Grounded Language Acquisition Enables Intuitive Reasoning

Josefina Sierra
Software Department
Technical University of Catalonia, Spain
jsierra@lsi.upc.edu*

Josefina Santibáñez
Education Department
University of La Rioja, Spain
josefina.santibanez@unirioja.es

Abstract

We describe an experiment which simulates a grounded approach to language acquisition in a population of autonomous agents without prior linguistic knowledge. The idea is to let the agents acquire at the same time a conceptualization of their environment and a number of linguistic conventions (i.e., a shared lexicon and a set of grammar rules) which allow them to express facts about their environment in a way that could be understood by other agents in the population.

The approach used to simulate the conceptualization and the language acquisition processes in each individual agent is based on general purpose cognitive capacities, such as visual perception, categorization, discrimination, evaluation, invention, adoption and induction. The emergence of a shared language in the population, and therefore the acquisition of a common set of linguistic conventions by the individual agents, results from a process of self-organization of a particular type of linguistic interaction, known as a *language game*, that takes place among the agents in the population.

By letting the agents acquire a grounded semantics at the same time they jointly construct a shared communication language we allow them not only to communicate facts about their environment, but to understand as well the meanings of such facts in an intuitive way. This enables the agents to reason about such facts in terms of their intuitive understanding of the properties and relationships stated by those facts about the objects in their environment.

Conceptualization: Basic Definitions

We consider the experimental setting proposed in *The Talking Heads Experiment* (Steels 1999): 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 conceptualize the perceptual information they obtain by looking at a white board and trying to characterize 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 that 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 experiment 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).

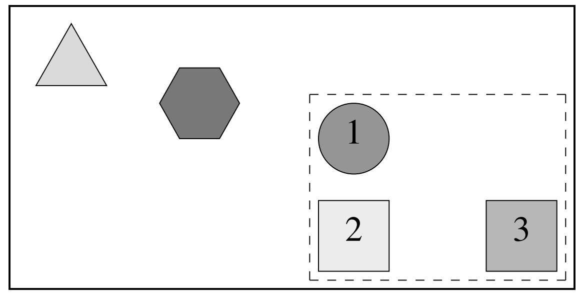

Figure 1: The area of the white board captured by the agents cameras is the lower right rectangle.

Perceptually Grounded Categories The data returned by the sensory channels are values from a continuous domain. To be the basis of a natural language conceptualization, these values must be transformed into a discrete domain. One form of categorization consists in dividing up each domain of output values of a particular sensory channel into regions and assigning a *category* to each region (Steels 1999). 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 categorized as [right].

Perceptually Grounded Categorizers At the same time the agents build categories in order to conceptualize perceptual information, they construct cognitive procedures (called *categorizers*) that allow them to check whether these categories hold or not for a given object.

*This work is partially funded by the DGICYT TIN2005-08832-C03-03 project (MOISES-BAR).

Categorizers give grounded meanings (Harnad 1990) to categories (i.e., symbolic representations) by establishing explicit connections between them and reality (perceptual input as processed by sensory channels). These connections are learned through language games (Wittgenstein 1953; Steels 1999), and allow the agents to check whether a category holds or not for a given object. Most importantly they provide information on the perceptual and cognitive processes the agent must go through in order to evaluate a given category.

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

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 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 procedure capable of finding the categorizers of a tuple of categories, applying them to an object, and observing their output. If $\vec{c} = (c_1, \ldots, c_n)$ is a category tuple and o is an object, $E(\vec{c}, o)$ is a tuple of Boolean values $(v_1, \ldots, v_n)$, where each v_i is the result of applying c_i^C (the categorizer of c_i) to 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.

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 {1,0}, 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 if the result of evaluating the pair of categories (c_1, c_2) is a Boolean pair which belongs to the subset of Boolean pairs $\{(1, 1), (1, 0), (0, 1)\}$.

The 16 Boolean functions of two arguments that can be constructed using this method are summarized by the following 10 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 area. The representation of logical categories as subsets of Boolean tuples is equivalent to the *truth tables* used for describing 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 propositional formulas, 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 Categorizers The categorizers of logical categories are cognitive procedures 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 behavior of their categorizers can be described therefore by constraints of the form $E(\vec{c}, o) \in S_l$, 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, $\vec{c}$ is a tuple of categories, and o is an object. For example, the constraint $E((c1, c2), o) \in \{(1, 1)\}$ describes the behavior of the categorizer for the logical category [and] (i.e., $c1 \wedge c2$).

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

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.

Grounded Language Acquisition

Language acquisition is seen as a collective process by which a population of autonomous agents constructs a

[1]We use the notation $[\text{cat}]^C$ to refer to the categorizer which is capable of recognizing whether category [cat] holds or not.

[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.

shared language that 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** (Sierra 2006). The goal of the experiment 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 four main steps of the *evaluation game,* which is played by two agents (a *speaker* and a *hearer*), can be summarized as follows.

1. Conceptualization 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 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 *conceptualization* of the topic, because it *characterizes the topic as the set of objects in the context that satisfy that formula.*

In general an agent can generate several conceptualizations 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 conceptualizations of the topic in an evaluation game.

2. Lexicalization The speaker chooses a conceptualization (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 conceptualizations of the topic, it tries to maximize the probability of being understood by other agents selecting the conceptualization whose lexicalization (i.e., 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 (Sierra 2006).

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, they 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 lexicalizes a conceptualization 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 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, the hearer conceptualizes the topic using a logical formula, and adopts an association between that formula and the sentence used by the speaker. Notice that the conceptualizations of speaker and hearer may be different, because different formulas can be used to conceptualize 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 in which it participates, 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 generated 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 speaker and hearer in the processes of generation and parsing, respectively, and allow the hearer to adopt an association between its conceptualization of the topic and the sentence used by the speaker.

Induction Besides inventing and adopting associations between sentences and meanings, the agents can use some *induction mechanisms* to extract generalizations from the grammar rules they have learnt so far. The induction mechanisms used in this paper are based on the rules for *simplification and chunk* in (Kirby 2002), although we have extended them so that they can be applied to grammar rules which have scores. 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 (Sierra 2006), 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

[4]The *communicative success* is the average of successful language games in the last 10 language games played by the agents.

rule in previous communication. Suppose an agent's grammar contains the following rules.

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

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

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

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

The induction rule of **simplification**, applied to 3 and 2, allows generalizing grammar rule 3 replacing it with 5.

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

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

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

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

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 *and* and *or* belong to the syntactic category $c2$.

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

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

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

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 for referring to the same perceptually grounded and logical 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 compound sentences in sufficiently similar ways as to avoid ambiguous interpretations.

The following **self-organization principles** help to coordinate the agents' grammars in such a way that the agents prefer using the grammar rules that are used more often by other agents (Steels 1999; Batali 2002).

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

Lexicon a1	**Lexicon a2**	**Lexicon a3**
s(up,1)→m	s(up,1)→m	s(up,1)→m
s(down,1)→va	s(down,1)→va	s(down,1)→va
s(right,1)→j	s(right,1)→j	s(right,1)→j
s(left,1)→sqj	s(left,1)→sqj	s(left,1)→sqj
s(light,1)→et	s(light,1)→et	s(light,1)→et
s(dark,1)→fh	s(dark,1)→fh	s(dark,1)→fh

Table 1: Expressions for perceptually grounded categories.

Suppose the hearer can interpret the sentence communicated by the speaker. If the hearer can obtain several meanings (formulas) 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 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 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 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*.

$$minimum(1,\ S + 0.1) \quad (11)$$

$$maximum(0,\ S - 0.1) \quad (12)$$

Experiments

We describe the results of some experiments in which three agents construct a conceptualization and a shared language that allows them to discriminate and communicate about subsets of the set of objects pasted on a white board in front of them. In particular, the agents characterize such subjects 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 a subset with respect to the rest of the objects in the context. These experiments have been implemented using the Ciao Prolog system (Bueno *et al.* 1997) .

Firstly the agents played 1000 evaluation games in which the Topic and the Context consisted of a single object. These games allowed them to construct perceptually grounded categories for *up, down, right, left, light* and *dark,* and a shared

Grammar a1
s([not,X],Q) → vu, 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) → o,{R is 1}
c1(xor,R) → qs,{R is 1}
c1(or,R) → e,{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) → fi,{R is 1}
c2(nor,R) → w,{R is 1}
c2(iff,R) → f,{R is 1}
c2(nand,R) → b,{R is 1}
c2(if,R) → wh,{R is 1}
Grammar a2
s([X,Y],R) → c1(X,P),s(Y,Q),{R is P*Q*1}
c1(not,R) → vu,{R is 1}
s([X,Y,Z],T) → c2(X,P),s(Y,Q),s(Z,R),{T is P*Q*R*1}
c2(and,R) → o,{R is 1}
c2(nif,R) → fi,{R is 1}
c2(or,R) → e,{R is 1}
s([X,Y,Z],T) → c3(X,P),s(Z,Q),s(Y,R),{T is P*Q*R*1}
c3(nor,R) → w,{R is 1}
c3(xor,R) → qs,{R is 1}
c3(iff,R) → f,{R is 1}
c3(nand,R) → b,{R is 1}
c3(if,R) → wh,{R is 1}
Grammar a3
s([X,Y],R) → c1(X,P),s(Y,Q),{R is P*Q*1}
c1(not,R) → vu,{R is 1}
s([X,Y,Z],T) → c3(X,P),s(Y,Q),s(Z,R),{T is P*Q*R*1}
c3(and,R) → o,{R is 1}
c3(nif,R) → fi,{R is 1}
c3(or,R) → e,{R is 1}
s([X,Y,Z],T) → c2(X,P),s(Z,Q),s(Y,R),{T is P*Q*R*1}
c2(nor,R) → w,{R is 1}
c2(xor,R) → qs,{R is 1}
c2(iff,R) → f,{R is 1}
c2(nand,R) → b,{R is 1}
c2(if,R) → wh,{R is 1}

Table 2: Logical categories, grammatical constructions and syntactic categories constructed by each individual agent.

vocabulary for referring to them. Then they played 5000 evaluation games in which the Topic and the Context could have one, two or three objects. These games allowed them to construct logical categories, a shared lexicon for referring to them, syntactic categories and grammatical constructions for expressing logical formulas. Tables 1 and 2 show the lexicons and grammars constructed by the individual agents.

Table 1 shows that all the agents have constructed categories for *up, down, right, left, light* and *dark*, and that all agents prefer the same expressions for referring to them.

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

All the agents have constructed logical categories for all **commutative connectives**: *and, nand, or, nor, xor* and *iff*. They also use the same expressions for referring to such connectives (*o, b, e, w, qs* and *f*, respectively). They do not agree however in the type of grammatical construction they use to express formulas constructed with such connectives. For example, agent a1 places the expression associated with the first argument of an exclusive disjunction (*xor*) in the second position of the sentence, whereas agents a2 and a3 place the same expression in the third 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 will have no difficulty in understanding each other. Because the difference in the positions of the expressions associated with the arguments of the connective in the sentence can only generate an interpretation which corresponds to a formula that 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 will be 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 have constructed the logical category *if*, which corresponds to *implication*, and all use the same expression (*wh*) for referring to such a connective. 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 third position of the sentence, and the expression associated with the consequent in the second position.

Agents a2 and a3 have constructed the logical category *nif*, whereas agent a1 does not have a grammar rule for expressing such a logical category. This is compensated by the fact that agent a1 has 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)→c3(X,P),s(Y,Q),s(Z,R),{T is P*Q*R*1}
c3(nif,R)→fi,{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}
c1(noif,R)→fi,{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 characterizing any subset of objects by the lack of the logical category *nif*. Because it will always prefer to conceptualize the topic using the second formula.

Finally, none of the agents has grammar rules for expressing formulas constructed using the logical category *oif*. But this does not prevent them from expressing meanings of the form [*oif*, A, B], because these meanings are logically equivalent to formulas of the form [*if*, B, A] and all the agents have grammar rules for expressing implications.

Intuitive Reasoning

During the process of grounded language acquisition the agents built categorizers 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 categorizers 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 categorizers 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 categorizer of $down$ returns false for a given object whenever the categorizer 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 for their sensory channels and check whether they satisfy a given formula. Or it may be grounded on the impossibility of firing simultaneously some categorizers 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 that constitute the basis of the logical approach to formalizing common sense knowledge and reasoning (McCarthy 1990).

1. Using the categorizers 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 categorizers 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 that 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 categorizers 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 that it is always true for every subset of objects of the white board shown in figure 1.

When the behavior of the categorizers of perceptually grounded and logical categories can be described by constraints, 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. It can also be proved that intuitive reasoning is closed under the operator of *logical consequence* when the behavior of the categorizers 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 that can be discovered using intuitive reasoning, it must be possible to show that such a formula is always true using intuitive reasoning as well. This is a consequence of the fact that the linear constraints describing the behavior of the categorizers of perceptually grounded categories constitute a logical model, in the sense of model theory semantics (McCarthy 1990), of the set of common sense axioms that can be discovered using intuitive reasoning.

Conclusions

We have described an experiment which simulates a grounded approach to language acquisition in a population of autonomous agents without prior linguistic knowledge. The conceptualization and the language acquisition processes in each individual agent are based on general purpose cognitive capacities, such categorization, discrimination, evaluation and induction. The emergence of a shared communication language results from a process of self-organization of the linguistic interactions (language games) that take place among the agents in the population.

By letting the agents acquire a grounded semantics at the same time they jointly construct a shared communication language we allowed them not only to communicate facts about their environment, but to understand the meanings of such facts and to reason about them in an intuitive way.

References

Batali, J. 2002. The negotiation and acquisition of recursive grammars as a result of competition among exemplars. In *Linguistic Evolution through Language Acquisition*, 111–172. Cambridge U.P.

Bueno, F.; Cabeza, D.; Carro, M.; Hermenegildo, M.; López-García, P.; and Puebla, G. 1997. The Ciao prolog system. Reference manual. Technical Report CLIP3/97.1, School of Computer Science, Technical University of Madrid (UPM). Available from http://www.clip.dia.fi.upm.es/.

Harnad, S. 1990. The symbol grounding problem. *Physica D* (42):335–346.

Kirby, S. 2002. Learning, bottlenecks and the evolution of recursive syntax. In *Linguistic Evolution through Language Acquisition: Formal and Computational Models*, 96–109. Cambridge University Press.

McCarthy, J. 1990. *Formalizing Common Sense. Papers by John McCarthy*. Ablex. Edited by Vladimir Lifschitz.

Sierra, J. 2006. Propositional logic syntax acquisition. In *Symbol Grounding and Beyond*, 128–142. Lecture Notes in Computer Science, volume 4211.

Steels, L. 1999. *The Talking Heads Experiment.* Antwerpen: Special Pre-edition for LABORATORIUM.

Wittgenstein, L. 1953. *Philosophical Investigations*. New York: Macmillan.