A cross-situational learning algorithm for damping homonymy in the guessing game

Joachim De Beule[1], Bart De Vylder[1] and Tony Belpaeme[2]
[1]Vrije Universiteit Brussel, Belgium
[2]University of Plymouth, United Kingdom
joachim@arti.vub.ac.be

Abstract

There is a growing body of research on multi-agent systems bootstrapping a communication system. Most studies are based on simulation, but recently there has been an increased interest in the properties and formal analysis of these systems. Although very interesting and promising results have been obtained in these studies, they always rely on major simplifications. For example, although much larger populations are considered than was the case in most earlier work, previous work assumes the possibility of *meaning transfer*. With meaning transfer, two agents always exactly know what they are talking about. This is hardly ever the case in actual communication systems, as noise corrupts the agents' perception and transfer of meaning. In this paper we first consider what happens when relaxing the meaning-transfer assumption, and propose a cross-situational learning scheme that allows a population of agents to still bootstrap a common lexicon under this condition. We empirically show the validity of the scheme and thereby improve on the results reported in (Smith, 2003) and (Vogt and Coumans, 2003) in which no satisfactory solution was found. It is not our aim to reduce the importance of previous work, instead we are excited by recent results and hope to stimulate further research by pointing towards some new challenges.

Introduction

The use of computer experiments and computational modeling has been a continuous source of interesting results in the fields of language evolution, acquisition and emergence — see for example (Steels, 1995; Cangelosi and Parisi, 2001; Kirby, 2002). A large number of these efforts is aimed at studying how autonomous agents can learn the meaning of words. This has long been the domain of semiotics and typically linguistics and philosophy have been used to set the agenda in these discussion. Recently computer modelling and artificial life have had an impact on the field, and it has added new vigour to the discussion. One of the major contributions has been to instigate a paradigm shift in language evolution studies, language development and language evolution is now seen as a complex dynamic system in which linguistic properties self-organise to the ecological, cognitive and physiological constraints (de Boer, 2000; Steels and Belpaeme, 2005).

The challenge is to let a population of language users agree on a set of words and their meanings. As in animal or human communication systems, there can be no central control. The agents have to reach consensus on what words mean through local interactions. For example, there is no central authority imposing that unsolicited email should nowadays be called "spam". The word "spam" is a consensus which arose between millions of language users. Our work focuses on the mechanisms that underlie this process and on the conditions under which we observe dynamics similar to how humans learn the meaning of words and come up with words for novel meanings.

In a related field, new technologies and web-applications that support or exploit the self-organization of communication systems are emerging and provide an additional field to which these results and techniques can be applied, e.g. (von Ahn and Dabbish, 2004). However, most of these models lack a solid theoretical foundation, and researchers are only recently taking up this challenge. For example, although many experimental studies have shown the successful emergence of a communication system in a population of multiple autonomous agents in the absence of any central control, only very recently a set of sufficient convergence criteria were formulated supporting this finding — see for example (De Vylder and Tuyls, 2005). And it was shown by (Baronchelli et al., 2005) that the emergence of a successful communication system through self organization scales to large populations, clearly a necessary property for large-scale web-applications.

However, these formal studies of word-meaning acquisition make very strong simplifications. More specifically, they consider only single-word utterances and assume a *meaning transfer*: when a speaker utters a word, the hearer knows what the intended meaning is. These simplifications greatly reduce both the complexity of the language emergence task and the difficulty of analyzing and understanding the dynamics involved.

In this paper we consider the effects of removing the meaning-transfer simplification. We propose a cross-situational learning mechanism that enables the agents to establish a successful and minimal communications system. Although we do not at this stage offer a formal proof, we support our claim with simulation results, and thereby improve on the notable work reported in (Smith, 2003) and (Vogt and Coumans, 2003). In their models the agents never seem to attain a one-to-one mapping between words and meanings, which hampered the agents' ability to communicate. The challenge in constructing artificial communication systems is to avoid homonymy (having words that have more than one meaning) and synonymy (having several words for the same meaning). If a communication contains a too high degree of homonymy of synonymy it can not be used effectively. We have solved this by introducing an improved learning algorithm that manages to avoid non-effective communication systems while at the same time being able to adapt to shifting semantics.

Learning the meaning of words using cross-situational learning

In *cross-situational learning* (Siskind, 1996) agents statistically infer the meaning of words by monitoring the cooccurence of words and sets of meanings. To illustrate this, suppose you do not know the word "banana" but you hear it often enough together with observing situations where a banana is involved. As you hear "banana" across different situations, there will come a time where you can unambiguously infer that it refers to yellow, sickle-shaped fruit.

We formalise this into a *language game* model in which N agents try to bootstrap a common lexicon for refering to (aspects of) some set O of objects. At each time-step t two agents are randomly selected from the population. One of them is the *speaker*, the other is the *hearer*. They are both presented with the current context $C_t \subset O$, which contains a random subset of objects from O. The speaker randomly selects one of the objects in C_t as the topic to which he wants to draw the hearer's attention. He therefore first calculates a semantic description for the topic object, in our case consisting of a single conceptual category like the objects ID (if it has one) or its color or size. Next, he calculates a word associated with the category and transmits it to the hearer. This may involve inventing a new word. We will refer to the set of possible conceptual categories as M.

The simplest language game is the *naming game* in which all objects are uniquely identifiable, like persons, but unlike chairs or appels. Such objects can be categorized with uni-referential categories. Because of this, each object can be uniquely named or labeled with a *proper name*. If in this setup both the speaker and the hearer know the topic (for example because the speaker points to it as is commonly done in many experiments and models), and if the probability of re-inventing an existing word is zero, then no homonymy can arise: an agent can always associate the correct meaning with an unknown and unique word. In contrast, synonymy may always arise whenever the population size is bigger than two because different agents will introduce different names for the same object. Most of the literature on language games has focused on synonymy damping mechanisms, such as *lateral inhibition* (Oliphant, 1996): a succesfully used association between a word and its meaning is strengthened and all competing synonyms are weakened. The dynamics involved can be studied for each object indepentently and are already well understood — see e.g. (De Vylder and Tuyls, 2005) and (Baronchelli et al., 2005).

In this paper we investigate what happens when homonymy is introduced, something which is much less understood and which has been largely neglected. There are several possible sources of homonymy. For example, homonymy may be introduced in the naming game when the speaker does not point to the topic and the context contains more than one object. Or homonymy is also introduced when there is a finite probability to re-invent an existing word. Homonymy also arises naturaly in a *guessing game*: when the objects involved are not uniquely identifiable but instead need to be categorized by multi-referential categories or concepts like color, size or position. These cases all share the property that a hearer is confronted with the problem of identifying one out of a set of possible meanings as the meaning of a (known or unknown) word, irrespective of whether these meanings are uni-referential categories (uniquely identifying an object) or multireferential categories (possibly applying to several objects in the context.) In the following, unless otherwise stated, we will therefore only consider the naming game case without pointing in which every uni-referential category in M identifies an object in O in a one-to-one fashion, but keep in mind that this is equivalent to a guessing game involving multi-referential categories. Likewise, the object to which a word via its associated category refers to will be called the meaning of that word.

An agent will be characterized by its production and interpretation behavior. The production behavior determines the probability with which the agent produces a word for some object. This also includes the mechanism that determines when to introduce a new word. The interpretation behavior determines how the topic is guessed given a context and a word. An interaction is considered a success when the hearer succeeds in identifying the topic chosen by the speaker. However, whether the topic is guessed correctly or not does not influence the agent: the agents do not receive feedback about the outcome of the game because this would again introduce some form of meaning transfer. The amount of corrective feedback that children receive when learning a language has been much debated on, and many mechanisms and heuristics have been proposed with which children can determine or guess the meaning of an unknown word (e.g. the whole object, exclusivity, relevance or joint attention principles.) However, our model can be applied independent of these: we assume that a hearer, using a further unspecified set of such mechanisms, is able to reduce the set of all possible word meanings M to some finite set M_t relevant in the current context. And for the same reasons that we can identify O (the set of all possible objects) with M (the set of all possible meanings), we can identify M_t with C_t. The fact that there can be more or less powerful meaning determination heuristics can be modeled by changing the size of C_t.

All agents have an equal chance of interacting and we assume a zero probability that the same word is invented twice. Still, our relaxed assumptions imply the possibility of homonymy and imply that the complete environment C_t has to be taken into account, i.e. it cannot be reduced to one single object.

To sum up, each time step t, the only information transmitted between a speaker and a hearer is the fact that the speaker produced a particular word w for one of the objects in the context C_t. This information is insufficient to determine the intended meaning of the word w (the object it refers to) and cross-situational learning is required. If the hearer is to learn a stable language, he could wait until the word w is observed again at time $t' > t$, concluding that the meaning of w should be in $C_t \cap C_{t'}$, and continuing with this strategy until its meaning is exactly known. However, as new words are introduced and as the meaning of existing words may change this strategy may fail. This because of inconsistent observations that reduce $C_t \cap C_{t'}$ to the empty set. Thus, a more intelligent cross-situational learning scheme is required. In the following section such a scheme is proposed.

Intelligent cross-situational learning

In this section we propose a learning scheme that is capable of estimating a word/meaning mapping that changes over time from incomplete information. The information consists of consecutive $\langle w, C_t \rangle$ pairs of a word w and a set of objects C_t of which one is apparently referred to by w. Previously, (Smith, 2003) proposed a Bayesian learning mechanism that estimates the probability of some meaning m occurring given the occurrence of a word w in a similar setting. Basically, it consists of storing all co-occurrences of words and meanings. However, such a mechanism on the one hand neglects the possibility that the learning target may change over time and on the other insufficiently uses available information. For example, if an agent at time t observes a word w_1 with a context $C_t = \{m_1, m_2\}$ and, at some later time step t', observes the same word with context $C_{t'} = \{m_1, m_3\}$, it seems logical to conclude that the meaning of w_1 should be m_1. However, only taking co-occurrences into account results in w_1 referring to m_1 with a probability of 0.5 and either to m_2 or m_3 with a probability of 0.25.

The mechanism we propose works independent for different words. Therefore, we will explain the learning scheme for a given word w^* given the subsequent contexts with which it appears. If w^* occurs first at time t with context C_t, the agent associates a probability distribution $s_t : M \to [0,1]$ with it, such that

$$s_t(m) = \begin{cases} \frac{1}{|C|} & \text{if } m \in C \\ 0 & \text{otherwise} \end{cases}$$

This implements the fact that all objects in C_t have an equal probability of being the meaning of the word w^*, while all other meanings have a zero probability.

Next, if w^* is observed again at time t' in a context $C_{t'}$, the probability distribution $s_{t'}$ is defined as follows. Let $\gamma = \sum_{m \in C_t} s(m)$ and $\delta = 1 - \gamma$. Furthermore, let $\gamma' = (1-\alpha)\gamma + \alpha$ and $\delta' = 1 - \gamma'$. Then

$$s_{t'}(m) = \begin{cases} \beta(\gamma) s_t(m) \frac{\gamma'}{\gamma} + (1 - \beta(\gamma)) \frac{\gamma'}{|C_t|} & \text{if } m \in C_t \\ s_t(m) \frac{\delta'}{\delta} & \text{if } m \notin C_t. \end{cases}$$

with $\beta(\gamma) = \sqrt{1 - (1-\gamma)^2}$, a definition which is motivated further on.

In words, γ is the total probability of all meanings consistent with the current observation (all objects in C_t). At time t', this probability is increased to $\gamma' \geq \gamma$ according to the parameter α. As such, α represents the strength with which the new information at time t is valued as more important than the information gathered before time t. Furthermore, this new probability γ' should be distributed among the consistent meanings such that if the new information is in agreement with the current state. If γ is close to 1, the relative probabilities among the consistent meanings should be more or less conserved (i.e. strong associations between words and meanings remain strong and weak ones remain weak). Therefore we require $\beta(1) = 1$. However, if the new information is not in agreement with the current state (γ low), then we want γ' to be more or less distributed evenly among all meanings in C_t. Therefore we also require $\beta(0) = 0$. Moreover, we want that all scores of objects in C_t increase if $\gamma < 1$, because this guarantees convergence to a unique interpretation if the contexts are random but always contain a certain object. It is easily verified that a necessary condition for this is that

$$\beta(\gamma) > \frac{\gamma}{\gamma'} = \frac{\gamma}{(1-\alpha)\gamma + \alpha}$$

for $\gamma < 1$. From this it follows that $\beta'(1) \leq \alpha$.[1] In order for the update mechanism to work for all values of α, we chose $\beta'(1) = 0$. The specific definition of β given meets all of these requirements. In any case, the total probability δ of inconsistent associations is weakened to $\delta' \leq \delta$.

We will refer to this updating mechanism which transformed s_t in $s_{t'}$ as a function u such that

$$s_{t'} = u(s_t, C_t).$$

[1] β' is the derivative

To illustrate this estimation function, assume that $\alpha = 0.3$ and that $C_t = \{m_1, m_2\}$ and $C_{t'} = \{m_1, m_3\}$. Then initially $s_t(m_1) = s_t(m_2) = 0.5$ and $s_t(m_3) = 0$. After observing $C_{t'}$ we have $s_{t'}(m_1) \simeq 0.61$, $s_{t'}(m_2) \simeq 0.35$ and $s_{t'}(m_3) \simeq 0.04$.

Agent Architecture

At time-step t, an agent can be described by a tuple $\langle W_t, \sigma_t, \phi_t \rangle$. W_t is the set of words the agent has encountered until that moment. Initially, for each agent holds $W_0 = \emptyset$.

$\sigma_t : W_t \times M \to [0,1]$ is a function which associates meanings with words, such that, $\sigma_t(w,m)$ gives the agent's estimation of the probability that word w means m. It might seem that the agents would have to know all the possible meanings M in advance, but this is not the case: as can be verified in the following, σ_t will always be zero for meanings not yet encountered.

$\phi_t : W_t \to [0,1]$ is a function which associates scores with words, which will be used to dampen synonymy.

An example of an agents' state and the way we represent it is the following:

	1.0	0.8	0.9
	w_1	w_2	w_3
m_1	0.1	0.7	0.4
m_2	0.5	0.2	0.4
m_3	0.4	0.1	0.2

Hereby we have $M = \{m_1, m_2, m_3\}$ and $W_t = \{w_1, w_2, w_3\}$. The lower-right matrix contains the values $\sigma_t(w,m)$ and the values above the words are the word scores $\phi_t(w)$.

Before explaining the production and interpretation behavior of an agent and the way he updates his internal state, we first define the interpretation function of an agent $\langle W_t, \sigma_t, \phi_t \rangle$ as the function $g : W_t \to M$ with

$$g(w) = \text{argmax}_{m \in M} \sigma(w,m).$$

If multiple meanings in M have a maximum value, one is chosen at random. Therefore $g(w)$ is possibly a stochastic value. For instance in the example agent above we have $g(w_1) = m_2$, $g(w_2) = m_1$ and w_3 interprets as m_1 or m_2, both with a probability of 1/2.

Production Suppose that the speaker at time t is given by $\langle W_t, \sigma_t, \phi_t \rangle$, the context is C and the topic he will speak about is $m^*(\in C)$. The production behavior of an agent will not depend on the context. The speaker searches for words $W' \in W_t$ which according to him interpret as m^*: $W' = \{w \in W_t | g(w) = m^*\}$.[2] As $g(w)$ can be stochastic, also can W'. If $W' = \emptyset$ (which is always the case if m^* is encountered for the first time) the speaker invents a new word $w^*(\notin W_t)$. If $W' \neq \emptyset$ then he chooses the word w^* from W' with the highest score: $w^* = \text{argmax}_{w \in W'} \phi_t(w)$. Again, if multiple words have a highest score, one is selected at random.

[2]The production resembles the *introspective oberverter* mechanism of (Smith, 2003).

Only if the speaker invented a new word he updates his internal state. Obviously, $W_{t+1} = W_t \cup \{w^*\}$. The word-meaning scores of known words stays the same, but for the new word we have

$$\sigma_{t+1}(w^*, m) = \begin{cases} 1 & \text{if } m = m^* \\ 0 & \text{otherwise} \end{cases}$$

Finally the new word gets score 1: $\phi_{t+1}(w^*) = 1$ (scores for other words are left unchanged.)

To illustrate the production consider the example agent inroduced before. If this agent is a speaker and he has to verbalize m_3 he will invent a new word, say w_4. The new agent's state then becomes (changed values in bold)

	1.0	0.8	0.9	**1.0**
	w_1	w_2	w_3	$\mathbf{w_4}$
m_1	0.1	0.7	0.4	**0**
m_2	0.5	0.2	0.4	**0**
m_3	0.4	0.1	0.2	**1**

If he has to verbalize m_2, he will look for words which interpret as m_2, hence there is 1/2 chance that he will use w_1 and 1/2 chance that he will have to choose between w_1 and w_3 according to the score function ϕ. As $\phi(w_1) > \phi(w_3)$ he will choose w_1 in this case. As mentioned before, in this case the state of the agent does not change.

To summarize, the speaker has produced a word w^* for meaning m^* in context C, thereby possibly changing its internal state. As will be described next however, the major state change occurs at the hearer side.

Interpretation Suppose that the hearer at time t is given by $\langle W_t, \sigma_t, \phi_t \rangle$, the context is C and the word received is w^*. If this word is unknown to him ($w^* \notin W_t$) then we obviously have $W_{t+1} = W_t \cup \{w^*\}$ and the word-meaning association scores for w^* are initialized as follows:

$$\sigma_{t+1}(w^*, m) = \begin{cases} \frac{1}{|C|} & \text{if } m \in C \\ 0 & \text{otherwise} \end{cases}$$

If the word is known to the agent ($w^* \in W_t$), the association scores involving w^* are altered according to the updating function u defined before:

$$\sigma_{t+1}(w^*, \cdot) = u(\sigma_t(w^*, \cdot), C).$$

We now describe the updating of the word-scores, using the auxiliary definition $\phi_t(w^*) = 1$ if $w^* \notin W_t$. First, the interpretation m' of w^* is determined as $m' = g(w^*)$ (with g using σ_{t+1}). Next, the set of synonyms S for w^* is determined as those words in $W_{t+1} \setminus \{w^*\}$ which also have interpretation m' (according to g). Finally, the score of w^* is increased: $\phi_{t+1}(w^*) = (1-\theta)\phi_t(w^*) + \theta$, and the scores of the synonyms are 'lateraly inhibited': $\phi_{t+1}(w) = (1-\theta)\phi_t(w)$

for $w \in S$. The other scores remain the same. In the following examples θ was set to 0.3.

To illustrate the interpretation and updating, consider the example agent introduced before. If he is a hearer and would hear the word w_4 with context $C_{t'} = \{m_1, m_3\}$ then w_4 is added with new entries for σ. In addition, synonyms of w_4 are inhibited. With 1/2 chance $g(w_4) = m_3$ and there are no synonyms. With 1/2 chance $g(w_4) = m_1$ in which case w_2 is a synonym and with 1/2 chance also w_3. Suppose both w_2 and w_3 are synonyms then the agent's state becomes

	1.0	**0.56**	**0.63**	**1.0**
	w_1	w_2	w_3	$\mathbf{w_4}$
m_1	0.1	0.7	0.4	**0.5**
m_2	0.5	0.2	0.4	**0**
m_3	0.4	0.1	0.2	**0.5**

If he would hear the word w_3 with context $C_{t'} = \{m_2, m_3\}$, the associated scores will be updated as follows. The total probability of the consistent meanings in $C_{t'}$ is $\gamma = 0.4 + 0.2 = 0.6$. With $\alpha = 0.3$ this will be transformed to $\gamma' = 0.72$, giving rise to $\beta \simeq 0.92$. With regard to the word scores we have that $\phi(w_3)$ will increase according to the lateral inhibition parameter θ, and since now w_3 and w_1 are synonyms $\phi(w_1)$ will be inhibited. The new state of the agent is given by

	0.7	0.8	**0.93**
	w_1	w_2	w_3
m_1	0.1	0.7	**0.28**
m_2	0.5	0.2	**0.47**
m_3	0.4	0.1	**0.25**

Experiments

Measures

In order to get insight in the way the population of agents comes to agree on an emerging language, we define some measures on the population's state. This population state can conceptually be summarized in a semiotic graph. Such a semiotic graph is a bipartite, directed graph in which nodes represent meanings and words and in which edges go from meaning nodes to word nodes or vice versa. An edge going from a meaning to a word node represents a possible production and an edge going from a word to a meaning node represents a possible interpretation (without context.) Each edge is weighted with the weights representing the probability to observe the associated production when a randomly chosen speaker produces the meaning represented by the starting node of the edge. Or, in the case of interpreting a word: each edge has a weight representing the probability to observe to associated interpretation when a randomly chosen hearer interprets the word represented by the starting node of the edge. The sum of the weights of the outgoing edges of a node is at most one, but may be lower. This because there is the possibility that an agent has not (yet) encountered a certain object or word, in which case no production or interpretation is done (only when building the semiotic graph, not during a game).

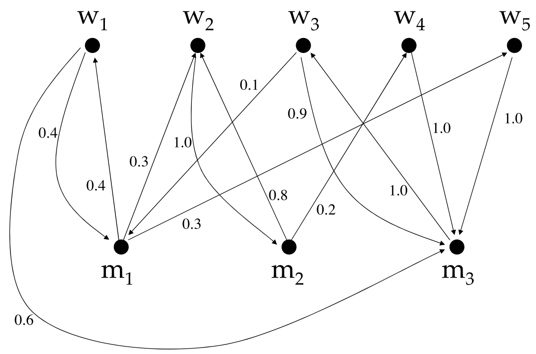

Figure 1: An example of a semiotic graph representing the state of the entire population at some point in time. The nodes marked m_i represent a meaning, those marked w_i represent words. The weights on the edges refer to probabilies averaged over the entire population (see text).

As a first measure we define the **return** of a graph as the probability that, starting from a random meaning node, one returns to that node after taking two steps (thus first going to a word node and then back to a meaning node), with the probability of taking an edge equal to the edge's weight. For example, in figure 1, the chance for returning to m_1 starting from it is $0.4 \times 0.4 = 0.16$, for m_2 it is $0.8 \times 1.0 = 0.8$ and for $m_3 = 1.0 \times 0.9 = 0.9$. Thus the return of this graph is the average 0.62. If a hearer would not take into account the context to guess the meaning of a given word, then the return equals the communicative success.

We will also define two other measures for the amount of synonymy and homonymy present in the graph. For this the notion of the **effective out-degree** of a node is needed, which is related to the number of edges starting in the node, but takes into account the weights of these edges. We assume that the weights of the outgoing edges are normalized such that their sum equals 1. If there are k edges each having an equal probability $1/k$, then the effective out-degree equals k. If, however, the k edges have differing probabilities, then the effective out-degree should be lower. Moreover, if one of the edges has a probability close to one, the effective out-degree should also be close to one (but still slightly higher). Therefore we define the effective out-degree of a node as the number of edges with equal probability needed for a hypothetical node to have the same associated Shannon information as the original node. More precisely, if a node has k outgoing edges with weights $x_i, 1 \le i \le k$, its Shannon information is given by

$$\sum_{i=1}^{k} -x_i \log(x_i).$$

A node with k' outgoing edges with equal probability $1/k'$ thus has a Shannon information of $\log(k')$. By definition this information should be equal to the information associated with the original node, from which it follows that

$$k' = \exp(\sum_{i=1}^{k} -x_i \log(x_i)).$$

The effective out-degree k' is not necessarily integer. For example, the effective out-degree of m_1 in figure 1 is 2.97 and of w_3 it is 1.38.

We now define the **synonymy** present in a semiotic graph as the average effective out-degree of the meaning nodes. The synonymy present in the graph in figure 1 is thus 1.87.

The **homonymy** present in a graph is defined as a weighted average of the effective out-degrees of the word nodes, where each node is weighed with its probability of being the result of a production. This probability is proportional to the sum of the node's incoming edges. The homonymy of the example graph is 1.26.

When no synonymy is present in the population, the synonymy measure will be 1. The same goes for the homonymy measure.

A final measure is the **number of words** present in the graph. This measure gives an indication of the parsimony of the language. Ideally, the number of words should equal the number of meanings.

Results

We will now present some results of a simulation involving $N = 20$ agents evolving a lexicon to communicate about 100 objects ($|O| = |M| = 100$). Each interaction the speaker and the hearer are presented with a context containing 5 objects ($\forall t \geq 0 : |C_t| = 5$.) The evolution over time of the return, the number of words used by the agents in the population, the homonymy and the synonymy are presented in figure 2.

As can be seen, the agents eventually reach a coherent successful language without synonyms or homonyms.

The return, which is related to the communicative success (and necessarily equivalent with it when is 1), starts at a small (chance level) value. As new words are introduced and as their meanings starts to settle, the return gradually increases until finally it becomes maximum at around $t = 105000$.

The maximum number of words present in the population is reached approximately at $t_{\max} \simeq 10000$ and is equal to 1194, which is of the order of $N\,|O|\,/2 = 1000$ as would be expected. The cross-situational learning mechanism as proposed in the previous sections allows the agents to eliminate homonyms, which partly explains why the number of words decreases steadily after $t_{\max}$.

In contrast to earlier findings in (Vogt and Coumans, 2003), our agents do reach a complete coherent language, where coherence is defined as the chance of two random speakers producing the same word for the same meaning. The main difference between Vogt and Coumans' agents and the ones defined in this paper is the use of a synonymy-damping mechanism which explains further why the number of words used by the agents eventually drops to the number of objects $|O|$.

Conclusion

In a nutshell we have identified three problems that arise when relaxing the meaning transfer assumption in a naming game. First, since during a single game a hearer can no longer determine the meaning of a word, a cross-situational learning mechanism is required. We have proposed a mechanism that is capable of estimating a word/meaning mapping that changes over time. Second, the conditions that determine when a speaker needs to invent a new word need to be extended beyond the obvious case of uncovered meaning: also when he does not know a word which he himself would *interpret* as the target meaning should he invent a new one. Third, for a coherent language to emerge some synonymy-damping mechanism is needed implementing a kind of lateral inhibition between competing synonyms.

For our study we have identified word-meanings and objects in the world. However, a more realistic setup should take perception and categorization into account. In this case, every object in the world is, through perception, mapped onto a set of agent-internal categories, for example related to the object's features, such as color or size. However the problem of determining the meaning of an utterance —as identified by (Quine, 1960) in his famous 'gavagai' thought experiment or by (Siskind, 1996) in his account of the lexical acquisition problem— persists and the mechanism we propose can still be applied.

Finally, our model is also independent of corrective feedback received by children when learning the meaning of words, or of principles of relevance, exclusivity, joint attention, whole object and the like. Although such principles do help to narrow down the set of possible meanings of an unknown word, the actual meaning still needs to be chosen from the remaining set. As we have shown, our learning model can be used for this, even if these mechanisms sometimes produce incorrect results or, equivalently, if the meaning of words changes over time.

References

Baronchelli, A., Felici, M., Caglioti, E., Loreto, V., and Steels, L. (2005). Sharp transition towards shared vocabularies in multi-agent systems. arXiv:physics/0509075 v1.

Cangelosi, A. and Parisi, D., editors (2001). *Simulating the Evolution of Language*. Springer Verlag, London.

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

De Vylder, B. and Tuyls, K. (2005). Towards a common lexicon in the naming game: The dynamics of synonymy reduction. In *Workshop on Semiotic Dynamics of Language Games*, Bagno Vignoni, Siena, Italy.

Kirby, S. (2002). Natural language from artificial life. *Artificial Life*, 8(2):185–215.

Oliphant, M. (1996). The dilemma of Saussurean communication. *BioSystems*, 37(1-2):31–38.

Quine, W. (1960). *Word and Object*. The MIT Press, Cambridge, MA.

Siskind, J. M. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition*, 61(1-2):39–91.

Smith, A. D. M. (2003). Intelligent meaning creation in a clumpy world helps communication. *Artificial Life*, 9(2):175–190.

Steels, L. (1995). A self-organizing spatial vocabulary. *Artificial Life*, 2(3):319–332.

Steels, L. and Belpaeme, T. (2005). Coordinating perceptually grounded categories through language. A case study for colour. *Behavioral and Brain Sciences*, 24(8):469–529.

Vogt, P. and Coumans, H. (2003). Investigating social interaction strategies for bootstrapping lexicon development. *Journal for Artificial Societies and Social Simulation*, 6(1).

von Ahn, L. and Dabbish, L. (2004). Labelling images with a computer game. In *ACM Conference on Human Factors in Computing Systems, CHI 2004*, pages 319–326.

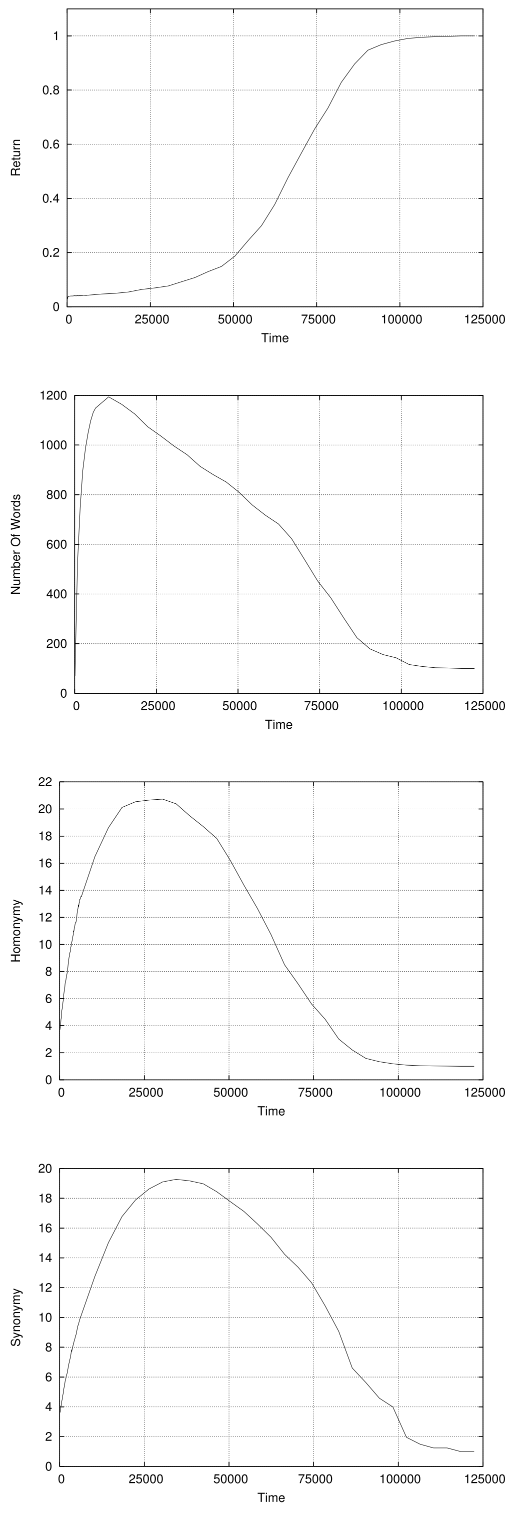

Figure 2: Evolution over time of the return, the number words used by the agents in the population, the homonymy and the synonymy as defined in the text. The graphs were obtained for $N = 20$ agents, $|O| = |M| = 100$ objects and context sizes $|C_t| = 5, t \geq 0$. The forgetting parameter α was set to 0.2 and the synonymy inhibition parameter θ was 0.3.