Submitted to *Journal of Italian Linguistics*

Verbs, nouns, and simulated language games

Domenico Parisi

Institute of Cognitive Science and Technology
National Research Council
parisi@ip.rm.cnr.it

Angelo Cangelosi

Centre for Neural and Adaptive Systems
School of Computing
University of Plymouth
acangelosi@soc.plym.ac.uk

Ilaria Falcetta

University of Rome La Sapienza
ilariafalcetta@libero.it

Abstract

The paper describes some simple computer simulations that implement Wittgenstein’s notion of a language game, where the meaning of a linguistic signal is the role played by the linguistic signal in the individual’s interactions with the nonlinguistic and linguistic environment. In the simulations an artificial organism interacts at the sensory-motor level with an environment and its behavior is influenced by the linguistic signals the individual receives from the environment (conspecifics). Using this approach we try to capture the distinction between (proto)verbs and (proto)nouns, where (proto)verbs are linguistic signals that tend to co-vary with the action with which the organism must respond to the sensory input whereas (proto)nouns are linguistic signals that tend to co-vary with the particular sensory input to which the organism must respond with its actions. Some extensions

of the approach to the analysis of other parts of speech ((proto)adjectives, (proto)sentences, etc.) are also described. The paper ends up with some open questions and suggestions on how to deal with them.

1. Simulated language games

The meaning of a linguistic signal is the manner in which the linguistic signal is used in the everyday interactions of speakers/hearers with the world and the role the linguistic signal plays in their overall behavior. This Wittgensteinian definition of meaning, while probably correct, poses a serious problem for the study of language in that, although linguistic signals as sounds or visual (written) forms are easily identified, observed, and described, the way in which linguistic signals are used by actual speakers/hearers in real life situations is very difficult to observe and describe with any precision, reliability, and completeness. Therefore, linguists, psycholinguists, and philosophers tend to replace meanings with such poor “proxies” as verbal definitions, translations (when studying linguistic signals in other languages), or the limited and very artificial uses of linguistic signals in laboratory experiments (e.g., the naming of pictures or the decision if a sequence of letters is a word or a nonword).

An alternative to such practices is to adopt Wittgenstein’s strategy of studying “language games”, i.e., simplified models of the very complex and diverse roles that linguistic signals play in our complicated everyday language which may be closer to the “games by means of which children learn their native language” (Wittgenstein 1953, 5e) and to languages “more primitive than ours” (Wittgenstin 1953, 3e). In this paper we adopt this Wittgensteinian strategy but with a significant change: our language games are simulated in a computer. We create artificial organisms which live in artificial worlds and which may receive and produce linguistic signals in such a way that these linguistic signals become incorporated in their overall behavior and in their interactions with the world. Simulated language games have two advantages when they are compared with the philosopher's language games. First, since simulated language games are "objectified" in the computer (the organisms' behavior can be actually seen on the computer screen) and they do not only exist in the philosopher's mind or in his/her verbal expressions and discussions with colleagues, they offer more degrees of freedom and more objectivity when one tries to describe, analyze, measure, and manipulate experimentally the meaning of linguistic signals conceived as their role in the overall behavior of the artificial organisms. Second, given the great memory and computing resources of the computer, which greatly execeed those of the human mind, one can progressively add new components to an initially very simple simulation in such a way that the language games may become more and more similar to actual languages.

Recently, computer models have been used to simulate the evolutionary emergence of language in populations of interacting organisms (Cangelosi & Parisi 2002; Knight *et al*. 2000; Steels 1997). Various simulation methodologies have been employed, such as communication between rule-based agents (Kirby 1999), recurrent neural networks (Batali 1994; Ellefson & Christiansen 2000), robotics (Kaplan, 2000; Steels & Vogt 1997), and internet agents (Steels & Kaplan 1999). Among these, artificial life neural networks (ALNNs: Parisi 1997) provide a useful modelling approach for studying language (Cangelosi & Parisi 1998; Cangelosi & Harnad in press; Parisi & Cangelosi 2002). ALNNs are neural networks that control the behaviour of organisms that live in an environment and are members of evolving populations of organisms. They provide a unifying methodological and theoretical framework for cognitive modelling because of the use of both evolutionary and connectionist techniques and the interaction of the organisms with a simulated ecology. All behavioral abilities (e.g., sensorimotor skills, perception, categorization, language) are controlled by the same neural network. This permits the investigation of the interaction between language and other cognitive and sensorimotor abilities.

2. Verbs and nouns

Among linguistic signals such as words one can distinguish among different classes of words based on some general properties of the use of these different classes of words (Brown & Miller 1999). The purpose of this article is to explore what neural network models can contribute to a better understanding of the nature of verbs and nouns and, possibly, other parts of speech. The distinction between verbs and nouns is perhaps the most basic and universal distinction among different classes of words in human languages and a neural network treatment of verbs and nouns, if successful, can then be extended to other parts of speech. Verbs and nouns may be distinguished on semantic or syntactic grounds. Semantically, verbs and nouns can be distinguished in terms of the different types of entities to which they refer. Verbs are said to refer to actions or processes while nouns refer to objects or static entities (cf., e.g., Langacker 1987). Syntactically, verbs and nouns are distinguished in terms of the different roles they play, or the different contexts in which they appear, in phrases and sentences. Given our simplified language games, in which almost no multi-component signals are used such as phrases and sentences, the work to be reported here tries to illuminate the semantics rather than the syntax of verbs and nouns.

We hypothesize that in the early stages of language acquisition in children, and perhaps also in the early stages of linguistic evolution in the lineage of Homo sapiens, words begin to differentiate into verbs and nouns with verbs referring to actions and nouns to objects. But what does it mean to refer to actions or to objects and, more generally, what it is for a word to refer? Heard sounds acquire meaning or reference (we use the two terms interchangeably) for an organism and therefore become linguistic signals for the organism when they influence the way in which the organism responds to the input from the environment. We imagine a basic situation in which the organism is exposed to visual input from the environment and the organism responds to this visual input with some motor action. Heard sounds are additional inputs to the organism which are physically produced by the phono-articulatory behavior of some nearby conspecific. If this additional input systematically influences how the organism responds to the visual input, with specific sounds having specific influences on the organism's behavior, we say that the sounds have become linguistic signals which have meaning or reference.

Our organisms see objects in the environment and they respond by moving their (single) arm in order to execute some action with respect to the objects. An organism's behavior is controlled by the organism's nervous system which is modeled using an artificial neural network. The neural network has two distinct sets of input units (sensory receptors). One set of input units encodes the content of the organism's retina (visual input). The other set of input units encodes the current position of the organism's arm (proprioceptive input). The network's output units encode muscle movements which result in changes in the arm's position. Intermediate between the input and the output units there are one or more layers of hidden units. All the network's units encode information in terms of the quantitative state of activation of the units. The neural network functions as a succession of input/output cycles of activity. In each cycle the pattern of activation of the input units is transformed into the patterns of activation of the successive layers of hidden units by the connection weights linking one unit to the next one until an output pattern of activation is generated which results in a micro-movement of the arm. A succession of micro-movements is an action of the organism with respect to the visually perceived objects. The organism may see a single object at a time or two objects at the same time and it may respond by moving its arm to reach an object or to push the object away from itself or to pull it toward itself.

Now we add language. Imagine that the organism's neural network includes a third set of input units which may encode various sounds (auditory input). These heard sounds tend to influence the way in which the organism responds to the visual input. When the organism hears one particular sound it responds to the visual input with some particular action which may be different (although it need not be) from the action with which the organism would have responded to that input in the absence of the sound (including no action at all). When a different sound is heard by the organism, the organism may respond with a different action.

We will describe a number of simple situations in which linguistic signals acquire their meaning in that they become part of the organism's total experience in its environment.

Imagine the following language game (Cangelosi & Parisi 2001; Parisi & Cangelosi 2002). The life of the organism is divided up into episodes which are composed of a number of successive input/output cycles. In each episode the organism sees one of two objects, O1 and O2, which vary in their shape. Together with this visual input the organism receives an auditory input, a heard sound presumably pronounced by some conspecific located nearby in the organism's environment. There are only two possible sounds, S1 and S2, but in any given episode the organism hears only one of these two sounds. At the beginning of each episode the endpoint of the organism's arm (the hand) is already positioned on the object. If we observe the organism's behavior, we see that the organism responds to the visually perceived object by pushing the object away from itself if it hears the sound S1 and by pulling the object toward itself if it hears the sound S2. This happens independently from whether the object is O1 or O2. In these circumstances, we say that the two sounds which are heard by the organism are (proto)verbs. (In fact they have a meaning which is equivalent to the meaning of the English verbs "push" and pull".) S1 and S2 co-vary with the action with which the organism responds to the visual input but they are indifferent to the content of the visual input, i.e., to whether the object which is seen and which is pushed or pulled is O1 or O2.

Imagine now another language game (Falcetta 2001). The organism sees both objects, O1 and O2, at the same time. The two objects are located one in the left half and one in the right half of the organism's visual field. Together with this visual input the organism hears one of two sounds, S3 and S4. At the beginning of each episode the organism's arm is in a randomly selected position but always away from the objects. (Notice that the organism does not see its arm. It is informed by the proprioceptive input about the arm's current position but it only sees the objects.) When the organism hears S3 it moves its arm and reaches object O1 whereas when it hears S4 it reaches object O2. In these circumstances, we say that the two sounds S3 and S4 are (proto)nouns.

Notice that, like S1 and S2, S3 and S4 influence the action produced by the organism. Assuming that in a given episode the object O1 is in the left hemifield and the object O2 in the right hemifield, if the organism hears S3 it moves its arm toward the left portion of the visual field and reaches the object which is there (O1) whereas if it hears S4 it moves the arm toward the right portion of the visual field and reaches O2. However, in this second language game the linguistic input has a different role in the overall experience of the organism. While in the first language game the two linguistic signals, S1 and S2, had the role of determining the particular action executed by the organism, pushing or pulling, independently from whether the object was O1 or O2, in this new language game there is a single action, reaching an object, and the two linguistic signals, S3 and S4, have the role of directing the action of the organism toward one particular object rather than toward the other.

Therefore, we characterize verbs as linguistic signals that co-vary with the actions of the organism whereas nouns are linguistic signals that co-vary with the particular objects which are involved in these actions.

Since in the second language game the organism is capable of only one action, i.e., reaching an object with its arm, there is no need for the language to specify which action to choose - which is the role of verbs. The organism has only to know which one of the currently perceived objects must be reached, and providing this information is the role of nouns. But consider a third, somewhat more complex, language game in which the organism is both capable of two distinct actions, pushing and pulling objects (as in our first language game) and it sees two different objects at the same time (as in our second language game). In the new language game the organism will need to hear two linguistic signals, one verb and one noun, in order to know what to do. The auditory input units will encode one of the two verbs S1 and S2 at time T0 and then one of the two nouns S3 and S4 at time T1, or viceversa. (In this language game the temporal order of the two words in each sequence is irrelevant but, whatever the temporal order, to be able to appropriately process this simple (proto)sentence the neural network will need a working memory which keeps a trace of the first word while hearing the second word.) In general, to have a (proto)sentence, one portion of the heard sounds must co-vary with the action to be executed and the other portion with the object on which the action is to be executed. Since actions can be executed on more than a single object (e.g., the action of giving involves two objects: the object given and the person receiving the object), (proto)sentences may include more than a single noun. (For the emergence of subjects or agents, cf. the last section. For the evolutionary emergence of compositionality, cf. Cangelosi 2001.)

We have defined nouns in terms of their role in directing the organism's action toward particular objects. Consider, however, that the organism’s action can also consist in what is called “overt attention”, i.e., movements of the organism’s eyes or head that allow the organism to visually access some particular object - the object which is specified by the noun. Normally organisms see many different objects at the same time and by hearing a noun they select one particular object as the object which is to be involved in the organism’s action while ignoring the other objects. However, in other cases the organism hears some particular noun without seeing the object which is indicated by the noun. In these circumstances the noun causes the organism to move its entire body (locomoting) or particular parts of its body (turning the head or the eyes) until it finds an object with the required properties and it can execute the expected action on the object.

To illustrate this role of nouns let us consider a fourth language game. The organism’s visual field is divided into three parts: a central portion with better seeing capabilities (fovea) and two peripheral portions, on the left and on the right of the central portion, with less good vision. The neural network which controls the organism’s behavior has two sets of output (motor) units, not just a single set as in the preceding language games. One set of motor units controls the organism’s arm, as in our previous simulations, while the second set of motor units controls the movements of the organism’s (single) eye. At the beginning of each episode the organism looks straight ahead but it can move its eye either to the right or to the left. In every episode the organism’s visual field contains three objects with different shapes, O3, O4, and O5, which are randomly distributed one in the visual field’s central portion and each of the other two in one of the two peripheral portions. Notice, however, that the organism can recognize the shape of an object if the object is located in the central fovea but not if it is located in the peripheral portions of the visual field.

The organism is capable of only one action using its arm: reaching an object. Hence, we don’t need verbs in this language game. In each episode the organism hears one of three linguistic signals (nouns): S3, S4, and S5. If the organism hears the linguistic signal S3 and the object O3 is in the fovea, the organism directly reaches the object with its arm. However, if O3 is not in the fovea the organism rotates its eye either to the left or to the right. The organism continues to rotate its eye until the object O3 is in the fovea, and at this point it reaches the object. The same is true for the other two objects, O4 and O5, and the other two linguistic signals, S4 and S5. The new language game makes it clear in what sense nouns control the movements of the organism's eye, head, or entire body that allow the organism to obtain visual access to some particular object contained in its environment so that the organism can execute some further action with respect to the appropriate object, i.e., the object specified by the noun.

In the language games we have described we can distinguish between verbs and nouns in that some particular linguistic signal co-varies either with the organism's action or with the particular object which is involved in the organism's action. In the former case we say that the linguistic signal is a verb whereas in the latter case it is a noun. But consider a fifth language game in which the organism lives in an environment which contains both edible and poisonous mushrooms (Cangelosi and Parisi, 1998). To survive and reproduce the organism must be able to approach (and eat) the edible mushrooms and to avoid the poisonous ones. Notice that each individual mushroom is perceptually different from all other mushrooms, including those belonging to the same category. Therefore, when it encounters a mushroom the organism must be able to both recognize (classify) the mushroom as either edible or poisonous and respond with the appropriate action to the mushroom (approaching and eating the edible mushrooms and avoiding the poisonous ones). When it encounters a mushroom the organism can hear one of two linguistic signals, S6 and S7, presumably produced by some nearby conspecific which wants to help our organism. Of these two linguistic signals, S6 co-varies with (all) edible mushooms and S7 co-varies with (all) poisonous mushrooms. Are S6 and S7 verbs or nouns? We think that the distinction cannot be made in this language game. S6 co-varies both with one type of action (approaching and eating the mushroom) and with one type of objects (edible mushrooms), and S7 co-varies with both the other type of action (avoiding the mushroom) and the other type of objects (poisonous mushrooms). Therefore, although S6 and S7 are linguistic signals since they influence the organism's behavior (for example they make the behavior more efficient), there is no ground for saying that they are either verbs or nouns because they co-vary simultaneously with both the action on the part of the organism and the type of objects to which the action is addressed. It might be that this type of language game, in which it is still impossible to distinguish between verbs and nouns, reflects a very primitive stage of language such as the language of our earliest language-using ancestors and the language of children between, say, 1 year and 1 year and a half of age.

In our model nouns co-vary with objects and verbs with actions. However, there are two types of objects, natural objects (e.g., trees) and artificial objects (e.g., knives). Organisms respond to natural objects with a variety of different actions depending on the circumstances but there is generally no particular action associated with each natural object. An organism may respond to a tree by cutting the tree, picking up fruits from the tree, recovering under the tree for shadow, etc. In contrast, organisms tend to respond to artificial objects with one particular action which is specific for each of them. A knife is normally used to cut, although a knife can also be bought, cleaned, put into a drawer, etc. Therefore, in a sense artificial objects are more associated with the specific actions than natural objects and, from this point of view, they resemble verbs. However, linguistic signals that co-vary with artificial objects are nouns in the same way as linguistic signals that co-vary with natural objects. In both cases the linguistic signal is used to direct the attention/action of the organism to some particular object in the environment.

3. Adjectives and, more generally, noun modifiers

Consider now a sixth, somewhat more complex, language game. In the preceding language games the different objects differed only in their shape. In the organisms' environment there was only one object for each shape, and therefore there were only two (or three, in the fourth language game) objects in all. In the new language game the organism's environment contains four objects. Two objects have one shape and the other two objects have a different shape. However, the two objects with the same shape differ in their color: one is blue and the other one is red.

In each episode the organism sees two objects and the two objects have the same shape but different color. Hence, providing the organism with the noun that refers to objects of a given shape (our second language game) is useless. The organism would not know which object to reach with its arm. However, we now introduce two new linguistic signals, S8 and S9. When the organism hears the sound S8 it reaches the blue object and when it hears the sound S9 it reaches the red object. In these circumstances S8 and S9 are (proto)adjectives. Notice that if the organism sees all four objects at the same time, it will need both a noun and an adjective in sequence (a (proto)noun phrase) to be able to identify the particular object which it is supposed to reach.

Adjectives have the same general role of nouns in the behavior of our organisms: they direct the attention of the organism to particular objects and guide the organism's action toward those objects. So what distinguishes nouns from adjectives? In our simulations nouns co-vary with (in common parlance, refer to) objects having particular shapes whereas adjectives co-vary with other properties of objects such as their color. In fact, shape appears to be more important for distinguishing among different nouns than other properties of objects. In psycholinguistic experiments both children and adults generalize invented words syntactically identified as nouns to other objects having the same color, size, or texture of an initial object more often than to objects with a different shape (Landau *et al.* 1988), although words syntactically identified as count nouns show this tendency more than words syntactically identified as mass nouns (Landau *et al*. 1992). Therefore, we hypothesize that, while both nouns and adjectives have the same general role of directing the attention/action of organisms to particular objects in the environment, nouns differ from adjectives because nouns direct the organisms' attention/action to objects with a given shape and adjectives to objects with a given color or size or some other property.

Of course, there is nothing special or metaphysical about shape as contrasted with color or size in object identification except that objects which differ in shape are more likely to require different actions on the part of organisms than objects differing in color or size. (This may explain why other properties of objects such as those that identify an object as an animal, e.g., texture, may also be important for nouns (Jones *et al*. 1991; 1998). Animals generally require different types of actions directed toward them in contrast to non-animals.) Shape rather than color or size tends to be unique to classes of objects that require specific types of actions. Trees tend to have a unique shape whereas they do not have a unique color or size. Only trees have the shape of trees but not only trees are green. All the objects which co-vary with (i.e. are designated by) a given noun share a particular shape which is not shared by other objects whereas even if they are all of the same color, like strawberries, this color is shared also by other objects not called "strawberries".

Now consider another language game. The organism sees two objects at the same time. The two objects can be either the same object (same shape) or two different objects (different shapes) but in any case they are located in different portions of the visual field. For example, an object can be located in the left portion and one in the right portion of the visual field. The organisms hears one of two sounds, S8 and S9. When it hears S8, the organisms reaches the object located in the left portion of the visual field whereas when it hears S9 it reaches the object located in the right portion of the visual field. Notice the difference between this language game and the second language game described above. In that language game the organism was also directed by language to go to the left portion or the right portion of the visual field. However, when the organism heard, for example, S3 it went to the left portion of the visual field if the object O1 was there but it went to the right portion of the visual field if the object O1 was in the right hemifield. In other words, the organism's behavior was guided by the shape of the objects and therefore S3 and S4 were classified as nouns. In this new language game, on the contrary, the organism reaches the object located in the left hemifield whethere the object is O1 or O2, i.e., independently from the shape of the object. Therefore the new linguistic signals, S8 and S9, cannot be nouns. Are they adjectives?

We introduce a new class of words called non-adjective noun modifiers. Both adjectives and non-adjective noun modifiers are noun modifiers but, while adjectives tend to co-vary with more or less permament properties of objects such as their color or size, non-adjective noun modifiers co-vary with more temporary properties of objects such as the object being located in the left or right portion of the organism's visual field. An object can be more or less permanently red or small but it is only temporarily placed, say, in the left portion of the organism's visual field. Hence, S8 and S9 are non-adjective noun modifiers. (Notice that non-adjective noun modifiers tend to be sequences of more than one word (phrases) whereas adjectives are single words. For example, the meaning of S8 is roughly equivalent to the meaning of the English phrase "on the left".)

To summarize, we have distinguished two large categories of linguistic signals: verbs and what we can call noun phrases. Verbs co-vary with the action with which the organism responds to the visual input largely independently from the content of the visual input. Noun phrases, on the other hand, direct the attention/action of the organism to particular visually perceived objects in the environment. Noun phrases can be simply nouns or they can be sequences of linguistic signals which almost always include a noun accompanied by a noun modifier, which can be either an adjective or a non-adjective noun modifier (itself a phrase in many cases). Noun modifiers have the same role of nouns in directing the attention/action of the organism to the particular object which is to be involved in the organism's action but they refer to different properties of objects. Nouns refer to the shape of objects or to other properties of objects that tend to be more highly correlated with the actions of the organism with respect to the objects. Adjectives refer to more or less permanent properties of objects which, however, are less highly correlated with the actions of the organism with respect to the objects. Non-adjective noun modifiers refer to more temporary or extrinsic properties of objects such as their current position in the organism's visual field or, more generally, in space (e.g., "on the desk").

Verbs also may be accompanied by verb modifiers which are similar to noun modifiers. These verb modifiers can be adverbs (single word) or adverbial phrases (sequence of words). Verb modifiers ask the organism to execute an action in the particular way which is indicated by the adverb or adverbial phrase. Consider this last language game. The language game is identical to our first language game in which the organism can either push or pull an object. What is new is that the organism can push or pull the object either slowly or quickly. The organism can hear two new signals, S10 and S11, together with the verbs S1 (pull) and S2 (push). When the organism hears S10, it pushes or pulls the object slowly whereas when it hears the S11 it pushes or pulls the object more quickly. S10 and S11 are (proto)adverbs.

4. Many open questions

We have described a number of simple simulated language games that are aimed at clarifying how heard sounds become linguistic signals and how different classes of sounds which play different roles in the organism's experience and interaction with the environment become different parts of speech. These language games are simulated in the sense that we can construct artificial organisms that behave in the ways we have described. Neural networks respond to the input, i.e., they behave, in particular ways because they have particular connection weights. In our simulations we use a genetic algorithm to find the appropriate connection weights which result in the desired behaviors. A genetic algorithm is a learning procedure which is inspired by evolution (Holland 1975). However, there is no assumption that the linguistic abilities (responding appropriately to linguistic signals) of our organisms are either entirely genetically inherited (which of course cannot be since different humans speak different languages) or entirely learned during life with no important genetically inherited basis (which cannot be since only humans have language). Simply, we have not addressed the problem of the origin of the linguistic abilities exhibited by our artificial organisms.

Of course, we have just scratched the surface of the problem of accounting for the differences among the parts of speech. Let us mention a list of open questions, with in some cases some hints as to how to address these questions in the present framework.

(1) We have simulated (some aspects of) the ability to understand language, i.e., to respond appropriately to heard sounds which are linguistic signals, but we haven't said anything about the ability to produce language, i.e., to execute the phono-articulatory motor behaviors which result in the physical production of the appropriate sounds/linguistic signals. To simulate the ability to speak it is necessary to add a further set of output units to the neural network of our organisms which will encode phono-articulatory movements resulting in the physical production of sounds. Aside from that, we believe that the basic categories of words remain the same: produced sounds are verbs if they co-vary with the actions of the speaker or of the hearer; they are nouns if they co-vary with the objects (mainly identified on the basis

of their shape) involved in the actions of the speaker or of the hearer; they are adjectives if they co-vary with other properties of objects; and so on.

(2) We have simulated verbal commands but language has many other pragmatic uses and is involved in different types of speech acts: acts of information, questions, expressions of intentions or desires, etc. To account for these other uses of language we will need more complicated language games and more complex social interactions among our simulated organisms.

(3) Many verbs to do not refer to actions and many nouns do not refer to concrete, perceptually accessible objects. Verbs sometimes co-vary with (i.e., refer to) processes rather than with actions (Langacker 1987). Actions are processes but many processes are not actions of organisms (e.g., the process of snowing). Verbs referring to processes which are not actions require that our artificial organisms possess an ability to abstract “change of state” (or even “lack of change of state” for verbs referring to states such as sleeping) in a succession of inputs even if the succession of input does not reveal an action. Furthermore, verbs and nouns may not all possess verbness and nounness to the same degree. There might be a continuum of verbness/nounness.

(4) Language is often used in situations in which the organism is not responding to external (in our case, visual) input with external motor behavior (in our case, the movements of the arm). The organism can respond to heard sounds without producing any external behavior, it can produce linguistic signals with no current input from the external environment, and it can even use language purely internally with no external input or external output of any kind (thinking). These uses of language all involve the self-generation of input by a neural network, both linguistic (imagined sounds) and nonlinguistic (imagined actions and their effects in the

environment) input. The ability to self-generate input is what defines mental life as distinct from behavior.

(5) Nouns and verbs, and of course the other parts of speech, have properties which are syntactic in nature, rather than semantic. These syntactic properties derive from their use in sequences of words which have sequential constraints (for example, in English verb objects follow verbs, do not precede them) and internal structure (cf. Cangelosi & Parisi 2002; Turner & Cangelosi 2002).

(5) Nouns can be morphologically "derived" from verbs and verbs from nouns.

(6) The kind of simple verb-noun sequences we have considered in one of our language games represent verb-object (proto)sentences. How verb subjects emerge in languages? Probably the emergence of subjects in action sentences (agents) is linked with the ability to recognize the same action as made by me and as made by other individuals (cf. the "mirror neurons" of Rizzolatti & Arbib 1998). In these circumstances one has to specify not only the object(s) on which the action is executed (the verb complement(s)) but also the author of the action, i.e., the agent (the verb's subject).

Acknowledgements
Angelo Cangelosi's work for this paper was partially funded by an UK Engineering and Physical Research Council Grant (GR/N01118).

Bibliographical References

Batali, John (1994), “Innate biases and critical periods: combining evolution and learning in the acquisition of syntax”, in Brooks, Rodney & Maes, Patti, eds., *Artificial Life IV*, Cambridge, Mass., MIT Press (1994:160-171).

Brown, Keith & Miller, Jim (1999), *Concise Encyclopedia of Grammatical Categories*, Amsterdam, Elsevier.

Cangelosi, Angelo (2001), “Evolution of communication and language using signals, symbols and words”, *IEEE Transactions on Evolutionary Computation*, 5:93-101.

Cangelosi, Angelo & Harnad, Stevan (in press), “The adaptive advantage of symbolic theft over sensorimotor toil: grounding language in perceptual categories”, *Evolution of Communication*, 4(1).

Cangelosi, Angelo & Parisi, Domenico (1998), “The emergence of a ‘language’ in an evolving population of neural networks”, *Connection Science*, 10:83-97.

Cangelosi, Angelo & Parisi, Domenico (2001), “How noun and verbs differentially affect the behavior of artificial organisms”, in Moore, Johanna D. & Stenning, Kennneth, eds., *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, Hillsdale, N.J., Erlbaum (2001:170-175).

Cangelosi, Angelo & Parisi, Domenico, eds. (2002), *Simulating the Evolution of Language*, London, Springer.

Ellefson, Michelle R. & Christiansen, Morten H. (2000), “Subjacency constraints without universal grammar: evidence from artificial language learning and connectionist modeling”, in *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*, Hillsdale, N.J., Erlbaum (2000:645-650).

Falcetta, Ilaria (2001), *Dalle reti neurali classiche alle reti neurali ecologiche: il significato come proprieta’ emergente delle interazioni senso-motorie tra organismo e ambiente*, Dissertation, University of Rome La Sapienza.

Holland, John H. (1975), *Adaptation in Natural and Artificial Systems*, Ann Arbor, Michigan, University of Michigan Press.

Jones, Susan S., Smith Linda B., & Landau Barbara (1991), "Object properties and knowledge in early lexual learning", *Child Development*, 62:499-512.

Jones, Susan S., Smith Linda B. (1998), "How children name objects with shoes", *Cognitive Development*, 13:323-334.

Kaplan, Frederik (2000), “Talking AIBO: first experimentation of verbal interactions with an autonomous four-legged robot”, in Nijholt, A., Heylen, D. & Jokinen, K., eds., *Learning to Behave: Interacting agents. CELE-TWENTE Workshop on Language Technology* (2000:57-63).

Kirby, Simon (1999), “Syntax out of learning: the cultural evolution of structured communication in a population of induction algorithms”, in Floreano, Dario *et al*., eds., *Proceedings of ECAL99 European Conference on Artificial Life*, New York, Springer (1999:694-703).

Knight, Chris, Studdert-Kennedy, Michael, & Hurford, Jim, eds., (2000) *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*, Cambridge, Cambridge University Press.

Landau, Barbara, Smith, Linda B., & Jones, Susan S. (1988), “The importance of shape in early lexical learning, *Cognitive Development*, 2:291-321.

Landau, Barbara, Smith, Linda B., & Jones, Susan S. (1992), “Syntactic context and the shape bias in children’s and adults’ lexical learning”, *Journal of Memory and Language*, 31:807-825.

Langacker, Ronald W. (1987), *Foundations of Cognitive Grammar. Volume 1: Theoretical Prerequisites*, Stanford, Cal., Stanford University.

Parisi, Domenico (1997), An Artificial Life approach to language, *Mind and Language*, 59:121-146.

Parisi, Domenico & Cangelosi, Angelo (2002), “A unified simulation scenario for language development, evolution, and historical change”, in Cangelosi, Angelo & Parisi, Domenico, eds., *Simulating the Evolution of Language*, London, Springer, 2002:255-276.

Rizzolatti, Giacomo & Arbib, Michael A. (1998), “Language within our grasp”, *Trends in Neurosciences*, 21:188-194.

Soja, Nancy N. (1992), “Inferences about the meanings of nouns: the relationship between perception and syntax”, *Cognitive Development*, 29-45.

Steels, Luc (1997), “The synthetic modeling of language origins”, *Evolution of communication*, 1:1-34.

Steels, Luc & Kaplan, Frederik (1999), “Collective learning and semiotic dynamics”, in Floreano, Dario et al., eds., *Proceedings of ECAL99 European Conference on Artificial Life*, New York, Springer (1999:679-688).

Steels, Luc & Vogt, Paul (1997), “Grounding adaptive language games in robotic agents”, in Husband, Paul & Harvey, Inman, eds., *Proceedings of the Fourth European Conference on Artificial Life*, Cambidge, Mass., MIT Press (1997:474-482).

Turner, Huck & Cangelosi, Angelo (2002), “Implicating working memory in the representation of constituent structure and the origins of word order universals”, paper presented at *4th International Conference on the Evolution of Language*, Boston.

Wittgenstein, Ludwig (1953), *Philosophical Investigations*, London, Blackwell.