

A Multi-Agent Experiment on the Acquisition of a Language System of Logical Constructions

Josefina Sierra-Santibáñez¹

Universidad Politécnica de Cataluña, Spain
(María.Josefina.Sierra@upc.edu**)

Abstract. This paper analyses an experiment which studies the acquisition of the linguistic competence required to communicate logical combinations of categories from the wisdom of the crowds perspective. The acquisition of such competence encompasses both the construction of a set of logical categories by each individual agent and of a shared language by the population. The processes of conceptualisation and language acquisition in each individual agent are based on general purpose cognitive capacities such as discrimination, invention, adoption and induction. The construction of a shared language by the population is achieved using a particular type of linguistic interaction, known as *the evaluation game*, which gives rise to a shared language system of logical constructions as a result of a process of self-organisation of the individual agents' interactions, when these agents adapt their languages to the expressions they observe are used more often by other agents.

1 Introduction

The wisdom of the crowds main thesis is that a diverse collection of independently deciding individuals is likely to make certain types of decisions and predictions better than individuals or even experts. This principle seems to work for many naturally occurring systems such as ant colonies, bird flocks or moving traffic flows, and it has been successfully applied to market prediction [1, 2] and multi-agent computer systems as well [3]. However not all crowds (groups) are wise, and it is therefore important to identify some criteria which separate wise crowds from irrational ones. Four such criteria are described in [4]: (1) diversity of opinion, enough variance in approach, thought processes and private information is necessary; (2) independence, agents' decisions should not be determined by other agents; (3) decentralisation, agents should be able to specialise and draw on local knowledge; and (4) aggregation, some mechanisms should be provided for turning individual decisions into collective ones.

Two additional important aspects of the wisdom of the crowds approach are also pointed out in [5]: the necessity of designing methods for describing how a group thinks as a whole; and the importance of disagreement and contest as mechanisms that enable the generation and selection of optimal decisions.

** The research leading to these results has received funding from the Spanish Ministry of Economy and Competitiveness under project TIN2011-27479-C04-03.

Experiments studying the effectiveness of the wisdom of the crowds approach often incorporate some functions which allow assessing the performance of a group in a given task, thus making it possible to establish a comparison between the collective performance and that of its individuals. Some of these functions are sometimes referred to as *collective intelligence quotient* (or *cooperation quotient*) and compared with the *individual intelligence quotient (IQ)*.

There are, however, other definitions of collective wisdom which not only focus on consensus-driven decision making, but on other aspects of it such as: shared knowledge arrived at by individuals and groups; shared intelligence that emerges from the collaboration, collective efforts and competition of many individuals; or collective learning over time. For example, [6] defines the collective intelligence phenomenon as 'the capacity of human communities to evolve towards higher order complexity and harmony, through such innovation mechanisms as differentiation and integration, competition and collaboration'. A step forward in this direction is *crowdsourced crisis mapping* [7, 8], which tries to bridge the gap between the creation and sharing of knowledge by global communities and the necessary action to solve social problems based on that information. Interesting projects addressing related issues such as the construction of a democratic political culture [9] or generalised access to education using crowd-based socio-cognitive systems [10, 3] are additional examples.

The rest of the paper is organised as follows. Firstly, we present the results of a multi-agent experiment in which a group of autonomous software agents try to construct at the same time a set of logical categories and a shared language. Then, we analyse such results from the wisdom of the crowds perspective, i.e. taking into account the definitions of wisdom of the crowds and criteria for distinguishing wise crowds from irrational ones introduced in this section.

The multi-agent experiment is not described in detail in this paper, although its main characteristics have been outlined in the abstract. A complete description of the evaluation game and the mechanisms the agents use for discrimination, induction and adaptation can be found in [11]. A summary of the main steps of the evaluation game, the induction rules and the adaptation strategies used by the agents are also given in appendixes A y B.

2 Results of the Experiment

As mentioned above, the multi-agent experiment analysed in this paper studies the acquisition of the linguistic competence required to communicate logical combinations of basic categories, such as 'up and to the left' (i.e. [and, up, left]), 'not up or to the right' (i.e. [if, up, right]) or 'either up or to the right, but not both' (i.e. [xor, up, right]). The acquisition of such competence encompasses both the construction of a set of logical categories by each individual agent and of a shared language by the population. In particular, the set of logical categories the agents can construct in this experiment is the set of Boolean functions of one or two arguments *not*, *and*, *nand*, *or*, *nor*, *xor*, *iff*, *if*, *nif*, *oif* and *noif*. Boolean functions *not*, *and*, *or*, *if* and *iff* correspond to the connectives of propositional

logic $\neg, \wedge, \vee, \rightarrow$ and \leftrightarrow respectively. The semantics of Boolean functions *nand*, *nor*, *xor*, *nif*, *oif* and *noif*, assuming they are applied to propositions A and B , can be defined by the following formulas $\neg(A \wedge B)$, $\neg(A \vee B)$, $(A \vee B) \wedge \neg(A \wedge B)$, $\neg(A \rightarrow B)$, $B \rightarrow A$, $\neg(B \rightarrow A)$ respectively.

The experiment involves a population of autonomous software agents which are made to interact with each other playing language games. The particular type of language game used in the experiment analysed in this paper is called the evaluation game. It is played by two agents, a speaker and a hearer. It requires the agents to communicate about subsets of objects of the set of all the objects in a given context. In order to do so, the speaker must construct a logical combination of categories that is true for the subset of objects it tries to communicate about and false for the rest of the objects in the context, i.e. a conceptualisation of the subset. Then, it should transform this conceptualisation into an utterance using its lexicon and grammar, and communicate that utterance to the hearer. The hearer then tries to parse the utterance, reconstruct its meaning and use it to identify the subset of objects the speaker had in mind. Depending on the outcome of the game speaker and hearer use different strategies to expand and adapt their internal languages in order to be more successful in future language games. All agents in the population play both the role of speaker and that of hearer in different language games.

In the experiment, the agents are initially endowed with a set of cognitive abilities for discrimination, invention, adoption and induction that are hypothesised to be necessary for seeing the emergence of possible language strategies to be successful in the evaluation game. Then, they are made to play a series of language games, where they configure possible strategies and try them out. The goal of the experiment is to find out whether the population as a whole succeeds in the evaluation game, i.e. communicates effectively, and to observe the conceptualisations and language strategies that emerge in the population as a result of the processes of collective invention and negotiation.

In the particular multi-agent experiment analysed in this paper we have performed several simulation runs. In each simulation the agents first play 700 evaluation games about subsets of objects which can be discriminated using a single category or the negation of a category. In this part of the simulation the population reaches a *communicative success* of 94% after playing 100 games (see figure 1). *Communicative success* is the average of successful evaluation games in the last ten games played by the agents. Next, the agents play 6000 evaluation games about subsets of objects which require logical combinations of one or two categories for their discrimination. In this part of the simulation the population reaches a communicative success of 100% after playing 3600 evaluation games. As it can be observed in figure 1, this level of communicative success is maintained until the end of the simulation. The results shown in the figure are the average of ten independent simulation runs with different random seeds.

At the end of a typical simulation run the set of logical categories and grammatical constructions built by each agent are not necessarily equal to the set of logical categories and grammatical constructions built by other agents. However

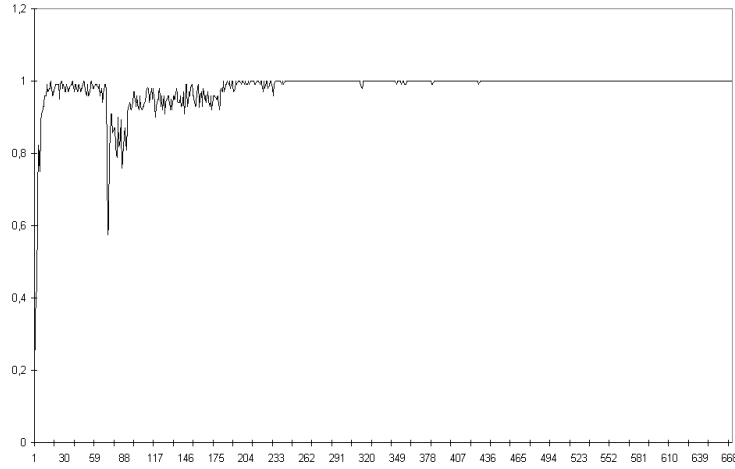


Fig. 1. Evolution of communicative success for a population of three agents.

they are compatible in the sense that they guarantee the unambiguous communication of logical combinations of one or two categories.

Let us focus now on the set of logical categories and grammatical constructions built by three agents at the end of a particular simulation run (see table 1). All the agents have constructed a grammar rule for expressing negations, and all of them use the same expression (i.e. *cp*) for referring to logical category *not*.

All the agents have constructed logical categories for all **commutative boolean functions of two arguments** (i.e. *iff*, *xor*, *and*, *nand*, *or* and *nor*) as well; and all of them prefer the same expressions for naming such categories (*j*, *wbt*, *y*, *nb*, *dol* and *ssq* respectively).

In order to express logical formulas constructed with binary Boolean functions, the agents use two types of grammar rules. Which rule is used for expressing a given formula depends on the Boolean function appearing in that formula and the syntactic category of the expression associated with such Boolean function. Syntactic category *c1* is used in grammatical constructions which place the expression associated with the first argument of a Boolean function in the second position of the sentence and the expression associated with the second argument of the Boolean function in the third position of the sentence. Syntactic category *c2* is used in grammatical constructions which place the expression associated with the first argument of a Boolean function in the third position of the sentence and the expression associated with the second argument of the Boolean function in the second position of the sentence. The expression associated with a Boolean function is always placed in the first position of the sentence in this experiment.

We now consider **non-commutative binary Boolean functions** (i.e. *if*, *nif*, *oif* and *noif*). All the agents have constructed logical category *nif*, which

Grammar a1
$s([\text{not}, X], Q) \rightarrow cp, s(X, P), \{Q \text{ is } P \cdot 1\}$ $s([\mathbf{X}, \mathbf{Y}, \mathbf{Z}], \mathbf{T}) \rightarrow c1(\mathbf{X}, \mathbf{P}), s(\mathbf{Y}, \mathbf{Q}), s(\mathbf{Z}, \mathbf{R}), \{\mathbf{T} \text{ is } \mathbf{P} \cdot \mathbf{Q} \cdot \mathbf{R} \cdot 1\}$ $c1(\text{nif}, R) \rightarrow ml, \{R \text{ is } 1\}$ $c1(\text{nor}, R) \rightarrow nb, \{R \text{ is } 1\}$ $c1(\text{or}, R) \rightarrow y, \{R \text{ is } 1\}$ $s([\mathbf{X}, \mathbf{Y}, \mathbf{Z}], \mathbf{T}) \rightarrow c2(\mathbf{X}, \mathbf{P}), s(\mathbf{Z}, \mathbf{Q}), s(\mathbf{Y}, \mathbf{R}), \{\mathbf{T} \text{ is } \mathbf{P} \cdot \mathbf{Q} \cdot \mathbf{R} \cdot 1\}$ $c2(\text{and}, R) \rightarrow j, \{R \text{ is } 1\}$ $c2(\text{xor}, R) \rightarrow dol, \{R \text{ is } 1\}$ $c2(\text{iff}, R) \rightarrow ssq, \{R \text{ is } 1\}$ $c2(\text{nand}, R) \rightarrow wbt, \{R \text{ is } 1\}$ $c2(\text{if}, R) \rightarrow why, \{R \text{ is } 1\}$
Grammar a2
$s([\text{not}, X], Q) \rightarrow cp, s(X, P), \{Q \text{ is } P \cdot 1\}$ $s([\mathbf{X}, \mathbf{Y}, \mathbf{Z}], \mathbf{T}) \rightarrow c1(\mathbf{X}, \mathbf{P}), s(\mathbf{Y}, \mathbf{Q}), s(\mathbf{Z}, \mathbf{R}), \{\mathbf{T} \text{ is } \mathbf{P} \cdot \mathbf{Q} \cdot \mathbf{R} \cdot 1\}$ $c1(\text{nif}, R) \rightarrow ml, \{R \text{ is } 1\}$ $c1(\text{nor}, R) \rightarrow nb, \{R \text{ is } 1\}$ $c1(\text{or}, R) \rightarrow y, \{R \text{ is } 1\}$ $s([\mathbf{X}, \mathbf{Y}, \mathbf{Z}], \mathbf{T}) \rightarrow c2(\mathbf{X}, \mathbf{P}), s(\mathbf{Z}, \mathbf{Q}), s(\mathbf{Y}, \mathbf{R}), \{\mathbf{T} \text{ is } \mathbf{P} \cdot \mathbf{Q} \cdot \mathbf{R} \cdot 1\}$ $c2(\text{and}, R) \rightarrow j, \{R \text{ is } 1\}$ $c2(\text{xor}, R) \rightarrow dol, \{R \text{ is } 1\}$ $c2(\text{iff}, R) \rightarrow ssq, \{R \text{ is } 1\}$ $c2(\text{nand}, R) \rightarrow wbt, \{R \text{ is } 1\}$ $c2(\text{if}, R) \rightarrow why, \{R \text{ is } 1\}$
Grammar a3
$s([\text{not}, X], Q) \rightarrow cp, s(X, P), \{Q \text{ is } P \cdot 1\}$ $s([\mathbf{X}, \mathbf{Y}, \mathbf{Z}], \mathbf{T}) \rightarrow c1(\mathbf{X}, \mathbf{P}), s(\mathbf{Y}, \mathbf{Q}), s(\mathbf{Z}, \mathbf{R}), \{\mathbf{T} \text{ is } \mathbf{P} \cdot \mathbf{Q} \cdot \mathbf{R} \cdot 1\}$ $c1(\text{nif}, R) \rightarrow ml, \{R \text{ is } 1\}$ $c1(\text{nor}, R) \rightarrow nb, \{R \text{ is } 1\}$ $c1(\text{or}, R) \rightarrow y, \{R \text{ is } 1\}$ $c1(\text{oif}, R) \rightarrow why, \{R \text{ is } 1\}$ $s([\mathbf{X}, \mathbf{Y}, \mathbf{Z}], \mathbf{T}) \rightarrow c2(\mathbf{X}, \mathbf{P}), s(\mathbf{Z}, \mathbf{Q}), s(\mathbf{Y}, \mathbf{R}), \{\mathbf{T} \text{ is } \mathbf{P} \cdot \mathbf{Q} \cdot \mathbf{R} \cdot 1\}$ $c2(\text{and}, R) \rightarrow j, \{R \text{ is } 1\}$ $c2(\text{xor}, R) \rightarrow dol, \{R \text{ is } 1\}$ $c2(\text{iff}, R) \rightarrow ssq, \{R \text{ is } 1\}$ $c2(\text{nand}, R) \rightarrow wbt, \{R \text{ is } 1\}$

Table 1. Logical categories and grammatical constructions built by each agent at the end of a particular simulation run. In principle, the agents can construct logical categories *not*, *and*, *nand*, *or*, *nor*, *if*, *nif*, *oif*, *noif*, *iff* and *xor*, although they do not necessarily construct all of them. Boolean functions *not*, *and*, *or*, *if* and *iff* have the standard interpretation ($\neg, \wedge, \vee, \rightarrow$ and \leftrightarrow respectively). The rest can be defined as follows: (A nand B) is equivalent to $\neg(A \wedge B)$, (A nor B) to $\neg(A \vee B)$, (A nif B) to $\neg(A \rightarrow B)$, (A oif B) to $(B \rightarrow A)$, (A noif B) to $\neg(B \rightarrow A)$, and (A xor B) to $\neg(A \vee B) \wedge \neg(A \wedge B)$.

corresponds to *the negation of an implication*, all of them use the same expression (i.e. *ml*) for referring to it, and all of them associate the expression *ml* with syntactic category *c1*.

None of the agents has constructed logical category *noif*. But this does not prevent them from characterising any subset of objects, because formulas [noif, A, B] and [nif, B, A] are logically equivalent and all the agents have constructed logical category *nif*.

Let us focus now on differences. Agents a1 and a2 have constructed logical category *if* (i.e. logical implication), whereas a3 has not. On the other hand, agent a3 has constructed logical category *oif*, while agents a1 and a2 have not. However the lack of only one of these two logical categories does not prevent any agent from characterising any subset of objects, because formulas [if, A, B] and [oif, B, A] are logically equivalent. Furthermore, the three agents can always understand each other. Because the word agents a1 and a2 use for referring to logical category *if* (namely *why*) is the same word agent a3 uses for referring to logical category *oif*; and the syntactic category agents a1 and a2 associate with such word (i.e. *c2*) is different from the syntactic category agent a3 uses for it (i.e. *c1*), which means that agent a3 does not invert the order of the expressions associated with the arguments of *oif* in the sentence whereas agents a1 and a2 invert the order of the expressions associated with the arguments of *if*.

3 Discussion

The experiment described in this paper constitutes an example of *collective learning and coordination*. As we have explained above the agents construct a language system of logical constructions that allows them to communicate logical combinations of categories. This language system includes a common vocabulary for logical categories, and a set of grammatical constructions which allow them to order the expressions associated with the components of logical formulas in sufficiently similar ways as to ensure unambiguous communication.

In this section we try to analyse the results of the experiment from the wisdom of the crowds perspective, focusing on the definitions of wisdom of the crowds and criteria for distinguishing wise crowds from irrational ones introduced in section one.

First of all, does the population make better decisions than individual agents in the experiment? It might be difficult to answer such a question without knowing in detail the mechanisms each agent uses to construct logical categories, invent new words and induce grammatical constructions (which are described in detail in [11]), but we think it does. The population is able to recognise that certain binary Boolean functions are redundant. For example, *nif* and *noif* can be used for discriminating the same subsets of objects, and the same happens with *if* and *oif*. Consequently, the language of the population contains only two words for the four logical categories (*if*, *nif*, *oif*, *noif*). The mechanisms the individual agents use for constructing a set of logical categories and grammatical constructions do not allow them to discover such redundancies for themselves. It

is the interaction with other agents, who use different formulas for conceptualising the same subset of objects, what generates the opportunity of first using the same word for two different categories, and then selecting a single meaning for that word as a result of the selection process that takes place among competing associations between expressions and meanings both in the individual languages constructed by each agent and in the language constructed by the population.

The population is also able to discover that the word order of the expressions associated with the arguments of commutative Boolean functions is irrelevant for language understanding. This cannot be observed in the grammars shown in table 1. But in other simulation runs we have performed the word associated with a commutative Boolean function such as *and* can be associated with syntactic category *c1* in an agent's grammar and with syntactic category *c2* in the grammar of a different agent of the same population. In the current experiment, the agents themselves are not aware of this fact. Because each agent uses a particular word order for expressing formulas constructed with each commutative Boolean function. But the external language spoken by the population shows that they perfectly understand each other in spite of using different word orders for the expressions associated with the arguments of commutative Boolean functions.

With respect to the four criteria proposed by [4], the agents in the population have some degree of *diversity of opinion*, in the sense that they can invent different words for referring to logical categories and order the constituents of sentences in different ways. But they basically use the same approach for constructing logical categories and for expressing logical formulas, i.e. they all use word order as the only syntactic mechanism for disambiguation and non-recursive formulas of one or two arguments for discrimination. They are *independent* of each other, because each agent chooses the words and grammatical constructions it uses for communication taking into account only the scores of the rules in its own grammar. The scores of such rules depend on the interaction history of the agent, which is always different from the interaction history of the others, thus providing each individual agent with a *different perspective* of the language used by the population. The *aggregation mechanism* used in the experiments is, as we explained above, the shared language system that emerges as a result of the self-organisation process of the interactions that take place among the agents in the population. The mechanisms the agents use to adapt their languages to the expressions they observe are used more often by other agents favour such self-organisation process, because each agent tries to use the same expressions as the others.

The necessity of designing methods for describing how a group thinks as a whole, pointed out by [5], is addressed in multi-agent experiments studying language emergence and evolution using a number of functions which evaluate the performance of the group as a whole. In the present experiment we have used *communicative success*, but other functions which compute the similarity of the agents' grammars, the discriminating capacity of the set of categories constructed by the population, or the complexity of the vocabulary and gram-

matical constructions of the language spoken by the population, can be used as well [12].

Finally, we think that the agent interaction mechanisms used to construct compatible conceptualisations and a shared language system in the experiment analysed in this paper might be appropriately adapted and applied to crowd-based socio-cognitive systems [3] addressing issues such as crowdsourced crisis mapping [7], the construction of a democratic political culture [9] or the generalisation of access to education [10]. Because in each of these domains sets of new concepts and linguistic constructions need to be constructed in order to accurately reflect the reality their users are dealing with and to enable communication, and the best way of constructing such new language systems is using mechanisms that enable meaning coordination.

References

1. Jahedpari, F., Padget, J., De Vos, M., Hirsch, B.: Artificial prediction markets as a tool for syndromic surveillance. In: *Crowd Intelligence: Foundations, Methods and Practices*. (2014)
2. Buckley, P.: A research agenda for prediction markets. In: *Crowd Intelligence: Foundations, Methods and Practices*. (2014)
3. Noriega, P., D’Inverno, M.: Crowd-based socio-cognitive systems. In: *Crowd Intelligence: Foundations, Methods and Practices*. (2014)
4. Surowiecki, J.: *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business*. Economies, Societies and Nations (2004)
5. Oinas-Kukkonen, H.: Network analysis and crowds of people as sources of new organisational knowledge. In: *Knowledge Management: Theoretical Foundation*. Informing Science Press (2008) 173–189
6. Por, G.: *Blog of collective intelligence* (2013)
7. Poblet, M., Casanovas, P.: Crowdsourced crisis mapping: how it works and why it matters. *The Conversation* (<http://theconversation.edu.au>) (2012)
8. Poblet, M., García-Cuesta, E., Casanovas, P.: It enabled crowds: Leveraging the geomobile revolution. In: *Crowd Intelligence: Foundations, Methods and Practices*. (2014)
9. Nesh-Nash, T., Deely, S.: The future of democratic participation: Online collaboration to build a democratic political culture. In: *Crowd Intelligence: Foundations, Methods and Practices*. (2014)
10. D’Inverno, M., Yee-King, M.: Pedagogical agents for social music learning in crowd-based socio-cognitive systems. In: *Crowd Intelligence: Foundations, Methods and Practices*. (2014)
11. Sierra, J., Santibáñez, J.: Experiments on the acquisition of cognitive and linguistic competence to communicate propositional logic sentences. In: *Papers from the AAAI Fall Symposium (FS-09-01)*. (2009) 153–158
12. Steels, L.: Modeling the cultural evolution of language. *Physics of Life Review* **8** (2011) 339–356
13. Vogt, P.: The emergence of compositional structures in perceptually grounded language games. *Artificial Intelligence* **167(1-2)** (2005) 206–242
14. Kirby, S.: Learning, bottlenecks and the evolution of recursive syntax. In: *Linguistic Evolution through Language Acquisition: Formal and Computational Models*, Cambridge University Press (2002) 96–109

15. Sierra, J., Santibáñez, J.: The acquisition of linguistic competence for communicating propositional logic sentences. In: Engineering Societies in the Agents World VIII, Lecture Notes in Computer Science, volume 4955 (2008) 175–191

A The Evaluation Game

The emergence of a shared language system of logical constructions in the population results from a process of self-organisation of the linguistic interactions that take place among the agents in the population. The particular type of linguistic interaction used in the experiment discussed in this paper is called *the evaluation game*. It is played by two agents, a speaker and a hearer, and its main steps can be summarised as follows.

1. Conceptualisation Firstly both agents, speaker and hearer, are given a description of a set of objects which constitute *the context* of the evaluation game. Then the speaker picks a subset of objects from the context which will be *the topic* of the evaluation game. The rest of the objects in the context are called *the background*.

The speaker tries to construct a *conceptualisation* of the topic, that is, a logical formula which is true for all the objects in the topic and false for all the objects in the background. It does so by finding a unary or binary tuple of categories such that its evaluation on the topic is different from its evaluation on any object in the background. Once it has found a discriminating category tuple, the speaker tries to find a logical category which is associated with the subset of Boolean values or pairs of Boolean values resulting from evaluating the topic on that category tuple, and constructs a conceptualisation of the topic applying this logical category to the discriminating category tuple.

In general an agent can build several conceptualisations for the same topic. For example, if the context contains objects 1, 2 and 3 such that object 1 is up and to the left, object 2 is down and to the left, and object 3 is down and to the right, and the topic consists of objects 1 and 3, then both formulas [iff, up, left] and [xor, up, right] can be used as conceptualisations of the topic.

2. Generation The speaker tries to generate a sentence for each of its conceptualisations of the topic using its lexicon and grammar. It tries to maximise the probability of being understood by other agents by selecting the sentence with the highest score, and communicates that sentence to the hearer. The algorithm for computing the score of a sentence from the scores of the grammar rules used in its generation is explained in [13].

The agents in the population start with an empty lexicon and an empty grammar. Therefore they cannot generate sentences for most formulas (conceptualisations) at the early stages of a simulation run. In order to let language to get off the ground, they are allowed to invent new sentences for those meanings (conceptualisations) they cannot express using their lexicon and grammar. As

the agents play language games they learn associations between expressions and meanings, and induce linguistic knowledge from such associations in the form of grammar rules and lexical entries.

3. Interpretation If the hearer can parse the sentence communicated by the speaker using its lexicon and grammar, it extracts a formula (a meaning) and uses that formula to identify the topic. At the early stages of a simulation run the hearers usually cannot parse the sentences communicated by the speakers, since they have no prior linguistic knowledge. In this case the speaker points to the topic, and the hearer adopts an association between its conceptualisation of the topic and the sentence used by the speaker. Note that the conceptualisations of speaker and hearer might be different, because different formulas can be used to conceptualise the same topic.

4. Adaptation The evaluation game is successful if the hearer can parse the sentence communicated by the speaker, and its interpretation of that sentence identifies the topic (i.e. the subset of objects the speaker had in mind) correctly. Depending on the outcome of the evaluation game, speaker and hearer take different actions. We have explained some of them already (*invention* and *adoption*), but they also *adapt their grammars* to communicate more successfully in future games.

Coordination of the agents' grammars is necessary, because different agents can invent different words to refer to the same categories, and because the invention process uses a random order to concatenate the expressions associated with the components of a given formula. In order to understand each other, the agents must use a common vocabulary and must order the constituents of sentences in sufficiently similar ways as to avoid ambiguous interpretations. The following *adaptation mechanisms* are used to coordinate the agents' grammars.

We consider the case in which the speaker can generate a sentence and the hearer can parse it. If the speaker can generate several sentences for its conceptualisations of the topic, the sentence with the highest score is chosen for communication and the rest of the sentences are kept as *competing sentences*. Similarly if the hearer can obtain several formulas (meanings) for the sentence communicated by the speaker, the formula with the highest score is selected as its interpretation of the sentence and the rest of the formulas are kept as *competing meanings*.

If the topic identified by the hearer is the subset of objects the speaker had in mind, the evaluation game succeeds. The speaker increases the scores of the grammar rules it used for generating the sentence communicated to the hearer and decreases the scores of the grammar rules it used for generating competing sentences. The hearer increases the scores of the grammar rules it used for obtaining its interpretation of the sentence and decreases the scores of the rules it used for obtaining competing meanings. This way the grammar rules which have been used successfully get reinforced, and the grammar rules which have been used for generating competing sentences or competing meanings are inhibited.

If the topic identified by the hearer is different from the subset of objects the speaker had in mind, the evaluation game fails and both agents decrease the scores of the grammar rules they used for generating and interpreting the sentence used by the speaker respectively. This way the grammar rules used without success are inhibited.

The scores of grammar rules are *updated* replacing the rule’s original score S with the result of evaluating expression 1 if the score is *increased*, and with the result of evaluating expression 2 if the score is *decreased*.

$$\text{minimum}(1, S + 0.1) \tag{1}$$

$$\text{maximum}(0, S - 0.1) \tag{2}$$

B Induction

Besides inventing expressions and adopting associations between sentences and meanings, the agents use some *induction mechanisms* to extract generalisations from the grammar rules they have learnt so far. The induction mechanisms used in this paper are based on the rules of *simplification and chunk* in [14], although we have extended them so that they can be applied to grammar rules which have scores attached to them [13]. The induction rules are applied whenever the agents invent or adopt a new association to avoid redundancy and increase generality in their grammars.

Instead of giving a formal definition of the induction rules used in the experiment, which can be found in [15], we give an example of their application. We use *Definite Clause Grammar* to represent the internal grammars constructed by the individual agents. *Non-terminals* have two arguments attached to them. The first argument conveys semantic information and the second is a *score* in the interval $[0, 1]$ which estimates the usefulness of the grammar rule in previous communication. Suppose an agent’s grammar contains the following rules.

$$s(\text{light}, S) \rightarrow \text{clair}, \{S \text{ is } 0.70\} \tag{3}$$

$$s(\text{right}, S) \rightarrow \text{droit}, \{S \text{ is } 0.25\} \tag{4}$$

$$s([\text{and}, \text{light}, \text{right}], S) \rightarrow \text{etclairdroit}, \{S \text{ is } 0.01\} \tag{5}$$

$$s([\text{or}, \text{light}, \text{right}], S) \rightarrow \text{ouclairdroit}, \{S \text{ is } 0.01\} \tag{6}$$

The induction rule of **simplification**, applied to 5 and 4, allows generalising grammar rule 5 replacing it with 7. In this case *simplification* assumes that the second argument of logical category ‘and’ can be any meaning that can be expressed by a ‘sentence’, because according to rule 4 the syntactic category of expression ‘droit’ is $s(\text{sentence})$.

$$s([\text{and}, \text{light}, B], S) \rightarrow \text{etclair}, s(B, R), \{S \text{ is } R \cdot 0.01\} \tag{7}$$

Simplification, applied to rules 7 and 3, can be used to generalise rule 7 replacing it with 8. Rule 6 can be generalised as well replacing it with rule 9.

$$s([\text{and}, A, B], S) \rightarrow \text{et}, s(A, Q), s(B, R), \{S \text{ is } Q \cdot R \cdot 0.01\} \tag{8}$$

$$s([or, A, B], S) \rightarrow ou, s(A, Q), s(B, R), \{S \text{ is } Q \cdot R \cdot 0.01\} \quad (9)$$

Induction rule **chunk I** replaces a pair of grammar rules such as 8 and 9 with a single rule 10 which is more general, because it makes abstraction of their common structure introducing a syntactic category $c2$ for binary connectives. Rules 11 and 12 state that the expressions *et* and *ou* belong to syntactic category $c2$.

$$s([C, A, B], S) \rightarrow c2(C, P), s(A, Q), s(B, R), \{S \text{ is } P \cdot Q \cdot R \cdot 0.01\} \quad (10)$$

$$c2(\text{and}, S) \rightarrow \text{et}, \{S \text{ is } 0.01\} \quad (11)$$

$$c2(\text{or}, S) \rightarrow \text{ou}, \{S \text{ is } 0.01\} \quad (12)$$

Suppose the agent of previous examples adopts or invents the following rule.

$$s([\text{if}, le, up], S) \rightarrow \text{siclairdroit}, \{S \text{ is } 0.1\} \quad (13)$$

Simplification of rule 13 with rules 3 and 4 would replace rule 13 with 14.

$$s([\text{if}, Q, R], S) \rightarrow \text{si}, s(Q, SQ), s(R, SR), \{S \text{ is } SQ \cdot SR \cdot 0.1\} \quad (14)$$

Then induction rule **chunk II**, applied to 14 and 10, would replace rule 14 with rule 15.

$$c2(\text{if}, S) \rightarrow \text{si}, \{S \text{ is } 0.1\} \quad (15)$$