

Constructing and Sharing Perceptual Distinctions

Luc Steels

VUB Artificial Intelligence Laboratory

Pleinlaan 2 1050 Brussels

email: steels@arti.vub.ac.be and

Sony Computer Science Laboratory

6 Rue Amyot

F-75005 Paris, France

In: van Someren, M. and G. Widmer (eds.) (1997)

Proceedings of the European Conference on Machine Learning (ECML).

Springer-Verlag, Berlin.

Abstract. The paper describes a mechanism whereby agents generate perceptual distinctions through a series of adaptive discrimination games and share these distinctions through adaptive language games. Results from computer simulations as well as experiments on robotic agents are presented.

1 Introduction

Machine learning research on concept formation has mostly concentrated on inducing definitions of concepts from examples (see e.g. [3]) It is assumed that the necessary conceptualisations have been performed by a teacher and that the learning system gets positive and negative examples of the concepts to be learned. There are however many situations where examples are not available. Indeed, it could be argued that as long as the examples are supplied by a teacher, the intelligence resides in the teacher and does not originate in the learner.

This paper focuses on how an autonomous agent might be able to *construct* appropriate conceptualisations himself. It investigates how an agent may construct new distinctions relevant for discriminating objects in its environment. The construction process is based on the random generation of a variety of distinctions which then undergo selectionist pressure based on how well they discriminate. The system is open in the sense that new objects may always enter into the environment forcing the agent to refine or expand the available repertoire of distinctions.

But if every agent creates his own distinctions, we are confronted with the problem how there could ever arise coherence in a group of agents. The second part of the paper addresses this question. It is proposed that language plays a critical role in establishing a shared set of perceptual distinctions. A mechanism is therefore proposed whereby distinctions are lexicalised and used in language-based interactions. As part of an interaction (further called a language game) the agents adapt their lexicon so as to maximise communicative success in future games. The agents thus not only construct their own distinctions but their own language for verbalising these distinctions as well, all without the intervention of a teacher.

Our work has proceeded in two steps. First software simulations have been constructed to explore possible mechanisms for the creation of distinctions and the formation of a language for communicating these distinctions in a group of agents. The resulting algorithms have then been ported to physical robotic agents to demonstrate that the perceptual distinctions can be grounded in real-world sensing and actuation.

The paper is in four parts. The first section reports the algorithm for the creation of perceptual distinctions. The second section reports on the language games. Then results on coupling the two mechanisms are described. Finally some robotic experiments for the development of grounded concepts are discussed. The paper builds further on earlier results reported in [5], [6] and [7]. These papers should be consulted for additional technical detail.

2 Creating perceptual distinctions through discrimination games

2.1 Principles

Let us assume that there is a set of objects, or more generally situations, which have characteristics that are sensed through sensory channels. The datastream on a sensory channel is either the direct output of a sensor (e.g. infrared or visible light) or is the result of a low-level perceptual routine. A sensory channel yields a value in the continuous range between 1.0 and 0.0.

A discrimination tree is associated with each sensory channel. The tree divides the continuous domain into subdomains and thus maps the sensory channel into a discrete set of categories represented as features. A feature consists of an attribute value pair. The attribute expresses a path in the discrimination tree and the value the end point of the path.

The goal of the learning process is to construct the appropriate discrimination trees starting from scratch. Which discrimination trees are appropriate depends on the task and the environment. In this paper, we focus on discrimination tasks. The agent must develop an adequate repertoire of features for discriminating between the objects present in its environment. For example, if one object can be distinguished from the other objects because it has a different color, then it is necessary to have enough color distinctions. If all objects in the context have the same color but different sizes, it is necessary to make use of size-distinctions, etc.

We propose that an agent engages in adaptive discrimination games. Each game involves the following steps:

1. A context is determined or presents itself. The context consists of a set of objects and some or all sensory channels have values for each of these objects.
2. One object in this context is chosen as the topic, i.e. the object to be discriminated from the others.
3. All available discrimination trees are used to construct the features for each object.
4. A distinctive feature set is computed. It consists of the minimal set of features that distinguish the topic from the other objects in the context.

The last step may fail because there may not be enough distinctions available. In that case, a channel is picked randomly from the channels for which the topic has a value,

and the discrimination tree associated with this channel is extended by subdividing the range and thus creating a new distinction. There is no guarantee that this distinction will be adequate for discriminating the object - that remains to be seen in future discrimination games.

It is also possible that there is more than one possible set of discriminating features. In this case a choice is made based on the following criteria:

1. The smallest set is preferred. Thus the least number of features are used.
2. 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.
3. 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.

A forgetting mechanism could be implemented that eliminates superfluous distinctions, although this is not yet done in the current implementation.

The algorithm is adaptive in the sense that an agent changes its internal structure to be more successful in future discrimination games. It is top-down because the most general distinctions are created before refinements are made. It is selectionist because distinctions are created and then subjected to selection pressure coming from success in discrimination. In this respect the proposed system is related to Edelman's proposals [1] except that selection comes from discrimination and not classification. It is adaptive because it keeps expanding to cope with the steady stream of new objects entering the environment, and thus not settle into an equilibrium like a Kohonen network for example [2].

2.2 Implementation Results

The discrimination games described above has been implemented and encapsulated in (software) agents. A set of sensory channels is created and an initial set of objects is randomly generated with arbitrary values for some of the sensory channels. Then the agent starts discrimination games based on random selection of a subset of objects as context and one of its members as topic. A typical sample output of the program is as follows:

```
a-5: o-0 <-> {o-8 o-1 }
Topic:
  o-0: ((s-3 v-1)(s-4 v-1)(s-5 v-0))
Context:
  o-8: ((s-0 v-1)(s-1 v-0)(s-3 v-1)
        (s-4 v-0)(s-5 v-0))
  o-1: ((s-0 v-1)(s-3 v-0)(s-4 v-1))
No distinctive features but refinements possible.
Refining attribute: s-5 => s-5-0, s-5-1
```

Agent *a5* tries to distinguish object *o0* from *o8* and *o1*. There are already some features available. For example *o0*, which acts as the topic, has the features $((s3\ v1)(s4\ v1)(s5\ v0))$. None of the features is however sufficient. Because $(s3\ v1)$ is shared by *o8*, $(s4\ v1)$ by

$o1$ and $(s5 v0)$ by $o8$. As a consequence of this failure, a new refinement is created by a further distinction on channel 5.

Figure 1. shows two snapshots of the progressive buildup of the discrimination trees by a single agent.




Fig. 1. Fig 1. Two snapshots of discrimination trees for one agent (a-4) at different consecutive time points. One can see that on the left no tree has been made yet for channel 4, whereas on the right such a tree exists. The trees are drawn with the top node to the left.

It can be shown experimentally that an agent will develop all the discriminations needed to distinguish between a set of objects. Figure 2. shows the decreasing *failure* in discrimination for a set of 5 objects. It starts from 1.0 (complete failure) and rapidly reaches 0.




Fig. 2. Figure 2. Graph showing the average failure in discrimination (series 2 on y-axis) against the number of discrimination games (x-axis - scale 1/10). The communicative success is shown as series 1.

3 Language games

We now turn to the question how different agents, which each create their own distinctions based on the adaptive mechanism described above, are capable to share distinctions. Some degree of sharing already happens because the agents are in the same environment. The feature repertoire is determined by the kinds of objects that are encountered, and therefore agents arrive at similar discriminations. However to encourage further congruence and to be able to share information, a lexicon formation mechanism has been implemented. This mechanism allows a group of distributed agents to develop a common lexicon through a series of adaptive language games.

3.1 Principles

Let a word be a sequence of letters drawn from a finite shared alphabet. An expression is a set of words. In the experiments reported here, word order does not play a role. A lexicon L is a relation between feature sets and words. A single word can have several associated feature sets and a given feature set can have several associated words. For each word-meaning pair, the use and success in use is recorded. Words that are used more and have more success are preferred. This establishes a positive feedback loop pushing the group towards coherence.

Each agent $a \in A$ is assumed to have a single lexicon L_a which is initially empty. A feature set of a word in L is denoted as $F_{w,L}$. The following functions can be defined:

- $cover(F, L)$ defines a set of expressions U such that $\forall u \in U, \{f \mid f \in F_{w,L} \text{ and } w \in u\}$
- $uncover(u, L)$ defines a feature set F such that $F = \{f \mid f \in F_{w,L} \text{ and } w \in u\}$.

A language game involves a dialog between two agents, a speaker and a hearer, within a particular contextual setting which consists of objects of which one is chosen as the topic. The agents perceive these objects through the sensory channels and construct features through the discrimination trees discussed earlier.

The scenario for a language game is as follows:

1. A speaker and hearer as well as a context consisting of a set of objects is randomly identified.
2. The speaker selects one object as the topic and points to this object so that the hearer shares the topic.
3. Both speaker and hearer identify possible distinctive feature sets using the perceptions on the sensory channels and the discrimination trees.
4. If there is at least one set of discriminating features, the speaker selects such a set and translates it to words using the cover function. Words that have been used the most and have been most successful in use are preferred. An important additional criterion for selecting the set of possible discriminating features is that those sets are preferred which can be expressed in language. This causes those features to be preferred that have been lexicalised and establishes a positive feedback loop pushing the agents towards coherence both for the lexicon and for the feature repertoire.

5. The hearer interprets this expression using the uncover function and compares it with his expectations.

As a side effect of such a language game, various language formation steps take place:

1. *No differentiation possible* (step 3 fails): In this case new features are created as discussed in the previous section.
2. *The speaker does not have a word* (step 4 fails): In this case at least one distinctive feature set S is detected but the speaker s has no word(s) yet to express it. The language game obviously fails. However the speaker may create a new word (with a probability $p_w = 0.05$) and associate it in his lexicon with S .
3. *The hearer does not have a word*: At least one distinctive feature set S is detected and the speaker s can construct an expression to express it, i.e. $cover(S, L_s) = W$. However, the hearer does not know the word. Because the hearer has a hypothesis about possible feature sets that might be used, he is able to extend his lexicon to create associations between the word used and each possible feature set. If there is more than one possibility, the hearer cannot disambiguate the word and the ambiguity is retained in the lexicon.
4. *The speaker and the hearer know the word*: In this case there are two possible outcomes:
 - (a) *The meanings are compatible with the situation*: The dialog is a success and both speaker and hearer achieve communicative success. Note that it is possible that the speaker and the hearer use different feature sets, but because the communication is a success there is no way to know this. Semantic incoherences persist until new distinctions become important and disambiguate.
 - (b) *The meanings are not compatible with the situation*: The same situation as before may arise, except that the feature set uncovered by the hearer is not one of the feature sets expected to be distinctive. In this case, there is no communicative success, neither for the speaker or the hearer.

3.2 Implementation Results

The language games described above have been implemented and encapsulated in software agents. The simulation experiments consistently show that a common repertoire of word-meaning pairs develops conjointly with the development of a feature repertoire. A typical language game is the following:

```
LANGUAGE GAME 878.
  Speaker: a53.
  Hearer: a49
  Topic: o-8
  Context: {o-3 o-7 o-1 o-4 o-9 o-6 }
  Speaker expresses: [(a53-s4 v0)]
  Hearer expects: [(a49-s4 v0)]
  a53: [(a53-s4 v0)] => ((G U))
  a49: => [(a49-s4 v0)]
success
```

Agent *a54* converses with *a49* about *o8*. Both speaker and hearer have identified the same distinctive feature set. It is translated into the expression ‘((G U))’ by the speaker which matches the expected distinctive feature uncovered from this expression by the hearer. The language game therefore ends in success.

Figure 2 shows for a group of agents both the decreasing failure in discrimination (which evolves towards 0.0) and the increased success in communication (which evolves towards 1.0). Figure 3. shows the discrimination trees developed by agents operating in the same environment. Some interesting observations can be made: We see that there are on the one hand strong similarities between the agents. For example, the discrimination trees for channel 4 look almost identical. At the same time we see that there are differences. For example, for channel 2 we see that agent 3 has an elaborate discrimination tree whereas agent 2 has almost no distinctions. Also when we inspect the lexicons of the different agents we see important similarities (how else could there be complete communicative success) but at the same time we see differences. These differences are maintained because the environment allows multiple possibilities for discrimination and the same word may therefore be associated with different features without being noticed. New objects coming into the environment may disambiguate words or may cause some of the agents to develop distinctions shared by others.




Fig. 3. Figure 3. Overview of different discrimination trees for two different agents.

Both the feature formation and language formation processes are open. The (distributed) lexicon adapts itself when new features are created. New words are created and sentences become more complex. The system is also open with respect to the number of agents: New agents may enter at any time in the population. The new agent will gradually take over words already present but is also a new source of novelty.

4 Grounding experiments

In our laboratory, we are conducting various experiments in sensory-motor intelligence with robotic agents. The agents 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 [4]. Experiments are going on to carry to port the mechanisms reported in this paper onto the physical robots operating in this ecosystem (see [7] for a more extended discussion). The various mechanisms described in the present paper are mapped on a physical counterpart as follows.

The sensory channels contain the real world data obtained from the physical sensors. An example of such data is given in figure 4. 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 these crossing points act as input to the discrimination games.




Fig. 4. Figure 4. Sensory data streams taken from physical robot. The channels include left and right infrared and visible light sensing and motor speeds.

The protocol for engaging in language games has been implemented 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 moving towards the topic, so that the other robot can observe in which direction it moves. The speaking robot halts when it is facing the object that it wants to see as the topic of the conversation. The listening robot detects the topic by consulting its own sensory map. Then the language game starts as described above.

Although many technical problems still remain to be resolved to enable larger scale experiments, the experiments so far have been successful in showing that the proposed

mechanisms do indeed lead to an adequate repertoire of perceptually grounded distinctions and a shared language for using these distinctions to identify objects in the environments.

5 Conclusions

This paper has proposed a mechanism for creating perceptually grounded distinctions driven by adaptive discrimination games. It has also proposed a mechanism for creating a shared set of words to express the features resulting from perception and thus to share the perceptual categories in a group of distributed agents. Both mechanisms are selectionist: There is a way to create variation (by the expansion of discrimination trees or the introduction of a new word-meaning pair) and to select appropriate variations based on their use (success in discrimination or communication).

The proposed mechanisms have two interesting characteristics from a machine learning point of view. First of all there is no distinction between the construction of new distinctions or new word-meaning pairs and the learning of distinctions/words. The agent always constructs and then tests whether the distinctions or word-meaning pairs have the desired effect. Thus the agents are not only capable to acquire existing distinctions or lexicalisations but also to create new ones. We can therefore explain both the origin *and* acquisition of concepts and words by the same mechanism. This contrasts with most inductive learning algorithms which require a teacher to make prior conceptualisations or lexicalisations.

Second, although the proposed mechanisms are selectionist they are not genetic (as in the case of genetic algorithms). Agents do not have to reproduce in order to induce a change. The selectionist forces do not operate on the agents but on the structures inside agents. The result is nevertheless an evolutionary process but it is a cultural as opposed to genetic evolution.

6 Acknowledgement

The author's research, including the simulation experiments reported in the paper, were financed by the Sony Computer Science Laboratory in Tokyo and Paris. Many people have contributed to realise the physical grounding experiments. Robot builders at the VUB AI Lab (soft and hardware) include Andreas Birk, Tony Belpaeme, Peter Stuer and Danny Vereertbrugge. Ruth was the first one to try the grounding experiment on real robots during a short term visit to the laboratory. These experiments were continued by Paul Vogt, a masters student from the university of Groningen, visiting the Brussels laboratory.

References

1. Edelman, G.M. 1987. *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books.

2. Kohonen, T. (1988) *Self-Organization and Associative Memory*. Springer Series in Information Sciences. Vol 8. Springer Verlag, Berlin.
3. Langley, P. (1987) Machine Learning and Concept Formation. *Machine Learning Journal* 2, pp. 99-102.
4. Steels, L. (1994) The Artificial Life Roots of Artificial Intelligence. *Artificial Life Journal* 1(1), pp. 89-125.
5. 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.
6. Perceptually grounded meaning creation. In: Tokoro, M. (ed.) (1996) *Proceedings of the International Conference on Multi-Agent Systems*. The MIT Press, Cambridge Ma.
7. Steels, L. and P. Vogt (1997) Grounding language games in robotic agents. (to appear)