Experiments on the Acquisition of the Semantics and Grammatical Constructions Required for Communicating Propositional Logic Sentences

Josefina Sierra[1,⋆] and Josefina Santibáñez[2]

[1] Software Department, Technical University of Catalonia, Spain
jsierra@lsi.upc.edu

[2] Education Department, University of La Rioja, Spain
josefina.santibanez@unirioja.es

Abstract. We describe some experiments which simulate a grounded approach to language acquisition in which a population of autonomous agents without prior linguistic knowledge tries to construct at the same time a conceptualisation of its environment and a shared language. The conceptualisation and language acquisition processes in each individual agent are based on general purpose cognitive capacities, such as categorisation, discrimination, evaluation and induction. The emergence of a shared language in the population results from a process of self-organisation of a particular type of linguistic interaction which takes place among the agents in the population.

The experiments, which extend previous work by addressing the problem of the acquisition of *both the semantics and the syntax of propositional logic,* show that at the end of the simulation runs the agents build different conceptualisations and different grammars. However, these conceptualisations and grammars are compatible enough to guarantee the unambiguous communication of propositional logic sentences.

Furthermore the categorisers of the perceptually grounded and logical categories built during the conceptualisation and language acquisition processes can be used for some forms of common sense reasoning, such as determining whether a sentence is a tautology, a contradiction, a common sense axiom or a merely satisfiable formula.

Keywords: Language acquisition, logical categories, induction, self-organisation.

1 Introduction

This paper addresses the problem of the acquisition of both *the semantics and the syntax* (i.e., lexicon and grammatical constructions) required for constructing and communicating concepts of the same complexity as propositional logic formulas. It describes some experiments in which a population of autonomous

⋆ Partially supported by the MICINN SESAAME-BAR (TIN2008-06582-C03-01) and DGICYT MOISES-BAR (TIN2005-08832-C03-03) projects.

A.V. Vasilakos et al. (Eds.): AUTONOMICS 2009, LNICST 23, pp. 236–251, 2010.

agents without prior linguistic knowledge constructs at the same time a conceptualisation of its environment and a shared language. The experiments show that at the end of the simulation runs the agents build different conceptualisations and different grammars. However these conceptualisations and grammars are compatible enough to guarantee the unambiguous communication of meanings of the same complexity as propositional logic formulas.

The research presented in this paper builds up on previous work on the acquisition of *the semantics of logical connectives* [1] by addressing the problem of the acquisition of both the semantics and the syntax of propositional logic. In [1] a grounded approach to the acquisition of logical categories (i.e., the semantics of logical connectives) based on the discrimination of a "subset of objects" from the rest of the objects in a given context is described. Logical categories are constructed by the agents identifying subsets of the range of the truth evaluation process (i.e., sets of Boolean pairs or Boolean values) which result from evaluating a pair of perceptually grounded categories or a single category on a subset of objects. Discrimination is performed characterising a "subset of objects" by a logical formula constructed from perceptually grounded categories which is satisfied by the objects in the subset and not satisfied by the rest of the objects in the context.

The complementary problem of the acquisition of *the syntax of propositional logic* by a population of autonomous agents without prior linguistic knowledge has been addressed independently as well. In [2] an approach to the acquisition of the syntax of propositional logic based on *general purpose cognitive capacities,* such as invention, adoption and induction, and on *self-organisation principles* is proposed. The experiments described in [2] show that a shared language (i.e., a lexicon and a grammar) expressive enough to allow the communication of meanings of the same complexity as propositional logic formulas can emerge in a population of autonomous agents without prior linguistic knowledge. This shared language, although simple, has some interesting properties found in natural languages such as recursion, syntactic categories for propositional sentences and connectives, and partial word order for marking the scope of each connective.

The acquisition of the syntax of subsets of logic has been addressed as well by other authors. In particular [3,4,5] study the emergence of case-based and recursive communication systems in populations of agents without prior linguistic knowledge. However none of these works deals with the problem of the acquisition of both the semantics and the syntax of logic.

The experiments described in this paper extend therefore previous work by using a linguistic interaction (*the evaluation game*) in which the agents must first conceptualise the *topic* (a subset of objects) using the mechanisms proposed in [1] for the acquisition of logical categories, and then construct a *shared language* (a lexicon and a grammar) using the invention, adoption, induction and self-organisation mechanisms proposed in [2].

The rest of the paper is organised as follows. Firstly we describe the mechanisms the agents use in order to conceptualise sensory information. Secondly we consider the process of truth evaluation and explain how logical categories can be discovered by identifying sets of outcomes of the truth evaluation process. Then we focus on the construction and emergence of a shared communication language describing the main steps of *the evaluation game:* conceptualisation, verbalising, interpretation, induction and coordination. Next we present the results of some experiments in which three agents without prior linguistic knowledge build a conceptualisation and a shared language that allows them to construct and communicate meanings of the same complexity as propositional logic formulas. Finally we consider the issue of common sense reasoning and summarise the main ideas we tried to put forward in this paper.

2 Conceptualisation: Basic Definitions

We assume an experimental setting similar to that proposed in *The Talking Heads Experiment* [6]: A set of robotic agents playing language games with each other about scenes perceived through their cameras on a white board in front of them. Figure 1 shows a typical configuration of the white board with several geometric figures pasted on it.

Firstly we describe how the agents conceptualise the sensory information they obtain by looking at the white board and trying to characterise subsets of objects pasted on it.

Sensory Channels. The agents look at one area of the white board by capturing an image of that area with their cameras. They segment the image into coherent units in order to identify the objects which constitute the context of a language game, and use some *sensory channels* to gather information about each segment, such as its horizontal and vertical position, or its light intensity. In the experiments described in this paper we only use three sensory channels: (1) H(o), which computes the horizontal position of an object o; (2) V(o), which computes its vertical position; and (3) L(o), which computes its light intensity. The values returned by the sensory channels are scaled with respect to the area of the white board captured by the agents cameras so that its range is the interval (0.0 1.0).

Perceptually Grounded Categories. The data returned by the sensory channels are values from a continuous domain. To be the basis of a natural language conceptualisation, these values must be transformed into a discrete domain. One form of categorisation consists in dividing up each domain of output values of a particular sensory channel into regions and assigning a *category* to each region [6]. For example, the range of the H channel can be cut into two halves leading to the categories [left] $(0.0 < H(x) < 0.5)$ and [right] $(0.5 < H(x) < 1.0)$. Object 3 in figure 1 has the value H(O3)=0.8 and would therefore be categorised as [right].

Perceptually Grounded Categorisers. At the same time the agents build categories in order to conceptualise sensory information, they construct as well cognitive procedures (called *categorisers*) which allow them to check whether these categories hold or not for a given object.

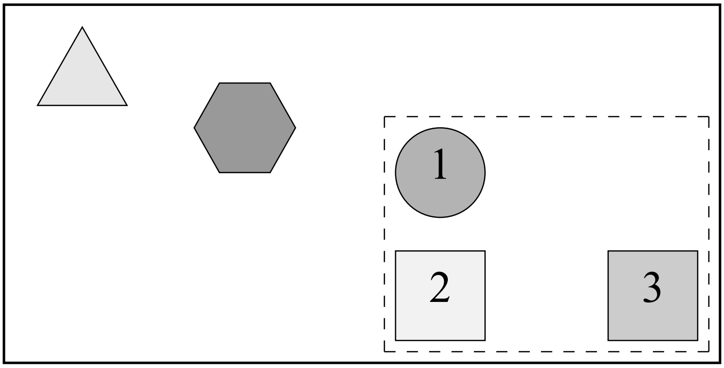

Fig. 1. The area of the white board captured by the agents cameras (i.e., the context of the language game) is the lower right rectangle

Categorisers give grounded meanings [7] to categories (i.e., symbolic representations) by establishing explicit connections between them and reality (external input processed by sensory channels). These connections are learned playing language games [8,6] and allow the agents to check whether a category holds or not for a given object. Most importantly they provide information on the sensory and cognitive processes an agent must go through in order to evaluate a given category.

The behaviour of the categorisers associated with the perceptually grounded categories used in this paper can be described by linear constraints[1]. For example, the behaviour of the categoriser associated with the category [left] can be described as follows: $[\text{left}]^C(x) \equiv 0.0 < \text{H}(x) < 0.5$.

2.1 Logical Categories

We consider now the process of truth evaluation and describe how logical categories can be constructed by identifying sets of outcomes of the truth evaluation process. Logical categories are important because they allow the generation of structured units of meaning, which correspond to logical formulas, and they set the basis for deductive reasoning.

Evaluation Channel. The *evaluation channel* (denoted by E) is a cognitive process capable of finding the categorisers of a tuple of categories, applying them to an object, and observing their output. If $\boldsymbol{c} = (c_1, \ldots, c_n)$ is a category tuple and o is an object, $E(\boldsymbol{c}, o)$ is a tuple of Boolean values $(v_1, \ldots, v_n)$, where each v_i is the result of applying c_i^C (the categoriser of c_i) to object o. For example, $E(([down], [right]), O1) = (0, 0)$, because $O1$ (object 1 in figure 1) is neither on the lower part nor on the right part of the white board area captured by the agents' cameras.

[1] We use the notation $[\text{cat}]^C$ to refer to the categoriser that is capable of determining whether category [cat] holds or not for a given object.

Logical Categories and Formulas. Although the evaluation channel can be applied to category tuples of any arity, we consider only unary and binary category tuples. The range of the evaluation channel for single categories is the set of Boolean values {0, 1}, and its range for category pairs is the set of Boolean pairs {(0,0), (0,1), (1,0), (1,1)}. By considering all the subsets of these ranges the agents can represent all the Boolean functions of one and two arguments, which correspond to the meanings of all the connectives of propositional logic (i.e., $\neg, \wedge, \vee, \rightarrow$ and $\leftrightarrow$), plus the meanings of other connectives (such as *neither* or *exclusive disjunction*) found in natural languages. For example, the propositional formula $c_1 \vee c_2$ is true for an object o if the result of evaluating the pair of categories (c_1, c_2) on object o is a Boolean pair which belongs to the subset of Boolean pairs $\{(1,1),(1,0),(0,1)\}$.

The sixteen Boolean functions of two arguments which can be constructed using this method are summarised in the following ten connectives in the internal representation of logical categories used by the agents: *and, nand, or, nor, if, nif, oif, noif, iff* and *xor*. Where *and, or, if* and *iff* have the standard interpretation ($\wedge, \vee, \rightarrow$ and $\leftrightarrow$), and the formulas *(A nand B), (A nor B), (A nif B)*, (A oif B), (A noif B) and *(A xor B)* are equivalent to $\neg(A \wedge B)$, $\neg(A \vee B)$, $\neg(A \rightarrow B)$, $(B \rightarrow A)$, $\neg(B \rightarrow A)$ and $\neg(A \leftrightarrow B)$, respectively.

The agents construct *logical categories* by identifying subsets of the range of the evaluation channel. The *evaluation game* creates situations in which the agents discover subsets of the range of the evaluation channel, and use them to distinguish a subset of objects from the rest of the objects in a given context. The representation of logical categories as subsets of Boolean tuples is equivalent to the *truth tables* used for defining the semantics of logical connectives.

Logical categories describe properties of propositions, therefore it is natural to apply them to perceptually grounded categories in order to construct structured units of meaning. For example, the formula [not, down] can be constructed by applying the logical category [not] (i.e., $\neg$) to the category [down]. The formula [or, up, right] can be constructed similarly by applying the logical category [or] to the categories [up] and [right][2].

If we consider perceptually grounded categories as propositions, we can observe that the set of concepts that can be constructed by the agents corresponds to the set of formulas of propositional logic, because: (1) a perceptually grounded category is a formula; and (2) if l is an n-ary logical category and F is a list (tuple) of n formulas, then $[l|F]$ is a formula[3].

Logical Categorisers. The categorisers of logical categories are cognitive processes that allow determining whether a logical category holds or not for a tuple of categories and an object. As we have explained above, logical categories can be associated with subsets of the range of the evaluation channel. The behaviour of their categorisers can be described therefore by constraints of the form

[2] Notice that we use prefix, Lisp like notation for representing propositional formulas. Thus the list [or, up, right] corresponds to the formula $up \vee right$.

[3] Where l is a logical category, F is a list of formulas and | is the standard list construction operator.

$E(\boldsymbol{c}, o) \in S_l$, where l is a logical category, S_l is the subset of the range of the evaluation channel for which l holds, E is the evaluation channel, $\boldsymbol{c}$ is a tuple of categories, and o is an object. For example, the constraint $E((c1, c2), o) \in \{(1, 1)\}$ describes the behaviour of the categoriser of the logical category [and] (i.e., $c1 \wedge c2$).

The evaluation channel can be naturally extended to evaluate arbitrary propositional logic formulas using the categorisers of logical and perceptually grounded categories. The following is an inductive definition of the evaluation channel $E(A, o)$ for an arbitrary formula A of propositional logic:

1. If A is a perceptually grounded category $[cat]$, then $E(A, o) = [cat]^C(o)$.
2. If A is a propositional formula of the form $[l|F]$, where l is a logical category, F is a list of formulas and S_l is the subset of the range of the evaluation channel for which l holds, then $E(A, o) = 1$ if $E(F, o) \in S_l$, and 0 otherwise.

3 Language Acquisition

Language acquisition is seen as a collective process by which a population of autonomous agents without prior linguistic knowledge constructs a *shared language* which allows them to communicate some set of meanings. In order to reach such an agreement the agents interact with each other playing language games. In a typical experiment thousands of language games are played by pairs of agents randomly chosen from a population.

In this paper we use a particular type of language game called the **evaluation game** [2]. The goal of the experiments is to observe the evolution of: (1) the communicative success[4]; (2) the internal grammars constructed by the individual agents; and (3) the external language used by the population. The main steps of *the evaluation game,* which is played by two agents (a *speaker* and a *hearer*), can be summarised as follows.

1. Conceptualisation. Firstly the speaker looks at one area of the white board and directs the attention of the hearer to the same area. The objects in that area constitute *the context* of the language game. Both speaker and hearer use their sensory channels to gather information about each object in the context and store that information so that they can use it in subsequent stages of the game. Then the speaker picks up a subset of the objects in the context which we will call *the topic* of the language game. The rest of the objects in the context constitute *the background.*

The speaker tries to find a unary or binary tuple of categories which distinguishes the topic from the background, i.e., a tuple of categories such that its evaluation on the topic is different from its evaluation on any object in the background. If the speaker cannot find a discriminating tuple of categories, the game fails. Otherwise it tries to find a logical category that is associated with the subset

[4] The *communicative success* is the average of successful language games in the last ten language games played by the agents.

of Boolean values or Boolean pairs resulting from evaluating the topic on that category tuple. If it does not have any logical category associated with this subset, it creates a new one. The formula constructed by applying this logical category to the discriminating category tuple constitutes a *conceptualisation* of the topic, because it *characterises the topic as the set of objects in the context which satisfy that formula.*

In general an agent can build several conceptualisations for the same topic. For example, if the context contains objects 1, 2 and 3 in figure 1, and the topic is the subset consisting of objects 1 and 2, the formulas [iff, up, left] and [xor, up, right] could be used as conceptualisations of the topic in an evaluation game.

2. Verbalising. The speaker chooses a conceptualisation (i.e., a discriminating formula) for the topic, generates a sentence that expresses this formula and communicates that sentence to the hearer. If the speaker can generate sentences for several conceptualisations of the topic, it tries to maximise the probability of being understood by other agents selecting the conceptualisation whose associated sentence has the highest score. The algorithm for computing the score of a sentence from the scores of the grammar rules used in its generation is explained in detail in [2].

The agents in the population start with an empty lexicon and grammar. Therefore they cannot generate sentences for most formulas at the early stages of a simulation run. In order to allow language to get off the ground, the agents are allowed to invent new sentences for those meanings 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. Once the agents can generate sentences for expressing a particular formula, they select the sentence with the highest score that verbalises a conceptualisation of the topic, and communicate that sentence to the hearer.

3. Interpretation. The hearer tries to interpret the sentence communicated by the speaker. If it can parse the sentence 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 cannot usually parse the sentences communicated by the speakers, since they have no prior linguistic knowledge. In this case the speaker points to the topic, the hearer conceptualises the topic using a logical formula, and adopts an association between that formula and the sentence used by the speaker. Notice that the conceptualisations of speaker and hearer may be different, because different formulas can be used to conceptualise the same topic.

At later stages of a simulation run it usually happens that the grammars and lexicons of speaker and hearer are not consistent, because each agent constructs its own grammar from the linguistic interactions it participates in, and it is very unlikely that speaker and hearer share the same history of linguistic interactions unless the population consists only of these two agents. In this case the hearer may be able to parse the sentence communicated by the speaker, but its interpretation of that sentence might be different from the meaning the speaker had in mind. The strategy used to coordinate the grammars of speaker and hearer when this happens is to decrease the score of the rules used by the speaker and the hearer in the processes of generation and parsing, respectively, and allow the hearer to adopt an association between its conceptualisation of the topic and the sentence used by the speaker.

Induction. Besides inventing expressions and adopting associations between sentences and meanings, the agents can 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 [5], although we have extended them so that they can be applied to grammar rules which have scores attached to them following the ideas of [9]. 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 experiments, which can be found in [2], we give an example of their application. We use Definite Clause Grammar for representing 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.

$$s(\text{light}, S) \rightarrow \text{clair}, \{S\ is\ 0.70\} \tag{1}$$

$$s(\text{right}, S) \rightarrow \text{droit}, \{S\ is\ 0.25\} \tag{2}$$

$$s([and, light, right], S) \rightarrow etclairdroit, \{S\ is\ 0.01\} \tag{3}$$

$$s([or, light, right], S) \rightarrow ouclairdroit, \{S\ is\ 0.01\} \tag{4}$$

The induction rule of **simplification**, applied to 3 and 2, allows generalising grammar rule 3 replacing it with 5. In this case *simplification* assumes that the second argument of the logical category *and* can be any meaning which can be expressed by a 'sentence', because according to rule 2 the syntactic category of the expression 'droit' is s (i.e., sentence).

$$s([and,light,B], S) \rightarrow etclair, s(B,R), \{S\ is\ R{\cdot}0.01\} \tag{5}$$

Simplification, applied to rules 5 and 1, can be used to generalise rule 5 again replacing it with 6. Rule 4 can be generalised as well replacing it with rule 7.

$$s([and,A,B], S) \rightarrow et, s(A,Q), s(B,R), \{S\ is\ Q{\cdot}R{\cdot}0.01\} \tag{6}$$

$$s([or,A,B], S) \rightarrow ou, s(A,Q), s(B,R), \{S\ is\ Q{\cdot}R{\cdot}0.01\} \tag{7}$$

The induction rule of **chunk** replaces a pair of grammar rules such as 6 and 7 by a single rule 8 which is more general, because it makes abstraction of their common structure introducing a syntactic category $c2$ for binary connectives. Rules 9 and 10 state that the expressions *et* and *ou* belong to the syntactic category $c2$.

$$s([C,A,B], S) \rightarrow c2(C,P), s(A,Q), s(B,R), \{S\ is\ P{\cdot}Q{\cdot}R{\cdot}0.01\} \tag{8}$$

$$c2(\text{and}, S) \rightarrow \text{et}, \{S\ is\ 0.01\} \tag{9}$$

$$c2(\text{or}, S) \rightarrow \text{ou}, \{S\ is\ 0.01\} \tag{10}$$

4. Coordination. The speaker points to the topic so that the hearer can identify the subset of objects it had in mind, and the hearer communicates the outcome of the evaluation game to the speaker. The game is successful if the hearer can parse the sentence communicated by the speaker, and its interpretation of that sentence identifies the topic (the subset of objects the speaker had in mind) correctly. Otherwise the evaluation game fails. 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 can *adapt their grammars* as well adjusting the scores of their grammar rules in order to communicate more successfully in future language games.

Coordination of the agents' grammars is necessary, because different agents can invent different expressions to refer to the same perceptually grounded or logical categories, and because the invention process uses 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 compound sentences in sufficiently similar ways as to avoid ambiguous interpretations.

The following **self-organisation mechanisms** help to coordinate the agents' grammars in such a way that they prefer using the grammar rules which are used more often by other agents [6,4].

We consider the case in which the speaker can generate a sentence to express the formula it has chosen as its conceptualisation of the topic using the rules in its grammar. If the speaker can generate several sentences to express that formula, it chooses the sentence with the highest score. The rest of the sentences are called *competing sentences*.

Suppose the hearer can interpret the sentence communicated by the speaker. If the hearer can obtain several formulas (meanings) for that sentence, the meaning with the highest score is selected. The rest of the meanings are called *competing meanings*.

If the topic identified by the hearer is the subset of objects the speaker had in mind, the evaluation game succeeds and both agents adjust the scores of the rules in their grammars. 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 the meaning which identified the topic the speaker had in mind and decreases the scores of the rules it used for obtaining competing meanings. This way the grammar rules that have been used successfully get reinforced, and the rules that 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 that have been used without success are inhibited.

The scores of grammar rules are *updated* as follows. The rule's original score S is replaced with the result of evaluating expression 11 if the score is *increased*, and expression 12 if the score is *decreased*.

$$minimum(1,\ S + 0.1) \tag{11}$$

$$maximum(0,\ S - 0.1) \tag{12}$$

4 Experiments

We describe the results of some experiments in which three agents try to construct at the same time a conceptualisation and a shared language which allow them to discriminate and communicate about subsets of objects pasted on a white board in front of them. In particular, the agents characterise such subsets of objects constructing logical formulas which are true for the objects in the subset and false for the rest of the objects in the context. Such formulas, which are communicated using a shared language, express facts about the relative spatial location and brightness of the objects in the subset with respect to the rest of the objects in the context. These experiments have been implemented using the Ciao Prolog system [10].

Figure 2 shows the evolution of the communicative success for a population of three agents. The *communicative success* is the average of successful language games in the last ten language games played by the agents. Firstly the agents play 700 evaluation games about subsets of objects which can be discriminated using only a single category or the negation of a perceptually grounded category. In this part of the simulation the population reaches a communicative success of 94% after playing 100 games and keeps it over that figure till the end of this part of the simulation. Next the agents play 6000 evaluation games about subsets of objects which require the use of perceptually grounded categories as well as unary and binary logical categories for their discrimination. In this part of the simulation the population reaches a communicative success of 100% after playing 3600 evaluation games and keeps it till the end of the second part of the simulation. The data shown in the figure correspond to the average of ten independent simulation runs with different random seeds.

We analyse now the conceptualisations and grammars built by the agents at the end of a particular simulation run. As we shall see the conceptualisations and grammars constructed by the individual agents are different, however they are compatible enough to guarantee the unambiguous communication of meanings of the same complexity as propositional logic formulas.

Table 1 shows the lexicon constructed by each agent in order to refer to perceptually grounded categories. We can observe that all the agents constructed the perceptually grounded categories (*up, down, right, left, light* and *dark*) and that all of them prefer the same expressions for referring to such categories.

We can observe in table 2 that all the agents constructed the logical category *not*. They all have a recursive grammar rule for expressing formulas constructed using negation and they use the same expression (*ci*) for referring to the logical

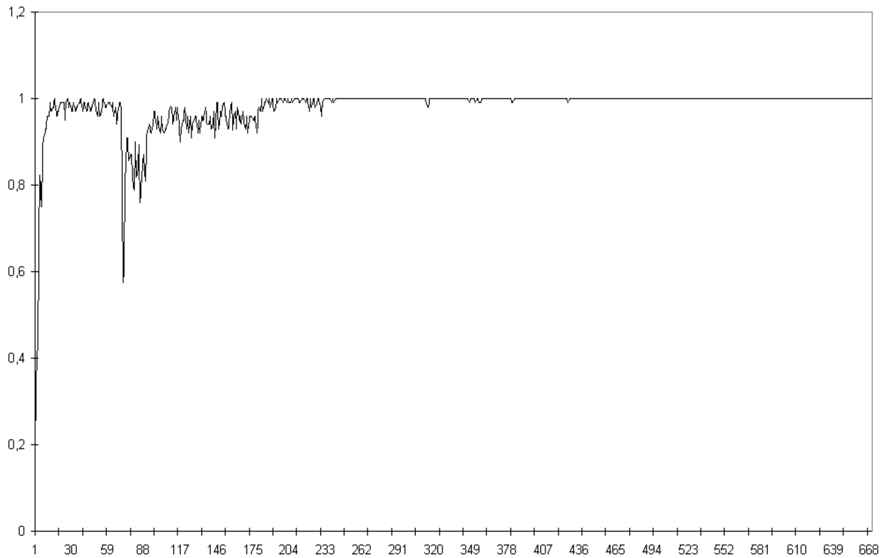

Fig. 2. Evolution of the communicative success for a population of three agents. Firstly the agents play 700 evaluation games which only require the use of perceptually grounded categories and negation for discrimination. Then they play 6000 evaluation games which require the use of all the perceptually grounded and logical categories for discrimination.

Table 1. Lexicon built by each agent to refer to perceptually grounded categories

Lexicon a1	**Lexicon a2**	**Lexicon a3**
s(up,1) → n	s(up,1) → n	s(up,1) → n
s(down,1) → b	s(down,1) → b	s(down,1) → b
s(right,1) → w	s(right,1) → w	s(right,1) → w
s(left,1) → dgq	s(left,1) → dgq	s(left,1) → dgq
s(light,1) → fdy	s(light,1) → fdy	s(light,1) → fdy
s(dark,1) → qyp	s(dark,1) → qyp	s(dark,1) → qyp

category *not*. There is a difference however: Agents a2 and a3 use a generic grammar rule based on a syntactic category for unary connectives, whereas agent a1 uses a specific grammar rule for expressing formulas constructed using negation.

We can also see in table 2 that all the agents constructed logical categories for all the **commutative connectives** (*and, nand, or, nor, xor* and *iff*), and that they use the same expressions (*ybd, d, j, sbr, wg* and *q*, respectively) for referring to such connectives.

Although in this particular simulation run all the agents use the same type of grammatical constructions to express formulas constructed using commutative connectives, this is not always the case. In a different simulation run agent a1

Table 2. Grammars built by the individual agents, including grammatical constructions, syntactic categories and lexicons for logical categories

Grammar a1
s([not,X],Q) → ci, s(X,P), {Q is P · 1}
s([X,Y,Z],T) → c1(X,P), s(Y,Q), s(Z,R), {T is P · Q · R · 1}
c1(and,R) → ybd, {R is 1}
c1(nor,R) → sbr, {R is 1}
c1(xor,R) → wg, {R is 1}
c1(iff,R) → q, {R is 1}
c1(if,R) → jdgq, {R is 1}
c1(or,R) → j, {R is 1}
s([X,Y,Z],T) → c2(X,P), s(Z,Q), s(Y,R), {T is P · Q · R · 1}
c2(noif,R) → oi, {R is 1}
c2(nand,R) → d, {R is 1}
Grammar a2
s([X,Y],R) → c1(X,P), s(Y,Q), {R is P · Q · 1}
c1(not,R) → ci, {R is 1}
s([X,Y,Z],T) → c2(X,P), s(Y,Q), s(Z,R), {T is P · Q · R · 1}
c2(nif,R) → oi, {R is 1}
c2(and,R) → ybd, {R is 1}
c2(nor,R) → sbr, {R is 1}
c2(xor,R) → wg, {R is 1}
c2(iff,R) → q, {R is 1}
c2(if,R) → jdgq, {R is 1}
c2(or,R) → j, {R is 1}
s([X,Y,Z],T) → c3(X,P), s(Z,Q), s(Y,R), {T is P · Q · R · 1}
c3(nand,R) → d, {R is 1}
Grammar a3
s([X,Y],R) → c1(X,P), s(Y,Q), {R is P · Q · 1}
c1(not,R) → ci, {R is 1}
s([X,Y,Z],T) → c2(X,P), s(Y,Q), s(Z,R), {T is P · Q · R · 1}
c2(nif,R) → oi, {R is 1}
c2(and,R) → ybd, {R is 1}
c2(nor,R) → sbr, {R is 1}
c2(xor,R) → wg, {R is 1}
c2(iff,R) → q, {R is 1}
c2(if,R) → jdgq, {R is 1}
c2(or,R) → j, {R is 1}
s([X,Y,Z],T) → c3(X,P), s(Z,Q), s(Y,R), {T is P · Q · R · 1}
c3(nand,R) → d, {R is 1}

used a grammar rule for expressing formulas constructed using *nor* (the negation of a disjunction) which placed the expression associated with the first argument of *nor* in the third position of the sentence, whereas agents a2 and a3 used a grammar rule which placed the same expression in the second position of the sentence. However, given that the expression associated with the connective of a logical formula is always placed in the first position of a sentence by the induction algorithm, the agents have no difficulty in understanding each other. Because the difference in the positions in the sentence of the expressions associated with the arguments of the connective can only generate an interpretation which corresponds to a formula which uses the same connective and which inverts the order of the arguments of such a connective with respect to the formula intended by the speaker. But such a formula is logically equivalent to the one intended by the speaker, because we are assuming that it is constructed using a commutative connective.

The results for **non-commutative connectives** are different however. All the agents constructed the logical category *if*, which corresponds to *implication*, and all of them use the same expression (*jdgq*) for referring to such a logical category. They also use the same grammatical construction for expressing implications, i.e., they all place the expression associated with the antecedent of an implication in the second position of the sentence, and the expression associated with the consequent in the third position.

Agents a2 and a3 constructed the logical category *nif*, whereas agent a1 does not have a grammar rule for expressing such a logical category. Instead of that, agent a1 constructed the logical category *noif* and a grammar rule that allows it to understand correctly the sentences generated by a2 and a3 in order to communicate formulas of the form [*nif*, A, B]. That is, whenever a2 and a3 try to communicate a formula of the form [*nif*, A, B], i.e., $\neg(A \rightarrow B)$, they use the grammar rules

s([X,Y,Z],T) → c2(X,P), s(Y,Q), s(Z,R), {T is P · Q · R · 1}
  c2(nif,R) → oi, {R is 1}

to generate a sentence. This sentence is parsed by a1 using the grammar rules

s([X,Y,Z],T) → c2(X,P), s(Z,Q), s(Y,R), {T is P · Q · R · 1}
  c2(noif,R) → oi, {R is 1}

interpreting the formula [*noif*, B, A], i.e., $\neg(B \leftarrow A)$, which is logically equivalent to the formula intended by the speaker. This is so because the grammar rules used by a1 not only use the same expression for referring to the logical connective *noif* than a2 and a3 for referring to *nif*, but they also reverse the order of the expressions associated with the arguments of the connective in the sentence.

On the other hand, given that the formulas [*nif*, A, B] and [*noif*, B, A] are logically equivalent, agent a1 will not be prevented from characterising any subset of objects because of the lack of the logical category *nif*. Since it will always prefer to conceptualise the topic using the second formula. The same holds for agents a2 and a3 with respect to the logical category *noif*.

Finally none of the agents constructed the logical category *oif* nor grammar rules for expressing formulas constructed using such a logical category. But this does not prevent them from characterising any subset of objects, because [*oif*, A, B] is logically equivalent to [*if*, B, A] and all the agents have grammar rules for expressing implications.

5 Intuitive Reasoning

During the processes of conceptualisation and grounded language acquisition the agents build categorisers for perceptually grounded categories (such as *up, down, right, left, light* and *dark*) and for logical categories (*and, nand, or, nor, if, nif, oif, noif, iff* or *xor*). These categorisers allow them to evaluate logical formulas constructed from perceptually grounded categories.

Intuitive reasoning is a process by which the agents discover relationships that hold among the categorisers of perceptually grounded categories and logical categories. For example, an agent may discover that the formula $up \rightarrow \neg down$ is always true, because the categoriser of *down* returns false for a given object whenever the categoriser of *up* returns true for the same object.

It may work as a process of constraint satisfaction in natural agents, by which they try to discover whether there is any combination of values of their sensory channels that satisfies a given formula. It is not clear to us how this process of constraint satisfaction can be implemented in natural agents. It may be the result of a simulation process by which the agents generate possible combinations of values of their sensory channels and check whether they satisfy a given formula. Or it may be grounded on the impossibility of firing simultaneously some categorisers due to the way they are implemented by physically connected neural networks.

In particular intuitive reasoning can be used to perform the following inference tasks which constitute the basis of the logical approach to formalising common sense knowledge and reasoning [11].

1. Using the categorisers of logical categories an agent can determine whether a given formula is a *tautology* (it is always true because of the meaning of its logical symbols) or an *inconsistency* (it is always false for the same reason).
2. Using the categorisers of logical and perceptually grounded categories an agent can discover that a given formula is a *common sense axiom*, i.e., it is always true because of the meaning of the perceptually grounded categories it involves. The formula $up \rightarrow \neg down$, discussed above, is a good example of a common sense axiom. Similarly it can discover that a particular formula (such as $up \wedge down$) is always false, because of the meaning of categories it involves. It can determine as well that certain formulas (such as $up \leftrightarrow left$) are merely *satisfiable,* but that they are not true under all circumstances.
3. Finally the categorisers of logical and perceptually grounded categories can be used as well to discover *domain dependent axioms.* These are logical formulas that are not necessarily true, but which always hold in the particular domain of knowledge or environment the agent interacts with during its development history. This is the case of formula $up \wedge light \rightarrow left$, which is not necessarily true, but it is always true for every subset of objects of the white board shown in figure 1.

The process of determining whether a formula is a tautology, an inconsistency or a common sense axiom by intuitive reasoning can be implemented using constraint satisfaction algorithms, if the behaviour of the categorisers of perceptually grounded and logical categories can be described by constraints. It can also be proved that intuitive reasoning is closed under the operator of *logical consequence* if the behaviour of the categorisers of perceptually grounded categories can be described by linear constraints. That is, if a formula is a logical consequence of a number of common sense axioms which can be discovered using intuitive reasoning, it must also be possible to prove that such a formula is always true using intuitive reasoning. This is a consequence of the fact that the linear constraints describing the behaviour of the categorisers of perceptually grounded categories constitute a logical model, in the sense of model theory semantics [11], of the set of common sense axioms that can be discovered using intuitive reasoning.

6 Conclusions

We have described some experiments which simulate a grounded approach to language acquisition, in which a population of autonomous agents without prior linguistic knowledge tries to construct at the same time a conceptualisation of its environment and a shared language.

These experiments extend previous work by using a linguistic interaction (*the evaluation game*) in which the agents must first conceptualise the *topic* (a subset of objects) using the mechanisms proposed in [1] for the acquisition of logical categories, and then construct a *shared language* (a lexicon and a grammar) using the invention, adoption, induction and self-organisation mechanisms proposed in [2] for the acquisition of the syntax of propositional logic.

The results of the experiments show that at the end of the simulation runs the agents build different conceptualisations and different grammars. However these conceptualisations and grammars are compatible enough to guarantee the unambiguous communication of meanings of the same complexity as propositional logic formulas.

We have also seen that the categorisers of the perceptually grounded and logical categories built during the conceptualisation and language acquisition processes can be used for some forms of common sense reasoning, such as determining whether a sentence is a tautology, a contradiction, a common sense axiom or a merely satisfiable formula – all this in a very restricted domain. However this form of intuitive reasoning requires the agents to be conscious of the fact that they use certain categorisers and of the behaviour of such categorisers.

References

1. Sierra, J.: Grounded models as a basis for intuitive and deductive reasoning: The acquisition of logical categories. In: Proceedings of the European Conference on Artificial Intelligence, pp. 93–97. IOS Press, Amsterdam (2002)
2. Sierra, J., Santibáñez, J.: The acquisition of linguistic competence for communicating propositional logic sentences. In: Artikis, A., O'Hare, G.M.P., Stathis, K., Vouros, G.A. (eds.) ESAW 2007. LNCS (LNAI), vol. 4995, pp. 175–192. Springer, Heidelberg (2008)

3. Steels, L.: The origins of syntax in visually grounded robotic agents. Artificial Intelligence 103(1-2), 133–156 (1998)
4. Batali, J.: The negotiation and acquisition of recursive grammars as a result of competition among exemplars. In: Linguistic Evolution through Language Acquisition, pp. 111–172. Cambridge University Press, Cambridge (2002)
5. Kirby, S.: Learning, bottlenecks and the evolution of recursive syntax. In: Linguistic Evolution through Language Acquisition: Formal and Computational Models, pp. 96–109. Cambridge University Press, Cambridge (2002)
6. Steels, L.: The Talking Heads Experiment. Special Pre-edition for LABORATORIUM. Antwerpen (1999)
7. Harnad, S.: The symbol grounding problem. Physica D (42), 335–346 (1990)
8. Wittgenstein, L.: Philosophical Investigations. Macmillan, New York (1953)
9. Vogt, P.: The emergence of compositional structures in perceptually grounded language games. Artificial Intelligence 167(1-2), 206–242 (2005)
10. Bueno, F., Cabeza, D., Carro, M., Hermenegildo, M., López-García, P., Puebla, G.: The Ciao Prolog system. reference manual. Technical Report CLIP3/97.1, School of Computer Science, Technical University of Madrid, UPM (1997), `http://www.clip.dia.fi.upm.es/`
11. McCarthy, J.: Formalizing Common Sense. Papers by John McCarthy. Edited by Vladimir Lifschitz. Ablex, Greenwich (1990)