

# Self-Organisation of Communicating Agents – Linguistic Diversity in Populations of Autonomous Agents

Efstathios Avdis\*

Department of Cybernetics  
University of Reading,  
PO Box 225, Reading RG6 6AY, UK

Kerstin Dautenhahn\*\*

Department of Computer Science  
University of Hertfordshire,  
College Lane, Hatfield AL10 9AB, UK

## Abstract

This paper presents the study of a population of communicating agents living on hilly landscapes. These simple language-learning agents are utilised as the basis of investigation of complex adaptive systems processes. Such agents are able to develop their own language by interacting with each other while traversing their environments. The model can then be used to examine phenomena of self-organisation of language and adaptation of agent behaviour to the environment.

## 1 Introduction

The use of Artificial Life techniques to investigate language and communication in Linguistics is concentrated on the experimentation with Artificial Languages (A-Languages). A-Languages use computational approaches to the problem of explaining the origin and evolution of language. In the field of the origins of language Luc Steels and the Origins of Language Group investigate linguistic phenomena [1, 2]. Mike Oliphant, Simon Kirby and Jim Hurford have applied mathematical and computational modeling techniques to traditional issues in the evolution of communication and language [3-7]. Ezequiel di Paolo has been studying the evolution of communication, mainly through the processes of action and social coordination and spatial organisation [8]. Aude Billard and Kerstin Dautenhahn [9] have been researching grounding communication in autonomous robots that use basic vocabularies. Daniel Livingstone and Colin Fyfe have studied linguistic diversity [11] in models of distributed language learning. Epstein and Axtell [12] have been developing Sugarscape, a computational system that synthesises agent technology and Cellular Automata (CAs) to study human social phenomena.

This paper summarises a research project in self-organisation of communicating agents<sup>1</sup>. Within human languages, it is possible that groups of people speaking the same language, such as nations, still exhibit linguistic diversity in the form of dialects. Noticeably, the spatial segregation of dialects follows the patterns of geographical constraints of the environment. Using this as a source of inspiration, the investigation of analogous phenomena in agent populations could help gain valuable insight into the complexity of communicating agents. The agent-based model is developed using the SWARM simulation system [11].

---


\* E-mail: shu96ea@reading.ac.uk

\*\* E-mail: K.Dautenhahn@herts.ac.uk

<sup>1</sup> The work was carried out by the first author under supervision of the second author.

## 2 The Model


The general aim of this research is to investigate complex adaptive systems phenomena. In particular, the model concerns a population of 49 agents living on different environments. The agents are controlled by simple behaviour schemes and use an internal model of human dialects, in order to facilitate the study of divergence and convergence in communicating media. The computational model visualises the emerging effects of dialect formation, as the agents are coloured in accordance to their individual representation of the internal language model (figure 1). It was quickly noticed that the population achieves uniformity in their communicating medium, and the model was used to explore the reasons why.



**Figure 1.** A screenshot from an agent population on a hilly landscape

### 2.1 The Environments

The agents live on discrete  $2\frac{1}{2}D$  environments (see figure 2) and are free to communicate. These environments represent landscapes mapped onto a 2D grid of cells that has the geometry of a torus. Five different landscapes have been used: four with different hills and valleys, whereas the fifth landscape is a completely flat environment, used as a reference in the analysis of the results.



**Figure 2.** The hilly landscape of figure 1 represents the above 3D landscape

The agents are free to move around anywhere in the world, but their motion is somewhat restricted. If they choose to step to a destination cell with the same height as their current position, they can do so in one simulation cycle. If the height of the destination cell has a different height than their current position, then they can move to the chosen destination in a number of cycles dependent on the height difference:


$$\text{Cycles to move to destination cell} = \text{height difference} + 1 \quad (1)$$

Thus, the higher the chosen cell is, the longer it takes the agent to move there. This gives a “climbing” effect, which couples the motion of the agents with the environment. When two agents find themselves in neighbouring cells, they interact. Clearly, the number of the agents’ interactions is dictated by their motion, which is heavily affected by the environment.

## 2.2 Agent behaviour

The agents have three *happiness* variables that allow them to get some feedback on their actions in the environment. These happiness values are used to control the agents’ behaviour. The *social* happiness is the average number of interactions per cycle in the last ten cycles. The *mobility* happiness is the average number of steps per cycle in the last ten cycles. The *generic* happiness is the average of the two.

The behaviour of the agents is determined by three behaviour rules. A rule is a triplet of three situation – response pairs,  $\{(S, R_0), (S, R_1), (S, R_2)\}$ . When the agent detects a situation  $S$ , it is free to choose between responses  $R_0$ ,  $R_1$  and  $R_2$ . For situation  $S_i$ ,  $i \in \{0, 1, 2\}$ , there is a probability that one of the responses  $R_j$ ,  $j \in \{0, 1, 2\}$ , is chosen. This probability is expressed by weight  $w_{ij}$ .



**Figure 3.** The structure of the behaviour rules

The agents can find themselves to be either “too lonely” or “too immobile”, which can be detected by thresholding the values of the two respective happiness values. When in one of these two situations, the agent decides to move randomly, higher in the landscape, or lower. If the agent decides to move to a higher/lower cell but cannot find one around its position, it stays put. If the agent finds itself to be both “too lonely” and “too immobile” then we have a conflict. The agent chooses which situation to respond to either randomly, or by selecting to satisfy itself feeling either “too lonely” or “too immobile”.


Three types of behaviour have been investigated: *random*, *fixed* and *dynamic*. The first type allows the agents to move around completely randomly, irrespective of what situation they find themselves in. In the second type, the agent chooses its responses equiprobably and the behaviour weights do not change. In the final type, the behaviour is dynamic, namely the probabilities are allowed to change according to a simple Reinforcement Learning scheme. Here, the weight of the chosen response is rewarded if the response does not cause the generic happiness to decrease, and penalised if it does.

### 2.3 The Language Neural Network

The agents use an Artificial Neural Network as a working model of a dialect. The network comprises two layers, the internal state layer and the signal layer. The internal state is modelled as a sparse bipolar vector ( $\pm 1$  at each node, only one being set to  $+1$  for any one state). This is fed forward through the ANN to determine an agent's signal for that state, each output being thresholded to a binary value. Signals are represented at the language layer as arbitrary bipolar vectors:

$$y_j = \sum_{i=0}^{M-1} x_i w_{ij} \quad (2)$$

Then  $Y_j = 1$  if  $y_j \geq 0$ ,  $Y_j = 0$  if  $y_j < 0$ , where the vector  $\mathbf{y}$  is the signal generated for internal state vector  $\mathbf{x}$ .



**Figure 4.** The language agent neural network, mapping each of 3 states to one of possible 8 signals

During learning, the output signal from a teacher ANN is presented at the output layer of the learner ANN and fed back to produce an internal state (3). The error between the teacher's internal state,  $\mathbf{x}$ , and the learner's internal state,  $\mathbf{x}'$ , is used for learning by the learner agent, (4), with learning only performed when the learner misclassifies the signal.


$$x'_i = \sum_{j=0}^{N-1} y_j w_{ij} \quad (3)$$

$$\Delta w_{ij} = \eta(x_i - x_i')y_j \quad (4)$$

This network has 3 output neurons and 3 input neurons, thus we have 3 states and  $2^3 = 8$  possible signals. Using these values, the fact that there is a greater range of signals available to the language agents than to be communicated provides redundancy in the signalling ability of the agents. Each agent in the population retains a *dialect*, i.e. a mapping of how the three states can be transmitted. In order to have a suitable visual output, the dialects have been coded in groups of 8. Each such group is called a *language* and it is represented with a single colour.

## 2.4 Learning

Consider the case when two agents meet in two neighbouring cells. These agents interact by teaching each other their dialect. Two types of learning can be used: We can have the learning agent updating its language NN before it takes its turn in teaching. This is called *asynchronous* learning (figure 5a). It can also be the case that the learning agent stores the weight updates in a latch, then teaches the other agent in a similar fashion. Each agent updates its NN after the whole interaction has taken place. This is called *synchronous* learning (figure 5b).



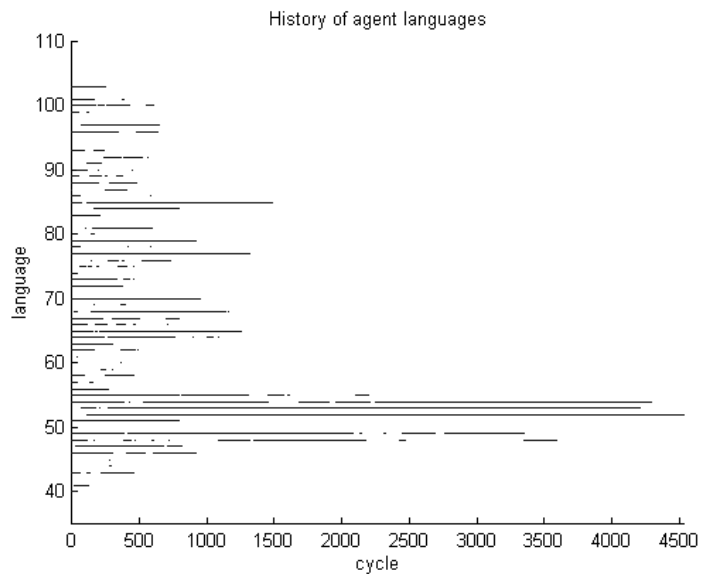
**Figure 5.** Asynchronous (a) and synchronous learning (b)

## 2.5 Analysis Tools

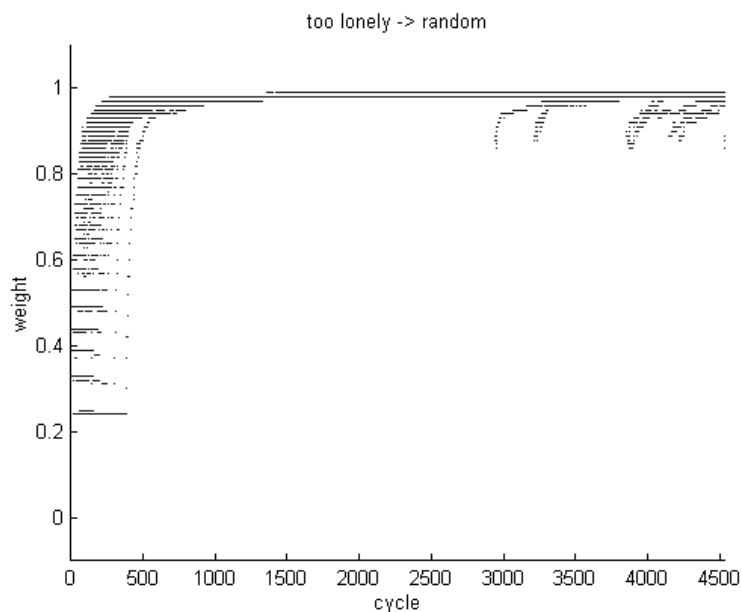
The model is observed by various graphs: a graph of the average interactions over the population of agents, a graph of the average generic happiness and a histogram of current languages. A history of the existing languages is also created by probing the agents and collecting the languages that exist in the population, at each cycle. For each cycle (x axis), a point is plotted at the existing languages' IDs (see figure 6). This together with the languages histogram gives a complete view of the self-organisation of the population's language.

In addition to the above, when the behaviour of the agents is dynamic, we track the changes in the three behaviour probabilities for all three situations. The graphs are produced in the same way as the language histories are, and they can be used to study the adaptation of the whole population to the environment (see figure 7).

Each simulation has been run long enough for the population to converge to a single language, with an additional time margin after convergence to verify that the population's winning language does not change.



**Figure 6.** History of existing languages



**Figure 7.** History of behaviour probability  $w_{00}$

### 3 Testing

This study uses two different settings for the initial conditions of the agents' languages, two different types of learning, three different types of behaviour and 5 environments. The test vectors have the following structure:

<initial conditions>/<learning>/<behaviour>/<landscape>

where each bracket is a place holder for the string codes shown in table 4. For example, a simulation using random language initialisation, asynchronous learning, random behaviour and a flat environment is represented by test vector `as/r/fl`.

<b>Initial conditions of languages</b>	<b>String code</b>
Random	no string
“cake”: 4 initial languages, with the same number of agents per language	c
<b>Learning</b>	
Asynchronous	as
Synchronous	syn
<b>Behaviour</b>	
Random	r
fixed probabilities	f
Dynamic	d
<b>Environment (landscape file)</b>	
Flat	fl
complex landscape 1	1
complex landscape 2	2
complex landscape 3	3
complex landscape 4	4

**Table 1.** Test vector notation

## 4 Results

The testing space is four-dimensional: (initial conditions, learning, behaviour, environment). The following table summarises selected runs of the tests:

<b>Test vector</b>	<b>Mean happiness</b>	<b>Cycles to reach mean happiness</b>	<b>Winning language</b>	<b>Last competing language</b>	<b>Cycles to converge</b>
<i>Varying environments</i>					
as/r/fl	0.53826	9	54	46	3400
as/r/1	0.50655	20	59	43	5650
as/r/2	0.38148	10	67	51 (2500 cycles)	2400, 2800
as/r/3	0.34025	10	68	64	3000
as/r/4	0.37371	11	51	50	4200
<i>Varying type of behaviour, flat environment</i>					
as/r/fl	0.53826	9	54	46	3400
as/f/fl	0.20751	21	93	89	5800
as/d/fl	0.51272	400	52	54	4300
<i>Varying type of behaviour, hilly environment</i>					
as/r/3	0.34025	10	68	64	3000
as/f/3	0.29044	20	71	55	3800
as/d/3	0.32383	200	78	86	4100
<i>Varying type of learning, flat environment</i>					
as/r/fl	0.53826	9	54	46	3400
syn/r/fl	0.53838	30	63	47	3900
<i>Varying type of initial language conditions, hilly environment</i>					
c/syn/d/3	0.32000	170	64	68	2200
syn/d/3	0.32238	100	83	81	4600

**Table 2.** Summary of selected test runs

## 5 Discussion

The most conspicuous phenomenon is the persistent convergence of the population's languages to a single one. In fact, the population converges at every single simulation run carried out. This is independent of the initial distribution of agents between languages, as this was disproved by carrying out runs with 4 initial languages, each being used by 12 agents. As a general observation, the fewer the number of initial languages, the faster the population converges. The model was also tested with one initial language. Although a different language was generated, it became extinct after a short period, and the winning language was the initial one.

Bearing in mind that a single language comprises a group of 8 different dialects, it is easy to see that an interaction between agents of different dialects can cause either one to change to a different language. Conversely, the uniformity of language can mean uniformity of dialect. As a matter of fact, this is what actually happens. Although it is hidden under the visual output, after the agents converge to the same language, they converge to the same dialect. It is at that point in the simulation when new languages of short-life span can be generated. This only happens between a very small number of agents and thus *most* of the time it is "absorbed" by the dominating language.

Let us discuss this point even further. Consider the case where the whole population but very few agents (in fact, it could well be just one agent), has converged. Chances are that the dominating language will prevail. However, here is where the landscape geography comes into play. It is possible that an agent of the minority language finds itself next to an agent of the majority language, and both of them are cut off from the rest of the population. If the interaction does cause yet another language to spring out, but instead the "rebel" agent converts the "conformist" agent. If this is suitably repeated, the positive feedback loop can lead to the minority language dominating, as was indeed the case with a number of simulation runs.

The agent happiness is a measure of how much the agents travel and interact with each other and it is heavily dependent on the landscape. For all landscapes, the mean happiness is lower when dynamic behaviour is used and even lower when the behaviour probabilities are fixed.

Other observations include landscape-specific phenomena. In the flat landscape (dynamic behaviour), the agents quickly adapt their behaviour so that situation "too lonely" mostly leads to random motion (see figure 7). Namely, there exists a positive feedback loop so that at the moment an agent chooses to modify its behaviour towards randomness, this behaviour is reinforced. The next time it is more probable that the response will be to choose a random destination cell, and so on and so forth. The limiting factor is that random motion cannot always guarantee more interactions, which explains the drift at the top right part of the aforementioned graph.

The concentration of the agents in valleys is disproportionately higher than on plains. The geography of the environments stalls the agents in the valleys, creating a "bottleneck" effect. When the behaviour is dynamic, the agents on plains converge much faster because they change their behaviour to random, which increases the number of interactions. These agents are also generally happier.

In the landscape of figure 2, a temporary spatial segregation of two groups of languages between the two deepest valleys was observed. However complicated its geography is, it is the landscape in which the second fastest convergence is observed. The complexity of this landscape keeps the happiness at low levels, which provides a drive for motion and perpetual adaptation of behaviour.

The use of the different types of learning affects how long the population takes to converge, with asynchronous learning shortening the time needed for convergence. Because it alters the way agents update their language network weights, it is quite reasonable that the winning dialect should be different in each case. Finally, it was noticed that between the last two competing languages, there is always only one state that is mapped differently.

## 6 Conclusions

This research uses a population of agents as a means of study of complex adaptive systems. In particular, the study is focused on the identification of factors that are responsible for convergence of the communicating medium in a population of agents, to a common instantiation.

The model has revealed that there is a high degree of coupling between the environment, agent behaviour and convergence to a particular language rather than another. The agent behaviour determines the agent's trajectory. The environment constrains the motion of the agent, which affects its happiness. When the model uses the dynamic behaviour scheme, these effects come in full swing, because the level of happiness dictates how the behaviour changes. The new behaviour controls the agent trajectories in the next cycles and so on and so forth.

The fact that the population languages converge to a single one is comparable to recent research by Arita & Koyama [13]. They discuss linguistic diversity and convergence of a group of communicating agents that use a vocabulary table as a language model. Their agents use genetic operations, whereas in this study the mechanism of convergence of the ANN language model is Hebbian-like learning. There is also an analogy between the aforementioned results and the convergence of linguistic mappings (*interpretants*) as discussed in Nehaniv [14]. In future work, the present study can be extended to incorporate the agents using the converged language model to communicate information about their environment, in hope that the results can be of more practical importance in the field of mobile robotics.

As a final remark, it was found that the self-organisation of the population's language is a result of the highly complex interaction between the model components. Complex landscape geography is no obstacle to language self-organisation, because agents adapt to it, but the steeper the landscape, the slower the convergence. If we were to select the most successful non-adaptable behaviour scheme, it would have to be one based on random motion. However, it would be wrong to analyse the system by separating it into independent dynamic processes. The *synthetic* approach of Artificial Life is the most valid method of investigation of such complex systems.

## 7 References

1. Steels, L., (1996), *Synthesising the origins of language and meaning using co-evolution, self-organisation and level formation*, in Hurford, J., Knight, C., & Studdert-Kennedy, M., (ed.), (1997), Evolution of Human Language, Edinburgh Univ. Press, Edinburgh.

2. Steels, L., (1997), *Language Learning and Language Contact*, in Daelemans, W. (ed.), (1997), Proceedings of the workshop on Empirical Approaches to Language Acquisition, Prague.
3. Kirby, S., *Learning, Bottlenecks and Infinity: a working model of the evolution of syntactic communication*, Proceedings of the AISB '99 Symposium on Imitation in Animals and Artifacts, pp. 55-63, University of Edinburgh.
4. Oliphant, M., *Rethinking the language bottleneck: Why don't animals learn to communicate?*, in Knight et al. (eds.), Proceedings of the Second International Conference on the Evolution of Language, London, April 1998.
5. Kirby, S., (1998), *Fitness and the selective adaptation of language*, in Hurford, J., Knight, C., & Studdert-Kennedy, M., (eds.), Approaches to the Evolution of Language, pp. 359-383, Cambridge University Press.
6. Hurford, J. & Kirby, S., *The evolution of incremental learning: language, development and critical periods*, Cognition, Edinburgh Occasional Paper in Linguistics EOPL-97-2.
7. Hurford, J., *Functional Innateness: explaining the critical period for language acquisition*, submitted to Noonan, M., Moravcsik, E. & Newmeyer, F., (eds.), Functionalism and Formalism in Linguistics, Volume 2 (SLCS 42), John Benjamins, Amsterdam.
8. Di Paolo, E., (1997), *Social coordination and spatial organisation: Steps towards the evolution of communication*, in Husbands, P. & Harvey, I., (eds.), Proceedings of the 4th European Conference on Artificial Life, ECAL97, MIT Press/Bradford Books 1997.
9. Billard, A. & Dautenhahn, K., (1997), *Grounding communication in autonomous robots: An experimental study*, Robotics and Autonomous Systems 24 (1998), pp. 71-79, Elsevier
10. Livingstone, D. & Fyfe, C., (1999), *Diversity in Learned Communication*, Proceedings of the AISB '99 Symposium on Imitation in Animals and Artifacts, pp. 139-146, University of Edinburgh.
11. SWARM homepage: <http://www.santafe.edu/projects/swarm/>, (June 1999).
12. Epstein, J.M. & Axtell, R., (1996), *Introduction*, Growing Artificial Societies: Social Science from the Bottom Up, pp. 4, The Brookings Institution.
13. Arita, T. & Koyama, Y., (1998), *Evolution of Linguistic Diversity in a Simple Communication System*, Artificial Life 4(1), pp. 109-124, UCLA, MIT Press
14. Nehaniv, C.L., (2000), *The Making of Meaning in Societies: Semiotic & Information-Theoretic Background to the Evolution of Communication*, Proceedings of the AISB 2000 Symposium: Starting from Society – the Application of Social Analogies to Computational Systems, Editors B. Edmonds and K. Dautenhahn, pp. 55-63