

Language evolution in large populations of autonomous agents: issues in scaling

Paul Vogt^{*†}

^{*}Language Evolution and Computation Research Unit, University of Edinburgh
40 George Square, Edinburgh, EH8 9EU, Scotland.
paulv@ling.ed.ac.uk

Federico Divina[†]

[†]Induction of Linguistic Knowledge / Computational Linguistics, Tilburg University
P.O. Box 90153, 5000 LE Tilburg, The Netherlands
divina@cs.vu.nl

Abstract

In this paper we discuss issues relating to modelling language evolution in large populations of autonomous agents that are situated in a realistic environment where they have to evolve and learn means to survive for extended periods of time. As we intend to build such a model in relation to the recently started New Ties project, we identify three major problems that are expected for such a model. The paper proposes some solutions and discusses future directions.

1 Introduction

Language evolution is a *hot* topic in today's sciences; especially in the field of computational modelling, see, e.g., (Cangelosi and Parisi, 2002; Kirby, 2002) for overviews. Typically, the computer models studied are simple, clear and provide useful insights into the origins and evolution of language. However, language is a complex phenomenon and this paper provides an outlook toward more complex models of language evolution.

The computational studies that have been proposed and studied so far have been very useful in investigating particular questions raised by theorists and empiricists in related disciplines, e.g., (De Boer, 2000) and sometimes these studies even have developed new hypotheses (Steels, 1998; Kirby and Hurford, 2002).¹ One limitation of today's state-of-the-art, however, is that most studies only focus on one, or possibly a few aspects of language evolution. This, in itself, is not problematic, but the models that are used to study these aspects typically discard all (or at least many) other aspects in their models, most no-

tably those aspects that have some additional form of complexity with it.

For instance, the studies presented in Vogt (2000) have investigated how a population of physical robots could develop a shared communication system that was perceptually grounded in their environment. However, the population in these studies was of size 2, the agents only communicated about 4 objects that were always present in a context, there was no population turnover, there was no grammatical structure in the communication system and there was no ecological function for the language. These studies have gradually been increased in terms of, e.g., larger population sizes and the number of objects – though without perceptual grounding (Vogt and Coumans, 2003), or evolving simple grammars – though still with small populations of size 6 (Vogt, 2005).

These issues, however, are not really points of critique, but merely an observation of the state-of-the-art. Refraining from complex models is very useful and justifiable. For instance, increasing the number of aspects that one includes in his studies will increase the complexity of one's models in terms of degrees of freedom of, e.g., learning, interactions, analysis and – very important – computational power. So, looking at one – or few – aspects of language evolution has many advantages and allows one to investigate structurally what happens inside his models. However, the limiting complexity can have a pitfall for our studies.

¹Note that these and other studies on which we base our arguments are selected for their high influence in the field. The critiques (or comments) made on these studies apply to all other modelling studies published so far. It should also be noted that although the critiques given are negative, this does not mean that we do not appreciate, like or even adhere to the studies discussed.

For instance, the assumption of using a population of size 2 (cf. Kirby, 2001) or ignoring generational turnover (cf. Steels, 2004) can have a huge impact on the qualitative results of the experiments (Smith and Hurford, 2003; Vogt, 2005). (Note that the studies that discovered such flaws themselves ignore other aspects that, undoubtedly, will lead to qualitatively different results as they too are limited in their set up.)

To what extent then, do we need to complicate our models in order to become more realistic and achieve results that are more likely to be alike real language evolution? The most perfect model of real human language evolution would be the result of reconstructing the real thing. This, however, is not what we want – even if we could do it. However, we should attempt to build models that are beyond our current level of complexity to allow testing hypotheses in large scale simulations that take into account more degrees of freedom in order to become more realistic with respect to the current models. Our aim with the recently started *New Ties* project² is to implement a benchmark simulation that allows a level of complexity far beyond the current state-of-the-art.

In the next section, we will briefly introduce the *New Ties* project and address some problems we think we will encounter. We will discuss how we think we can tackle some of these problems in Section 3. Finally, Section 4 concludes.

2 Identifying the problems

The *New Ties* project aims at setting up a large scale multi-agent simulation in which the population is to learn and evolve a social culture and individual capabilities that enables them to (co-)operate viably in their environment.³ The environment will be modelled loosely after the *Sugarscape* environment (Epstein and Axtell, 1996), which will have a spatial grid, different food sources, tokens, different types of terrain and a large population of agents (Gilbert et al., 2005). We assume that the agents have capacities that will loosely reflect the capacities of early homo sapiens. The agents, which are genetically specified, are supposed to develop a repertoire of behaviours that allow them to survive for extended periods of time. The aim is to have these behaviours develop through individual adaptation, cultural and genetic evolution. The environment will be constrained in such a way that the most efficient way to survive is to develop

co-operation. We allow the agents to evolve language such that they can improve on co-operation.

Although eventually the aim is to have the population evolve a drive and means to evolve language, we will start by assuming that they have this drive and means. This leaves us with the non-trivial problem of having the agents develop a shared communication system. Before identifying some of the problems, it is important to realise that each agent starts its lifetime without any knowledge about the world, so it has no representations of meaning and language. It is also important to mention that each agent acts autonomously; there is no form of telepathy or central control regarding the behaviour of agents. We have identified three major problems we have to deal with in *New Ties*:

1. At each moment in time we aim to deal with a population of around 1,000 agents or more. No known experiment in language evolution has had such a large population size. It is expected that having all agents interact with all other agents leads to an unrealistic scenario and requires a huge number of interactions to arrive at a shared language. However, the agents are distributed spatially across the environment and we do not expect them to travel fast, so the likelihood they will meet every other agent during a lifetime is expected to be low. Nevertheless, we do want them to mix to some extent, but we also believe that learning language in small communities is both realistic and more efficient. So the problem is, how do we control communication?
2. There are a relatively large number of different objects (food items, tokens, agents, roads, places etc.), which are perceived by agents through a fixed, but relatively large number of feature channels. In addition, there are many actions that can be performed. How do we allow agents to categorise the different objects/actions such that they become sufficiently similar to allow language learning (cf. Smith, 2003), and such that these categories are not predefined (i.e. there is typically no one-to-one relationship between object and category)?
3. The contexts in which communication take place are acquired by the agents autonomously. As a result, they may differ from one individual to another (see Fig. 2). In addition, the languages of two individuals may differ, for instance because one of the individuals is still a ‘child’. In brief: if a speaker communicates about one object in its context, how will the hearer infer its reference?

²See <http://www.newties.org>.

³In order to deal with the computational complexity of such a large scale simulation, 50 computers will be connected through the Internet in a peer-to-peer fashion.

And, how do the agents infer the effectiveness of the interaction? These problems are loosely related to what is known as the *Theory of Mind* (ToM).

The next section will present some directions we propose as solutions to this problem.

3 Proposed solutions

3.1 Large populations

In order to deal with large populations, we decided not to treat it as a problem. Instead, we regard it as an asset with which we can learn about how different dialects and languages may evolve. Nevertheless, we do not want each agent in the population to communicate with all other agents, as we believe this will give us huge convergence problems. In addition, we do not want each agent to communicate unrestrictedly with another agent, as this may lead to unlimited chat sessions among agents who happen to be near to each other.

When an agent S sees another agent A in its visual field, it will evaluate, for each object o_i in the context, the function:⁴

$$f(A, o_i) = \nu_1 \cdot SB(A) + \nu_2 \cdot strA(o_i) + T_0 \quad (1)$$

where ν_1 and ν_2 are weights, $SB(A)$ is the *social bond* of S with A , $strA(o_i)$ is the *attention strength* of object o_i , and T_0 is a *talkativeness* parameter.

In order to favour communication with close kin and ‘friends’, we introduce a *social bond* variable $SB(A)$, which is based on the social network an agent constructs (Fig. 1). $SB(A)$ is a function that is proportional to the number of interactions between two agents (it is assumed that agents can recognise each other) and the effectiveness of such interactions (cf. Gong et al., 2004). The relation between parents and offspring will be treated separately. It is assumed that kinship innately promotes $SB(A)$ and may be regulated genetically.

The attention strength $strA(o_i)$ is based on a (possibly large) range of aspects occupying an agent with respect to one of the objects o_i in the agent’s context. For instance, if the agent is hungry, has no food, but sees that another agent carries a food item F , $strA(F)$ gets a high value. The function is part of

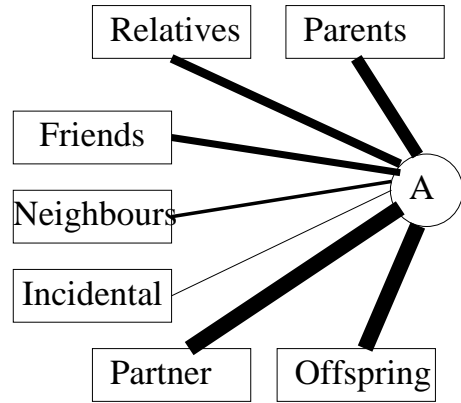


Figure 1: An illustration of an agent’s (A) social network. The thickness of the lines indicate the *social bond* $SB(A')$ of A with another agent A' . See the text for more details.

the agent’s ToM and is defined in Section 3.3, Eq. (2).

The *talkativeness* T_0 is a bias of the agent to communicate. This bias may be genetically specified, but may also be based on learning that communication is useful.

The agent determines which object in the context yields the highest attention strength and the result of $f(A, o_i)$ will be forwarded to an *action decision* mechanism that evaluates which action the agent will take. In this action decision mechanism, the action *communicate* will compete with other possible actions, such as *move-forward* or *eat*. If the agent now communicates about o_i , $f(A, o_i)$ will temporarily remain low for o_i afterwards in order to prevent unrestricted communication.

Given these mechanisms, we expect that there will emerge a self-regulating communication drive, which has a bias to communicate in small communities, but does not exclude communication outside such communities.

3.2 Categorising objects

Categorisation of objects will be based on the *discrimination game* model (Steels, 1996a) and implemented using a form of 1-nearest neighbourhood classification (Cover and Hart, 1967). The aim of the discrimination game is to categorise the object such that it is distinctive from other objects that are in the agent’s context. This need not require that the category is a perfect exemplar of the object. Each object has a number of perceptual features, e.g., shape, colour, weight, location. Objects of a different type

⁴For the current presentation we are only discussing objects to communicate about. Actions will be treated similarly, but are set aside for the time being.

object	f_1	f_2	f_3	f_4	f_5
A_1	0.3	0.7	0.2	0.4	0.1
A_2	0.8	0.2	0.2	0.4	0.5
T_1	0.0	0.0	0.4	0.3	0.3
T_2	0.0	0.0	0.4	0.2	0.5
T_3	0.1	0.5	0.4	0.3	0.3
F_1	0.0	0.0	0.8	0.7	0.8
F_2	0.9	0.3	0.8	0.9	0.8
F_3	0.0	0.0	0.8	0.8	0.6

Table 1: An example context with objects having 5 features f_i . The objects include 2 agents (A_i), 3 tokens (T_i) and 3 food sources (F_i). The features can be perceived as direction (f_1), distance (f_2), shapes (f_3), colours (f_4) individual characteristics (f_5).

may have exactly the same features when they are on the same location. Some objects of the same type, e.g., agents, have the same features in some dimensions, but differ in others to allow identifying individuals.

Table 1 shows an example context containing 8 objects, each perceived with 5 features. Each object has a direction and distance with respect to the perceiving agent (features f_1 and f_2). The objects T_1, T_2, F_1 and F_3 have $f_1 = f_2 = 0.0$ indicating they are carried by the agent in a ‘bag’. All objects of the same type have the same shape (f_3) and often the same colour (f_4). Although for most objects colours are fixed, the colour of food sources (F_i) change over time, indicating the freshness and nutrition of the food. Although individual characteristics (f_5) may be the same for different individuals of the same type, they are typically distinct; this is more so the case for agents (A_i) than for tokens (T_i) or food sources. Across different types, similar individual characteristics can serve as a perceptual basis for analogies.

Each object is categorised by finding for each feature the nearest *categorical feature* $c_{i,j}$, which when combined forms a category (cf. Vogt, 2005). Here i refers to the feature dimension, and j is the j -th categorical feature in that dimension. Suppose an agent’s repertoire of categories (or *ontology*) includes categorical features $c_{3,1} = 0.2$, $c_{3,2} = 0.5$, $c_{4,1} = 0.3$, $c_{4,2} = 0.4$, $c_{4,3} = 0.7$ and $c_{4,4} = 0.85$. Then objects A_1 and A_2 are mapped onto categorical features $c_{3,1}$ and $c_{4,2}$, and the agent can form the category $\mathbf{c}_1 = (c_{3,1}, c_{4,2})$. In principle, all possible combinations of categorical features can be used as a category, so categories $\mathbf{c}_2 = (c_{3,1})$ and $\mathbf{c}_{4,2} = (c_{4,2})$ are also valid categories. In order to prevent a combinatorial explosion of the search space for categories, we are designing heuristics to prevent searching all

possible combinations, such as looking for distinctive categories of the lowest dimension, or by taking combinations that form groups of objects.

Similar to the categorisation of A_1 and A_2 , T_1 , T_2 and T_3 are categorised using categorical features $c_{3,2}$ and $c_{4,1}$; the food source F_1 has categorical features $c_{3,2}$ and $c_{4,3}$; and F_2 and F_3 are categorised using $c_{3,2}$ and $c_{4,4}$. As mentioned, the aim of the discrimination game is to find categories that distinguish an object (or group of objects) from the rest of the context. In this example, only F_1 has distinctive categories. When trying to categorise F_3 , for example, the discrimination game fails, and the ontology has to be expanded (recall that initially, each agent’s ontology is empty). This is done by taking the features of F_3 as exemplars for new categorical features, yielding $c_{3,3} = 0.8$ and $c_{4,5} = 0.8$. Of course when additionally considering all different feature dimensions, the agent may have had categorical features that would distinguish each object from another.

In the language that will be constructed, agents map categories to words. The agent can use a combination of categories to distinguish the object it wants to communicate, thus forming a multiple word utterance. We intend to use this possibility as a means to develop a simple grammar.

3.3 Theory of mind and language games

Probably the biggest problem that this project has to deal with is what we loosely call the *Theory of Mind* (ToM). When a speaker communicates something to another agent, the hearer has to infer what the speaker refers to. When the language is well developed, this may not need to be problematic, but when the communication system of an agent is undeveloped or when the agents speak a different language, this is arguably one of the biggest problems in science. Nevertheless, humans seem to deal with this problem of referential indeterminacy relatively easy. It is commonly accepted that humans have developed (either phylogenetically or ontogenetically) ToM, which relates to the ability to form theories about other individual’s intentions (Bloom, 2000).

Although eventually we intend to evolve some aspects of ToM in New Ties, we shall begin by implementing them directly. The ToM will become an integral part of the language games we will develop. The language game, based on (Steels, 1996a), implements the interaction between two (or more) individuals as illustrated in Table 2. In essence, the agents start by perceiving the context of the game and categorise the objects they see using the discrimination game (DG)

t	speaker	hearer
n	-perceive context -categorisation/DG -focus attention -produce utterance -update lexicon1 -send message	
$n + 1$		-receive message -perceive context -categorisation/DG -focus attention -interpret utterance -update lexicon1 -respond
$n + 2$	-evaluate effect -respond	
$n + 3$		-evaluate effect -respond
$n + 4$	-update lexicon2	-update lexicon2

Table 2: An example scheme for playing a language game between a speaker and hearer. The game may take up to 5 time steps t . See the text for details.

as explained above.

That the contexts of agents typically differ is illustrated in Fig. 2. The context of agent A_1 contains 4 of the 5 food items of type `Food1`, agent A_2 , the contents of A_2 's bag (2 more food items of `Food1`) and the contents of its own bag (1 `Token`, 1 `Food1` and 1 `Food2`).⁵ The context of agent A_2 contains 2 `Tokens`, 2 `Food1` and agent A_1 from the visual field, the contents of A_1 's bag and the contents of its own bag. Due to the physical nature of the environment, we can (and will) not make sure that the contexts of different agents are the same. However, we can introduce aspects of ToM that give the agents cues what the other can see and what the other's intentions are. This will be part of the *focus attention* mechanism. In this mechanism we will assume an attention strength $strA(o_i)$ for object o_i , which is calculated using a formula such as:

$$strA(o_i) = w_1 P_{A'}(o_i) + w_2 V_{A'}(o_i) + w_3 N(o_i) + w_4 I_S(o_i) + w_5 I_{A'} + \dots \quad (2)$$

where w_j are weights and the other arguments are functions that estimate certain aspects of both agents' intentions and knowledge of the current situation. $P_{A'}(o_i)$ is the normalised frequency with which the other agent A' has communicated about object o_i in

⁵Note that obscured objects are not perceived by the agents.




Figure 2: An example situation of a language game. An agent's context is defined by those objects that in its visual field, the contents of the bag of any agent that is in their visual field, including their own. Objects that are obscured by another object are not visible, nor are objects outside the visual field.

the presence of the evaluating agent – the self S . $V(A')(o_i)$ is a function that estimates the likelihood that o_i is in the visual field of A' . $N(o_i)$ is a novelty (or *saliency*) detector, indicating how novel o_i is in the context of S . We assume $N(o_i) = 1$ if o_i first enters the context or when it is shown by another agent; after which it decays following $N(o_i) = e^{-\beta t}$, with β a positive constant and t the time period that o_i is in the context. If an agent explicitly shows an object, the object will also get a high novelty value. $I_S(o_i)$ is a function that calculates the drive of S to communicate about this object, based on its internal states. For instance, if S is hungry, it has a large drive to communicate about food items. Finally, $I_{A'}(o_i)$ is a function that estimates the drive of A' to communicate about object o_i .

The speaker of the language game will use $strA(o_i)$ to select the topic it wants to talk about. If any of the strengths is below a certain threshold, or – at least – lower than any other action to take, then the language game will not proceed. If, however, for some object o_i the value $strA(o_i)$ exceeds any other object attention strength or action value, the language game will proceed with the utterance production of the speaker.

The agents construct and maintain a lexicon in their

	m_1	\dots	m_N
w_1	P_{11}	\dots	P_{1N}
\vdots	\vdots	\vdots	\vdots
w_M	P_{M1}	\dots	P_{MN}
	m_1	\dots	m_N
w_1	σ_{11}	\dots	σ_{1N}
\vdots	\vdots	\vdots	\vdots
w_M	σ_{M1}	\dots	σ_{MN}

Table 3: Two association matrices constructed and maintained as part of an agent’s lexicon. The upper matrix (lexicon1) associates m_j with w_i using conditional *a posteriori* probabilities P_{ij} . The lower matrix (lexicon2) associates categories (or meanings) m_j with words w_i using an association score σ_{ji} .

memories, which is represented by two association matrices as illustrated in Table 3. One of the matrices (referred to as lexicon1 in Table 2) keeps an *a posteriori probability* P_{ij} , which is based on the usage frequencies of associations. The other matrix (lexicon2) keeps an *association score* σ_{ij} , which indicates the effectiveness of an association based on past experiences. The reason for this twofold maintenance is that studies have revealed that when strong attentional cues (such as joint attention or corrective feedback, discussed shortly) guide learning, the lexical acquisition is much faster than when such cues are absent (Vogt and Coumans, 2003). The games using strong attentional cues work fast, because the update mechanism reinforces the score σ_{ij} more strongly than the update of usage based probabilities P_{ij} , which – in turn – work more effectively when the strong attentional cues are absent.

The probabilities are conditional probabilities, i.e.

$$P_{ij} = P(m_j|w_i) = \frac{u_{ij}}{\sum_j u_{ij}} \quad (3)$$

where u_{ij} is the co-occurrence frequency of meaning m_j and word w_i . This usage frequency is updated each time a word co-occurs with a meaning that is either the topic (in case of the speaker) or that is in the context (in case of the hearer). The update is referred to in Table 2 as ‘update lexicon1’. If this principle would be the only mechanism, the learning is achieved across different situations based on the covariance in word-meaning pairs (Vogt and Smith, 2005).

The association score σ_{ij} is updated following:

$$\sigma_{ij} = \eta\sigma + (1 - \eta)X \quad (4)$$

where η is a learning parameter (typically $\eta = 0.9$), $X = 1$ if the association is used successfully in the language game, and $X = 0$ if the association is used wrongly in the language game, or – in case of a successful language game – if the association is competing with the used association (i.e. same word, different meaning; or same meaning, different word). The latter implements lateral inhibition. The update of association scores is referred to in Table 2 as ‘update lexicon2’.

Given these two matrices, the speaker, when trying to produce an utterance, calculates an association strength $strL(\alpha_{ij})$ for each association α_{ij} of a word w_i with meaning m_j , where the meaning is the category that the speaker wants to communicate. This is done using Eq. (5).

$$strL(\alpha_{ij}) = \sigma_{ij} + (1 - \sigma_{ij})P_{ij} \quad (5)$$

This formula neatly couples the two variables. When σ_{ij} is high, the influence of P_{ij} is low, and when σ_{ij} is low, P_{ij} will have more influence. This implements a bias toward basing a choice on known effectiveness vs. estimated probabilities. The speaker will select the association that has the highest strength and utter its word. If no association can be found, e.g., because the lexicon is still empty, the speaker may invent a new word and adds the association to its lexicon.

When the hearer receives an utterance, after it perceives a context and categorises its objects using the DG, it will estimate the attention strength of objects in the context using Eq. (2). Then it calculates for each association of which the word matches the utterance and the meaning matches one of the categorised objects using Eq. (5). The hearer then interprets the utterance using the following equation:

$$\rho_{ij} = \omega_L \cdot strL(\alpha_{ij}) + \omega_A \cdot strA(o_i) \quad (6)$$

where ω_L and ω_A are weights. This equation is based on the model presented in Gong et al. (2004).

Based on its choice, the hearer will respond with some action, which still needs to be specified. An example response could be that the hearer will give the speaker food. The speaker will then (time step $n+2$ in Table 2) evaluate the effect of the language game. If this is what the speaker intended, it can signal the effect to the hearer as response.⁶ In turn, the hearer will evaluate this signal and – if necessary – respond as well. If this finishes the language game, the agents

⁶Here we assume that agents can notice if an action has a positive or negative effect.

can update `lexicon2` using Eq. (4) with $X = 1$ for the used association and $X = 0$ for competing ones if the game is successful. If the game had noticeably failed, then `lexicon2` is updated with $X = 0$ for the used association.

There are many reasons why the language game may fail. For instance, the hearer could not interpret the utterance, or its response does not match the speaker’s intention. In the first case, the hearer can signal a failure as response. In the latter case, the speaker can signal a failure. In both cases, the game will need to be repaired in order to allow significant learning.

For now, we will assume that the initiative to repair the game lies with the speaker. For example, the speaker can ignore the failure when the hearer was not the direct addressee, or when the social bond is low and the speaker wishes to proceed with another action. The speaker can also decide to do one of the following things in order to provide the hearer with additional cues about which object is the reference of the game:

- show an object from the bag;
- point to an object in the context by looking in its direction;
- show an action;
- go to the object;
- ...

Using these cues, the hearer tries to reinterpret the utterance with a strong additional bias to the shown object, and the game is re-evaluated. We will implement a mechanism to prevent this continuing forever; for instance by allowing only one or two reinterpretations.

If the hearer did not have an interpretable association of the utterance in the first place, it will adopt the utterance and add a new word-meaning association to its lexicon. The initial value of $\sigma_{new,j}$ will be based on existing associations with word w_j – if any – and the attention strength of object o_{new} according to

$$\sigma_{new,j} = k \cdot (1 - \max_i(\sigma_{i,j})) \cdot strA(o_{new}) \quad (7)$$

where we assume that $\sigma_{new,j}$ relates to meaning m_{new} that is a distinctive category of object o_{new} . (Note that there may be more than one such association.) The association(s) will be added to `lexicon1` with an initial usage of $u_{new,j} = 1$.

To summarise, we intend to extend the familiar language game model in order to include aspects of ToM. The language game is largely based on

the *guessing game*, which uses corrective feedback to guide meaning inference, and a game that uses cross-situational statistical learning (Vogt and Smith, 2005). The cues as formalised in Eqs. (2) – (7), together with the repair mechanisms, are the core mechanisms of the ToM. Initially we intend to hard-wire the ToM into the New Ties project, but at some stage we wish to evolve this – for instance by evolving the various weights of Eqs. (2) and (6).

4 Conclusions

In this paper we identify three major problems regarding modelling language evolution in large populations of autonomous agents, such as proposed in the New Ties project. The problems and proposed solutions can be summarised as follows:

1. How can we control communication in large populations? We intend to treat this as a minor problem by limiting communication based on the spatial location of agents and the social networks they develop. In addition, to provide well structured learning environments for the young agents, we will treat close kinship relations as an extra drive to communicate.
2. How can we categorise a large number of objects such that they are learnable in language? To solve this problem, we propose a model based on Steels’ discrimination games (Steels, 1996b) where perceptual features are categorised following the implementation of Vogt (2005). To deal with overlapping classes of objects we intend to develop heuristics that group categorical features that are similar across different objects.
3. How do we deal with issues relating to the Theory of Mind? This problem is identified as the hardest problem. In order to deal with it, we propose to design mechanisms that allow an individual to use perceptual and interpretational information to provide cues concerning the objects that the other agent is likely to communicate about. These mechanisms will be integrated in the familiar language game models used earlier in Vogt and Coumans (2003) similar to the way proposed by Gong et al. (2004). In addition, social strategies are proposed in order to repair failures in language games.

We are currently investigating stability conditions of social networks in relation to the scaling of populations. In addition, we are implementing a simple version of the ToM to prove the principle.

We believe that, although the problems we identified are hard, we can scale up models of language evolution successfully much in the way we discussed. If we succeed, we expect that this experiment will provide an exciting benchmark for many large scale experiments regarding the evolution of language.

Acknowledgements

The NewTies project is partially supported by an EC FET grant under contract 003752.

References

- P. Bloom. *How Children Learn the Meanings of Words*. The MIT Press, Cambridge, MA. and London, UK., 2000.
- A. Cangelosi and D. Parisi, editors. *Simulating the Evolution of Language*. Springer, London, 2002.
- T. M. Cover and P. E. Hart. Nearest neighbour pattern classification. *Institute of Electrical and Electronics Engineers Transactions on Information Theory*, 13:21–27, 1967.
- B. De Boer. Self-organization in vowel systems. *Journal of Phonetics*, 28:441–465, 2000.
- J. M. Epstein and R. Axtell. *Growing artificial societies: social science from the bottom up*. MIT Press, Cambridge, MA., 1996.
- N. Gilbert, S. Schuster, M. den Besten, and L. Yang. Environment design for emerging artificial societies. In *this issue*, 2005.
- T. Gong, J. Ke, J. W. Minett, and W. S-Y. Wang. A computational framework to simulate the co-evolution of language and social structure. In Jordan Pollack, Mark Bedau, Phil Husbands, Takashi Ikegami, and Richard A. Watson, editors, *Artificial Life IX Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*, pages 214–219. The MIT Press, 2004.
- S. Kirby. Spontaneous evolution of linguistic structure: an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2):102–110, 2001.
- S. Kirby. Natural language from artificial life. *Artificial Life*, 8(3), 2002.
- S. Kirby and J. R. Hurford. The emergence of linguistic structure: An overview of the iterated learning model. In A. Cangelosi and D. Parisi, editors, *Simulating the Evolution of Language*, pages 121–148, London, 2002. Springer.
- A. D. M. Smith. Intelligent meaning creation in a clumpy world helps communication. *Artificial Life*, 9(2):559–574, 2003.
- K. Smith and J. R. Hurford. Language evolution in populations: extending the iterated learning model. In W. Banzhaf, T. Christaller, J. Ziegler, P. Dittrich, and J. T. Kim, editors, *Advances in Artificial Life: Proceedings of the 7th European Conference on Artificial Life*, Lecture Notes in Computer Science, pages 507–516. Springer-Verlag, Heidelberg, 2003.
- L. Steels. Emergent adaptive lexicons. In P. Maes, editor, *From Animals to Animats 4: Proceedings of the Fourth International Conference On Simulating Adaptive Behavior*, Cambridge Ma., 1996a. The MIT Press.
- L. Steels. Perceptually grounded meaning creation. In M. Tokoro, editor, *Proceedings of the International Conference on Multi-Agent Systems*, Menlo Park Ca., 1996b. AAAI Press.
- L. Steels. The synthetic modeling of language origins. *Evolution of Communication*, 1(1):1–34, 1998.
- L. Steels. Constructivist development of grounded construction grammars. In W. Daelemans, editor, *Proceedings Annual Meeting of Association for Computational Linguistics*, 2004.
- P. Vogt. Bootstrapping grounded symbols by minimal autonomous robots. *Evolution of Communication*, 4(1):89–118, 2000.
- P. Vogt. The emergence of compositional structures in perceptually grounded language games. *Artificial Intelligence*, 2005. In press.
- P. Vogt and H. Coumans. Investigating social interaction strategies for bootstrapping lexicon development. *Journal for Artificial Societies and Social Simulation*, 6(1), 2003. <http://jasss.soc.surrey.ac.uk>.
- P. Vogt and A. D. M. Smith. Learning colour words is slow: a cross-situational learning account. *Behavioral and Brain Sciences*, 2005. In press.