Grounding Language About Actions: Mobile Robots Playing Follow Me Games

Paul Vogt
Vrije Universiteit Brussel
Artificial Intelligence Laboratory
Pleinlaan 2
1050 Brussels
Belgium
paul@arti.vub.ac.be

Abstract

This paper presents a new experiment that has been carried out in the context of the research on the origins of language that is going on at the Free University of Brussels. Two mobile robots ground time series of motor commands into categories. The categorization is based on the theory of phase space reconstruction and learning by generation and selection. Besides the categories, the robots try to develop a lexicon. Lexicon formation occurs within the concept of language games in which agents can communicate about a topic and try to develop a coherent communication system. Again this is learned by generation and selection. The paper concludes that although the categorization works fine, the success of the lexicon formation is still unsatisfying.

1. Introduction

How does a communication system develop? This is a question that interests researchers already for many centuries. Recently the question can be investigated empirically by using robots. This paper presents one of these recent investigations. Before the above question is discussed, it should be clear what communication system the robots try to develop and how. The communication system they might construct could be tactile (as in (Barnes et al., 1997)), but a language-like lexicon may be more useful if the robots need to work with generic tasks. Such communication systems have already been developed (Billard and Hayes, 1997, Yanco and Stein, 1993). But in the work of Billard and Hayes a part of the lexicon has been preprogrammed. The work of Yanco and Stein heavily depends on a human instructor. From a scientific point of view it would be more interesting if the robots could ground the language system themselves.

This paper reports a new experiment in which two robots try to develop a symbolic ontology about actions they can perform in a coordinated task. Additionally, they try to construct a language-like communication system that can help the robots in performing this task. The paper reports preliminary results of this experiment.

In the past few years, research at the Artificial Intelligence Laboratory of the Free University in Brussels focuses on the origins and evolution of language and meaning, for an overview see (Steels, 1997). The approach is that both language and meaning are viewed as complex dynamical adaptive systems as proposed by Luc Steels (1996a, 1996b). Language, for instance, is complex because there is a huge expressibility in the system. Furthermore, there is no language user that has complete competence of or control over the language. Language is a dynamic phenomenon because the language itself, its users and environment change over time. The fact that language users change the language according to its changing needs makes the language adaptive. From the hypothesis that language is a complex dynamical system it follows that the language could arise through self-organization, see e.g. (Prigogine and Strengers, 1984). The models that are used are similar to those used by (Werner and Dyer, 1991, McLennan, 1991, Kirby and Hurford, 1997) and (Oliphant, 1996). An overview of research in this field can be found in (Hurford et al., 1998).

In previous experiments a lexicon has been grounded by mobile robots on perceptual sensory information (Steels and Vogt, 1997, Vogt, 1998a). In these experiments it has been observed that the feedback systems used were not sufficient for learning a completely coherent lexicon, so that the performance in the communication stayed poor. Furthermore, other work (De Jong, 1997) suggested that it might be useful to ground a lexicon within a coordination task. In this paper work is presented on a preliminary study in which following a robot is the coordination task. While following each other, the robots try to develop a lexicon about the actions the robots perform through the motor commands the robots give to their motors.

When physical robots have to develop a lexicon, they have to ground symbolic meaning from non-symbolic sensorimotor information. This problem has been addressed as the *symbol grounding problem* by (Harnad, 1990). Many attempts to solve this problem have been made on real robots. The problem is especially hard for categorizing time series of sensorimotor information, relating to this problem it is interesting to note the work of (Tani and Nolfi, 1998, Billard and Hayes, 1997) and (Rosenstein and Cohen, 1998). Tani and Nolfi use a mixture of Recurrent Neural Networks to categorize the sensorimotor flow in a navigation learning task. Billard and Hayes use a dynamic reccurent associative memory architecture (DRAMA) to categorize the sensorimotor flow of the robots. Rosenstein and Cohen use the method of delays for the categorization. The symbol grounding problem has also been tackled on real robots at the AI Lab in Brussels (Steels and Vogt, 1997, De Jong and Vogt, 1998) and (Belpaeme et al., 1998). The work presented here combines the work of (Billard and Hayes, 1997, Rosenstein and Cohen, 1998) and (Steels and Vogt, 1997).

The remainder of the paper is organized as follows. In the next section the experimental setup is described. The categorization is explained in section 3. In section 4 the lexicon formation is explained. Section five presents the experimental results. The results are followed by a discussion in section 6.

2. Follow me games

The experiment described here is based on three previous experiments. (1) The experimental setup is inspired by the work of (Billard and Hayes, 1997), where one student robot learned a lexicon through imitating a teacher robot. In the experiment of Billard and Hayes, the imitation task of the student was to follow the teacher. (2) (Rosenstein and Cohen, 1998) describe an experiment where robots grounded concepts (or *categories* or *meaning*, as will be referred to here) from time series. (3) The creation and learning part of the categorization and the lexicon formation is based on the work by (Steels, 1996a, Steels, 1996b) and (Steels and Vogt, 1997).

In the experiment two LEGO robots are used (see figure 1). Both robots are equipped with a bright halogen light on the back, two light sensors on the front, two front and two back bumpers, an infrared module, a radio link, two motors each connected to a wheel and a sensorimotor board (SMBII). The light and the light sensors are used so that one robot can follow the other using phototaxis. The bumpers and infrared module are used for obstacle avoidance. The radio link is used for both linguistic and non-linguistic communication. The motors drive the wheels and therefore the motion of the robot. The robots are completely autonomous, though for practical reasons and data logging some processing is done off-board.

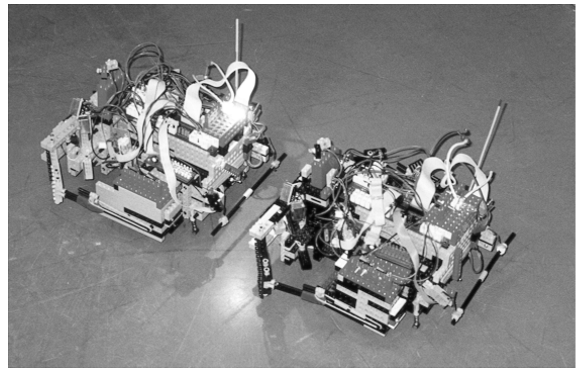

Figure 1: The two robots that are used in the experiments. In a follow me game the leading robot that carries a bright light is the speaker. The second robot (the hearer) can follow the speaker using phototaxis.

The SMBII processes the sensory input to the actuator output at a rate of 40 Hz. Every process cycle, further called a *time step* is thus equal to $\frac{1}{40}$ second. The SMBII is programmed in the Process Description Language PDL (Steels, 1994). PDL is a behavior-based programming language for robots and is embedded in C. One can define processes of behaviors like *phototaxis*, *forward-movement* etc.. These processes are processed by the system as if they were parallel processes, resulting in emergent behavior.

The task of the robots is to develop categories and a lexicon, so that they can communicate actions like going left or going backward. They do so by playing a series of so-called *follow me games.* The experiment is divided in two parts: a development part and a testing part. This section describes how the follow me games are organized.

In the work of (Billard and Hayes, 1997) the student learns through imitating the teacher robot, who has its lexicon preprogrammed. The task of the student is to learn the lexicon from the teacher by following this robot. In the follow me games both robots take turns in playing the 'teacher' and 'student' role. Furthermore, no part of the lexicon has been preprogrammed; the lexicon, like the robots' ontology of categories is completely developed from scratch. Another difference with the work of Billard and Hayes is the architecture in which the learning of associations take place. They use a neural network architecture DRAMA. In this experiment the categorization uses prototypical exemplars as representation and for the lexicon formation *symbolic* networks of associations are constructed. The learning of both the ontology and lexicon is based on generation, selection

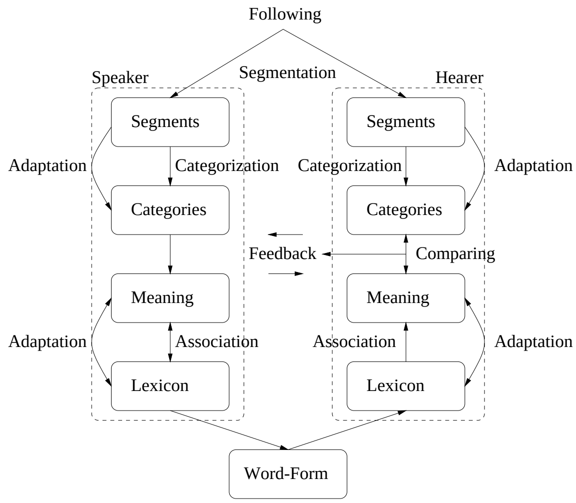

Figure 2: A schematic outline of a language game. See the text for further explanation.

and self-organization.

In a follow me game, the speaker takes the role of the teacher robot and the hearer takes the role of the student. The robots try to communicate every action they perform. Such a communication act is called a *language game* (Wittgenstein, 1958) as in previous experiments (see e.g. (Steels, 1996a) and (Steels and Vogt, 1997)). A schematic overview of a language game is given in figure 2. Note that several language games could be played in one follow me game.

In a language game two robots try to communicate about something (in these experiments about an action). When the language game is started the speaker first segments the time series of the motor commands (or time series for short) during a significant change. Secondly the speaker tries to categorize this segment[1]. Then it tries to name the category. The speaker looks in its lexicon for *word-forms* that are associated with a matching category. Naming yields one word-form, which is transmitted via the radio link to the hearer.

The hearer, while following the speaker, also segments and categorizes its actions and tries to understand the speaker. The hearer is following the speaker using phototaxis. Every time the hearer receives a word-form transmitted by the speaker, it will segment and categorize a part of the time series starting one second before the expression arrived until a few seconds after the expression arexemplarrived. This time span is taken for two reasons. First, the speaker starts to categorize and name the segment one second after the change started (see next section), so the hearer may have already started the relevant action. Secondly, the hearer might be a bit far behind the speaker, so it may still have to start the relevant action. The importance of a time span has already been observed by (Billard and Hayes, 1999). When the hearer categorized the segments it observed during this time span, it will try to understand the received word-form. It does so by looking in its lexicon for matching word-forms. When a matching word-form is found, its associated meaning is matched with one of the categorized meanings. If there is a match and the hearer is still following the speaker, the language game is successful. If one of the robots is incapable of either categorizing or naming, then it adapts its ontology or lexicon in order to be successful in the future.

A language game is in principle successful when both robots associate the communicated word-form to an action (or meaning) that is similar for both robots. As the hearer is following the speaker by phototaxis, the hearer wiggles behind the speaker. So the hearer might be going to the left when the speaker is going to the right. But one of the actions the hearer must always perform in order to follow the speaker is going to the right. By introducing sensory channels that work as a low pass filter, the unwanted fluctuations may be filtered. The resulting time series will be taken into consideration as well. The categorization and naming will be explained in more detail hereafter.

3. Categorization of time series

The categorization of time series is modeled with so-called *identification games*. In an identification game, which is played individually, a robot's task is to match a segment with a previously categorized prototype that the robot has stored in its memory. The identification game is successful when a matching category can be found. When the identification game is a failure, the robot may construct a new prototype where the current segment is used as an .

The robots segment the time series when a significant change is detected. This happens when the derivative of the series changes in sign or if it changes significantly because the time series might be subject to noise. The segmentation stops when the derivative approaches zero or when it changes sign. The segment is the time series that is in between the start and end point.

To filter out the fluctuations caused by phototaxis, new sensorimotor channels are introduced. These channels keep track of the *moving average* of the motor commands with different time windows T and thus work as low-pass filters. The application of these filters yields to a certain extend the net movement of the wiggling robots.

The segmentation of these sensory channels is the

[1]Note that it would be more realistic to categorize sensorimotor couplings than solely motor commands. However, the speaker and hearer use different sensors (infrared vs. white light) to which they respond in their actions. And since the robots take turns in being the speaker and hearer, it seemed more convenient to categorize motor commands only.

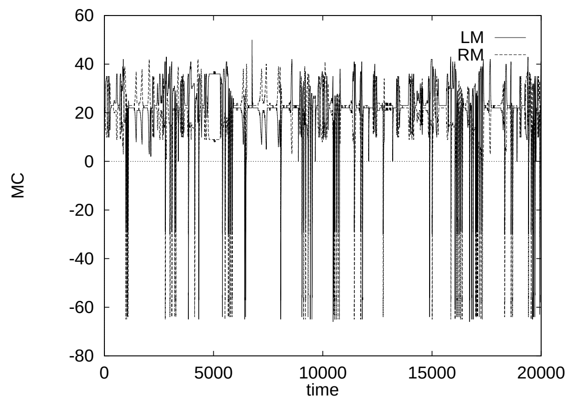

(a)

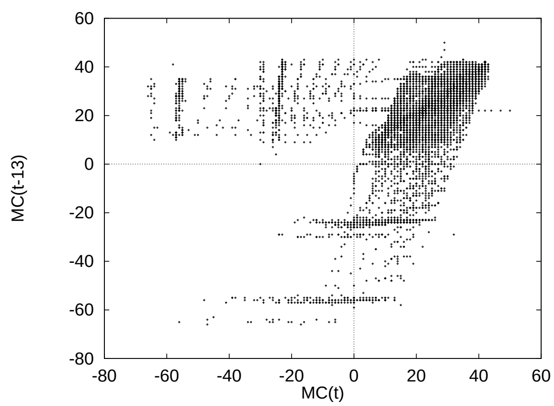

(b)

Figure 3: Plot (a) shows a time series of the motor commands for 20,000 time steps (approximately 500 seconds). The phase space of the two motor values with a delay of 13 time steps (approximately 0.33 seconds) is shown in plot (b). This phase plot should actually be 4 dimensional, but for illustrative purposes only the left motor is plotted. The data is used to model the parameters m and ν of the delay vectors.

same as defined above. The filters may cause segments over a different period than the segments of the original sensory data. In the experiment, only the hearer will consider the channels with the moving average, since it is this robot that is wiggling. Because these sensory channels work as a low-pass filter a nice side effect may emerge: the hearer might learn to follow the speaker more smoothly (i.e. without wiggling) when applying a prototype of these sensory channels to control its motors.

Obviously, the behaviors of the robots are very dynamic. As a consequence, the time series of the motor commands are nonlinear. The phase space plot of the motor commands in figure 3 shows that the trajectory of the time series typically goes to certain points in the phase space. These points in the phase space constitute the attractor of the system. It is beyond the scope of this paper to prove that this system is mathematically speaking an attractor, so it is assumed that it is. If the system has deterministic properties, the theory of *phase space reconstruction* can be applied. Since the robots are controlled by given behavior processes in a more or less stable environment, their future motion is in principle determined from the system's initial state. According to the dynamical systems theory such a system may be called deterministic. The theory of state space reconstruction solves the problem of how time series can be reconstructed by state vectors that uniquely correspond to the time series. This problem is technically solved by the *method of delays*.

The categorization of the segments is analogous to the method used by (Rosenstein and Cohen, 1998). Rosenstein and Cohen introduced the method of delays for categorization. With the method of delays a deterministic nonlinear time series can be described by a series of delay vectors of dimension m. According to Takens' Theorem (Takens, 1981) a time series can be mapped uniquely to a set of delay vectors $x_n = (s_{n-(m-1)\nu}, s_{n-(m-2)\nu}, \ldots, s_{n-\nu}, s_n)$ if the embedding dimension $m > 2D$. In this equation s_i is the original time series, ν is the time lag and D is the dimension of the system's attractor. The set of delay vectors correspond to the phase space of the time series, see figure 3. The choice of parameters m and ν are important, see e.g. (Rosenstein et al., 1994) and (Kantz and Schreiber, 1997) for a discussion. Here they are calculated for every two-dimensional sensory channel separately (original series and moving averages). It is interesting to note that the principal component analysis also uses the properties proven by Takens.

The identification game can now be defined as follows.

1. Let $P = \{P_i\}$ be a set of prototypical categories, where $P_i = (p_1^i, \ldots, p_{M_i}^i)$ is a series of motor commands, p_j^i are motor commands and M_i the length of the category i.

2. Let
$$\mathbf{X}_n^{P_i} = (p^i_{n-(m-1)\nu}, \ldots, p^i_{n-\nu}, p^i_n)$$
with $n \geq m\nu$ be the delay vectors of P_i. Delay vector $X_n^{P_i}$ is an m dimensional vector of motor commands $p^i_j \in P_i$, with $m\nu \leq M_i$. Two adjacent elements of $X_n^{P_i}$ are points of P_i that have a delay with each other of time-lag ν.

3. Suppose that the robot has generated a segment $S = (s_1, \ldots, s_N)$ with delay vectors
$$\mathbf{X}_n^S = (s_{n-(m-1)\nu}, \ldots, s_{n-\nu}, s_n)$$
and s_j motor commands.

4. Segment S can now be categorized with those categories P_i for which
$$\| \mathbf{X}_n^S - \mathbf{X}_n^{P_i} \| \leq \varepsilon \text{ for } m\nu \leq n \leq \min(N, M)$$
where ε is sufficiently small.

Note that the categories may be of different lengths due to the dynamic segmentation procedure. Consequently, segments should be categorized with prototypes of different length. Because the time span of a delay vector is at least $m\nu$, the categorization cannot start before the segment has this length.

In the method described above the calculations are performed on every delay vector $\mathbf{X}_n^S$ of a segment after $n \geq m\nu$. However a computationally more efficient system can be developed. Given the properties of Takens' Theorem, the system can also compare only a few delay vectors evenly spread over the coinciding domain of S and P_i. So, instead of calculating the difference of every delay vector of S and P_i, only five delay vectors are compared that are evenly distributed along the set of delay vectors.

The categorization of the segments may result in several categories or in none, in which case the identification game fails. If the game ends in one or more categories, it is successful. The category that has the most successful association with a word-form in the robot's lexicon will be selected during the naming phase as explained in the next section. For every category a score μ between 0 and 1 is kept, and each time a category is used in the naming phase the score is increased: $\mu := \lambda \cdot \mu + (1 - \lambda)$, where learning rate $\lambda = 0.9$. When the category is identified but not used in the naming phase, the score is decreased: $\mu := \lambda \cdot \mu$. When the identification game fails, the robot may generate a new category where the segment S is taken as an exemplar for the new prototype. If a category does not get successfully associated with a word-form, it will be forgotten. Once in a while unassociated categories will be pruned. The most successful categories will be used over and over again. The described procedure results in the self-organization of a set of categories that relates to those movements that the robots usually make. Rare movements will normally not be used successfully and therefore they will be forgotten.

4. Lexicon formation

The lexicon is developed during the naming phase of a language game. This phase is based on the naming game model as introduced in (Steels and Kaplan, 1998). The naming game model consists of four steps: production, understanding, feedback and adaptation. The lexicon is a set of word-meaning associations. A word(-form) is an arbitrary string of characters and a meaning is a category. Each association has a score σ_{wm}. The steps can be described as follows.

1. *Production.* If the speaker has categorized a segment, it will look for relating word-meaning (WM) associations in its lexicon. The association for which $\sigma_{\text{wm}} + \alpha \cdot \mu_m$ is highest will be selected and transmitted to the hearer[2]. α is a weighting parameter, set to 0.01 in the experiment.

2. *Understanding.* If the hearer receives the word-form it will look in its lexicon for associations with this word-form. The associated meanings are matched with the categories that are related to the segments within the relevant time frame. The association for which $\sigma_{\text{wm}} + \alpha \cdot \mu_m$ is highest will be selected. The related segment will be considered to be the action that the speaker performed.

3. *Feedback.* The language game is considered to be *successful* when the hearer understands and is still behind the speaker. This is the case when the intensity of the hearer's light sensor I is greater than or equal to a certain threshold Θ. If $I < \Theta$ and the hearer did find a matching word-form there is a *mismatch* in meaning between the two robots. If the hearer could not understand at all, the language game is a *failure*. The outcome of the language game is transmitted to the speaker via the radio link as part of non-linguistic communication.

4. *Adaptation.* The final step is the learning phase. A language game can have four different outcomes:

   (a) The speaker could **not produce**. The speaker did not have an association with one of the categories. In this case, the speaker may create a new word-form to associate with one of the categories.

   (b) The hearer could **not understand**. The hearer did not have a WM association with the received word-form that corresponds to one of the

[2] $\alpha \cdot \mu$ is introduced to bootstrap the lexicon and to decrease the influence of the success of the category. In (Vogt, 1998a), α was set to 1. If α is set to 0, the lexicon does not get off the ground.

categories that were valid in the relevant time frame. If the hearer is still behind the speaker ($I \geq \Theta$), then the hearer associates the word-form with the observed categories if there are not too many.

(c) There was a **mismatch in meaning**. The hearer found a matching word-form, but it was not behind the speaker anymore. The association score is now adapted as follows: $\sigma := \lambda \cdot \sigma$, where $\lambda = 0.9$.

(d) The language game was a **success**. The winning association score is adapted as $\sigma := \lambda \cdot \sigma + (1 - \lambda)$. The other associations with the used word-form are lateral inhibited with $\sigma := \lambda \cdot \sigma$.

Once in a while unsuccessful entries in the lexicon will be pruned if $\sigma = 0$ and if the entry already exists for a while. Again the learning method uses generation and selection of WM associations and coherence results from self-organizing properties of the system. Successful associations are remembered whereas the unsuccessful ones are forgotten. The system is open in that new actions can be classified and named and the ones that are not used for a long period may be forgotten.

5. The experiment

The experiment consists of two parts: a development and a testing part. Section 5.1 will discuss the development part of the experiment. Section 5.2 will discuss the testing phase, in which the resulting lexicon is tested on its instructive power.

Before the results are discussed, some details of the implementation are given. The segmentation of the sensory channels start when a significant change is observed. This change is not measured over the period $[t_s, t_s + 1]$, but over $[t_s, t_s + 9]$ because most changes occur very smoothly, so that in 1 time step usually only very small changes can be detected which might be caused by noise. The starting point of a segment is taken at time t_s. Similar arguments are valid for the end of a segment. The segmentation stops when a change in sign, or no large changes are detected over $[t_e - 9, t_e]$. The end of the segment is taken at t_e.

An exception rule is introduced when the robot hits an obstacle. If a robot bumps into a wall, the motor commands are set to a large negative value immediately, causing the sign of the derivative to become negative. Directly after that the motor values increase, changing the sign of the derivative again. This all happens in a time period smaller than $m\nu$, so the segment cannot be categorized. Therefore it is preferable to classify the complete action in which the second change of sign is ignored.

If the length of a segment either exceeds 40 time steps or if it ended and the length was 9 time steps larger than the time lag ν of the corresponding sensory channel, the segment is considered to be valid, because shorter segments cannot be categorized. At the moment that a segment becomes valid, it is categorized. Five delay vectors of the segment are compared with the corresponding delay vectors of the available categories. If the segment is sufficiently close to a category, this category is selected for consideration in the naming phase. Every time the speaker successfully categorizes a segment it will start the production.

Prior to the experiment, a data set of 20,000 motor commands has been recorded for modeling the parameters for the delay vectors (figure 3). First, the box counting dimension D_2 of the attractor is calculated. This is done for the original series, as well as for moving averaged series with periods T of respectively $10, 20, 30$ and 40. This resulted in $D_2 = 0.49$ for the original series and respectively $D_2 = 0.55, 0.59, 0.63$ and 0.67 for the other series. Therefore the embedding dimension was set to $m = 2$ for all sensory channels. The time lag is calculated using the autocorrelation method and yielded $\nu = 13$ for the original series and respectively $\nu = 16, 20, 26$ and 34 for the other sensory channels. How these parameters can be calculated are described in (Kantz and Schreiber, 1997) or (Rosenstein et al., 1994). The method for calculating D_2 is taken from (Takens, personal communication). It is beyond the scope of this paper to present these methods here.

5.1 The development phase

In the development phase the task is only to acquire categories and a lexicon about the categories while the hearer is following the speaker through phototaxis. In order to speed up the experiment (recording 500 language games takes about one day), a data set of 245 follow me games consisting of approximately 500 language games was recorded and replayed over and over off-board.

The measures that are used to investigate the success of the system are the *identification success* and *communication success*. The identification success (IS) measures the average success of the categorization as a moving average over the past 50 language games (LGs). Note that a robot may play more identification games than language games, because the hearer may consider several segments within a time frame of one language game. The communicative success (CS) is the average success of the language games over the past 50 games.

If we look at figure 4 we see that the identification success increases very fast to a value around 90 % where it remains. Categories that emerge look like the ones as in figure 5. Every robot effectively categorized around 40 categories. If a closer look is taken at the categories, six main groups can be found in relation to their corresponding action: (*going-backward*), (*going-forward*), (*going-*

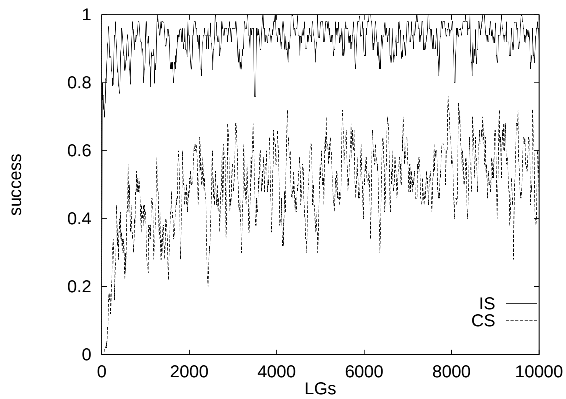

Figure 4: This graph shows the success of the categorization (IS) and communication (CS) in the development phase.

left), (*going-right*), (*going-left going-right*) and (*going-right going-left*). Some of the categories have a combination of actions like (*going-left going-right*), but these categories may also additionally be grouped in their resultant action like (*going-forward*) or simply (*going-right*). The robots obviously do not know these labels.

Figure 4 also shows the communicative success of the development phase. The CS increases rapidly towards a value around 40 % in the first 1,000 LGs. After that, the CS increases slowly to approximately 55 % after 10,000 LGs. Note that it is still unknown whether or not the robots will be successful in following each other by using language only. The language games were successful if the hearer had an appropriate word-meaning association and when the hearer was still following the speaker, but the hearer was following the speaker by *phototaxis* not using language. Failure occurred when the hearer could not understand or when it failed to follow the speaker.

5.2 The testing phase

The goal of the testing phase is to find out how well the robot could perform in a following task, where the robots only use the lexicon. Since the communicative success of the language games in the development phase is only around 55 %, it is difficult to let the robots follow each other on the lexicon alone. Another possibility was to let the hearer follow the speaker on phototaxis. When the hearer could understand a word-form, it could execute the associated category. Success could then be measured as before, when the hearer is still behind the speaker. But the time span of a category is not very large and although the hearer might still be behind the speaker, it is not certain that it performed the right action. So another test has been developed.

In the testing phase, the robots started with a part of the lexicon (i.e. categories and word-meaning associations) that they developed in the development phase. The lexical entries that the robots used were those that

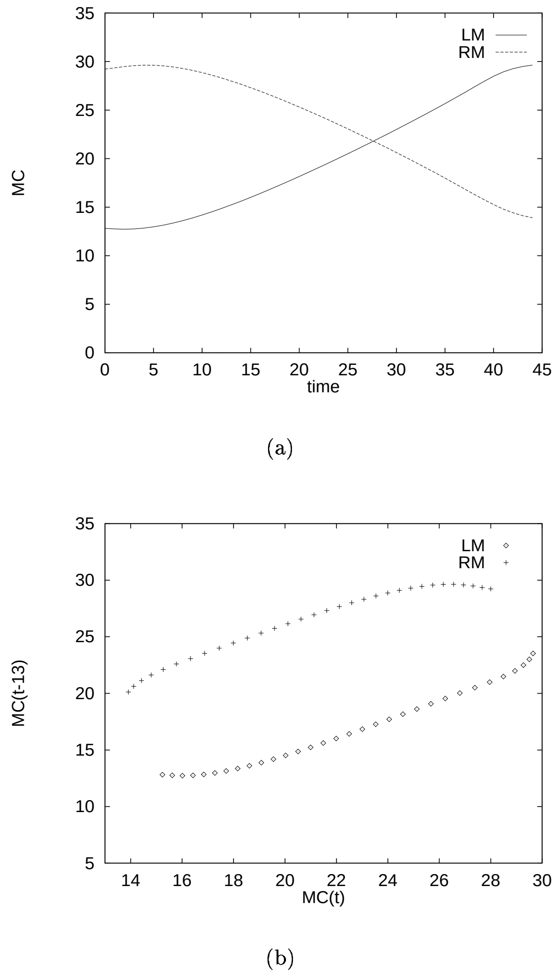

Figure 5: (a) The time series of a category that corresponds to an action that can be labeled as (*going-left going-right*). The duration of this category is approximately 1.125 seconds. Plot (b) shows the phase space of this category.

have nonnegative association scores and were available for both robots. The robots played a series of 1,000 language games. The language games in the testing phase are slightly different from the ones in the development phase. Prior to the testing phase, each category is classified by hand by the experimenter and labeled in their corresponding action. Some actions can be classified in two ways (e.g. (*going-forward*) and (*going-right*)), because both actions might suffice to follow each other.

In the testing phase both robots changed the role of speaker and hearer after every language game. The speaker selected randomly a category and produced a word-form, which was transmitted to the hearer. The hearer tried to understand the word-form yielding an associated category. Production and understanding was as described in section 4. After the naming part both robots executed the categories. The language game was success-

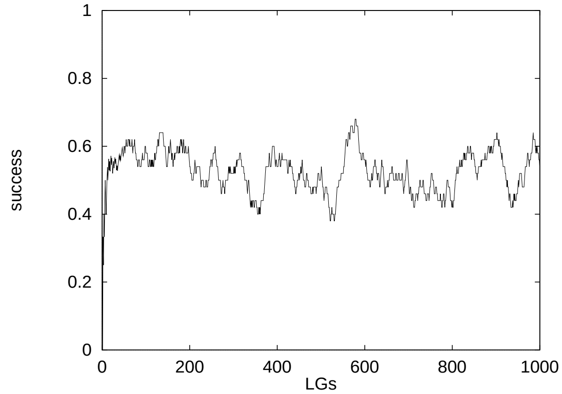

Figure 6: This graph shows the communicative success in the testing phase.

ful if both robots had a similar label on the category that was used. The lexicon was adapted only by adjusting the scores; no new categories of word-meaning associations were generated. So there were two possible outcomes of the games: a complete success and mismatch in meaning. The scores were adjusted as described in section 4. Note that this method reduces to reinforcement learning.

In figure 6 the communicative success of the testing phase is shown. The CS varies slightly above 50 %. Although the scores are adjusted according to the outcome of the game, the system does not seem to be learning. This is not surprising, since the word-forms are not associated with all categories. And through competition between ambiguous associations the lexicon keeps fluctuating between the boundaries of its capabilities. Furthermore, since the robots do not acquire new elements, nor do they have for every word-form a coherent meaning, the robots cannot converge to a coherent lexicon.

6. Discussion

This paper presented a preliminary research on how robots could ground a lexicon for actions in a coordination task. The experiment consisted of a series of follow me games in which the robots played a series of language games while one robot was following the other by phototaxis. The robots were able to generate and select categories and lexical associations. Categorization was implemented as identification games that made use of the theory on phase space reconstruction like in (Rosenstein and Cohen, 1998). The lexicon was developed using a naming game model as proposed by (Steels, 1996a).

The categorization of segments from nonlinear time series works well. The robots can rapidly categorize the time series to a high degree as was the case in (Vogt, 1998a). The method of Rosenstein and Cohen seems to be working well. Still further investigation is needed to see how reliable the results are and, if necessary, how the categorization can be optimized. Categorization might be optimized by modeling the delay method more carefully using other methods for choosing the embedding dimension and time lag as for instance is discussed in (Rosenstein et al., 1994).

In the current experiments only motor signals are categorized. It is, however, more plausible to categorize *sensorimotor* signals (signals that combine both sensory and motor information). This has not yet been done because the speaker uses a different sensorimotor coupling to control its actions than the hearer. It is wishful that the robots use the same word-form when doing a certain action independent on their role in the game (i.e. whether they are speaker or hearer). The problem of the different sensorimotor couplings might be overcome to abstract away from the raw sensory data. This abstraction can be fed into a general (behavior-based) control structure, allowing categorization on the abstracted sensory data.

The success of the naming stays well behind the categorization. Like in (Vogt, 1998a) this will probably have to do with the noise in the feedback system and the poor ability to coordinate their actions. In the development phase, success of the language games is dependent on whether or not the hearer is still following the speaker. The following occurs through phototaxis, which is a reactive and imprecise behavior. So it is very difficult for the hearer to decide which action caused the right behavior, thus making it difficult to decide which category should be associated with a particular word-form.

The difference in success between categorization and naming is remarkable, because similar results have been obtained in (Vogt, 1998a) that concerned perceptual grounding. At the beginning of the paper it has been mentioned that qualitatively good coordination and feedback is thought to improve the robots' performance in language formation. The improvement of language performance has not been observed in this experiment. However, it must be made clear that the stability of the robots used in the experiments is poor, so that they are not capable of coordinating their actions properly nor providing proper feedback. Therefore similar experiments are planned in the future using more reliable robots.

Such experiments could combine the development- and testing phase. Thus investigating the performance of lexicon formation in an experiment where the hearer tests a certain word-meaning association that is already successful during the development phase. This would be more realistic since we humans continuously test newly learned behaviors and adjust our capabilities according to the outcome of such a test. Simulations of such an experiment on *situation concepts* have already been carried out in the context of the evolution of communication

(De Jong, 1999).

Acknowledgments

The author would like to thank Professor Luc Steels for the opportunity to work in his fine laboratory and for his supervision on this research. The members of the AI Lab in Brussels are thanked for the technical support and useful suggestions. Special thanks are for Professor Floris Takens (University of Groningen) for our discussion concerning the method of delays. The author is sponsored by an assistantship from the Vrije Universiteit Brussel.

References

Barnes, D. P., Aylett, R. S., Coddington, A. M., and Ghanea-Hercock, R. A. (1997). A hybrid approach to supervising multiple co-operant autonomous mobile robots. In *Proceedings of the International Conference on Advanced Robotics*.

Belpaeme, T., Steels, L., and van Looveren, J. (1998). The construction and acquisition of visual categories. In Birk, A. and Demiris, J., editors, *Learning Robots, Proceedings of the EWLR-6, Lecture Notes on Artificial Intelligence 1545*. Springer.

Billard, A. and Hayes, G. (1997). Robot's first steps, robot's first words ... In Sorace and Heycock, editors, *Proceedings of the GALA '97 Conference on Language Acquisition - Edingburgh*, Human Communication Research Centre. University of Edinburgh.

Billard, A. and Hayes, G. (1999). Drama, a connectionist architecture for control and learning in autonomous robots. *Adaptive Behaviour*, 7(1).

De Jong, E. D. (1997). Coordination developed by learning from evaluations. In *Notes of the VIM'97 Workshop on Collaboration between human and artificial societies.*, Universita' di Salerno.

De Jong, E. D. (1999). Autonomous concept formation. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence IJCAI'99*.

De Jong, E. D. and Vogt, P. (1998). How should a robot discriminate between objects? In Pfeifer, R., Blumberg, B., Meyer, J.-A., and Wilson, S., editors, *From animals to animats 5, Proceedings of the fifth internation conference on simulation of adaptive behavior*, Cambridge, Ma. MIT Press.

Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42:335–346.

Hurford, J., Knight, C., and Studdert-Kennedy, M. (1998). *Evolution of Human Language*. Edingburgh University Press, Edingburgh.

Kantz, H. and Schreiber, T. (1997). *Nonlinear time series analysis. Cambridge Nonlinear Science Series 7*. Cambridge University Press, Cambridge.

Kirby, S. and Hurford, J. (1997). Learning, culture and evolution in the origin of linguistic constraints. In Husbands, C. and Harvey, I., editors, *Proceedings of the Fourth European Conference on Artificial Life*, Cambridge Ma. and London. MIT Press.

McLennan, B. (1991). Synthetic ethology: an approach to the study of communication. In Langton, C., Taylor, C., and Farmer, J., editors, *Artificial Life II, Vol. X of SFI Studies in the Sciences of Complexiy*, Redwood City, Ca. Addison-Wesley Pub. Co.

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

Prigogine, I. and Strengers, I. (1984). *Order out of Chaos*. Bantam Books, New York.

Rosenstein, M. and Cohen, P. R. (1998). Concepts from time series. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, Menlo Park Ca. AAAI Press.

Rosenstein, M., Colins, J., and Luca, C. D. (1994). Reconstruction expansion as a geometry-based framework for choosing proper delay times. *Physica D*, 73:82-98.

Steels, L. (1994). Building agents with autonomous behavior systems. In Steels, L. and Brooks, R., editors, *The 'artificial life' route to 'artificial intelligence'. Building situated embodied agents*, New Haven. Lawrence Erlbaum Associates.

Steels, L. (1996a). Emergent adaptive lexicons. In Maes, P., editor, *From Animals to Animats 4: Proceedings of the Fourth International Conference On Simulating Adaptive Behavior*, Cambridge Ma. The MIT Press.

Steels, L. (1996b). Perceptually grounded meaning creation. In Tokoro, M., editor, *Proceedings of the International Conference on Multi-Agent Systems*, Menlo Park Ca. AAAI Press.

Steels, L. (1997). The synthetic modeling of language origins. *Evolution of Communication*, 1(1):1–34.

Steels, L. and Kaplan, F. (1998). Stochasticity as a source of innovation in language games. In *Proceedings of Alive VI*.

Steels, L. and Vogt, P. (1997). Grounding adaptive language games in robotic agents. In Husbands, C. and Harvey, I., editors, *Proceedings of the Fourth European Conference on Artificial Life*, Cambridge Ma. and London. MIT Press.

Takens, F. (1981). Detecting strange attractors in turbulence. *Lecture Notes in Mathematics*, 898:366-381.

Tani, J. and Nolfi, S. (1998). Learning to perceive the world as articulated: An approach for hierarchical learning in sensory-motor systems. In Pfeifer, R., Blumberg, B., Meyer, J.-A., and Wilson, S., editors, *From animals to animats 5, Proceedings of the fifth internation conference on simulation of adaptive behavior*, Cambridge, Ma. MIT Press.

Vogt, P. (1998a). The evolution of a lexicon and meaning in robotic agents through self-organization. In La Poutré, H. and van den Herik, J., editors, *Proceedings of the Netherlands/Belgium Conference on Artificial Intelligence*, Amsterdam. CWI Amsterdam.

Vogt, P. (1998b). Perceptual grounding in robots. In Birk, A. and Demiris, J., editors, *Learning Robots, Proceedings of the EWLR-6, Lecture Notes on Artificial Intelligence 1545*. Springer.

Werner, G. and Dyer, M. (1991). Evolution and communication in artificial organisms. In Langton, C., Taylor, C., and Farmer, J., editors, *Artificial Life II, Vol. X of SFI Studies in the Sciences of Complexiy*, Redwood City, Ca. Addison-Wesley Pub. Co.

Wittgenstein, L. (1958). *Philosophical Investigations*. Basil Blackwell, Oxford, UK.

Yanco, H. and Stein, L. (1993). An adaptive communication protocol for cooperating mobile robots. In Meyer, J.-A., Roitblat, H., and Wilson, S., editors, *From Animals to Animats 2. Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 478–485, Cambridge Ma. The MIT Press.