

The Processing of Verbs and Nouns in Neural Networks: Insights from Synthetic Brain Imaging

Angelo Cangelosi

Centre for Neural and Adaptive Systems and School of Computing
University of Plymouth (UK)

Domenico Parisi

Institute of Cognitive Sciences and Technologies
National Research Council (Italy)

Abstract

The paper presents a computational model of language in which linguistic abilities evolve in organisms that interact with an environment. Each individual's behavior is controlled by a neural network and we study the consequences in the network's internal functional organization of learning to process different classes of words. Agents are selected for reproduction according to their ability to manipulate objects and to understand nouns (objects' names) and verbs (manipulation tasks). The weights of the agents' neural networks are evolved using a genetic algorithm. Synthetic brain imaging techniques are then used to examine the functional organization of the neural networks. Results show that nouns produce more integrated neural activity in the sensory processing hidden layer, while verbs produce more integrated synaptic activity in the layer where sensory information is integrated with proprioceptive input. Such findings are qualitatively compared with human brain imaging data that indicate that nouns activate more the posterior areas of the brain related to sensory and associative processing while verbs activate more the anterior motor areas.

1. Language processing in natural and artificial neural networks

Artificial neural networks have been frequently used to build models of language processing abilities in adults and children. They have been employed to study the acquisition of lexicon and meaning, the processing of morphology and syntax, reading and speech production (cf. Christiansen et al., 1999). However, much connectionist work on linguistic tasks tends to study language in isolation from other cognitive abilities and from the sensory-motor interactions of the organism with the environment. This is an obstacle to considering the important issue of the “grounding” of symbols on sensory-motor experience through which linguistic symbols acquire their meaning. Furthermore, in most connectionist models the issue of the neural plausibility and significance of the network architecture and functioning is not addressed, and this makes it impossible to compare the simulation results with such neural data as neuroimaging data.

Computational models have also been successfully employed for investigating the evolution of language through simulation (Cangelosi & Parisi, 2002; Kirby 2002). These models use various approaches: artificial neural networks (e.g. Batali, 1994; Cangelosi & Harnad, 2000), rule-based systems (Kirby, 2001), and robotics (Steels & Kaplan, 1999). Neural networks have proven particularly useful because they can focus on the influence of both cognitive and neural mechanisms on language development and evolution. For example, evolutionary neural networks, or networks viewed in an Artificial Life perspective (Cangelosi & Harnad, 2000; Parisi 1997), are used to control the behavior of organisms that live in an environment and communicate between themselves. This approach provides a unifying theoretical framework for cognitive and neural modeling because of the use of both evolutionary and connectionist techniques. It offers the benefits of (i) studying some of the neural mechanisms at the basis of the evolution of linguistic abilities, (ii) understanding the interaction between the neural control of language and that of other behavioral and cognitive

abilities, and (iii) grounding language directly on the organism's sensorimotor and cognitive abilities.

Some studies using evolutionary neural networks have focused on the evolution of word classes and syntax. Cangelosi (2001) simulated the emergence of learned compositional languages, obtaining a strong evolutionary tendency to evolve a predicated-argument structure resembling a "verb-noun" rule. In Cangelosi & Parisi (2001), an explicit verb-noun language evolves in a population of simulated organisms. The model highlights the beneficial effects of language on non-linguistic behavior and the asymmetric evolution of nouns before verbs. The analysis of categorical perception effects in the internal representations of networks (Cangelosi & Harnad, 2000) demonstrated that verbs produce more distinct internal representations of objects and events because of their more direct involvement with the motor actions with which the organism responds to the sensory input. However, in such studies no issues are considered regarding the architecture of the network, and there is no attempt at identifying internal representations of verbs and nouns in different parts of the network architecture.

On the contrary, the neuropsychological and neurocognitive literature on language processing in the brain is quite extensive (Gazzaniga 2000). Neuropsychological experiments use a variety of experimental approaches such as the investigation of patients with psycholinguistics deficits and brain imaging studies on normal adults and language-impaired patients. For example, various studies have analyzed the neural correlates of the processing of various word classes and the verb-noun dissociation in patients. Cappa & Perani (2003) recently analyzed the literature on the neural processing of verbs and nouns. They found a general agreement on the fact that the left temporal neocortex plays a crucial role in lexical-semantic tasks related to the processing of nouns whereas the processing of words related to actions (verbs) involves additional regions of the left dorsolateral prefrontal cortex. For

example, in the well known neuropsychological study on verbs and noun processing, Damasio & Tranel (1993) reported that most of the patients with selective disorders of noun retrieval had lesions in the left temporal lobe. Instead, verb impairment was associated with damage on the left prefrontal cortex. In a PET study, Martin and colleagues (1995) compared color naming (nouns) and action naming (verbs). They observed a selective activation for color naming of the left fronto-parietal cortex, the middle temporal gyrus, and the cerebellum. Perani, Cappa *et al.* (1999) also used PET for the processing of concrete and abstract verbs and nouns in Italian. Results indicated that left dorsolateral frontal and lateral temporal cortex were activated only by verbs. In the comparison of abstract and concrete words, only abstract word processing was associated with selective activation of the right temporal pole and amygdala and the bilateral inferior frontal cortex. Finally, in evoked potential studies it was reported that there is selective activation of the frontal lobes for action words (Preissl, Pulvermueller *et al.*, 1995). This difference is related to the semantic content of words rather than to grammatical differences, since no difference was observed between action verbs and nouns with a strong action association (Pulvermueller, Mohr & Schliechert, 1999).

Brain simulation models, such as those of computational neuroscience, have rarely focused on complex linguistic behavior, except for a few studies (e.g., Just et al. 1999). This is due to the complexity of the various linguistic functions (speech processing, lexical and semantic knowledge, syntax) to be included in a model. However, brain simulation models have been commonly developed for a variety of behavioral and cognitive abilities, such as vision, memory, and motor control. More recently, in such models the method of synthetic brain imaging (Arbib et al. 2000; Horwitz et al. 1999) has permitted a more strict integration of experimental data and computational models and a direct comparison of performance in artificial and natural brains.

This paper presents a computational model of language in which linguistic abilities evolve in organisms that interact with an environment. Therefore, the model makes it possible to study the interaction of cognitive and linguistic abilities. The aim of the simulation is to study the consequences of architectural and functional constraints (Feldman et al., 1988) in evolutionary neural networks that learn to process (understand) different classes of words (nouns and verbs). The model uses synthetic neuroimaging techniques to examine the internal organization of the neural networks and compares the results with data reported in the literature on language processing in the brain. The comparison is a limited one and there is no attempt at building a plausible model of language processing in the human brain.

2. Synthetic brain imaging (SBI)

The development of functional brain imaging techniques, such as PET (positron emission tomography) and fMRI (functional magnetic resonance imaging), has permitted significant advancements in our understanding of the neural basis of human and animal cognition and sensorimotor abilities. However, functional imaging data are complex to interpret. To help appreciate the cognitive and neural significance of such rich data, neural modeling methods can be useful. This methodology is referred as *synthetic brain imaging (SBI, thereafter)* (Arbib et al. 2000; Horwitz et al. 1999). It can be used to compare directly PET and fMRI imaging data in empirical studies and in computational models. The aim is to develop more detailed and neurally-plausible models of behavioral functions.

Experimental brain imaging using PET and fMRI is based on differences in brain activity between tasks. This activity is measured through changes in blood flow (rCBF, regional Cerebral Blood Flow) which are the results of synaptic and neural activation in different parts of the brain. Synthetic brain imaging is also based on differences in artificial neural network “activity” between tasks. For example, following Arbib’s et al. (2000) assumptions, in

synthetic PET the involvement of a region in a specific task is correlated with local integrated synaptic activity. Synaptic activity is computed by summing together the synaptic inputs arriving to the artificial neurons of a specific area, ignoring the positive/negative sign of the connections. Synthetic fMRI, instead, can be correlated with the activation of neurons. The activation values of artificial neurons in a specific portion of the neural network are summated over a period of time and compared with experimental fMRI data. However, in the literature on SBI the methods for calculating synthetic PET/fMRI data from synaptic/neural activity can vary depending on the types of computational model adopted (Tagamets & Horwitz, 2001).

SBI has been used in studies with animal and humans for developing more detailed models of cognitive and behavioral abilities. Arbib *et al.* (2000) used synthetic PET to test their FARS model of parietal-premotor interactions underlying primate grasp control. Tagamets and Horwitz first applied synthetic PET (Tagamets & Horwitz, 1998) and then synthetic fMRI (Horwitz & Tagamets, 1999) to a large-scale model of working memory for determining the interregional connection patterns and strengths.

In a study on language, Just *et al.* (1999) developed a computational model of sentence comprehension and analyzed the simulation results using synthetic neuroimaging. The results accounted for fMRI data that vary as a function of sentence complexity in Broca and Wernicke areas and in the dorsolateral prefrontal cortex. They found good agreement between the number of activated voxels in human brain imaging experiments and the predictions of their model for sentence types of different complexity.

The use of SBI methods is still in an exploratory stage. There is ongoing debate on the general role of synthetic PET and fMRI data and on how to exactly compute such data. In addition, some researchers (e.g. Arbib *et al.*, 2000) propose the use of SBI only for “realistic” neural network models, such as those based on leaky integrator models. However, SBI appears to be a useful and promising research methodology, as we will try show in this paper.

3. A neural network model of the evolution of nouns and verbs

To study the processing of different word classes (nouns and verbs) in artificial neural networks, an artificial life model of object manipulation was developed. The simulated agent consists of an organism with a retina and a two-segment arm (e.g. Marocco et al., 2002; Schlesinger & Barto, 1999). The object manipulation scenario was chosen because of the important role of hand use and tool manipulation abilities in language evolution hypotheses (e.g. Corballis, 2002; Rizzolatti & Arbib, 1998).

The simulations presented in this paper are based on a previous model of the evolution of verbs and nouns (Cangelosi & Parisi, 2001; cf. also Parisi et al., in press). Most of the model parameters, which are summarized in the following sections, are the same as in the above studies. The major differences concern the neural network architecture.

3.1 *The behavioral and linguistic tasks*

The artificial agent has to perform a number of different tasks. At the beginning of each task the agent is grasping an object with its hand and it has to either push the object away from itself or pull the object toward itself. Two different objects are used, a vertical bar (object A) and a horizontal bar (object B). To perform the task the agent may respond to visual stimuli, linguistic stimuli, or both. Furthermore, the agent receives a proprioceptive input regarding the current position of the two segments of its arm.

The lifetime of an individual agent is divided into 11 tasks. In the first task (No_language), the retina encodes a 5x5 image of either object A or B in different positions and there are no linguistic stimuli. When object A is presented, the agent must push the object toward itself. When object B is presented, the object must be pulled away. These are the two default actions with which the agent must respond to the two objects. In the next five tasks (Vision+Language), the agent receives both the visual input from the retina and a linguistic input. In the remaining 5 tasks (Language_only), the retina input is shut off and only the

linguistic input is received by the agent. (The proprioceptive input is received in all epochs.) In the epochs with linguistic input five different linguistic inputs can be received by the agent: (1) noun of object A or B; (2) verb describing the default action (push for object A; pull for B); (3) verb describing the opposite action (pull for object A; push for object B); (4) both noun of object and verb for default action; (5) both noun of object and verb for opposite action. When the verb is received as input, the agent must do whatever the verb meaning indicates, thus overriding the default action if the verb describes the opposite action.

The use of the terms “noun” and “verb” in the present model is not intended to correspond to the full blown grammatical categories of verbs and nouns. Nevertheless, the “nouns” and “verbs” of the present simulation seem to capture some of the fundamental properties of nouns and verbs of real human languages (Parisi et al., in press). Nouns co-vary with the input stimuli and can act as attentional cues. Verbs co-vary with the motor action performed by the agents. These simple properties may have been those of primitive proto-nouns and proto-verbs which, through cultural and linguistic evolution, have given way to the complex, full blown grammatical categories of nouns and verbs. In fact, the analysis of the internal (neural) organization of our agents shows that there are functional communalities between the verbs/nouns in the model and the verbs/nouns in human subjects.

3.2 Neural network

The behavior of the agent is controlled by a feedforward neural network (Figure 1). The visual scene is perceived through a retina of $5 \times 5 = 25$ cells. An object is encoded as a pattern of 25 bits with three horizontally/vertically aligned 1s and twenty-two 0s. The proprioceptive input is encoded in 4 units, which represent proprioceptive information about the two pairs of muscles (extensor and contractor) of each of the two arm’s segments. The linguistic input is encoded in 4 input units which localistically encode two nouns and two verbs. Depending on

the task, one or two of these units can be on (noun or verb only; both noun and verb) but in the No_language task all four linguistic input units are turned off.

The network has two successive layers of hidden units (Figure 1). The first hidden layer is divided into two separate sensory-processing modules, one for processing the proprioceptive input and the other one for processing the visual input from the retina. Both modules project to a second hidden layer which has the role of integrating sensorimotor knowledge because of the convergence of motor (proprioceptive) and visual information¹. The four linguistic input units project to both the visual hidden units in the first layer and to the sensorimotor hidden units of the second layer.

The output layer contains four motor units. These units control the extension/contraction of the four arm muscles (a pair of extension/contraction muscles per arm segment). At each time step, the output activation corresponds to the force that is applied to each muscle.

The activation rule for all hidden and output neurons is a standard logistic function.

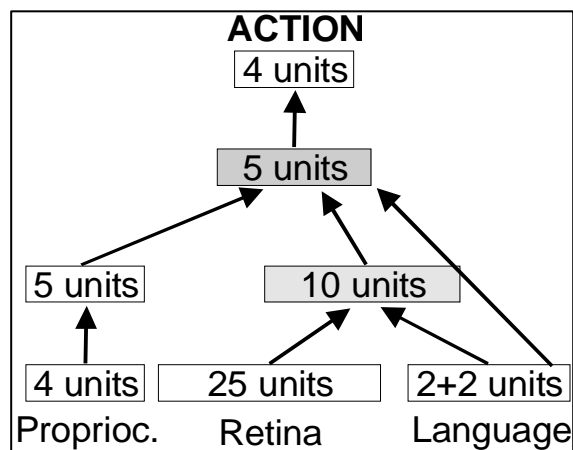


Fig 1 – Neural network architectures with specialized hidden layers. The first hidden layer works as a sensory processing module. The second layer works as a sensorimotor integration module.

¹ This architecture has been previously compared with another network without specialized hidden layers. Analyses on the hidden layer activity showed that this network actually uses different processing strategies in the two hidden layers (Cangelosi, in press).

3.3 Evolutionary algorithm

A genetic algorithm is used to evolve the genotype (connection weights) of the agents' neural networks. During an individual agent's lifetime, the weights do not change. The motor behavior and the language understanding of the agents improve (are learned) during evolution because the best agents are selected for reproduction in a succession of generations.

In each generation, a population of 80 agents is used. The genotypes encode the connection weights of the network (as real numbers). At generation 0, weights are initialized by selecting randomly a value in the range ± 1 . A single task in the lifetime of an individual consists of a total of 18 subtasks (2 objects x 9 positions) each subtask lasting for 20 input/output cycles. These 20 cycles are necessary for the arm to move an object from its initial position to the target. The fitness formula computes the total number of subtasks successfully completed by each agent in all epochs.

At the end of their life, agents are ranked on the basis of their fitness and the 20 agents with highest fitness are selected for reproduction. Each individual asexually generates 4 offspring with the same genotype of its single parent except for the addition of some random changes to some of the weights.

In the first 1000 generations the agents are only exposed to a single task, i.e., the No_language task. These initial generations are needed for evolving a set of connection weights that allows agents to respond appropriately to the two objects (before language is introduced). After generation 1000, agents have a longer lifetime and they are exposed to all 11 tasks, that is, the single No-language task and the 10 linguistic tasks. (The 11 tasks are experienced in a random order by each individual). In the fitness formula for the 10 tasks with language, a subtask is successful if the agent pushes or pulls the object according to the linguistic input.

The genetic algorithm is used here for simulating the process of acquiring simple linguistic abilities such as that of understanding (proto)nouns and (proto)verbs. However, the present model is not intended to support a view of language as based only on biological evolution. A simulation model meant to address the issue of the biological vs cultural origins of language should include both genetic and learning algorithms (cf. Cangelosi & Parisi 2001; Kirby, 2002). For example, Munroe and Cangelosi (2003) combine the genetic algorithm with the error backpropagation learning algorithm for neural networks. Their model makes it possible to analyze different types of Baldwin effects in the evolution of language.

4 Results: Synthetic brain imaging

The simulation was replicated 10 times with different randomly generated initial conditions. In all replications, agents evolved an ability to respond appropriately to the different input conditions. At the end of generation 2000, over 90% of the objects are responded to with the appropriate pushing or pulling behavior in all tasks. Because in this paper we focus on the synthetic brain imaging (SBI) analyses, other interesting results are not reported here (e.g., differences in the timing of evolution for verbs and nouns) but they can be found in Cangelosi & Parisi (2001) and Cangelosi (in press).

For the analysis of the internal organization of the agents' neural networks using SBI, the methods proposed by Arbib *et al.* (2002) were used. One first analysis is based on the amount of synaptic input to a layer of neurons and will be called ISA: Integrated Synaptic Activity (*synthetic PET* in Arbib *et al.* 2000). A second analysis is based on the neural activity of neurons and will be called INA: Integrated Neural Activity (*synthetic fMRI* in Arbib *et al.* 2000)

The overall integrated synaptic activity value $ISA_A(1/2)$ for the difference of activity between two different tasks (task 1 and task 2) in receiving region (hidden layer) A was computed using the following formula:

$$ISA_A(1/2) = |rISA_A(1) - rISA_A(2)| \quad (1)$$

where the regional synaptic activity for each task $rISA_A$ is

$$rISA_A = \sum_{i=1}^N w_{B \rightarrow A(i)} \quad (2)$$

with N equal to the number of units in the hidden layer, and the activity of each individual synapses $w_{B \rightarrow A(i)}$ is

$$w_{B \rightarrow A(i)} = \text{activation_sending_neuron} \times |\text{weight}| \quad (3)$$

Similar calculations were used for the integrated neural activity $INA_A(1/2)$. In this case, the data refer to the activity within each layer, instead of the receiving synaptic activity. Formulae (4) and (5) were used for the difference of regional $rINA_A$ values in the two tasks, based on the summation of the activity (activation value) of each neuron $a(t)$.

$$INA_A(1/2) = |rINA_A(1) - rINA_A(2)| \quad (4)$$

$$rINA_A = \sum_{i=1}^N a(t) \quad (5)$$

In the simulation there are 11 different tasks. The ISA and ISA values for all possible pairs of task subtractions were calculated for each hidden layer (e.g., No_language–Noun_only, No_language–Verb_only, Verb_only–Noun_verb). Subsequently, all subtractions were averaged for the two linguistic classes of verbs and nouns, e.g., the average of the values No_language–Noun_only and Verb_only–Noun_verb contribute to the ISA value for Noun processing.

Data were statistically evaluated using repeated measure analysis of variance (MANOVA) separately for the ISA and INA data. A 2x2 design utilized as independent variables HIDDEN_LAYER (levels H1 and H2) and WORD_CATEGORY (levels Noun and Verb). The dependent variable of each MANOVA was the respective subtraction ISA or INA data of 20 networks. These are the set single best individuals of the last generation in each of the 20 replications of the simulation. The means and standard errors for the integrated synaptic activity ISA values are reported in Figure 2a. The integrated neural activity INA data are in Figure 2b.

In the ISA MANOVA, the two main factors and their interaction were all significant: HIDDEN_LAYER $F(1/19)=13.59$ $p<0.002$, WORD_CATEGORY $F(1/19)=6.43$ $p<0.02$, interaction $F(1/19)=8.92$ $p<0.008$. The comparison of means (Fig. 2a) indicates that the processing of nouns is not differentiated between the two layers. Instead, the synaptic activity for verbs is significantly differentiated in the two layers, with the second hidden layer (sensorimotor integration) being more active.

In the INA MANOVA, only the word category and the interaction were significant: HIDDEN_LAYER $F(1/19)=1.11$ $p=3.04$, WORD_CATEGORY $F(1/19)=15.96$ $p<0.001$, interaction $F(1/19)=24.37$ $p<0.0001$. Figure 2b shows that during noun processing, the first hidden layer (sensory processing) is significantly more active than the second layer. No relevant differences appear for verbs, although activity in the sensorimotor integration layer is higher.

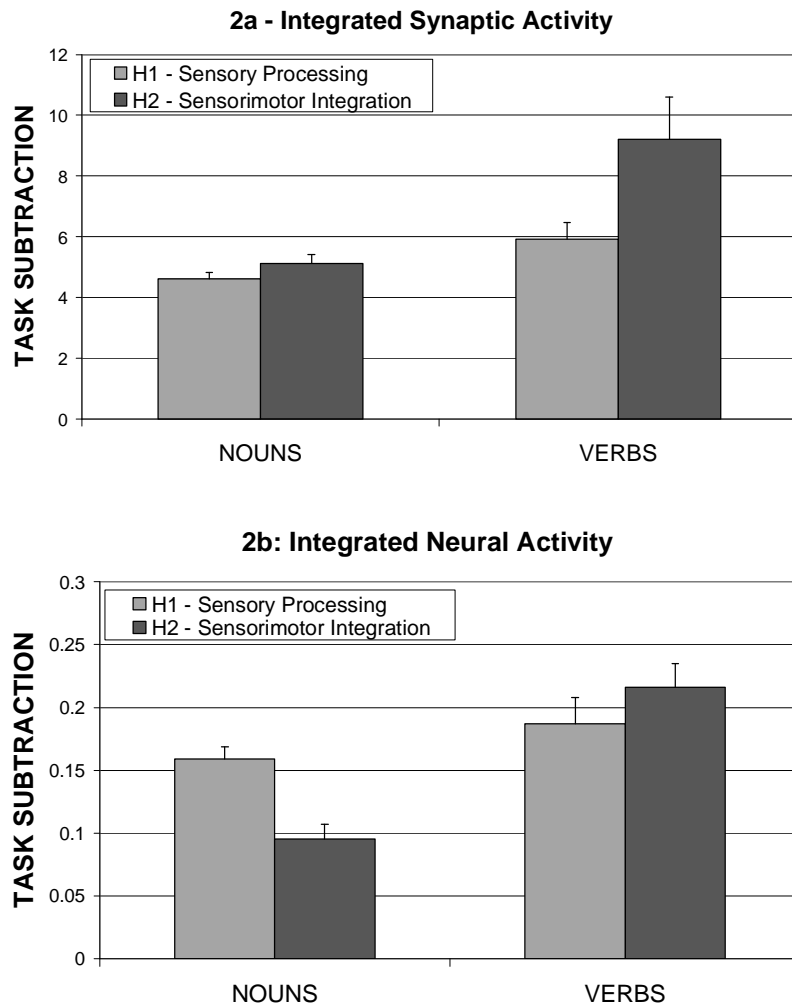


Fig 2 – Data on the Integrated Synaptic Activity ISA (2a) and Integrated Neural Activity INA (2b) for verb and noun processing in network with specialized architecture. See text for explanation.

5. Discussion

Our simple neural network model of the evolution of verbs and nouns exhibits some surprising functional and architectural similarities with data on the neural processing of different word classes in the brain. In the computational model, nouns produce more neural activity (INA data – Fig 2b) in the sensory processing layer, while verbs produce more synaptic activity (ISA data – Fig. 2a) in the layer where sensory information is integrated with proprioceptive input to plan the action. In the human brain, nouns activate more the (posterior) areas of the brain related to sensory and associative processing, while verbs activate more the (anterior) motor areas (Cappa & Perani, 2003).

As previously stated, the aim was not to build a realistic and plausible model of language processing in the brain. The scope of the simulation was to establish if the architectural constraint of dividing the hidden layer into two different processing layers actually resulted in a functional specialization for the two word categories. The SBI data confirm that this is the case. This reinforces previous analyses of activation pattern representations in architectures with different modular organization of the hidden layer. In another study of linguistic evolutionary networks (Cangelosi & Parisi, 2001), the categorical perception effects differ for the two word categories: verb processing produces more compressed within-category representations and larger between-category differences compared to noun processing. In addition, the enhanced categorical perception effects for verbs always only appear in the sensorimotor integration layer (Cangelosi, in press).

The present study demonstrates the usefulness of using SBI to study the functional and architectural specializations of artificial neural networks and to determine how these specializations might correspond to those reported in the literature on cognitive and neural processing of language. The methodology can be used for various research purposes. In the field of language evolution modeling, simulations can focus on how functional changes in network architecture will affect the evolution of cognitive and linguistics abilities (Deacon, 1997). For example, the connectivity pattern of the networks can be evolved leading to changes in the size and position of the sensorimotor integration layer. The researcher can then monitor the effects of such changes, e.g., in favoring, or inhibiting, the emergence of nouns and verb processing abilities. The simulation results can subsequently be extended to build broader explanatory hypotheses on the evolution of language.

One further direction of research is to investigate in more detail the SBI data. The results presented in Section 4 only use data averaged over all 11 language processing tasks. Therefore, potentially important differences in synaptic and neural activity among tasks may

have been overlooked. For example, detailed analyses of the synthetic synaptic/neural data for individual pairs of task subtractions could reveal significant differences between tasks in which language is received in isolation and those in which language is added to the visual input. Moreover, the ISA data on synaptic activity could be decomposed into two values, that of inhibitory synapses and that of excitatory synapses. This could reveal a differential positive/negative contribution of layers in the processing of nouns and verbs.

Finally, our artificial life approach to neural networks (Parisi, 1997) has the advantage of studying the simultaneous interaction between different cognitive abilities and language. This is achieved in a system where symbols are directly grounded in the individual's sensorimotor behavior and interaction with the environment (Harnad, 1990). The tool of SBI can make it possible to reach a deeper understanding of the cognitive and neural mechanisms which interact to ground linguistic and symbol manipulation abilities on simpler sensorimotor behavior (Cangelosi & Harnad, 2000).

Acknowledgements

Angelo Cangelosi's research was supported by the UK Engineering and Physical Research Council (EPSRC Grant GR/N01118).

References

- Arbib, M.A., Billard, A., Iacoboni M., & Oztop E. (2000). Synthetic brain imaging: grasping, mirror neurons and imitation. *Neural Networks*, 13, 975-997
- Arbib, M.A., Fagg, A.H., & Grafton, S.T. (2002). Synthetic PET imaging for grasping: From primate neurophysiology to human behavior. In F. Sommer and A. Wichert (Eds.), *Explorative analysis and data modelling in functional neuroimaging*. Cambridge, MA: The MIT Press.
- Batali, J. (1994). Innate biases and critical periods: Combining evolution and learning in the acquisition of syntax. In R. Brooks and P. Maes (Eds.), *Artificial Life IV* (pp. 160-171). Cambridge, MA: MIT Press.
- Cangelosi, A. (in press). Neural network models of category learning and language. *Brain and Cognition*
- Cangelosi, A. (2001). Evolution of communication and language using signals, symbols, and words. *IEEE Transactions on Evolutionary Computation*. 5(2), 93-101

- Cangelosi, A., & Harnad, S. (2000). The adaptive advantage of symbolic theft over sensorimotor toil: Grounding language in perceptual categories. *Evolution of Communication*, 4(1), 117-142
- Cangelosi, A., & Parisi D. (2001). How nouns and verbs differentially affect the behavior of artificial organisms. In J.D. Moore and K. Stenning (Eds.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 170-175), London: LEA.
- Cangelosi, A., & Parisi, D. (Eds.) (2002). *Simulating the Evolution of Language*. London: Springer-Verlag.
- Cappa, S.F., & Perani, D. (2003). The neural correlates of noun and verb processing. *Journal of Neurolinguistics*, 16 (2-3), 183-189
- Christiansen, M.H., Chater, N., & Seidenberg, M.S. (Eds.) (1999). Connectionist models of human language processing: Progress and prospects (Special Issue). *Cognitive Science*, 23(4), 415-634.
- Corballis, M.C. (2002). *From Hand to Mouth: The Origins of Language*. Princeton Univ. Pr.
- Deacon, T.W. (1997). *The Symbolic Species: The Coevolution of Language and Human Brain*, London: Penguin.
- Feldman, J., Fianty M., & Goddard, N. (1988). Computing with structured neural networks. *IEEE Computer*, 21, 91-104.
- Gazzaniga, M.S. (Ed.) (2000). *The New Cognitive Neurosciences*. Cambridge, MA: MIT Press.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42, 335-346.
- Horwitz, B., & Tagamets M.-A. (1999). Predicting human functional maps with neural net modeling. *Human Brain Mapping*, 8, 137-142.
- Horwitz, B., Tagamets, M.-A., & McIntosh, A.R. (1999). Neural modeling, functional brain imaging, and cognition, *Trends in Cognitive Science*, 3, 91-98.
- Just, M.A., Carpenter P.A., & Varma, S. (1999). Computational modeling of high-level cognition and brain function. *Human Brain Mapping*, 8, 128-136.
- Kirby, S. (2001). Spontaneous evolution of linguistic structure: An iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2), 102-110.
- Kirby, S. (2002). Natural language and artificial life. *Artificial Life*, 8, 185-215.
- Marocco, D., Cangelosi, A., & Nolfi, S. (2002). The role of social and cognitive abilities in the emergence of communication: Experiments in evolutionary robotics. *EPSRC/BBSRC International Workshop Biologically-Inspired Robotics* (pp. 174-181).
- Martin, A., Haxby, J.V., Lalonde, F.M., Wiggs, C.L., & Ungerleider, L.G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. *Science*, 270, 102-105.
- Munroe, S., & Cangelosi, A. (2003). Learning and the evolution of language: The role of cultural variation and learning cost in the Baldwin Effect. *Artificial Life*, 8, 311-339
- Parisi, D. (1997). An Artificial Life approach to language. *Brain and Language*, 59, 121-146.
- Parisi, D., Cangelosi, A., & Falcetta, I. (in press). Nouns, verbs and language games. *Journal of Italian Linguistics*
- Perani, D., Cappa, S. F., Schnur, T., Tettamanti, M., Collina, S., Rosa, M. M., & Fazio, F. (1999). The neural correlates of verb and noun processing: A PET study. *Brain*, 122, 2337-44.
- Preissl, H., Pulvermueller, F., Lutzenberger, W., & Birbaumer, N. (1995). Evoked potentials distinguish between nouns and verbs. *Neuroscience Letters*, 197, 81-83.
- Pulvermueller, F., Mohr, B., & Schliechert, H. (1999). Semantic or lexico-syntactic factors: what determines word-class specific activity in the human brain? *Neuroscience Letters*, 275, 81-84.

- Rizzolatti, G., & Arbib, M. (1998). Language within our grasp. *Trends in Neuroscience*, 21, 188-194.
- Schlesinger, M., & Barto, A. (1999). Optimal control methods for simulating the perception of causality in young infants. In M. Hahn & S.C. Stoness (Eds.), *Proceedings of the 21st Annual Conference of the Cognitive Science Society* (pp. 625-630). New Jersey: LEA.
- Steels, L., & Kaplan, F. (1998). Situated grounded word semantics. In *Proceedings of IJCAI-99* (pp 862-867). Los Angeles: Morgan Kauffman Publishing.
- Tagamets, M.A., & Horwitz, B. (1998). Integrating electrophysiological and anatomical experimental data to create a large-scale model that simulates a delayed match-to-sample human brain imaging study. *Cerebral Cortex*, 8, 310-320.
- Tagamets, M.A., Horwitz, B. (2001). Interpreting PET and fMRI measures of functional neural activity: The effects of synaptic inhibition on cortical activation in human imaging studies. *Brain Research Bulletin*, 54, 267-273.