Meaning Creation by Object Discrimination

Andrew D. M. Smith
andrew@ling.ed.ac.uk
Language Evolution and Computation Research Unit

ABSTRACT

The emergence and maintenance of a shared system of syntax among a population of language users has been an important topic for much recent work on human language evolution. Most of this work, however, takes as innately given the existence of particular arbitrary meanings within the users' brains, and also assumes that these meanings can be accurately transmitted to other users in tandem with the transmission of signals. This paper looks at the creation, evolution and development of meanings themselves, presenting a model in which agents use abstract feature detectors to differentiate a particular object from others. The model results in the creation of categories and meanings reflecting the agents' environment. Having developed such a model, we will then look at models in which the use of language can trigger the dissemination and maintenance of a shared set of meanings throughout a population, without relying on explicit meaning transmission.

1. INTRODUCTION

1.1. Language Evolution

Interest in understanding the origins and evolution of human language has increased dramatically in recent years, since Pinker & Bloom (1990) argued that "a specialisation for grammar evolved by a conventional neo-Darwinian process". This means, in their view, that humans have evolved an innate, genetically specified module in the brain (the Chomskyan Language Acquisition Device or LAD), which is a genetic coding of the principles and parameters of Universal Grammar. Although children clearly learn their particular language through cultural interaction within their community, this is only made possible in the first place by the existence of the LAD. It constrains the set of possibilities and provides humans with prior knowledge about the structure of language, allowing children to succeed in language learning where general learning techniques would fail.

Others have argued that it is not necessary to assume the evolution of a highly specialised language module, and have developed other theories to account for the evolution of language. Carstairs-McCarthy (1999) argues that language is a "by-product of a change in the anatomy of the vocal tract", coupled with an "expectation that different vocalizations should mean different things", while Bickerton (1998) argues that combining the social intelligence needed for a system of reciprocal altruism with a primitive protolanguage would have "led directly to a cascade of consequences that would, in one rapid and continuous sequence, have transformed protolanguage into language substantially as we know it today".

1.2. The Evolution of Syntax

Arguments such as these have provided us with many useful insights, but it is very difficult to evaluate them without some kind of plausible cognitive model about how language works, and how it evolved.

Many researchers have, therefore, developed computer models of language users (e.g. Hurford (1991), Batali (1994), Oliphant (1996), Kirby & Hurford (1997)), producing working simulations of the development and maintenance of certain aspects of language-like systems, which can be observed and investigated. Kirby (in press) has shown that in a population of users who have the ability to analyse their language, syntax arises inevitably, as the users take advantage of coincidental matches between parts of utterances and parts of corresponding meanings to create general rules. The simple fact that general rules can produce more utterances than idiosyncratic ones and so the utterances produced by them are, other things being equal, more likely to be produced (and therefore heard), accounts for the emergence of a compositional, recursive syntax. A similar result has been shown by Batali (in press), in a simulation where users hypothesise mappings between strings and meanings, given a set of examples. The users in Batali's simulation are explicitly given the ability to combine and modify phrases, rather than the ability to generalise, but again they take advantage of coincidental correspondences between utterances and meanings to develop syntax.

These and similar simulations concentrate on the evolution of syntax, and provide evidence for how it can evolve from holistic expressions, given a pre-existing structured (predicate-argument or an equivalent) semantic representation which is defined as being compositional and recursive. In order to make this a fully cogent analysis, we need to provide some explanation of how such a structured semantics could arise, either before syntax, or possibly simultaneously. Following on from Steels (1996) and (1999), this paper looks at the creation of meanings, to see whether we can provide the missing explanation.

2. OBJECT DISCRIMINATION

2.1. An imaginary world

In order to explore ways in which meanings could be created from agents' interaction with the world around them, we set up a simple, idealised world, based on that described by Steels (1999), in which there are a number of *objects*, which in turn are described in terms of features. In this world, there also exist a number of *agents*, who we hope will be able to interact with their environment and create meanings to describe it. We would like the agents' meaning system to be flexible enough to cope with a variety of different circumstances, but which doesn't have any pre-defined, innate meanings.

The features which define the objects in the world are completely devoid of any intrinsic meaning. Each object has a fixed number of features, the values of which are generated pseudo-randomly with a uniform distribution and are normalised to lie between 0 and 1. The agents have channels which are sensitive to the objects' features, and we set the world up so that they have exactly one sensory channel corresponding to each of the objects' features. The agents can, through their sensory channels, sense the features, but they cannot measure the actual value of the feature. Instead, they can detect whether or not a particular feature falls into a certain range of values. In terms of language, this might be analogous to being able to look at a tree and say that it was tall, without being able to give its exact height.

Crucially, the agents are also able to refine their sensory channels, by splitting the range to which the channel is sensitive into discrete segments. When the channels are repeatedly refined, the resulting structure is that of a binary tree, and they can be referred to as *discrimination trees*. Each node on a discrimination tree can be regarded as a *category*. Figure 1 shows an example of a discrimination tree, showing both the values each node is activated for, and a conceptualisation of one set of possible meanings of the nodes on the tree. In this example we have superimposed a pre-defined meaning 'height' onto the discrimination tree, but in actual fact the categories have no intrinsic meaning at all.

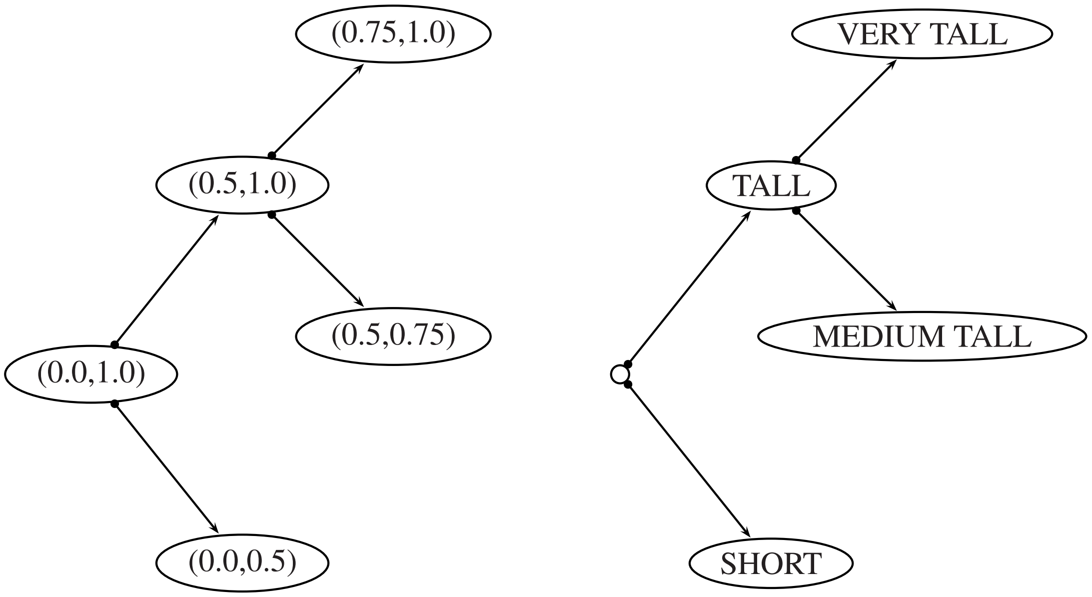

Figure 1: A sensory channel represented as a binary tree, after two refinements. The numbers at each node give the range of values for which it is activated (left), and the same nodes are superimposed with the meaning 'height' (right).

Object	Value	Category	
		Abstract	'Height'
A	0.43276	0-0	SHORT
B	0.87423	0-11	VERY TALL
C	0.12098	0-0	SHORT
D	0.50012	0-10	MEDIUM TALL
E	0.61419	0-10	MEDIUM TALL

Figure 2: Categorisation of objects, given the discrimination tree 0 in figure 1. The categories are shown both in abstract terms, and with the 'height' meaning superimposed.

The agents can then categorise objects using their discrimination trees, as shown in figure 2. In order to keep an abstract representation of the categories, we use the notation $channel$-$path$, where $channel$ is the number of the sensory channel, and $path$ is the path from the root of the discrimination tree to the category. A path is represented by successive instances of 0s (when the lower branch of the tree is taken) or 1s (when an upper branch is taken).

2.2. Discrimination games

Steels (1999) describes discrimination games, where an agent attempts to discriminate one object from a group of others, and then uses the result of the game to develop its representation of the environment. In a discrimination game, a random agent is given a random set of objects *(the context)*, from which one it must discriminate one particular randomly-chosen object *(the topic)*. In attempting the discrimination, the agent investigates all its sensory channels, and uses them to categorise the objects, as in figure 2. If the topic is uniquely identified by any category, then the game is a success - if not, a failure.

For example, given the objects and sensory channel in figure 2, and the agent's task was to discriminate topic A from the context BDE, then the game would be a success: the topic's category 0-0 is a *distinguishing category* because it is different from the categories of all the other objects in the context (0-11, 0-10 and 0-10). If the agent had had many sensory channels, and in performing the task had arrived at more than one distinguishing category, then only one of them is chosen, according to the following criteria:

1. the category chosen is as abstract as possible — with the fewest number of segmentations in the tree;
2. the most frequently used category is preferred — so that a minimal set of categories is used.

If, however, the task had been to discriminate topic D from the context ABE, again assuming that the agent only has the sensory channel in figure 1, then the game would fail, as the topic's category 0-10 is not different from all the others in the context. If the game is a failure, then the agent refines one of its trees, at the node which represents the topic's category. For instance, given the above game of discriminating topic D from context ABE, the tree in figure 1 would be refined to that in figure 3. There is, however, no guarantee that this refinement will actually prove sufficient to discriminate any objects in subsequent games, and any node can continue to be refined if necessary in the future.

It is important to note that the trees are not created intelligently, but evolve blindly, without any preconceived strategy, yet they form a hierarchical structure which can distinguish objects from each other, as we shall see.

2.3. Starting simply

Let us consider a simple world containing one agent and twenty objects. Each of the objects has just a single feature, and so the agent has one sensory channel, which is initially not refined at all. In order to explore the dynamics of the discrimination game system, the following measures are taken at regular intervals throughout the simulation:

1. the percentage of successful games (discrimination success rate);
2. the percentage of objects which can be distinguished from all other objects in the world (object discriminability).

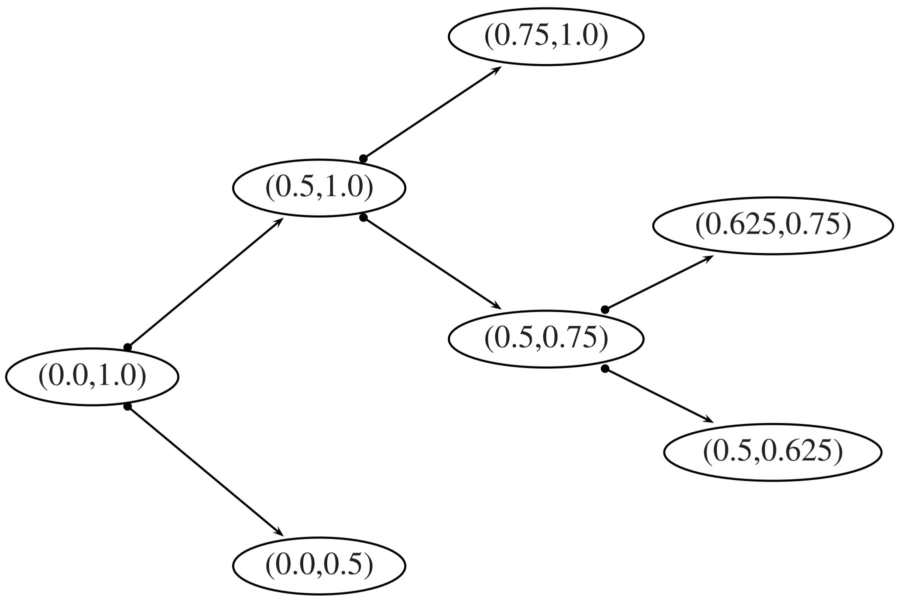

Figure 3: The discrimination tree refined after the agent failed to discriminate object D from the context ABE.

The discrimination success rate is fairly self-explanatory; object discriminability can be thought of as the percentage of objects which can be uniquely identified in the world by the agent, so that they can always be identified, whatever the context.

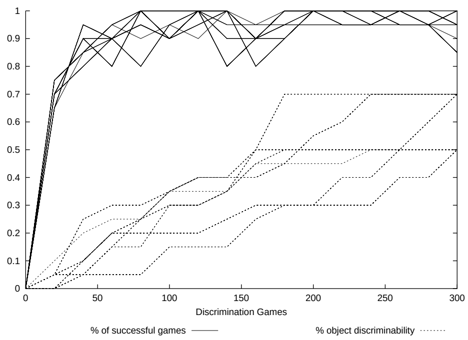

Figure 4: One agent and 20 objects with one feature — context fixed at two objects.

Initially, the context size is fixed at two objects, i.e. the topic must be distinguished from just one other object, and measurements are taken after every ten games. Figure 4 shows twenty different simulations on this world plotted on top of each other, with the percentage of successful discrimination games shown with solid lines, and object discriminability shown with dashed lines. We can see that in all the simulations the percentage of successful games rises rapidly from zero to 80 or 90 % within about 30 games (although does not always reach complete success), but that the discriminability of the model is much more variable across simulations, but varying between 70 and 100% after 500 games.

If we take a closer look at the dynamics of the simulations, we can see how the agent is evolving its sensory channel to take account of the objects' values. At the start of the simulation, the unique sensory channel 0 is unrefined, and so the agent will categorise all objects as 0-, leading to the first game always being a failure. Equally necessary is that after the inevitable failure, the channel will be refined into two sections, with ranges (0–0.5) and (0.5–1). There is now a possibility that the next game will be a success, as long as the topic is categorised on one node and the context on the other. The probability of success is:

$$\sum_{i=1}^{k} p(i) \times 1 - p(i), \tag{1}$$

where k is the number of terminal categories (leaf nodes on the tree), and $p(i)$ is the probability of the topic occurring being categorised by category i. Given the uniform distribution of the objects' feature values, and the fact that they are normalised between 0.0 and 1.0, the probability of an object being categorised by a particular category is exactly equivalent to the range of the category, or the difference between the lower and upper bounds. After one refinement, therefore, the probability of success is

$$\left(\frac{1}{2} \times \frac{1}{2}\right) \times 2 = 0.5.$$

It can easily be seen that the more refined the tree is, the higher the probability of success is in the next game. As refinements only occur after a discrimination game failure, we have the position where every failure increases the probability that the next game will succeed. In other words, *the agent is learning from its mistakes.*

The discriminability measure, however, is much more volatile. Discriminability is defined as the percentage of objects which can be distinguished from all other objects in the world, so we only get discriminability when one of the categories in the tree uniquely identifies an object. So the *maximum* amount of discriminability at any time is given by:

$$\text{Max. Discriminability} = \frac{\text{No. of categories}}{\text{No. of objects}}. \tag{2}$$

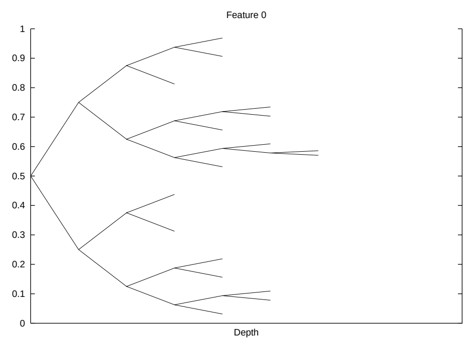

Figure 5: The agent's discrimination tree after 300 discrimination games.

Figure 5 shows the discrimination tree after 300 games on a representative simulation from those shown in figure 4. The pattern of the tree's refinement reflects the distribution of the feature values of the objects in the world - if they are very close together, then there is likely to be deep refinement; if they are evenly spread, then the refinement will also be much more shallow and even.

Although the tree is refined quite deeply in parts, and the probability of success in the next game is high (93%), there are still only 17 terminal categories, and so the *maximum* discriminability is 85% (equation 2). Every refinement of a tree after a discrimination game failure increases the number of categories by one, and so increases the potential discriminability, but also, as we have seen above, it reduces the chance of any further refinement, by increasing the likelihood that the next game will succeed.

2.4. Increasing the size of the context

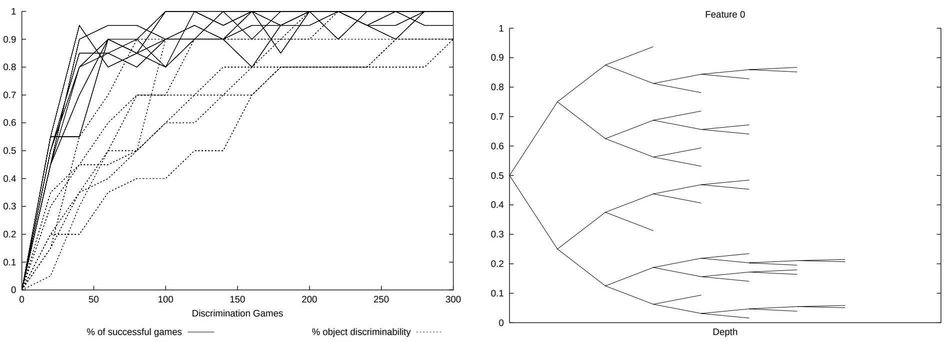

Figure 6: One agent and 20 objects with one feature — context fixed at five objects. Success rate and discriminability (left); discrimination tree from a representative run (right).

Figure 6 shows what happens in the same world of one agent and twenty objects, but when the context size of each discrimination game is increased to five objects (i.e. the topic and four others from which it must be distinguished). We can see that the discriminability is much increased compared to figure 4, and the success rate is lower early in the simulation, but eventually rises and ends up at a similar level to figure 4. This is mirrored in the discrimination tree itself, which is much more deeply refined than that in figure 5, and has more terminal categories. Both these results suggest that when the context is increased, the success rate early in the simulation is reduced, which therefore means that the agent develops its discrimination trees to a greater extent. Mathematically, we can see that this is true, as we modify equation 1 to take account of the size of the context:

$$\sum_{i=1}^{k} p(i) \times (1 - p(i))^{z-1}, \tag{3}$$

where z is the size of the context, and $z-1$ is therefore the number of objects which must be in a different category than the topic for the game to succeed. Clearly, for any particular number of terminal categories k, the success rate in the next game is inversely proportional to the context size z. The low success rate with high z, however, is countered by the successive refinements (and therefore increases in k) which each failure brings.

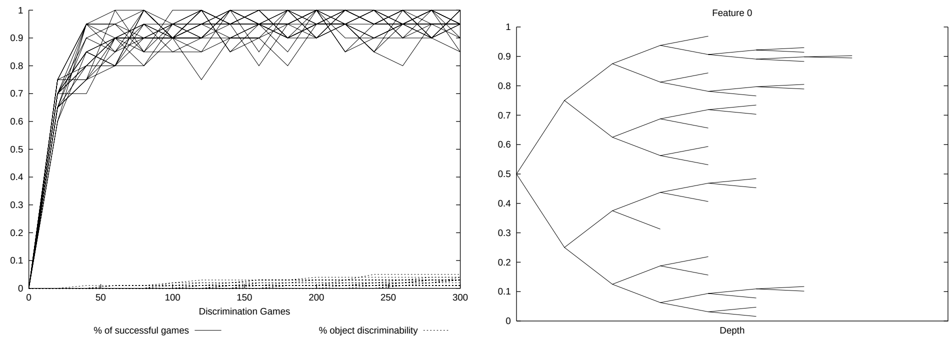

Figure 7: One agent and 100 objects with one feature — context fixed at two objects. Success rate and discriminability (left); discrimination tree from a representative run (right).

2.5. Increasing the number of objects

Instead of increasing the size of the context, we can change the world in which the agent finds itself by increasing the number of objects which it meets. Figure 7 shows the results for a world of 100 objects, with the context size reduced back to two (the topic and one other), and with one feature available. Comparing this to figure 4, we can see that the success rate follows a similar path, but the discriminability is massively reduced, not even reaching 10% in any of the simulations. If we look at the discrimination tree in figure 7 in conjunction with equation 2 (page 6), it is immediately obvious that the number of terminal categories on the tree (26 in this example), although higher than in previous scenarios, is never going to be enough to produce high discriminability. However, despite the inevitable low discriminability, the games can succeed because the proportion of objects in the world involved in each individual game is very low.

Indeed, a world with 100 objects is still very unrealistically small, and we can see that as we increase the number of objects still further, the chance of achieving any discriminability at all with one feature becomes vanishingly improbable.

2.6. Increasing the number of features

In addition to changing the number of objects, both in the context and in the world as a whole, we can introduce more features into the simulation.

Figure 8 shows the results when each object is defined with five features, and so the agent has five discrimination trees to use, any one of which could identify the topic. This time, the discriminability is both much reduced (compared to the baseline in figure 4), but also much more variable (compared to figure 7).

The formula for the probability of success needs to be altered, to take account of the increase in features. Strictly speaking, formula 3 shows $P_{dg}(sc)$, the chance of success in a discrimination game dg on a particular sensory channel sc, reproduced below:

$$P_{dg}(sc) = \sum_{i=1}^{k} p(i) \times (1 - p(i))^{z-1}.$$

We now have more than one sensory channel, and only one of them needs to discriminate the topic in order for the game to succeed, so the game only fails if *none* of the sensory channels can discriminate.

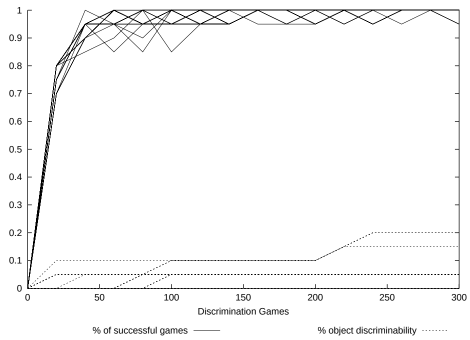

Figure 8: One agent and 20 objects with five features — context fixed at two objects.Success rate and discriminability.

The probability of a discrimination game failure on a sensory channel sc is $1 - P_{dg}(sc)$, and so the probability of success on any channel is:

$$P_{dg} = 1 - \prod_{j=1}^{s} 1 - P_{dg}(j) \tag{4}$$

where s is the number of sensory channels.

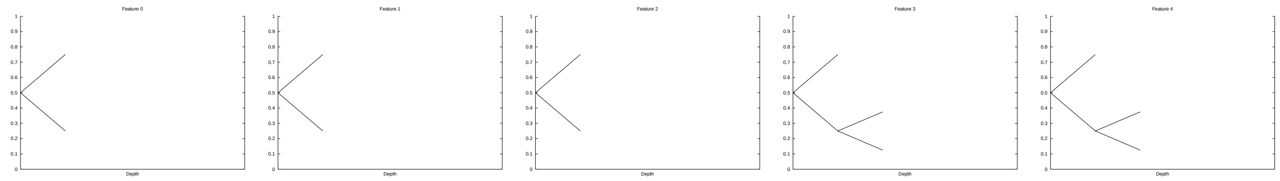

Figure 9: One agent and 20 objects with five features — context fixed at two objects. Discrimination trees from a representative run.

We would expect, therefore, that increasing the number of features will, as long as the features are refined at least once (and so $P_{dg}(sc)$ is not 0), increase the success rate, and therefore reduce discriminability. Figure 9 shows the five discrimination trees stored by the agent on a representative run. We can see that the trees themselves are much less refined than in previous simulations, and the success rate is indeed higher than the other simulations, rarely being below total success after 300 games.

2.7. How frequently is an object encountered?

All the simulations presented so far have looked at a world in which all the objects are equally likely to occur in a discrimination game. It is more reasonable to assume that certain objects would be encountered more regularly than others in reality, and we would expect the agents to be able to discriminate these objects better than infrequently-met objects.

Figure 10 shows the result of simulations where a quarter of the objects were fifty times more likely to be chosen in the discrimination game than the others. The feature values of these objects were also modified, to make them within the range 0-0.25, for illustrative purposes. We can clearly see that the

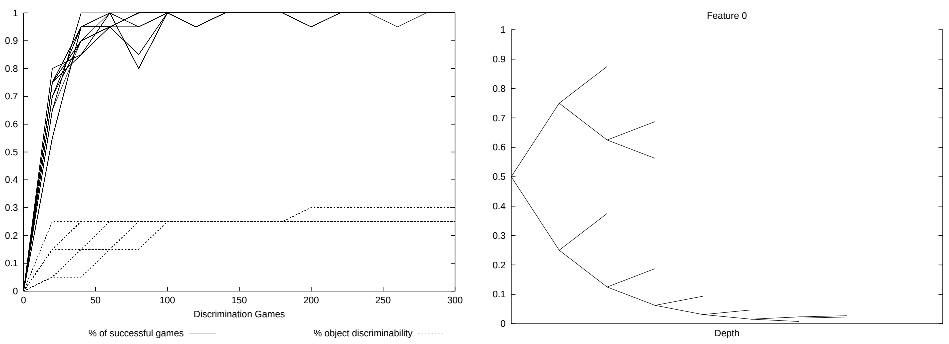

Figure 10: One agent and 20 objects with one feature — context fixed at two objects. A quarter of the objects were fifty times more likely to be chosen than the others. Success rate and discriminability (left); discrimination tree from a representative run (right).

agents' tree is much more developed in the lower quarter of the meaning space, where the frequently-encountered objects are situated, and that very little refinement has occurred on the rest of the tree. The success rate is also very high, because the agent has successful identified all the frequently-occurring objects, so that whenever they are the topic, the game succeeds.

2.8. Different representations of meaning

We have seen that, when the number of features increases, the success rate also increases, and the features themselves are much less refined (see section 2.6.) Another effect of the increase in features is that the agents themselves have more flexibility about which channels should be refined. If each agent has only one sensory channel, then that channel will always be chosen and will inevitably, in the limit, produce a faithful translation of the world into categories which uniquely identify certain objects. When the agent has more than one channel, however, then there is a choice to be made, as any of the channels could be refined, and so there are many different possible ways in which the agents can 'view the world.'

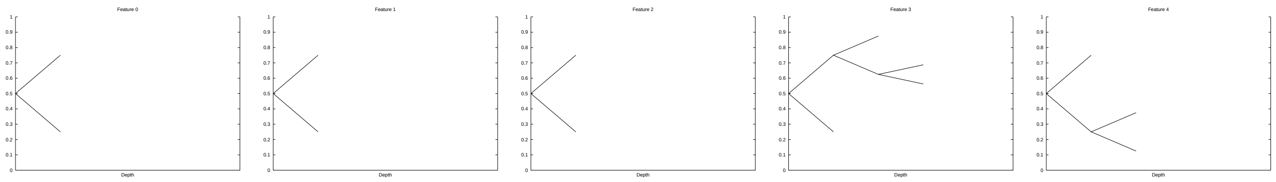

Agent One

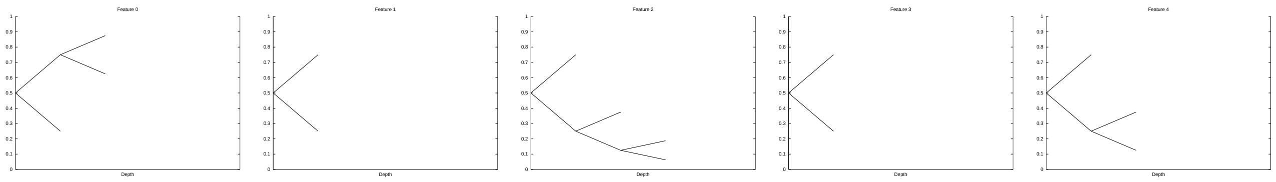

Agent Two

Figure 11: Two agents' representations of a world of twenty objects.

Figure 11 shows two agents' representations of the same world. Their representations of the world are quite different, reflecting the features which that agent actually refined, and which it then uses in the discrimination games. Each agent has a different feature which is refined deepest, and the distribution of the refinements is also different, as expected. In this run all the trees are refined at least once, and two of the features produces the same tree for both agents, though this is not inevitable, and does not happen on all runs. Neither of these representations is 'better' than the other, however. Despite the differences, the agents are just as likely to be successful in the discrimination games, and can therefore describe their environment with similar accuracy.

The agents in figure 11 have only five channels, and are only considering twenty objects. As we scale the model up, increasing the number of features available to the agents, we find inevitably that as the number of possible representations of the world increases, the actual representations recorded by the agents diverge even more strongly, producing meaning structures which are incompatible with each other.

2.9. So far, so good

The model we have looked at so far has been reasonably successful. From the very simple structure of a discrimination tree, and given the ability to refine the tree when discrimination fails, the agents have been able to develop representations of meaning, which allow the agents to successfully distinguish the objects from one another. The representations are *flexible* and *dynamic*, developing in those areas of the meaning space where there are frequently-encountered objects, and not wastefully areas where they need not, so that they evolve to reflect the environment in which the agent finds itself. They are *variable*, with different agents producing different, sometimes incompatible, but equally valid, representations of the same environment. We have also seen that by changing simple parameters in the simulations, particularly by modifying the number of objects, the number of features/sensory channels and the size of the context, we can get varying levels of object discriminability, and therefore varying levels of *abstractness* in the meanings developed.

We now turn our attention to communication, and investigate whether and how the meaning representations can be shared among the agents.

3. COMMUNICATION

In order to simulate communication between the agents, we have to endow them with the ability to create signals, or words, which they will use to express meanings. This also implies that they will have a lexicon of associations between words and meanings, which they develop and change as they use language. We will also assume, clearly, that the agents can express and understand these words without any difficulty, i.e. the signals can be transmitted between agents.

The transmission of meanings, however, is not so straightforward. Many simulations in this field (e.g. (Batali 1998), (Kirby in press)) assume that signals and meanings are both transmitted between agents with equal ease, and that the agents can then analyse correspondences between the two. The problem we face is that in human language, meaning is *not* explicitly transmitted, but is somehow derived by the hearer. Steels (1999) calls this the *gavagai*-problem, after Quine (1960), who illustrated the issue by imagining an anthropologist, who observes a speaker of an unknown language uttering the word "*gavagai*" while pointing at a white rabbit scurrying past. The anthropologist now has a problem, and cannot decide whether the utterance "*gavagai*" means *'rabbit'*, *'whiteness'*, *'what was that?'*, *'look over* *there!'*, or even *'dinner!'*. We can see that there are countless possible meanings, and the anthropologist has no way of knowing which is meant for sure, without hearing the word used in other contexts.

The aim of the next part of this continuing work is to investigate and explain whether and under what conditions a shared system of meaning can emerge between agents, without the explicit transmission of meanings.

References

Batali, J. 1994. Innate biases and critical periods: Combining evolution and learning in the acquisition of syntax. In R. Brooks & P. Maes (eds). *Artificial Life 4: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pp.160–171. Redwood City, CA: Addison-Wesley.

——. 1998. Computational simulations of the emergence of grammar. In Hurford *et al.* 1998.

——. in press. The negotiation and acquisition of recursive grammars as a result of competition among exemplars. In Briscoe in press.

Bickerton, D. 1998. Catastrophic evolution: the case for a single step from protolanguage to full human language. In Hurford *et al.* 1998.

Briscoe, E. (ed). in press. *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge: Cambridge University Press.

Carstairs-McCarthy, A. 1999. *The Origins of Complex Language*. Oxford University Press.

Hurford, J. R. 1991. The evolution of critical period for language acquistion. *Cognition* 40: 159–201.

——, M. Studdert-Kennedy & C. Knight (eds). 1998. *Approaches to the Evolution of Language*. Cambridge University Press.

Kirby, S. in press. Learning, bottlenecks and the evolution of recursive syntax. In Briscoe in press.

—— & J. R. Hurford. 1997. Learning, culture and evolution in the origin of linguistic constraints. In P. Husbands & I. Harvey (eds). *Fourth European Conference on Artificial Life*, pp.493–502. Cambridge, MA: MIT Press.

Oliphant, M. 1996. The dilemma of saussurean communication. *BioSystems* 37: 31–38.

Pinker, S. & P. Bloom. 1990. Natural language and natural selection. *Behavioral and Brain Sciences* 13: 707–784.

Quine, W. V. O. 1960. *Word and Object*. MIT Press.

Steels, L. 1996. Perceptually grounded meaning creation. In M. Tokoro (ed). *Proceedings of the international conference on multi-agent systems*, Cambridge, MA. MIT Press.

——. 1999. *The Talking Heads Experiment*, volume I. Words and Meanings. Antwerpen: Laboratorium. Special pre-edition.