a-3: ((SIZE TALL)) -> ((Z U)) <- a-1: nil
   New word: a-1: ((SIZE TALL)) -> (Z U)
   New word: a-1: ((WEIGHT LIGHT)) -> (Z U)
```

When later the word "(z u)" is used to identify the same object within the same context, the communication will succeed and so no disambiguation is possible. Disambiguation only takes place when a-1 uses "(z u)" in a situation where the description *(weight light)* is not appropriate. We often see also a struggle going on between different words competing for the same meaning which will eventually resolve itself in favor of one word as illustrated in fig. 1. For example, some agents use "(j u)" in the same circumstances as others "(z u)". Some agents

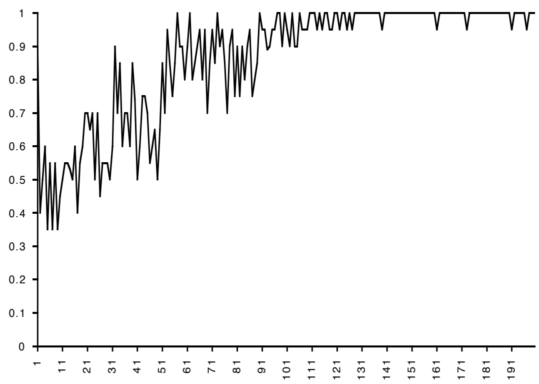

Figure 3: This figure plots the adaptation of the lexicon after six new meanings have been added. Coherence is reached again after 2000 language games.

know about both meanings but preferentially use the one that had the most success in the past. After about 4000 language games the lexicon stabilises as all distinctions that need to be made have been lexicalised.

3.2 Multiple word utterances

When a new feature is introduced, the lexicon gets extended as agents now use it as a distinctive feature (fig 3.). Thus we have introduced color with possible values *blue*, *red*, *yellow* and *white*. The words "(z i)" are adopted for *yellow*, "(d e)" for *blue*, "(k a)" for *white*. The first associations appear where a combination of features are associated with a single word. For example *(size tall) (color white)* is associated with "(z o)".

Multiple word utterances emerge naturally as the set of features expands and the lexicon adapts to cope with it. Here is an example dialog. a_1 identifies himself using the distinctive feature set $\{(size\ tall)(color\ white)\}$. Because a_1 uses "(k a)" for ($color\ white$) and "(v o)" for ($size\ tall$), a multiple word utterance is made: "((k a) (v o))". a_3 expects either $\{(size\ tall)(color\ white)\}$ or $\{(weight\ light)(color\ white)\}$ as possible distinctive feature sets. a_3 understands both words in isolation, takes the union of the feature sets and sees that it fits with the expectations. The language game therefore succeeds.

```
Dialog 2301 between a-1 and a-3 about a-1.
    Context: {a-1 a-2 a-3 a-6 }
    Distinctive:
      (((size tall) (color white))
       ((weight light) (color white)))
 a-1: ((size tall) (color white))
        -> ((k a) (v o))
         -> a-3: (((size tall)) ((color white)))
```

Unknown word-meaning pairs can also be guessed by the hearer in multiple word utterances. For example, in the following dialog a_3 knows the meaning of "(k a)" as being *(color white)*, and infers from this that "(f o)" must be either *(size medium)* or *(weight heavy)*. Both are added to his lexicon and later language games will make it clear which one was intended.

```
Dialog 2464 between a-5 and a-6 about a-3.
    Context: {a-1 a-5 a-3 a-2 a-6 }
    Distinctive:
      (((size medium) (color white))
       ((weight heavy) (color white)))
 a-5: ((size medium) (color white))
         -> ((k a) (f o)) <- a-6: nil
    New word a-6: ((size medium)) -> (f o)
    New word a-6: ((weight heavy)) -> (f o)
```

3.3 Entrance of new agents

We have already seen in the previous paragraphs that the (distributed) lexicon adapts itself when new features become available to distinguish between different objects. New words are created and sentences become more complex. Another way in which the lexicon formation process can be seen to be adaptive is because new agents may enter at any time in the population. The new agent will gradually take over words already present but is also a new source of novelty, particularly in the present experiments because a new agent means that new distinctions become relevant. Such an experiment has been carried out (fig 4). A new agent a_6 is created and added to the population. The agent has a random assignment of the various features already in use but new distinctions may now have to be lexicalised. After a few hundreds of conversations, a_6 has acquired his own version of most of words used in the group. More interestingly new distinctions become lexicalised, such as *(weight average)* which is expressed as "(m e)". As the lexicon develops, a rich tapestry of meanings and words emerges where no agent shares exactly the same language but the global system nevertheless manages to achieve quasi-total communicative success.

4 Conclusions

This paper reported computational experiments to study the spontaneous formation of a lexicon between distributed agents in a shared environment. It was shown that self-organization is an effective mechanism for achieving coherence and that many properties of natural languages, in particular synonymy, ambiguity and multiple-word sentences, occur as a side effect of the proposed lexicon formation process. Further work is going on to study the global dynamical properties of the proposed mechanism, to test out applications on robotic and software agents, and to link lexicon formation with other linguistic components.

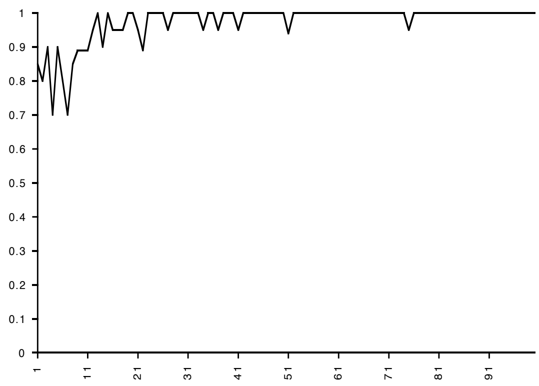

Figure 4: This figure plots 2000 language games illustrating adaptation after a new agent comes in. The x-axis plots the number of language games (scale 1/20). The y-axis shows the communicative success.

5 Acknowledgement

The research and writing of this paper has been financed by the Belgian Federal government FKFO project on emergent functionality (FKFO contract nr. G.0014.95) and the IUAP 'Construct' Project (nr. 20) of the Belgian government, with additional support from the external researcher program of the Sony Computer Science Laboratory in Tokyo.

References

[1] Batali, J. (1994) Innate Biases and Critical Periods: Combining Evolution and Learning in the Acquisition of Syntax. In: Brooks, R and P. Maes (eds.) (1994) Proceedings of Artificial Life IV. The MIT Press, Cambridge MA. p. 160-171.

[2] Bickerton, D. (1995) Language and Human Behavior. University College London Press. London.

[3] Chomsky, N. (1980) Rules and Representations. Brain and Behavior Science. Vol 3. pp. 1-15.

[4] Clark, E. (1990) The Lexicon in Acquisition.

[5] Deneubourg, J-L. (1977) Application de l'ordre par fluctuations a la description de certaines etapes de la construction du nid chez les termites. Insectes Sociaux, Tome 24, 2, p. 117-130.

[6] Langton, C. (ed.) (1995) Artificial Life. An overview. The MIT Press, Cambridge, Ma.

[7] MacLennan, B. (1991) Synthetic Ethology: An Approach to the Study of Communication. In: Langton, C., et.al. (ed.) Artificial Life II. Addison-Wesley Pub. Co. Redwood City, Ca. p. 631-658.

[8] Pinker, S. (1994) The language instinct. Penguin Books. London.

[9] Prigogine, I. and I. Stengers (1984) Order Out of Chaos. Bantam Books, New York.

[10] Steels, L. (1996a) The Spontaneous Self-organization of an Adaptive Language. In: Muggleton, S. (ed.) Machine Intelligence 15. Oxford University Press, Oxford.

[11] Steels, L. (1996b) A self-organizing spatial vocabulary. Artificial Life Journal, **3**(2).

[12] Steels, L. (1996c) Perceptually grounded meaning creation. ICMAS 1996, Kyoto.

[13] Steels, L. (1996d) Synthesising the origins of language and meaning using co-evolution and self-organisation. In: [15]

[14] Werner, G. and M. Dyer (1991) Evolution of Communication in Artificial Organisms. In: Langton, C., et.al. (ed.) Artificial Life II. Addison-Wesley Pub. Co. Redwood City, Ca. p. 659-687.

[15] Hurford, J. (ed.) (1996) Evolution of Human Language. Edinburgh Univ. Press. Edinburgh.

[16] Maynard-Smith, J. and E. Szathmary (1994) The major transitions in evolution. Oxford University Press, Oxford.

[17] Wittgenstein, L. (1974) Philosophical Investigations. Translated by G. Anscombe. Basil Blackwell, Oxford.

[18] Yanco, H. and L. Stein (1993) An Adaptive Communication Protocol for Cooperating Mobile Robots. In: Meyer, J-A, H.L. Roitblat, and S. Wilson (1993) From Animals to Animats 2. Proceedings of the Second International Conference on Simulation of Adaptive Behavior. The MIT Press, Cambridge Ma. p. 478-485