Combining Different Interaction Strategies Reduces Uncertainty When Bootstrapping a Lexicon

Thomas Cederborg

VUB AI Lab - Brussels
thomas@arti.vub.ac.be

Abstract. When bootstrapping a new language, the agents in a population need to be able to agree on the meaning of the individual words. In order to do so, they need to overcome the problem of referential uncertainty, which captures the idea that the meaning of words can not realisticly be transferred directly between agents nor through the environment. One way to reduce the amount of uncertainty, is to allow the agents, based on their current knowledge of the language system and the environment, to choose the interaction script they play based on a motivational system. We show the impact of this idea through a computational model on the time needed for a population of agents to converge on a shared language system and how the motivational system allows the agents to self-regulate this process.

1 Introduction

In recent years, there has been an increasing amount of computational and mathematical models within the language game framework that investigate various aspects of the origins and evolution of language, see [1] for an introduction. In the research at the VUB AI-Lab and Sony CSL Paris we frame our models in a larger language game paradigm. In short this means all experiments consist of multiple agents that engage in pairwise interactions (language games) in a shared environment. By playing these language games the agents gradually bootstrap and maintain a communication system in order to solve a communicative task. It is important to note that there is no central control and the agents can never access the internal states of others. Different issues have already been tackled by other researchers in this field, including a detailed study of shared lexicons [2], spatial language [3], color [6], hierarchy and recursion [4] and case grammar [5].

In this research we investigate the impact of using different language game scripts for solving the problem of referential uncertainty (also known as the Gavagai problem in linguistics [14]). As an example of this problem imagine one agent describing an object, let's say a bowling ball, with the word "gavagai". Another agent, even knowing the word was used to refer to a bowling ball, still cannot with complete certainty know what part of the object was described, it could still mean; black, round, large, heavy, shiny, hollow, etc. This problem has received quite some attention in the last few years, for example [7] [8]. The key in

G. Kampis, I. Karsai, and E. Szathmáry (Eds.): ECAL 2009, Part I, LNCS 5777, pp. 458–465, 2011.

the research presented here, is that agents autonomously choose which language game to play based on the environment and their estimates of their own and their interaction partners linguistic capabilities.

2 Experimental Setup

In the experiments we used a population of ten agents from which, at the beginning of a language game, two are randomly drawn. The context (or scene) consists of ten randomly selected objects, each composed of different features, such as shape and color. The agents perceive the scene and both build the exact same worldmodel. This simplification allows us to study several communicative problems separately from the problems arising from having different world views. The communicative task the agents need to solve is to bootstrap and align a set of form to meaning associations. These meanings however, can not only refer to the instantiated features (for example “red”, “blue”, “round”, we shall call these descriptive words) but also to the more general dimensions (for example, “color” or “shape”). This second type of words the agents will use to ask questions, which is why we call them question words.

One of the two agents starts the game (the initiator) by first deciding what game to play. This choice is based on the scene and the agent’s own lexicon, including the scores of his lexical constructions. How these scores are updated is explained in more detail below but in general words that are used in successful games will be rewarded, thus this score hints at the chance of being understood in future interactions.

The initiator can choose from among three different types of games, each having its own functional goal. The first is what we call a Guessing Game [9] [7] [8], specifically with regard to the problem of referential uncertainty. The goal of the Guessing Game is to learn the descriptive words. The interaction pattern starts with the initiator refering to one object by uttering the highest scored word for one of the objects features. The respondent either points to an object if he believes it fits the word or does nothing if he doesn’t know the word or cannot map it to any object in the context. If the right object is pointed to, the initiator nods and otherwise shakes his head and points to the correct object.

Several previous experiments have examined what happens if this is the only interaction pattern available to the agents. Many different methods of analysis have been tried, both involving single word and multiple word versions [9]. We illustrate the principle of the game with an example in which the purpose of the game is to negotiate a word for red and the communicative goal is to get the other agent to point to a specific red object.

1) initiator: “red”.
2) respondent: silent (does not understand the word “red”).
3) initiator: points to a small triangular red object.
4) respondent: makes three hypotheses about what the word “red” means: small, triangular or red.

The other two interaction patterns both allow questions to be asked. The goal of the first Question Game (type 1) is to learn question words (i.e. words for dimensions). In this game the initiator points to an object and utters a question word (such as "color"). If the respondent is able to interpret this word, he will try to describe the object using a corresponding descriptive word, otherwise he remains silent. The initiator nods when the description is satisfactory and otherwise he shakes his head and gives the description he was hoping for. In the following example the initiator has a word for red with a score that is high enough for him to assume that the respondent knows it. The purpose of the game is to negotiate a word for the color dimension, and the communicative goal is to get the object described by the word "red".

1) initiator: points to a big round red object.
2) initiator: "color"? (assuming the word "red" is known by the other agent).
3) respondent: silent (does not understand the word "color").
4) initiator: "red".
5) respondent: learns how to describe an object when hearing "color".

In question game type 2 the initiator uses a question word he thinks the respondent knows to ask for the name of a feature. In the following game the initiator has a word for color with a high enough score to assume that the respondent will know it. The purpose of the game is to negotiate a word for yellow, and the communicative goal is to get the object described by the word the respondent has for yellow.

1) initiator: points to a big triangular yellow object.
2) initiator: "color"? (assuming the word "color" is known by the other agent).
3) respondent: "yellow".
4) initiator: nods (does not have a word for it).
5) initiator: Learns the word "yellow" (knowing it is not referring to big or triangular).

The question games will obviously result in incorrect learning if the assumptions made by the agents are incorrect. If the respondent in the second example is sure that "red" means round, he will also conclude that "color" must mean shape. The initiator of the third example will also make incorrect inferences if his assumption about the respondent's knowledge of "color" is wrong. The agents will have to hear a descriptive word being used by other agents (describing objects that fit with what the agent thinks the word mean) before feeling confident enough to use it in a question game. With every word there is associated a score with a value between 0 and 1 of the same type as described in[9].

The motivation of the agents is to increase the expected communicative success in situations where they are not allowed to choose the topic of discussion. This can be seen as them having some reason, other than language learning, to describe an object. The different types of games they can play receive their utility from this goal. The communicative goals of all the games they play (for example the communicative goal: getting an other agent to describe a red object using color) should be seen as subgoals of this. When deciding what game to play and what the topic should be the agents try to find an appropriate challenge level by simultaneously trying to avoid anxiety and boredom. This approach draws inspiration from psychology and the concept of flow [12], a mental state where the challenge (and therefore the experience) of some task is optimal. A typical example is that of a mountain climber who has managed to find a mountain that is very difficult and challenging (and thus not boring) but that does not contain any obstacles that are so difficult that he does not know how to start dealing with them (something that would have produced anxiety) and has been explored before in the context of language learning, in which it was called the autotelic principle [10].

The utility U of playing a specific game about a specific feature is assigned by a utility function and the game with the highest utility will be played. We will call the score of the best question construction for the type of the feature Sc and the score of the best descriptive construction Sd.

For the guessing game the agents have no help other than pointing (and thus have no good way of communicating what they were actually describing). They are thus always anxious in this game and always prefer to talk about the feature they know the best as this is their only way of reducing anxiety within the guessing game framework. Thus the utility of playing a guessing game about a given feature is simply the score of the highest scoring construction with the same value as that feature. We get $U = Sd$.

For question game type 1 the uncertainty and thus the anxiety comes from 2 sources, the uncertainty of the question construction used and the uncertainty about the descriptive construction used. We get:

$$U = ((Sc^3) * 32) * (1/(1 - Sd)). \tag{1}$$

Here the term $(1/(1 - Sd))$ can be seen as confidence in the descriptive construction used to learn the question construction. If no question construction is available the utility function is set to that corresponding to a score of 0.5.

For question game type 2 there are 2 competing factors, the will to learn a descriptive construction that should increase with a low score for this construction and a fear of failure as a result of the question construction. We get:

$$U = (2 - Sd) * (1/(1 - Sc)). \tag{2}$$

Here the first term is from wanting to learn descriptive words not known (avoiding boredom) and the second term anxiously avoiding to use bad question constructions to learn them. When using color to learn green a high score of color (to reduce risk of failing due to the other agent not understanding the question) is preferred and a low score of green is preferred (wanting to learn what he does not already know well). This is a good example of how the will to avoid anxiety can conflict with the will do avoid boredom. If a game is found that manages to avoid both of these (finding a question word that is sure to be understood and is useful for learning the name of a feature that is completely unknown) we can see this as the agents having achieved flow in the sense of [12]. U is set to 0 if boredom is very high (the word that can be learnt is very well known) or if anxiety is to high, for example asking a question that you do not know the answer to using a question word that other agent almost certainly will not understand.

When an agent has completed a language game he will analyze what has happened and can update his language in three ways; he can adopt a construction (when a new word is heard), he can change the score of a construction (based on how likely it is that other agents share this construction) and finally he can delete a construction (when the score is below a threshold of 0.5) [1]. When a descriptive word is used about an object the construction with its form is increased if it's meaning matches the object described and decreased if it does not match. If the score of one descriptive word is increased the score of constructions that are it's homonyms or synonyms are decreased. This is called lateral inhibition and is a standard feature in language game experiments, see [11] for an early implementation. When a new descriptive construction is heard the agent will adopt multiple constructions if he is not sure what the word means, one for every possible meaning. If during a question game type 2 the agent knows exactly what feature the new word is describing he only needs to adopt one construction.

The scores of question constructions are updated in similar ways. If a question game (type 1 or 2) is successful the score of the question construction that was used is increased as it is almost impossible to succeed in such a game if the meaning of the question construction is not shared amongst the agents. If a question game type 1 fails the score of both the question construction and the descriptive construction used is decreased. Decreasing the score of the descriptive construction helps the agents realize that they are moving to fast from guessing to question games. This is important as the agents are autonomously regulating their own progress and moving to fast can seriously damage the language they develop. Using "red" to negotiate a word for color is only useful if all agents understand "red" the same way. If they do not understand its meaning the same way but try to use it anyway the question games will keep failing, the score of "red" will decrease and the agents will stop playing the question game (until they have settled on the meaning of "red").

3 Results

There are striking patterns in how the agents choose to play games. In the beginning they always play guessing games as the other games require some prerequisites to play. When they start seeing features that they have a word for (invented or heard) they will focus on this feature. The feedback of wanting to talk about what you know and learning the words others talk about soon makes them talk about some features more than others. The behavior will then be even more coordinated as they will attempt to find a name for color or shape and now they have similar vocabularies and thus they are all more or less able to start naming the same things (if all know red they can all learn color but if

[1] I will not go through the details of every update rule for every possible situation, only list the most important cases.

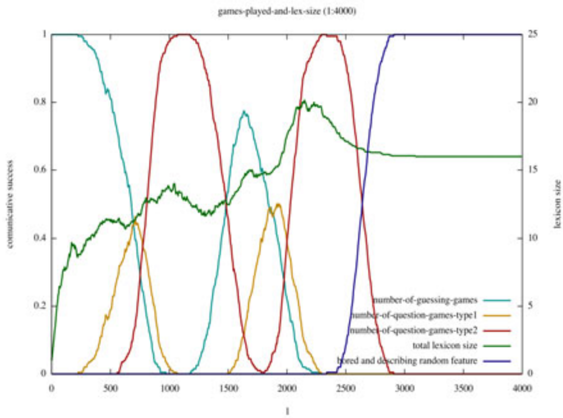

Fig. 1. The figure shows the results of an experiment with 2 different types of features. A pattern can be seen of first guessing games (light blue) followed by learning a question word (yellow) followed by using the question word for further learning (red). The pattern is then repeated when the question word can not be used for learning new things any more ("shape" can be used for learning "round", "square" etc but when all shapes are named a new question word must be invented to move forward). Finally the agents get bored (dark blue) when they are presented with a scene with only features that hey have good names for. The average lexicon size is shown in green.

they all know round they can all learn shape). When they have a word for color they start using it to learn the names of all other colors and now they are very coordinated as can be seen in figure 1. As they learn all the color words they get bored and start the whole process again with for example shape and so on resulting in the wave pattern seen.

To measure how successful an experiment is we need to introduce a formal measure of how good the language they have developed is. Since the agents themselves not only chose the game to play and the topic to talk about, but many times do this in a way that reduces risk of failure (they often choose games that are as easy as possible), a measure of how often individual games succeed is not a very good measure of performance. The communicative goals are only subgoals of learning a language. An agent that makes incorrect approximations about how useful a communicative subgoal is in learning a language can systematically cause it to play very easy games that are very successful but that does not teach it anything. It will have its subgoals met more often than an agent who chooses games where learning takes place (but that fails sometimes) but the former agent should be seen as less successful than the latter.

I introduce lexical coherence as the measure of success for the remainder of this paper. This is defined as the probability that when 2 agents are drawn at random and a feature of the world is picked at random they will both have at least one construction with this feature at the meaning pole and that the highest scoring such constructions for the two agents will have the same form. Put in

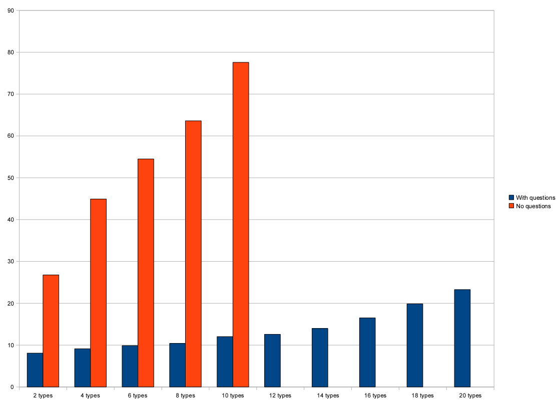

Fig. 2. The blue columns are convergence times for agents asking questions and the red are convergence times for agents playing only the guessing game. All experiments involve 10 agents. Convergence times are measured in number of interactions per feature and agent to reach 98% coherence. It is clear that the questions give the agents a significant boost in performance.

another, and less formal way, it is the probability that two agents will prefer the same word for a feature.

We can get a good overview of how performance improves in figure 2. The convergence times of the agents asking questions are shown in blue and those not asking questions are shown in red. A comparison is made in five different worlds with 2, 4, 6, 8 and 10 types of features (2 types would correspond to a world with shapes and colors, 4 would be a world with shapes, colors, texture and size, etc). An average has been taken over 20 runs in the case of guessing games. For the cases being compared (2,4,6,8 and 10 types) 50 runs have been averaged in the question games. For the question games even more complex worlds have been tested up to 20 different types. For these five experiments fewer runs have been used to average the results (between 5 and 20).

4 Conclusion

The problem of referential uncertainty is an important obstacle when bootstrapping a language in a population of agents from scratch and has been investigated by several different computational approaches before. We have introduced another approach on how this uncertainty can be reduced, by giving increased autonomy to the agents to choose among different interaction scripts based on a motivational system. We have shown that such a system allows the population to achieve significant improvements in the time needed to converge on a language system. It was also demonstrated that agents can self-regulate their interactions and decide when they are ready to move from one type of interaction to another even if they have no global coordination, no way of directly observing the abilities of other agents and only partial information about the interaction history of the population. This suggests that when investigating a linguistic phenomena where different learning stages are required, it might not be necessary for the experimenter to guide the experiment using global information or information that is private to the agents.

References

1. Steels, L.: Evolving grounded communication for robots. Trends in Cognitive Science 7, 308–312 (2003)
2. Baronchelli, A., Loreto, V., Steels, L.: In-depth analysis of the naming game dynamics: the homogeneous mixing case. Int. J. Mod. Phys. C 19, 785 (2008)
3. Steels, L., Loetzsch, M.: Perspective alignment in spatial language. In: Coventry, K.R., Tenbrink, T., Bateman, J.A. (eds.) Spatial Language and Dialogue. Oxford University Press, Oxford (2007) (to appear)
4. De Beule, J.: The emergence of compositionality, hierarchy and recursion in peer-to-peer interactions. In: Smith, A.D.M., Smith, K., Ferrer-i-Cancho, R. (eds.) Proceedings of the 7th International Conference on the Evolution of Language, pp. 75–82. World Scientific, Singapore (2008)
5. van Trijp, R.: The emergence of semantic roles in Fuid construction grammar. In: Smith, A.D.M., Smith, K., Ferrer-i-Cancho, R. (eds.) Proceedings of the 7th International Conference on the Evolution of Language, pp. 346–353. World Scientific, Singapore (2008)
6. Steels, L., Belpaeme, T.: Coordinating perceptually grounded categories through language: A case study for colour. Behavioral and Brain Sciences 28(4), 469–489 (2005)
7. Wellens, P., Loetzsch, M., Steels, L.: Flexible word meaning in embodied agents. Connection Science 20, 173–191 (2008)
8. Smith, A.D.M.: The inferential transmission of language. Adaptive Behavior 13(4), 311–324 (2005)
9. De Beule, J., De Vylder, B., Belpaeme, T.: A cross-situational learning algorithm for damping homonymy in the guessing game. In: Rocha, L.M., et al. (eds.) Artificial Life X, pp. 466–472. MIT Press, Cambridge (2006)
10. Steels, L., Wellens, P.: Scaffolding language emergence using the autotelic principle. In: IEEE Symposium on Artificial Life, pp. 325–332 (2007)
11. Steels, L., Kaplan, F.: Spontaneous lexicon change. In: COLING-ACL 1998, ACL, Montreal, pp. 1243–1249 (1998)
12. Csikszentmihalyi, M.: Flow: The Psychology of Optimal Experience. Harper Perennial (1991)
13. Steels, L.: The autotelic principle. In: Iida, F., Pfeifer, R., Steels, L., Kuniyoshi, Y. (eds.) Embodied Artificial Intelligence. LNCS (LNAI), vol. 3139, pp. 231–242. Springer, Heidelberg (2004)
14. Quine, W.: Word and Object. MIT Press, Cambridge (1960)