Perceptual grounding in robots

Paul Vogt

Vrije Universiteit Brussel, Artificial Intelligence Laboratory
Pleinlaan 2
1000 Brussels, Belgium
e-mail: paul@arti.vub.ac.be

Abstract. This paper reports on an experiment in which robotic agents are able to ground objects in their environment using low-level sensors. The reported experiment is part of a larger experiment, in which autonomous agents ground an adaptive language through self-organization. Grounding is achieved by the implementation of the hypothesis that meaning can be created using mechanisms like feature generation and self-organization. The experiments were carried out to investigate how agents may construct features in order to learn to discriminate objects from each other. Meaning is formed to give semantic value to the language, which is also created by the agents in the same experiments. From the experimental results we can conclude that the robots are able to ground meaning in this self-organizing manner. This paper focuses on the meaning creation and will only discuss the language formation very briefly. The paper sketches the tested hypothesis, the experimental set-up and experimental results.

1 Introduction

In the past few decades a lot of research has been done on the grounding problem. The grounding problem investigates how symbolic structures may arise from non-symbolic sensory information extracted from the physical world (Harnad, 1990). This paper discusses how robotic agents may ground symbolic meaning from sensory observations. Approaches to this problem normally use rule-based systems, connectionist models or genetic learning mechanisms. Our approach is a bottom-up approach, conform the behavior-oriented approach to artificial intelligence (Steels, 1994). This means that agents use sensory information to conceptualize meaning from the bottom-up, which has also been suggested by Harnad. The experiments are embedded in our research on the origins of language (Steels, 1996a).

The experiments are based on the hypothesis that language is a complex dynamical system and that a lexicon and meaning emerges through interaction within a population of agents and their environment (Steels, 1997). Cultural interaction between agents causes the language to spread through the population. Individual adaptation lets the lexical and semantic symbols be categorized. Because the system is a distributed complex dynamical system, there is no agent that has a complete view of the language, nor is there an agent that can control the language. The major goal in studying the origins of complex systems is to find the boundary conditions and mechanisms, like self-organization, that give rise to the complexity, see e.g. (Prigogine and Strengers, 1984).

In our research on the origins of language, see e.g. (Steels, 1996c), agents can adapt a *lexicon*, which are words that they associate with semantic symbols, which in turn discriminate between concepts (be it objects or other relations). The lexicon is propagated through the population of agents by means of social interaction. This is done by a distributed set of agents that have the ability to communicate with each other, and a self-organizing learning method is used to adapt the lexicon appropriately. No linguistic information is pre-programmed and the lexicon emerges through a cultural evolution. This is opposed to the linguistic theory introduced by Chomsky (see e.g. (Chomsky, 1980)), who claims that linguistic structures are innate.

As there is no linguistic information pre-programmed there is also no pre-programmed meaning. In the experiments individual agents ground the semantic symbols. The robots ground the symbols, which compose the linguistic meaning of a concept, i.e. the meaning of a word. They do so by classifying sub-symbolic low-level sensory data into distinctions, conform (Harnad, 1990). The grounded symbols (or feature sets) distinguish one object from some other observed objects. An adaptive self-organizing learning mechanism is incorporated in the system in order to select those symbols that are useful for the lexicon formation (Steels, 1996b). Learning is achieved by generating features, which relate to observed sensory information of an object. These features are then classified based on their usefulness by means of a selectionist approach. This method is based on natural selection during the interaction with the physical world and within a lifetime (as suggested by e.g. (Foley, 1997)(Edelman, 1987)). The useful features will be used more frequently than unuseful ones. The useful features remain within the system, whereas the unuseful ones will be forgotten. The method differs from genetic algorithms because the agents do not inherit information from generation to generation. All learning and selection takes place within a lifetime of an agent. They also do not reproduce features for selection, but they create new features.

This paper reports on an experiment where robotic agents can develop a lexicon about objects in their environment using these mechanisms. In (Steels and Vogt, 1997) a similar experiment is reported, in that paper the focus was however on the co-evolution of language and meaning, conform (Steels, 1997). Here the focus is on the evolution of meaning in robotic agents. Especially the physical implementation of the hypothesis is discussed in detail.

In the next section the mechanisms for meaning creation and learning are explained. In section 3 the whole experimental set-up is described. Section 4 reports on the experimental results. And finally, in section 5, conclusions and future research are discussed.

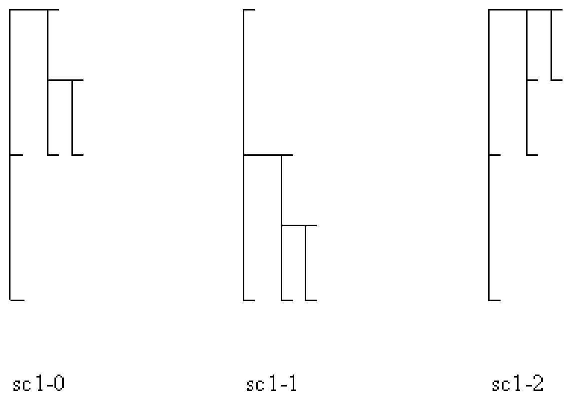

Fig. 1. The robots can construct a binary tree of features that divides the sensor space.

2 The evolution of meaning

At the AI-Lab in Brussels we investigate how meaning may evolve using selectionistic mechanisms (Steels, 1996b). These mechanisms are mainly based on adaptive self-organization. The mechanisms for the meaning creation are (1) *feature generation* and (2) *self-organization* . The generation principle is based on the generation of features that may represent sensory information of an object. Self-organization by means of a natural selection-like learning method causes the meaning to be coherent and useful.

The meaning creation is introduced by Luc Steels (Steels, 1996b). The principles for meaning creation are valid for all kinds of concepts like physical objects, spatial relations, internal states and movements. It emerges in an individual agent under the notion of *discrimination games*: Agents construct binary feature trees of the sensory space representing objects (figure 1). Observed objects are related with feature nodes that are sensitive for observed sensor values. These features are used to discriminate one object from other objects. Those features that are capable to distinguish one object from another constitute the feature sets and represent the symbolic meaning. The following description is adapted from (Steels, 1996b).

Suppose we have a system that consists of a set of objects $O = \{o_1, \ldots, o_n\}$ and a set of agents $R = \{r_1, \ldots, r_m\} \subset O$. Each agent r has a set of sensory channels $S_r = \{sc_1, \ldots, sc_p\}$ and a set of features $F_r = \{f_1, \ldots, f_q\}$. A sensory channel sc_j is a real-valued function from $O \to \Re$, which is (in-)directly derived from the agent's sensor. A feature f is a pair $(p\ V)$, where p is the attribute name (usually sc_j) and $V = [v_{min}, v_{max}] \in SC_j$. Here SC_j is the range in which sc_j is sensitive.

In a discrimination game, the procedure is as follows: Suppose an agent has observed a context that consists of a set of objects: $C = \{o_i \mid o_i \in O\}$. Observing an object in the robot's world using (low-level) sensors results in a set of sensor values for every sensory channel. These values can be related to features that are sensitive to these values. So every object $o_i \in C$ can be related to a set of features $F_{r,o_i} \in F_r$. An object o_i can be described by a set of *descriptions*:

$$\Lambda^{o_i} = \{\Lambda_k^{o_i} \mid \Lambda_k^{o_i} \subseteq F_{r,o_i}\}$$

where $\Lambda_k^{o_i}$ is any possible description of an object o_i constrained as follows:

$$\Lambda_k^{o_i} \subseteq \{f_j = (p_j V_j) \mid f_j \in F_{r,o_i} \wedge (l \neq j \Rightarrow p_l \neq p_j)\}$$

Note that: (1) The set on the right hand side of this relation may result in several sets since the o_i may be related to several features on one sensory channel, since the features are build up in a tree. The relation holds for every possible set governed by the right hand side. And (2) $\bigcup_k \Lambda_k^{o_i} = F_{r,o_i}$.

The agent now identifies a *topic* $t \in C$ of the discrimination game. Then the agent tries to distinguish the topic from the other objects in the context by constructing a *set of distinctive feature sets*. A set of distinctive feature set can be defined as:

$$D_t^C \equiv \{\Lambda_j^t \mid \forall o_i \in C \backslash \{t\} : (\neg \exists \Lambda_k^{o_i} : (\Lambda_k^{o_i} = \Lambda_j^t))\}$$

So, the agent constructs a set of features that relates to t and which distinguishes t from the other objects in the context. A distinctive feature set is associated with a use-score and a success-score. A discrimination game results in a *set of* distinctive feature sets (D_t), which can be empty, but which can also consist of several distinctive feature sets. All distinctive feature sets resulting from the discrimination game are stored in the memory of the robot to be associated words in the language games.

If the agent thus determined a set of distinctive feature sets, the following is possible:

- The agent was not able to discriminate $t \in C$, i.e. $D_t = \emptyset$. The discrimination game was a failure.
- The agent was able to discriminate $t \in C$, i.e. $D_t \neq \emptyset$. The discrimination game was a success. If there is more than one distinctive feature set, the agent will choose the most general distinctive feature set as follows:
  - Choose the set with the smallest amount of features.
  - If there is still more than one distinctive feature set, then choose the set that is least segmented in the feature space. I.e. that set that has the least average depth in the feature tree.
  - If there are still several sets left, then choose the most successful one. I.e. the set for which the measure $m = success/use$ is the highest. (Steels, 1996b).

If the discrimination game was a failure, then the agent constructs a new feature. This can be done in two ways: (1) If there is still a sensory channel unused, then use this sensory channel to construct a feature in the range of this channel. (2) Choose an arbitrary feature that is 'activated' by the topic, and

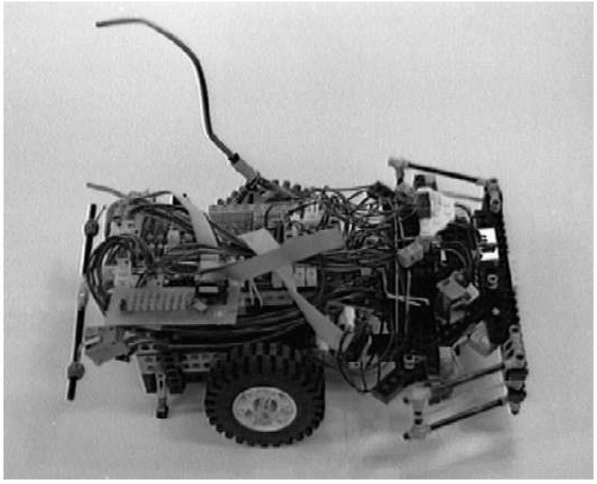

Fig. 2. The LEGO robots used at the AI-Lab in Brussels.

divide the range of this feature in two equal halves, thus constructing two new features. This way a binary tree of features is created (figure 1).

If the discrimination game was a success, the use- and success-scores are updated as follows: The use parameter is incremented for all distinctive feature sets. The success-score is incremented for the distinctive feature set, which represents the object the best. Self-organization is achieved by the way the agent selects a distinctive feature set in using these scores. So, learning to ground the meaning of objects is achieved by generating perceptual features, then trying to distinguish an object from others and increasing the strength of the feature sets that were successful in the discrimination task.

3 The experiments

3.1 The robots and their environment

At the AI-Lab in Brussels we have several mobile LEGO robots (figure 2) which can explore an ecosystem of approximately 4 by 6 meters (see figure 3). On the robots several sensor modules are mounted for sensing and two motors to drive them. The process unit of the robots is a Motorola MC68332 micro controller with 128 kB ROM and 256 KB RAM located on a Vesta board. The Vesta board is integrated with the SMB2 sensory-motor board, which is dedicated to low-level sensory-motor processing (Vereertbrugghen, 1996). The robots have two white light and two modulated light sensors, each on the left and the right front side of the robot. Three infrared (IR) sensors are mounted on the left-, middle- and right front side. Four IR emitters are mounted on the robot in such a way that they emit IR in four perpendicular directions (to the front, the back, the right

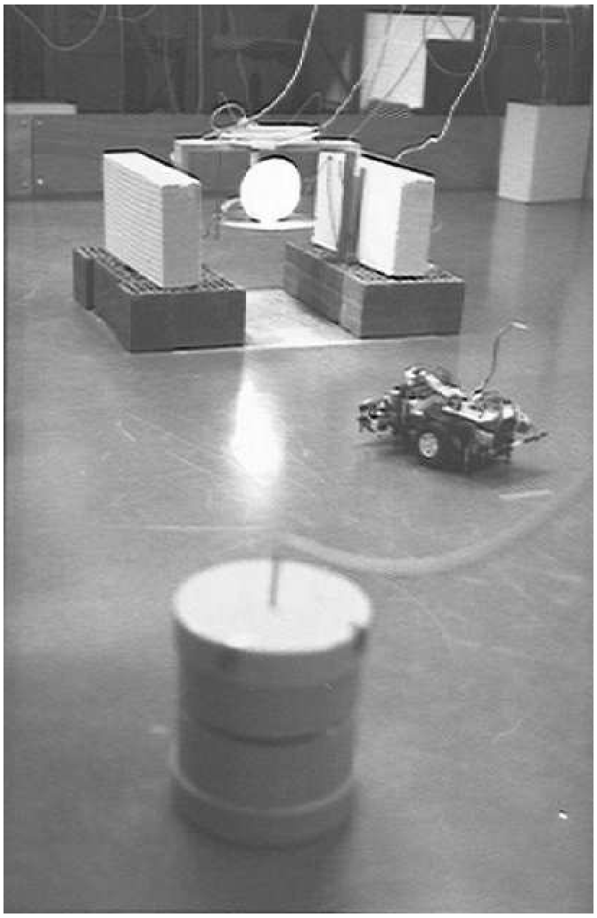

Fig. 3. The ecosystem of the AI-Lab in Brussels. In the foreground we see a competitor that emits modulated light, and in the background the charging station emitting white light is visible.

and the left sides). Four bumpers are used for obstacle avoidance. Secondary NiMH batteries supply the power (Birk, 1997).

Communication between the robots is done by a radio-link that can transmit and receive messages up to 40 Kbit/s. This radio-link is part of the sensory-motor board, and is therefore controlled by the processor of this board. The radio-link is unreliable in the sense that it is not guaranteed that a message arrives, but when it arrives the message contains no errors (Vereertbrugghen, 1996). The radio-link is used for both linguistic and extra-linguistic communication.

The robots are programmed in PDL, a language that is designed for behavior-oriented control as described in (Steels, 1994). The system directly couples sensory data streams to the actuators, thus controlling behavior dynamically. Behavior is implemented as a motivational coupling between sensory input and motor output.

In the robot's environment (see figure 3) there are several objects for which the robots need to ground a semantic representation. These objects can be described by the light source they emit. There is one light source that emits white light, there are two light sources that emit light modulated at a particular frequency, and the robots emit infrared to make themselves visible for the other robot. The whole system, i.e. robots and environment, is used for the meaning and language experiments described in this paper. In the next section this experiment is described in more detail.

3.2 The language experiments

In order to investigate the origins of language on physical robots we constructed an experiment in which robots engage in a series of *language games*. At first the robots have no knowledge of language at all; they only have knowledge about how to communicate (Steels and Vogt, 1997). In the experiment described here, there are only two robots that play a series of *naming games*. In naming games (a variant of the language games) the agents indicate objects by their name (Steels, 1996a). The objects about which the robots can communicate are, apart from the robots themselves, an object that emits white light and two objects that emit light modulated at a particular frequency.

In a language game, one robot decides to be the speaker, while the other becomes the hearer. In order to construct a coherent context, both robots must know what objects are present in their near surroundings and where these objects are. This can be done while the robots are standing close to each other and facing one another. So, when the robots are facing each other at close distance they have to build a map of their surroundings, thus constructing a context of objects. The speaker chooses one object randomly from the context as the topic. Then the speaker uses extra-linguistic communication (as described later) in order to let the hearer identify this topic as well. For this object each agent constructs a set of distinctive feature sets in a discrimination game as described in section 2. Now that both agents have a set of distinctive feature sets they can start the communication as described in more detail in (Steels and Vogt, 1997)

The experiment is implemented so that the robots can be in three different modes: (1) default exploration mode, (2) the speaker mode, and (3) the hearer mode. In the default exploration mode, the robots drive around in the ecosystem in forward direction with touch-based obstacle avoidance. Each robot emits pulses of IR, so that they can detect the presence of another robot. When no IR is emitted, but the robot senses IR, it can infer that there must be another robot nearby. The robot that detects the other then broadcasts a radio message that it wants to communicate. The other agent then replies a confirmation to the first one, and enters the hearer mode. When the first robot receives the confirmation, it enters the speaker mode. This type of communication is extra-linguistic, so it does not add to the language being formed.

The hearer starts to emit a continuous signal of IR, and the speaker uses IR-taxis (or phototaxis towards infrared) to drive towards the hearer as described by (Braitenberg, 1984). The motor values are adjusted so the robot drives towards the IR source. When the speaker is close to the hearer, it stops and uses *IR-orientation* to orient towards the hearer. With IR-orientation, as with IR-taxis, the robot orients towards the highest gradient of difference between the outer IR sensors, but without forward drive. After the speaker has completed this, the hearer orients towards the speaker, so they now stand in front of each other.

One by one the speaker and the hearer start to scan their surroundings. They do so by rotating around their axis while recording their sensory input, thus constructing a map like in figure 4. As we can see in figure 4(a), the complementary sensor pairs show intersections at, for instance, time/position 100, 140 and 160. The first one is an intersection of modulated light; the second one is from white light. IR intersects at the beginning and the end of the graph. The intersections occur when the robot passes an object, so this information could be used to decide that an object is at that particular place. If such an intersection is detected, the robot makes a record of all sensory channels. These values may be direct sensor values or real valued functions of direct sensor values. The construction of such sensory channels is very important for the success of the discrimination games. This construction will be discussed in the next section.

If the speaker has chosen a topic, it has to use extra-linguistic means in order to let the hearer identify the topic as well. We have decided to use pointing as method for this process. Pointing is implemented as follows: The speaker orients towards the topic, while it emits IR in four perpendicular directions. The hearer observes the pointing and perceives a graph as shown in figure 5. In this figure, one can observe three relative high peaks of IR. Each peak is at the boundary of a particular quadrant. Counting these peaks, the hearer can determine the quadrant in which the topic must be. Because there are not so many objects in the environment, the hearer can determine the topic with reasonable reliability. When the speaker does not rotate, but it does emit IR, the hearer is the topic. When the speaker does not emit IR nor does it rotate, then the speaker is the topic.

The figure, however, represents a graph where the speaker was exactly facing the hearer before it started to rotate. The experiments showed that this method was very unreliable (Steels and Vogt, 1997). The hearer often observed a very different graph than shown and was therefore not able to identify the same topic the speaker pointed at. In the experiments discussed here, this pointing was simulated by means of radio communication. The speaker transferred the quadrant in which the topic was observed to the hearer by radio link. The hearer then maps this quadrant to its own topic. For this mapping, the hearer must first mirror the received quadrant to a quadrant from its own point of view.

When both agents identified the topic and related features to sensory channels for all objects in their surroundings, they can play a discrimination game. The yielded set of distinctive feature sets is then used for the language formation (as described in (Steels and Vogt, 1997)).

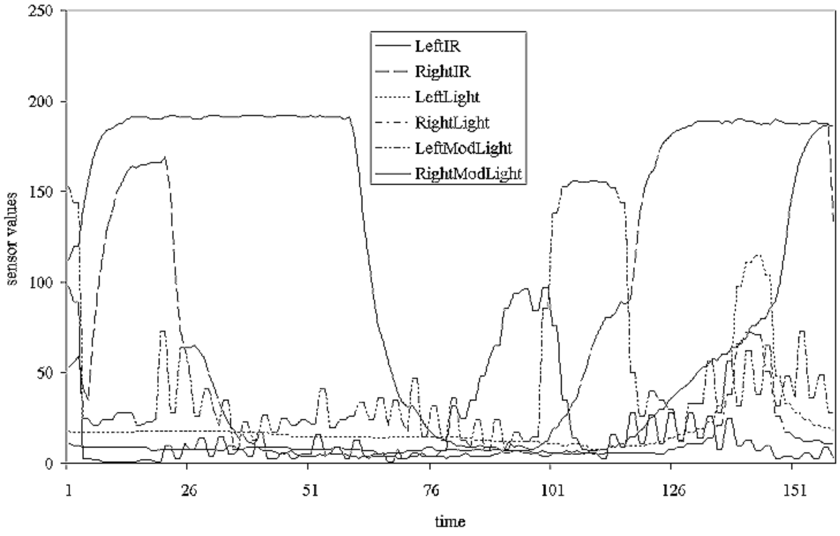

(a)

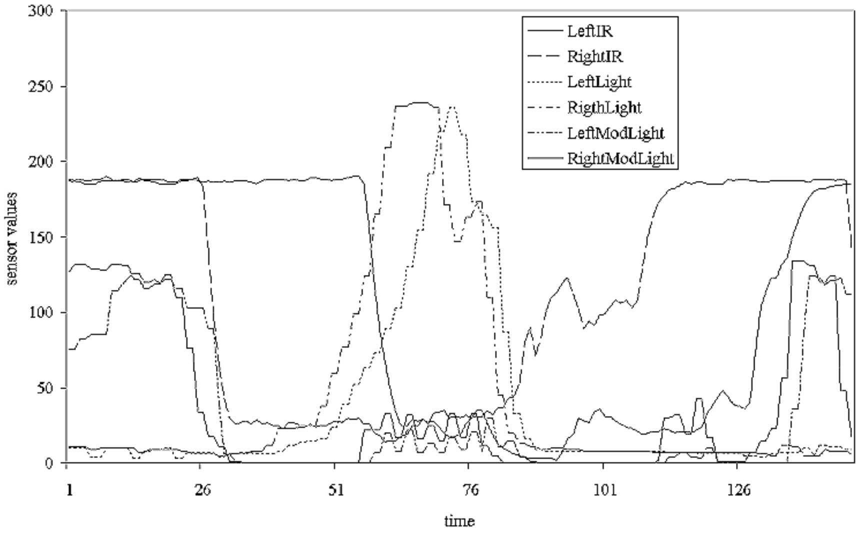

(b)

Fig. 4. The graph that two opposing robots may sense during the perception. The robots infer that there is an object when two complementary sensors intersect the graph.

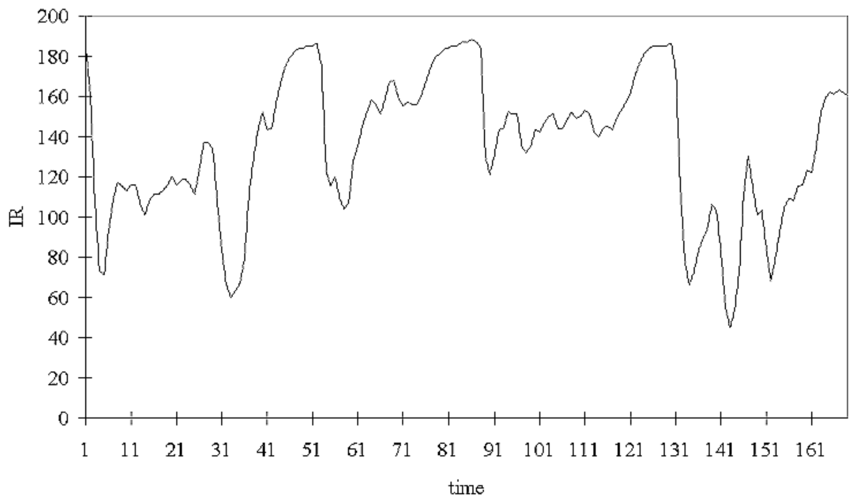

Fig. 5. The graph that the hearer observes when the speaker is rotating 360 degrees, while it is emitting IR in four perpendicular directions.

The speaker and hearer modes can be specified using Finite State Automata where the agents use the radio-link to synchronize the states in which the automata should be, if necessary. Hereto they send messages for asking to communicate, to say that they have aligned etc. The next section reports on the results of the meaning creation and the main features that were important for the successful implementation of the mechanisms.

4 Experimental results

We have done several experiments on the meaning and language formation. This section reports on only one of these experiments. The experiment resulted in 145 naming games, which were the result of two days experimenting. This low amount is due to the following reasons: (1) The robots can only work for approximately 30 minutes when the batteries are charged, and they can complete the process of one naming game roughly every 2 minutes. (2) The internal states (or memory) of the robots and the discourse of the naming games in both robots were monitored by a serial cable. This was extremely difficult since we have two robots, each connected to a cable, that move in an environment full of obstacles. And (3) producing much more games appeared not to be possible because after more games the internal memory started to exceed the 256 kB RAM. This, of course, is the reason why there are not more than 145 games.

As we shall see, the grounding of distinctive features was successful in that the robots revealed distinctive feature sets, which represented the objects rather coherently. This already happened after a short while of learning. The language formation was less clear, because the formation of a language takes much more time. Reports on the language formation can be found in (Steels and Vogt, 1997). This section focuses on the meaning creation.

As was mentioned before, the construction of sensory channels was very important in the development of this experiment. The right choice of these channels must enable the agents to select the right features in order to discriminate the objects coherently. I.e. the selected features, which compose the distinctive feature set(s), must be general enough to be selected in other situations as well. The task for the robots was to ground three objects[1], which are made visible by means of white light, light modulated at a certain frequency and IR. Therefore three sensory channels were constructed for discriminating the objects: *sc0 - white light, sc1 - modulated light,* and *sc2 - IR.* All these channels produce values as if the object is seen in the front.

Consider the first discrimination game, where the robot has no features yet. The discrimination game was executed by robot r2, which scanned the object o0 at time/position 67 and object o1 at position 102. The sensory channels for o0 revealed values 69 for sc0, 54 for sc1 and 37 for sc2, for o1 this was 13, 2 and 189 respectively. Obviously, the discrimination ended in failure because there were no features. A new feature is created for sc0, which expects a positive value between 0 and 255.

Discrimination game by r2
Objects r2:
  o0 [67] [69, 54, 37]
  o1 [102] [13,2,189]
Topic r2: o0
Failure r2. No feature sets.
New features r2: r2-sc0 [0,255]

Let us now look at another early discrimination game that fails, but where a feature is divided in two equal halves creating two new features.

Discrimination game by r2
Objects r2:
  o0 [64] [47, 1, 9]
  o1 [103] [76, 46, 182]
  o2 [105] [90, 73, 185]
Topic r2: o0. Feature sets:
  o0 {r2-sc0,r2-sc0-0,r2-sc1,r2-sc2}
  o1 {r2-sc0,r2-sc0-0,r2-sc1,r2-sc2}
  o2 {r2-sc0,r2-sc0-0,r2-sc1,r2-sc2}
Failure r2. No distinctive feature sets.
New features r2: r2-sc1-0 [0,127.5] r2-sc1-1 [127.5,255]

[1] The agents need not ground the feature *self*, because this is not an observable feature. It is assumed that this feature is true in every context at time/position 0.

Now we see a discrimination game that was successful. Note that the context is the same as in the previous discrimination game. The robots play several language- and discrimination with the same context (with a maximum of 10).

Discrimination game by r2
Objects r2:
  o0 [64] [47, 1, 9]
  o1 [103] [76, 46, 182]
  o2 [105] [90, 73, 185]
Topic r2: o2. Feature sets:
  o0 {*r2-sc0,r2-sc0-0,r2-sc0-0-0,r2-sc1,r2-sc1-0,*
    r2-sc1-0-0,r2-sc2,r2-sc2-0}
  o1 {*r2-sc0,r2-sc0-0,r2-sc0-0-1,r2-sc1,r2-sc1-0,*
    r2-sc1-0-0,r2-sc2,r2-sc2-1}
  o2 {*r2-sc0,r2-sc0-0,r2-sc0-0-1,r2-sc1,r2-sc1-0,*
    r2-sc1-0-1,r2-sc2,r2-sc2-1}
Distinctive feature sets r2:
  {{*r2-sc0,r2-sc1-0-1,r2-sc2*},
   {*r2-sc-0,r2-sc1-0-1,r2-sc2-1*},
   {*r2-sc0,r2-sc1-0-1*},
   {*r2-sc1-0-1,r2-sc2*},
   {*r2-sc0-0,r2-sc1-0-1*},
   {*sc1-0-1,r2-sc2-1*},
   {*r2-sc0-0-1,r2-sc1-0-1*},
   {*r2-sc1-0-1*}}
Success r2. {*r2-sc1-0-1*} *[63.75,127.5]*

The 10th discrimination game already showed the above result. Robot r2 finds all combinations of features that include *sc1-0-1*. This feature is the only one that distinguishes o2 from o0 and o1. The game yields {*r2-sc1-0-1*} on the range *[63.75,127.5]*, because this set is the smallest one, which is the preferred condition when there are several distinctive feature sets. The use of all features is increased, whereas the success of only {*r2-sc1-0-1*} is incremented.

This last game illustrates an almost typical view of what the robot observes when it is scanning its surroundings. All sensory channels always yield a noise value, and it is very difficult to say what object in the environment is meant. In order to detect what object is meant, a monitoring program keeps track of what sensory channel is used to decide when the robot observes an object during the perception. In this case the object was recorded because sc2 showed an intersection, therefore o2 is the other robot.

As can be seen in figure 6, discriminative success of the successful distinctive feature sets increase quite rapidly. This is because in the first few discrimination games these sets were used, they were immediately successful. In later games some of these distinctive feature sets were used in a discrimination game, but these were not preferred, thus decreasing the success rate. Some feature sets

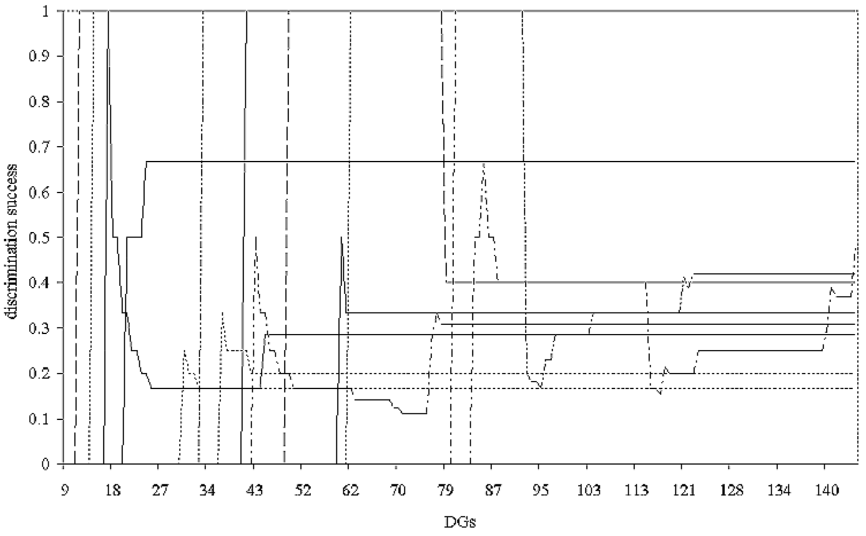

Fig. 6. The discriminative success, i.e. *success/use* of some successful feature sets for the first 145 discrimination games. In the beginning all feature sets are built up with increasing success scores. Some of the score decrease in later games, which is due to competition between distinctive feature sets.

reach a score of 1 immediately. Most of these sets are used only one time after 145 discrimination games.

If we look for which objects the distinctive feature sets are used, we can see that every object has a rather coherent set of distinctive feature sets (table 1). The white light object is recognized with only 16.7% singular distinctive feature sets, but the modulated light objects with 75.0% and the infrared with 72.0%. These figures are, however, taken from only 145 discrimination games, and may therefore not be very significant. Although there is some overlap, most of these distinctive feature sets reveal clear preference for a particular object. To have some overlap and no full coherence in the meaning is likely. It is thought that meaning in the mind is represented by fuzzy sets that reveal a family resemblance (Aitchison, 1994), so that underlying features may reveal some overlap in representing particular objects. Although the experiments have not yet revealed this, it is also thought that the co-evolution of language with meaning will cause all feature sets that now represent one real world object to be associated with one *global* linguistic meaning. So there is normally only one word representing one object, conform (Steels, 1997).

Concluding, we saw that robotic agents were able to learn to discriminate objects using sensory features. The agents build up their own tree of features, and are able to use distinctions in order to recognize the objects. A mechanism of selection for the best used features causes the robot to learn the feature-tree

White light use: 12	Modulated light use: 20	Infrared use: 50
{sc0-0-0}:2	{sc1-0-1}:1	{sc0-0-2}:1
{sc1-0-0}:1	{sc1-0-2}:1	{sc0-1-0}:2
{sc2-1,sc1-0-1}:1	{sc2-0-1}:2	{sc0-0-3,sc1-0-4}:1
{sc2-0-0}:1	{sc2-0-0}:3	{sc1-0-2}:5
{sc2-1-0}:4	{sc1-0-3}:1	{sc1-1-0}:2
	{sc2-0-3}:2	{sc0-0-4}:1
	{sc2-0-4}:2	{sc0-0-5}:9
	{sc2-1,sc1-0-2}:1	{sc2-1,sc1-0-0}:2
	{sc1-0-4}:1	{sc2}:16
	{sc0-0-1,sc2-1-1}:1	{sc2-1,sc0-0-6}:1
	{sc1-0-5}:1	{sc2-1-0}:3
	{sc2-1-1}:3	{sc0-0-0}:5
		{sc0-0-7,sc2-1-1}:1

Table 1. The coherence of distinctive feature sets vs. the objects. The number behind the distinctive feature sets indicates the successfulness of the set for the particular object. Note that the notation of the sensory channels is different than before in order to save some space. Feature *sc1-0-0* means now the lower half of sensory channel 1, with identifier 0, *sc2-1-1* is sensitive on the upper half of sensory channel 2, with identifier 1, etc.

in a self-organizing way. Although it has not been shown here, it should be clear that each agent constructs its own set of features, and therefore the agents may have different representations of these features in their 'mind'.

5 Conclusions and future research

This paper discussed the implementation and the results of experiments that has been done on learning to ground semantic features in robotic agents. The experiments were held in combination with language formation experiments, which are reported in (Steels and Vogt, 1997).

We saw that robots were able to construct a feature tree in order to discriminate and categorize sensory information from objects into symbolic feature sets. This categorization was made possible by implementing a self-organizing learning mechanism to create and select features adaptively according to the rules given in section 2. The robots could learn which features were useful to discriminate between objects, thus building coherent classes of feature sets in order to represent these (aspects of) objects, so a lexicon could be formed using these items as semantic features. The feature sets have been grounded completely in distinctions. Real identification of objects (or concepts) though would be another step, which according to Harnad is a necessary, but extremely difficult step in the grounding problem (Harnad, 1990).

Although the results are promising, a lot of research still has to be done. The sets of discriminative features that are given in table 1 must be classified higher up in the hierarchy of meaning. This can be done by the language formation, where agents construct word-meaning pairs that represent discriminative feature sets. It is shown (Steels and Vogt, 1997), that the agents construct ambiguous words (i.e. words with more than one meaning) in the sense that they represent different feature sets. This ambiguity though, is not necessarily visible in the physical world outside an agent, because the different features may represent the same object (see table 1). Therefore we could call this kind of ambiguity a *representational ambiguity*, as opposed to lexical ambiguity and could overcome the identification problem. This principle representational ambiguity, however, has not revealed in the experiments sufficiently, because the formation of a coherent lexicon may take more than a thousand language games. Due to technical problems experiments at such a scale were not yet executed.

Currently we are working on the improvement of these robotic experiments, so we will be able to execute these large-scale experiments. We are also preparing experiments to ground and lexicalize spatial relations, internal states and actions. Work is being carried out in order to implement grounded language formation using active vision.

Concluding, we could say that this experiment might be an important step towards a new theory of meaning and language, where the formation is based on a natural selection-like principle. The basic building blocks of these languages are not innate structures (conform (Chomsky, 1980)), but rather the result of (1) life-long evolution inside an agent based on its interactions with its environment (Foley, 1997) and (2) the social interaction between different agents in a group.

6 Acknowledgements

This experiment was done as a project for my master's thesis in Cognitive Science and Engineering at the University of Groningen. It was carried out as a practical period at the VUB Artificial Intelligence Laboratory in Brussels.

I would like to thank my supervisors Luc Steels (VUB) and Petra Hendriks (Univ. of Groningen) for their support and advice during the project. Furthermore I would like to thank the members of the AI-Lab for their support: Tony Belpaeme, Andreas Birk, Bart de Boer, Edwin de Jong Peter Stuer and Dany Vereertbrugghen for advice and support.

The project was partly financed by an EU Erasmus fellowship.

References

Aitchison, J. (1994). *Words in the mind: an introduction to the mental lexicon.* Blackwell Publishers, Cambridge Ma., second edition.

Birk, A. (1997). Autonomous recharging of mobile robots. In *Proceedings of ISATA97, Dedicated Conference on Robotics, Motion and Machine Vision.*

Braitenberg, V. (1984). *Vehicles, Experiments in Synthetic Psychology.* The MIT Press, Cambridge MA.

Chomsky, N. (1980). Rules and representations. *The behavioral and brain sciences*, 3, pp. 1-61.

Edelman, G. (1987). *Neural Darwinism.* Basic Books Inc., New York.

Foley, W. (1997). *Anthropological linguistics : an introduction.* Blackwell, Cambridge Ma.

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

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

Steels, L. (1994). Artificial life route to artificial intelligence. *Arificial Life Journal*, 1(1).

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. (1996c). The spontaneous self-organization of an adaptive language. In Muggleton, S., editor, *Machine Intelligence 15*, Oxford. Oxford University Press.

Steels, L. (1997). Synthesising the origins of language and meaning using coevolution, self-organisation and level formation. In Hurford J., C. K. and Studdert-Kennedy, M., editors, *Evolution of Human Language*, Edingburgh. Edingburgh University Press.

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.

Vereertbrugghen, D. (1996). Design and implementation of a second generation sensor-motor control unit for mobile robots. Master's thesis, VUB.