Grounding adaptive language games in robotic agents

Luc Steels (1,2) and Paul Vogt (1)
(1) VUB AI Laboratory, Pleinlaan 2, 1050 Brussels,
steels@arti.vub.ac.be, paul@arti.vub.ac.be
(2) SONY Computer Science Laboratory
6 Rue Amyot, Paris
In: Harvey, P. and P. Husbands (eds.)
Proceedings of European Conference on Artificial Life.
The MIT Press, Cambridge Ma

Abstract

The paper addresses the question how a group of physically embodied robotic agents may originate meaning and language through adaptive language games. The main principles underlying the approach are sketched as well as the steps needed to implement these principles on physical agents. Some experimental results based on this implementation are presented.

1 Introduction

In the past five years, a large number of robotic agents, i.e. physical systems capable of sensori-motor control, have been built in order to investigate a bottom-up approach to artificial intelligence (see the overview in [8]). Important results have been achieved, particularly by using behavior-oriented architectures [14] and learning methods based on neural networks [6] or genetic algorithms [3]. Nevertheless, it is still largely an open question how these robots may reach sufficient complexity in order to qualify as cognitive agents. Most of the experiments have focused on 'low level' tasks like obstacle avoidance or navigation, and these have been difficult enough to preclude any work on cognitive tasks.

One approach for pushing ahead, taken for example in the COG project[1], is to increase the complexity of the robots themselves by adding many more sensory channels and many more degrees of freedom. Another approach, which we are exploring, is based on the hypothesis that communication, if not full-fledged language, is a necessary stepping stone towards cognitive intelligence. This implies that we cannot restrict ourselves to individual robots but must perform experiments how groups of robots may build up communication systems of increased complexity. In the spirit of the bottom-up approach, these communication systems must be developed by the robots themselves and not designed and programmed in by an external observer. They must also be grounded in the sensori-motor experiences of the robot as opposed to being disembodied, with the input given by a human experimenter and the output again interpreted by the human observer.

Some initial experiments have been reported in the Alife literature on how communication itself may arise to aid cooperation between agents [5],[15]. In this paper, we assume that there is already communication and focus instead on the grounding problem, as in [18]: How the evolving language is anchored into the sensory and motor data streams generated through normal behavior. We also address the problem of the origin of meaning: How the distinctions that the robots lexicalise may arise in the first place.

The work reported here builds further on earlier software experiments that show how agents may develop a shared vocabulary through a series of adaptive naming games [9] and how agents may generate distinctions to discriminate between objects in their environment [10],[12]. These papers can be consulted for formal descriptions of the mechanisms. This paper focuses in particular on how the software experiments have been carried to real robots.

The rest of the paper is in three parts. The first part explains the adaptive language games including the mechanisms that cause the build up of distinctions and of lexicons to express these distinctions. The second part discusses how adaptive language games have been mapped onto physical robots. The third part gives some results of concrete experiments. Conclusions and ideas for future research end the paper.

2 Adaptive Language Games

At the heart of our approach is the notion of a language game [17]. A language game involves two agents, a speaker and a hearer, as well as a context which consists of agents, objects and situations. Different kinds of language games can be played depending on the goals

that the participating agents want to achieve. The game being pursued in the experiments reported here is for the speaker to identify an object in a certain context using linguistic means. We call this game the naming game. Initially extra-linguistic means, such as pointing, can be used to bootstrap the language. Other language games would allow the speaker to get the hearer to perform a certain action, to ask the hearer for more information, etc.

2.1 The basic scenario

To play a naming game both participants follow a specific scenario, which consists of the following six steps:

1. *Making Contact*: Two agents must make contact with each other. One assumes the role of speaker, the other of hearer. The agents are physically close together so that there is automatically a shared context.

2. *Topic identification*: Each agent perceives the surrounding environment through its sensors and identifies a set of objects which constitute the context. The speaking agent chooses one object in this context as the topic of the conversation. He then draws attention to this topic using extra-linguistic means, for example by pointing. The hearer thus also identifies the topic.

3. *Perception*: Each agent then categorises the sensory experience of the different objects in terms of features, and identifies a distinctive feature set which distinguishes the topic from the other objects in the context. It will often be the case that more than one distinctive feature set is appropriate.

4. *Encoding*: The speaker chooses one distinctive feature set (for example the smallest one) and encodes this into an expression. Encoding means that the smallest set of words, which expresses all the features in the distinctive feature set, is searched for in the lexicon.

5. *Decoding*: The hearer decodes the expression which means that he looks up all the words in his lexicon and reassembles a feature set covering all the words. Words are ambiguous in the lexicon (the same word may have different meanings), so that there is typically more than one possible feature set resulting from the decoding process.

6. *Feedback*: The hearer compares the decoded feature sets with the distinctive feature sets that he was expecting. If one of the distinctive feature sets is equal to the decoded feature set, the language games ends in success and the hearer gives a positive feedback. Otherwise the game ends in failure and the hearer signals failure.

This scenario assumes that (1) both agents have a perceptual apparatus for categorising sensory experiences and identifying distinctive feature sets and (2) a lexicon that associates features or feature sets with words and vice-versa. However we are precisely interested in the problem how (1) and (2) may originate. Initially the agents have no repertoire of perceptual distinctions and no lexicon. They build these up as a side activity of each language game using the methods described in the following two subsections.

2.2 Originating distinctions

Each agent has a series of sensory-motor channels which are the direct output of sensors, the result of automatic low-level sensory processes, or the dynamically evolving contents of internal states such as left and right motor command streams. These sensory-motor channels are given by the hardware or low-level routines. For each channel there is a discrimination tree which divides the output of a channel into distinct regions. It is assumed that the discrimination trees are binary. Each end-node of a tree constitutes a feature. The feature is denoted by a string *agent-channel-region-subregion-subsubregion-...*, as in a1-s0-0-1, which refers to a feature associated with channel s0 in agent a1. Initially there are no discrimination trees.

As part of the perception phase, the agent engages in a discrimination game. He categorises the sensori-motor states for each object based on his discrimination trees. The result is a set of features for each channel that contains active data, and this for each object. The different sets are then used to find the possible distinctive feature sets that distinguish the feature set of the topic from the feature sets of the other objects. If this fails, i.e. if no distinctive feature set can be built using the existing discrimination trees, a new distinction is created by a further subdivision of one of the end-nodes of a discrimination tree which was active in the categorisation process. The choice which of these nodes is expanded is arbitrary. The agent keeps track of which features are used and the success in discrimination. A forgetting process eliminates those end-nodes which turn out not to be useful.

Thus our approach is selectionist (as in [2]): There is a generator of diversity and a separate selectionist process which maintains or eliminates features from the feature population. Earlier software experiments [10] have shown that this method stabilises on a successful repertoire of discriminations. Moreover new objects, new sensori-motor channels, or new agents may at any time enter, causing the discrimination trees to be expanded and adapted as the need arises. Note that each agent builds up his own discrimination trees. There are similarities due to the fact that the agents operate in the same environment but this does not guarantee complete

coherence. More coherence is reached when the lexicalisation of a feature is an additional selectionist criterion for its further survival, as discussed in more detail in [12].

2.3 Originating a lexicon

A lexicon consists of a set of word-meaning pairs, where the meaning consists of a feature set. Each agent has his own lexicon and an agent cannot directly inspect the lexicon of another one. Each agent maintains how often a word-meaning pair has been used and how successful it has been in its use. While encoding, a speaker will prefer word-meaning pairs that have been used more often and were more succesful in use.

A discrimination game results in a series of possible distinctive feature sets of which one is chosen by the speaker. This feature set is encoded by the speaker and then decoded by the hearer. Several things can go wrong in this process and each failure results in appropriate actions:

1. *The speaker does not have a word for a certain feature set.* In this case, the speaker is allowed to construct a new word (formed by a random combination drawn from a given prior alphabet) and associate that in his lexicon with the feature set. This happens with a low probability because a word may already exist in the population for this feature set.

2. *The hearer may lack a word used by the speaker.* In this case, the hearer can infer possible feature sets that might be meant by that word, based on the distinctive feature sets that he is expecting. In the simplest situation, there is only one feature necessary to distinguish the topic from the objects, so that the meaning is unequivocally known. It could also be that some words are known but not others. The meaning of the missing words must then be reconstructed from the remaining unknowns. Because there may be more than one distinctive feature set, it is inevitable that ambiguity creeps into the lexicon of the hearer. These ambiguities are weeded out by future use and success in use which determine what word-meaning pairs will become most common.

3. *Some of the feature sets decoded by the hearer do not match with the expected distinctive feature sets.* This means that there are some word-meaning pairs which are not shared by some of the agents. For the successful word-meaning pairs, both success and use is incremented, whereas for the others only the use is incremented, so that their future use diminishes.

4. *The feature set decoded by the hearer does not match with any of the expected distinctive feature sets.* In that case, the hearer extends the lexicon, using the same procedure as for situation 2 above.

Note that the approach is again selectionist. Agents create or infer word-meaning pairs. Which pairs 'survive' depends on use and success in use, and this is determined by how many agents have adopted the same word-meaning pairs. Typically we see a phase transition when one word starts to dominate for the expression of a particular meaning. This phase transition is due to the positive feedback loop inherent in the system: The more a word is used, the more success it will have in use and the more it will be used even more. Software simulations reported in [9] have shown that a group of agents indeed converges towards a common lexicon after a sufficient number of adaptive naming games. Moreover new agents may enter at any time, and due to the adaptive nature of the discrimination games, new features may enter the repertoire of possible meanings.

Given these results we now turn to the challenge of implementing these algorithms on physically embodied robots.

3 Physical implementation

As is well known by now, software simulations do not at all guarantee that the methods will also work in real world settings. Indeed, the problems encountered during the physical implementation of the language games have been enormous. Robots are basically parallel distributed computer systems which operate in real-time and whose communication links are very unreliable. We must therefore achieve overall reliability despite unreliable components and processes. Second we must have sufficiently robust and autonomous robots (also autonomous in terms of energy) to permit hundreds, and even thousands, of consecutive language games. Next, we must find equivalents of all the different steps in the scenario: Robots must be able to recognise each other, approach each other, and establish the necessary contact to start a language game. They must be able to point or in other ways draw attention to the topic. Their perceptual capabilities must be the basis of the discrimination games and finally they must realise the language games themselves. In addition, it remained to be seen whether the proposed discrimination mechanisms were adequate for handling the inherently noisy real world data coming from actual sensors and whether the lexicon would stabilise despite possible (and actual) failures at all steps of a game.

3.1 The robots and the ecosystem

The robots used in the experiments are Lego-vehicles built for our laboratory's experiments in self-sufficient robots (see figure 1) [7]. Each robot (size: 30 x 20 x 15 cm) has three infra-red sensors (mounted on the left-front, middle-front and right-front side), four infrared emitters (mounted on front, left, right, and back side), two visible light sensors (mounted on left- and right-front

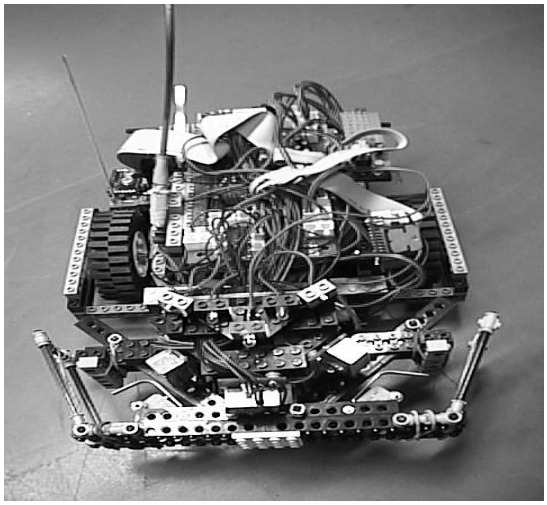

Figure 1: The robots used in the experiment are Lego vehicles. which are autonomous with respect to sensing, actuating, computation, and energy.

side), two modulated light sensors (mounted on left- and right-front side), various touch sensors mounted on all sides, and a battery sensor. There is a left and right motor. The overall processing capacity resides in a Motorala MC86332 micro controller with 128 kB ROM and 256 kB RAM located on a Vesta board. Its CPU is 16.78 MHz at 5V. The Vesta board is extended with a second board dedicated to low level sensory-motor processing and buffering [16].

The robots are programmed using a behavior-oriented architecture [7]. The sensors, actuators and internal states constitute continuous data streams and the behavior is based on continous dynamical systems implementing direct couplings between sensors and actuators. An example of such a coupling realises photo-taxis by minimising the difference between the left and right visible photosensors, as in Braitenberg vehicles. The couplings are modulated by motivational states. Thus the photo-taxis is modulated by a decreasing battery level, so that the robot drives towards the charging station when its energy resources are getting low.

The robots are equiped with a radio-link that is designed for communication among themselves at a reasonable speed, and for central monitoring of internal states. It is a module that extends the sensory-motor board. It has a build in power supply, a transmission and reception module, and an antenna. The module can transmit and receive messages up to 40 Kbit/s [16]. This radio-link is used for some of the extra-linguistic exchanges, as well as the linguistic communication itself. 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.

The robots are located in an ecosystem which contains a charging station on which a visible light source is located. Robots can recharge their batteries by sliding into the charging station. There are also 'competitors' in the environment in the form of boxes in which a (modulated) light source is mounted. This light source takes energy from the global energy flowing into the ecosystem. Robots can dim a light by pushing against its box and thus assure that there is enough energy in the charging station. After being dimmed, the light source regenerates, thus requiring the robots to alternate between recharging and work. The biological motivation for this setup is explained in [4].

We now turn to the physical implementation of the different steps in the language game scenario. The objects that can be the topic of a conversation are: obstacles, the robot itself, other robots, the charging station, and the competitors.

3.2 Making Contact

The robot can be in three modes: Regular exploration, being the speaker, and being the hearer. Any robot which is in the first mode may at any time randomly decide to become a speaker, when he 'sees' another robot in the environment. The robots used in these experiments do not have vision. They can however recognise each other because each robot emits infrared as part of its obstacle avoidance behavior. This infrared light is modulated so that the infrared of one robot does not confuse the infrared of another one. A robot detects another one when there is an infrared source which is not his own (see figure 2).

A robot which has adopted a speaker mode and which detects a possible hearer in the environment emits a request for entering into communication. On receiving this request, the other robot may switch from an exploration mode to a hearer mode. The hearer confirms that he wants to play a hearer role and halts while continuing to emit infrared. On receiving the confirmation, the speaker switches off its infrared and uses infrared-taxis to approach the hearer. Infrared-taxis means that the speaker moves up the infrared gradient as shown in figure 3. Movement stops when the gradient starts to fall off. The speaker broadcasts an "aligned" signal and turns on its infrared.

On receiving the alignment signal, the hearer also tries to position himself so that he faces the speaker. He turns off his own infrared emission and performs the infrared-orientation behavior while not moving forward. When maximum infrared is detected, the hearer emits in turn an "aligned" signal. The speaker turns off its infrared emission. The two robots are now in a situation as depicted in figure 4. They are facing each other and ready

Figure 4: Two robots have approached each other and are now facing each other. Note the other objects in the environment surrounding the robots, which will be the subject of the conversation.

for starting a language game.

3.3 Topic Identification

The next problem is how both robots could get a shared perception of the environment. This has been handled as follows: The speaker and the hearer take turns in scanning the environment by making a 360 degree turn. During this scan all the sensory data are recorded giving a panoramic view as shown in figure 5. There is no direct sensing of the degree of turning. The robot recognises that he has turned 360 degrees when the same sensory data are perceived as at the start of turning. The time dimension is later used as a spatial dimension.

The next important issue is what counts as an object. The robot has no explicit notion of an object and no sophisticated visual sensing that could detect an object by matching it against a background for example. We notice that the robot is facing an object precisely at the point where two sensors of the same type (for example left and right visible light sensors) cross each other, simply because sensors come in pairs and are mounted on each side. Consequently these crossings are taken to be the positions of the object and the states of *all* sensory streams at those points will play a role in formulating a distinctive feature set to categorise the object. For example, another robot will not only be recognisable because he emits infrared light, but also because he reflects visible light, although less than the charging station.

Through this procedure, each robot constructs a series

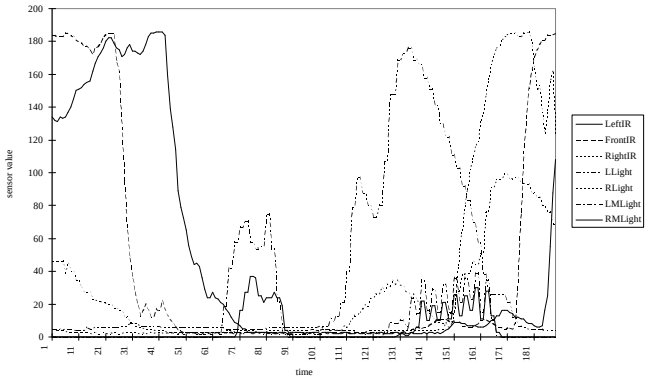

Figure 5: The result of a 360 degree scan for a single robot, in which data from 7 sensory input streams are recorded (no significant data appear on touch sensing or energy sensing).

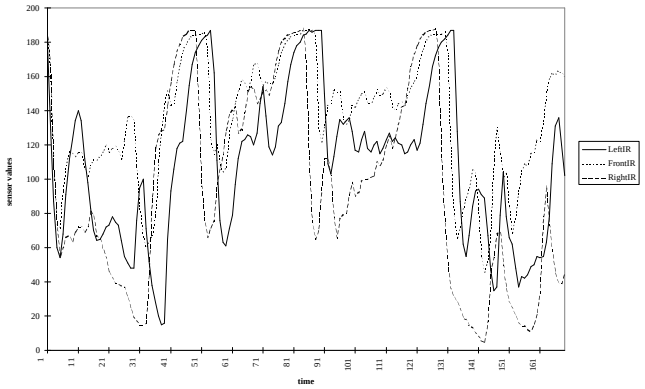

Figure 6: The speaker points towards the topic. The figure shows the infrared detection by the hearer. Each crossing of left and right IR sensors (beyond a certain threshold) indicates that one ray has passed. This happens around point 51, 85, and 131. Note that the data stream is also influenced by reflection from objects around the speaker and the hearer.

of objects and associated sensory data values. To this the robot adds himself as a possible topic of the conversation. The speaker then selects randomly one object from this list to be the topic of the conversation and proceeds by drawing the attention of the hearer to this object. This is again quite difficult to achieve because the robots have no physical device for pointing. We have opted for a procedure in which the speaker orients himself towards the topic. By convention, the speaker talks about himself when he does not engage in any movement for drawing attention to another object. The hearer can follow the turning and estimate the direction because each robot emits 4 infrared rays mounted on the front, left, back, and right side. Thus by counting the number of passing infrared rays, whose focal points are seen when left and right infrared is crossing, the quadrant in which the topic is located can be calculated (figure 5). For example, when three passing rays are measured, the speaker is pointing direction east which means that the topic is west of the hearer, i.e. to his left side.

Each of these various steps may (and does) go wrong. Sometimes one of the robots turns more than 360 degrees and loses track of its position. The hearer may not be able to detect well the turning of the speaker towards the topic and thus miss the topic. However the general success rate is high enough (about 75 %) to allow subsequent language games. At the moment we obtain 3 to 3,4 language games per minute.

3.4 Categorisation

As discussed in the previous subsection, the robots have a panoramic view of their environment and a list of objects with sensory states for each one. Moreover the topic of the conversation is now known by both robots. The next step is for each to derive a distinctive feature set which allows a discrimination of the topic from the other objects in the context. This proceeds along the lines outlined in section 2.1. The robots build up discrimination trees if there are not enough features to allow discrimination following the procedure describe in section 2.2.

3.5 Encoding, Decoding and Feedback

The encoding and decoding steps proceed exactly as outlined earlier in section 2.3. The result of encoding is transmitted through the radio link. The robots use random combinations of letters to form new words when needed. Feedback is based on the same procedure as outlined in section 2.1: When the distinctive feature set decoded by the hearer matches with an expected feature set for the topic, the language game succeeds otherwise it fails. The hearer provides feedback by a signal through the radiolink.

4 Results

We have conducted different experiments with the present implementation. Each experiment consists of a series of language games. The results of one such experiment are now reported.

First we look at the discrimination games between two robots r1 and r2. An object is detected at time/position 176 with the values 9 for channel s0, 0 for s1 and 192 for s2. The discrimination ends in failure but leads to the construction of a new feature detector which expects a positive value for channel 0 (i.e. a value between 0 and 255).

```
Discrimination game by r2
Objects r2:
  o1 [176] [s0:9,s1:0,s2:192]
Topic r2: o1
Failure r2. No feature sets.
New feature detectors r2: r2-s0 [0,255]
```

Here is another discrimination game when the build up of discriminators is already further advanced. Two objects are seen o1 and o2, with both positive values for s0 and s2. This is not enough to discriminate so a new feature detector is created by further subdividing channel 2.

```
Discrimination game by r2
Objects r2:
  o1 [151] [s0:1,s1:0,s2:59]
  o2 [217] [s0:7,s1:0,s2:3]
Topic r2: o1
Feature sets r2:
  o1 {r2-s0,r2-s2}
  o2 {r2-s0,r2-s2}
Failure r2. No distinctive feature sets.
New feature detectors r2: r2-s2-0 [0,127.5]
    r2-s2-1 [127.5,255]
```

Here is a discrimination game involving three objects which is successful:

```
Discrimination game by r2
Objects r2:
  o1 [45] [s0:8,s1:0,s2:5]
  o2 [58] [s0:4,s1:156,s2:2]
  o3 [166] [s0:6,s1:0,s2:187]
Topic r2: o2
Feature sets r2:
  o1 {r2-s0,r2-s2-0}
  o2 {r2-s0,r2-s1,r2-s2-0}
  o3 {r2-s0,r2-s2-1}
Distinctive feature sets r2:
  {{r2-s1}}
Success r2.
```

The set of features of r2 at this point is as follows. Each feature is followed by the range on the channel and the score (use and success) of the feature.

```
r2-s0 [0,255] 125/3
  r2-s0-0 [0,127.5] 111/0
  r2-s0-1 [127.5,255] 92/0
r2-s1 [0,255] 125/11
  r2-s1-0 [0,127.5] 91/0
  r2-s1-1 [127.5,255] 72/0
r2-s2 [0,255] 136/14
  r2-s2-0 [0,127.5] 72/0
  r2-s2-1 [127.5,255] 120/5
```

When games continue, there is further refinement and the features that are most useful increase their use and success, as can be seen from figure 7.

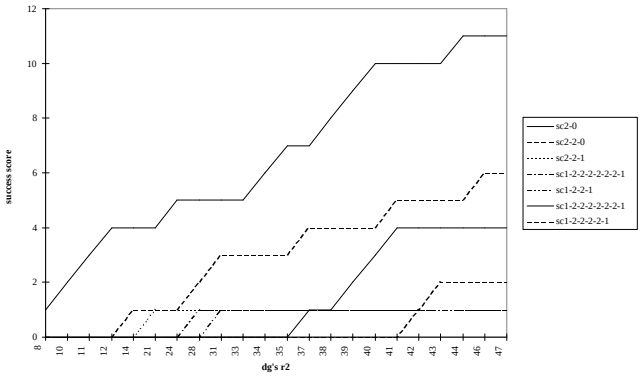

Figure 7: This figure plots in one robot the evolution in the success score of a feature over a period of 45 discrimination games. Features that are relevant in the environment gradually get a higher score.

We now look at the language games. An example of a complete successful language game (after 43 discrimination games and language games) is the following:

```
This is dialogue nr 43
Speaker: r2. Hearer: r1.
Objects r2:
  o1: [138] [s0:2,s1:0,sc-2:183]
Topic r2: self
Distinctive feature sets r2:
  {{r2-self}}
Objects r1:
  o1: [9] [s0:1,s1:0,s2:186]
  o2: [185] [s0:2,s1:12,s2:188]
Topic r1: o1
Distinctive feature sets r1:
  {{r1-s0,r1-s2-2-2-2-2-2-1},
   {r1-s1,r1-s2-2-2-2-2-2-1},
   {r1-s0-0,r1-s2-2-2-2-2-2-1},
   {r1-s1-0,r1-s2-2-2-2-2-2-1},
   {r1-s2-2-2-2-2-2-1},
   {r1-s2-2-2-2-2-2-1}}
Encoded expression r2: (a b)
Decoded expression r1:
  {{r1-self},{r1-s2-1},{r1-s2-2-0},
   {r1-s2-2-2-2},{r1-s2-2-2-2-2-1},
   {r1-s0},{r1-s2-2-2-2-2-2-1},
   {r1-s2-2-2-2-2-2-1}}
Success
```

The game ends in success because the feature sets decoded by r1 match with one of the distinctive feature sets r1 was expecting. The lexicon of r1, r2 are at this point as follows. The meaning, the word and the score (use/success) is printed out:

```
The lexicon of r1:
r1-self == (a b) 10/1
```

```
r1-s2-1 == (a b) 1/1
r1-s2-2-0 == (a b) 3/1
r1-s2-2-2-2 == (a b) 0/0
r1-s2-2-2-2-2-1 == (a b) 0/0
r1-s0 == (a b) 0/0
r1-s2-2-2-2-2-1 == (a b) 1/1
r1-self == (a c) 0/0
r1-s1-2-2-0 == (a d) 0/0

The lexicon of r2
r2-self == (a b) 14/3
r2-s2 == (a b) 2/0
r2-s2-1 == (a b) 2/0
r2-s1-0 == (a b) 4/0
r2-s2-2-1 == (a b) 0/0
r2-s1-2-0 == (a b) 2/1
r2-s1-2-2-0 == (a c) 1/0
```

We see that r2 uses "(a b)" for itself and r1 has coupled the same word to features it uses for recognising r2. r2 has coupled features for r1 to "(a c)" and this is also the name r1 has adopted for itself. Finally "(a d)" is being used as name for the competitors (the boxes in which a modulated infrared is housed).

Overall there is now a context coherence of 88.5% (the agents recognise the same context). The agents successfully recognised each other as the topic 31% of the time. Recognition of other objects was still low after 45 games but increasing.

5 Conclusions

The paper reports on experiments with physically embodied robotic agents which are relevant for two fundamental questions in the origins of cognition, namely (1) how can a set of perceptual categories (a grounded ontology) arise in an agent without the assistance of others and without having been programmed in (in other words not innately provided), and (2) how can a group of distributed agents which each develop their own ontology through interaction with the environment nevertheless develop a shared vocabulary by which they can communicate about their environment.

The proposed solution centers around coupled adaptive discrimination games and adaptive language games. Agents engage in interactions with the environment or with others and change their internal structure in order to be more successful in the next game. Both systems are selectionist: Structure is created by random processes and eliminated based on selectionist criteria centering around use and success in use.

Although we feel that this experiment represents an important milestone, there are obviously many things which can and should be done next, and some of this work is already going on in our laboratory. First, we have done other software simumations showing how spatial categories may become lexicalised [11]. These experiments are currently being ported to physical robots. Second, we are doing experiments in which vision is the primary source of sensory experiences. One of these experiments is based on two robotic heads that are located near the robotic ecosystem and give comments on the dynamically evolving scene they see before them. The use of vision allows for a much broader repertoire of objects and features and enables us to study how syntactic conventions may arise. The first results of this experiment are reported in [13]. Third, we are investigating other language games, including games where one robot attempts to entice the other robot to perform certain actions. It is clear to us that an exciting new area of bottom-up AI research is opening up and that through language and ontological development a possible road is opening up for evolving cognitive agents in a bottom-up fashion.

6 Acknowledgement

The experiments reported in this paper were conducted at the VUB AI Laboratory. Tony Belpaeme, Andreas Birk, Peter Stuer and Dany Vereertbrugghen have participated in creating and maintaining the robots and the ecosystem at the VUB AI Laboratory used in this experiment. Their contributions are supported by the Belgian government through an IUAP project (1990-1996). We are also indebted to Ruth Aylett from the university of Salford and a number of VUB undergraduate students, who performed the first attempts to develop a physical implementation of the language games.

Luc Steels developed the perception and language game algorithms. His research was financed by the Sony Computer Science Laboratory in Paris. Paul Vogt achieved the successful physical implementation as part of his masters thesis project for Cognitive Science and Engineering at the University of Groningen on a 'stage' at the VUB AI laboratory. This visit was financed by an EU Erasmus fellowship.

References

[1] Brooks, R. and L. Stein (1994) Building Brains for Bodies. Autonomous Robots (1) 7-25.

[2] Edelman, G.M. (1987) *Neural Darwinism: The Theory of Neuronal Group Selection.* New York: Basic Books.

[3] Husbands, P, I. Harvey, D. Cliff (1995) Circle in the Round: State Space attractors for Evolved Sighted robots. In: Steels, L. (ed.) (1995) The Biology and Technology of Intelligent Autonomas Agents. Nato ASI Series, vol 144. Springer Verlag, Berlin. pp. 222-257.

[4] McFarland, D. (1994) In: Proceedings of the third Simulation of Adaptive Behavior Conference. The MIT Press, Cambridge Ma.

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

[6] Pfeifer, R. (1996) Building Fungus Eaters: Design Principles of Autonomous Agents. In: Proceedings of Fourth International Conference on Simulation of Adaptive Behavior. The MIT Press, Cambridge Ma.

[7] Steels, L. (1994) The Artificial Life Roots of Artificial Intelligence. *Artificial Life Journal* 1(1), pp. 89-125.

[8] Steels, L. (1994) A case study in the behavior-oriented design of autonomous agents. In: Proceedings of the third Simulation of Adaptive Behavior Conference. The MIT Press, Cambridge Ma.

[9] Steels, L. (1996a) Emergent Adaptive Lexicons. In: Maes, P. (ed.) (1996) From Animals to Animats 4: Proceedings of the Fourth International Conference On Simulation of Adaptive Behavior, The MIT Press, Cambridge Ma.

[10] Steels, L. (1996b) Perceptually grounded meaning creation. In: Tokoro, M. (ed.) (1996) Proceedings of the International Conference on Multi-Agent Systems. AAAI Press, Menlo Park Ca. p. 338-344.

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

[12] Steels, L. (1997) Constructing and Sharing Perceptual Distinctions. In: van Someren, M. and G. Widmer (eds.) (1997) Proceedings of the European Conference on Machine Learning. Springer-Verlag, Berlin.

[13] Steels, L. (1997) The origins of syntax in visually grounded robotic agents. In: Polack, M. (ed.) Proceedings of IJCAI 97. AAAI Press, Menlo Park Ca.

[14] Steels, L. and R. Brooks (eds.) (1995) Building Situated Embodied Agents. The Alife route to AI. Lawrence Erlbaum Ass. New Haven.

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

[16] Vereertbrugghen, D. (1996) Design and Implementation of a Second Generation Sensor-Motor Control Unit for Mobile Robots. Licenciaatsverhandeling VUB Dept of Computer Science. 1996.

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

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