Lexicon Convergence in a Population With and Without Metacommunication

Zoran Macura and Jonathan Ginzburg

Department of Computer Science,
King's College, London
The Strand, London WC2R 2LS
{zoran.macura, jonathan.ginzburg}@kcl.ac.uk

Abstract. How does a shared lexicon arise in population of agents with differing lexicons, and how can this shared lexicon be maintained over multiple generations? In order to get some insight into these questions we present an ALife model in which the lexicon dynamics of populations that possess and lack metacommunicative interaction (MCI) capabilities are compared. We suggest that MCI serves as a key component in the maintenance of a linguistic interaction system. We ran a series of experiments on mono-generational and multi-generational populations whose initial state involved agents possessing distinct lexicons. These experiments reveal some clear differences in the lexicon dynamics of populations that acquire words solely by introspection contrasted with populations that learn using MCI or using a mixed strategy of introspection and MCI. Over a single generation the performance between the populations with and without MCI is comparable, in that the lexicon converges and is shared by the whole population. In multi-generational populations lexicon diverges at a faster rate for an introspective population, eventually consisting of one word being associated with every meaning, compared with MCI capable populations in which the lexicon is maintained, where every meaning is associated with a unique word.

1 Introduction

A key feature of natural language is metacommunicative interaction (MCI)—utterance acts in which conversationalists acknowledge understanding or request clarification. The need to verify that mutual understanding among interlocutors has been achieved with respect to any given utterance—and engage in discussion of a clarification request if this is not the case—is one of the central organising principles of conversation [1,2]. However, hitherto there has been little work on the emergence and significance of MCI meaning. Communicative interaction is fundamental to evolution of grammar work, since it is interactions among communicating agents that leads an initial 'agrammatical' system to evolve into a grammar (with possible, concomitant phylogenetic modification; see e.g. [3,4]).

P. Vogt et al. (Eds.): EELC 2006, LNAI 4211, pp. 100–112, 2006.

However, given an I-language[1] perspective, the communicative aspect as such is not internalised in the grammar (though see [6]). Consequently, such models of evolution of grammar cannot explain the existence of forms whose meaning is intrinsically MCI oriented.

What significance does MCI have for linguistic interaction within a community? Pretheoretically, MCI is redundant in so far as the communication channel, i.e. that which mediates between speaker and addressee, is perfect or close to that. The need for MCI not only arises when the communication channel is noisy, it also arises when there is ambiguity in the referents of the communicative interaction.

Moreover, languages are ever changing. Utterances used in slightly various contexts can rapidly lead the language itself to change so much as to become unrecognisable in only a few generations [7].

Given this, acknowledgements, clarification requests (CRs) and corrections are a key communicative component for a linguistic community. They serve as devices for allaying worries about miscommunication (acknowledgements) or for reducing mismatches about the linguistic system among agents (CRs and corrections). That is, they serve as a device for ensuring a certain state of equilibrium or lack of divergence gets maintained within a linguistic community. The plausibility of this speculation can be assessed by converting it into more concrete questions such as the following:

(1) Given a community A where clarification requests do not get expressed, and community B where they do, how do the two communities evolve with respect to vocabulary drift. How does this vocabulary drift change once a gradual turnover of community members is introduced?

In previous work, we have shown how language converges for different types of populations in a mono-generational model [8]. In this paper we modify the set up in two significant ways: (a) the lexicon is continually dynamic (in our previous set up once a word is acquired, its meaning does not change) (b) there is generational turnover. As will become evident, this changes the results in a dramatic and quite unexpected way.

In the next section we describe the computational model, including how gradual turnover of agents is implemented. In Sect. 3 we present the experiments and assess the validity of the proposed model. Finally, in Sect. 4, we conclude.

2 The Model

The model we propose here is an extension of the model described in [8], with the main extensions being the introduction of a dynamic lexicon, and the implementation of a gradual turnover of agents. In our previous work we have shown how

[1] Following Chomsky (as clarified by Hurford), a distinction is sometimes made between ‘I language’ — language as represented in the brains of the population and ‘E-language’ — language that exists as utterances in the arena of use. Ginzburg and Sag [5] dispute the dichotomy particularly given the need for a view of language that accommodates MCI.

language converges for different types of populations within a single generation. In this type of model as there is no generational turnover of agents the transmission of language is horizontal, where the communication is between adult agents of the same generation (e.g. [6]). In multi-generational models such as the iterated learning model (e.g. [9,10]) language is vertically transmitted from one generation to the next, where the adult agents are allowed to speak to the child agents only. So in these models there is no horizontal communication (i.e. between adults of the same generation).

We present a model which implements both horizontal (adult-adult) and vertical (adult-child) language transmission (see [11] for a similar approach). The model contains an ALife environment in which the lexicon dynamics of populations that possess and lack MCI capabilities are compared. The environment is modelled loosely after the Sugarscape environment [12], in that it is a spatial grid containing different plants. This environment is resemblant to the mushroom environment in [13]. Plants can be perceived and disambiguated by the agents. Unlike the environments in [12,13], plants are not used as a food resource but only as topics for conversations. Agents walk randomly in the environment and when proximate to one another engage in a brief conversational interaction concerning plants visible to them.[2]

In the next section we look at the communication protocol in more detail, followed by a closer look at the implementation of generational agent turnover.

2.1 Communication

Agents can talk about the plants in the environment by making syntactically simple utterances—essentially one consisting of a single word. Every agents has an internal lexicon which is represented by an association matrix (see [10,14] for similar approaches). The lexicon stores the association scores for every meaning–representation pair (i.e. plant–word) based on individual past experiences. Agents don't have an invention capability therefore are only able to talk about the plants that they have a representation for.

Communication is a two sided process involving an intrinsic asymmetry between speaker and addressee: when talking about a plant in his field of vision, the speaking agent necessarily has a lexical representation of the plant (a word with the highest association score for the plant chosen as the topic), which he sends to the hearing agent. There is no necessity, however, that the addressee agent is able to interpret this utterance. If unable to do so (meaning that the hearing agent doesn't have the word in her lexicon, or that the plant it associates with the word is not in her context) the way that the agent tries to ground it depends on the agent's type.

Three types of communicative agents exist in the model; agents capable of making a clarification request (CR agents), agents incapable of doing so (introspective agents), and hybrid agents that use both CRs and introspection.

[2] An agent's field of vision consists of a grid of fixed size originating from his location. Hence proximate agents have overlapping but not identical fields of vision.

An introspective agent learns the meanings of words through disambiguation across multiple contexts. Upon hearing a word the agent looks around her and for every plant in her context (field of vision) she increases its association score with the word heard. This strategy is akin to the cross-situational statistical learning strategy used by inferential agents in [10], and to selfish learners in [14].

A CR agent on the other hand can resort to a clarification request upon hearing a word. If hearing the word for the first time (no associations with the word in her lexicon) or if there are no plants in her context, a clarification request is raised. Otherwise the agent checks the plants in her context and if there is a mismatch between her internal state and the context (agent thinks that the word heard refers to a plant not in her context) she again resorts to raising a clarification request. The speaking agent answers this clarification request by pointing to the plant intended, after which the hearing agent increases the association score of the word heard with the pointed plant. However, if the perceived plant is in her context then the hearing agent only reinforces its association score with the word heard without resorting to a clarification request.

A hybrid agent has a capability of either using the CR strategy or the introspective strategy. The agent only resorts to a clarification request if she cannot ground the word heard (there are no plants in her context or there is a mismatch between her internal state and the context). When hearing an unknown word and having some plants in the context the agent follows the introspective strategy.

After updating her lexicon[3] the hearing agent chooses the plant with the highest association score for the word heard. If this perceived plant matches with the speakers intended plant then the conversational interaction is deemed as a success. Neither agent is given any feedback on the outcome of their conversational interaction (see [10] for a similar approach). Note that there is no lateral inhibition of all competing associations after a conversational interaction, as is the case for guessing game models such as [6,14]. Another significant difference, specifically between a guessing game strategy (e.g. [14]) and the CR strategy, is in the way feedback is provided. In a guessing game, agents verify whether the intended and perceived meanings match by evaluating 'corrective feedback' provided to them by the system. On the other hand, in our model, feedback is given only on the initiative of the hearing agent. In other words, a hearing agent is given feedback only when it explicitly asks for it (by raising a clarification when there is an uncertainty in the meaning of the word heard).

2.2 Generational Turnover

A typical approach when modelling a multi-generational population is the introduction of agent turnover. The iterated learning model [9] is an example of a multi-generational model where the language transmission is vertical (i.e. from one generation to the next). In such models the adult agents are always the speakers and child agents are always the hearers. The agents play a number of

[3] Only the hearing agents update their lexicons after a conversational interaction.

language games, which defines the length of a generation. At the end of a generation, the adults are removed from the model, the children become the new adults, and new children are introduced. This way of implementing generational turnover in the iterated learning model and other multi-generational models (e.g. [15]) is very rigid.

We propose a multi-generational model which is more realistic and resembles closer a human community (e.g. a tribe). In order to extend the mono-generational model described in [8] into a multi-generational model, there is a need to introduce a gradual agent turnover. This is done by introducing mortality. Every agent has a maximum age which is set randomly when the agent is born, and it lies in the range of $\pm 20\%$ from agent to agent. Upon reaching his maximum age the agent dies. Thus it is very unlikely that the whole adult population dies out at the same time, as the adult agents are of different ages and have different maximum ages.

In order to keep the population size stable, we also introduce natality. So for every agent that dies a new infant agent is born to a random adult agent in the model. The infant agent inherits the parent's type (introspective, CR or hybrid). Infants have an empty lexicon, with no knowledge of the meaning space or the word space. Each infant follows the parent around and is only able to listen to the parent's dialogues with other agents. In fact an infant only hears the dialogues in which her parent is the speaker. So the assumption here is that an infant learns only the words uttered by her parent. An infant cannot be a speaker and learns exclusively by introspection. The reason for this restriction is that infants start without any knowledge of language, and before they can actively engage in conversations they need to have at least some knowledge of the language. Every infant agent has an adulthood age which is set randomly and is about a sixth of the agent's lifespan. The adulthood age was experimentally determined and it gives the infant enough time to reach a good enough grasp of the language, enabling her to actively participate in conversations with other agents. When reaching the adulthood age an infant stops following her parent and becomes an adult, meaning that she is able to walk around independently, engage in conversations with other adult agents and become a parent. An infant can die only if her parent reaches the maximum age and dies.

This multi-agent model implements both vertical and horizontal language transmission as adult agents can communicate with each other as well as parent agents can communicate with their children. There is no clear distinction of when a generation starts and ends, like in the other multi-generational models, because there is continual agent turnover which makes calculating the results more intricate (see Sect. 3).

3 Experimental Results

This section describes different setups and experiment results for the model described in Sect. 2. In order to test the questions raised in (1) we ran several experiments in which agents posses distinct lexicons, and clarification requesting (CR) and introspective capabilities.

Before creating a population of agents, the environment is created containing 20 different plants (which represent 20 different meanings). There are six instances of every plant and they are randomly distributed in the environment.

The population in the simulations described here is made up of 40 agents that are also randomly distributed in the environment at the start. 20% of the initial population is made up of infants (i.e. 8 infant agents). Agents form two different communities each of whose members initially share a common lexicon. The initial community lexicons are distinct from each other (in that no meaning has the same representation associated with it). Agents can be either of the same or different type within the community. Apart from the differences in the initial lexicons and types between the agents, all other properties are the same.

Once the simulation starts the agents begin walking randomly in the environment. At every tick (time step) an agent's age increase and the agent walks one step in a random direction. After moving an agent looks for other agents (that fall into his field of vision). If he sees another agent then two of them enter a dialogue where the 'see-er' is the speaker and the 'seen' is the addressee. After a dialogue the agent continues walking in a random direction. In one tick every agent goes through this process. When an agent reaches his maximum age he dies and a new infant is born.

The performance of the model is based upon these behaviours which are collected at regular intervals in a simulation run:

- *Lexical Accuracy*: the population average of correctly acquired words. A word is said to be correctly acquired if it is associated with the same meaning as in either of the two initial lexicons.
- *Meaning Coverage*: the average number of meanings expressible by the overall population. There is no requirement that the meanings have correct associations.
- *Word Coverage*: the average number of words expressible by the population (correctness not taken into account).
- *Communicative Success*: the percentage of successfully completed conversations. A successful conversation is when the intended meaning by the speaker matches the perceived meaning by the hearer.
- *Method of Acquisition*: the percentage of conversational interactions that follow the introspective strategy or the CR strategy.
- *Distinct Lexicons*: the total number of distinct lexicons in the population. A lexicon is distinct only if there is no other lexicon in the population with which it shares all plant-word associations, so even if two or more lexicons have 19 out of 20 same plant-word associations they are regarded as distinct.
- *Lexical Convergence*: the percentage of agents sharing a lexicon. Agents share a lexicon if and only if all the plant-word associations are the same in their respective lexicons. Lexical convergence of 1 implies that all the agents use the same words for every plant in their lexicons.

We ran four types of experiments with different population make-ups, namely introspective populations, CR populations, hybrid populations and mixed populations (made up of both introspective and CR agents in a 1:1 ratio). For all different experiments, 10 trial runs were carried out for statistical analysis.

Firstly, mono-generational experiments were carried out (Sect. 3.1) in order to see how lexicon changes within a single generation with infant agents and no mortality. Then multi-generational experiments (Sect. 3.2) were carried out to view the lexicon change on a longer timescale with a gradual turnover of agents.

3.1 Mono-generational Experiments

Mono-generational experiments were ran to see how the introduction of infant agents into the model affects the performance of different populations based on the behaviours described above. In these experiments the population was made up of 40 agents in total, 20% of which were infants. There was no mortality and the experiments were stopped after 100,000 ticks. Results were collected at every 1,000 ticks.

As can be seen in Fig. 1(a) the lexical accuracy for every population raises rapidly and reaches nearly 100% by 20,000 ticks. What this means is that out of 40 possible words[4] the populations are correctly acquiring around 99% of them. Communicative success (Fig. 1(b)) also rises similarly to the lexical accuracy. As agents acquire and strengthen their plant-word associations their lexicons become more similar and their communications more successful. The introspective population is slower than the others as learning less frequent words takes more time, but performancewise doesn't differ much from the CR or hybrid populations.

The meaning and word coverage for each population reaches 100% and there is no difference in the time it takes between the different populations. The graphs are very similar to the graphs in Fig. 1, thus not shown here.

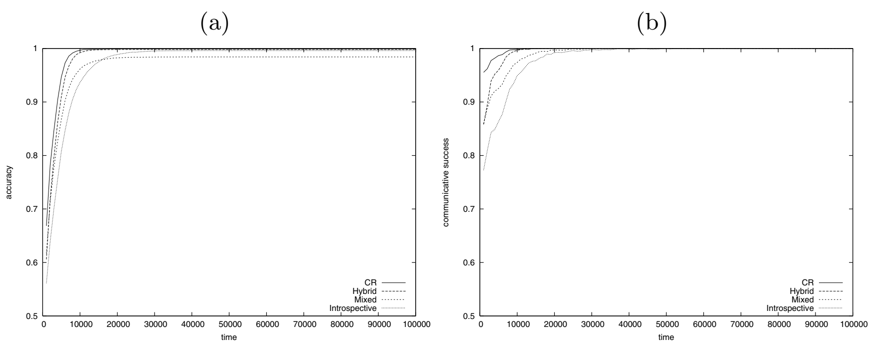

Fig. 1. (a) Lexical accuracy against time for mono-generational populations. There is a sharp initial increase in accuracy as new words are acquired correctly. (b) Communicative success also increases sharply and eventually reaches 100%.

[4] There are two initial communities, each with a distinct lexicon—for every plant the two communities are using a distinct word. As there are 20 different plants in the environment, the total number of distinct words in the population is 40.

The percentage of conversational interactions where introspective or clarification strategy have been used can be seen in Fig. 2(a). In an introspective population all the interactions follow the introspective strategy. In CR and hybrid populations clarifications are raised 45% of the time while 55% of the time they use introspection. In a mixed population the level of clarifications drops down to around 25% as half of the population is made up of introspective agents.

Figure 2(b) shows the number of distinct lexicons in the population. There is a sharp increase initially in the number of distinct lexicons. At the beginning of the simulation there are two distinct lexicons. As the agents speak they acquire novel plant-word associations so their lexicons diverge and the number of lexicons increases. Between 10,000 and 20,000 ticks there is a peak of 38 distinct lexicons indicating that only two agents in the population share a lexicon while everyone else has a distinct lexicon. As the time increases, agents have more conversations and the plant-word associations in their lexicons are strengthened, thus more and more agents use the same word for a given plant. This increases the lexical similarities between agents so the number of distinct lexicons starts to decrease. Eventually one lexicon becomes predominant in the population, where every agent uses the same word for a given plant. The CR and hybrid populations are fastest in converging to a shared lexicon, while it takes considerably longer for an introspective population to achieve this. A mixed population is a bit slower than the CR and hybrid, but the peak of distinct lexicons is smaller than in other populations (i.e. 35 distinct lexicons).

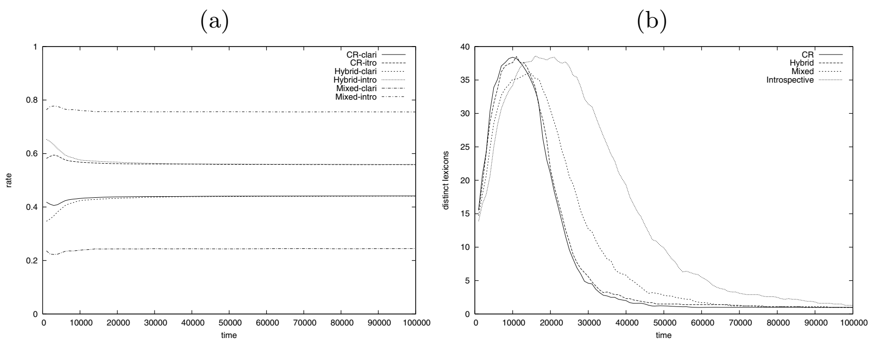

Fig. 2. (a) Percentage of conversational interactions that follow either the introspective or the CR strategy. (b) Number of distinct lexicons in the populations raise sharply reaching a peak of nearly one distinct lexicon per agent. As more dialogues occur within the population so does the number of lexicons drop, eventually stabilising at a single lexicon that is shared by the whole population.

The lexical convergence of different populations is shown in more detail by Fig. 3. The general trend is similar for different populations, where at the beginning there are many distinct lexicons shared by few agents (represented by peaks on the right side of the graphs). As time increases more and more agents start sharing a lexicon (represented by smaller peaks going from right to left), up to the point where every agent shares a single common lexicon (represented by peaks on the left side). What can be seen from Fig. 3(d) is that in a CR population there are considerably fewer peaks in the middle of the graph. This means that there are fewer competing lexicons, and that the convergence is faster than in other populations. Introspective population shown by Fig. 3(a) is on the other hand much slower in reaching a shared lexicon. The convergence of hybrid (Fig. 3(c)) and mixed (Fig. 3(d)) populations are similar.

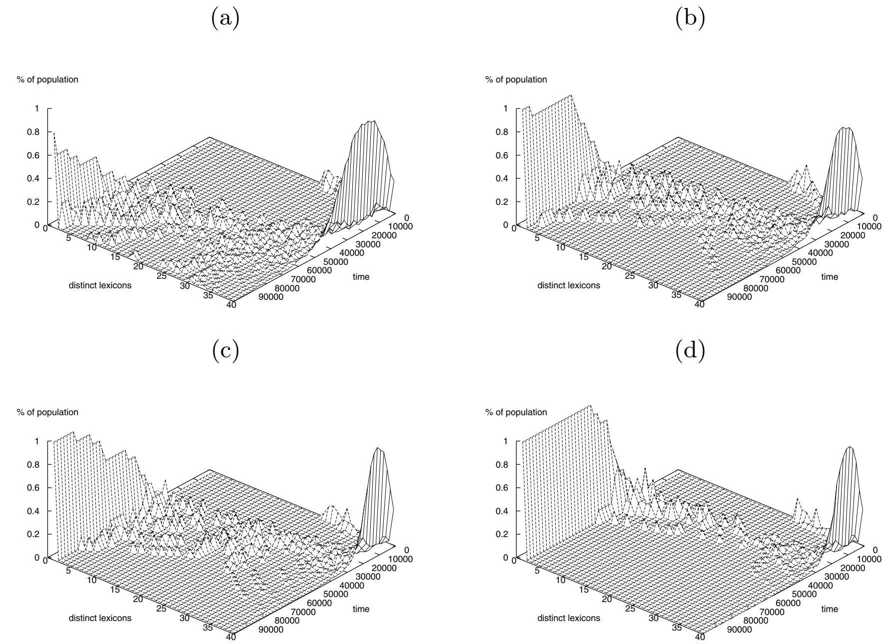

Fig. 3. Lexical convergence for different populations. The y-axis shows the number of distinct lexicons in a population while the z-axis indicates the percentage of agents sharing a distinct lexicon. Average results are shown for (a) Introspective population (b) Mixed population (c) Hybrid population and (d) CR population.

3.2 Multi-generational Experiments

The population in these experiments is kept constant to around 40 agents at any moment in time and the ratio of adults to infants is roughly 3:1. The agent life span is limited to around 30,000 ticks (±20%). This should reduce convergence and raise issues of generational variation. Results were taken at every 20,000 ticks. The simulation is stopped when it reaches 2 million ticks, which means after around 70 generations.

There are some clear differences between the mono-generational results and the multi-generational ones. The lexical accuracy initially drops very sharply for every population (Fig. 4(a)). At the beginning of the simulation there are a total of 40 words in the population (20 words from each community). As the words compete with one another there is a point when one word becomes dominant for a given plant and the majority of agents start using it. Thus the other competing words for the same meaning are used less frequently. The fact that the infant agents only learn the words uttered by their parents makes it very unlikely that the infrequently uttered words will pass to the next generation. After around three generations (100,000 ticks) the lexicon stabilises for every population except for the introspective.

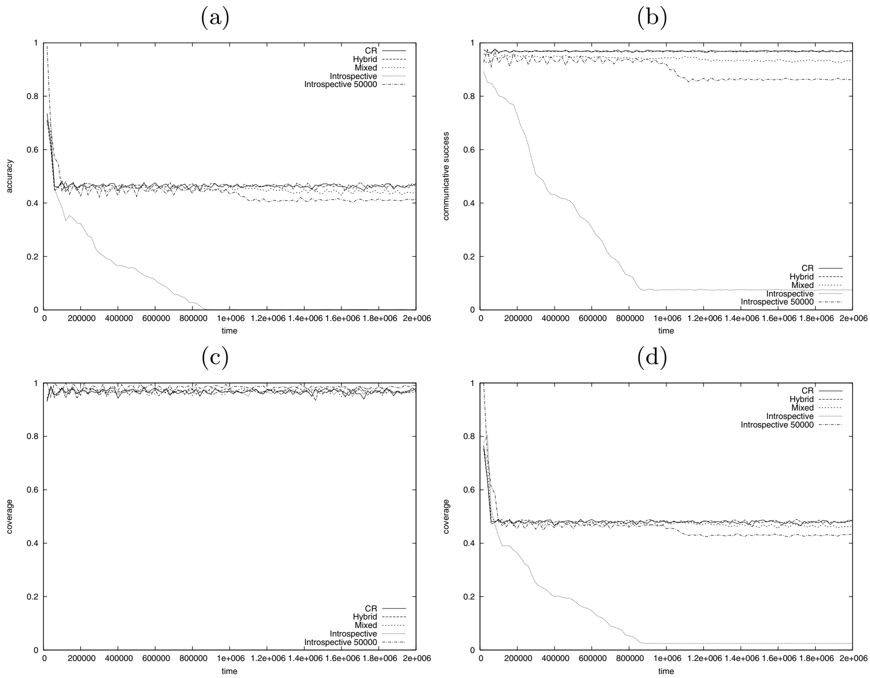

Fig. 4. Results for multi-generational populations showing (a) Lexical accuracy, (b) Communicative success, (c) Meaning coverage (d) Word coverage

The reason for this stabilisation can be explained by looking at Fig. 4(c) and Fig. 4(d). Fig. 4(c) shows that the meaning coverage for different populations is stable (all of them are able to express nearly every meaning). The word coverage however drops rapidly along with the lexical accuracy, as seen in Fig. 4(d). This is an indication that only the dominant words are surviving. Once the word coverage drops to 50% the lexicon stabilises. Around 20 different plants (Fig. 4(c)) and 20 different words (Fig. 4(d)) are expressible by the population at this stage, so every plant is associated with one word. These words can be successfully passed on to the next generation as they are used with greater frequency, causing the lexicon to stabilise.

This is not the case for the introspective population. The lexicon keeps diverging very rapidly and eventually reaches nearly 0% convergence (meaning that very few words have the association with the same plants as in the initial lexicon). Looking again at Fig. 4(c) and Fig. 4(d) explains why this happens. The word coverage also drops very sharply, where in the end only one word is known by the whole population. As the meaning coverage is comparable with other populations it can be derived that every plant in the population is associated with the single word expressible by the population, causing the lexical accuracy to decrease.

This in turn affects the communicative success of the introspective population (Fig. 4(b)). As for the other populations the communicative success is constant throughout the simulation, with the CR and hybrid populations doing slightly better than the mixed population. Thus even though the lexicon is diverging at a fast rate initially, the agents are still able to communicate successfully about different plants.

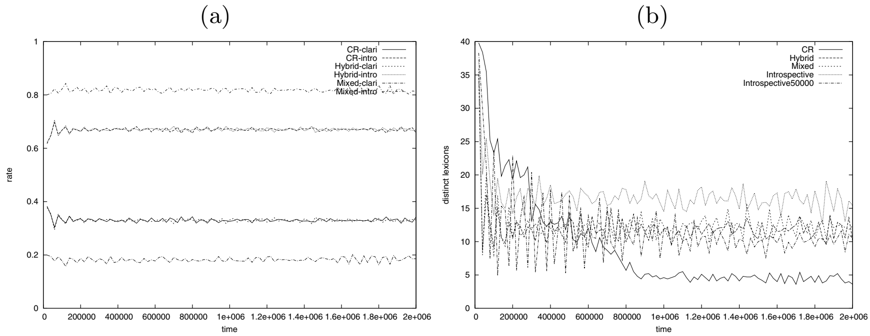

Fig. 5. (a) Percentage of conversational interactions that follow either the introspective or the CR strategy. (b) Number of distinct lexicons in the populations drop to around 10 lexicons then oscillate around that value.

The percentage of conversational interactions where introspective or CR strategy has been employed in shown by Fig. 5(a). The results are similar to mono-generational results with the clarification frequencies for all populations (except introspective populations) being slightly lower. The ascending order of CR frequency is: introspective 0%, mixed 20%, hybrid and CR 32%. It can be seen in Fig. 4 that the populations in which CRs can be expressed (CR, hybrid and mixed) perform much better than the ones in which CRs can't be expressed (introspective populations).

None of the populations converge to a single common shared lexicon as was the case in the mono-generational model (Fig. 5(b)). The reason is that infants make up around 20% of the population. As infant agents tend to have incomplete lexicons which differ form other agents, the number of distinct lexicons is higher than in mono-generational experiments.

Figure 6(b) shows a high degree of convergence to a common lexicon on the adult part in a CR population. The infant lexicons are represented by peaks on the right side of the graph and are used by about 20% of the population. The majority of the population shares a common lexicon represented by the peaks on the left side of the graph (between 20 and 35 agents). As for the introspective population (Fig. 6(a)) it looks as the population has converged to a common lexicon. This is true, but as we have shown one word is used for representing every plant so the majority of agents converge to the same lexicon containing only this single word.

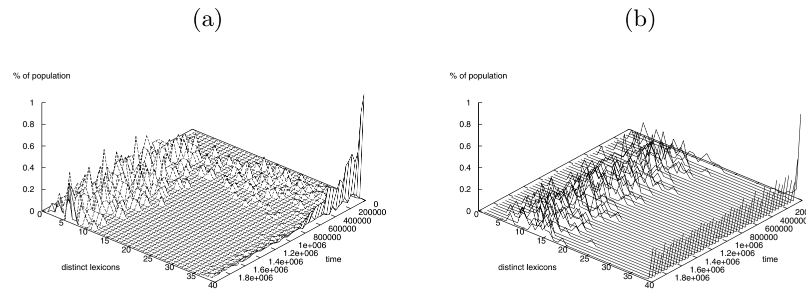

Fig. 6. Lexical convergence for (a) Introspective population and (b) CR population. Results for hybrid and mixed populations are similar to those of CR populations thus not shown here.

4 Conclusions and Future Work

In this paper we have discussed how metacommunicative interaction (MCI) serves as a key component in the maintenance of a linguistic interaction system. We ran a series of experiments on mono-generational and multi-generational populations in which lexicon dynamics of the populations that posses and lack MCI capabilities were compared. We have shown that in a mono-generational model all the populations converge to a common lexicon, although the introspective population was the slowest to achieve this.

Limiting life span of agents in the multi-generational model raised some clear differences in the lexicon dynamics between the MCI capable and incapable populations. The main effect demonstrated is that in the introspective populations the lexicon diverges continually, ending up with a situation where every agent in the population uses the same word to represent every plant in the environment. On the other hand MCI capable populations are able to maintain the lexicon, and the adult agents converge to a common lexicon.

While this confirms our initial theorising, much work remains to butress it as a fundamental dividing line between MCI-ful and MCI-less populations. In our current experiments we are seeing that increasing the maximum age of agents in introspective populations to 50,000 improves the lexicon stability and convergence (see *Introspective 50000* results in Fig. 4). Another issue concerns the influence of topography (e.g. variety of plants in the environment), as increasing the variety affects the performance of all populations. Further work needs to be done in order to get more insight into both of these issues. A more far reaching goal is to see whether using a syntactically complex language where the meaning space is potentially unbounded changes the results.

Acknowledgements

This work is supported by an EPSRC quota award to Macura. We would like to thank three reviewers for EELC06 for their very helpful comments.

References

1. Schegloff, E.A.: Repair after next turn: The last structurally provided defense of intersubjectivity in conversation. American Journal of Sociology **97**(5) (1992) 1295–1345
2. Clark, H.H.: Using Language. Cambridge University Press, Cambridge (1996)
3. Briscoe, E.J.: Grammatical acquisition: Inductive bias and coevolution of language and the language acquisition device. Language **76**(2) (2000) 245–296
4. Kirby, S.: Syntax without natural selection: how compositionality emerges from vocabulary in a population of learners. In Knight, C., Studdert-Kennedy, M., Hurford, J.R., eds.: The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form. Cambridge University Press, Cambridge (2000) 303–323
5. Ginzburg, J., Sag, I.A.: Interrogative Investigations: the form, meaning and use of English Interrogatives. Number 123 in CSLI Lecture Notes. CSLI Publications, Stanford: California (2000)
6. Steels, L.: The origins of ontologies and communication conventions in multi-agent systems. Autonomous Agents and Multi-Agent Systems **1**(2) (1998) 169–194
7. Deutscher, G.: The Unfolding of Language: an evolutionary tour of mankind's greatest invention. Metropolitan Books, New York (2005)
8. Ginzburg, J., Macura, Z.: Lexical acquisition with and without metacommunication. In Lyon, C., Nehaniv, C.L., Cangelosi, A., eds.: Emergence of Communication and Language. Springer Verlag, London (in press)
9. Kirby, S., Smith, K., Brighton, H.: From ug to universals: linguistic adaptation through iterated learning. Studies in Language **28**(3) (2004) 587–607
10. Smith, A.D.M.: The inferential transmission of language. Adaptive Behavior **13**(4) (2005) 311–324
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. Epstein, J.M., Axtell, R.: Growing Artificial Societies: Social science from the bottom up. MIT Press (1996)
13. Acerbi, A., Parisi, D.: Cultural transmission between and within generations. Journal of Artificial Societies and Social Simulation **9**(1) (2006)
14. Vogt, P., Coumans, H.: Investigating social interaction strategies for bootstrapping lexicon development. Journal of Artificial Societies and Social Simulation **6**(1) (2003)
15. Vogt, P.: Stability conditions in the evolution of compositional languages: issues in scaling population sizes. In: Proceedings of the European Conference on Complex Systems. (in press)