Autonomous Agents and Multi-Agent Systems, 1, 169–194 (1998)
© 1998 Kluwer Academic Publishers. Manufactured in The Netherlands.

The Origins of Ontologies and Communication Conventions in Multi-Agent Systems

LUC STEELS steels@arti.vub.ac.be
Sony Computer Science Laboratory Paris and Artificial Intelligence Laboratory, Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium

Abstract. The paper proposes a complex adaptive systems approach to the formation of an ontology and a shared lexicon in a group of distributed agents with only local interactions and no central control authority. The underlying mechanisms are explained in some detail and results of some experiments with robotic agents are briefly reported.

Keywords: origins of language, self-organization, distributed agents, open systems

1. Introduction

Agents cooperating in a multi-agent setting need a shared ontology and a shared set of communication conventions [12]. The question addressed in this paper is where these conventions might come from. One approach is to agree upon a domain ontology and a set of conventions in advance, and embed them in all future agent communication protocols. This is the approach underlying the standardisation efforts associated with Ontolingua [2] and KQML [1]. There are several reasons however why this may not be the best way to proceed.

1. It is hard to imagine how there could ever be a world-wide consensus about the ontologies and associated languages for every possible domain of multi-agent application.
2. Multi-agent systems are typically open systems. This means that the conventions cannot be defined once and for all but are expected to expand as new needs arise.
3. Multi-agent systems are typically distributed systems. There is no central control point. This raises the issue how evolving communication conventions might spread to agents which are already operational.
4. In the case of robotic agents, the ontology needs to be grounded in the sensori-motor capabilities of the agent.

This paper explores an alternative to top-down design and global enforcement, namely self-organised emergence. I discuss mechanisms by which a group of agents develop a shared lexicon for communicating a description, mechanisms by which agents develop their own ontology grounded in perception (but possibly grounded in other domains, e.g. social relations), and co-evolutionary couplings so that the ontology and the lexicon become tightly coordinated.

The main features of the proposed approach are:

1. There is no central controlling agency. Coherence arises in a bottom up, self-organised fashion.
2. The language community is open. New agents may enter at any time. They progressively adopt the conventions of the group and the group adopts new conventions that might be introduced by the new agent.
3. Conventions are adaptive. New meanings may enter at any time and the group develops the appropriate lexicalisations as needed.
4. The ontologies are adaptive. New stimuli from the environment may require the formation of new distinctions.

These features are achieved without giving up the basic principles ofan (autonomous) agent approach:

1. The agents have only limited knowledge. They cannot inspect the internal states of other agents.
2. The agents engage only in local interactions with other agents. No agent has a complete overview of what is happening.
3. The agents are autonomous. They acquire their own knowledge and decide for themselves how to communicate or divide up their world.
4. There is no global synchronisation. The system can operate in a fully distributed parallel fashion.

The proposed principles have been implemented in software simulations [7, 6, 8] andhave been integrated in robotic agents, in which case the ontology is based on an embodied physical interaction with the environment [5, 9]. This research is strongly related to a growing body of work on the origins of (natural) languages, extensively reviewed in [10].

The rest of the paper is intended as a survey paper of our experiments with more details available in the cited papers. The basic idea is presented briefly in the next section (section 2). Section 3 then focuses on lexicon formation and section 4 on ontology creation. Section 5 shows results of experiments in grounding.

2. The language game approach

An interaction between two agents can profitably be modeled as a game. When the interaction involves language, it is a *language game*. The games that we have studied concretely, assume that the speaker wants to identify an object to the hearer given a particular context of other objects. In other games, the speaker could demand the hearer to perform an action, request information, transmit an intention, etc. In order to perform a communication, the speaker must conceptualise the objects so as to find a description which distinguishes the topic from the other objects in the context. This requires an ontology, i.e. a set of distinctions. Then the speaker must find words to encode the distinctive features thus found, and transmit these words to the hearer. Next, the hearer receives the transmitted message, decodes it into one or more possible interpretations, and checks whether the interpretations are compatible with the present situation. The game succeeds if this is the case.

Failure may be due to (1) missing categories in the ontology of the speaker or hearer, or (2) missing or wrong linguistic conventions. In each case the agent engages in a repair action. New categories are created by extending the ontology, in other words by creating a new distinction or refining an existing distinction. New linguistic conventions are created by creating a new word or by adopting the word used by the speaker. Agents record the success of words and prefer words that had the most success. This causes coherence to emerge because the probability that a word is used increases if more agents adopt it. Agents also record the success of using a distinction. If a distinction is used often and has been successfully lexicalised it has a higher chance to remain in the population of possible distinctions.

The coordination of ontology creation and lexicon formation in a single agent and in a multi-agent system happens by co-evolution. There is an information flow and selectionist pressure in both directions. The ontology creation produces distinctions which are lexicalised. Lexicalisation is successful if the word is also used by other agents. Feedback is established from the lexicon to the ontology because the agents prefer distinctions that have been successfully lexicalised. This causes convergence of the ontology without a central control agency.

The coming sections contain more details of these various mechanisms followed by results from computational and robotic experiments showing that indeed a common lexicon and an ontology grounded in perceptual experiences emerges.

3. Lexicon formation

We begin by studying how adaptive language games lead to a shared lexicon associating form with meaning, focusing on one specific example, known as the naming game [7]. Similar systems have been proposed and studied by [3, 11, 4].

We assume a set of *agents* $\mathscr{A}$ where each agent $a \in \mathscr{A}$ has contact with a set of *objects* $\mathscr{O} = \{o_0, \ldots, o_n\}$. At this point meanings are taken to correspond to pointers to objects but later they are replaced by distinctions that distinguish the objects from each other. A *word* is a sequence of letters drawn from a finite alphabet. The agents are all assumed to share the same alphabet. A *lexicon* $\mathscr{L}$ is a time-dependent relation between meanings, words, and a score. Each agent $a \in A$ has his own set of words $W_{a,t}$ and his own lexicon $L_{a,t} \subset \mathscr{O}_a \times W_{a,t} \times \mathscr{N}$, which is initially empty. An agent a is therefore defined at a time t as a pair $a_t = \langle W_{a,t}, L_{a,t} \rangle$. There is the possibility of synonymy and homonymy: An agent can associate a single word with several meanings and a given meaning with several words. It is not required that all agents have at all times the same set of words and the same lexicon.

3.1. Operation of the naming game

We assume that environmental conditions identify a context $C \subset \mathcal{O}$. The speaker selects one object as the topic of this context $f_s \in C$. He signals this topic using extra-linguistic communication (for example, through pointing). Based on the interpretation of this signalling, the hearer constructs an object score $0.0 \leq e_o \leq 1.0$ for each object $o \in C$ reflecting the likelihood that o is the speaker's topic. If there is absolute certainty, one object has a score of 1.0 and the others are all 0.0. If there is no extra-linguistic communication, the likelihood of all objects is the same. If there is only vague extra-linguistic communication, the hearer has some idea what the topic is, but with less certainty. In our experiments, the object-score is determined by assuming that all objects are positioned on a 2-dimensional grid. The distance d between the topic and the other objects determines the object-score, such that

$$e_{object} = \frac{1}{1 + \left(\frac{d}{\alpha}\right)^2} \tag{1}$$

α is the object-focus factor. The higher the object-focus, the sharper the speaker's topic is distinguished from the other meanings.

Next the speaker retrieves from his lexicon all the associations which involve f_s. This set is called the association-set of f_s. Let $o \in \mathcal{O}$ be an object, $a \in \mathcal{A}$ be an agent, and t a time moment, then the association-set of o is

$$A_{o,a,t} = \{\langle o, w, u \rangle \mid \langle o, w, u \rangle \in L_{a,t}\} \tag{2}$$

Each of the associations in this set suggests a word w to use for identifying o with a score $0.0 \leq u \leq 1.0$. The speaker orders the words based on these scores. He then chooses the association with the largest score and transmits the word which is part of this association to the hearer.

Next the hearer receives the word w transmitted by the speaker. Uncertainty is modeled by assuming that the hearer recognises a set of possible words W related to w. These are all the words in the word-set of the hearer $W_{h,t}$ that are either equal to w or related with some distance to w. This distance gives a score

$$m_{w_1} = \frac{1}{1 + \left(\frac{d}{\beta}\right)^2} \tag{3}$$

β is the form-focus factor. The higher this factor, the sharper the hearer has been able to identify the word produced by the speaker.

For each word w_j, the hearer then retrieves the association-set that contains it. He constructs a *decision-matrix* which contains for each object a row and for each word-form a column. The first column contains the object-scores e_o, the first row the form-scores m_{w_j}. Each cell in the inner-matrix contains the association-score for the relation between the meaning (i.e. the object pointer) and the word-form in the lexicon of the hearer:

		w_1	w_2	...
		m_{w_1}	m_{w_2}	...
o_1	e_{o_1}	$u_{\langle o_1, w_1 \rangle}$	$u_{\langle o_1, w_2 \rangle}$	...
o_2	e_{o_2}	$u_{\langle o_2, w_1 \rangle}$	$u_{\langle o_2, w_2 \rangle}$	...
...	...	...	...	...

Obviously many cells in the matrix may be empty (and then set to 0.0), because a certain relation between a meaning and a word-form may not be in the lexicon of the hearer. Note also that there may be objects identified by lexicon lookup which are not in the initial context C. They are added to the matrix, but their object-score is 0.0.

The final state of an inner matrix cell of the decision-matrix is computed by taking the sum of (1) the object-score e_o on its row, (2) the word-form score m_w on its column, and (3) the association-score $a_{\langle o, w \rangle}$ in the cell itself. One meaning-word pair will have the best score and the corresponding meaning is the topic f_h chosen by the hearer. The association in the lexicon of this meaning-word pair is called the winning association. This choice integrates extra-linguistic information (the object-score), word-form ambiguity (the word-form-score), and the current state of the hearer's lexicon (the association-score).

3.2. Adaptation

The hearer then indicates to the speaker what topic he identified. In real-world language games, this could be through a subsequent action or through another linguistic interaction. When a decision could be made and $f_h = f_s$ the game succeeds, otherwise it fails. The following adaptations take place by the speaker and the hearer based on the outcome of the game.

1. The game succeeds This means that speaker and hearer agree on the topic. To re-enforce the lexicon, the speaker increments the score u of the association that he preferred, and hence used, with a fixed quantity δ. And decrements the score of the n competing associations with δ. 0.0 and 1.0 remain the lower and upperbound of u. An association is competing if it associates the topic f_s with another word. The hearer increments by δ the score of the association that came out with the best score in the decision-matrix, and decrements the n competing associations with δ. An association is competing if it associates the wordform of the winning association with another object. These changes implement an excitation-exhibition dynamics similar to the one used in Kohonen networks, except that the change is constant.

2. The game fails There are several cases:

1. The Speaker does not know a word.

It could be that the speaker failed to retrieve from the lexicon an association covering the topic. In that case, the game fails but the speaker may create a new word-form w' and associate this with the topic f_s in his lexicon. This happens with a word creation probability w_c.

2. The hearer does not know the word.

In other words there is no association in the lexicon of the hearer involving the word-form of the winning association. In that case, the game ends in failure but the hearer may extend his lexicon with a word absorption probability w_a. He associates the word-form with the highest form-score to the meaning with the highest object-score.

3. There is a mismatch between f_h and f_s.

In this case, both speaker and hearer have to adapt their lexicons. The speaker decrements with δ all the associations that have a word-form for f_h, and the hearer decrements with δ all associations that have f_h as meaning.

3.3. Tracing the game

Here are some traces for an experiment with 20 agents and 20 possible topics 1. After 500 games, the following dialogs are seen. Each time the speaker is given, the hearer, the topic, the possible repair actions, and then a list with the topic, the word used by the speaker, an arrow, the word heard by the hearer, and the interpretation by the hearer. If any of these are missing a question mark is printed.

```
500. Speaker: a8, Hearer: a14, Topic: o1
  Repair a8:
    Extend Lexicon o1 failed
  o1 ? ⇒ ? ? [FAILURE]
501. Speaker: a4, Hearer: a19, Topic: o5
  Repair a4:
    Extend Lexicon o5 failed
  o5 ? ⇒ ? ? [FAILURE]
502. Speaker: a20, Hearer: a16, Topic: o17
  Repair a16:
    Extend Lexicon o17 GEGO
  o17 GEGO ⇒ GEGO ? [FAILURE]

503. Speaker: a17, Hearer: a6, Topic: o2
  Repair a17:
    Extend Lexicon o2 failed
  o2 ? ⇒ ? ? [FAILURE]
```

```
504. Speaker: a10, Hearer: a4, Topic: o15
  Repair a10:
    Extend Lexicon o15 failed
  o15 ?⇒ ? ? [FAILURE]
505. Speaker: a17, Hearer: a1, Topic: o11
  Repair a1:
    Extend Lexicon o11 GUBO
  o11 GUBO ⇒ GUBO ? [FAILURE]
```

At this point moist games are still failing. The following table summarises for the group the most dominating word-meaning pair and their frequency:

meaning	*form*	*frequency*	*meaning*	*form*	*frequency*
o2	gota	0.15	o3	pitu	0.15
o4	dopu	0.20	o5	gabi	0.20
o6	gu	0.15	o7	gigu	0.10
o8	potu	0.10	o9	toga	0.25
o10	gu	0.25	o11	gubo	0.15
o12	depe	0.25	o13	ka	0.10
o14	bati	0.20	o15	beke	0.50
o16	tu	0.10	o17	to	0.25
o19	butu	0.50	o20	de	0.10

We see that the associations are still very weak. Only two words ("butu" and "beke") reach 50% spread in the population. Continuing the simulation, here is another series of traces as we are approaching 1000 games:

```
994. Speaker: a18, Hearer: a7, Topic: o9
  Repair a7: Store o9 TOGA
  o9 TOGA ⇒ TOGA ? [FAILURE]
995. Speaker: a19, Hearer: a7, Topic: o12
  Repair a19: New word o12 failed
  o12 ?⇒ ? ? [FAILURE]
996. Speaker: a15, Hearer: a5, Topic: o1
  Repair a15: New word o1 failed
  o1 ?⇒ ? ? [FAILURE]
997. Speaker: a5, Hearer: a6, Topic: o2
  Repair a6: Store o2 BITI
  o2 BITI ⇒ BITI ? [FAILURE]
998. Speaker: a13, Hearer: a11, Topic: o11
  o11 GUBO ⇒ GUBO o11 [SUCCESS]
```

```
999. Speaker: a4, Hearer: a17, Topic: o15
  Repair a17:
    Store o15 BEKE
  o15 BEKE ⇒ BEKE ? [FAILURE]
```

There is now a higher success rate. An overview of the lexicon is as follows:

meaning	*form*	*frequency*	*meaning*	*form*	*frequency*
o1	dato	0.10	o2	gota	0.25
o3	pitu	0.55	o4	dopu	0.35
o5	gabi	0.70	o6	gu	0.40
o7	gigu	0.30	o8	totu	0.20
o9	toga	0.40	o10	kebi	0.20
o11	gubo	0.35	o12	depe	0.35
o13	bu	0.10	o14	du	0.25
o15	beke	0.70	o16	tu	0.25
o17	ke	0.35	o18	gaba	0.20
o19	butu	0.70	o20	bopo	0.20

Some words (like "beke" for $o15$ and "gabi" for $o5$) have strongly established themselves, but for most words no consensus has emerged yet. For some meanings, a different word is used. For example, $o20$ changed from "de" to "bopo". 3000 games later, the production language is like this:

meaning	*form*	*frequency*	*meaning*	*form*	*frequency*
o1	dato	1.00	o2	biti	0.80
o3	pitu	0.60	o4	dopu	1.00
o5	gabi	1.00	o6	gu	0.85
o7	koti	0.50	o8	totu	0.65
o9	toga	0.90	o10	ku	0.80
o11	gubo	0.55	o12	ge	1.00
o13	bu	0.85	o14	ba	0.60
o15	beke	1.00	o16	tu	0.95
o17	ke	0.75	o18	gaba	0.95
o19	butu	1.00	o20	ki	0.95

Here is a sample trace after 8000 games. All games succeed:

```
8000. Speaker: a9, Hearer: a8, Topic: o11
  o11 GUBO ⇒ GUBO o11 [SUCCESS]
```

```
8001. Speaker: a16, Hearer: a18, Topic: o19
  o19 BUTU ⇒ BUTU o19 [SUCCESS]
8002. Speaker: a15, Hearer: a18, Topic: o16, o8
  o16 TU ⇒ TU o16 [SUCCESS]
8003. Speaker: a9, Hearer: a16, Topic: o18
  o18 GABA ⇒ GABA o18 [SUCCESS]
8004. Speaker: a7, Hearer: a14, Topic: o5
  o5 GABI ⇒ GABI o5 [SUCCESS]
8005. Speaker: a15, Hearer: a11, Topic: o7
  o7 GIGU ⇒ GIGU o7 [SUCCESS]
```

3.4. Success and coherence

The naming game model can be viewed as a complex dynamical system. The agents have a certain local behavior (an agent can only interact with one single agent, not with all agents at the same time), which is determined by their internal lexicons. Behavior changes because agents adapt their lexicon. In order to 'see' the global order in the system, we need macroscopic variables. These macroscopic variables are invisible to the agents because no agent has a complete overview of the behavior of the group. The first such variable quantifies the *average success* after n games. When average success approaches total success, this must mean that the conventions are sufficiently shared to speak of the emergence of a shared lexicon. But, because a word may have many meanings and the same meaning may be expressed by multiple words, communicative success does not necessarily mean complete coherence. An agent can very well know a word but prefer not to use it himself.

Given the preferred lexicon for a single agent, it is straightforward to determine the lexicon of the group as being the set of word-meaning associations that are preferred by most agents. The *production coherence* of the lexicon is equal to the average frequency of the most preferred word-meaning association.

It is also instructive to look at the evolution of the average association-scores competing for the preferred expression of a particular word. This is done through *competition diagrams* as the one shown in figure 2. The diagram shows that there is a winner-take-all situation. This is due to the positive feedback loop between score and use. The higher the score of a word, the more it is used, and the more its chances increase to be successful in further use. Such a winner-take-all situation takes place for every meaning so that a global shared lexicon emerges.

3.5. Open systems

Once total game success is reached, the lexicon does not change anymore. The only source of possible innovation is the introduction of new words, which only happens when an agent does not have a word yet, or the progressive adoption of one word by the group, which stops as soon as a winner-take-all situation has occurred.

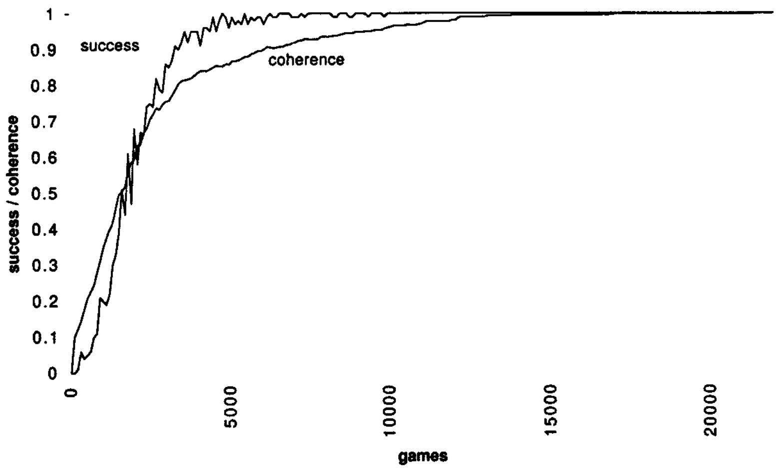

Figure 1. This figure shows the evolution of both average success and production coherence for a group of 20 agents and 20 meanings. Total production coherence climbs less fast once the population has reached total average success.

A lexicon is even resistent (up to a certain degree) to changes in the population. This is investigated by introducing an in- and outflux in the population. When agents leave, they take their lexicons with them. When new agents enter, they have to acquire the lexicon of the other agents in the group. They may occasionally create a new word (with a small probability the word creation probability w_c) but

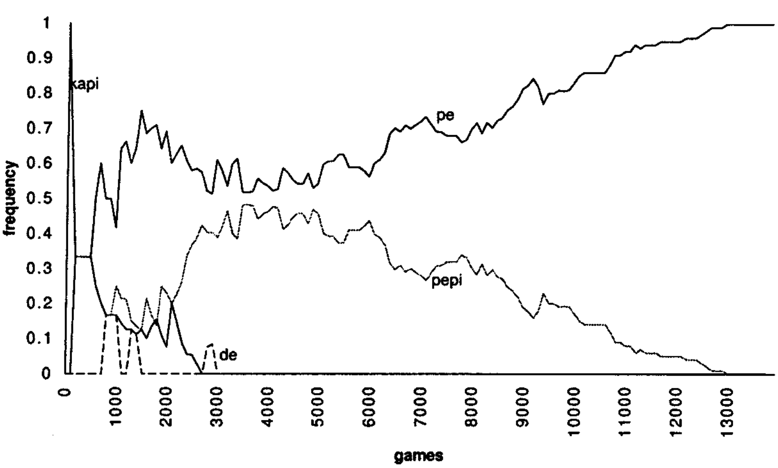

Figure 2. The form-competition diagram shows for a single meaning the frequency of each competing form. We clearly see a winner-take-all situation emerging.

this new word quickly gets damped against the dominance of the preferred word. Acquisition of an existing lexicon by a new agent happens without any addition or change to the model, as shown in figure 3 which plots also the language change. Change is quantified by comparing the state of the lexicon at two time points and counting the number of preferred word-meaning pairs that changed. We see that the lexicon changes rapidly in the beginning as the population moves towards total average game success. Thereafter the lexicon remains stable. Figure 3 shows what happens when a flux is introduced in the population. When new agents come in, game success and coherence drops because the new agent has to acquire the lexicon of the group. But if there are not too many agents coming in, the group will maintain a high rate of success. More importantly, the lexicon itself does not change at all. It is transmitted culturally from one generation to the next. When the rate of population renewal is too high, the lexicon disintegrates, as also shown in figure 3 There is rapid lexicon change because the new agents start to create new word-meaning associations, but these conventions cannot propagate fast enough in the population.

An influx *and* an outflux of meanings is investigated in the next experiment 4. Not only are meanings taken out but new meanings enter at regular intervals. In a first phase, the system is closed and a shared lexicon emerges. In phase 2, a relatively small meaning flux is introduced (1/1000 games). The population copes with the change. New words are created and propagate in the population. Next (phase 3 in figure 4) a much higher meaning flux is imposed (1/1000 games). Production coherence decreases and average success plummets. The system restores itself afterwards when the flux of meaning is brought back to 1/100 games.

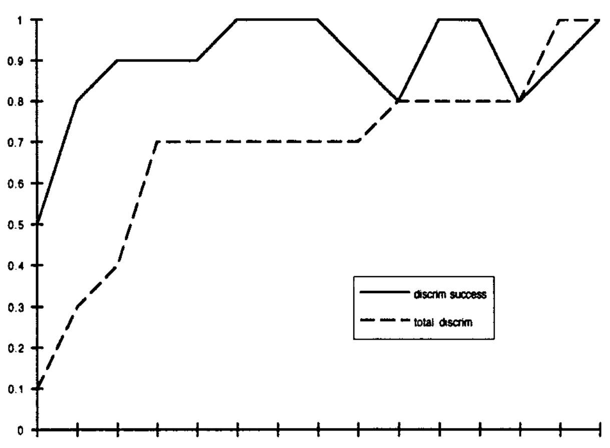

Figure 3. This graph shows both lexicon change and average success. In a first phase, the lexicon forms itself in a closed population. In a second phase, an in- and outflow of agents (1 in/outflow per 100 games) is introduced, the lexicon stays the same and success is maintained. In the third phase the flux is increased to 1 per 10 games and the lexicon disintegrates.

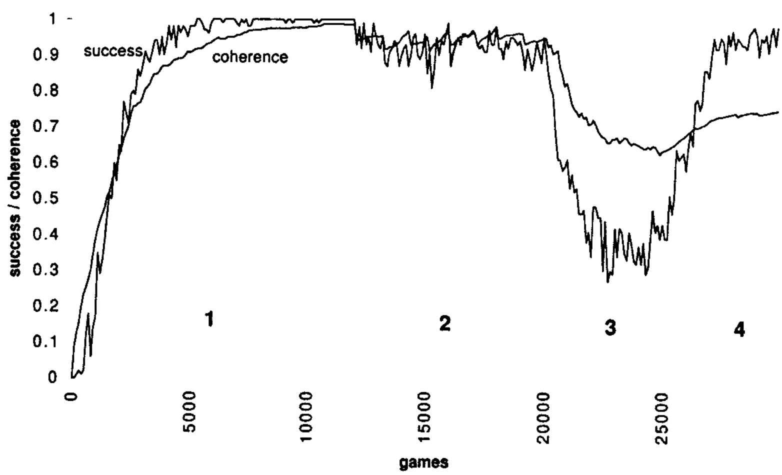

Figure 4. Both average success (per 100 games) and production coherence is shown. For a small rate (1/1000) the group is able to cope (20 agents and 20 meanings). For a large rate (1/100) success and production coherence drops, but is restored when the rate of change moves back to an earlier level (1/1000).

4. Ontology creation through discrimination games

The experiments discussed in the previous section show that a population of distributed agents is able to autonomously develop a shared lexicon through self-organisation. The system is resistent to fluxes in the set of agents and the set of meanings, within certain parameter ranges. We now turn to the problem of the ontology. In the experiments earlier on, it is assumed that objects could be identified through direct pointers. In real-world language games this is not possible. Agents must conceptualise reality and use as meaning a set of features that distinguishes the topic from the other objects in the context. This raises two issues: (1) where does the ontology come from used to make these distinctions, and (2) how can autonomous agents ever develop a shared ontology even though no agent can inspect the brain state of another one.

Our approach is similar to the naming game approach. We define another type of adaptive game (called a discrimination game) which is played between an agent and the world. An agent has an evolving repertoire of distinctions which are binary categorisations dividing up data arriving at sensory channels. The agent attempts to perform a discrimination with these distinctions, i.e. find a feature set that distinguishes the topic from the other objects in the context. If that fails the agent extends the repertoire by creating a new node in one of the binary decision trees or by starting a new tree on a sensory channel that had not been explored yet. There is again a positive feedback loop between success and survival in the population of distinctions.

Let there be a set of objects $\mathcal{O} = \{o_1, \ldots, o_m\}$ and a set of sensory channels $S = \{\sigma_1, \ldots, \sigma_n\}$, being real-valued partial functions over $\mathcal{O}$. Each function σ_j defines a value $0.0 \leq \sigma_j(o_i) \leq 1.0$ for each meaning o_i.

An agent a has a set of feature detectors $D_a = \{d_{1,s}, \ldots, d_{a,m}\}$. A *feature detector* $d_{a,k} = \langle p_{a,k}, V_{a,k}, \phi_{a,k}, \sigma_j \rangle$ has an attribute name $p_{a,k}$, a set of possible values $V_{a,k}$ a partial function $\phi_{a,k}$, and a sensory channel σ_j. The result of applying a feature detector $d_{a,k}$ to an meaning o_i is a feature written as a pair $(p_{a,k}\; v)$ where p is the attribute name and $v = \phi_{a,k}(\sigma_j(o_i)) \in V_{a,k}$ the value.

The *feature set* of a for o_i is defined as $F_{a,o_i} = \{(p_{a,k}\; v) \mid d_{a,k} \in D_a, d_{a,k} = \langle p_{a,k}, V_{a,k}, \phi_{a,k}, \sigma_j \rangle, v = \phi_{a,k}(\sigma_j(o_i))\}$. Two features $(a_1\, v_1)$, $(a_2\, v_2)$ are *distinctive* iff $a_1 = a_2$ *and* $v_1 \neq v_2$. A distinctive feature set D^C_{a,o_t} is a set of features distinguishing an meaning o_t from a set of other meanings C. $D^C_{a,o_t} = \{f \mid f = (p\, v) \in F_{a,o_t}$ *and* $\forall o_c \in C$ *either* $\setminus \exists f' = (p'\, v') \in F_{a,o_c}$ *with* $p = p'$ *or* $\exists f' \in F_{a,o_c}$ *with* f *and* f' *distinctive*$\}$. Clearly there can be several distinctive feature sets for the same o_t and C, or none.

A discrimination game $d = \langle a, o_t, C \rangle$ involves an agent a, a topic $o_t \in C \subseteq \mathcal{O}$. C is called the context. The outcome of the game is twofold. Either a distinctive feature set could be found, $D^C_{a,o_t} \neq \emptyset$, and the game ends in success, or no such feature set could be found, $D^C_{a,o_t} = \emptyset$, and the game ends in failure.

As part of each game the repertoire of meanings is adjusted in the following way by the agent:

1. $D^C_{a,o_t} = \emptyset$, i.e. the game is unsuccessful. This implies that there are not enough distinctions and therefore $\exists o_c \in C$, $F_{a,o_t} \subseteq F_{a,o_c}$. There are two ways to remedy the situation:
   (a) If there are still sensory channels for which there are no feature detectors, a new feature detector may be constructed. This option is preferred.
   (b) Otherwise, an existing attribute may be refined by creating a new feature detector that further segments the region covered by one of the existing attributes.
2. $D^C_{a,o_t} \neq \emptyset$. In case there is more than one possibility, feature sets are ordered based on preference criteria. The 'best' feature set is chosen and used as outcome of the discrimination game. The record of use of the features which form part of the chosen set is augmented. The criteria are as follows:
   (a) The smallest set is preferred. Thus the least number of features are used.
   (b) In case of equal size, it is the set in which the features imply the smallest number of segmentations. Thus the most abstract features are chosen.
   (c) In case of equal depth of segmentation, it is the set of which the features have been used the most. This ensures that a minimal set of features develops.

The whole system is selectionist. Failure to discriminate creates pressure to create new feature detectors. However the new feature detector is not guaranteed to do the job. It will be tried later and only thrive in the population of feature detectors if it is indeed successful in performing discriminations.

The discrimination game defined above has also been implemented. To test the mechanism, we create a set a sensory channels and an initial set of objects which have arbitrary values for some of the sensory channels. A typical example is the following list of objects and associated values for channels:

```
o-0: [sc-3:0.73] [sc-4:0.82] [sc-5:0.07]
o-1: [sc-0:0.89] [sc-3:0.02]
      [sc-4:0.56] [sc-6:0.48]
o-2: [sc-0:0.74] [sc-1:0.92] [sc-2:0.22]
      [sc-3:0.56] [sc-8:0.52] [sc-9:0.03]
o-3: [sc-2:0.36] [sc-3:0.09] [sc-4:0.14]
o-4: [sc-1:0.47] [sc-2:0.61] [sc-3:0.69]
      [sc-5:0.67] [sc-6:0.14] [sc-9:0.43]
...
```

A feature detector is a function assigning a value to a certain attribute. The name of the attribute indicates its nature. It is of the form $sc_i - n_1 - \ldots$ where i is the sensory channel followed by which one of the two segments has been chosen. For example, sc-5 is the name of an attribute whose feature detector operates on sc-5. sc-5-1 is a feature that identifies the second segment of sc-5. sc-5-1-0 identifies the first segment of the second segment of sc-5, etc. (sc-5-1-0 v-0) is a feature combining this attribute with the value v-0.

In normal operation, the agent continuously goes through a loop performing the following activities:

1. A context is delineated. The context consists of the objects currently in the field of attention of the agent.
2. One object in this context is chosen randomly as topic.
3. The feature sets of the topic and the other objects in the context are derived.
4. An attempt is made to find possible discriminating feature sets.

We now show some typical situations for an agent a-5, which starts from no features at all. In the first game, a-5 tries to differentiate the object o-5 from o-3. The agent does not have a way yet to characterise the topic and creates a new attribute operating on sc-5.

```
a-5: o-5 ↔ {o-3 }
Topic: NIL
Not enough features topic
New attribute: sc-5
```

The next game to distinguish o-5 from o-9 and o-1 is already successful, because o-5 is again the topic. The context contains objects that do not have any response for sc-5, and thus no features can be constructed:

```
a-5: o-5 ↔ {o-9 o-1 }
Topic: ((sc-5 v-1))
Topic: (NIL NIL)
Success: ((sc-5 v-1))
```

The next game is also sucessful because o-6 has value v-0 for sc-5, o-2 has nothing and o-5 has v-1.

```
a-5: o-6 ↔ {o-2 o-5 }
Topic: ((sc-5 v-0))
Topic: (NIL ((sc-5 v-1)))
Success: ((sc-5 v-0))
```

In the following game the attributes are not sufficiently distinctive and therefore a new attribute is created. As long as there are possibilities to focus on additional sensory channels, existing attributes will not be refined. The new attribute operates on sc-3.

```
a-5: o-7 ↔ {o-1 o-2 }
Topic: ((sc-1 v-1))
Topic: (NIL ((sc-1 v-1)))
No distinctive features but new
  one possible: (sc-2 sc-3 sc-8)
New attribute: sc-3
```

When uncovered sensory channels are no longer available, more refined feature detectors for existing attributes start to be made. In the following example, o-0 fails to be distinguished from o-8 and 0-1, even though a set of features is available to characterise each object. A refinement of the attribute operating over sc-5 is chosen.

```
a-5: o-0 ↔ {o-8 o-1 }
Topic: ((sc-3 v-1)(sc-4 v-1)(sc-5 v-0))
Topic: (((sc-0 v-1)(sc-1 v-0)(sc-3 v-1)
      (sc-4 v-0)(sc-5 v-0)))
    ((sc-0 v-1)(sc-3 v-0)(sc-4 v-1)))
No distinctive features but refinements possible.
Refining attribute: sc-5 ⇒ sc-5-0, sc-5-1
```

After a sufficient number of discrimination games the set of features stabilises. For the set of objects given above, the following is a stable discrimination tree. For each attributes the possible values are listed with their corresponding regions as well as the number of times a feature has been used.

```
sc-5:
  v-0: [0.00 0.50] 358.
    sc-5-0:
      v-0: [0.00 0.25] 31.
        sc-5-0-0:
          v-0: [0.00 0.12]
            sc-5-0-0-0:
              v-0: [0.00 0.06] ; v-1: [0.06 0.12] 3.
            v-1: [0.12 0.25]
          v-1: [0.25 0.50] 22.
        v-1: [0.50 1.00] 309.
```

```
sc-1:
  v-0: [0.00 0.50] 651. ; v-1: [0.50 1.00] 628.
sc-3:
  v-0: [0.00 0.50] 713. ; v-1: [0.50 1.00] 733.
sc-8:
  v-0: [0.00 0.50] 15. ; v-1: [0.50 1.00] 8.
sc-2:
  v-0: [0.00 0.50] 99. ; v-1: [0.50 1.00] 112.
sc-0:
  v-0: [0.00 0.50] ; v-1: [0.50 1.00] 42.
sc-4:
  v-0: [0.00 0.50] 223.
        sc-4-0-0:
          v-0: [0.00 0.25]
          v-1: [0.25 0.50] 1.
            sc-4-0-0-1:
              v-0: [0.25 0.37] 5.; v-1: [0.37 0.50] 5.
  v-1: [0.50 1.00] 215.
    sc-4-1:
      v-0: [0.50 0.75] 1.
      v-1: [0.75 1.00] 2.
        sc-4-1-1:
          v-0: [0.75 0.87] 5. ; v-1: [0.87 1.00] 2.
sc-6:
  v-0: [0.00 0.50] 2. ; v-1: [0.50 1.00]
```

We see that more abstract features, like (*sc-1 v-0*), are used more often. For some, like (*sc-5 v-0*), there is a deep further discrimination. For others, like (*sc-5 v-1*), there is none. Some features, like (*sc-6 v-1*), have not been used at all and could therefore be eliminated. Another experiment with the same objects but for a different agent a-6 yields a different discrimination tree. In one example, some sensory channels (such as sc-6) were not used, sc-4 was no longer refined, etc. Usually there are indeed many different possibilities and an important question for further study is how optimal the discrimination trees obtained with the proposed mechanism are.

When new objects enter the environment, the agent should construct new distinctions if they are necessary. This is effectively what happens. If new sensory channels become available, for example because a new sensory routine has become active, then it will be exploited if the need arises.

Figure 5 shows a typical example where an agent builds up a repertoire of feature detectors, starting from scratch. We start from a set of 10 objects and gradually add new objects in a probabilistic fashion, to reach a total of 50 objects. We see that the feature repertoire is extended occasionally. The average discrimination success remains close to the maximum (1.0) because new objects are only encountered occasionally and the feature detectors already constructed are general. Figure 6 shows how the system copes with new objects.

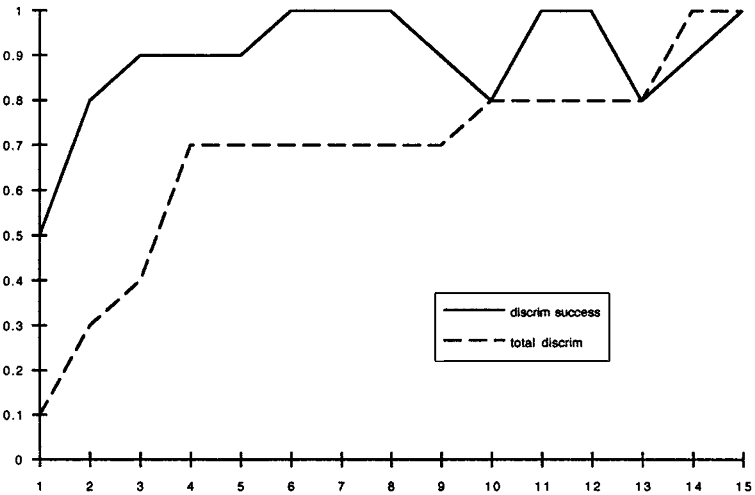

Figure 5. The graph shows the evolution of the discriminatory capacities of a single agent. The total number of objects (10) is fixed. There are 5 sensory channels. The average success in discrimination games as well as the global success is shown on the y-axis. The number of discrimination is mapped on the x-axis (scale 1/10). All objects can be discriminated after 150 discrimination games.

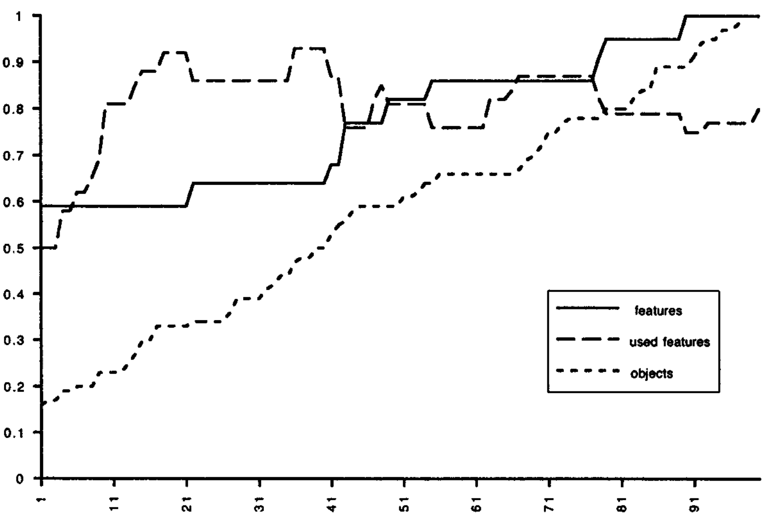

Figure 6. The graph shows on the y-axis the number of objects (as a percentage of the total reached at the end, i.e. 50), the discriminatory success which remains close to the maximum, and the number of features (as a percentage of the total reached at the end, i.e. 35). The x-axis plots the number of discrimination games (scale 1/10).

When performing multi-agent experiments, each of the agents is running the same ontology creation mechanisms. Even if they are in the same environment they will end up with different ontologies. Similarities are uniquely due to the fact that the agents share the same context. The coupling to lexicon formation discussed in the next section pushes the ontologies towards greater coherence because it is a collective activity with feedback between words and meanings.

5. Grounding experiments

The self-organised coherence in lexicons and ontologies has been well-established in software experiments. Based on this success, we decided to see whether the mechanisms would also work on physically instantiated robotic agents. This is even more challenging because it forces us to test the robustness of the proposed mechanisms in real world settings and to see whether ontology creation can handle the rich variation present in real-world data.

5.1. Language games on mobile robots

A first experiment developed in collaboration with Paul Vogt (reported more extensively in [5] was conducted on fully mobile robots. The robots are small Lego-vehicles which have a variety of sensors (infrared, visible light, sound, touch, etc.), actuators for moving around in the environment, batteries, and on board processors. The robots operate in a physical ecosystem in which they have opportunities to recharge their batteries but also competitors which have to be countered by performing work (figure 7).

The observational channels contain the real world data obtained from the physical sensors. An example of such data is given in figure 8. The sensors are always located on the body in pairs, for example left infrared and right infrared sensor, left and right visible light sensing, etc., so that the robot has a center of perception (as most animals). An object is in this center of perception when the left and right sensory data cross over. Thus if the robot turns left towards the visible light emitted by the charging station, it will be centered on the charging station when the left visible light peak decreases and crosses the increasing right visible light peak. The sensory values at each crossing point act as input to the discrimination games.

The protocol for engaging in language games has been implemented on the physical robots by a combination of physical gestures and communications through a radio link between the robots. Two robots engage in a communication when both are facing each other. Then each robot makes a 360 degree turn to develop a panoramic sensory view of the environment. The pointing is implemented by a gesture: The speaking robot emits 4 infrared beams while turning towards the topic, so that the other robot can observe in which direction it moves. The speaking robot stops turning when it is facing the object that it wants to see as the topic of

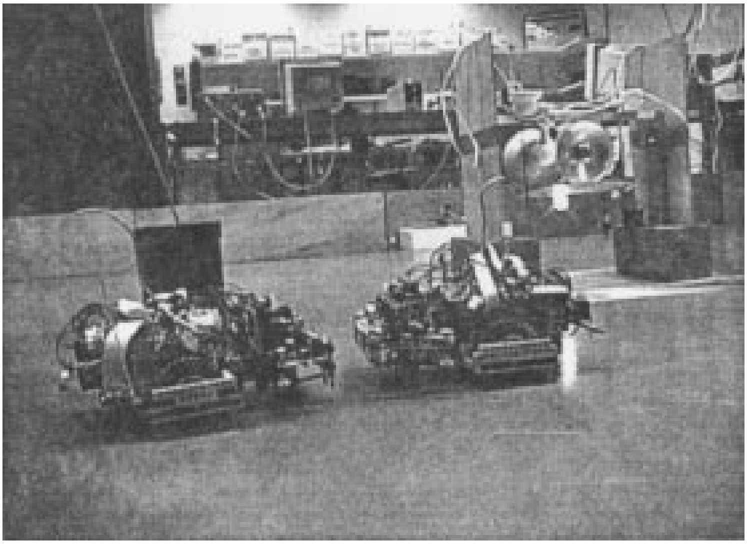

Figure 7. Two robots have approached each other and are now facing each other. The robots are equiped with a dozen low-level sensors. The discrimination trees are based on output from these sensory channels. Note the other objects in the environment surrounding the robots (charging station, competitors, obstacles), which are the subject of conversations.

the conversation. The listening robot detects the topic by consulting its own sensory map. Then the language game starts as described above.

An example of a language game between two robots (r1 and r0) at the earliest stages is as follows. Three objects are encountered by r1 and 7 by r0. For each of these objects, the position is given (for example 46 for o1), as well as the data (for o1 this is [0,2,12,3]) followed by the features that have been extracted based on the

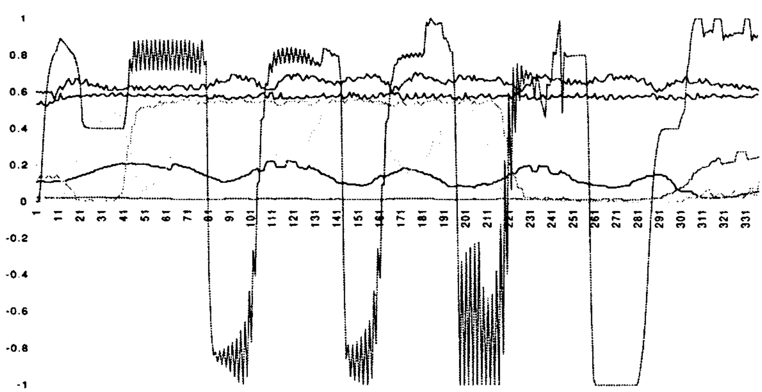

Figure 8. Sensory data streams taken from physical robot. The channels include left and right infrared and visible light sensing and motor speeds.

discrimination trees developed so far (for o1 sc1-127,sc2-127,sc3-127). Although the speaker has a distinctive feature set for the topic namely {sc0-1}, it has no words yet for it. The game therefore fails.

```
20. Speaker: r1. Hearer r7. Topic: o0.
  Objects seen by speaker:
    o0 = [0] [1,0,0,0] → {sc0-1}
    o1 = [46] [0,2,12,3] → {sc1-127,sc2-127,sc3-127}
    o2 = [96] [0,1,0,193] → {sc1-127,sc3-127}
  Objects seen by hearer:
    o0 = [0] [1,0,0,0] → {sc0-1}
    o1 = [6] [0,86,12,169] → {sc1-127,sc2-127,sc3-127}
    o2 = [7] [0,81,9,168] → {sc1-127,sc2-127,sc3-127}
    o3 = [9] [0,82,12,171] → {sc1-127,sc2-127,sc3-127}
    o4 = [20] [0,37,29,167] → {sc1-127,sc2-127,sc3-127}
    o5 = [67] [0,1,4,195] → {sc1-127,sc2-127,sc3-127}
    o6 = [72] [0,0,4,217] → {sc2-127,sc3-127}
  {sc0-1} ? ⇒ ? ? [failure]
```

Another language game much further in the process (after about 1000 games) is as follows: Both speaker and hearer have a distinctive feature set (sc0-1 and sc3-190 respectively) to distinguish the topic from the other objects. The speaker uses the word "(c d)" which is recognised as compatible by the hearer with what he expects. The game succeeds.

```
1010. Speaker r0. Hearer r1. Topic: o0
  Objects seen by speaker:
    o0 = [0] [1,0,0,0] → {sc0-1}
    o1 = [7] [0,152,10,190] → {sc1-149,sc2-10,sc3-189}
    o2 = [45] [0,1,7,181] → {sc1-1,sc2-7,sc3-186}
  Objects seen by hearer:
    o0 = [0] [1,0,0,0] → {sc0-1}
    o1 = [6] [0,3,0,115] → {sc1-3,sc3-115}
    o2 = [19] [0,3,0,208] → {sc1-3,sc3-209}
    o3 = [36] [0,3,0,29] → {sc1-3,sc3-30}
    o4 = [118] [0,4,0,192] → {sc1-4,sc3-190}
  Topic perceived by hearer: o4
  {sc0-1} (c d) ⇒ (c d) {sc3-190} [success]
```

The graphs in 9 show the evolution of the success rate in the lexicon of a single agent. We have now demonstrated in a large number of experiments, that even in these very difficult circumstances coherence and successful communication emerges. The circumstances are difficult because every step in the process may fail: A robot may loose its orientation in constructing a panoramic view, the pointing may fail, the data is to some extent erratic, they may loose radio contact during the communication, etc.

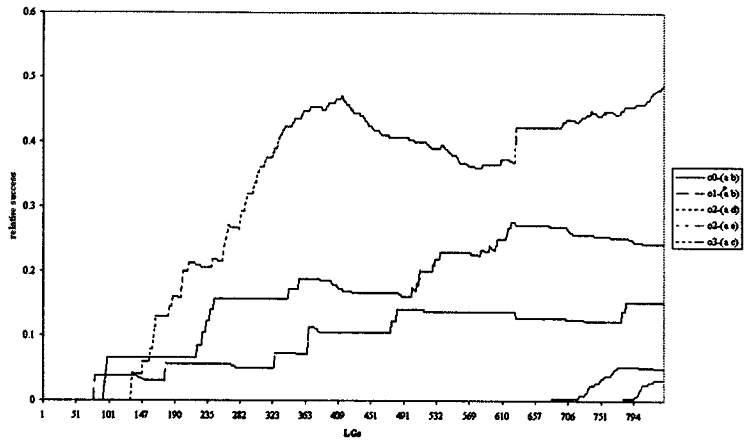

Figure 9. This figure shows for the different associations how much success each one has had in the games played so far.

5.2. The talking heads experiment

A second experiment in physical grounding of language formation processes is known as the *talking heads experiment*. It is reported more extensively in [Steels, 1997b]. The experiment is based on two robotic heads which can track moving objects based on visual inputs. The heads watch a static or dynamic scene. A typical example of a scene as seen through one of the heads is contained in 10. The segments recognised by low level sensory routines are surrounded by a bounding box. These segments act as the objects of a language game. Low level visual processing extract data for each segment, such as the area of the bounding box, the ratio of the segment area compared to the bounding box area, the average light intensity within a bounding box, etc. Based on these data distinctions are created such as large–small, rectangular–not rectangular, dark–light. Then the creation of a lexicon expressing distinctive feature combinations necessary to identify an object proceeds as outlined in earlier sections.

Here are some examples of interactions. In the first one the speaker fails to conceptualise the scene and creates a new category by dividing the sensory channel called fill-ratio into two segments associated with the values v-81 and v-82.

```
0. Speaker: Head-16. Hearer: Head-17. Topic: o4.
  Repair Head-16:
    Extend categories: FILL-RATIO [-1.0 1.0]: v-81 v-82
  ? ?⇒ ? ? [failure]
```

In the next game, there is another failure and a new distinction is created now on

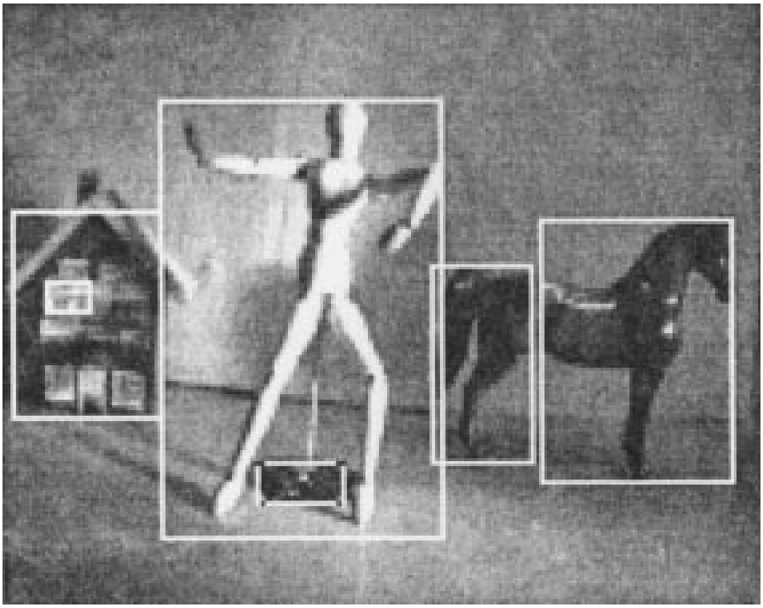

Figure 10. View through the camera of one of the heads. The consecutive bitmaps are segmented and here visualised by a bounding box around each segment.

the visibility channel:

```
1. Speaker: Head-16. Hearer: Head-17.
  Repair Head-16:
    Extend categories: VISIBILITY [-1.0 1.0]: v-83 v-84
  ? ? ⇒ ? ? [failure]
```

In game 4, a set of distinctive features has been found but there is no word yet. The speaker creates a new word:

```
4. Speaker: Head-16. Hearer: Head-17. Topic: o12.
  Repair Head-16:
    Extend word repertoire: “(d u)”
    Extend lexicon: ((visibility v-88)) (d u)
  ((visibility v-88)) (d u) ⇒ ? ? [failure]
```

In the following game, the speaker is able to find a distinctive feature set and a word, but the hearer is missing the required distinctions:

```
6. Speaker: Head-16. Hearer: Head-17. Topic: o10.
  Repair Head-17:
    Extend categories: VISIBILITY [-1.0 1.0]: v-91 v-92
  ((visibility v-88)) (d u) ⇒ (d u) ? [failure]
```

The first successful game happens after 47 games:

```
47. Speaker: Head-16. Hearer: Head-17. Topic: o25.
  ((visibility v-109)(area v-108)(fill-ratio v-81)) (k i)
    ⇒ (k i) (fill-ratio v-125)(intensity v-134)(area v-
132))
  [success]
```

A snapshot of the lexicon of one agent is as follows:

meaning	*form*
((visibility v-88))	(d u)
((fill-ratio v-82))	(t e)
((fill-ratio v-81)(area v-86)(visibility v-87))	(l e)
((intensity v-90))	(n a)
((fill-ratio v-81)(area v-86)(intensity v-89))	(p u)
((intensity v-89))	(m i)
((fill-ratio v-81)(area v-108)(visibility v-109))	(k i)

Figure 11 shows the increased success in communication as the agents continue to build up a shared lexicon and the increase in complexity of the lexicons.

Although the physical embodiment of the Talking Heads experiment is quite different from the mobile robots, we see the same phenomena: steady increase and adaptation of a perceptually grounded ontology, and progressive build up and self-organised coherence of a shared lexicon. The Talking Heads experiment is somewhat easier because visual perception provides a richer source for building an ontology and the communication and perceptual conditions are more stable.

6. Conclusions

This paper has discussed mechanisms for the creation of ontologies in the form of discrimination trees of perceptually grounded categories and the formation of a lexicon expressing a feature structure using these categories. The mechanisms exploit three principles known from biology: self-organisation, selectionism, and co-evolution. Self-organisation appears when there is a positive feedback loop between an emergent structure (in this case a shared lexicon) and future behavior. Selectionism occurs when a system generates spontaneous variation which is amplified or filtered under environmental pressure. In the present case, the spontaneous variation occurs through the (relatively) random expansion of the discrimination trees which will be positively selected for if they are relevant in future games. Co-evolution occurs when two selectionist systems are coupled in the sense that selectionist pressure flows from one to the other and vice-versa. Because

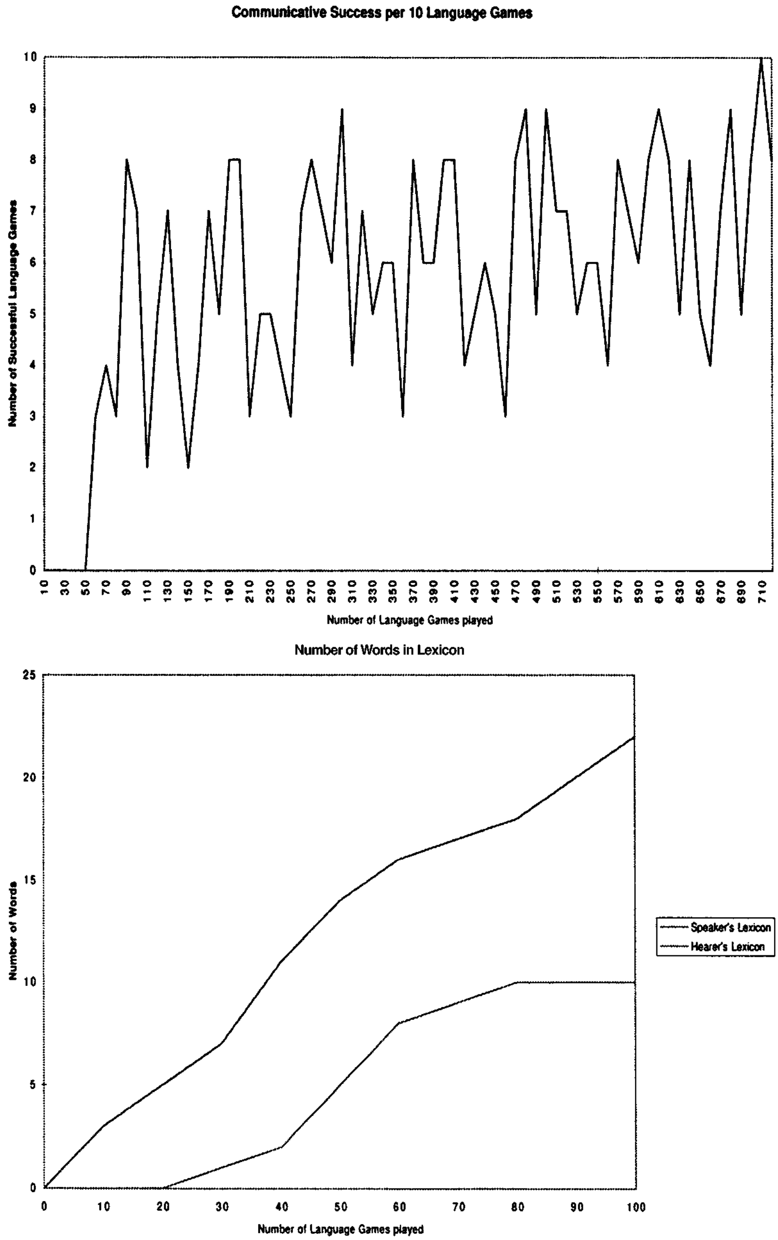

Figure 11. Graph showing the increase in average communicative success (top) as well as the increase in the number of words in the vocabularies of two robotic heads (bottom).

agents prefer words that have shown more success in the past, the more successful words will propagate in the population. Because the success of a word feeds back to the survival of the distinctions underlying this word, a shared ontology emerges. The sharing is always incomplete and dynamic. It is incomplete because agents may have success in communication even though they use different categories or they have different meanings for the same word which are nevertheless compatible with the situations in which they find themselves. The sharing is dynamic because new distinctions or new words may be created as required by the circumstances.

The mechanisms proposed here are generally applicable both to software agents and to robotic agents. It is sufficient to identify the observational channels, and to set up the appropriate feedbacks from the environment (for example, initially some form of pointing to establish a shared context).

Although results obtained with the presented mechanisms are very encouraging, many open issues remain. The issue of syntax and its origins has not been discussed even though some progress in this area has been made (see [9]). Syntax becomes necessary when the meaning to be conveyed is more complex and when the agents want to press more information in a single expression and thus optimise communication and make it more reliable. It is also clear that natural languages have a much more flexible way to match meaning against a lexicon, occasionally using analogical reasoning. This implies that a flexible inference machinery is integrated in lexicon lookup. These and other issues are the subject of intense current research.

Acknowledgment

The software simulations discussed in this paper have been developed at the Sony Computer Science Laboratory in Paris by Luc Steels, Angus McIntyre and Frederic Kaplan. Several robot builders at the VUB AI Lab (soft and hardware) have contributed to the grounding experiments, including Andreas Birk, Tony Belpaeme, Peter Stuer and Danny Vereertbrugghe. The grounding experiments on mobile robots were carried out by Paul Vogt. The physical implementation of the Talking Heads experiment was developed by Tony Belpaeme, Luc Steels and Joris Van Looveren.

References

1. *Arpa Knowledge Sharing Initiative*. *Specification of the kqml agent-communication language*, July 1993.
2. M. Genesereth and S. Ketchpel, "Software agents," *Communications of the ACM*, vol. **7-37**, pp. 48–53, 1994.
3. B. MacLennan, *Synthetic ethology*: *An approach to the study of communication.* In C. Langton, editor, *Artificial Life II*, Addison-Wesley Pub. Co., Redwood City, CA, 1991.
4. M. Oliphant, "The dilemma of saussurean communication," *Biosystems*, vol. **1:2-37**, pp. 31–38, 1996.
5. L. Steels and P. Vogt, "Grounding adaptive language games in robotic agents," In I. Harvey and P. Husbands, editors, *Proceedings of the 4th European Conference on Artificial Life*, Cambridge, MA. 1997.

6. L. Steels, "Emergent adaptive lexicons," In Mataric M. Meyer J.-A. Pollack J. and Wilson S. W. Maes, P., editor, *From Animals to Animats 4*: *Proceedings of the Fourth International Conference On Simulation of Adaptive Behavior*, Cambridge, MA, 1996. The MIT Press.
7. L. Steels, "Self-organizing vocabularies," In C. Langton, editor, *Proceeding of Alife V*, Nara, Japan, 1996.
8. L. Steels, "Constructing and sharing perceptual distinctions," In M. can Someren and G. Widmer, editors, *Proceedings of the European Conference on Machine Learning*, Springer-Verlag, Berlin, 1997.
9. L. Steels, "The origins of syntax in visually grounded robotic agents," In M. Pollack, editor, *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, Morgan Kauffman Publishers, Los Angeles, 1997.
10. L. Steels, "The synthetic modeling of language origins," *Evolution of Communication Journal*, vol. **1-1**, pp. 1–34, 1997.
11. G. M. Werner and M. G. Dyer, "Evolution of communication in artificial organisms," In C. G. Langton, C. Taylor, and J. D. Farmer, editors, *Artificial Life II*, *Vol. X of SFI Studies in the Sciences of Complexity*, Addison-Wesley Pub., Redwood City, CA, 1991.
12. M. Wooldridge and N. R. Jennings, "Intelligent agents: Theory and practice," *Knowledge Engineering Review*, vol. **2-10**, 1995.