In *Cognitive Processing*, 6, 202-207 (2005)

# Evolving Cognitive Systems: Adaptive Behaviour and Cognition Research at The University of Plymouth (UK) [1]

**Angelo Cangelosi**
Adaptive Behaviour & Cognition Research Group
School of Computing, Communications and Electronics
University of Plymouth (UK)
email: acangelosi@plymouth.ac.uk
**http://www.tech.plym.ac.uk/soc/research/abc**

## Introduction

The “Adaptive Behaviour and Cognition” research group of the University of Plymouth was established in 2002 within the then School of Computing (now School of Computing, Communications and Electronics). The overarching research aim of the group is to investigate the evolution and organization of natural and artificial cognitive systems using adaptive behaviour and cognitive modelling methods. The group, lead by Angelo Cangelosi, mainly consists of academic staff and PhD students from the School of Computing, Communications and Electronics. It also has few members from the School of Psychology. The group is part of the “Centre for Interactive Intelligent Systems”, and has close links with the Spatial Language Group within the University. Members of the group also closely collaborate with other research groups in the United Kingdom and abroad, such as the Institute of Cognitive Sciences and Technology of the CNR National Research Council in Rome (Domenico Parisi, Stefano Nolfi), the university of Genoa (Alberto Greco) and the University of Quebec at Montreal (Stevan Harnad). The research activity of the group members has been financially supported by various grants from the UK research councils and charitable foundations (Engineering and Physical Sciences Research Council, Biology and Biotechnology Sciences Research Council, The Nuffield Foundation, The British Academy, EOARD office of US Air Force Research Labs).

Adaptive behaviour is the area of artificial intelligence that uses computer and robotic modelling methodologies to simulate behaviour in autonomous systems (e.g. virtual agents, interactive robots, multi-agent systems). It embraces a variety of modelling techniques such as artificial life, cognitive robotics, evolutionary computation, neural networks. The application of adaptive behaviour methodologies to cognitive science permits the investigation of the evolution and organization of behaviour in natural and artificial cognitive systems. The study of cognitive systems has been identified as one

---

[1] Cangelosi A. (2005). Evolving cognitive systems: Adaptive behaviour and cognition research at the University of Plymouth. *Cognitive Processing*, 6, 202-207 – Springer copyright: contact publisher for permission to reprint.

of the priorities areas for research in Europe. Various national and international programmes have focuses on such a theme, such as the Foresight Cognitive Systems initiative in the UK, and the IST Cognitive Systems Programme in the European Union Framework VI. The main goal is to design and build artificial information processing systems which include various capabilities for perception, learning, reasoning, decision-making, communication and action. This has both a technological and scientific bearing. In technology, this research effort supports the development of novel methodologies of the design of intelligent systems, such as autonomous robots. In science, the process of constructing intelligent systems can help the understanding of the organisation of natural cognitive systems, such as humans and animals.

The great majority of studies on adaptive behaviour have typically focused on simple behaviours such as sensorimotor coordination, development of animal perception systems (e.g. rats' whiskers), navigation and action selection, and group behaviour. Only recently there has been an increased interest in applying adaptive behaviour techniques to model higher-order cognitive abilities, such as language. This is the research area where the "Adaptive Behaviour and Cognition" group in Plymouth has specialised. The group has played an important role in supporting this shift of focus from simple adaptive behaviour to complex cognitive capabilities. In particular, we have focused on modelling the relationship between language, cognition and action[2], and the design of artificial cognitive systems able to ground language in their own cognitive and sensorimotor capabilities. This is achieved through the evolution of cognitive systems able to develop autonomously shared communicate and linguistic abilities. In addition, great emphasis is put on the design of psychologically-plausible connectionist models of language where neural networks are trained to ground linguistic terms (such as spatial prepositions and quantifiers) into perceptual and cognitive representations. This is achieved through the use of training data from ad-hoc psycholinguistic experiments.

In the past years, the activity of the group has centred on a variety of research projects, including these four main topics: (i) The cognitive symbol grounding hypothesis; (ii) Communication in cognitive robots; (iii) Simulating the evolution of language; (iv) Grounding spatial terms and quantifiers in perception. These will be briefly summarised in the next sections, to highlight the contribution given by the group in the understanding of the relationship between language, cognition and action, and the design of psychologically-plausible cognitive systems. Details on other research activity of the group, including agent modelling for egress behaviour and its applications to evolutionary architecture (Holden & Cangelosi 2004, 2005), can be found in the group's web page: http://www.tech.plym.ac.uk/soc/research/ABC

## *The cognitive symbol grounding hypothesis*

Computational cognitive models that focus on linguistic and symbol-manipulation abilities can use symbols that are either grounded or ungrounded. Models using

---

[2] The group organised the 9th Neural and Psychology Workshop (September 2004) with a special theme on "Modelling Language, Cognition and Action" (see Cangelosi, Bugmann and Borisyuk 2005)

ungrounded symbols require the interpretation of an external user, such as the researcher, to identify and understand the meaning associated to symbols. On the other hand, the cognitive models based on grounded symbols use words that are inherently significant to the cognitive system. The issue of intrinsically linking the symbols used by a cognitive agent to their corresponding meanings has been called by Harnad "the symbol grounding problem". The various modelling approaches to the symbol grounding problem all have in commons some core features. First, each symbol is directly grounded into an internal categorical representation. Internal representations include perceptual categories (e.g. the concept of red colour, square shape, and female face), sensorimotor categories (e.g. the concept/action of grasping, pushing, and pulling), social representations (e.g. individuals, social groups and relationships) and other categorizations of the organism's own internal states (e.g. emotional states, motivations). Secondly, these categories are connected to the external world through our perceptual, motor and cognitive interaction with the environment. Categorical representation of the organism's internal states can also be mediated by our sensorimotor and cognitive system. This view of the symbol grounding process is called "Cognitive Symbol Grounding Hypothesis" (Cangelosi 2005). This is consistent with growing theoretical and experimental evidence on the strict relationship between symbol manipulation abilities and our perceptual, cognitive and sensorimotor abilities. For example, some studies have explicitly supported the fact that symbols are grounded in our ability to form categories through our categorical perception abilities (e.g. Cangelosi & Harnad 2000).

Our group has focuses on the symbol grounding problem in connectionist and in embodied agent simulations. The grounding of language in autonomous cognitive systems requires two mechanisms. The first is the direct grounding of the agent's basic lexicon. This assumes the ability to link perceptual (and internal) representations to symbols through supervised feedback. For example, an agent can learn that the symbol "horse" is grounded in its direct experience with this animal. The second mechanism implies the ability to transfer the grounding from the basic symbols to new symbols obtained by logical (e.g. syntactic) combinations of the elementary lexicon. The same agent can learn, without direct experience, that there is a hypothetical animal, the "unicorn", which is grounded in the linguistic description of "horse with a horn". This indirectly provides the perceptual grounding of "unicorn", when the categories "horse" and "horn" have been acquired through direct sensorimotor interaction with horses and horns. Various connectionist simulations by Cangelosi, Greco, Harnad and Riga have studied the symbol grounding transfer in neural networks. For example, in the model by Riga et al. (2004; see also Cangelosi et al. 2000), a network has to categorize abstract images consisting of combinations of three different shapes (square, cross, dots) and three different colours (red, green, blue). A modular dual-route neural network was used, in which the hidden layer was organized into two separated groups to produce the functional division of the hidden layer into a group dedicated to categorizing shapes and a group to classifying colours. This type of models is trained through a series of sequential stages: Prototype-sorting, Entry-Level naming, Entry-Level imitation and naming, and Higher-Level learning and Grounding transfer test. In the Prototype-sorting and Entry-Level naming stage the neural nets are initially trained to categorize and name the colour and shape of objects perceived on the retina. The first two stages consist in the direct grounding of basic categories (i.e. sensorimotor toil in Cangelosi & Harnad's model). In the third training phase (Higher-Level learning), networks acquire new higher-order categories solely through symbolic descriptions (i.e. symbolic theft in Cangelosi & Harnad's model). New categories are built by combining previously-grounded names. Each description contains the name of a shape, of a colour and the name of an object that is new to the network. The grounding test is performed at the end of training. Novel retinal stimuli, depicting previously unseen objects, were presented to the networks for the first time, in order to check if grounding had been "transferred" from directly grounded names to higher-order categories. All networks are able to pass this test, demonstrating the autonomous transfer of grounding through linguistic instructions. This symbol grounding mechanism has also been extended to simulation of embodied robotic agents, such as in the experiments by Cangelosi & Riga (submitted) described in the next section. The robotic model permits a gradual validation of the symbol grounding transfer mechanisms with more realistic scenarios and stimulus sets.

## *Communication in cognitive robots*

Simulations based on grounded adaptive agents permit investigations on the evolutionary origins of language with a reduced level of complexity in the representation of the agents' body and their physical environment. However, the increasing evidence on the role of embodiment in language (e.g. Glenberg's work on the grounding of language in action) requires a more accurate representation of the physical properties of the cognitive agents and their interaction with physical entities. This is why, more recently, we have started to use cognitive robotic systems to model language. This work has been in close collaboration with the Laboratory of Artificial Life and Robotics (Institute of Cognitive Sciences and Technology, CNR Rome), one of the leading international institutions in evolutionary robotics. The first robotic model of language evolution was developed by Davide Marocco (a researcher at CNR) during his study visit at the Adaptive Behaviour and Computation group. We used an evolutionary robotics model based on a physics dynamic simulator which permits the accurate and realistic modelling of the physical properties of the robots and environment (Marocco et al. 2003). A group of robots evolve for their ability to manipulate three-dimensional objects, whilst they can also develop a shared lexicon to talk about their tasks. This study focused on the role of social factors (e.g. communication between parent-child or between peers) and cognitive mechanisms (use of language before/after the evolution of object manipulation skills) in the emergence of linguistic production and comprehension abilities. The model supported the hypothesis that the ability to form categories from direct interaction with the environment constitutes the ground for subsequent evolution of communication and language. This model was also extended to simulate the emergence of simple syntactic categories such as action names (verbs). Comparisons between the two simulation experiments indicated show an evolutionary pattern resembling that of nouns and verbs, as observed in previous agent models on the evolution of syntactic categories (Cangelosi & Parisi 2004). Results also support the language origin hypothesis that nouns precede verbs both in phylogenesis and ontogenesis.

In a subsequent simulation on communicating robots, we used an epigenetic robotic approach to model language learning and grounding in sensorimotor abilities (Cangelosi & Riga, submitted). Epigenetic robotics is based on the training of robots through a prolonged developmental process during which complex cognitive and perceptual structures emerge as a result of the interaction of an embodied system with a physical and social environment. In our study, an imitator robot initially learns a series of basic actions and their corresponding names. Learning happens through the imitation of the movement of a demonstrator robot, already programmed to perform actions and name them. This first training stage consists in the direct grounding of the basic lexicon. The imitator robot then acquires new, higher-order behavioural abilities following linguistic interaction with the demonstrator robot or human users. Higher-order behaviour consists of composite actions (e.g. combination of the basic actions CLOSE_LEFT_ARM and CLOSE_RIGHT_ARM to perform a GRAB action). The robot is able to learn these new actions only through linguistic interactions by transferring the grounding of the basic words to the newly acquired higher-order words. This is an important demonstration of autonomous acquisition of new knowledge through linguistic instructions, with potentially valuable technological implications for the design of interactive intelligent systems.

Ongoing research is focusing on the scaling up of these robotic models. For example, new simulations are explicitly addressing the use of more structured lexicons, gradually introducing syntactic structures (Chourkadis & Cangelosi 2005). The first step consists in the ability to use arguments for the learned actions. For example, through the use of three types of objects (e.g. round spheres for balls, flat objects for books, long cylinders for sticks) it is possible to train robots to apply the same action (and word) to different objects, such as “Grab(Ball)”, “Grab(Book)”. In other studies, we are working on the scaling up of the robotic arm ability to manipulate objects (Massera et al. 2005). The use of objects with different shapes will permit the construction of a variety of linguistic categories whose representation might vary depending on the interaction between the robot’s own embodiment properties and the object motor affordances. This aims at building a computational model able to replicate the action compatibility effects (ACE) studied by Glenberg. Finally, future models might look at other aspects of communicative and social behaviour, such as the role of imitation, motivations, emotions and intentionality in linguistic communication.

## *Simulating the evolution of language*

Investigations on the evolution and emergence of language have greatly benefited from the use of computational models. Adaptive behaviour and artificial life provide useful modelling methodologies (Cangelosi & Parisi 2002) for dealing with the complexity of language and its evolutionary origins. In adaptive behaviour models, populations of autonomous agents interact via language games to exchange information about the environment. Their coordinated communication system emerges from the direct interaction between agents. Amongst the various adaptive behaviour approaches to language evolution, we have developed a methodology based on an integrative vision of language where the agent’s linguistic abilities are strictly dependent on, and grounded in, other behaviours and skills (sensorimotor, cognitive, neural, social and evolutionary factors). We call this the grounded adaptive agent modelling of the emergence of language (Cangelosi 2004; Cangelosi, Riga et al. 2005). This approach is consistent with the psychologically-plausible theories of the grounding of language in the organism’s perceptual and action systems (e.g. theories of Barsalou and Glenberg; see also Joyce et al. 2003, Coventry & Garrod 2004; Cangelosi 2005).

Numerous simulation studies of language evolution have been developed within the Adaptive Behaviour & Cognition group. In some pioneering models on the emergence of language in populations of evolving neural networks (Cangelosi 2001), we have simulated the evolution of shared animal-like lexicons and of proto-compositional languages. In the original model by Cangelosi and Parisi (1998), a population of agents evolves repertoires of communication signals (e.g. labels for naming edible and poisonous "mushrooms"). This type of model studies communication systems based on simple signal-object associations. Subsequently, a model for the emergence of compositional lexicons was developed. This focused on the emergence of languages based on symbol-symbol relationships, in addition to the simple signal-object associations. This second model was subsequently expanded by Munroe and Cangelosi (2003) to address the role of cultural variation and of learning costs in the Baldwin Effect for the evolution of language. Results showed that when there is a high cost associated with language learning, agents gradually assimilate in their genome some explicit features (e.g. lexical properties) of the specific language they are exposed to. When the structure of the language is allowed to vary using a process of cultural transmission, Baldwinian processes cause, instead, the assimilation of a predisposition to learn, rather than any structural properties associated with a specific language. The analysis of the mechanisms underlying such a predisposition in terms of categorical perception (see Cangelosi 2005; Cangelosi et al. 2000), supported the evolutionary hypothesis on the Baldwinian inheritance of general underlying cognitive capabilities that serve language acquisition. This is in opposition to the thesis that argues for assimilation of structural properties needed for the specification of a fully blown language acquisition device.

More studies have been dedicated to the evolution of proto-syntactic categories, such as that of nouns and verbs. In our simulation, we use a restricted meaning for such grammatical terms, with "nouns" referring to names used by agent to refer to perceptual objects, and "verbs" being the names of actions. In particular, the goal has been to investigate the role of the organisation of sensorimotor knowledge (and the corresponding neural substrates) in the evolutionary transition towards compositional languages. This is an important issue in the field of language origins, because it can shed light on the evolutionary emergence and diversification of word categories, such as nouns, verbs and functions words. For example, in a recent grounded agent model of the evolution of nouns and verbs (Cangelosi & Parisi 2004), two techniques are used to analyse the internal representations of language and categories: (i) categorical perception analyses of the internal similarity space and (ii) synthetic brain imaging, a new technique also utilized to compare neural models with empirical brain imaging studies. The simulation also uses two different architectures for the agents' neural network controller (Cangelosi 2004). One uses a fully-distributed architecture, whilst the second has a brain-inspired modular organisation in which different regions (hidden layers) separately perform sensory processing and sensorimotor integration. Analyses show that the neural processing of verbs is consistently localized in the regions of the networks that perform sensorimotor integration, while nouns are associated with sensory processing areas.

In another grounded adaptive agent model of language evolution, Cangelosi and Harnad (2000) use the "mushroom metaphor" foraging task to test the language origin hypothesis of Symbolic Theft. The model considered two ways of learning categories: (i) "sensorimotor toil" when new categories are acquired through real-time, feedback-corrected, trial and error experience in sorting them; and (ii) "symbolic theft" when new categories are acquired by hearsay from propositions. In competition, symbolic theft always beats sensorimotor toil.  We hypothesize that this is the basis of the adaptive advantage of language. Entry-level categories must still be learned by toil, however, to avoid an infinite regress (the "symbol grounding problem"). Evolutionary simulations showed that populations using the symbolic theft strategy always beat the sensorimotor toilers when in competition. In addition, the analysis of the similarity space of internal categorical representations show that the compression/expansion effects of categorical perception can also arise from symbolic theft alone. In fact, internal representations are better organised for categories acquired via symbolic theft, compared to those via sensorimotor toil. The picture of natural language and its origins that emerges from this analysis is that of a powerful hybrid symbolic/sensorimotor capacity, infinitely superior to its purely sensorimotor precursors, but still grounded in and dependent on them. It can spare us from untold time and effort learning things the hard way, through direct experience, but it remain anchored in and translatable into the language of experience.

## *Grounding spatial terms and quantifiers in perception*

Expressions involving spatial prepositions in English convey to a hearer where one object (located object) is located in relation to another object (reference object). For example, in the sentence "the coffee is in the cup", the coffee is understood to be located with reference to the cup in the region denoted by the preposition *in*. Spatial descriptions can also specify the relative position of objects in relation to the speaker and/or hearer point of view, such as in experiments on frames of reference. In a similar fashion, the selection and use of linguistic quantifiers, such as *few*, *a few*, *some*, *many*, has the role of highlighting important contextual factors and the pragmatics of communication between speaker and hearer. For example, in the two sentences "A few went to the cinema. They liked the movie" and "Few went to the cinema. They preferred the restaurant", the selection of the quantifier *a few* vs. *few* indicate the focus of attention (we use *a few* when we refer to those that went to the cinema, *few* when we want to talk of those that did not go). Understanding the meaning of such terms is important as they are among the set of closed class terms which are generally regarded as having the role of acting as organising structure for further conceptual material. Furthermore, from the semantic point of view, these terms have the virtue of relating in some way to visual scenes being described. Hence, it should be possible to offer more precise semantic definitions of these, as opposed to many other expressions, because the definitions can be grounded in perceptual representations.

In spatial language and cognition research, most approaches to spatial prepositions have focused on the geometric factors ("where" object are) and assumed that they only require coarse grained properties of the objects involved as constraints on their use. Yet there is now much evidence (Coventry & Garrod 2004) that "what" objects are influences how one talks about “where” they are. To establish empirically the relative extent to which various constraints on object knowledge (“what” information, such as object properties and function) and on object localisation (“where” information related to geometric factors) influence the comprehension of a range of linguistic terms, a new connectionist model of language grounding was developed, in parallel with experimental investigations upon which the model mapped. More generally, the use of a grounded connectionist approach aims at bridging the gap between theories of meaning which capture meaning in terms of symbol-symbol relations (e.g. those based on Latent Semantic Analyses) versus those which ground language directly in perceptual representation (Coventry & Garrod, 2004).

We have developed a computational model for the perceptual grounding of spatial prepositions consisting of three main modules: Vision Processing, Elman Network, and Dual-Route Network (Cangelosi et al. 2005). The first module uses a series of Ullman-type visual routines to identify the constituent objects of a visual scene. The visual input consists of 60-second movies showing a located object (e.g. teapot) pouring a liquid (water) into a reference object (cup). The Elman network module utilises the output information from the vision module to produce a compressed neural representation of the dynamics of the scene, such as the movement of liquid flow between the reference and located objects (Joyce et al. 2003). This compressed representation is given in input to the dual-route vision-language neural network (Cangelosi et al., 2000) to produce a judgement regarding the appropriate spatial terms describing the visual scene. The dual-route network is the core component of the model. It integrates visual and linguistic knowledge to produce a verbal description of the visual scene such as “Teapot over Cup”. Some output nodes encode the names of objects and others the four spatial prepositions. The nodes with the highest activation permit the selection of the words corresponding to the two objects present in the scene and the selection of the preposition that best described their spatial relation. In addition, the activation values of the linguistic output nodes correspond to rating values given by subjects for the prepositions under, below, above, over. The multi-layer perceptron was trained via error backpropagation, using rating data collected during experiments (Coventry et al., submitted). Some of the ratings are also used for the generalisation test. Simulation results clearly indicate that the networks produce rating values similar to that of experimental subjects. In addition, the model has accurately predicted the ratings in new experiments requiring subjects to imagine the end result of the teapot/liquid/cup scenes (Cangelosi, Coventry et al. 2005).

This model is currently being extended to deal with further linguistic terms, namely vague quantifiers such as *some, few, a few, lots of*. The hypothesis is that this grounded connectionist approach will permit the identification of the main mechanisms responsible for quantification judgment and their linguistic expression. Vague quantifiers like *a few* and *several* exhibit many of the same context effects that have been observed for spatial prepositions. For example, relative size of objects and their expected frequency have both been shown to affect the comprehension of quantifiers. “A few cars” is associated with a smaller number than “a few crumbs”. In addition, new experiments (Coventry et al. 2005) have demonstrated that the rating of vague quantifiers is affected by the extent to which objects are grouped together and the degree of spacing between objects. The issue we are exploring with the new model is that these context effects originate from visual processing constraints.

Therefore, information regarding specific numbers of objects in a scene cannot be derived very easily from visual processing of that scene. New simulations have focused on the grounding of numerosity judgments as a basis for linguistic quantifier use (Rajapakse et al. 2005, in press). Preliminary experiments show that the part of the model trained to judge "psychological number" uses some of the same factors known to play a major role in the production of quantification judgments in human subjects. This simulation further supports the ongoing development of a psychologically-plausible model of language which uses the contextual factors such as object properties and their functionality.

## *Summary*

The research activity of the Adaptive Behaviour and Cognition group in Plymouth has focuses on the modelling of artificial cognitive systems able to ground language in their own cognitive and sensorimotor capabilities. This is achieved through the design and evolution of cognitive systems, such as simulated agents and robots, that are able to develop linguistic knowledge through embodied interaction with their environments and communication with other agents. Numerous experimental and modelling studies have permitted a deeper understanding of the relationship between language, cognition and action, such as the role of sensorimotor representation in supporting the emergence of linguistic structure.

**Angelo Cangelosi**
Adaptive Behaviour and Cognition Research Group

University of Plymouth (UK)
acangelosi@plymouth.ac.uk
http://www.tech.plym.ac.uk/soc/research/ABC

## *Cited References from the Adaptive Behaviour and Cognition Research Group*

Cangelosi A. (2001). Evolution of communication and language using signals, symbols and words. *IEEE Transactions on Evolutionary Computation*, 5(2), 93-101

Cangelosi A. (2004). The sensorimotor bases of linguistic structure: Experiments with grounded adaptive agents. In S. Schaal et al. (Eds.), Proceedings of the Eighth International Conference on the Simulation of Adaptive Behaviour: From Animals to Animats 8, Cambridge MA, MIT Press, pp. 487-496

Cangelosi A. (2005). Grounding symbols in perceptual and sensorimotor categories: Connectionist and embodied approaches. In H. Cohen & C. Lefebvre (Eds), *Categorization in Cognitive Science*, Elsevier

Cangelosi A., Bugmann G., Borisyuk R. (Eds.) (2005). Modelling Language, Cognition and Action: Proceedings of the 9th Neural Computation and Psychology Workshop. Singapore: World Scientific.

Cangelosi A., Coventry K.R., Rajapakse R., Joyce D., Bacon A., Richards L., Newstead S. (2005), Grounding language in perception: A connectionist model of spatial terms and vague quantifiers. In A. Cangelosi, G. Bugmann, R. Borisyuk (eds.), 2005, pp. 47-58

Cangelosi A., Greco A. & Harnad S. (2000). From robotic toil to symbolic theft: Grounding transfer from entry-level to higher-level categories. Connection Science, 12(2), 143-162.

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. (1998). The emergence of a "language" in an evolving population of neural networks. *Connection Science*, 10(2), 83-97

Cangelosi A., Parisi D. (Eds.) (2002). Simulating the Evolution of Language. London: Springer.

Cangelosi A., Parisi D. (2004). The processing of verbs and nouns in neural networks: Insights from synthetic brain imaging. *Brain and Language*, 89(2), 401-408

Cangelosi A., Riga T. (submitted). An epigenetic robotic model for the sensorimotor grounding of language. *Cognitive Science*

Cangelosi A., Riga T., Giolito B., Marocco D. (2005). The emergence of language in grounded adaptive agents and robots. In K. Hasida & K. Nitta (eds.), *New Frontiers in Artificial Intelligence: Joint Proceeding of the 17th and 18th Annual Conferences of the Japanese Society for Artificial Intelligence*. LNCS, Springer Verlag

Chourkadis E., Cangelosi A (2005). Grounding transfer in autonomous robots. *22nd Annual Workshop of the European Society for the Study of Cognitive Systems*, London.

Coventry K.R., Cangelosi A., Newstead S., Bacon A., Rajapakse R. (2005). Grounding natural language quantifiers in visual attention. *XXVII Annual meeting of the Cognitive Science Society*, Stresa, July 2005

Coventry K.R., Cangelosi A., Rajapakse R., Bacon A., Newstead S., Joyce D., Richards L.V. (2005). Spatial prepositions and vague quantifiers: Implementing the functional geometric framework. In C. Freksa, B. Nebel, M. Knauff & B. Krieg-Bruckner (Eds.), *Spatial Cognition, Volume IV*. Lecture notes in Artificial Intelligence. Springer Verlag.

Coventry, K.R., Garrod, S. C. (2004). Saying, Seeing and Acting. The Psychological Semantics of Spatial Prepositions. Lawrence Erlbaum Associates. Essays in Cognitive Psychology Series.

Coventry K.R., Richards L., Joyce D., Cangelosi A. (submitted). Towards a psychological plausible model of spatial language processing. *Journal of Memory and Language*

Holden R., Cangelosi A. (2004). A general, computationally intelligent model for egress simulation. *Proceedings of the 10th International Conference on Fire Science and Engineering: Interflam 2004*. London: Interscience Communication Publishing, pp. 387-398

Holden R., Cangelosi A. (2005). A multi-agent based fitness function for evolutionary architecture. *IEEE International Conference on Integration of Knowledge Intensive Multi-Agent Systems: KIMAS'05: Modelling, Evolution, and Engineering*, Boston

Joyce D.W., Richards L.V., Cangelosi A., Coventry, K. R. (2003). On the foundations of perceptual symbol systems: Specifying embodied representations via connectionism. In F. Dretje, D. Dorner & H. Schaub (Eds.), *The Logic of Cognitive Systems. Proceedings of the Fifth International Conference on Cognitive Modelling*, pp147-152. Universitats-Verlag Bamberg, Germany.

Marocco D., Cangelosi A., Nolfi S. (2003). The emergence of communication is evolutionary robots. *Philosophical Transactions of the Royal Society London – A*, 361: 2397-2421

Massera G., Nolfi S., Cangelosi A. (2005). Evolving a simulated robotic arm able to grasp objects. In A. Cangelosi, G. Bugmann, R. Borisyuk (eds.), 2005, pp. 203-208

Munroe S., Cangelosi A. (2002). 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., Cangelosi A., Falcetta I. (2002). Verbs, nouns and simulated language games. *Italian Journal of Linguistics*, 14(1): 99-114

Rajapakse R.K., Cangelosi A., Coventry K., Newstead S. and Bacon A. (2005). Grounding linguistic quantifiers in perception: Experiments on numerosity judgments. *2nd Language & Technology Conference: Human Language Technologies as a Challenge for Computer Science and Linguistics*. April 21-23, 2005, Poznañ, Poland

Rajapakse R.K., Cangelosi A., Coventry K.R, Newstead S., Bacon A. (in press), Connectionist modeling of linguistic quantifiers. *ICANN 2005 International Conference on Artificial Neural Networks,* Warsaw, Poland September 11-15, 2005

Riga T., Cangelosi A., Greco A. (2004). Symbol grounding transfer with hybrid self-organizing/supervised neural networks. *IJCNN04 International Joint Conference on Neural Networks*. Budapest, July 2004.