Formation of a Common Spatial Lexicon and its Change in a Community of Moving Agents

Peter BODÍK, Martin TAKÁČ
Institute of Informatics, Faculty of Mathematics, Physics and Informatics
Comenius University
email: *bodikp@cs.berkeley.edu, Martin.Takac@fmph.uniba.sk*

Abstract. This paper investigates factors influencing the establishment of a common spatial lexicon in a community of agents moving in a simulated environment. The model avoids some traditionally criticized features of other models of the emergence of a common lexicon such as the use of only cued representations, pre-defined fixed meanings shared by all agents, explicit meaning transmission and nonverbal feedback about the outcome of a game. While each agent forms its own concepts for distances and directions, coherent lexicon emerges enabling agents to localize objects in the environment based on their spatial description. Factors necessary for language change are then investigated in an experiment where agents join/leave the community and the results are compared to those of the related model of Steels [14].

1. Introduction

Human language is a complex communication system that has evolved during thousands of years among humans. *How* exactly did language evolve is still a big puzzle and often a topic of heated debates. Theoretical approaches of linguists, psychologists and biologists to the problem have been complemented with *mathematical models of language*, e.g. [9], where typically the language dynamics is described by a set of equations and the properties of the language are studied using mathematical methods. However, for the mathematical proofs to be feasible it is often necessary to make crucial simplifications decreasing linguistic relevance of the models.

Yet another recent approach to modeling of language evolution is offered by computer simulations and artificial life [6]. Artificial life has been used so far to simulate complex dynamic systems, where verbal theorizing often leads to incorrect predictions because our intuitions about the links between local interactions and global behavior are notoriously unreliable. This makes human language an ideal topic for exploration using A-life models.

Except insights to the origins of human language, computational models can yield also interesting technical applications. Coordination, negotiation and language emergence among various types of artificial agents as well as man-machine interaction are important topics of today.

This paper describes a computational model of emergence of spatial concepts and lexicon in a community of moving agents. In the next section we summarize relevant existing models, pinpoint the differences and discuss the methodological issues leading us to assumptions of our model. Section 3 describes architecture of the model in detail, section 4 brings results on factors influencing the formation of the lexicon, section 5 deals with a flux of agents. Section 6 focuses on factors causing language change and compares them to the relevant model of Steels and the last section concludes.

2. Existing Models and Methodological Issues

Computational models of language origins can be generally categorized to those emphasizing the role of innate factors and their genetic evolution, e.g. [2,16], and those based on learned factors and cultural transmission, e.g. [8,12] (although hybrid models exist as well [7,10]). Pioneering work in investigating cultural mechanisms of language emergence has been done by Steels et al. [12]. Their models typically consist of agents, each with its own knowledge of language (lexicon), engaging in local interactions (e.g. pointing to an object and emitting a word). Agents adapt their lexicons according to the outcome of the interactions, which creates positive loop between success and use and leads to a global coherence.

This approach based on mapping of meanings to signals is suitable for modeling the emergence of a communication system resembling that of vervet monkeys [3]. However this is yet far from human language. According to Gärdenfors [4], signaling systems of animals are based on *cued representations* – those standing for something actually present or triggered by the current situation. He states that primary function of human language is to speak about things “not here and now” which requires *detached representations* independent of the outside context. Thus the necessary step to more complex models of language would be to endow agents with internal needs, motives, drives or goals.

This was the primary motivation for our model of moving agents, however it has soon shown too a complex step. Thus we decided to start with much simpler model of agents moving in a spatial environment and creating cognitive maps [15] of objects they discovered. Agents are born with general notion of *distance* and *direction* which they refine during their lifetime thus each creating its own spatial concepts. Referents of meanings are thus not in external world but in each individual agent. This implies that no two agents must have the same set of meanings (although they turn out to be quite similar due to using the same perceptual apparatus for experiencing in a shared environment). This complies with sociognitive semantics of Gärdenfors [5]. Agents are able to talk about objects in their vicinity and later, when the basic lexicon has established, about any object they remember in their cognitive maps.

3. Architecture of Our Model

All experiments presented in this paper were conducted as simulations of a multi-agent system. The whole system consists of agents, objects, and a square playground, where the agents and objects are situated. Each experiment proceeds in turns; in every turn an agent can move and/or communicate with another agent (play a *language game*).

3.1 Lexicon

The lexicon stores all concepts of the agent (see below) along with corresponding words in the form of “word/meaning” pairs. *Meaning* is the internal representation (the concept) and *word* is a string of characters that is used to communicate the meaning to other agents. The agent uses its lexicon a) to find a meaning of a word (word to concept), and b) to find a word to express a concept (concept to word).

Each word/meaning pair has its score – a real number expressing its success in communication. The score is modified during language games; increased after a successful communication and decreased otherwise. To express a meaning, the agents use the word with the highest score, the preferred word.

3.2 Concepts

Agents have three types of concepts: objects, distances and directions. All concepts are stored in the agent's lexicon along with the associated words. The concepts of distance and direction are innate in the agents. However, in the beginning they only have a very general notion of the *spatial* concepts. For example, the most general direction concept is "0 to 360 degrees". During an experiment the agents play *spatial games* (see below) where they describe objects using distances and directions (e.g. 20 meters to the north). If an agent can't disambiguate among several objects, it needs more specific spatial concepts. The agent can thus *divide* an existing spatial concept into two *more specific* concepts. The spatial concepts thus form a *discrimination tree* [11] (see Figure 1).

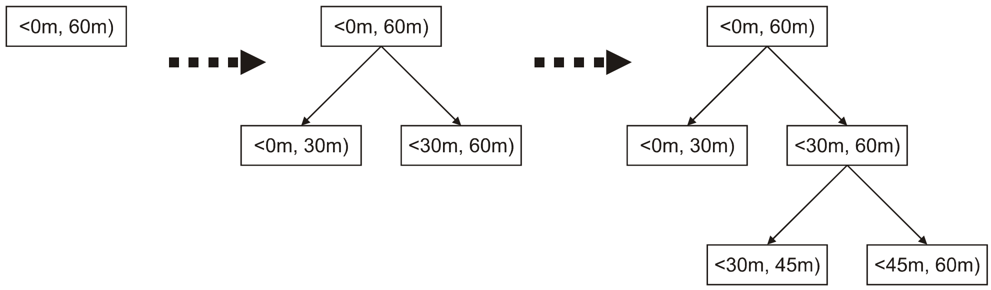

Figure 1. Discrimination Tree

To describe a position of an object, an agent can use just the distance *or* the direction (a one-word description) or both (a two-word description). Also, the agent can use the most specific concepts (e.g., "20 to 22 meters") or prefer more general ones (e.g., "15 to 30 meters"). In [11] agents always used leaves of a discrimination tree for a description. In our work we decided to use the *Succinct-and-General* rule which we believe is more natural. It means that to identify object uniquely the agent prefers *one-word* description of the position over *two-word* descriptions and general concepts over specific ones.

3.3 Language Games

During the experiments, the agents can communicate using three types of language games: pointing games, spatial games and evaluation games.

Pointing game involves two agents (A and B) and an object O in their vicinity. Agent A searches the lexicon for a word w describing O (a preferred word). If A has no word for O, the game fails and A creates a new word for O. Otherwise, A *points* to O and utters w. If agent B has word w associated to O in its lexicon, the game succeeds. Otherwise, the game fails and B learns the new word w. Finally, agent B updates the score of (w,O) in its lexicon. This game is used to develop a lexicon for objects.

Spatial game also involves two agents (A and B) and an object O, but the agents describe O by its relative position (e.g. 20 meters to the north), instead of pointing. Agent A first selects one or two spatial concepts (direction and/or distance) that unambiguously describe object O (if A has no such concepts, more specific ones are created). A then finds preferred words for the selected concepts (again, if there are no words for a concept, the game fails and the word for the concept is created). Finally, agent A utters preferred word for O and one or two words describing O's position (words w_1 and w_2). Agent B then finds object O on the playground (using its name) and describes its position from A's point of view (using spatial concepts c_1 and c_2). If (w_1,c_1) and (w_2,c_2) are in B's lexicon, the game is successful. Otherwise it fails and B learns the new word/meaning pairs. In the end, B updates the score of (w_1,c_1) and (w_2,c_2) in its lexicon. This game is used to develop a lexicon for spatial concepts.

Evaluation game is used only to assess the quality of the emerged language and does not change the lexicons of agents. It involves two agents – A and B – and an object O. A selects the object, describes its relative position and utters the words for the spatial concepts. B hears the words (position of the object), decodes them and tries to locate an object with such position. If B finds object with the corresponding position and if the object is the one A referred to, the game is a success. Otherwise it's a failure. The average success rate of the last 200 evaluation games is called localization success.

3.4 The Experiments

The experiments were conducted with 10 agents and 9 objects on the playground. 20% of the games played were evaluation games, the remaining 80% were spatial games. A pointing game was played only if a spatial game failed due to object word misunderstanding (agent B didn't understand the word for object O). We used multi-generation experiments; every *n* turns a random agent was removed from the experiment and a new one was introduced. The speed of flux of agents, *n*, is called the *flux rate* (FR).

4. Factors Influencing the Emergence of Language

During the experiments, a shared lexicon for objects emerged followed by a shared spatial lexicon. The pointing games were necessary only in the beginning and after the object lexicon emerged, only the spatial games were played (see Figure 2).

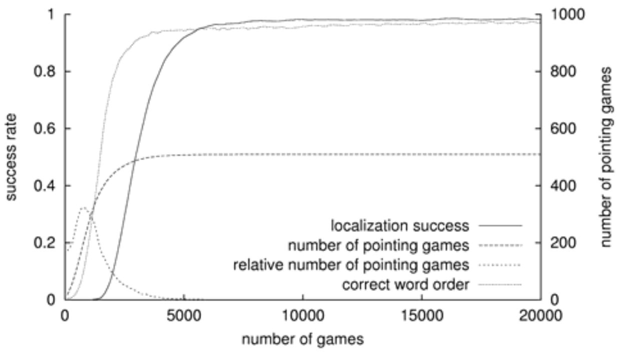

Figure 2. Use of pointing games

On Figure 3 is a part of a spatial lexicon of a single agent. We can see that the discrimination tree for distances is asymmetrical; for longer distances more specific concepts were needed to precisely describe the position of an object. The tree for directions was always symmetrical, because agents used all the directions equally. The most general concept of distance was not used during the experiment, however, from the other concepts, more general ones were used much more frequently than more specific ones (we can see that from the success/use statistics of the used words). Also, the most specific concepts were created much later; word fht59 was created approx. in game 5900 and word gss33 in game 3300.

We examined the following factors influencing the emergence of language: conceptualization phase, random vs. fixed word order, feedback during language games. In the following we describe the factors and the results of our experiments. More detailed description of the model, experiments and results can be found in [1].

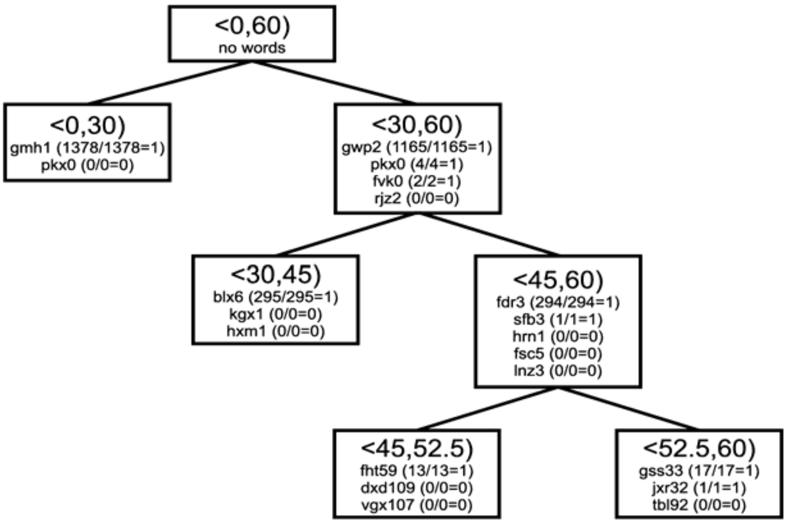

Figure 3. An Example of a Discrimination Tree for Distances

Conceptualization Phase

An agent is created with only one general concept of distance and one of direction. New spatial concepts are created later, during its contact with the environment. The conceptualization is done by looking at an object and trying to describe its position using spatial concepts. The agent can create its spatial concepts before he starts playing language games experiment (*early conceptualization*) or on the way during the games (*on-line conceptualization*). We found out that early conceptualization speeds up the growth of localization success, but the community using on-line conceptualization have also reached high localization success, albeit slower.

Random vs. Fixed Word Order

During spatial games, the agents utter two- or three- word sentences. The word order in these sentences could be fixed (the agents know which of the words is word for an object, distance and direction) or random (the agents use a simple algorithm to determine the word order). We expected that the random word order would hinder the emergence of spatial language. However, after some number of games, the agents were able to successfully determine the correct word order and came up with a shared language.

Feedback During Language Games

After a pointing or a spatial game, the hearer knows the result of the game automatically (by looking to its lexicon). He then can give the speaker (nonverbal) feedback about the

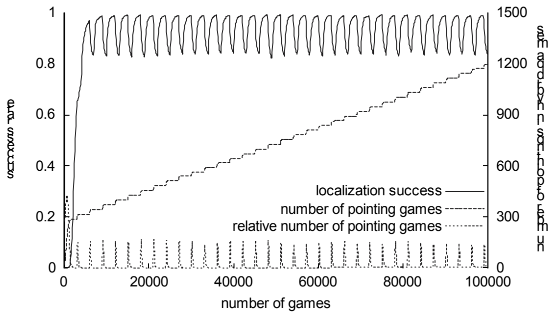

Figure 4. Experiment 1c, Flux Rate 3000

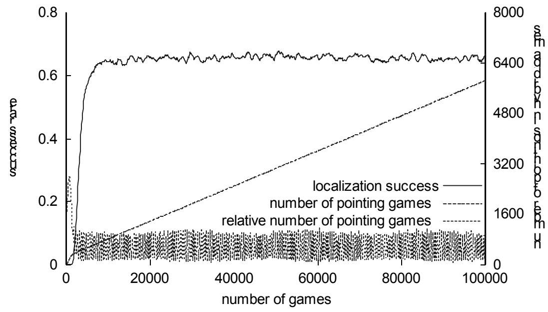

Figure 5. Experiment 1f, Flux Rate 500

outcome of the game, if it was successful or not, and the speaker can update its lexicon accordingly. However, our experiments have shown that the feedback was not necessary and a shared language emerged without feedback during language games as well.

Scalability

A typical experiment consisted of 10 agents and 9 objects. We were interested whether the results remain valid if we make the environment more complex (e.g., more objects) or included more agents. The only effect of increasing number of agents or objects was that a shared language emerged after a longer time.

5. Flux of Agents

To understand the dynamics of an experiment with an inflow/outflow of agents, we conducted a set of experiments with a varying flux rate. The flux rates were as follows: experiment 1a) 10000, 1b) 5000, 1c) 3000, 1d) 2000, 1e) 1000, 1f) 500, 1g) 200, 1h) 100, 1i) 50.

The results of experiments c and f are presented in Figures 4 and 5. We can see from the graphs, that after an old agent leaves the experiment (and a new one enters), there is a drop in the localization success. This drop is caused by the fact that the new agent doesn't *share* the language of the other agents. First, it has to learn the object lexicon (notice the increase of the number of pointing games) and then the spatial lexicon.

We can divide the experiments into three groups: the *first group* contains experiments a to d (flux rate at least 2000). In these experiments, a shared language evolved and the agents were able to sustain it. The new agent had enough time (played enough games) to learn the established language from the other agents. The *second group* (e, f, g) contains experiments with flux rates between 200 and 1000. In these experiments, the localization success reached about 80%, 60% and 20% respectively and remained on this level during the whole run. The agents didn't have enough time to learn and the language couldn't develop fully. The *last group* (h and i) contains experiments with flux rate less than 200. The new agents were entering the environment very quickly and didn't have enough time to establish a shared language. The localization success was 0% during the whole experiment.

In the following two experiments, we look at how the language evolves when we change the flux rate during the experiment. Both experiments start with flux rate of 2000. After 20000 games the flux rate is changed to 500 (experiment 2a) or 50 (experiment 2b) and changed back to 2000 after another 50000 games (Figures 6 and 7).

In the first phase of the experiment (games 0 to 20000), the flux was slow enough for the language to evolve. In the second phase (games 20001 to 70000), the flux became

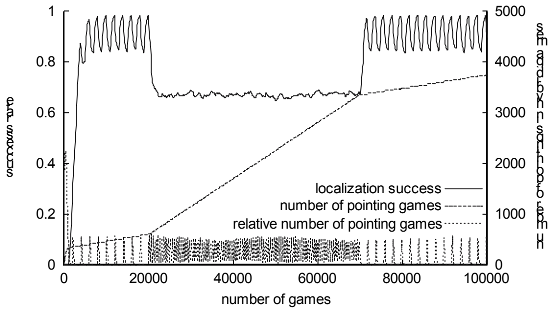

Figure 6. Experiment 2a, Varying Flux Rate; FR 2000, 500, 2000

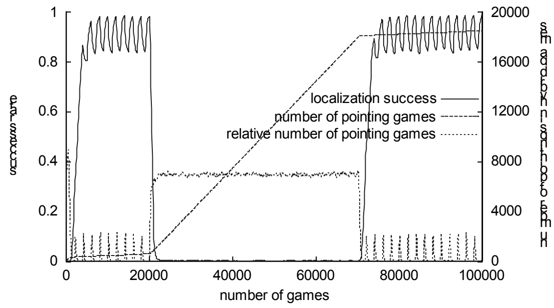

Figure 7. Experiment 2b, Varying Flux Rate; FR 2000, 50, 2000

faster (flux rate of 500 and 50). Even though the language was fully evolved, the agents could not sustain it during the faster flux. During flux rate of 500, the quality of language deteriorated to about 60%. During flux rate of 50, there was no shared lexicon at all (localization success fell to 0%). During the third phase of the experiment (games 70001 to 100000), the flux rate was slowed down again to 2000. Despite the low quality of language in the previous phase, the slower flux rate allowed the agents to build the language anew. Localization success have risen again to about 99% in both experiments.

These experiments have shown that the evolved language was able to sustain a reasonable flux. Higher rate of population change lead to deterioration of the language.

6. Spontaneous Language Change

One of the interesting characteristics of human languages is its spontaneous change. Human languages are constantly evolving; some words that were used 100 years ago are no longer used in the present and new words are created every day. An A-life model of spontaneous language change was proposed by Steels in [14]. In his experiments both stochasticity in communication of agents and flux of agents (old agents leaving and new agents entering the experiment) are necessary for a change in language. In our model we have not incorporated stochasticity, yet we have observed a language change.

Let us now look at two experiments and examine how the language changes during each one of them. Experiment 3a starts with flux rate of 2000 (phase 1), after 20000 games the flux rate is changed to 50 (phase 2) and changed back to 2000 after another 160000 games (phase 3). Experiment 3b runs with flux rate 2000 for 5 million games.

6.1 Constant Change at Flux Rate 50

We have learned from the experiments that the faster is the flux, the higher is the probability a word for concept will change. The most extreme case is experiment 3a with flux rate 50. A shared language emerged during the first phase (FR 2000), but the established words were forgotten soon after the second phase began (with FR 50). New agents were flowing in so quickly (one every 50 games) that new words were being constantly invented but soon forgotten.

We took a snapshot of the preferred words for objects of all 10 agents after 60000 games. The agents were in the middle of the second phase with FR 50 and many different words were used for every single object. Also, a lot of concepts (35 out of 90) didn't have *any* word assigned to them and thus the agents had many possibilities to invent new ones. With so many different words used for every object, the new ones could easily outcompete them. The original words were soon forgotten as the agents left the experiment. The competition of words for object 1 is depicted on Figure 8 (the graph represents relative score of each word used between games 42700 and 54000 among all agents). During about 10000 games, 35 words were invented for object 1 and only two of them survived (both invented in game 53500). Both of these words remained in the population of agents for quite short time and were soon replaced by new ones.

Another snapshot of the preferred words was taken after 200000 games. The flux rate was set to 2000 after 180000 games, the lexicon settled down and only one or two words were used for every object. Most of the words that survived were created between games 178000 and 181000 – close to the end of phase 2. None of the words created before game 178000 survived.

The flux rate in this experiment was extremely fast and the shared language that had evolved after 20000 games deteriorated very quickly in the second phase. After the flux rate slowed down to 2000, a new shared language emerged and localization success stayed close to 100%. This experiment is, however, not a very realistic simulation of language change in real languages, because the quality of language is very low and changes occur very often.

6.2 Gradual Change at Flux Rate 2000

In experiments with flux rate 2000, the words changed very rarely. In an experiment with 10 objects and flux rate 2000 only 13 changes of words for objects took place during 5 million games. Competition of words during one of the changes can be seen on Figure 9. Let's look at this particular change of word in detail.

Up to game 720800, there was only one word used for object 8 – word *svz0*. However, a few hundred games ago a new agent entered the experiment and invented a new word in game 720800 – word *nlf720*. The next agent that came (2000 games later) learned the new word and there were thus two agents using this new word. After 2000 games came another agent, also learned the new word, but switched back to the old one as the majority of agents still used it. An agent which came in game 726 thsd. made up a new word, *fwz726*, but changed it to *nlf720* after a while. In game 728 thsd. the new word was used by

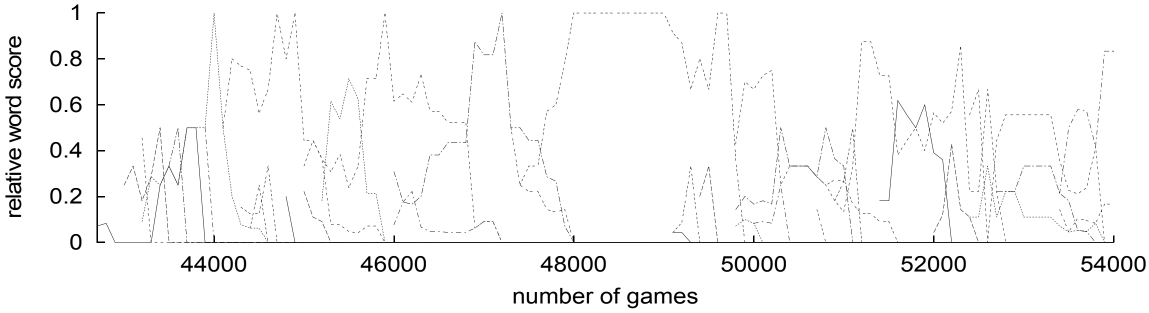

Figure 8. Experiment 3a, Evolution of Relative Word Scores, Flux Rate 50

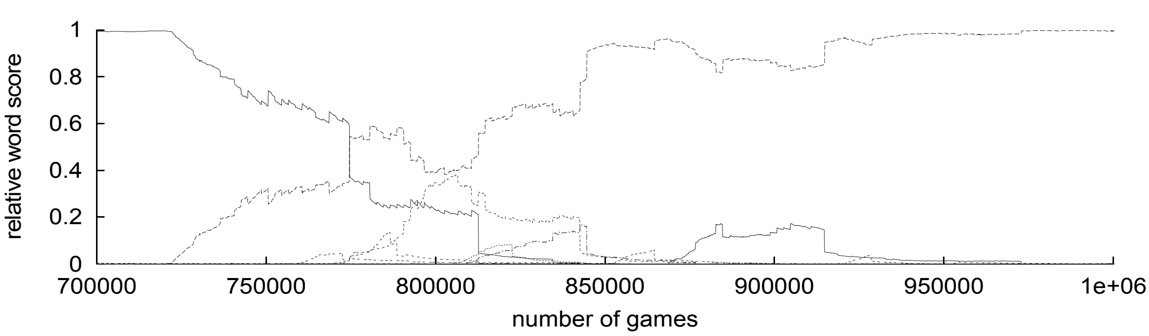

Figure 9. Experiment 3b, Evolution of Relative Word Scores, Flux Rate 2000

4 agents, after 743 thsd. games by 5 agents, and after 746 thsd. by 6 agents. Then some new words appeared, *wnt752*, *shr758*, *vck768*, and *wdp772*, of which *vck768* was the strongest and was used by 5 agents after 800 thsd. games. Eventually these four words disappeared. A few thousand games later, *nlf720* regained its popularity, was used by 7 agents in game 816 thsd. and by all agents in game 845 thsd. The whole transition from *svz0* to *nlf720* took more than 120000 games. During the change, 60 new agents entered the experiment and the whole population thus changed 6 times.

We noticed that it's only the *young* agents that can learn the new word. An *old* agent which have been using an old word for a long time is unlikely to change because it's hard to outcompete the established word. Also, it sometimes happens that a new agent learns a new word, but after communicating with the *old* agents, it starts using the old word.

The second experiment is a much more realistic simulation of a language change in real languages. With flux rate 2000, a shared language emerged and localization success was close to 100% during the whole experiment. The emerged language was very stable (as opposed to experiment 3a) and most of the time the new agents didn’t invent any new words. However, very rarely a change of word for an object did occur and took several generations of agents to accomplish.

6.3 Classes of Concepts

According to results of the experiments we can divide all concepts into three classes: class A – concepts of objects, class B – concepts for *very specific* spatial concepts, mostly concepts of distance, and class C – the remaining spatial concepts (not included in B). The words in all three classes change at different rates; the results are summarized here: **Class A** – all 9 object words changed at flux rates 60 and 80, only 6 changed at FR 100 and none changed at slower FR. **Class B** – most of the words changed even at flux rate 3000. **Class C** – all of the words changed at flux rates 60 to 300, some at flux rates 500 and 1000 and none of the words changed at flux rates 2000 and 3000.

We noticed that the most specific distance concepts (class B) changed faster than the more general ones (class C). The most specific distance concepts were usually used only within a smaller group of agents (4 to 6 agents) and it was easier for the new words to spread in this smaller group. These results suggest that the size of the population is a crucial factor influencing the frequency of word changes. We conducted a set of experiments where we varied the number of agents (4, 10, 15, 20, and 30 agents) and recorded the number of words for objects that changed. We executed 50 experiments for each population size and each experiment ran for 100000 games with flux rate of 1000. The results are summarized here: 4 agents – 259 words changed, 10 agents – 33 words changed, 15 agents – 5 words changed, 20 agents – 1 word changed, and 30 agents – no words changed.

7. Discussion

We have presented a model of a population of artificial agents situated and moving in a 2D environment and using spatial concepts of distance and direction to describe the positions of objects around them. During conceptualization each agent has created its own discrimination trees representing meanings. In the experiments a shared spatial language emerges without the necessity of early conceptualization, fixed word order and/or feedback to speaker. Meanings were not transmitted directly and pointing was necessary only in the early stage of the experiment.

We used multi-generation experiments with agents *flowing* in and out of the community to examine the change of language. In our model the flux of agents alone was sufficient to cause a change in language in contrary to Steels [13,14]. The reason can be that agents in the experiment of Steels used strong lateral inhibition of competing words and meanings thus forcing winner-take-all situation where it was hard for new forms to outcompete the established words. We does not undermine the role of stochasticity in the change of real languages. However we focused on a different type of language change corresponding to invention of neologisms and population migration. Another factor influencing the change of language is the size of the group that is using the word. We have found that rarer words shared by a smaller portion of the population were more prone to change, which corresponds to emergence of local dialects and professional/group jargons.

This model is a first step toward the detachment of representation and more "off-line" communication independent from "here and now" events. We see the direction for further research in extending the model with internal needs/goals for agents together with planning mechanisms providing them with motivation and topics for autonomous communication.

Acknowledgements

This research is supported by the grant VEGA 1/0172/03.

References

[1] Bodík, P.: Emergence of Language in Spatial Language Games. Diploma Thesis, Comenius University, April 2003. http://pauli.fmph.uniba.sk/~8bodik/thesis.pdf

[2] Cangelosi, A., Parisi, D.: The emergence of a language in an evolving population of neural networks. Technical Report NSAL--96004, National Research Council, Rome, 1996.

[3] Cheney, D. L., Seyfarth, R. M.: How Monkeys See the World. In *Inside the Mind of Another Species*. Univ. of Chicago Press, 1990.

[4] Gärdenfors, P.: Language and the Evolution of Cognition. In Rialle, V. Fisette, D. (eds.): *Penser l'esprit: Des sciences de la cognition? une philosophie cognitive*, Presses Universitaires de Grenoble, Grenoble, 1996.

[5] Gärdenfors P.: *Conceptual Spaces,* p. 200, The MIT Press, Cambridge, MA, 2000.

[6] Kirby, S.: Natural Language from Artificial Life. In *Artificial Life*, 8(2), 2002.

[7] Kirby, S., Hurford, J.: Learning, culture and evolution in the origin of linguistic constraints. In Husbands, P., Harvey, I. (eds.), *4th European Conference on Artificial Life*, MIT Press, 1997.

[8] Kirby, S., Hurford, J.: The emergence of linguistic structure: an overview of the iterated learning model. In Cangelosi, A., Parisi, D. (eds): *Simulating the Evolution of Language*, Springer, 2002.

[9] Nowak, M. A., Komarova, N. L., Niyogi, P.: Evolution of universal grammar. In *Science*, 291:114–118, 2001.

[10] Smith, K.: The Importance of Rapid Cultural Convergence in the Evolution of Learned Symbolic Communication. In Kelemen, J., Sosík, P., (ed.), ECAL01, Springer, 2001.

[11] Steels, L.: Constructing and Sharing Perceptual Distinctions. In: van Someren, M., Widmer, G. (eds.): *Proceedings of the European Conference on Machine Learning*. Springer-Verlag, 1997.

[12] Steels, L.: Language as a complex adaptive system. In *Proceeedings of PPSN-VI*, Springer Verlag, 2000.

[13] Steels, L., Kaplan, F.: Stochasticity as a source of innovation in language games. In Adami, C., Belew, R., Kitano, H., Taylor, C. (eds): *Artificial Life VI*, MIT Press, Los Angeles, 1998.

[14] Steels, L., Kaplan, F.: Spontaneous lexicon change. In *COLING-ACL98*, pp. 1243–1249, Montreal, 1998.

[15] Tolman, E. C.: Cognitive maps in rats and men. In *Psychological Review* 55, pp. 189---208, 1948.

[16] Werner, G., Dyer, M.: Evolution of Communication in Artificial Organisms. In: Langton, C., et.al. (ed.) *Artificial Life II.* Addison-Wesley Pub. Co. Redwood City, Ca. p. 659-687. 22, 1991.