

Scaffolding Language Emergence Using the Autotelic Principle

Luc Steels^{1,2}

¹Sony Computer Science Laboratory
6 Rue Amyot
75005 Paris - France
Email: steels@arti.vub.ac.be

Pieter Wellens²

²VUB AI Laboratory
Pleinlaan 2
1050 Brussels - Belgium
Email: pieter@arti.vub.ac.be

Abstract—The paper focuses on the problem how a community of distributed agents may autonomously invent and coordinate lexicons and grammars. Although our earlier experiments have shown that a communication system can indeed emerge in a socio-cultural dynamics, it relies on the control of complexity by the experimenter, so that agents first acquire words, then simple constructions, and then more complex ones. This paper addresses the question how agents could themselves regulate the complexity both of the mechanisms they bring to bear to the language task and on the semantic complexity of what they want to express. We make use of the autotelic principle, coming from psychology. It requires monitoring challenge and skill (based on actual performance) and maintaining a 'flow' regime balancing the two. We show in computational experiments that the autotelic principle is capable to explain autonomous scaffolding towards greater complexity in the emergence of language.

I. INTRODUCTION

Research into the origins of grammar has been making substantial progress lately (see recent overviews in [10], [17], [4], a.o.). In our own work we have focused on a socio-cultural approach which relies on four important ideas [18]. First, it relies crucially on a peer-to-peer social dynamics within a community of agents, similar to models of opinion dynamics [8]. Each agent is capable to generate structure and then aligns this progressively with other agents based on local feedback in language games. Many simulations [17], as well as theoretical research [1] have now shown that this leads to self-organised coherence of lexicons, grammars, and ontologies without prior innate structure nor global control. Second the socio-cultural approach takes a functional rather than structural view of language, in line with cognitive linguistics [5]. Syntax is argued to be motivated by attempts to solve some aspect of the communication problem, for example, avoid uncertainty in who is playing what role in an event (as in "John gives Mary a book" versus "Mary gives John a book") or avoid combinatorial search in parsing. Third, the socio-cultural approach to language emergence sees an important role for intelligence both in the invention of new language forms that solve certain problems in communication and in the learning of new conventions invented by somebody else by guessing the communicative intentions and understanding the relation between structure and function. Learning is not simply a matter of imitation but of understanding at some level the

function of language structure in communication. And finally, the language faculty is not seen as a genetically fixed network of highly specialised modules but as a dynamically assembled collection of processes that have generic functions. The agent recruits new cognitive mechanisms as needed or expands the resources of mechanisms that make up his language faculty [19].

In socio-cultural models, it is crucial to model embodied situated communications because this is the only way to get a realistic measure of communicative success. Our own multi-agent experiments go as far as using real robots to increase the realism of the experiments and examine how constraints coming from embodiment may shape the emergent languages [13]. Whereas earlier work focused on understanding how a social peer-to-peer dynamics can lead to a shared language inventory, our current focus, as illustrated in the present paper, is part of understanding how exactly intelligence can or must intervene in the invention and intentional learning of grammar.

Intelligence has many facets, but it surely involves the ability to make models and solve problems by thinking through a solution, rather than trial and error in the world. We operationalise this in our experiments (including the one reported here) through re-entrance: In order to model the expected impact of their utterance on the hearer and thus find out whether there is ambiguity, risk of misinterpretation, combinatorial search, or other problems, speakers re-enter the sentence they produced and use themselves as models of the hearer [14]. Similarly hearers first parse and interpret the sentence as well as they can and solicit additional feedback if necessary, but then they use themselves as a model of the speaker in order to reconstruct missing or different rules in their lexicon and grammar through abduction [15].

A second facet of intelligence is that learning may take place through diagnosis and repair. Diagnosis means to detect what may be problematic, risky, or mistaken, and repair means to activate strategies or recruit tools that expand behavior to deal with the problem. For example, somebody acquiring the skill of mountain climbing progressively discovers that at some point he is going to need better shoes, warmer clothes, hooks, and then later on perhaps ropes or a sleeping bag. We will apply the same approach in bootstrapping languages. Agents have various diagnostics which they apply to their

language production and comprehension processes including the models they make of the other. For example, they may detect that there is a word in the input they do not know or that there is part of the meaning for which they do not have a word yet. Diagnostics are assembled and then repair strategies evoked to deal with them. A repair strategy typically expands or changes the language inventory but it may also involve the recruitment of new cognitive mechanisms. We will show concrete examples of this approach in the remainder of the paper.

A third facet of intelligence is the ability to scaffold the growth of complexity in acquiring a complex skill. In social learning, care givers play a very important role to simplify the world and give challenges to a child, but even young children are able to regulate the complexity of their interactions with the world, if only by ignoring many details that they cannot handle yet, and there is clear evidence that growth processes take place in the developing brain, including recruitment of connections between neural subsystems under performance pressure. All this appears necessary because many learning methods are only effective if their inputs are scaffolded and if they start out with limited resources that are gradually expanded [7]. The question how complexity growth can be regulated autonomously by agents in a distributed fashion is the key topic of the present paper. Before turning to that subject, we first introduce the kind of diagnosis and repair strategies that agents are able to recruit and what kind of dynamics results when they are applied (cf. [21]).

II. DIAGNOSIS AND REPAIR STRATEGIES

We assume that the meaning to be expressed by the speaker consists of a set of predicates and arguments. For example, “the big ball next to the red box” starts from: ‘big(obj1), ball(obj1), next-to(obj1,obj2), red(obj2), box(obj2)’. The hearer must reconstruct this meaning which will be an expression with variables: ‘big(?x), ball(?x), next-to(?x,?y), red(?y), box(?y)’ which is then matched against a world model derived from perception to find bindings for these variables and hence a viable interpretation, for example, ?x may map to obj1, ?y to obj2. The challenge of the agents is to invent a communication system that will allow them to make these mappings without a prior agreed upon lexicon nor grammar. Agents play description games. They describe to each other a scene which has a potential set of possible descriptions and the game succeeds if the meaning reconstructed by the hearer is compatible and unambiguous with the scene they are both looking at. For example, if the speaker says “ball” and the hearer indeed detects a ball in the scene, the game is a success. If there is more than one ball or the hearer does not know the word ball, the game would be a failure.

New lexical items appear in a language because parts of meaning could not be covered adequately with existing words. To implement this strategy, the speaker has a diagnostic ‘uncovered meaning’ which is repaired by inventing a new word and associating it with that meaning. This may either result in

a holistic coding or in compositional coding if already parts of the meaning could be covered with existing words. The hearer has a diagnostic ‘unknown word’ when an unknown word is found in the input. He can repair by grabbing the inferred uncovered meaning and associating it with this new word. We also assume that agents update the scores of the words in their lexicon following the lateral inhibition dynamics introduced in [12].

Let us see what effect this has when we supply agents with increasingly more complex meanings and expect them to handle it. More complex means that they have to express increasingly more objects in the scene (starting from 1), or that there is increasingly more sharing of variables between predicates. Results for experiments with the purely lexical strategy for 5 runs with 5 agents playing 4000 games are shown in figure 1. We see that quite quickly agents reach a high level of communicative success and the lexicon becomes optimal after about 1500 games. Note that at first the lexicon contains more words than necessary, simply because in a distributed population some agents may invent a new word independently of the other ones, but the lateral inhibition will efficiently dampen the incoherence. Since there are about 13 predicates in the example domain, an optimal lexicon is around 13 words.

However when the challenge is increased to include co-referential variables, communicative success starts to drop. This is simply because agents do not express what belongs to what. For example, if the hearer gets “ball left-of box” he derives ‘ball(?x), left-of(?y,?z), and box(?u)’ where ?x,?y,?z, and ?u are variables. He cannot derive from purely lexical means which object is left of the other one, except with a holistic coding. So when the scores of the existing lexical items is very low because they do not reach success, the lexicon of the agents begins to increase again as they use now holistic coding to cope with the more complex meanings. Indeed if we continue the experiment (not shown in figure 1) we will see that communicative success starts to climb back up, but only at the expense of a much larger lexicon which is slower to get off the ground, more difficult to learn, and more expensive in terms of memory resources.

A way to reach communicative success more effectively is to recruit another strategy: The speaker, while modeling the parsing and interpretation process of the hearer, could detect that there are variables which are co-referential. If that is the case, he could introduce a grammatical construction that signals that these two variables are to be considered co-referential thus avoiding potential ambiguities for the hearer (as explained in more detail in [18]). The hearer in turn could also have a diagnostic to detect co-referential variables and then abduct a similar construction. For example, suppose the meaning of “big” is ‘big(?x)’ and of “ball” ‘ball(?y)’. In “big ball” both variables refer to the same object so that ?x is co-referential with ?y. A grammatical construction maps a combination of predicates on the semantic side with a combination of words on the syntactic side and then establishes co-reference between variables. Grammatical constructions are not in terms

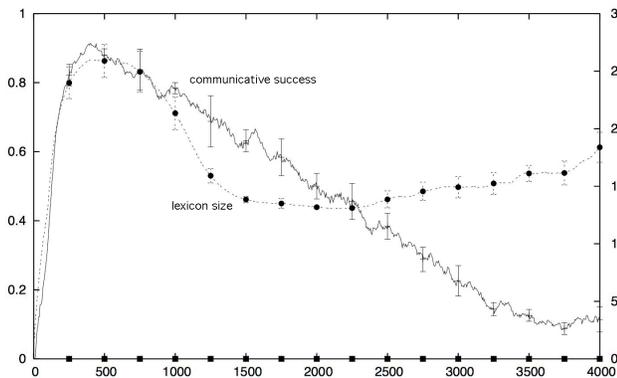


Fig. 1. Experiments where 5 agents use a purely lexical language to bootstrap a communication system. The x-axis shows the number of games. The y-axis shows communicative success and lexicon size. The strategy is initially successful but cannot cope when the complexity of meanings begins to increase.

of specific predicates and specific words but in terms of semantic and syntactic categories (like noun, adjective, etc.) so that the construction has wider applicability. Of course, in our experiments no categories are supplied to the agents because a theory of grammar emergence has to explain how such categories can ever arise in a population. Instead, they invent their own syntactic and semantic categories and impose them on the linguistic material. They then start to adjust and coordinate these categories the same way they coordinate perceptual categories through language [20].

Human language users clearly try to re-use as much as possible words, categories, or constructions in order to minimise the size of the language inventory, and thus avoid pressure on memory or ease transmission. This principle can be integrated in the repair strategy. Thus if an agent has already a syntactic pattern in a construction for “red ball” (e.g. ADJ+NOUN) and now needs to handle “green ball” or “red box”, the agent will reuse the existing pattern by categorising “green” as Adjective and “box” as Noun. So the use of a newly introduced category will progressively spread in the lexicon of the population, and we can measure this spread by counting the number of words that belong to a particular category. In case of syntactic ambiguity (a single word belonging to more than one syntactic category) the number of category associations will be larger than the number of words in the lexicon.

Figure 2 shows the impact of adding these diagnostic and repair strategies to the lexical ones discussed earlier. Again we, as experimenters, regulate the complexity of the input. The lexicon shows the same overshoot in the beginning and then stabilises around 13 words as competing words are resolved and lexical coherence reached. The necessary grammatical constructions are built early on. They are similar to the Adj-Noun construction discussed above, but without any significant syntax, just a mapping of semantic to syntactic categories. The graph also shows ‘grammaticality’, the running average of the number of utterances that make use of grammar. Overall, we see that the agents are able to cope with increasing

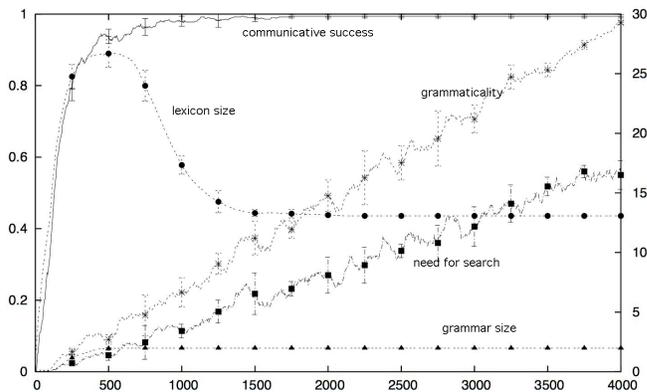


Fig. 2. Experiments where 5 agents use grammatical constructions in addition to a lexicon. They are now able to maintain communicative success even as complexity of meanings increases. The graph shows also the grammar size, as well as grammaticality (the number of sentences which required grammatical as opposed to only lexical rules) and the amount of search, which is steadily increasing.

complexity in their meaning space but it comes at a price. The amount of search required during parsing steadily increases because there are multiple ways in which constructions can be applied. If nothing is done, the combinatorial search problem becomes so large that parsing can no longer terminate within a reasonable time period. So there is again a ceiling of meaning complexity that the agents can handle efficiently.

Multiple hypotheses in parsing arise unavoidably as soon as the same syntactic pattern is re-used as part of a bigger structure and as soon as the same syntactic pattern is used with different levels of detail. For example, it is possible to build a noun phrase with just an article and a noun (“the box”) but also with an article, an adjective and a noun (“the big box”), or two noun phrases combined with a preposition (“a small box next to the orange ball”), and so on. Unless there is additional syntax, “a” or “the” in the latter example can both be combined with either “box” or “ball”, and “big” or “orange” can equally be combined with both nouns. Ignoring word order, the phrase can also be parsed as “(an orange ball) next to (the small box)”. Clearly languages introduce syntactic means to restrict the set of possible combinations which otherwise would quickly run out of hand. In English, this additional syntax is usually based on word order, but other languages may use other syntactic devices such as agreement between number and gender. For example, in the French sentence “une belle fille voit un beau garçon” (a beautiful girl sees a beautiful boy) feminine versus masculine gender of article, adjective and noun unmistakably specifies which words belong together in a single noun phrase and so the word order is strictly speaking not needed.

An earlier paper [21] reported an additional strategy in the form of a diagnostic with which agents can detect that combinatorial search is happening and a repair whereby they tighten the grammatical rules involved by adding more syntax. For example, the adjective is now forced to come before the noun and this then gives a clue that the two words form part of the same structure and their meanings have co-referential

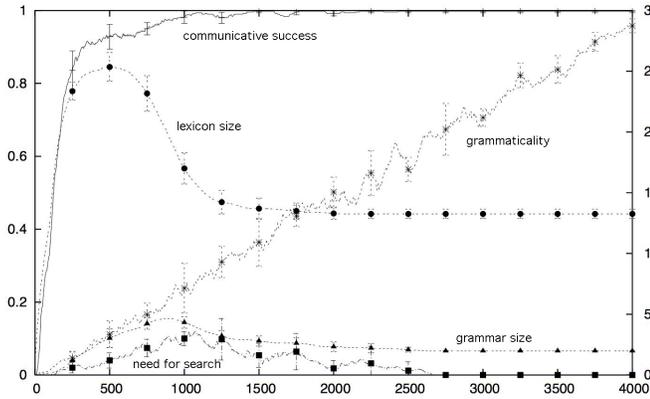


Fig. 3. Experiments where 5 agents now tighten grammatical constructions with additional syntax in order to avoid combinatorial search. We see a drastic reduction in the search needed.

variables. Of course, the specific syntactic constraint that is introduced by an agent is open and so agents have to co-ordinate their syntactic conventions through socio-cultural dynamics the same way they co-ordinate the lexicon, the ontologies, and the other rules of the grammar. Results in figure 3 show that this helps the agents to cope. There is a rapid climb of communicative success in the beginning and overshoot in lexicon size before it becomes optimal. We also see the emergence and coordination of grammatical constructions with the same characteristic curve: an overshoot (in the sense of more constructions are circulating in the population than strictly needed) because there are different ways to add syntax to a construction (e.g. Adj-Noun versus Noun-Adj order) followed by a stabilisation and optimisation phase due to lateral inhibition. The most important point, seen in the bottom graph, is that the search space is now completely under control and parsing has become deterministic.

This series of experiments is clearly an important breakthrough in the emergence of grammar, but it uses a fundamental assumption which is not tenable in the long run, namely we, the experimenters, regulate the challenge by increasing the semantic complexity, and we also regulate when additional cognitive mechanisms (in other words repair strategies) are to be integrated in the language faculty of the agents. In this paper we take away these assumptions by showing that agents can autonomously regulate the input/output complexity and the recruitment of strategies to build and learn lexicons and grammars. The key idea is to incorporate a general motivational system based on the autotelic principle.

III. THE AUTOTELIC PRINCIPLE

Several motivational theories have been studied in psychology and used as inspiration in artificial systems. Some are based on extrinsic motivation based on reward and punishment such as in re-enforcement learning [9]. In this case, an outside force could potentially regulate complexity as in the experiments discussed so far. Other theories are based on intrinsic motivation, such as curiosity-driven learning, which

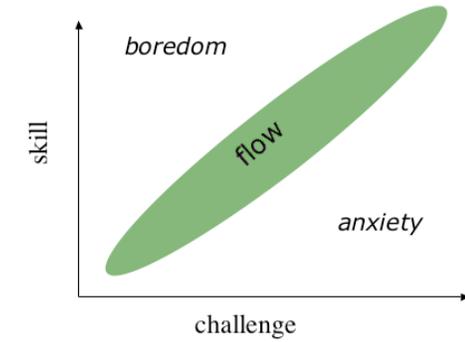


Fig. 4. Three emotional states related to motivation. The flow experience arises in the zone when there is an appropriate balance between challenge and skill. When the skill is consistently too high for the challenge, there is an experience of boredom. When the skill is consistently too low, there is an experience of anxiety. To maintain flow, the challenge needs to be increased to avoid boredom. Alternatively when the challenge is too high, learning possibly based on new cognitive mechanisms needs to be recruited to avoid anxiety.

is based on the desire to improve prediction [11]. Agents will first consider simple situations before they move on to more complex ones because they choose the situation in which they have the highest chance to increase their predictive success. Here we use another principle also based on intrinsic motivation, originally identified by the humanistic psychologist Csikszentmihalyi, and based on studying the activities of painters, rock climbers, surgeons, and other people who showed to be deeply involved in some very complex activity for the sake of doing it, i.e. without direct reward in the form of financial or status compensation [6]. He called these activities autotelic. "Autotelic" signifies that the motivational driving force ("telos") comes from the individual herself ("auto") instead of from an external source, administered by rewards and punishments. Autotelic activities induce a strong form of enjoyment which Csikszentmihalyi has characterised as "flow" and he has argued that this flow experience arises when there is a balance between challenge and skill.

According to Csikszentmihalyi, those who manage to anchor their motivation systems into the flow experience not only have very satisfactory lives, but they often reach unusual levels of excellence in their field. Conversely, when the challenge is too high for the available skill and when there is at the same time no hope to develop appropriate skills by learning, frustration and even anxiety sets in and a person may get paralysed and eventually develop symptoms of withdrawal and depression. When the challenge is too low for the available skill, boredom sets in and the long term reaction may be equally negative. The optimal regime is somewhere in between the two, when there is a match of challenge and skill. So to remain psychologically healthy, an individual needs to be able to decrease challenge when it is too high so as to get an opportunity to increase skills, but it is equally important that the individual can increase challenge when the skill has become higher than required to cope with the challenge. Moreover the environment should generate new opportunities for the individual to grow, so that if she wants to increase

challenge she can do so. Flow is not a steady state. When an activity is done a lot, skill normally increases so that it becomes boring and new challenges need to be found. Consequently an individual seeking the flow experience is always ‘on the move’.

The autotelic principle is not only relevant for understanding the motivational dynamics of a single individual, like a mountain climber trying to climb mountains of progressively greater complexity with increasingly more complex tools. It is also relevant for understanding the regulation of the complexity in interactions between two individuals, for example between tennis players. A player with low ability should play against one at the same level (if she wants to have fun anyway) and select more skillful opponents if her own skill has sufficiently grown to deal with a higher challenge. When this is done in a group, the result will be a progressive increase in the complexity that all players can handle, and we envision that something similar happens when a group of agents is bootstrapping an emergent communication system.

In an earlier paper, we already proposed an operationalisation of the autotelic principle for the domain of robotics [16] and we now apply it to the domain of language. Even though Csikszentmihalyi’s basic idea remains valid, we need to substantially expand and alter it to make it work from a computational point of view. First we need more precise operational notions for challenge, performance, confidence, and resources. Then we show the dynamics governing the whole system.

IV. OPERATIONALISING THE AUTOTELIC PRINCIPLE

The overall behavior of an agent is based on a chain of subsystems where output of one is input to the other. Thus in language production, the perceptual system produces segmentation and features for the conceptual system. The conceptual system then plans what to say based on the output of the perceptual system. The language system then turns this conceptualisation into a sentence and the speech system transforms a sentence into an acoustic signal. In language comprehension the same sort of subsystems are chained in the other direction. Each of these systems is in turn complex. For example the language system includes lexicon lookup, morphological analysis, semantic and syntactic categorisation, application of grammatical constructions, assembly of syntactic structure. We also assume that learning is totally integrated with execution. Every subsystem is not only capable to establish a particular mapping (for example decomposition of words into word-stems and morphemes) but also to learn the mapping.

To operationalise the autotelic principle, each subsystem of the agent is now given (i) a way to gauge the *challenge* of the output it produces for the next subsystem in the chain (for example, the perceptual system is able to compute a complexity measure for a particular scene, the conceptual model is able to compute the complexity of the conceptualisation it produces, etc.), and (ii) a way to set a particular *challenge level* of what it will accept as input (for example, conceptualisation

may accept only scenes of a certain complexity, the language system may accept only conceptualisations that are below a threshold of structural detail, parsing may try to deal with every word in the sentence or just focus on some of the words that the lexicon can recognise, etc.). Challenge comes usually in stages and they are cumulative in the sense that more difficult challenges rely on handling first simpler ones.

Next the agent is given the capacity to monitor his own *performance* (agents have never a total view, neither of the global communicative success nor of the inventories used by other agents). Each agent has an overall measure of performance based on success or failure in the language game as a whole and specific measures of performance tied to the subtask each subsystem must deal with. For example, the conceptual system has a performance measure how well it was able to come up with a discriminating conceptualisation, the language system has a measure whether it was able to come up with a sentence covering the conceptualisation, etc. The overall and local performance measures are integrated and tracked by each subsystem in each agent.

Each subsystem of the agent also tracks the *confidence* it has that it is able to reach the challenge level it posed itself. For example, suppose that the language subsystem has set for itself the challenge to deal with referring to single objects based on one-place predicates, as in “ball”. As the agent has a growing inventory of words to refer to more and more objects and is achieving increased success in doing so, his confidence with respect to that challenge increases, and eventually he believes that he can handle the challenge (even though this may mean that occasionally new words will have to be learned).

According to the recruitment theory, agents dynamically recruit cognitive mechanisms and link them into the language faculty if needed [19]. In this case, agents recruit diagnostic and repair strategies which they use to build up their inventories and learn those used by others. We call these the *resources* available to the agent. The resources increase whenever agents decide to add more strategies to their repertoire.

The dynamics of an autotelic agent can now be stated as follows:

- 1) If the confidence of an agent grows, meaning that he is able to reach a particular challenge level (based on a steady series of successful games that have increased that confidence), he decides to increase the challenge level. For example, the language system may raise the stakes from accepting to describe a single object to describing two objects.
- 2) If performance is falling, the agent will first go back to an earlier level of challenge that he knew how to handle. This is necessary to stabilise the system before trying something else.
- 3) If the agent goes into a limit cycle with decreased performance, decreased challenge, increased performance, increased challenge, decreased performance, decreased challenge, etc. this means that the current set of diagnostics and repair strategies is inadequate to deal with the higher challenge and hence the agent will

recruit new ones, in other words resources get increased. Recruitment amounts to a search process, because some mechanisms may be tried which do not have the desired effect at all.

Note that because performance depends also on the skill of other agents, the autotelic principle automatically handles the regulation of complexity growth in the population as a whole. Agents can only move to higher levels of challenge when others are able to cope as well. Even if they themselves reached a higher level of performance, they automatically lower challenge levels so that other agents have a chance to get bootstrapped into a new level of complexity. This kind of challenge level control (and thus control of input and/or output) is spontaneously done by humans in the case of ‘motherese’ spoken to children or ‘pidgin’ spoken to foreigners unfamiliar with the language.

V. EXPERIMENTAL RESULTS

Let us now illustrate this dynamics using experimental runs from our implementations. We use the same set of meanings as used in the experiments reported earlier. But now an experimenter no longer controls complexity of meaning nor what mechanisms are to be recruited when. Instead the agents do this themselves by decreasing and increasing their challenge levels and by recruiting at critical moments new diagnostic and repair strategies. The experiments are with a population of ten agents which can describe 7000 different situations.

The first experiment illustrates what happens when agents have no way to get out of a limit cycle that traps their challenge growth. Agents can only recruit the lexical strategy. Averaged results for the total population are shown in figure 5. It tracks the following quantities: *Communicative success* (between 0.0 and 1.0): This is the overall success in the language game for all agents. Failure is either due to the need to execute a repair strategy (for example invent a new word) or because the meaning could not be mapped unambiguously onto the scene. *Lexicon size*: The total number of different words in the lexicons of all agents. *Resources*: The number of repair strategies that are used. This is constant in this experiment. *Confidence*: This shows how close agents are to reach the challenge level that they have set themselves.

The initial challenge level is set by the agents to 2.0 (they talk at least about one object and avoid any uncertainty due to unexpressed co-referentiality of variables). We see that they progressively reach the confidence that they can handle this challenge level after about 4200 games. The lexicon grows to about 19 words and then starts to become more optimal due to lateral inhibition. Communicative success also rapidly increases from 0.0 and then stays at a maximum of 1.0. So agents feel confident enough to increase the challenge level. They now start talking about at most two objects avoiding interpretational ambiguity. The confidence of some agents is indeed beginning to increase but quite quickly, they run into problems with a 30 % drop in performance (around game 4500). Due to failures, the lexicon begins to disintegrate and

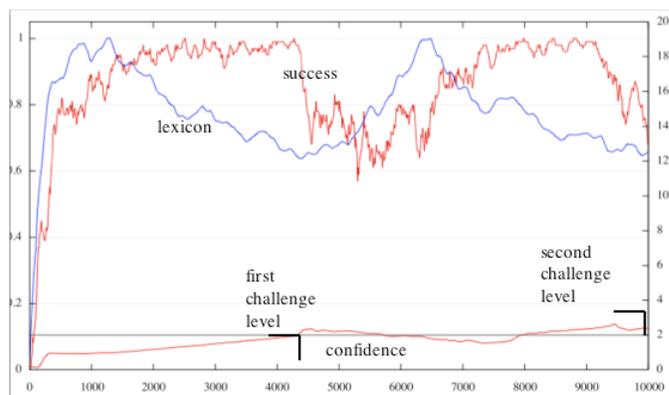


Fig. 5. A first experiment with autotelic agents. We see that their skill wavers around the first challenge level, showing a limit cycle behavior as they decrease and increase challenge.

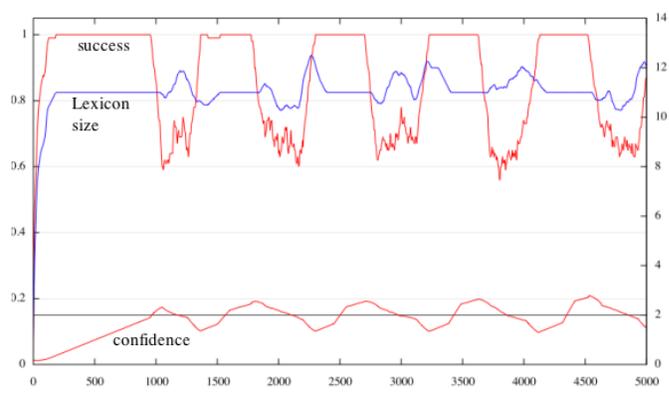


Fig. 6. Experiment with two agents using a lexical strategy to invent a language. As this strategy is not enough as complexity increases they keep falling back and are trapped in a limit cycle.

becomes less coherent (visible because the number of words increases again). Agents fall below their original confidence level and set the challenge level back to a lower level. As they now only handle simpler cases again, communicative success rises once more and the lexicon optimises. This leads to greater confidence, an increase in challenge level, and the same behavior repeats itself.

The limit cycle behavior is even more clearly visible when we only look at two agents (figure 6). We see that the confidence keeps going up and down. Agents reach success, increase the challenge (this is not shown) but then have to move back to an earlier challenge level so that they need to deal with simpler situations and hence can regain success. Each time it takes a while before their lexicon stabilises again. The diagram in figure 7 shows the same dynamics plotted as a trajectory in the performance/confidence space.

In the next experiment (figure 8), agents are able to escape the limit cycle by recruiting other repair strategies. The graph shows two additional quantities: *Grammar size*: The total number of constructions in the grammars of all agents. and *Category spread* (moving up to the size of the lexicon): The spreading of syntactic categories. Recruitment takes place after

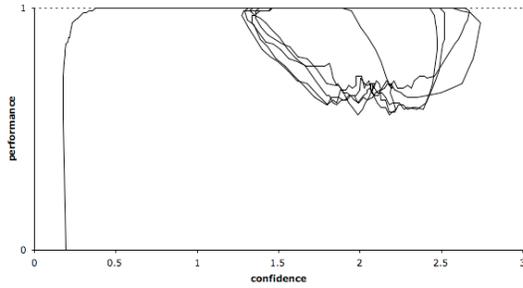


Fig. 7. This diagram plots the trajectory of the agent dynamics in the performance vs. confidence space.

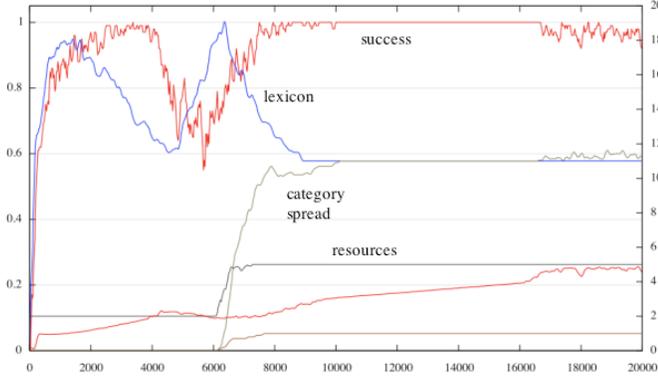


Fig. 8. A second experiment where autotelic agents have now grammatical repair strategies. Agents are able to escape the limit cycle and move confidently towards a higher challenge level by introducing grammatical constructions and syntactic categories.

about 6000 games. We see an increase in the resources used (to reach five repair strategies) and consequently a spreading in the population of the categories until all words in the lexicon have them. Note how the lexicon stabilises to an optimal size and how communicative success increases to become maximum. At game 16000, agents feel confident enough that they can increase their challenge levels once more. They now attempt to talk about at most three objects, still avoiding uncertainty due to unexpressed co-referentiality of variables. As no other repair strategies are recruitable we see again a limit cycle behavior, wavering around the second challenge level. The syntactic categories of the agents are beginning to destabilise, in the sense that ambiguities are arising (one word having more than one category), and communicative success is decreasing.

Figure 9 zooms in on the lower part of this graph which shows clearly the critical point when agents recruit the grammatical strategy and thus move to a higher plane of complexity.

VI. CONCLUSIONS

The paper brought together a number of threads that we explored separately in our earlier experiments: the use of a social dynamics in which agents invent language forms and coordinate them by adjusting weights using lateral inhibition, the use of diagnostics and repair strategies to organise the

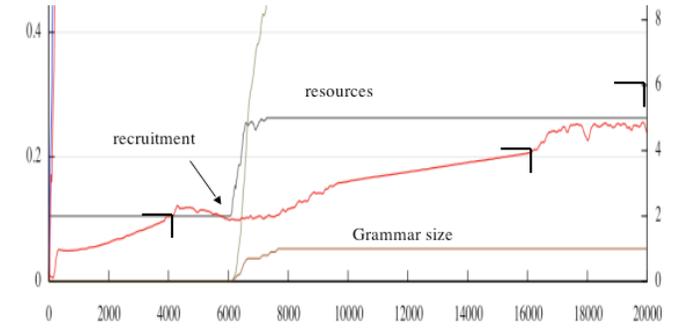


Fig. 9. This figure zooms in on the previous one. The different challenge levels are indicated with hooks. Agents can overcome the first two challenge levels but are still on their way to the third. Recruitment of new strategies leads to an increase in the resources, an increase in the size of the grammar (which was empty) and a spreading of the syntactic categories.

invention and adoption of progressively more complex language in a group of agents, the autotelic principle which regulates growth of developmental complexity, and the recruitment theory which argues that the language faculty is a dynamic configuration of cognitive mechanisms self-organised by agents in order to cope with challenges in the tasks and domains they encounter. We showed how these various principles interact to explain how a group of agents could progressively bootstrap a language system starting with the emergence of a lexicon for talking about single objects to the emergence of grammatical constructions with syntax. The experiments discussed here are only the beginning. We now have the agent software architecture to easily implement new diagnostics that focus on other functional pressures on language and new repair strategies that introduce lexical, grammatical or syntactic material to deal with them. As we further scale up, we expect to see greater difficulty in uneven skill levels between the agents. Indeed if one agent has already a fully developed grammar and now is confronted with an agent that is still learning the first words, the more mature agent needs to adjust challenge levels, but only in interaction with the absolute beginner. We suspect that this requires that agents can individually gauge the performance of other agents.

ACKNOWLEDGMENT

The research reported here has been conducted at the Artificial Intelligence Laboratory of the Vrije Universiteit Brussel (VUB) and at the Sony Computer Science Laboratory in Paris. Partial funding for the Sony CSL work has come from the EU FET-ECagents project 1940. We are very much indebted to Joachim De Beule for help in the technical foundations for the experiments reported here, Martin Loetzsch for helping to shape the Babel system underlying this implementation, and Remi Van Trijp for discussions on the linguistic foundations of this work.

REFERENCES

- [1] Baronchelli, A., M. Felici, E. Caglioti, V. Loreto and L. Steels (2006) Sharp Transition towards Shared Vocabularies in Multi-Agent Systems. *J. Stat. Mech.*, (P06014).

- [2] Boyd, R., and Richerson, P.J. (1985). Culture and the evolutionary process. University of Chicago Press, Chicago.
- [3] Cangelosi, A., C. Nehaniv, and C. Lyon (eds) (2006) Emergence of Communication and Language. Springer Verlag, Berlin.
- [4] Croft, W. and D.A. Cruse (2004) Cognitive Linguistics. Cambridge University Press, Cambridge.
- [5] Csikszentmihalyi, M. (1990) Flow. The Psychology of Optimal Experience. Harper and Row, New York.
- [6] Elman, J. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, v. 48. p. 71-89
- [7] Fortunato, S., V. Latora, A. Pluchino and A. Rapisarda (2005) Vector Opinion Dynamics in a Bounded Confidence Consensus Model *International Journal of Modern Physics C*, Vol. 16, No. 10 (2005) 1535-1551.
- [8] Kaelbling, L., M. Littman, A. Moore (1996) Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research* 4 (1996) pp. 237-285
- [9] Minett, J. W. and Wang, W. S-Y. (2005) Language Acquisition, Change and Emergence: Essays in Evolutionary Linguistics. City University of Hong Kong Press: Hong Kong.
- [10] Oudeyer, P-Y., Kaplan, F. and Hafner, V. (2007) Intrinsic Motivation Systems for Autonomous Mental Development. *IEEE Transactions on Evolutionary Computation*, 2007. to appear.
- [11] Steels, L. (1996d) A self-organizing spatial vocabulary. *Artificial Life Journal*, 2(3),1996.
- [12] Steels, L. (2003) Evolving grounded communication for robots. *Trends in Cognitive Science*, 7(7):308-312 July 2003.
- [13] Steels, L. (2003b) Language-reentrance and the inner Voice'. *Journal of Consciousness Studies*, 10(4-5):173-185 2003.
- [14] Steels, L. (2004) Constructivist Development of Grounded Construction Grammars Scott, D., Daelemans, W. and Walker M. (eds) (2004) *Proceedings Annual Meeting Association for Computational Linguistic Conference*. Barcelona. p. 9-19.
- [15] Steels, L. (2004b) The Autotelic Principle. In Fumiya, I., Pfeifer, R., Steels, L. and Kuniyoshi, K. (eds.) *Embodied Artificial Intelligence*, Lecture Notes in AI (vol. 3139), pages 231-242, Springer Verlag. Berlin, 2004.
- [16] Steels, L. (2005) The emergence and evolution of linguistic structure: from lexical to grammatical communication systems. *Connection Science*, 17(3-4):213-230.
- [17] Steels, L. (2005) What triggers the emergence of grammar? A. Cangelosi and C.L. Nehaniv (ed.) (2005) *Proceedings of the EELC satellite workshop at AISB 2005*. University of Hertfordshire, Hatfield, UK.
- [18] Steels, L. (2006) The Recruitment Theory of Language Origins. In: Lyon, C., C. Nehaniv, and A. Cangelosi (eds) *Emergence of Communication and Language*. Springer Verlag, Berlin. p. 129-151.
- [19] Steels, L. and T. Belpaeme (2005) Coordinating Perceptually Grounded Categories through Language. A Case Study for Colour. *Target article. Behavioral and Brain Sciences*. 28 (4).
- [20] Steels, L. and P. Wellens (2006) How grammar emerges to dampen combinatorial search in parsing. Vogt, P., Y. Suga, E. Tuci, C. Nehaniv (2006) (eds.) *Symbol Grounding and beyond*. LNAI 4211 Springer-Verlag, Berlin.