Modelling language evolution in a complex ecological environment*

ILK Research Group Technical Report Series no. 06-01

Federico Divina[1] and Paul Vogt[1,2]

[1] Induction of Linguistic Knowledge / Computational Linguistics section
Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands
[2] Language Evolution and Computation Research Unit, School of Philosophy, Psychology and Language Sciences, University of Edinburgh, U.K.
f.divina@uvt.nl,paulv@ling.ed.ac.uk

Abstract. This paper introduces a new model that has been developed for simulating language evolution in the New Ties project. This project aims at evolving a cultural society by integrating evolutionary, individual and social learning in large scale simulations. The model presented here introduces a more natural implementation of language games with various novel features. Some preliminary results are presented and discussed. As the project is still in its developmental stage no staggering results were obtained yet but we conclude that the current model is very promising for future research into the evolution of language and social behaviour.

1 Introduction

Human language is thought to have evolved from an interaction between three adaptive systems: biological evolution, individual learning and cultural evolution [1]. This evolution is thought to be constrained and driven by the embodiment of humans and their situatedness in the ecology of our world. The New Ties project[3] aims at merging these aspects in a large scale simulation to evolve a cultural society of simulated agents who are situated in a complex environment. One important aspect of this simulation is to evolve language that allows the social learning of skills.

Although a lot has been achieved with computational modelling of language origins and evolution (see, e.g., [2–4] for overviews), such models necessarily have to simplify a great deal with respect to the real world, even if processed in the real

* This research and the New Ties project is supported by an EC FET grant under contract 003752. We thank all members the New Ties project for their invaluable contributions. Opinions and errors in this manuscript are the authors' responsibility, they do not necessarily reflect those of the EC or other New Ties members.

[3] New Ties stands for New Emerging World models Through Individual, Evolutionary and Social learning. See http://www.new-ties.org.

world using real robots [4]. Of course, simplification is very useful to gain insights from simulations that only look at one particular aspect of language evolution. Such aspects vary from the evolution of sound systems [5], syntax [6], grounded lexicons [7–9] to grounded grammars [10, 11]. The problem of simplifications are that results achieved may not hold in more complex simulations. For instance, Vogt [11] has shown that grammatical structures can emerge under completely different conditions than those reported by Kirby et al. [6] if the meanings are perceptually grounded and acquired from scratch, and if the language is acquired using a slightly more complex learning mechanism.

The New Ties project aims to combine various aspects of language evolution models in a world that contains many agents who need to survive in a complex environment that has quite some aspects similar to our own world. Agents are to acquire behaviours that allow them to survive by combining evolutionary learning (i.e., genetic evolution), individual (reinforcement) learning and social learning. One aspect of social learning involves language learning to allow cultural evolution of language. In turn, this evolved language will be used to transfer acquired skills culturally, which thus is the second aspect of social learning involved. From the evolution of language point of view, the New Ties project will allow us to investigate many questions concerning language evolution in a realistic scenario. The sorts of questions we may ask include, for example: Under what environmental constraints will language evolve? What type of learning and interaction mechanisms are required for a language to evolve? Will dialects emerge? And if so, what are the dynamics underlying its emergence? How can learning mechanisms evolve biologically?

In the next section we will briefly present the New Ties project. In Section 3 a more detailed description of model that allows the population to evolve language is provided. This model implements language games [12], modified in various ways to become more natural. Some preliminary results are presented in Section 4. Section 5 concludes the paper.

2 New Ties

As mentioned, the objective of the New Ties project is to set up a simulation in which a large population of agents (i.e., 1,000+) are to evolve a cultural society using evolutionary, individual and social learning. Sub-objectives include investigating the interaction between these three adaptive systems and evolving a communication system that aids social learning. The software will be made publicly available to allow other researchers to use the platform and test their own agents.

The New Ties world is a grid world in which each point is a location. The world is set up with places of varying roughness, tokens, edible plants, building bricks, and agents (see [13] for a full description). Agents are provided with sensors and actuators that allow them to see and act. The sensors are configured such that an agent sees a number of perceptual features about the various objects, events and other aspects in their world. The actuators allow the agents to

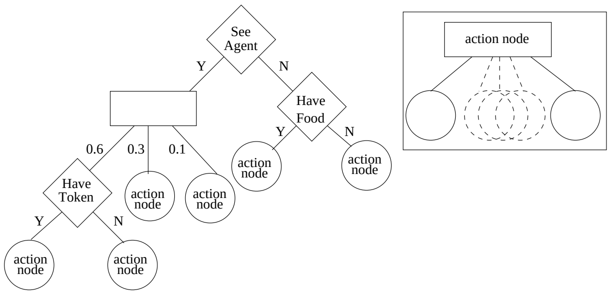

Fig. 1. A simplified example of a DQT. The diamond boxes are test nodes in which a category (or conjunction of categories) is tested. The rectangular boxes are bias nodes from which one branch is selected with a probability function relating to an innate bias and a learnt bias (the learnt bias is adapted using reinforcement learning). Each leave of the tree ends with an action node, which further expands as a bias node and all possible actions consistent with the test nodes (see box). Note that in practise, the DQT will be far more complex than this one.

move forward, turn left or right, pick up and put down objects, give and take objects to/from other agents, hit other agents (possibly to death), eat, build roads and barriers, mate, shout and talk. (Shout is directed to all agents within vicinity, while talk is only directed to one selected agent within the visual field.) Each action costs energy, the amount of which depends, for instance on the weight carried by the agents or the roughness of the terrain. When the energy reaches a level less or equal to zero, the agents die (they can also die of old age). Eating plants increases the agent's energy level, the amount of which depends on the 'ripeness' of the plant.

The world can be configured to set *challenges* for the agents (i.e. pressures to evolve particular types of behaviour that allow the agents to survive in this world). For instance, the world can be configured such that food only grows in distinct places at different times, as regulated by seasons. This would require the agents to either evolve a sort of trading system to share food or to become nomadic [13].

Agents develop their own control system mainly using individual learning. This control system is a *decision Q-tree* (DQT), see Fig. 1, which is a decision tree that can change using reinforcement learning [14]. The agents are 'born' with an innate tree, whose structure is specified by the agent's genome. Using this DQT, an agent reacts on the input (i.e. categories representing the current visible situation). The resulting action is then evaluated yielding some reward, which changes the Q-values attached to each node in the decision tree. Through exploration, the agents can insert or delete nodes. Insertion of new nodes may be guided by social learning (i.e. by inserting nodes to align parts of the tree communicated by other agents).

The genome carries, apart from the initial DQT, a number of biases influencing the behaviours of agents regarding certain aspects [14], such as the tendency to be aggressive or social. The social bias regulates, for instance, the frequency with which agents communicate or help each other with learning language (see Section 3). The genome is subject to mutation and cross-over as in standard GAs, but the reproduction cycle is asynchronous. When agents achieve the adulthood life stage, a male can ask a female to mate. If the female accepts, the agents mate and produce offspring. The child is then initialised with a DQT and biases specified by the genome. During their childhood, agents have the innate tendency to follow the first agent they see (usually the mother), which allows the agent to learn from interacting with their mother.

Interaction is achieved by the predefined production and interpretation mechanisms, which resemble the language game model [12], though modified to be more natural and stochastic as described in the next section. Communication will be about parts of the path traversed through the DQT to decide on an action that time step. This action may be a talk or shout action, but also something else. The information communicated allows the receiver to reconstruct a part of the speaker's DQT, which may then be used to guide the insertion of new nodes. As this aspect of social learning is still under development, it will not be further addressed in this paper.

In order to facilitate the processing of these complex simulations, a peer-to-peer network consisting of 50 PCs or more is constructed, where each PC processes a dynamically changing part of the world in parallel [15]. It is clear that such simulations will generate huge amounts of data, for which data mining techniques are being developed to analyse these and detect emergent patterns of cultural behaviour. The current state of the project is that each adaptive system (biological evolution, individual learning and social learning) has been developed to a large extent independently. These systems are now being tested before they will be integrated. The remainder of this paper reports on some initial tests relating to the language evolution mechanisms as part of the social learning system.

3 Language evolution in New Ties

Figure 2 shows the basic architecture of the New Ties agents. The architecture consists of four modules, which are processed in sequential order from top to bottom. In addition, each agent has a short term memory (STM) and long term memory (LTM). The input to an agent includes perceptual input regarding all objects an agent can see in its visual field and all messages that were sent within its audible range. The perceptually observed objects constitutes the agent's *context.* The output are actions and messages sent by the agent.

The perceptual input is first send to the categorisation module, where each feature of each object is categorised with its nearest category. Categories are represented as 1-dimensional prototypes (each dimension relating to a feature dimension), some of which are innately specified, while some can be acquired during development by playing discrimination games similar to the one described in [11]. All categories are stored in the LTM. For the experiment described

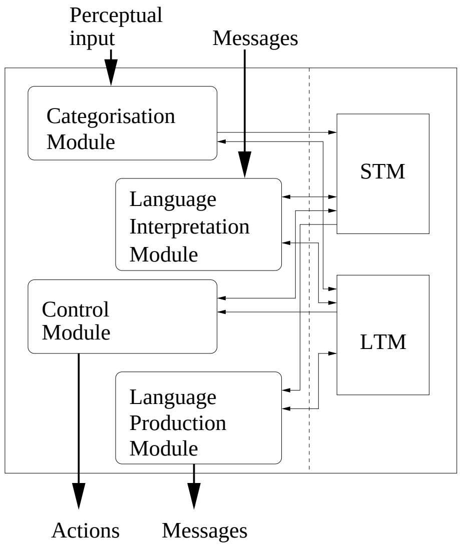

Fig. 2. The basic agent architecture of New Ties; see text for details.

$$
\begin{array}{c|ccc}
 & m_1 & \dots & m_N \\
\hline
w_1 & \sigma_{11} & \dots & \sigma_{1N} \\
\vdots & \vdots & \vdots & \vdots \\
w_M & \sigma_{M1} & \dots & \sigma_{MN}
\end{array}
\qquad\qquad
\begin{array}{c|ccc}
 & m_1 & \dots & m_N \\
\hline
w_1 & P_{11} & \dots & P_{1N} \\
\vdots & \vdots & \vdots & \vdots \\
w_M & P_{M1} & \dots & P_{MN}
\end{array}
$$

Fig. 3. A simplified illustration of the lexicon. The lexicon consists of two matrices associating meanings m_j with words w_i. The left matrix stores association scores σ_{ij} and the right matrix stores co-occurrence probabilities P_{ij}. See the text for details.

in this paper, all categories are predefined. The categorisation yields for each object a set of categories that describes the object. If the object is an agent, the categorisation includes the action performed by that agent. All resulting category sets are stored in the STM for further processing by other modules.

3.1 Interpretation

All messages an agent receives are processed by the language interpretation module (LIM). A message can consist of multiple words. For each word an agent receives, the LIM will do the following: The agent will start by searching its lexicon (stored in the LTM) for entries that match the word. As described in [16] the lexicon is represented by two association matrices (Fig. 3), one that maintains association scores σ_{ij} and one that maintains a posteriori probabilities $P_{ij} = P(w_j|m_i)$ of finding word w_j, given meaning m_i. The association scores contain information about the association's effectiveness as evaluated through feedback.

However, since we assume feedback is not always provided, nor is it always accurate, the agents also maintain the co-occurrence probabilities allowing for *cross-situational statistical learning* [17]. This latter method assumes that across different situations a word always co-occurs with its meaning, thus allowing an agent to learn the proper meaning without the need for feedback. The problem with this is approach, however, is that cross-situational learning is a slow learner, it requires consistent input and it is difficult to scale up in terms of population size [18]; all conditions most likely not met.

If an agent searches its lexicon, it selects the association matching the heard word and of which the association strength $strL_{ij}$ is highest. This association strength is a coupling between the two scores σ_{ij} and P_{ij}:

$$strL_{ij} = \sigma_{ij} + (1 - \sigma_{ij})P_{ij} \tag{1}$$

This coupling assures that the association strength relies more on the association score σ_{ij} if it is high (i.e., it has been effective in previous interactions); otherwise $strL_{ij}$ relies more on the co-occurrence probability P_{ij}.

The speaker may have pointed to an object that closely relates to the message's meaning. (This object may be the entire meaning, but it also may be that only a few features of this object constitutes the meaning or even that a few features constitute only a part of the meaning.) If an object is pointed to, the LIM verifies if this is consistent with (a part of) the interpreted meaning. If this is the case, the interaction (aka the language game) is considered a *success*. If the object does not relate to the interpretation, there are two possibilities: First, the association score σ_{ij} exceeds a threshold Θ, in which case it is assumed that the interpretation is correct, but the speaker got it wrong. In that case, the hearer may send a hard-wired 'error' message as a form of feedback. This message is send with a probability β based on the social bias (see above) and the social bond the agent has with the speaker. This social bond is a score based on the number of interactions the agent has had with the speaker and is maintained in the LTM. Second, if $\sigma_{ij} \leq \Theta$, the interpretation is assumed to be wrong, so the association score σ_{ij} is lowered and the word is associated with all categories relating to the pointed object.

If no object is pointed to and $\sigma_{ij} > \Theta$, then the game is considered a success. Otherwise, the agent either sends a 'did not understand message' (dum) or nevertheless assumes success. This decision is also made with a probability proportional to β. The higher β, the higher the chance a 'dum' message is send. The rationale behind this is that the more social an agent is, the more inclined it is to put more effort in correctly understanding the speaker.

If the language game was considered a success for all words in a message, the interpretation of the message is added to the context in the STM for further processing by the DQT. With a probability proportional to β, the hearer will send the speaker an 'understood message' (um) as part of the feedback. Also the interpretation's association scores σ_{ij} are increased using the following equation:

$$\sigma_{ij} = \eta \cdot \sigma_{ij} + 1 - \eta \tag{2}$$

where $\eta = 0.9$ is a learning parameter. In addition, the hearer lowers all competing association scores σ_{kl} ($k = i$ or $l = j$, but not both) following:

$$\sigma_{kl} = \eta \cdot \sigma_{kl} \tag{3}$$

In the cases where a 'dum' message was sent (i.e., the game is considered to have failed), the association scores of the interpretation are also lowered using this formula. In addition to these updates, in all cases the co-occurrence probability of this word with all meanings (i.e., categories) in the current context is increased. (Note, by the way, that the agent may not have seen the object relating to the word's meaning, so interpretations and/or adaptations may be wrong.)

If the agent did not find an interpretation in the current context, the agent will check if the speaker pointed at an object. If so, the agent associates the word with all categories relating to that object. Otherwise, the agent associates the word with all categories relating to all objects in the current context. In this latter case, the initial association scores are inversely proportional to the maximum association score a meaning already has with other words. This way, there is a bias that the word's meaning is one that has not yet an association with another word. This could be seen as a loose implementation of the *principle of contrast* [19] observed in children's acquisition of word-meanings.

3.2 Production

When the LIM has finished processing, the control module will process the DQT using all categories resulting from the categorisation and language interpretation modules as input. By traversing the tree following the results of the test nodes (each test node checks for the truth value of a category) and following the decisions made at the bias nodes (Fig. 1), the controller reaches a leave node, which is some action to be performed. Irrespective of whether or not the action is to talk or shout, the language production module (LPM) is started, because even if the action is not to talk or shout, the LPM may nevertheless decide to communicate about something.

First, however, the LPM checks if the agent has produced an utterance in the previous time step. If so, the LPM will evaluate the success of that interaction. To this aim, it will see if either an understood message (um) or an error message was received. If an 'um' was received, the game of the previous time step is considered successful and the used association score is increased following Eq. (2) and competing ones are inhibited following Eq. (3). If an error message was received, the game is considered to have failed and the used association score is lowered using Eq. (3).

Next, the speaker will decide whether to produce a message or not. If the control module has decided the agent should talk or shout, the LPM will produce a message. Otherwise, the agent may decide – with a probability proportional to β – to produce a message if the agent received a 'dum' message or if the agent produces a different action than in the previous time step (i.e. in the case of a novel action). In the case a 'dum' message was received, the agent will produce a (possibly different) message about the same meaning as before, but now will include pointing. (Note that this message may be different if the agent has meanwhile changed its lexicon.) Otherwise, a meaning is selected as follows.

First a task complexity C_t is chosen. The task complexity is a value that indicates how many words the message will contain. For now C_t is a value between 1 and 5, and it is determined by generating a random number following a Gaussian distribution with a mean equal to the average age of other agents in the context in tens of years and a standard deviation of 0.75. If the number is larger than 5, task complexity is set to 5. This way, the agent will tend to speak short sentences to younger agents and longer sentences to older agents. The rationale behind this formula is that short sentences are easier to interpret by less skilled language users than longer sentences. In addition, this may form an interesting basis for bootstrapping compositionality [20].

Once C_t is set, the LPM decides on the actual meaning to be communicated. This meaning is selected from the *trace* of the DQT. The trace is the list of test nodes (i.e. categories) that have been traversed through the DQT in order to reach the final action node. This action node is also included. The agent selects a position (Pos) within this list and then selects the categories using one of the following three options at random:

1. the first C_t categories starting at Pos,
2. the first $C_t - 1$ categories starting at Pos plus an action node, or
3. the action node.

Note that an action node may be a 1-, 2-, or 3-place predicate, such as *move-forward*, *pick_up(o)* or *give(a,o)*, where o is an object and a is an agent. The position Pos is also determined generating a random number following a Gaussian distribution with a mean equal to the average age of an agent and some standard deviation. This number is then scaled to the size of the trace.

So, the younger the communication partner the higher the position of the meaning in the tree. The reason for this mechanism is that children in the model are born with an innate DQT constructed from the genomes of its parents. Hence, the likelihood that the child will have the communicated structure in its own DQT is fair, which increases the chance of the child being able to learn the language well.

The meaning thus selected consists of at most 5 categories. For each category, the LPM searches its lexicon for associations of which the meaning matches the category and for which the association strength $strL_{ij}$ (Eq. 1) is highest. The associated word is then appended to the message. If a category has no entry in the lexicon yet, a new word is created as a random string and the new association is added to the lexicon. It is important to realise that agents are 'born' with an empty lexicon.

Once a message is thus constructed, the LPM decides whether the agent will point to an object that directly relates to the message's meaning. This decision is, again, made with a probability proportional to the socialness bias β. So,

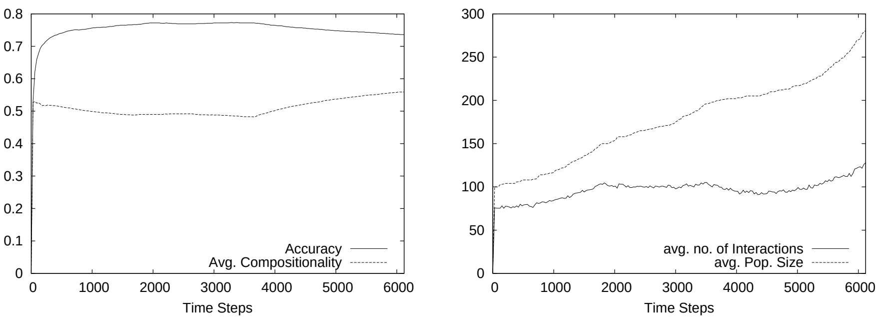

Fig. 4. These graphs show the evolution of accuracy and compositionality (left) and population size and number of games per 30 time steps (right) of one simulation run.

the more social the agent is, the more likely it is to provide its communication partner with hints as to what it is referring to.

4 Preliminary results

In this section we present some preliminary results obtained with a somewhat simplified version of the model presented above. The reasons for these simplifications is that we are still in a developmental stage and we first want to test some mechanisms using a stripped model. The simplifications involve the following: 1) Instead of taking the meanings from the DQT trace, speakers communicate about a number of features of one object in their visual field. The number of features is determined by the task complexity C_t. 2) Pointing is implemented by transferring only the perceptual features of the object the speaker is communication about. Ideally, it will transfer all perceptual features of an object relating to the meaning being conveyed, allowing the hearer to search for the object that most closely resembles these features. The current simplification is in effect explicit meaning transfer, which will be avoided in future studies. 3) There is no evolution on the genome and there is no individual learning. As a result, all agents have the same DQT and the socialness gene is not subject to evolution, but generated randomly. In the simulation reported here, agents only move around, eat, reproduce without changing their genomes and communicate with each other.

Figure 4 shows the results of one single simulation that was run for slightly more than 6,000 time steps, which comes to about 18 years in New Ties time.[4] The left graph shows communicative accuracy, which is the average number of successful interactions (i.e. the hearer interpreted the speaker's message correctly) during an interval of 30 time steps, and compositionality, which measures

[4] We intended to show a longer simulation, but failed to complete that before the deadline of SAB. This simulation took about 4 days to run on a single PC. We expect include state-of-the-art results in the final camera-ready version.

the proportion of interactions that used more than one word averaged over an interval of 30 time steps. Accuracy rapidly rises to a level around 75% during the first 1,000 time steps, but starts to decrease from around 4,000 time steps. This occurs at the same moment compositionality increases from around 50% to 55%. This sudden change is marked by the fact that the first agents start to reach an age higher than 10, which is a turning point that triggers task complexity to increase. Before that stage the mean fed into the Gaussian used to calculate the task complexity cannot be larger than 1 as the mean is the age in tens of years (so 1 corresponds to an age of 0-10, 2 to 11-20, etc.).

Note that since the population is still increasing, not all agents have reached the age of 10 yet. The rightmost graph of Fig. 4 shows the average population size and the average number of interactions, both during time periods of 30 time steps. This graph shows that the population size is increasing steadily from 100 agents to approximately 275 agents. This increase is due to the reproduction of agents, while only few agents died from lack of food as there was plenty of food around in this set up, and no agents had died yet of old age, since no agent has reached old age.

The average number of interactions shows that many agents produce a message each time step. Over the entire period of 6,000 time steps, the average was around 100 interactions, so in total around 600,000 interactions were performed. Of these interactions, feedback was provided in only 12%, while pointing was carried out in only 28%. So in 60% of all interactions there was no explicit meaning transfer. Nevertheless, the population was quite successful in communication. These results show that the current coupling of cross-situational learning with cases where feedback is evaluated works well, as was previously shown in a simulation of the Talking Heads experiment [21].

5 Conclusions

This paper presents the model concerning language evolution in the New Ties project. This project aims at setting up large scale simulations to study the evolution of cultural societies by combining evolutionary, individual and social learning techniques. Language evolution is modelled through agents' interactions (or language games), whose mechanisms are predefined, but work stochastically. Interactions are initiated by novel actions and some mechanisms are only carried out stochastically following a socialness bias. The learning mechanisms allow each individual agent to invent new words and acquire words from other agents.

The preliminary results show that the main principles work well, though some mechanisms were simplified, so these still need to prove their value. For instance, pointing is currently simplified by exchanging the exact meaning, allowing the hearer's to verify or acquire the intended meaning. If instead the speaker points to an object that only relates to a part of the meaning as is intended in future simulations, the hearer will have more difficulties identifying the exact meaning conveyed by the speaker. In the easy case, where the speaker communicates one word referring to one feature of an object it points to, there are on average 10 different features of the object this word may refer to. As there is no way an agent can point at individual features, the hearer can only depend on cross-situational learning of this word. The context size in this case is 10, which is relatively large, given the finding that learning speed in cross-situational learning under ideal conditions increases almost log-exponentially with the context size [22]. The ideal condition, however, that the input to the hearer is consistent is not met because different agents will use different words for expressing the same meaning, which makes cross-situational learning hard for larger populations [18]. Possible solutions could include the use of mutual exclusivity [23], which relates to the principle of contrast used here.

In the current simulation, all agents talk about aspects of objects, whose meanings are transmitted explicitly. It is interesting to investigate how the language will evolve when the agents will communicate about parts of their own controller which may differ from one agent to another as the result of genetic evolution or reinforcement learning. In that case, the meaning conveyed may relate to more objects (e.g. *eat(food)*), so one cannot point to both food and the action eat. This will thus increase the complexity of learning the words' meanings. It is hoped that by varying the task complexity (in addition to the cross-situational learning) agents will be able to learn the meaning of single words that will allow learning the more complex expression in a similar way as shown in [20, 11]. For instance, if the hearer knows the word for food, but not the word for eat, she may learn the right association if she hears "eat food" and sees the action *eat(food)*.

Once the project has further developed and all different techniques, such as evolutionary, individual and social learning, have been integrated with this language evolution model, the resulting platform could be used to investigate many interesting aspects of (human) evolution. These aspects include the dynamics of cultural evolution, the interaction between evolutionary, individual and social learning, social interaction mechanisms and cognitive capacities required for social behaviour and language to evolve, and the effect environmental constraints have on the evolution of a cultural society. Since the New Ties project has very ambitious aims, we will invite other researchers to use our platform to investigate their own ideas.

References

1. Kirby, S., Hurford, J.R.: The emergence of linguistic structure: An overview of the iterated learning model. In Cangelosi, A., Parisi, D., eds.: Simulating the Evolution of Language, London, Springer (2002) 121–148
2. Briscoe, E.J., ed.: Linguistic evolution through language acquisition: formal and computational models. Cambridge University Press, Cambridge (2002)
3. Cangelosi, A., Parisi, D., eds.: Simulating the Evolution of Language. Springer, London (2002)
4. Vogt, P.: Language evolution and robotics: Issues in symbol grounding and language acquisition. In Loula, A., Gudwin, R., Queiroz, J., eds.: Artificial Cognition Systems, Idea Group (2006)
5. De Boer, B.: The origins of vowel systems. Oxford University Press, Oxford (2001)

6. Kirby, S., Smith, K., Brighton, H.: From UG to universals: linguistic adaptation through iterated learning. Studies in Language **28**(3) (2004) 587–607
7. Marocco, D., Cangelosi, A., Nolfi, S.: The emergence of communication in evolutionary robots. Philosophical Transactions: Mathematical, Physical and Engineering Sciences **361**(1811) (2003) 2397–2421
8. Steels, L., Kaplan, F., McIntyre, A., Van Looveren, J.: Crucial factors in the origins of word-meaning. In Wray, A., ed.: The Transition to Language, Oxford, UK, Oxford University Press (2002)
9. Vogt, P.: Anchoring of semiotic symbols. Robotics and Autonomous Systems **43(2)** (2003) 109–120
10. Steels, L.: The emergence and evolution of linguistic structure: From lexical to grammatical communication systems. Connection Science **17**(3-4) (2005) 213–230
11. Vogt, P.: On the acquisition and evolution of compositional languages: Sparse input and the productive creativity of children. Adaptive Behavior **13(4)** (2005) 325–346
12. Steels, L.: Emergent adaptive lexicons. In Maes, P., ed.: From Animals to Animats 4: Proceedings of the Fourth International Conference On Simulating Adaptive Behavior, Cambridge Ma., The MIT Press (1996)
13. Gilbert, N., Schuster, S., den Besten, M., Yang, L.: Environment design for emerging artificial societies. In: Proceedings of AISB 2005: Socially inspired computing joint symposium. (2005)
14. Griffioen, A., Schut, M., Eiben, A., Bontovics, A., Hévízi, G., Lõrincz, A.: New ties agent. In Edmonds, B., Gilbert, N., Gustafson, S., Hales, D., Krasnogor, N., eds.: Proceedings of the Socially Inspired Computing Joint Symposium (AISB 2005), AISB (2005)
15. Craenen, B., Paechter, B.: Peer-to-peer networks for scalable grid landscapes in social agent simulations. In Edmonds, B., Gilbert, N., Gustafson, S., Hales, D., Krasnogor, N., eds.: Proceedings of the Socially Inspired Computing Joint Symposium (AISB 2005), AISB (2005)
16. Vogt, P., Divina, F.: Language evolution in large populations of autonomous agents: issues in scaling. In Edmonds, B., Gilbert, N., Gustafson, S., Hales, D., Krasnagor, N., eds.: Proceedings of AISB 2005: Socially inspired computing joint symposium, AISB (2005) 80–87
17. Vogt, P., Smith, A.D.M.: Learning colour words is slow: a cross-situational learning account. Behavioral and Brain Sciences **28** (2005) 509–510
18. Vogt, P., Coumans, H.: Investigating social interaction strategies for bootstrapping lexicon development. Journal for Artificial Societies and Social Simulation **6(1)** (2003) http://jasss.soc.surrey.ac.uk.
19. Clark, E.V.: The lexicon in acquisition. Cambridge University Press (1993)
20. De Beule, J., Bergen, B.: On the emergence of compositionality. In Cangelosi, A., Smith, A., Smith, K., eds.: Proceedings of Evolang 6, World Scientific Publishing (2006)
21. Divina, F., Vogt, P.: Perceptually grounded lexicon formation using inconsistent knowledge. In Capcarrere, M.S., Freitas, A.A., Bentley, P.J., Colins, C.G., Timmins, J., eds.: Proceedings of ECAL'05, Berlin, Springer (2005) 644–654
22. Vogt, P., Smith, A.D.M.: Quantifying lexicon acquisition under uncertainty. In Lenaerts, T., Nowe, A., Steenhout, K., eds.: Proceedings of Benelearn 2004. (2004)
23. Smith, A.D.M.: Mutual exclusivity: Communicative success despite conceptual divergence. In Tallerman, M., ed.: Language Origins: perspectives on evolution, Oxford, Oxford University Press (2005) 372–388