The evolution of a lexicon and meaning in robotic agents through self-organization

Paul Vogt

Vrije Universiteit Brussel, Artificial Intelligence Laboratory
Pleinlaan 2
1050 Brussels, Belgium
e-mail: paul@arti.vub.ac.be
http://arti.vub.ac.be/ paul

Abstract. This paper discusses interdisciplinary experiments, combining robotics and evolutionary computational linguistics. The goal of the experiments is to investigate if robotic agents can originate a language, in particular a lexicon. In the experiments two robots engage in a series of so-called language games. Starting from the assumption that the robots know how to communicate and are able to detect some sensory information from the environment, the agents ground conceptual meaning and develop a lexicon. The experiments show that the robots are able to form a shared communication system. The paper investigates the influence of using non-linguistic information in the formation of the lexicon, which takes the form of pointing (1) to indicate the topic of the language game, and (2) to give feedback on the outcome of the game.

1 Introduction

In the last decade there has been an increasing interest in a dynamical approach to evolutionary linguistics (Hurford et al., 1998). In this approach language is seen as a complex dynamical system, like one can see in thermodynamics or in biological systems. Our underlying theory assumes that language emerges through self-organization within a community of culturally interacting agents (Steels, 1996a). In addition, meaning emerges from interactions with the environment, individual adaptation and self-organization (Steels, 1996b).

Our research group investigates various aspects of the emergence of language ranging from lexicons, meaning, phonetics, syntax and language change, for an overview see (Steels, 1997). Also research is underway on the evolution of communication (De Jong, 1997). Our language system is comparable to the ones described in (Oliphant, 1996), (McLennan, 1991) and (Werner and Dyer, 1991). Most experiments were carried out in computer simulations, whereas some other were grounded in physical robotic experiments like in the emergence of phonetics (De Boer, 1998), the 'talking heads' experiment (Belpaeme et al., 1998) and in mobile robots (Steels and Vogt, 1997). This paper discusses experiments of the latter kind. Other research on grounding communication in mobile robots can be found in (Billard and Dautenhahn, 1997) where a student robot learns a lexicon by imitating a teaching robot, using a connectionist model, and in

(Yanco and Stein, 1993) where lexicons are learned from human instructions by reinforcement learning.

The experiments are centered on the notion of *language games* (Wittgenstein, 1967). The robotic language game integrates two variants on the basic language game: *discrimination games* (Steels, 1996b) and *naming games* (Steels, 1996a). In discrimination games the task for the agents is to categorize (or ground) sensory information in terms of distinctive categories, resulting in a symbolic representation of meaning, conform (Harnad 1990). In the naming games the agents learn to communicate words that are associated with the learnt meaning. Depending on the outcome of the game (success or failure) the robots' lexicons are adapted accordingly as explained in (Steels and Kaplan, 1998). The lexicon is formed through its use in communication and the success of the communicative acts.

This paper discusses the experiments and in short the improvements made in contrast to previous work (Steels and Vogt, 1997), (Vogt, 1998) and (De Jong and Vogt, 1998). The most important improvement came from the innovation in the naming game model so that it can deal with stochasticity (Steels and Kaplan, 1998). The experiments discussed compare influences that non-linguistic information can have on the lexicon formation. The next section discusses briefly the implementation of the language games. In section 3, results of the experiments are reported. And finally, section 4 presents some concluding remarks.

2 Language Games

In the experiments the aim of the two robots is to develop a communication system with which they communicate about the same objects in their environment. Or, to be more precise, to communicate about similar distinctions that the robots observe in particular situations. In order to do so, the robots engage in a series of language games as described in (Steels and Vogt, 1997). It must be stressed that the robots do not have any preprogrammed knowledge about meaning and lexicon; they only know how to communicate and how to acquire meaning and a lexicon. In the course of the experiment, the robots acquire meaning and word-meaning associations without supervision of an (external) observer.

The experiments involve two mobile LEGO robots that are equipped with sensors that can sense light intensities from different light sources that exist in their environment, see figure 1. The environment is especially set up for robot experiments and it consists of different light sources: there is one white light sources, there are two light sources that emit light at a modulated frequency and the two robots have infrared sources. A radio link is used for non-linguistic communication to synchronize the robots' processes in the cooperation of the language game, and to send sensed object information to a PC. A PC is used in order to increase memory and speed, and also to enable experiments in which one can compare different influences of parameters and methods more reliable.

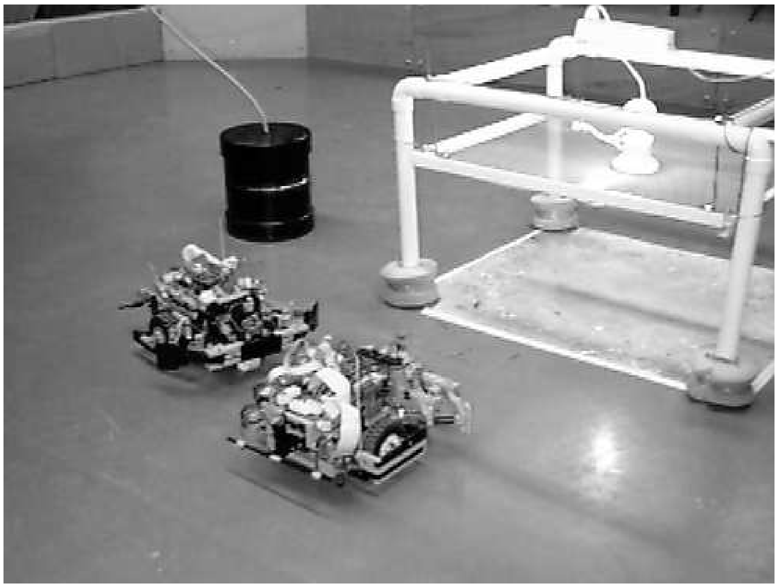

Fig. 1. This picture shows the robots used in the experiment as they are engaged in a language game and the objects the robots can detect, from which they learn to name related features. Besides the robots themselves, there is a black cylindrical object that emits light modulated at a particular frequency (upper left), and there is an object emitting white light (right). The robots make themselves visible with infrared emitters. Both robots are equipped with sensors to detect all these types of light sources.

The basic scenario of a language game is described in (Steels and Vogt, 1997). The robots come together at close distance ending up facing each other. One by one they rotate 360 degrees to observe their surroundings. Doing so, they segment the sensor information in time, extracting *objects* from the perception as in (De Jong and Vogt, 1998), thus constructing a *context* of objects. Every object is defined by the values sensed for several *sensory channels*, which are functions over a particular sensor. The term objects here must not be confused with the objects as we humans detect them, but they are rather segmentations of the sensor information. The context, including the observed sensor information is then sent to the PC.

At the PC, the language game is further processed. The robots in the PC are implemented as artificial agents. First, the speaker of the game chooses an object to be the *topic* of the conversation. Then it points at this object by giving the hearer the angle at which the speaker observed the object. This is analogous to the mechanism described in (Steels and Vogt, 1997), where the speaker orients towards the topic, while the hearer observes this behavior. From this, the hearer identifies the topic by looking which object in its own context is closest to the angle the speaker pointed at. As mentioned in (Steels and Vogt, 1997), this method is very imprecise. Various questions may rise: Is pointing at all is

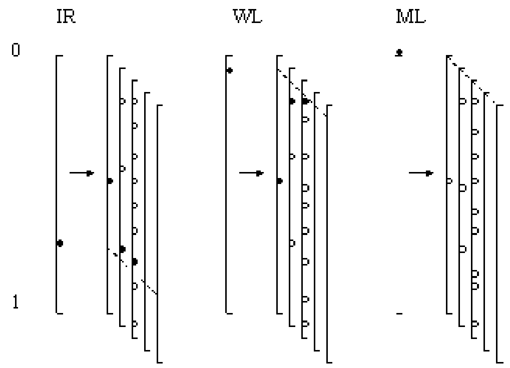

Fig. 2. The hierarchical trees of prototypical features for sensory channels infrared (IR), white light (WL) and modulated light (ML) sensors are shown on the right of every arrow. The features are placed as a dot on every layer. The sensor values that are detected by the agent (shown on the left of every arrow) are related to the closest feature on every sensory channel in every layer of the hierarchy (dark bullets), as schematically shown. Note that the ML sensor does not relate to a feature since the sensor value is zero. A new feature can be created on one of the sensory channels that have a non-zero value, which becomes the feature's value. The feature is added to one of the hierarchical layers that have still place for a feature and where the value of the feature has a minimum distance to its neighbors.

necessary? Should one use the information as a fact, or only as an indication? Other questions are: must we use this information prior to the linguistic communication takes place, or can it be used to give feedback as well? These are the central questions that will be investigated in the experiments reported here. The basic idea is now that the hearer calculates for every object in its context the likelihood for that object to be the topic, depending on the distance between the observed and pointed angle. The likelihood (or *topic score*) for object o is calculated as follows (Steels and Kaplan, 1998):

$$\varepsilon_o = \frac{1}{1 + (\frac{d_o}{\alpha})^2}$$

In this equation, d_o is the distance between the observed angle of object o, and the angle that the speaker pointed at, α is the tolerance factor of the pointing. The topic score $\varepsilon_o \geq 0.5$ if the difference between the angles is smaller than the tolerance factor. In the naming game, the hearer will consider all objects that have a non-zero topic score as the possible topic. All objects with a non-zero topic score, and that have a meanings associated with the word that the speaker utters (see below) are said to be in the hearer's *scope*. Since the robots are standing opposite of each other during a language game, they may obscure objects that are standing behind the observing robot, causing the other robot not to be able to detect this object. In order to overcome this problem, the objects that are standing within a certain range behind the robot may not be considered in the language game. Another solution may be that the topic score ε_o is multiplied by e.g. 0.1 for these objects.

All objects can be related to a set of *features*. A feature relates sensory channels with a certain value to a sub-symbolic representation that the agent has constructed (see figure 2). Initially the agents only have one feature for every sensory channel. From this point, the agents individually start to play discrimination games by trying to distinguish the topic from all other objects in the context. Since the hearer might consider more than one object to be the topic, it might play several discrimination games. The discrimination games, first described in (Steels, 1996b) and developed further in (Vogt, 1998) and (De Jong and Vogt, 1998), aim to find those features of an object that distinguish that object from the other objects in the context. This yields a so-called set of *distinctive feature sets*, which we may use as the meaning of that object. If the distinctive feature set becomes lexicalized, we may refer to this set just as a *feature set*, because it is not necessarily distinctive any more, since they may not be valid in every language game. If discrimination is not possible, new features can be constructed. The way this is done in the experiments presented here differs from previous experiments.

In (De Jong and Vogt, 1998) we introduced the Simple Prototype method. In this method a feature was created as an example of the topic at the end of a failed discrimination game, taking one sensory channel as the attribute of the feature and the sensed value of that sensory channel as the feature's value. All features on one sensory channel were spread over the whole range of the sensor in one layer, so without hierarchy as in (Steels, 1996b). In these experiments, the prototypical features are organized in a hierarchy, as can be seen in figure 2. In addition, no forgetting is incorporated (as was the case in (De Jong and Vogt, 1998)), but feature sets are only remembered when they are associated with a word.

After the discrimination, the naming can start. The naming process is adapted from the naming game model described in (Steels and Kaplan, 1998). The speaker chooses the most effective distinctive feature set and looks for a word-meaning association in the lexicon that has the highest association score. Each agent has a lexicon that is a set of word-meaning associations that each has an *association score* σ. A distinctive feature set is effective if it has a minimal amount of features, if it is not very deep in hierarchy, if the features that constitute the set are used frequently and if the set has been lexicalized. For all these influences there are measures to calculate the effectiveness. If no association could be found with this distinctive feature set, the next most effective feature set is chosen until an association is found or no more distinctive feature sets are available.

The hearer decodes the expression using a decision matrix as was introduced in (Steels and Kaplan, 1998). This matrix contains all the associations that have a meaning consistent with the meaning of objects in the hearer's scope, which may be derived from the pointing as explained above. Together with a meaning score, which is calculated from the effectiveness of the feature set, the hearer chooses the association for which the sum of the topic score, the meaning score and the association score is highest.

In previous experiments the feedback of a language game was given by the hearer as a success or failure signal when the hearer *understood* the speaker, i.e. when the hearer had an appropriate association in its lexicon (Steels and Vogt, 1997). In the experiments reported here, this method is compared with giving feedback by means of pointing. The hearer then points at the object which the hearer now identified as the topic. In turn, the speaker calculates the *success score* β that is calculated in the same way as the topic score ε, comparing the angle of the speaker's topic with the angle the hearer pointed at.

According to the outcome of the language game, the lexicon is adapted. This is done as follows:

- *The speaker had no word associated with any meaning.* Then the speaker might create a new word with probability P_s. This is already done before the word is 'uttered', so if a new word is created, it is uttered immediately.
- *The naming game was a complete success*, i.e. $\beta \geq 0.5$. In this case, the hearer pointed at an object that the speaker 'recognizes' as the topic within a certain tolerance factor. Then the association scores (are updated as follows:
  - $\sigma = \sigma' + \beta \cdot \delta$ for the winning association.
  - $\sigma = \sigma' - \beta \cdot \delta$ for all other associations in the matrix. where σ' is the old association score, and δ is some constant.
- *There was a mismatch in meaning between the hearer and the speaker*, $\beta < 0.5$. Both the hearer and the speaker update the used association score as follows:
  - $\sigma = \sigma' - (1 - \beta) \cdot \delta$. Note that in all cases (is bounded between 0 and 1.
- *The hearer has no meaning associated with the expression.* In this case the hearer may associate the word with the distinctive feature sets that resulted from the discrimination for the object that has the highest topic score ε if and only if $\varepsilon_{max} \geq 0.5$. Two other constraints are set for the actual association. Associations are made
  - when the features constituting the distinctive feature set are used frequently, or
  - with a certain probability P_h.

In order to be able to follow the process of the meaning and lexicon formation three monitoring measures are used. The measures are defined as follows:

- *Discrimination success* is the average success over the past 100 discrimination games. Note that the hearer may play several discrimination games in one language game.
- *Communicative success* is the average success over the past 100 language games.

– *Objective success* is the *actual* success of the language games taken as the average over the past 100 language games that have communicative success. Success is assigned when, in case of communicative success, the robots were really communicating about the same object. It is assumed that the robots actually communicate about the same object, if both objects were detected using the same sensory channel; this is information that can only be observed by an objective observer.

In this section, we briefly saw how the experiments are set up and how the language game is organized. The script of the language games on the robot is extremely complex and is controlled by some processes that are sometimes depending on imprecise and incomplete information. This results in the observation that most of the attempts to complete a language game will fail. In addition, as we shall see, the concept (or meaning) formation is always successful, and the successfulness of the lexicon formation depends on the way non-linguistic information is used.

3 The experiments

This section will discuss several experiments that have all been carried out with the same robot data that was recorded beforehand. The data set that was used consists of 930 language games, recorded in approximately 25 hours. This comes to approximately 1 language game every 1.6 minutes, which is pretty slow. The slowness is due to imprecise sensors and the unreliability of the radio-link, see (Vogt, 1998). Speeding up the process would increase the unreliability of the language games. We could, for example, decrease thresholds for the infrared sensors, which are used frequently in the course of a language game for deciding whether or not another robot is in the vicinity of the processing robot. This, however, will cause language games in which the robots' distance is too large to observe a coherent context, and even more language games in which rotation during the perception is either significantly less or significantly more than 360 degrees. This already happens quite a lot.

All experiments consisted of a run in which 10,000 language games were played. The first 930 games were run through the data in recording order, afterwards every game was chosen randomly. In all the games the role of the speaker was assigned arbitrary to an agent, as well as the topic of the speaker. The tolerance factor α was taken to be 1 in order to allow a deviation of 90 degrees in the accuracy of the pointing. The measure δ that is used to update the association score σ was set to 0.1, the word creation probability P_s was set to 0.1 and the word association probability P_h was set to 0.4.

The following subsection discusses one of the experiments in detail. In this experiment, the formation of meaning and word-meaning is investigated where the method of using non-linguistic information is held constant. This experiment is compared with the experiment in which feedback is given as a success or failure signal. The other experiments, discussed in section 3.2, are variants on this game

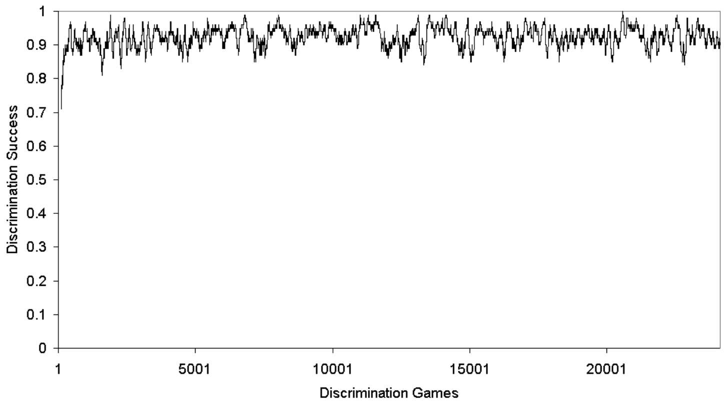

Fig. 3. The discrimination success of one of the robots.

to show influences of including pointing as initial non-linguistic information for the hearer, either with or without linguistic information, or considering one or more objects in the hearer's scope. Also a comparison is made to show the influence of pointing back as feedback as opposed to a success or failure feedback by the hearer, i.e. positive feedback is only given by the hearer when it understood the speaker.

3.1 The concept and lexicon formation

This subsection discusses the results of one of the experiments in detail. In this experiment, after the speaker has chosen a topic, it will point at this object and the hearer calculates the topic score for every object in the hearer's context. The speaker encodes an expression and the hearer constructs the decision matrix and it will decode the speaker's expression. Feedback is given by pointing back at the topic the hearer determined, and the speaker then evaluates the outcome of the game.

If we look at figure 3 we can see that the agents learn to discriminate the objects very fast. Already within 200 discrimination games, the agent has a discrimination success higher than 80% and will not drop under this value anymore. Note that the amount of discrimination games (approximately 25,000 games) is higher than the total amount of language games that have been played. This is of course because the hearer plays a discrimination game for every object with a non-zero topic score. Although this figure varies per experiment, the development and success rate is more or less the same for every experiment. The long run average success rate is approximately 90% for every robot. The amount of effective feature sets (i.e. those feature sets that were successfully lexicalized) did not exceed the number of 50, which is a significant improvement compared to the experiments discussed in (De Jong and Vogt, 1998).

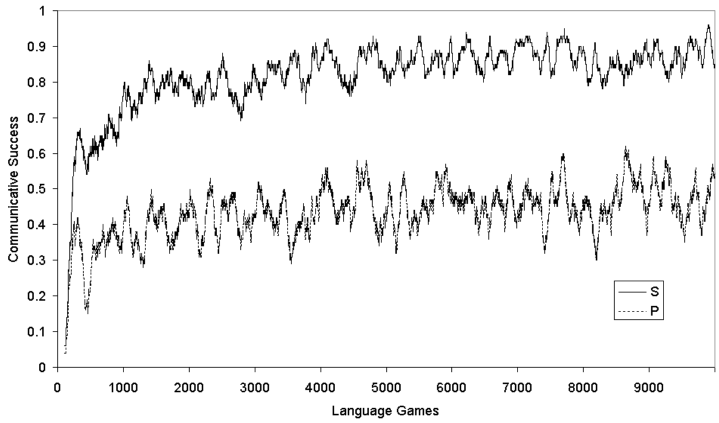

Fig. 4. The communicative success with feedback given by pointing (P) and as a signal (S).

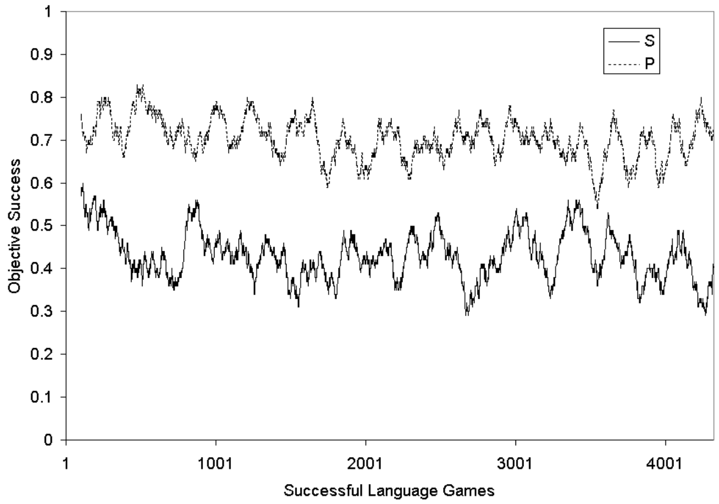

Fig. 5. The objective success with feedback given by pointing (P) and as a signal (S).

Figure 4 shows the evolution of the communicative success for this experiment (dashed line). We can see that the communicative success is lower than in the case when feedback was given by a success or failure signal (continuous line). In addition, we can see that the communicative success rates increase rapidly in the beginning and, although there is still a lot of variation, it becomes more stable after approximately 5,000 games. On the other hand, if we look at the objective success (figure 5), we see a higher success rate in the case were feedback is given by pointing. So, the robots have a higher success rate in actually naming the objects appropriately, although the cost of increasing reliability is a lower communicative success.

It seems that giving feedback by pointing has a strong influence on the successfulness of the language game. There are significantly less language games successful because the agents use an action (pointing) in order to agree that both robots communicate about the same object. As we can see in figure 5, the agents still do not agree in every language game, but it is significantly better than in the case where feedback is given by a signal. That the communicative success does not exceed the 50% rate on the average is probably due to imprecise sensing by the robots as well as unreliable data transmission of the context using the radio link.

WM r0	aab	aac	aaf	aam	WM r1	aab	aac	aaf	aam
sc0-0-1	1.00	0.00	0.00	0.00	sc3-0-127	1.00	0.00	0.00	0.00
sc3-0-127	0.00	0.90	0.00	0.00	sc3-1-97	0.20	0.00	0.00	0.00
sc2-1-208	0.00	0.30	0.00	0.00	sc0-0-1	1.00	1.00	0.10	0.00
sc2-3-84	0.00	0.00	0.10	0.00	sc2-1-183	0.00	0.00	0.27	0.00
sc2-4-17	0.00	0.00	0.07	0.00	sc2-3-64	0.00	0.10	0.00	0.00
sc1-2-216	0.00	0.00	0.00	0.10	sc1-2-91	0.00	0.00	0.00	0.10
sc1-3-38	0.00	0.00	0.00	0.01	sc1-3-4	0.00	0.09	0.00	0.00
sc1-5-119	0.00	0.10	0.00	0.10	sc1-4-10	0.00	0.00	0.00	0.10
					sc1-4-173	0.00	0.00	0.00	0.16

Table 1. A snapshot of the lexicon of the two robots after 10,000 language games. Note that sc0 (the 'self') in one robot should be associated with sc3 (the infrared) in the other robot, which indeed happens.

A part of the resulting lexicon of this experiment is shown in table 1. If we look at both lexicons we see that both robots have the same words associated with more or less similar meanings. Although there is ambiguity and homonymy, the lexicon is pretty alike. Some of the ambiguity is what I have called *representational ambiguity* (Vogt, 1998), as opposed to lexical ambiguity. Representational ambiguity means that the same word is used to indicate different meaning, but since all these feature sets act on the same sensory channel they may indicate the same real world object. For instance the word **aaf** is in both agents associated with several feature sets that all are related to the same sensory channel (*sc1*). On the other hand, robot r1 has associated word **aab** with both the infrared-channel (*sc3*) and the self-channel (*sc0*), which is a lexical ambiguity. Homonymy can be seen for robot r1 that has 3 different words associated with meaning *sc0-0-1*, but it prefers word **aab**. Although it is not shown here, the lexicon also revealed a lot of inappropriate associations, where one robot has associated a word with a different sensory channel than the other robot did.

3.2 Comparing different influences

In this subsection, different types of non-linguistic information are compared. To do so, both the types of non-linguistic topic information that the speaker gives to the hearer and the ways of giving feedback are varied. The variation of the non-linguistic information given to the hearer prior to the linguistic communication considers the following five mechanisms:

- *Random* (R): Both the speaker and the hearer choose the topic randomly. The hearer only considers the topic in its scope; i.e. the topic score of the chosen object is 1, whereas all other objects have a topic score of 0.
- *Randomly constrained* (RC): Again both the speaker and the hearer choose the topic randomly. But now the agents do not consider the objects that are behind the agent in a range of 36 degrees. This is done because the agent may obscure these objects for the other robot.
- *Pointing* (P): In this experiment, the speaker points at the topic it has chosen, and the hearer identifies the object with the highest topic score as the topic. The hearer only considers this object as the possible topic, like in (Steels and Vogt, 1997).
- *Scope* (S): In this experiment, as in the random case, the hearer does not receive information from the speaker, but it considers in the matrix all objects it detected in its scope, so all objects have an equal topic score ($\varepsilon = 1$). If a new word-meaning association has to be made by the hearer, then one object is chosen randomly to be the topic.
- *Pointing and scope* (PS): Here, the speaker points at the topic. The hearer then calculates the topic score for every object, as described in section 2. All objects that have a non-zero topic score are considered in the scope of the hearer. This is the same experiment as discussed in the previous subsection.

For all the above cases two ways of giving feedback are explored. The first one is the case where the hearer gives a success or failure signal (S), and the second is where the hearer points at the object it determined as the topic (P). In the first case, a language game ends in success if the hearer *understood* the speaker, in all other cases the game ends in failure. In other words, the language game is successful if the success score $\beta > 0$. So, the agents cannot determine any mismatch in meaning (see section 2). In the second case, the game ends in success if the angle the hearer points at lies within the tolerance factor around the angle that the speaker has observed its topic. For the experiments R, RC and P we also look at the a priori chances that the agents choose the same topic. I.e. the robots individually choose a topic, and we see whether they have chosen the same topic without having linguistic communication.

In table 2, we see the results after 10,000 language games for each experiment. In the first two columns we see the description of the experiments, in the third and fourth column we see the average communicative success CS with its standard deviation s_{CS} and in the fifth and sixth column the average objective success OS with its standard deviation s_{OS}.

	TC	F	CS	s_{CS}	OS	s_{OS}
1	R	-	-	-	33.77	4.45
2	R	S	76.23	12.08	34.85	5.81
3	R	P	14.45	4.10	58.11	5.67
4	RC	-	-	-	33.41	4.53
5	RC	S	81.03	10.06	41.17	5.73
6	RC	P	17.80	3.82	62.91	4.48
7	P	-	-	-	34.38	4.41
8	P	S	78.97	11.97	51.55	4.81
9	P	P	56.90	8.42	63.09	5.10
10	S	S	82.24	10.06	41.55	8.25
11	S	P	30.28	8.56	79.12	5.97
12	PS	S	81.83	10.08	42.82	5.14
13	PS	P	42.82	5.14	70.37	4.67

Table 2. The comparison of the influences of using non-linguistic information. All figures are percentages. TC stands for Topic Choice and F stands for Feedback.

First we can note that the random case R yields the lowest communicative and objective success in all three cases, except for the a priori chance (row 1). We further see that whenever pointing is used as feedback (P), the communicative success drops significantly in comparison with the cases where feedback is given by a signal (S). But, on the other hand, the objective success increases significantly. All a priori chances (-) that both agents choose the same topic are the same within significance independent of the method used.

In the experiment in which the topic was chosen randomly, but which has a constraint of a dead angle (RC), the results are slightly better than in the R case. The a priori chances of having the same topic are almost exactly the same as in the random case (33.41% vs. 33.77%). If pointing is used to identify the topic and when the feedback is given by a signal (row 8), then the objective success increase significantly compared to row 5 (51.55% vs. 41.17%). This is, however, not the case where pointing was the feedback mechanism (row 9 compared to row 6). Row 9 shows a slight increase in the objective success rate (63.09% vs. 62.91%). But we see a significant increase in the communicative success as compared to the R (row 3) and RC (row 6) cases (56.90% vs. 14.45% and 17.80% respectively).

Although we would expect that the case in which both pointing and considering a scope using additional linguistic information was used to determine the topic by the hearer (row 13) would reveal the highest objective success (Steels and Kaplan, 1998), the experiments showed that this was not the case. The experiment in which every object has an equal topic score and where the topic was determined only by the meaning- and association scores revealed the highest objective score (row 11), 79.12% vs. 70.37%. The communicative success, however, is relatively low compared to the PS case (30.28% vs. 42.82%). As Steels and

Kaplan showed in their simulations, if the stochasticity is high (i.e. when there is no or very little information on the topic), then the lexicon cannot be build up. They, however, use a somewhat different measure to show this: *coherence*. Coherence, as they use it, indicates the average number of agents that prefer the same words for every meaning. Since the internal meaning representation may vary a lot in the robotic agent (as opposed to the simulated agent) the coherence measure cannot be calculated in the experiments presented here in the same way as in their experiments. It is obvious that the robotic experiments have a high level of stochasticity. Steels and Kaplan argued that the hearer needs non-linguistic topic information prior to the linguistic communication in order to bootstrap the construction of a shared lexicon. Our experiments, however, show that this is not a necessary condition. It is not obvious yet why this is the case. It seems that the feedback using pointing is sufficient to construct a shared lexicon if the hearer considers more than one object in its scope. But Steels and Kaplan also give feedback with object information (i.e. in their experiment, the naming game is a success if the topic of both agents indicate the same object). In addition, initial pointing only appears to increase the stochasticity of the system. But, it is hard to say what system performs best, the PS system or the S system, since the results overlap within the range of the standard deviations. More thorough statistical analysis is needed.

The question why the communicative success remains so low in the feedback by pointing systems was already discussed in the previous subsection. It is partly due to the stronger feedback condition and it is probably for a great deal due to imprecision in the perception task of the robots, where they do not always rotate exactly 360 degrees. So, the position of objects cannot always be approximated with an angle properly. A lot of uncertainty, is also due to inconsistencies arising from the assumption that both robots determine the same context with the objects observed under similar angles.

4 Conclusions

In this paper the emergence of meaning and lexicon in robotic agents has been discussed. The processes of creation and acquisition are selectionistic and self-organizing. The lexicon is propagated through social interactions between the agents. The whole model is based on the interactions of the individual agent with its environment including other agents, and a natural selection-like ontogenetic development.

First we saw that the agents are able to ground meaning from simple sensory perception from the bottom-up, conform (Harnad, 1990). Although a different method for feature generation is used than discussed in (Steels, 1996b), (Vogt, 1998), (Ward, 1997) and (De Jong and Vogt, 1998), the agents learn to discriminate reliably within 200 discrimination games. So, we may conclude that the actual implementation of the feature generation is not very important for the success of the discrimination as such. But, as discussed by De Jong and Vogt, the method may make significant differences with respect to the task for which the results are used. It appears that this method is promising in that it uses only a few feature sets in the lexicon. Furthermore, the experiments reveal that the discrimination success do not converge to exactly one. This is because there are some objects within one context that are described by the same sensor characteristics, so distinctions can not be made. This should put pressure on the generation of new sensory channels, which has not been implemented in this experiment.

Another interesting result is that although the objective success is higher when feedback is given by means of pointing, the overall communicative success from the point of view of the robots drops significantly. Although this result has to be investigated further, it is thought that, besides the stronger conditions on accepting the success of a naming game, imprecise sensing and executing the perception task by the robots is the main reason for why the communicative success stays well behind.

Comparing influences of using non-linguistic information prior to the actual linguistic communication on the topic choice revealed a surprising result. It appears not to be necessary to indicate the topic in advance, but adopting word-meaning associations to random objects in a context and selection in later games considering the whole context in the scope appear to do well. This has been observed before in the talking heads experiment, and might be due to the arbitrary generation of variation and selective pressures in similar situations (Van Looveren, personal communication). On the other hand, according to (Steels and Kaplan, 1998), these results should not have been possible. In Steels and Kaplan's experiment no (coherent) lexicon arose when stochasticity was high concerning initial topic indication. In addition, using initial topic information also seems to be very promising when it is used in the naming game model that considers more than one object in its scope. If the hearer, however, only considers one possible topic, the objective success rates stay well behind. All these results have to be investigated in greater detail in the near future.

So, although a lot of progress has been booked in relation to our previous work (Steels and Vogt, 1997), there is still a lot of work to do. First, the sensing capabilities and radio link reliability are being improved. Secondly, the basic naming game model used here has to be investigated in more detail to explain the results reported here. Furthermore, a more thorough statistical analysis has to be carried out on the data in order to understand better what the system does.

5 Acknowledgements

The author would like to thank Aude Billard, Bart de Boer, Edwin de Jong, Frederic Kaplan, Joris van Looveren and Luc Steels for the discussions and useful suggestions made on various topics. The Vrije Universiteit Brussel sponsors this research as an assistantship.

References

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

Billard, A. and Dautenhahn, K. (1997). Grounding communication in situated, social robots. In *Proc. Conference "Towards Intelligent Mobile Robots"*, Manchester.

De Boer, B. (1998). Emergence of sound systems through self-organisation. In Knight, C. and Hurford, J. R., editors, *The Evolution of Language (selected papers from the 2nd International Conference on the Evolution of Language, London, April 6-9 1998).*

De Jong, E. (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. and Vogt, P. (1998). How should a robot discriminate between objects? In *Submitted to SAB 98.*

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

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

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.

Steels, L. (1996a). Emergent adaptive lexicons. In P., M., 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 M., T., 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.

Vogt, P. (1998). Perceptual grounding in robots, lecture notes on artificial intelligence. In Birk, A. and Demiris, J., editors, *Proceedings of the 6th European Workshop on Learning Robots, Lecture notes on Artificial Intelligence.* Springer.

Ward, D. (1997). Analysis of perceptual grounding through parameter variation. In *Notes of the VIM'97 Workshop on Collaboration between human and artificial societies.*, Universita' di Salerno.

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. (1967). *Philosophische Untersuchungen*. Suhrkamp, Frankfurt.

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.