

The Role of Environment Structure in Multi-Agent Simulations of Language Evolution

Mark Bartlett, Dimitar Kazakov*

* Department of Computer Science
University of York
Heslington, York YO10 5DD

bartlett@cs.york.ac.uk, kazakov@cs.york.ac.uk

Abstract

This paper presents a multi-agent system which has been developed in order to test our theories of language evolution. We propose that language evolution is an emergent behaviour, which is influenced by both genetic and social factors and show that a multi-agent approach is thus most suited to practical study of the salient issues. We present a hypothesis that the original function of language in humans was to share navigational information, and show experimental support for this hypothesis through results comparing the performance of agents in a series of environments. In particular, we study how the degree to which language use is beneficial varies with a particular property of the environment structure, that of the distance between resources needed for survival.

1 Introduction

The use of computer simulations to study the origin and evolution of language has become widespread in recent years (Briscoe, 2002; Kirby, 2002; Steels, 1999). The goal of these studies is to provide experimental support for theories developed by linguists, psychologists and philosophers regarding the questions of how and why language exists in the form we know it. As verbal language by its very nature leaves no historical physical remains, such as fossils, computer simulation is one of the best tools available to study this evolution, allowing us to construct a model which simulates the relevant aspects of the problem, and to abstract away any unnecessary detail. Experiments conducted using such an approach can be performed much more quickly than those involving teaching language to apes or children and allow researchers to study language in situations which would be impossible with live subjects.

Our multi-agent system has been designed to test the theory that language could potentially have evolved from the neural mechanisms our ancestors used to navigate (Kazakov and Bartlett, 2002; Hauser et al., 2002; O’Keefe and Nadel, 1978). Specifically we wish to explore the feasibility of the hypothesis that the original task of speech in humans was to inform others about the geography of the environment. To this end, we have constructed an artificial-life environment in which a society of learning agents uses speech to direct others to resources vital to survival, using the same underlying computational mechanism as they use to navigate. We discuss the altruistic nature of this activity and evaluate the benefits this brings to the community by measuring the difference in

performance between populations of agents able to communicate and those unable to do so. We also evaluate the performance of a population of agents engaging in a non-communicative act of altruism, namely sharing stores of resources. Our simulations are carried out in a series of environments in which the distance between the two resources needed for survival is varied to assess what impact this feature of the terrain may have on any benefit that language may bring. We also study how the distance between resources of the same type may affect performance. Our aim in this paper is to assess the conditions in which language use may be beneficial and hence could be selected for by evolution.

We have chosen to carry out our simulations within a multi-agent setting as the nature of these systems allows us easily to capture much of the behaviour we wish to program into our models. We assume that language has both innate components (such as the willingness to speak) and social components (such as the words used in a language). Multi-agent systems allow both these aspects to be modelled relatively simply. Though the genetic nature of language could be modelled equally well using only genetic algorithms, indeed some researchers have done this (Zuidema and Hogeweg, 2000; Oliphant, 1997), in order to model the social aspects of language which rely on phenotypical behaviour, situatedness, grounding and learning, it is necessary to employ an agent model.

Simulations of the evolution of language using the multi-agent paradigm can also be of interest to the designer of any general-purpose agent-based application. In a dynamic environment that is expected to change considerably during an agent’s lifetime, the faculty of learning could be essential to its success. In an evolutionary MAS

setting, sexual reproduction and mutation can be used to explore a range of possible learning biases, from which natural selection would choose the best. One would expect Darwinian evolution to advance in small steps and select only very general concepts. One could also implement Lamarckian evolution, that is, use a MAS in which the parent's individual experience can be reflected in the learning bias inherited by their offspring. Lamarckian evolution is faster but brings the risks of inheriting concepts that were relevant to the parents but are not reflected in the new agent's lifetime. This work explores what could be a third way of evolving a learning bias, through communication. Language uses concepts that are specific enough to be useful in the description of an agent's environment, yet general enough to correspond to shared experience. In this way, language serves as a bias inherited through social interaction rather than genes. In the current work, language serves as a way to share a bias for navigating in an environment.

2 Altruism and Sharing

Language use of the kind presented in this work is an inherently altruistic act; by informing another agent of the location of resources, an agent increases both the survival prospects of his rival and the probability that the resource mentioned will be exhausted before he is able to return to it. In our previous work (Turner and Kazakov, 2002) which focussed on resource sharing as a form of altruism, we found that through the mechanism of kin selection (Hamilton, 1964; Dawkins, 1982) it is possible to promote altruistic acts in an agent population if the degree of relatedness between the agents is known and the amount of resource given is able to be measured and controlled. It was also found that a poor choice of sharing function (which is used to decide the amount of resource shared) could lead to a population in which all agents had enough energy to survive, but insufficient to reproduce. This can lead to extinction of a large population even in the circumstances where a smaller population would be able to survive.

Quantifying the cost of the act is much harder to achieve when the altruistic act is to share knowledge of a resource location. In place of the one-off payments of resources that occur when resources are shared, informing an agent of the location of a resource can involve a long-term, sizeable drain on the resource, especially if the receiver goes on to inform others. Conversely, the act may result in no cost at all in the case where the receiver is in such need of the resource that it is unable to reach it before dying. This point also illustrates another problem with using language to share, it is also hard to estimate the benefits that the altruistic act will bring to the listener. When resources are given directly to another agent, the amount of help is easy to quantify, whereas when language is used the reward is delayed for several turns and

may never be achieved, as explained previously.

While altruism based on the underlying assumption of kin selection is used here due to the desire to provide a feasible explanation for a phenomenon in an artificial life setting, the idea of kin selection provides a useful metaphor for any community of agents. In a collection of agents solving a variety of tasks with limited resources, the degree to which two agents should cooperate (and share resources, if needed) should intuitively be related to the proportion of tasks they have in common: if we view the relatedness of two agents as the proportion of tasks they share, the ideas behind kin selection and neo-Darwinism provide inspiration for how agents might best cooperate to achieve their collective goals.

3 The Multi-Agent System

In order to simulate language evolution, we have created a multi-agent system. The York Multi-Agent System (Kazakov and Kudenko, 2001) is a Java based application which allows for artificial life simulations to be conducted in two dimensional environments. It is particularly well suited to studying learning and evolution. Agents in this system have a behaviour based on drives, the most important of which are hunger and thirst. At each time-step, after the values of the drives are increased to reflect the cost of living, an agent will attempt to reduce whichever of its drives is greater. If either drive reaches its maximum then the agent dies and is removed from the simulation. Agents can also die of old age. This can occur starting when the agent is 300 cycles old.

If two agents with sufficiently low hunger and thirst values share a location, they may mate to produce a third agent. Mating costs an agent a one-off payment of one third of its food and water reserves, with the amount subtracted from both parents going to the child as its initial levels.

Agents attempt to reduce their hunger or thirst by finding and using food and water resources respectively. A resource can be utilised by an agent if they share a square in the environment, and this decreases the appropriate drive to a minimum and reduces the amount of resource remaining at this location. When a resource is entirely depleted by use, it is removed for several turns after which it reappears at its previous location with its resources renewed. When resources are required, agents will first utilise any in the square they occupy or will look for resources in adjacent squares. If this fails, they attempt to generate a path to a known resource as will be explained below and, finally if all else fails, they will make a random exploratory move. Additionally, in some experiments performed, agents are able to request resources or directions from other agents. Sharing resources uses a progressive taxation policy with 50% of any resources in excess of the minimum needed for reproduction being given to the requesting agent, while directions are shared in the man-

ner which will be described below. Agents always share when they are asked for help (assuming that they are able to) and they will do so as if they were clones of a single genotype: there is no account taken of the degree of kinship when deciding how much help to give.

4 Representation of Knowledge

The information used to navigate and communicate about the location of resources is stored in the form of rules which contain ordered sequences of landmarks, or paths, which are to be passed while travelling from the current location to the target resource. The landmarks used for this purpose are items within the environment whose function as landmarks and whose names are assumed to be known by all agents throughout the whole simulation. Informing another agent of a resource through the exchange of a path is an altruistic act for the reasons discussed in previous sections.

Agents can acquire knowledge of new paths in 2 ways. Agents may obtain paths through linguistic methods as noted above or they may find resources through exploration, in which case they will be able to construct a path linking this new resource to the previous location they visited by recalling the landmarks they have seen on their journey. In either case, the agent will store the new path internally in the same data structure in a form equivalent to

$$goto(resource) \rightarrow goto(locX), l_1, l_2, l_3, \dots, l_i$$

in the case of a rule acquired through communication and in a form similar to the following for rules acquired through exploration.

$$goto(resource) \rightarrow goto(locY)$$

$$goto(locY) \rightarrow goto(locX), l_1, l_2, l_3, \dots, l_i$$

where l_1, \dots, l_i are the landmark names received, in the case of linguistic acquisition, or the landmarks seen during the journey, in the case of exploration, and $locX$ is the location the agent gained the knowledge, in the case of linguistic acquisition, or the previous location visited, in the case of exploration. $locY$ is the location of the resource itself in the latter example. $locX$ and $locY$ are stored as a list of landmarks visible from the location in question (range of sight is limited to 2 squares in all directions). The first rule given above can be understood as stating that to go to the resource is equivalent to first going to $locX$ and then passing the given landmarks in order. The second rule set can similarly be understood as stating that to go to the resource it is sufficient to go to $locY$ which can be achieved by going to $locX$ and then passing the landmarks in the stated order. Note that the path that can be generated to take an agent from $locX$ to the resource is the same using the either first rule or the rule set given above, the only difference being that in the latter case the agent

has already visited the actual resource and thus can store a description of the environment at that point. A further rule of similar form is added to an agent's knowledge base when new paths are acquired through exploration that enable paths to be traversed in reverse direction. It should be noted that this description is very impoverished, containing no information on factors such as direction, absolute position or distance (aside from that which can be inferred from the number of landmarks mentioned in the rule).

These rules can be viewed in two ways, and are used as such by the agents. Firstly, they can be seen as procedural rules which capture the spatial relationship between locations and can be used for navigation by agents to resources and secondly, as grammar rules forming a regular language which can be used by the agent to share knowledge with others. When viewed as a grammar, the rules form a proto-language whose structure mirrors that of the landscape.

To access the data stored in the grammar to reach a resource, the agent will need to generate a sequence of landmarks from the grammar which will form a path that the agent can follow. This is done by using the grammar to generate a list consisting of only landmark names. The starting rule for this expansion is a grammar rule in which the left-hand side is $goto(resource)$, where $resource$ is the type of resource the agent wishes to locate. The rule used is the one that leads to the resource which the agent most recently visited, which implies that rules the agent has acquired through exploration are used in preference to those gained linguistically. Symbols of the form $locX$ are used as non-terminals in the parsing and landmark names as terminals. Expansion takes place using an A* search with a metric based on the length of the path (as measured by the number of landmarks).

Agents utilise the grammar when they are called upon to share their knowledge linguistically with others, but instead of following the path it is passed to the other agent, who will store it in the form of the first rule given above. There are a few minor but important changes to how the agents use the grammar when generating a path for another agent which ensure that hearsay and misinformation are not passed onto other agents. These changes involve only allowing rules which have been formed through exploration to be used when communicating information to another agent and then only the subset of these rules that refer to resources that have been recently visited. These changes, which maximise the chances that a resource to which an agent is directed is actually there, were found to significantly improve the performance of communicating agents in our earlier research (Kazakov and Bartlett, 2004).

The deliberately impoverished representation of geographical knowledge that is used here might be found to be of use in other agent applications where navigation is needed. A particular strength of the representation is that it allows for a certain amount of generalisation to occur. As locations are not precisely specified as coordinates, but

rather they are defined in terms of the landmarks visible from that position, positions close to each other may be viewed as the same location by the algorithm, which allows a degree of abstraction about an agent’s position. For example, in a domain involving robots navigating a building, it may not be important exactly where in a room a robot is for a navigational planner, rather which room the agent is in may be much more significant at this level. Landmarks too can provide for a form of generalisation. In the current work, landmarks have unique names, but there is no reason why this must be the case. If different landmarks had the same name, generalisation of any structural regularity of the environment would be possible, indeed inevitable, without an exhaustive search. Returning to the robot example, an identical series of visual or radio beacons leading to each door of a given room could be used, allowing a generalised concept of a path heading into that room to be formed. Other features, such as the minimal amount of information required to build up a series of routes, and the ability to associate scores rating the usefulness of particular rules, may also make this representation potential useful in other domains.

The representation of knowledge and its utilisation are discussed in more detail in our previous work (Bartlett and Kazakov, 2003).

5 Study of the Role of Environment

Intuitively, the benefits to a community of the ability to use language in this scenario might seem obvious. However it is not inevitable that the capacity to share resource locations should lead to an increase in a population’s ability to survive. It is conceivable that spreading knowledge of resources could lead to much faster depletion of those resources and hence starvation, or that a population could be created in which all individuals had sufficient resources to survive but insufficient to reproduce.

The benefits of a particular altruistic policy to a community of agents can be measured through conducting experimentation in which a population of agents is placed within an environment and the simulation is allowed to run for a period equivalent to several generations. The simulation is performed three times with different populations of agents, one of which is able to communicate, one of which is able to share resources and a final population which is entirely selfish. Agents in populations in which language use is not allowed still use the information gathered about the environment for personal navigation to make for fairer comparisons between the populations. Any survival or reproductive benefit associated with the use of a strategy should be manifested in the population size, in this way we study language evolution as a form of multi-agent learning which improves a population’s performance as measured by its ability to utilise resources.

In the present work, the environments in which agents

were placed had two of their structural properties changed in order to evaluate the effect that these had on populations implementing the various altruistic policies. A basic environment shape was defined, with four food resources equally spaced across the top of the environment and four water resources across the bottom as shown in figure 1. This map was then transformed along two dimensions to produce a series of maps. The distance between resources of the same kind was one factor that altered between experiments with the other factor being the distance between the different types of resource. Three values of each distance were chosen based on indications from our earlier work (Bartlett and Kazakov, 2004), 1, 3 and 5 squares for the distance between resources of the same kind and 3, 5 and 7 squares for the distance between resources of differing types. Each pair of distances were used in a map leading to a total of 9 environments. Figures 1 and 2 are two examples of the maps formed. Note that the number of resources in each environment is kept as a constant, as is their relative positioning, and that the different distances are thus achieved by altering the size of the environment. To ensure valid comparisons between experiments, the same number of agents were used in each environment.

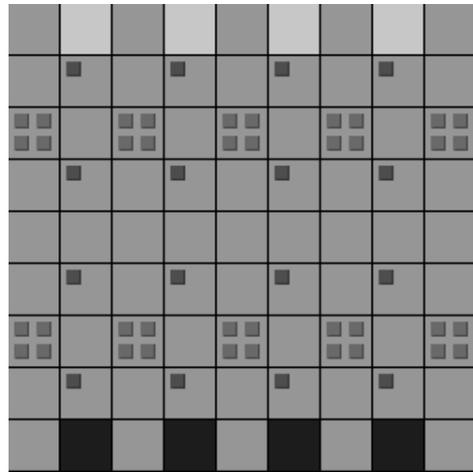


Figure 1: An environment with a space of one square between resources of the same type and seven squares between resources of different kinds. The coloured environment squares show where resources are located (food at the top, water at the bottom), and the smaller squares show the agents (in groups of four) and landmarks (uniformly spaced).

6 Results and Discussion

The results, as shown in figure 3, demonstrate that the population in which language is used performs significantly better than the other populations in most cases studied. As would intuitively be expected, the greater the

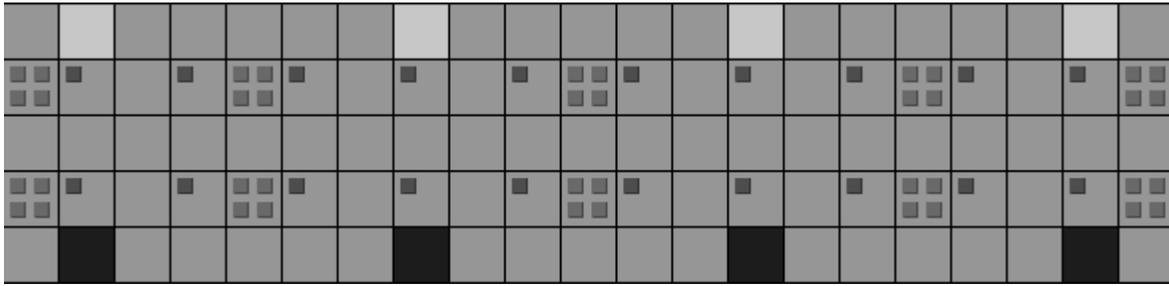


Figure 2: An environment with spacing of 5 squares between like resources and three between differing ones. The same rendering as figure 1 is used.

distance between food and water the fewer agents manage to survive regardless of strategy. However, though the resulting size of the language-using population decreases as this distance increases, the relative benefits of language increase, as can be seen by the ratio of the size of the language-using population to the selfish population (which can be viewed as a baseline for comparison).

The size of the final population also appears to be dependant on the spacing between resources of the same kind. Again the general rule appears to be that the greater the distance, the lower the population which manages to survive though the effect is not as pronounced as changing the distance between food and water. Language use also appears to increase in effectiveness with increasing distance for this variable.

The reason for the decrease in the size of population as the food-water distance increases is quite obvious, and is simply due to the greater difficulty in an agent managing to find both resources before dying of hunger or thirst. This reveals why language use is more beneficial in situations where the distance is greater. By using language, the complexity of the search for resources is simplified, as agents are able to gain information on resource location without having to go through the process of exploration. In effect, the use of language manages to parallelise the searching for resources. The effect of language is likely to be particularly pronounced in later generations when agents born into non-communicating populations have to locate resources through exploration as their parents did, but communicating agents will most likely be able to gain this knowledge linguistically almost immediately from their parents.

It is less immediately obvious that the distance between resources of the same type should influence the population dynamics though there are actually two ways in which the spacing between resources is likely to do so. Firstly, as stated earlier in this paper, resources are exhaustible with use. This means that occasionally agents will need to migrate from one source of food to another (similarly with water). At this point the distance between food sources becomes relevant. When the resources are separated by one square, it is almost inevitable that a move from an exhausted food source will lead to a po-

sition in which a new resource can be seen, whereas the task of locating a new resource becomes a matter of quite extensive exploration when the next closest resource is 5 squares away. Secondly, the difficulty of locating a resource initially in the environment increases in complexity as the distance between resources of the same kind increases. To see why this is so, consider two environments with the same distance between food and water but different distances between food and food (and therefore water and water). The number of squares from which food is visible in both environments is constant, being one square in each direction away from the food. However the number of squares in the environments are different and hence the environment in which food is more spaced out will need a greater proportion of the squares to be searched in order for resources to be found. Again, the reason that language performs better is due to the fact that it manages to reduce the amount of searching that must be done to locate resources.

There are two other effects that may also explain some of the benefits that language use brings. Firstly, populations using language have a tendency to cluster into a smaller part of the environment: the majority of agents tend to become reliant on a subset of the resources available. This occurs as the process of agents informing others of the resources of which they know results in most agents becoming aware of the same resources and routes between them. By clustering themselves into a reduced part of the environment, the probability of agents who are able to mate meeting is increased and hence the expected birthrate will also be increased. Secondly, the process of being able to learn routes through language produces an effect that results in shorter, more efficient routes becoming used. As communicating agents are able to gain information through both exploration and language, they are able to learn of several different routes to the same location. The details of the navigational planning algorithm ensure that agents will tend to take the shortest route which they are able to, and hence agents in communicating populations will tend to travel shorter routes than those agents in other populations, which reduces their chances of dying while travelling between resources and increases the time for which they are able to mate.

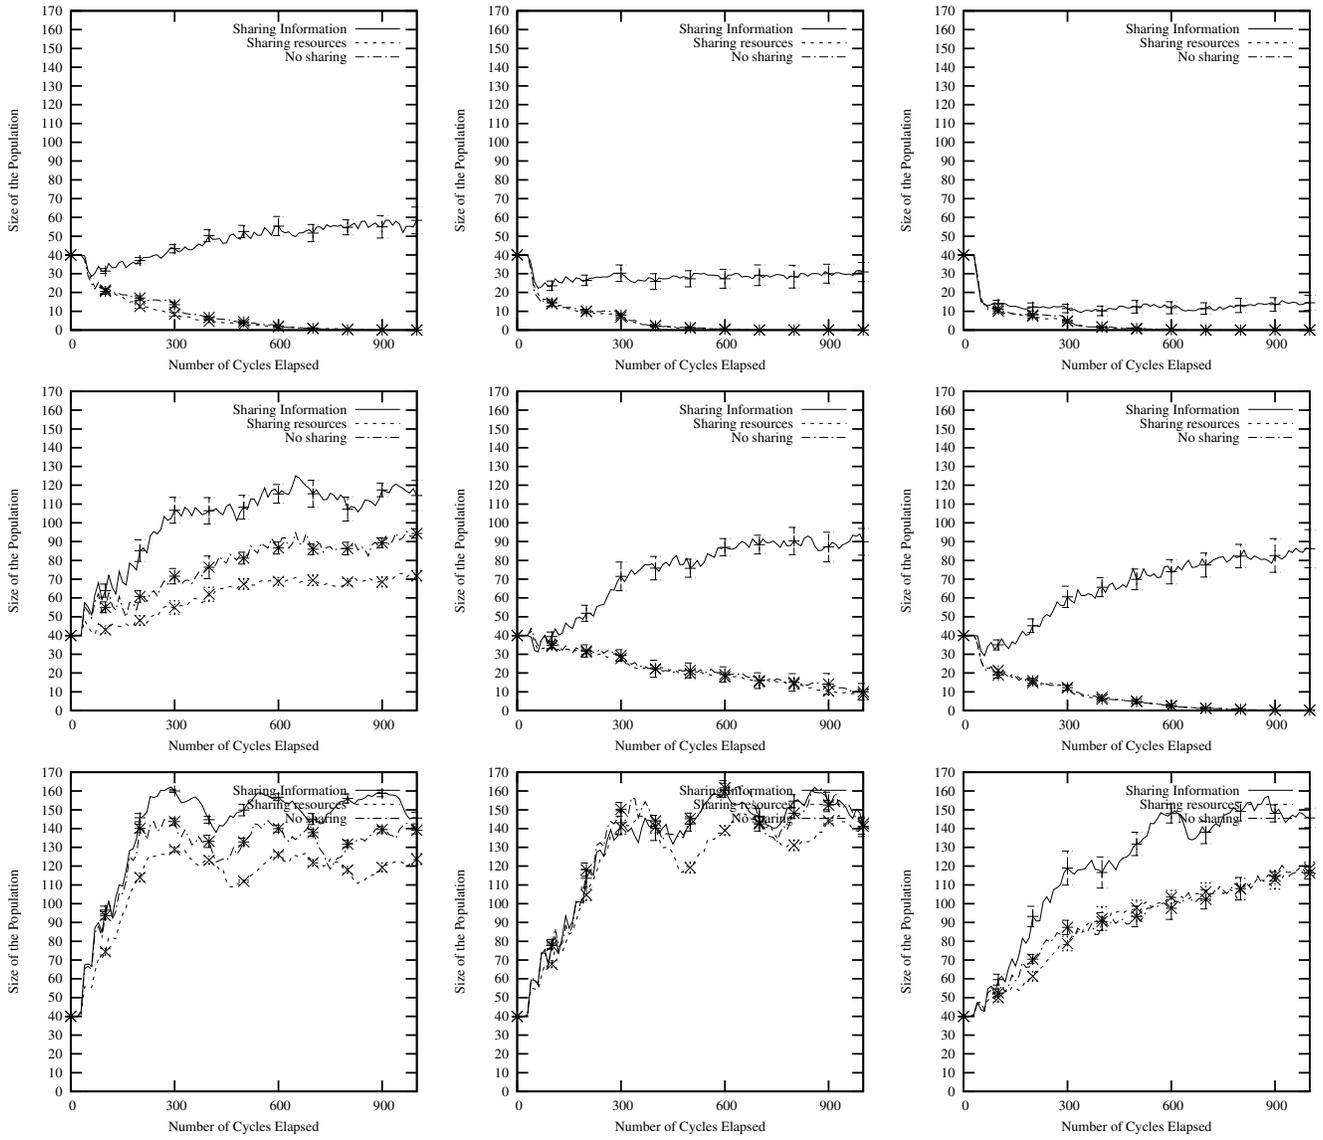


Figure 3: Comparison of sharing information, sharing resources and selfish behaviour. The distance between resources of different kinds increases from bottom to top, from 3 to 5 to 7 squares. The distance between resources of the same kind increases from left to right, from 1 to 3 to 5 squares. All results are averaged over 10 runs.

In all the experiments conducted, populations which shared resources performed no better than the baseline behaviour of not sharing at all, and in some cases did significantly worse. This serves to illustrate that the improved performance of the language-using community is not just down to the altruism of the communicating agents. The poor performance of the population who share resources can probably be best explained through the fact that sharing resources can lead to a population in which resources tend to become more evenly spread throughout the population. This creates a population in which fewer agents are able to reproduce, and those agents that are created through mating begin life with lower resource levels and thus a worse chance of surviving. One conclusion that might be drawn from the evidence produced is that agent populations would perform better if they concentrated on gathering more resources and were less concerned about how those resources were distributed between them.

A further point that should be mentioned with respect to the performance of agents who share resources is that they take noticeably longer to recover from the depletion of a resource than do agents who are not sharing resources. This can be seen in the graphs towards the bottom left-hand corner of figure 3. The periodic increase and decrease in population size seen in these graphs is characteristic of resources becoming depleted (Kazakov and Bartlett, 2004). It can be seen that in these experiments, the population size of the agents who share resources falls for longer than the other populations. This is caused by exploration to locate new resources being delayed by agents, who are content to draw energy from others rather than find a new permanent source of resources. When agents do eventually begin to explore, their collective resource levels are much reduced by this period of inactivity and it takes agents longer to regain a state conducive to procreation.

In passing, it can perhaps be noted that the fact that the periodic increases and decreases in population size are not seen in some graphs is due to the fact that in the experiments in which this is the case the population size is much smaller in relation to the environment size. This means resources are not being used to their full potential and thus the time at which they become depleted varies from resource to resource and from trial to trial, blurring the cyclic behaviour to the point where it cannot be seen in the graphs.

7 Conclusion and Future Work

This paper has outlined a multi-agent system built specifically to study the phenomenon of language evolution and has presented a range of experiments that have been carried out in order to study the influence of certain features of the environment upon the usefulness of language to a community of speakers. Specifically, we have carried out investigations in which two different types of the distance

between resources was varied. Our results show that in all the cases considered, language use was never detrimental to a population and became increasingly advantageous compared with other policies as distances between resources increased. Sharing resources, on the contrary never offered any benefit over selfish behaviour and actually proved to be a disadvantage when resources were most easily discovered.

It is our future intent to assess the effect that other features of the environment may have upon the usefulness of language to a population, for example varying the spacing of landmarks or introducing impassible terrain. Through experimentation we hope to assess the way in which features of the landscape such as these affect the size of the population and the segmentation of the linguistic community, and draw parallels between this and speciation and cultural division in the natural world.

When completed, this research should shed light on when exactly language evolution of this type should be expected to provide some benefit and hence its use be selected for by natural selection. Our intuition is that we will find the conditions in which language use is encouraged to be similar to those found at the time when primitive language use is believed to have begun in humans. Indeed the fact that in the present research, language use was found to be more beneficial as distance between resources increased, may be considered alongside the fact that language may well have first arose in early humans at the same time as they moved from a jungle habitat to the savannahs of east Africa. Such a change would have taken them from an environment in which food and water were quite abundant into one in which resources could frequently require many miles of travel to reach. Our research suggests that this is exactly the kind of environment in which language users would be likely to have a reproductive advantage.

References

- M. Bartlett and D. Kazakov. Social learning through evolution of language. In P. Liardet, P. Collet, C. Fonlupt, E. Lutton, and M. Schoenauer, editors, *Proceedings of 6th International Conference on Artificial Evolution, EA'03*, pages 340–351, Marseille, France, 2003.
- M. Bartlett and D. Kazakov. Replacing resource sharing with information exchange in co-operative agents. 2004. Submitted to the Third Conference on Autonomous Agents and Multi-Agent Systems, New York.
- E. J. Briscoe, editor. *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge University Press, 2002.
- R. Dawkins. *The Extended Phenotype*. Oxford University Press, 1982.

- W. D. Hamilton. The genetic evolution of social behaviour I. *Journal of Theoretical Biology*, (7):1–16, 1964.
- M. D. Hauser, N. Chomsky, and W. T. Fitch. The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298:1569–1579, November 2002.
- D. Kazakov and M. Bartlett. A multi-agent simulation of the evolution of language. In M. Grobelnik, M. Bohanec, D. Mladenic, and M. Gams, editors, *Proceedings of Information Society Conference IS'2002*, pages 39–41, Ljubljana, Slovenia, 2002. Morgan Kaufmann.
- D. Kazakov and M. Bartlett. Benefits of sharing navigational information in a dynamic alife environment. 2004. Submitted to the Ninth International Conference on the Simulation and Synthesis of Living Systems (ALIFE9), Boston, Massachusetts.
- D. Kazakov and D. Kudenko. *Machine Learning and Inductive Logic Programming for Multi-agent Systems*, pages 246–270. Springer, 2001.
- S. Kirby. chapter Learning, Bottlenecks and the Evolution of Recursive Syntax. Cambridge University Press, 2002.
- J. O'Keefe and L. Nadel. *The Hippocampus as a Cognitive Map*. Oxford University Press, 1978.
- M. Oliphant. *Formal Approaches to Innate and Learned Communication: Laying the Foundation for Language*. PhD thesis, Department of Cognitive Science, University of California, San Diego, 1997.
- L. Steels. The spontaneous self-organization of an adaptive language. In K. Furukawa, D. Michie, and S. Muggleton, editors, *Machine Intelligence 15*, pages 205–224, St. Catherine's College, Oxford, 1999. Oxford University Press. Machine Intelligence Workshop: July 1995.
- H. Turner and D. Kazakov. Stochastic simulation of inherited kinship-driven altruism. *Journal of Artificial Intelligence and Simulation of Behaviour*, 1(2):183–196, 2002.
- W. H. Zuidema and P. Hogeweg. Selective advantages of syntactic language - a model study. In *Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society*, pages 577–582, Hillsdale, USA, 2000. Lawrence Erlbaum Associates.