International Journal of Signs and Semiotic Systems

January – June 2011, Vol. 1, No. 1

Table of Contents

Solutions and Open Challenges for the Symbol Grounding Problem

Angelo Cangelosi, University of Plymouth, UK

ABSTRACT

This article discusses the current progress and solutions to the symbol grounding problem and specifically identifies which aspects of the problem have been addressed and issues and scientific challenges that still require investigation. In particular, the paper suggests that of the various aspects of the symbol grounding problem, the transition from indexical representations to symbol-symbol relationships requires the most research. This analysis initiated a debate and solicited commentaries from experts in the field to gather consensus on progress and achievements and identify the challenges still open in the symbol grounding problem.

Keywords: *Cognitive Robotics, Developmental Robotics, Embodiment, Symbol Grounding, Symbol Relationships*

MAIN TEXT

The grounding of symbols, i.e. the process by which symbols (such as words) are linked to the agent's own semantic representation of the world, is an issue of crucial importance to the field of cognitive science, raging from philosophy and semiotics, to artificial intelligence, and to cognitive robotics. For example, in cognitive and developmental robotics, in the last decades there has been a tremendous increase in new models of the evolution of communication and the developmental acquisition of language. These models directly regard the issue of symbol grounding, since robotic agents are trained to use symbols that often refer to entities and states in the robot's own world.

In the literature on artificial intelligence, cognitive science and philosophy, there has been extensive discussion on the symbol grounding problem. Different views on the importance of the symbol grounding have been proposed. On one extreme, we find the so called "symbolic" approaches that practically ignore the cognitive significance of such an issue (Fodor, 1983; Chomsky, 1965). This is the case of classical GOFAI (Good Old Fashion Artificial Intelligence) approaches, based on symbolic methods such as logic, where the meaning or a linguistic symbol (word) is simply represented by another symbol (meaning), without having to deal of how meanings have formed and have been associated to words. On the other hand, new "embodied" approaches to cogni-

DOI: 10.4018/ijsss.2011010104

tion acknowledge the importance of grounding symbols through the agent's interaction with the environment. However, even in this field, there have been suggestions that the problem has practically been solved (Steels, 2008), thus appearing to suggest that the symbol grounding problem is not important anymore.

In this target article, I am going to suggest a different analysis of the current progress on symbol grounding, indentify specifically which aspects of the problem have been already addressed, and which other issues still require investigation. This analysis will initiate a debate and solicit commentaries from experts in the field, to gather consensus on actual progress, achievements and open challenges in the symbol grounding problem.

To assess better the current state of the art on the Symbol Ground Problem, I will use the definition and discussion of the problem originally given by Stevan Harnad in his seminal 1990 article "The Symbol Grounding Problem". Here Harnad explains that the symbol grounding problem refers to the capability of natural and artificial cognitive agents to acquire an *intrinsic link* (*autonomous*, we would say in nowadays robotics terminology) between internal symbolic representations and some referents in the external word or internal states. In addition, Harnad explicitly proposes a definition of a symbol that requires the existence of logical links (e.g. syntactic) between the symbols themselves. It is thanks to these inter-symbol links, its associated symbol manipulation processes, and the symbol grounding transfer mechanism (Cangelosi & Riga, 2006) that a symbolic system like human language can exist. The symbol-symbol link is the main property that differentiates a real *symbol* from an *index*, as in Peirce's semiotics. These symbolic (e.g. syntactic) links also support the phenomena of productivity and generativity in language, and contribute to the grounding of abstract concepts and symbols (Barsalou, 1999). Finally, an important component of the symbol grounding problem is the social and cultural dimension, that is, the role of social interaction in the sharing of symbols (the external/social symbol grounding problem, as in Cangelosi, 2006; Vogt, 1997; Vogt & Divina, 2007).

To summarise, we can say that there are *three sub-problems* in the development of a grounded symbol system:

(i) How can a cognitive agent autonomously *link symbols to referents* in the world such as objects, events and internal and external states;
(ii) How can an agent autonomously create a set of *symbol-symbol relationships* and the associated transition from an indexical system to a proper symbol system;
(iii) How can a society of agents autonomously develop a *shared set of symbols*.

When we look at the current cognitive agent models of language learning (Cangelosi & Parisi, 2002; Wagner et al., 2003; Lyon et al., 2007), I believe that much progress has been achieved in the first and third sub-problems (individual and social grounding of symbols), whilst I believe that for the second sub-problem (autonomous creating of symbol-symbol relationships) much still needs to be done. I therefore agree with Steels (2008) that much has been done on the robotic and cognitive modelling of the symbol grounding problem, but only when we consider the two sub-problems (i) and (iii), "we now understand enough to create experiments in which groups of agents self-organize symbolic systems that are grounded in their interactions with the world and others" (Steels, 2008, p. 240).

For example, when we consider the well known "Talking Heads" model of the emergence of communication in robotic agents (Steels, 1999; Kaplan & Steels, 2002), we have a demonstration that each individual robot can autonomously create categorical representations of the world (a white board with a variable combination of coloured shapes). Using the technique of discrimination trees, each agent creates categories defined by specific patters of boundaries of color, shape and spatial values. In addition, these perceptual categories become associated with words that the agents use to communicate with each other in "language games" (Steels, 2001). That is, these categories constitute the internal meanings of the agent's communication lexicon. What is crucial in this model is that the association of a specific perceptual category with a word is mediated by the social interaction with the other agents. This therefore leads to the cultural evolution of a shared set of words (symbols). Such a modelling approach has been extended to human-robot communication systems, such as for the training of the AIBO pet dog to name objects (Kaplan & Steels, 2000).

In a multi-agent model of the evolutionary emergence of communication, similar mechanisms for the grounding of shared lexicons have been utilised (Cangelosi, 2001). Simulated cognitive agents live in a 2D world to perform foraging tasks and learn to differentiate between edible mushrooms (that when eaten increase the chances of survival and reproduction) and toadstools (which cause a decrease of fitness). Agents autonomously learn to categorize foods in edible and inedible categories through their own neural network that controls the agent's perceptual and motor system. In addition, as agents are allowed to communicate with each-other, this cause the gradual evolution of a shared lexicon. The activation patterns in the agents' neural network (hidden units) constitute the distributed semantic representations of the food categories, and at the same time they support the linguistic comprehension and production capability of the agents. In fact, analyses of the agent's internal representations suggest that the evolutionary pressure to evolve a shared lexicon is explained by the categorical perception effects of these representations.

The above examples are only some of the numerous cognitive agents and robotics models of the evolution and acquisition of language that have demonstrated the mechanisms for the autonomous grounding of symbols and the emergence of shared lexicons (Loula et al., 2010; Marocco et al., in press; Roy, 2005; Vogt, 2005; Vogt & Divina, 2007). At the same time, as Steels also acknowledges, it is also true that we do not yet have a full understanding of all the mechanisms in grounding, such as on the nature, role and making of internal symbolic representations.

As for the sub-problem (ii), i.e. the transition from a communication systems based on indices (e.g. labels, animal communication, early child language learning) to that of a full symbolic system (e.g. adult human languages), I believe that the problem has not really been solved at all, and much needs to be done.

Most computational models of syntactic learning and evolution, based on the grounding principles as defined above, have only loosely addressed the issue of the autonomous development of symbol-symbol relationships. In very few cases, grounded models of the evolution of communication use limited syntactic constructs that partially resemble grammatical constructs (Cangelosi, 2001; Cangelosi et al., 2000). In other cases, robotics models of the evolution of language use pre-defined, symbolic approaches to represent symbol-symbol relationships, e.g. by assuming the pre-existence of semantic and syntactic categories in the agent's cognitive system. This is however in contrast with the grounding principles, as the aim of a cognitive agent model of the evolution of language is to discover the evolutionary, neural and cognitive mechanisms that lead to the emergence of syntactic symbol-symbol representations.

I invite my colleagues to comment on the state of the art on the symbol grounding problem in cognitive science and robotic models of communication and language, and on their view on the importance of the symbol grounding problem. I suggest below some open challenges for future research that I believe are crucial for our understanding of the symbol grounding phenomena, and I welcome suggestions for other important, unsolved challenges in this field.

(1) Is the symbol grounding problem, and the three sub-problems as identified above, still a real crucial issue in cognitive robotics research? And if the problem appears

to have been solved, as some have suggested, why is it that so far we have failed at building robots that can learn language like children do?

(2) What are the processes that lead to the transition from indexical communication system to a full symbolic system such as language? Is there a continuum between indices (labels) and symbols (words), or is the transition qualitative and sudden? What known phenomena in language origins theories, and in developmental studies, should be included in developmental and evolutionary robotics model of language?

(3) Notwithstanding the importance of the grounding problem, there are still various approaches in the agent/robot language learning/evolution literature that practically ignore the process of grounding and use a symbolic-only approach to the definition of meanings and words. Do these symbolic approaches really give an important contribution to our understading of human cognition, of should all models of language learning be based solely on grounding mechanisms?

(4) Does cognitive development really plays an important role in symbol grounding and acquisition, or is it just an epiphenomenon of no crucial importance to the understanding of human cognition? Some key findings and experiments show that infants have strong specific biases that allow them to learn very easily language. And most attempts at building robots without these biases have failed so far to learn realistically complex concepts/semantic categories. Is the symbol grounding problem just a matter of using and identifying such biases in robotics language models?

(5) What kind of robotic experiment would constitute a real breakthrough to advance the debate on symbol grounding, and what kind of principle and ideas are still unexplored?

(6) What are the properties and differences of internal representations beyond both indexical and symbolic systems? Or are representation issues not really crucial, as a pure sensorimotor modelling approach would not require any internal representation capability?

(7) How can we model the grounding of abstract concepts such as beauty, happiness, and time. Or is the grounding approach inconsistent with the study of higher-order symbolic capabilities?

(8) What are the grounding components in the acquisition and use of function words (such as verb preposition "to", as in verbs "to go", "if", "the"), of number concepts or words, and of morphology and other syntactic properties.

(9) How can we model the grounding phenomena studies through empirical investigations of language embodiment (Barsalou, 1999; Glenberg & Kaschak, 2002; Pecher & Zwaan, 2005)?

ACKNOWLEDGMENTS

The author wishes to acknowledge the support of the EU FP7 projects ITALK and ROBOTDOC. A shorter version of this target article, and some of the commentaries here included (Harnad, Steels, Madden et al. Sloman, Muller, Cowley), have appeared in the *AMD Newsletter* (vol. 7, April 2010) of the IEEE-CIS Autonomous Mental Development Technical Committee http://www.cse.msu.edu/amdtc/amdnl/

REFERENCES

Barsalou, L. (1999). Perceptual symbol systems. *The Behavioral and Brain Sciences*, *22*, 577–609.

Cangelosi, A. (2001). Evolution of communication and language using signals, symbols and words. *IEEE Transactions on Evolutionary Computation*, *5*(2), 93–101. doi:10.1109/4235.918429

Cangelosi, A. (2006). The grounding and sharing of symbols. *Pragmatics & Cognition*, *14*(2), 275–285. doi:10.1075/pc.14.2.08can

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. doi:10.1080/09540090050129763

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

Cangelosi, A., & Riga, T. (2006). An embodied model for sensorimotor grounding and grounding transfer: Experiments with epigenetic robots. *Cognitive Science*, *30*(4), 673–689. doi:10.1207/s15516709cog0000_72

Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: The MIT Press.

Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: Bradford Books, MIT Press.

Glenberg, A., & Kaschak, M. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, *9*(3), 558–565. doi:10.3758/BF03196313

Harnad, S. (1990). The Symbol Grounding Problem. *Physica D. Nonlinear Phenomena*, *42*, 335–346. doi:10.1016/0167-2789(90)90087-6

Loula, A., Gudwin, R., El-Hani, C., & Queiroz, J. (2010). Emergence of self-organized symbol-based communication in artificial creatures. *Cognitive Systems Research*, *11*(2), 131–147. doi:10.1016/j.cogsys.2008.10.002

Lyon, C., Nehaniv, C. L., & Cangelosi, A. (Eds.). (2007). *Emergence of Communication and Language*. London: Springer. doi:10.1007/978-1-84628-779-4

Marocco, D., Cangelosi, A., Fischer, K., & Belpaeme, T. (in press). Grounding action words in the sensorimotor interaction with the world: Experiments with a simulated iCub humanoid robot. *Frontiers in Neurorobotics*.

Pecher, D., & Zwaan, R. A. (2005). *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thinking*. Cambridge, MA: Cambridge University Press. doi:10.1017/CBO9780511499968

Roy, D. (2005). Grounding words in perception and action: Insights from computational models. *Trends in Cognitive Sciences*, *9*(8), 389–396. doi:10.1016/j.tics.2005.06.013

Steels, L. (1999). The Talking Heads Experiment: *Vol. 1. Words and Meaning*. Antwerp: Laboratorium.

Steels, L. (2001). Language games for autonomous robots. *IEEE Intelligent Systems*, *16*(5), 16–22.

Steels, L. (2008). The symbol grounding problem has been solved. So what's next? In de Vega, M. (Ed.), *Symbols and Embodiment: Debates on Meaning and Cognition* (pp. 223–244). Oxford, UK: Oxford University Press.

Steels, L., & Kaplan, F. (2000). AIBO's first words, the social learning of language and meaning. *Evolution of Communication*, *4*(1), 3–32. doi:10.1075/eoc.4.1.03ste

Steels, L., & Kaplan, F. (2002). Bootstrapping grounded word semantics . In Briscoe, T. (Ed.), *Linguistic evolution through language acquisition: formal and computational models* (pp. 53–73). Cambridge, UK: Cambridge University Press. doi:10.1017/CBO9780511486524.003

Vogt, P. (2002). Physical symbol grounding. *Cognitive Systems Research*, *3*(3), 429–457. doi:10.1016/S1389-0417(02)00051-7

Vogt, P. (2005). The emergence of compositional structures in perceptually grounded language games. *Artificial Intelligence*, *167*(1-2), 206–242. doi:10.1016/j.artint.2005.04.010

Vogt, P., & Divina, F. (2007). Social symbol grounding and language evolution. *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, *8*(1), 31–52. doi:10.1075/is.8.1.04vog

Wagner, K., Reggia, J. A., Uriagereka, J., & Wilkinson, G. S. (2003). Progress in the simulation of emergent communication and language. *Adaptive Behavior*, *11*(1), 37–69. doi:10.1177/10597123030111003

Angelo Cangelosi is Professor of Artificial Intelligence and Cognition at the University of Plymouth (UK), where he leads the Centre for Robotics and Neural Systems. He obtained a PhD in psychology and computational modeling in 1997 from the University of Genoa, whilst also working as visiting scholar at the National Research Council (Rome), the University of California San Diego (USA) and the University of Southampton (UK). He has produced more than 140 scientific publications, and has been awarded numerous research grants from UK and international funding agencies. He is coordinator of the EU FP7 project "ITALK" (italkproject.org) and the "RobotDoC" Marie Curie doctoral network in developmental robotics (robotdoc.org).

RESEARCH COMMENTARY

Research Commentaries on Cangelosi's "Solutions and Open Challenges for the Symbol Grounding Problem"

SYMBOL GROUNDING PROBLEM: TURING-SCALE SOLUTION NEEDED

Stevan Harnad,
University of Quebec at Montreal, Canada,

and University of Southapton, UK

Toys. The symbol grounding problem is the problem of causally connecting symbols inside an autonomous system to their referents in the external world without the mediation of an external interpreter. The only way to avoid triviality, however, is to ensure that the symbols in question, their referents in the world, and the dynamic capacities of the autonomous system interacting with the world are nontrivial. Otherwise a toy robot, with exactly two symbols – go/stop – is "grounded" in a world where it goes until it bumps into something and stops.

DOI: 10.4018/ijsss.2011010105

Turing. From the very outset, the symbol grounding problem – which was inspired and motivated by Searle's Chinese Room Argument – was based on the Turing Test, and hence on a system with *full, human-scale linguistic capacity*. So it is the words of a full-blown natural language (not all of them, but the ones that cannot be grounded by definition in the others) that need to be connected to their referents in the world. Have we solved that problem? Certainly not. Nor do we have a robot with Turing-scale capacities, either symbolic or sensorimotor (with the former grounded, embodied, in the latter). Designing or reverse-engineering an autonomous system with this Turing-scale robotic and linguistic capacity, and thereby causally explaining it, is the ultimate goal of cognitive science. (Grounding, however, is not the same as meaning; for that we would also have to give a causal explanation of consciousness, i.e., feeling, and that, unlike passing the Turing Test, is not just hard but hopeless.)

Totality. Grounded robots with the sensorimotor and learning capacities of subhuman animals might serve as waystations, but the gist of the Turing methodology is to avoid being fooled by arbitrary fragments of performance capacity. Human language provides a natural totality. (There are no partial languages, in which you can say this, but not that.) We are also extremely good at "mind-reading" human sensorimotor performance capacity for tell-tale signs of mindlessness; it is not clear how good we are with animals (apart perhaps from the movements and facial expressions of the higher mammals).

Terms. There are certain terms (or concepts) I have not found especially useful. It seems to me that real objects, plus (internal or external) (1) analogs or iconic copies of objects (similar in shape) and (2) arbitrary-shaped symbols in a formal symbol system (such as "x" and "=", or the words in a language, apart from their iconic properties), systematically interpretable as referring to objects, are entities enough. Peirce's "icon/index/symbol" triad seems one too many. Perhaps an index is just a symbol in a toy symbol system. In a formal symbol system the links between symbols are syntactic whereas the links between internal symbols and the external objects that they are about are sensorimotor (hence somewhat iconic). And inasmuch as symbols inside a Turing-scale robot are linked to object categories rather than to unique (one-time, one-place) individuals, all categories are abstract (being based on the extraction of sensorimotor invariants), including, of course, the category "symbol." The rest is just a matter of degree-of-abstraction. Even icons are abstract, inasmuch as they are neither identical nor co-extensive with the objects they resemble. There are also two sorts of productivity or generativity: syntactic and semantic. The former is just formal; the latter is natural language's power to express any and every truth-valued proposition.

Talk. Yes, language is fundamentally social in that it would never have bothered to evolve if we had been solitary monads (even monads born mature: no development, just cumulative learning capacity). But the nonsocial environment gives enough corrective feedback for us to learn categories. Agreeing on what to call them is trivial. What is not trivial is treating symbol strings as truth-valued propositions that describe or define categories.

REFERENCES

Harnad, S. (2003). Can a Machine Be Conscious? How? *Journal of Consciousness Studies*, *10*(4-5), 69–75.

Harnad, S. (2005). *To Cognize is to Categorize: Cognition is Categorization*. In C. Lefebvre & H. Cohen (Eds.), *Handbook of Categorization*. Amsterdamn, The Netherlands: Elsevier. Retrieved from http://eprints.ecs.soton.ac.uk/11725/

Harnad, S. (2008). The Annotation Game: On Turing (1950) on Computing, Machinery and Intelligence. In R. Epstein & G. Peters (Eds.), *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer*. New York: Springer. Retrieved from http://eprints.ecs.soton.ac.uk/7741/

Harnad, S., & Scherzer, P. (2007). First, Scale Up to the Robotic Turing Test, Then Worry About Feeling. In *Proceedings of the 2007 Fall Symposium on AI and Consciousness,* Washington, DC. Retrieved from http://eprints.ecs.soton.ac.uk/14430/

Picard, O., Blondin-Masse, A., Harnad, S., Marcotte, O., Chicoisne, G., & Gargouri, Y. (2009, December 11-12). Hierarchies in Dictionary Definition Space. In *Proceedings of the 23rd Annual Conference on Neural Information Processing Systems: Workshop on Analyzing Networks and Learning With Graphs,* Vancouver, BC, Canada. Retrieved from http://eprints.ecs.soton.ac.uk/18267/

WORK ON SYMBOL GROUNDING NOW NEEDS CONCRETE EXPERIMENTATION

Luc Steels, Vrije Universiteit Brussel, Belgium, and Sony Computer Science Laboratory, France

A few years ago, I wrote a paper called "The symbol grounding problem has been solved. So What's Next?" (Steels, 2008) based on a talk I gave in 2005 at the very interesting Garachico workshop on "Symbols, Embodiment, and Meaning" organized by Manuel de Vega, Arthur Glenberg and Arthur Graesser. Angelo Cangelosi, who himself has made several contributions in this area (Cangelosi & Harnard, 2000), responded to this paper calling for further discussion within the developmental robotics community and I of course welcome such dialogs which are all too rare in our field. Cangelosi decomposed the symbol grounding problem in three subproblems: (1) How can a cognitive agent autonomously link symbols to referents in the world such as objects, events and internal and external states? (2) How can an agent autonomously create a set of symbol-symbol relationships and the associated transition from an indexical system to a proper symbol system? and (3) How can a society of agents autonomously develop a shared set of symbols? I will take the same decomposition to argue again why I believe the symbol grounding problem has been solved.

Let's start with (2). AI has developed over the past decade a large array of mechanisms for dealing with symbol-symbol relationships. It suffices to open up any textbook, particularly in machine learning, and you will find effective algorithms that will induce symbolic relationships and inferential structures from sets of symbols (usually texts but they can also be more formalized representations of knowledge as in the semantic web). The success of this body of techniques is undeniable when you look at search engines such as Google which are entirely based on them. This problem can therefore be considered solved, or at least sufficiently solved that applications can be built on a large scale.

The value of the philosophical discussion initiated by Searle and Harnad was to point out that there is more to symbol usage than symbol-symbol linking, in particular there is problem (1): How can symbols be grounded in sensory-motor experiences. Searle and Harnad are right, but also this question was already addressed by early AI systems such as Shakey, built in the early seventies by Nilsson et al. (1984). The Shakey experiment was unfortunately burried in technical papers which were not accessible to non-experts and its significance has therefore been missed by philosophers. Shakey had a cognitive architecture designed to trigger embodied grounded robot behaviors through natural language commands. Given the state of the art in computing, robotics, visual perception, etc., the response of the robot was slow and the complexity of what could be handled limited, but the fundamental principle how to ground symbols through feature extraction, segmentation, pattern recognition, etc. was clearly demonstrated and many subsequent experiments have confirmed it.

Philosophers undoubtly object that it is the designers that take care of the grounding. Although the robot autonomously carries out grounding, it does not autonomously acquire the competence to do so. This criticism is right, but I believe this problem has been solved as well. What we need to do is to set up cooperative interactions between agents in which symbols are useful. For example, I am in a store and want to get from you (the salesperson) a brown T-shirt (and not the white, green, yellow, or blue one). And suppose that I cannot point to the T-shirt because they are not on display. Then if I have a way to categorize colors and associate names with these categories, and you have a way to see what category I mean and use that to find the T-shirt back in the world, then I can get the T-shirt I want and you are happy because you can sell it to me. Similar situations can be set up with robots, for example because they need to draw attention to an object in the world, or get the other one to do an action, or inform each other about an event that took place, etc. The key point is of course that in these experiments neither the categorizations of reality nor the symbols themselves should be supplied by designers. They should be invented by the agents themselves. And this is precisely what the language games experiments, from the first one reported in Steels (1995), have shown. The key idea is to establish a coupling between success in a task (where communication is critical or at least useful) and the grounding and linguistic processes that are potentially able to handle symbols.

Interestingly enough, the solution to (2) implies a solution to (3). Producers and interpreters must negotiate a tacit agreement how they are going to conceptualize the world for language and how they are going to name and express concepts otherwise the interaction will not be successful. The solution, again demonstrated for the first time in Steels (1995) but since then applied on a grand scale in dozens of robotic experiments, is to monitor success in communication and adapt or shift concepts and symbolic conventions based on that so that conceptual and symbolic inventories gradually align. When this is done systematically by all agents for every language game they play, a process of self-organization towards a shared system is set in motion, without the need for the intervention of a human designer to coordinate symbol use. And this solves subproblem (3).

What should we do now? Work on the specifics, particularly if you are interested in mental development: For example, how can concepts involved in the representation of time arise and get symbolized as tense and aspect? How can a system for categorizing the roles of participants in events (agent, patient, etc.) emerge? Why and how can modality systems to express the attitude of speakers towards information arise? And so on. Much work is left to do but chasing philosophical chimera is not one of them.

REFERENCES

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. doi:10.1075/eoc.4.1.07can

Nilsson, N. J. (1984). *Shakey The Robot* (Tech. Rep. 323). Menlo Park, CA: AI Center, SRI International. Retrieved from http://www.ai.sri.com/shakey/

Steels, L. (1995). A self-organizing spatial vocabulary. *Artificial Life Journal, 2*(3), 319-332.

Steels, L. (2008). The Symbol Grounding Problem has been solved. So What's Next? In A. Glenberg, A. Graesser, & M. de Vega (Eds.), *Symbols, Embodiment and Meaning. Oxford University Press* (pp. 506-557).

SYMBOL GROUNDING: SHALL WE MOVE ON TO THE HARD BITS NOW?

Tony Belpaeme, University of Plymouth, UK

Even though there is a wealth of experimental data available on categorisation and lexicon acquisition, the symbol grounding discussion itself largely remains a fact-free effort. Too many theories and models are the result of arm chair pondering, and the models that do operationalise a theory unfortunately remain limited in scope. One such limitation is the artificial nature of the models' perception and its resulting symbols (Cangelosi et al., 2000; Vogt, 2002). These models are aimed at explicating the processes underlying symbol grounding and as such it would appear justified to take artificial perceptual input. Their value indeed lies in their illustrative nature, but since they use artificial environments they lack empirical power and are often only appreciated by scholars familiar with both computational modelling and the symbol grounding discussion. This is to some extent remedied by taking non-artificial input to study symbol grounding (Seabra Lopes & Chauhan, 2007). While such studies illustrate similar phenomena as former models, they do lead to new insights, for example how a cultural component in symbol grounding can explain the universal nature of certain categories (Belpaeme & Bleys, 2005; Steels & Belpaeme, 2005). However, they are again perhaps limited, this time by being constrained to a single and often simple domain.

The next push in operationalising symbol grounding needs to be in ever increasing and more realistic sensorimotor input and looking at how the social world supports symbol grounding (Belpaeme, Cowley, & MacDorman, 2009). The first is to some extent aided by the availability of better robotic hardware, such as the humanoid iCub robot. However, artificial perception is after half a century of concerted effort still not sufficiently developed to bring symbol grounding to bear upon cognitive systems research.

At the same time, as with many scientific studies, the symbol grounding discussion has not been aided by the too narrow focus on embodiment: were we perhaps too eager in ridiculing non-grounded systems? There is ample evidence that the structure of language does reflect semantics: without the need for grounding, statistical approaches, as used in latent semantic analysis (Landauer & Dumais, 2007) or statistical machine translation, perform feats which hard-line symbol grounding champions would deem impossible. The implicit presence of semantics in the structure of language deserves more attention and I would very much like to see where and how symbol grounding and linguistic structure meet to create semantics beyond perceptually grounded language.

REFERENCES

Belpaeme, T., & Bleys, J. (2005). Explaining universal colour categories through a constrained acquisition process. *Adaptive Behavior*, *13*(4), 293–310. doi:10.1177/105971230501300404

Belpaeme, T., Cowley, S. J., & MacDorman, K. (Eds.). (2009). *Symbol grounding*. Amsterdam, The Netherlands: John Benjamins.

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. doi:10.1080/09540090050129763

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: the Latent Semantic Analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, *104*(2), 211–240. doi:10.1037/0033-295X.104.2.211

Seabra Lopes, L., & Chauhan, A. (2007). How many words can my robot learn? An approach and experiments with one-class learning . *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, *8*(1), 53–81. doi:10.1075/is.8.1.05lop

Steels, L., & Belpaeme, T. (2005). Coordinating Perceptually Grounded Categories through Language. A Case Study for Colour. *The Behavioral and Brain Sciences*, *24*(8), 469–529.

Vogt, P. (2002). The physical symbol grounding problem. *Cognitive Systems Research*, *3*(3), 429–457. doi:10.1016/S1389-0417(02)00051-7

COMMENTARY ON SOLUTIONS AND OPEN CHALLENGES FOR THE SYMBOL GROUNDING PROBLEM

Carol J. Madden, Inserm, France
Stéphane Lallée, Inserm, France
Peter Ford Dominey, Inserm, France

One recent perspective that has been successful in accounting for event representation in humans is that of physicalist models (Wolff, 2007), in which causes and effects are understood in terms of physical dynamics in the world, such as momentum and impact forces. Furthermore, nonphysical causation (e.g., forcing someone to decide) is understood by analogy to these physical force primitives. This approach can be borrowed for any artificial system having a body and the ability to perceive kinematics and dynamic forces. Of course in the robot, these embodied perceptual and physical primitives will surely be transduced into symbols (steels' c-representations). While this is not necessarily problematic, it should be noted that in human cognition, we can easily follow the link back to grounded perception and action; simulating or imagining what it is like to see a red fire truck or pick up a hot potato, with rich perceptual and motor information (and corresponding activation in perceptual-motor areas). In robot models, once an input is transduced to a symbol, there often remains no real link to the continuous/embodied perceptual-motor inputs. Introducing this type of simulation capacity could provide groundbreaking benefits to meaning representation and event understanding (Madden et al., in press).

To avoid being trapped at the indexical level, robots need to not only perceive objects and force primitives (e.g., motion, contact), but also integrate these representations to comprehend events. This requires that representational symbols participate in a kind of grammatical or rule-based processing, just as their labels (words) participate in a linguistic grammar, and their referents in the world participate in role-dependent interactions. However, rule-based symbol processing necessitates some sort of categorical variation in the symbols themselves. Therefore, the robot will need different kinds of symbols that correspond to the different types of meaningful things that can occur in the world, such as entities, states, actions, and relations. These different categories of symbols can then be processed in a representational system according to relational rules. Ideally, the differences between symbol types could even be grounded in the different perceptual processes from which they arise.

In our research, the robot can observe, learn, and perform a series of actions during goal-directed interaction with humans. While the objects and actions are recognized through perceptual primitives, it is clear that this yields no real understanding of the events. Through this new approach, instead of representing sequences of actions to build goal directed events, the robot will represent both actions (A) and states (S), so that events are built from rule-based SAS chains of enabling-states, actions, and resulting-states (Lallée et al., submitted). Recoding the perceptual input into these two different kinds of symbols (states, actions) is a first step that opens the door to representing causality and reasoning about goals. For example, a resulting state can be achieved if the enabling state is true and the action is performed, and this resulting state may serve as the enabling state for another stored SAS event.

Furthermore, this type of representation could be a precursor to understanding intentionality. If a robot is allowed to interact with its environment at random and observe the changes, the robot might observe that a given motor action always produces a given change in state, such as an object being "covered". Eventually, the robot will learn the causal connection between its own action and that resulting state. It can remember this when a task later requires an object to be covered, and initiate the remembered motor action itself to cover the object. By initiating the causal sequence itself, the robot demonstrates a primitive representation of intentionality.

REFERENCES

Lallée, S., Madden, C. J., Hoen, M., & Dominey, P. F. (Manuscript submitted for publication). Linking Language with Embodied and Teleological Representations of Action for Humanoid Cognition. *Frontiers in Neurorobotics*.

Madden, C. J., Hoen, M., & Dominey, P. F. (in press). A cognitive neuroscience perspective on embodied language for human-robot cooperation. *Brain and Language*.

Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology. General*, *136*, 82–111. doi:10.1037/0096-3445.136.1.82

SYMBOLS AND THE PROBLEM OF COACTION

Stephen J. Cowley, University of Hertfordshire, UK

Angelo Cangelosi asks: "How can an agent autonomously create a set of symbol-symbol relationships and the associated transition from an indexical system to a proper symbol system?" Having coded both human-human-robot and infant-caregiver encounters, I answer: Agents can be designed to mimic humans provided that engineers address the *problem of coaction*.

Symbol grounding aims to show how language can be grounded into machines. First, therefore, we face the 'systematic ambiguity' of 'symbol' (Belpaeme & Cowley, 2007). On the one hand, *symbol* can denote an agent's inner representations and, on the other, index how, in real-time, situated, embodied signals (symbols) are created and construed. For Cangelosi, the 'social and cultural grounding of language' (2007, p. 635) depends on how the 'external symbols' that are grounded differ from formal descriptions. Thus, living language is 'a heterogeneous set of artifacts that are implicated in cultural and cognitive activities' (2007, p. 631). Infants use external symbols to ground public language into brains and actions. In *extended symbol grounding* (MacDorman, 2007; Cangelosi, 2007), I argue, people rely on concerting their actions to develop *coactional* capacities and skills (Cowley & MacDorman, 2006; Cowley, 2007).

Infants are outsider-agents born into a population that relies on evolved public language. Adults are thus insider-agents who, by definition, have solved the first and third parts of Cangelosi's puzzle: as functioning bodies, they use public symbols to embody inner control systems (symbols) that function, in part, to track aspects of the world. Physical signals serve in linking situations, objects, public 'words' and how bodies self-produce signals. To manage this is to be a *symbol sharer*. How do infants achieve this status? Regardless of whether they really use inner symbols, they depend on a phenotype, situated encounters and developmental history. This can be modelled by artificial (symbol using) agents who solve the problem of coaction: *they find ways of linking inner representations with external symbolic patterns by drawing on how, in real-time, actions are co-regulated.*

Human infants do not merely learn from interacting. Although much depends on both perception and engaging with physical objects, humans are sensitive to affective rewards. These arise during human-human *coaction*. In Wegner and Sparrow's (2007) terms, during coaction, each party is influenced by, or acts within, the *context* of the other's actions. Moreover, at least one party uses this to come up with action that could not otherwise have arisen. Caregiver-infant dyads create ways of moving, and routines, that parties are motivated to enact (and avoid). They depend on contingencies, rewards and infant control of expressive movement. This also applies to language—provided that we follow Cangelosi (2007) in defining this in terms of cultural and cognitive activities that use artefacts (including 'words'). Since living language is both symbolic and based in physical signals, use of its dynamics can suffice to bring movements under the influence of (unheard) cultural patterns. For a baby, 'words' are *not* needed: in the early months, external symbols are contextual or, in Elman's (2009) terms, *cues to meaning*. A baby needs neither grammar nor perception of phonetic invariants ('phonemes' or 'words'): it uses developing motivations to co-ordinate what is perceived, its own movements, and a history of rewarding contingencies. The rewards of living in a human world make experience of interaction and coaction sufficient for external symbols to be grounded into both the baby's brain and how he or she acts.

Symbols exert effects as caregivers seek to control infants: in learning about rewards, babies rely on anticipatory dynamics. For evolutionary reasons, these shape dyadic expression ('primary intersubjectivity') that links talk, gesturing (and, later, babbling). Here too reward and context give a baby control over movement and vocalization. *Human symbol grounding* (Cowley, 2007) integrates routines with unheard verbal patterns. By 3-4 months babies pick up on cultural norms; by 6-8 months, they share formats (e.g. this little piggy; nappy changing) and by 9-12, they act as if they are using intentions. While learning may be reward-based, adult construals sensitise infants to situations. Far from representing utterance-types ('linguistic forms'), dyads use co-regulation (and sensorimotor control) to reshape perception. That is nature's trick: external symbols become linked with what we *hear* and *see* as we redeploy our bodies and brains (Anderson, 2008). Once we perceive 'words' children answer, and ask, questions ("Do you understand?", "What's that?", "Where's the car?"). A shared *perspective* on language alters perception, cultural skills and first-person phenomenology. No qualitative neural change makes words 'real' (i..e. heard). Children rely on external symbolic patterns and a history of coordinating while acting and talking about talk. By integrating the results with action, humans develop new modes of control: they become symbol sharers.

Agents can use real-time coordination (an 'indexical system') in coming to use situations in making and construing physical symbols (a 'proper symbol system'). To solve the problem of coaction, engineers need to build agents whose inner control systems shape actions that are perceived as external symbols which map onto a situation in ways that, in real-time, co-regulate what humans and/or other machines do. The problem looks soluble. When robots coordinate with us, naïve humans attempt to establish coactional routines: they seek novelty by using the context of robot doings (Cowley, 2008). However, engineers need to change their thinking. First, robots need to mimic (or show) learning from human or human-like coaction attempts (not just actions). Once achieved, they will have to use how humans respond to *their* coactional moves. This is tricky because rewarding contingency detection must link situation appraisal with a developing (pseudo) motivation. To the extent that humans find the outcomes flexible and rewarding, solving the problem of coaction may set machines on the royal road to language. Like babies, artificial symbol sharers will link dynamics and symbols—how we speak and move—in assessing and managing *situations*. In principle, such machines might even perceive 'objects' that afford opportunities for 'acting'. To achieve *that* would show the viability of external symbol grounding.

REFERENCES

Anderson, M. L. (2008). Circuit sharing and the implementation of intelligent systems. *Connection Science*, *20*(4), 239–251. doi:10.1080/09540090802413202

Belpaeme, T., & Cowley, S. J. (2007). Extending symbol grounding. *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, *8*(1), 1–6. doi:10.1075/is.8.1.02bel

Cangelosi, A. (2007). Adaptive agent modeling of distributed language: investigations of the effects of cultural variation and internal action representations. *Language Sciences*, *25*(9), 633–649. doi:10.1016/j.langsci.2006.12.026

Cowley, S. J. (2007). How human infants deal with symbol grounding. *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, *8*(1), 83–104. doi:10.1075/is.8.1.06cow

Cowley, S. J. (2008). Social robotics and the person problem. *Proceedings of the AISB: Computing and Philosophy*, *12*, 28–34.

Cowley, S. J., & MacDorman, K. F. (2006). What baboons, babies, and Tetris players tell us about interaction: A biosocial view of norm-based social learning. *Connection Science*, *18*(4), 363–378. doi:10.1080/09540090600879703

Elman, J. L. (2009). On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive Science*, *33*(4), 547–582. doi:10.1111/j.1551-6709.2009.01023.x

MacDorman, K. F. (2007). Life after the symbol system metaphor. *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, *8*(1), 143–158. doi:10.1075/is.8.1.09mac

Wenger, D., & Sparrow, B. (2007). The Puzzle of Coaction . In Ross, D., Spurrett, D., Kinkaid, H., & Stephens, L. G. (Eds.), *Distributed Cognition and the Will* (pp. 17–41). Cambridge, MA: MIT Press.

THERE SEEMS NO SYMBOLIC REPRESENTATIONS IN THE BRAIN AND WHY

Juyang Weng, Michigan State University, USA

The symbol grounding problem revisited by Angelo Cangelosi is extremely important to clarify, but the way the problem is commonly raised is very misleading. I do not think that all the fields on intelligence (natural and artificial) have understood much of the problem, let alone an effective and well understood solution. I will first address the Angelo's specific questions and then will briefly explain how our "5-chunk model" of the brain-mind deals with the "symbol" phenomenon. I will use a brief form of the questions to make this writing easier to read.

1. My answer to Question 1 (What does the brain link?): Subparts of numeric vectors. The brain and our brain-inspired Epigenetic Developmental (ED) networks (Weng, 2010) seem not to have the luxury to *directly* deal with symbols during their life time. Instead, they link non-symbolic numeric motor vectors (actions that correspond to our reference to symbols, e.g., saying "car") to an attended subpart of the high-dimensional sensorimotor vectors (e.g., a single car event in an image sequence that contains multiple car events and other events).
2. My answer to Question 2 (How does the brain link?): Incrementally and recursively. The brain and our ED networks create links from subparts of sensorimotor vectors (a single car image patch in background and the temporal context of the attended event) to subparts of other sensorimotor vectors (e.g., the next predicted new car image patch and the next motor action) via their internal representations. Such internal representations are not purely symbolic and are emergent from their living experiences.
3. My answer to Question 3 (How do different humans share a symbolic system?): Through communication. The human society shares a symbolic system through calibration of learned behaviors. For example, if a child's behavior indicates that he does not understand ``car'' correctly, the teacher corrects his corresponding behaviors in various educational ways. An ED network may share a symbolic system with humans through interactive learning.

The principled solution I discussed here seems what the brain uses, but I am afraid it is now very difficult for many researchers to confidently accept. To be confident, sufficient width and depth of knowledge in several disciplines are necessary, including biology, neuroscience, psychology, electrical engineering, computer science and mathematics.

It is known that cortical regions are typically inter-connected in both directions (Moran & Desimone, 1985; Callaway, 1998; Felleman & Van Essen, 1991). However, computational models that incorporate both directions have resisted full analysis (Deco & Rolls, 2004; Hinton, 2007; Weng & Prokhorov, 2008; Weng et al., 2008). Informed by the neuro-anatomic studies in neuroscience (Felleman & Van Essen, 1991; Callaway, 1998), our ED computational model in Weng (2010) provides a basis for discussing how the brain-mind deals with the symbol problem.

Before the 5-chunk brain-mind model (Weng, 2010), Weng and his coworkers presented a cortex inspired developmental network model for spatial and temporal abstraction. The network does not have any symbolic representation in its internal (inside its closed "skull") representation, but it has shown its capabilities of dealing with video sequences (Luciw & Weng, 2010), different events in a complex background (Luciw, Weng & Zeng, 2008), and text streams (Weng, Zhang, Chi & Xue, 2009) like those from the Wall Street Journal. We argued: "The brain faces a major challenge. It does not have the luxury of having a human teacher to implant symbols into it, as the brain is not accessible directory to the external human teacher. Thus, it must generate internal representations from the two signal sources: the sensors and the effectors (motors)" (Weng, Zhang, Chi & Xue, 2009).

The sensors for a human brain includes, rods and cones in the retina, hair cells in the cochlea, and various types of somatic receptors under the skin. The effectors of the brain include muscles and glands. The signals from them are all numerical in nature—numeric patterns as instances. None of them receives any abstract symbol in the sense of "one symbol corresponding to one meaning" as used in computers.

Many researchers may ask: Without symbols, how can the brain abstract? Luciw and Weng (2010), Luciw, Weng and Zeng (2008), and Weng, Zhang, Chi and Xue (2009) explained some principled mechanisms for abstraction, inspired by the cerebral cortex. In their models, the weight vector of each feature neuron has three parts—bottom-up, lateral, and top-down—none of which is pure in any linguistic term. Therefore, no internal neuron in their model is pure in any symbolic meaning.

The 5-chunk brain-mind model in Weng, Zhang, Chi and Xue (2009) gives an overall model that addresses how numeric, non-symbolic neurons not only are grounded in the human symbol-based communicative environment, but also reason abstractly. The theory of Finite Automaton (FA) has been the basis of many probabilistic models, such as HMM, POMDP, and Bayesian Nets, and many reasoning engines, such as CYC, ACT-R, and Soar. As explained in Weng (2010), given any FA, an ED network which does not have any symbol in its internal representation can duplicate the functions of the FA. The original symbolic FA is not grounded because its inputs and outputs are purely symbolic so it needs a human as a medium between itself and the real physical world. In contrast, the learned ED is fully grounded as its inputs are directly from the real physical world and its outputs directly apply to the real physical world. Minsky (1991) correctly criticized that prior artificial neural networks do not reason well. This time period is over because of the correspondence between FA and ED (Weng, 2010).

Therefore, there seems no computer-like abstract symbol in the real world either. For example, any word is always sensed as an instance of numeric pattern, an image, a sound, etc. Similarly, no two instances of a word are vocalized exactly the same. Human individuals can communicate with one another about an abstract symbol (e.g., "car") because their learned behaviors are successfully calibrated through social interactions.

REFERENCES

Buschman, T. J., & Miller, E. K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science*, *315*, 1860–1862. doi:10.1126/science.1138071

Callaway, E. M. (1998). Local circuits in primary visual cortex of the macaque monkey. *Annual Review of Neuroscience*, *21*, 47–74. doi:10.1146/annurev.neuro.21.1.47

Deco, G., & Rolls, E. T. (2004). A neurodynamical cortical model of visual attention and invariant object recognition. *Vision Research*, *40*, 2845–2859. doi:10.1016/S0042-6989(00)00140-1

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, *1*, 1–47. doi:10.1093/cercor/1.1.1-a

Fox, M. D., Corbetta, M., Snyder, A. Z., Vincent, J. L., & Raichle, M. E. (2006). Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(26), 10046–10051. doi:10.1073/pnas.0604187103

Hinton, G. E. (2007). Learning multiple layers of representation. *Trends in Cognitive Sciences*, *11*(10), 428–434. doi:10.1016/j.tics.2007.09.004

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews. Neuroscience*, *2*, 194–203. doi:10.1038/35058500

Ji, Z., Weng, J., & Prokhorov, D. (2008). Where-what network 1: "Where" and "What" assist each other through top-down connections. In *Proceedings of the IEEE International Conference on Development and Learning,* Monterey, CA (pp. 61-66).

Luciw, M., & Weng, J. (2010). Where What Network 3: Developmental top-down attention with multiple meaningful foregrounds. In *Proceedings of the IEEE International Joint Conference on Neural Networks,* Barcelona, Spain.

Luciw, M., Weng, J., & Zeng, S. (2008). Motor initiated expectation through top-down connections as abstract context in a physical world. In *Proceedings of the IEEE International Conference on Development and Learning,* Monterey, CA (pp. 1-6).

Minsky, M. (1991). Logical versus analogical or symbolic versus connectionist or neat versus scruffy. *AI Magazine*, *12*(2), 34–51.

Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastrate cortex. *Science*, *229*(4715), 782–784. doi:10.1126/science.4023713

Weng, J. (2010). A 5-chunk developmental brain-mind network model for multiple events in complex backgrounds. In *Proceedings of the IEEE International Joint Conference on Neural Networks,* Barcelona, Spain.

Weng, J., Luwang, T., Lu, H., & Xue, X. (2008). Multilayer in-place learning networks for modeling functional layers in the laminar cortex. *Neural Networks*, *21*, 150–159. doi:10.1016/j.neunet.2007.12.048

Weng, J., Zhang, Q., Chi, M., & Xue, X. (2009). Complex text processing by the temporal context machines. In *Proceedings of the IEEE 8th International Conference on Development and Learning,* Shanghai, China (pp. 1-8).

SOME SHIFTS FOR DISCUSSING SYMBOL GROUNDING

Alberto Greco, University of Genova, Italy

To answer the question whether the symbol grounding (SG) problem has been solved, the best thing to do is to revert to the latest definition that Stevan Harnad, who first pointed out this as a "problem", has given: it "is the problem of causally connecting symbols inside an autonomous system to their referents in the external world without the mediation of an external interpreter" (Harnad, 2010). I think that in discussing this topic we need to make some shifts.

(1) In the new Harnad's formulation, the problem has lost a strong bias present in the 1990 formulation, that was inclined to consider *sensorimotor* events as the prototypical grounding sources. So the first shift we need is one from a molecular perspective to a molar one (to use the old Tolman's terminology), like the one that made possible, even to behaviourists, to go from stimulus to situation and perhaps to environment, and from response to behaviour. Many of the questions asked by Cangelosi can be answered if "looking for grounding" is not considered as looking only for strict sensorimotor connections of symbols, but as abstractions at different layers of events involving the individual. Abstract words like "freedom" or "beauty" *must* be grounded if they are not pure sounds when uttered, and grounding comes from many situations, sensorimotor in origin, that become more and more abstract scripts with experience.

Sloman (2010) is right that Kant had already said everything. Kant was attracted by Hume's idea of concepts built from experience according to association laws, but he was stuck with abstract concepts (number, causality, substance, etc.); since they don't look normal abstractions from experience, so they cannot be simply learnt, his answer was that they are particular concepts called categories. Such concepts are named "a priori synthetic judgments": synthetic because they are constructed from experience but "a priori" because this construction is made by a function. (Kant named it consciousness, we could name it central processor.) that works according to general principles. Such principles go beyond (transcend) experience: so they are built "genetically", they are what cognitivist named "functional architecture" and in a robot are part of firmware. Stimuli reflect the environmental structure, but our knowledge is built according to constraints determined by our cognitive system functioning and resources (so Kant was right that we never know "things in themselves").

(2) A second change of perspective is rethinking what *symbols* are. I think that looking for grounding components for syntactic labels or function words is misleading, simply because such words are *not* symbols.

Let me elaborate more on this point. The question to ask after the latest Harnad's formulation of the SG problem is: are there systems that are able to manipulate symbols in a non-formal way, without having to rely on the builder/user interpretation? the answer seems yes, but now Harnad clarifies that the question was originally intended in terms of full human-scale systems, and we have to admit that at such level the problem is far to be solved. So the answer to the first subproblem posed by Cangelosi is both yes (at a simple scale) and no (at the full Turing scale).

But if the symbol grounding problem can be considered solved only at a full scale, then the second subproblem posed by Cangelosi is not a separate problem but is very connected with the first and may become the main problem. This is because, as Harnad (1990) himself had pointed out, labels (words) assume the function of symbols only when they take place in propositions, otherwise they are "inert taxonomies". Now, a proposition can only be built when there are at least two words and at least one predicate is present or implied. The need for a syntax comes already with only two words. The simplest and more natural syntactic "rule" would be simple juxtaposition, but toddlers learn very early to use more sophisticated ways of connecting different symbols. As we know, after having uttered their first word expressing a whole sentence, toddlers combine two words using a peculiar syntax (Braine, 1976), where one word assumes the role of "pivot" and the other changes (as in "mama comes", "mama good" etc.). From the very start, word ordering tends to be determined by the intended focus (the concept that has focus comes first) (Halliday, 1967). More complex relationships can only be expressed by using special "non-content" words like articles, prepositions, etc., even if at the beginning they are perceived and used as holistic forms (e.g. expressions like "the ball" or "in other words" may be intended as "theball" or "inotherwords" by toddlers) (Rapaport, 2007).

So the problem that Cangelosi describes as one of symbol-symbol relations can mean two different things: how *relationships* are grounded or how *words for relationships* are. One can ask whether words like "with" or "how" are symbols themselves or they just work as a "gluing" material for symbols properly considered. I think that syntactic elements are procedural patterns, acquired via social interaction, about how to glue together properly considered symbols. Relationships, however, are abstracted and grounded from sensorimotor (and beyond) events.

(3) Harnad (2010) seems to consider meaning a broader concept than grounding. But is it possible a non-grounded meaning? what has to be added to grounding for obtaining meaning? In Harnad's view, it seems that this "plus" is something like "feeling". But I think that the third bias that must be overcome is considering grounding sources only as external, i.e. strictly originated by the canonical five senses. Proprioception is often forgotten, although many concepts stem from our bodily states (e.g. happiness from well-being) and are grounded just like visual perceptions.

The SG problem, from the very start, originated in a context (the Chinese Room, the Merry-Go-Round situation, etc.) where the focus was on *understanding*. I think that the only empirical way for testing grounding is just understanding. Meaning and grounding, then, are double-linked. Grounding can be tested as invariant meaning across changes, be they syntactic or pragmatic (context, expectations, etc.) changes. When interpretation changes, grounding is different. Changes that do not affect understanding need not grounding: changing word ordering, making syntax errors like toddlers or people beginning to learn a foreign language, are factors that often do not affect understanding. In these cases looking for grounding of single words is worthless.

REFERENCES

Braine, M. D. S. (1976). Children's First word combinations. *Monographs of the Society for Research in Child Development*, 41.

Halliday, M. (1967). Notes on Transitivity and Theme in English (Part 2). *Journal of Linguistics*, *3*, 206.

Harnad, S. (1990). The Symbol Grounding Problem. *Physica D. Nonlinear Phenomena*, *42*, 335–346. doi:10.1016/0167-2789(90)90087-6

Harnad, S. (2010). Symbol Grounding Problem: Turing-Scale Solution Needed. *AMD Newsletter, 7*, 1, 4.

Sloman, A. (2010). Symbol Grounding is Not a Serious Problem. Theory Tethering Is. *AMD Newsletter*, *7*(1), 7–8.

THE POSSIBLE ROLE OF INTRINSIC PROPERTIES OF THE LANGUAGE COGNITIVE SYSTEM IN EXPLAINING OUR LANGUAGE CAPACITIES

Barbara Giolito, Università Vita-Salute San Raffaele, Italy

One of the most interesting questions in the study of language concerns the symbol grounding problem, in particular the relationship between indices and symbols. The capability of using indices seems to be explainable, from an evolutionary point of view, in terms of the relations between the agents and their environment. The capability to refer to external objects by means of indexical labels is useful: agents start to produce sounds in the presence of some relevant objects or events and the evolutionary process selects such ability because of its usefulness. Some kinds of animal posses such a capability, and also human being can use labels. But human being are also able to use symbols, that is, they can create symbol-symbol relationships. The attempt to explicate the transition from an indexical system to a symbol system by means of cognitive and Artificial Intelligence models is still not fully satisfying: for instance, robots have not been able to develop a language really comparable to human language. But robots don't have a control system as much complex as the human brain is: as Chomsky pointed out, the possibility to develop the human language could depend on the presence of predetermined structures in our minds/brains. Such structures have been in all probability naturally developed, because of their utility, by evolutionary processes: Artificial Life models can be useful in showing such aspects of the language development, in particular about its possible bases. Nevertheless, the actual complexity of language could exceed the possibilities of a fully evolutionary explanation in terms of relationships between words and objects/events: at least for the most complex aspects of language, such as its symbolic level, the mere relationships between words and objects/events could not represent a completely satisfying source of explanation. To explain the more complex aspects of language it could be necessary to analyse internal - mind/brain - structures. Language could be interpreted as an artefact created by the relationship between the human beings and their environment, but once created language - as a complex cognitive system - could exceed the properties directly deriving from such relationship.

A metaphor of such kinds of supervening structures is the development of economy: the use of money has created a system - the economic world - that is characterized by properties not immediately reducible to the properties of mere exchanges money-goods. But there is a difference between the economy system and language system: whereas the economy can be explained as a social artefact, it could be that language is an artefact naturalistically explainable. The language system could exceed the mere interactions between the human beings and their environment because of the creation of a language system much more complicated in comparison with such relations: but such a system could be explainable from a naturalistic point of view, in case we suppose it is implemented in brain structures.

In conclusion, Artificial Intelligence models - in particular, Robotics and Artificial Life - could be useful in showing the evolutionary bases of language development, but their actual incapacity in explaining some - symbolic - aspects of language could depend on the fact that, for such an explanation, we need a different point of view: a point of view concentrating on the properties of the language cognitive system. For example, properties like the rules analysed in the theory of Principles and Parameters, developed by Chomsky, could be essential for the human ability of creating sets of symbol-symbol relationships and for their ability of developing shared sets of symbols: rules like these seem to be essential in individuating the relevant aspects in the innumerable data derived from the relationship between human beings each others and from the relationship between human beings and their environment.

SIMULATING THE GROUNDING OF LANGUAGE WITH NEURO-ROBOTS

Domenico Parisi, National Research Council, Italy

If the symbol grounding problem is the problem of showing that the meaning of words or sentences in the brain is not represented by "symbols" decided by us but is represented as neural patterns of activation that are evolved or learned during the interactions of the individual with the world, I think the problem is solved. It is not clear what a "symbol" might be since all which exists in the brain is neurons, activation level of neurons, synapses between neurons, and strength of these synapses. Words are acoustic stimuli or movements of the phono-articulatory effectors and they are "symbols" in the sense that their auditory processing evokes specific activities in the rest of the brain (language understanding) or these activities cause the movements of the phono-articulatory effectors (language production). The activities in the rest of the brain which are evoked by heard words and which cause the production of words, are the "meaning" of the words. I see two open questions here. First, is there a separated "lexicon" in the brain, or the "meaning" of words is a process which takes place in all of the brain when one hears or produces a word and which can go in different directions according to context and has no specific termination? Is the "meaning" of words derived from the perceptual properties of the things to which the words refer or from the actions with which we respond to these things? My preferences are in both cases for the second alternative but I think these are important research questions that should be addressed with both experiments and computer models. An example of an open question is how we can account for abstract words if the meaning is words is derived from the action with which we respond to the thing to which the words refer. Notice that since I think that the brain should replace the mind in the study of behaviour and that both things and the words that refer to them are represented in the brain as the actions with which we respond to them, the most appropriate simulation models are neuro-robotic models, that is, the construction of either simulated or physically realized robots that interact with a physical environment and are governed by a neural network.

Another claim of symbol grounding is that the meaning of words is autonomously acquired as a consequence of the individual's interaction with the physical and social world. I think this also has been established since organisms are not programmed by anyone but they are what they are as a result of autonomous processes including evolution, development, and learning. However, it is still necessary to actually show how the meanings of words are acquired and I think that this should be done not only by studying natural organisms but also by constructing artificial organisms (robots) that initially are not able to respond appropriately to words and sentences or to produce words and sentences in the appropriate circumstances and then they acquire these abilities. An important research issue here is the respective role of evolution, development (changes during life mainly based on the inherited genotype), and learning (changes mainly due to the particular experiences of the individual) in this acquisition and, especially, how these three processes of change interact. Most computational and robotic models have not addressed these interactions yet.

I think that organisms link words and referents because words and their referents co-vary in their experience and this co-variation is captured by the brain. This can be shown with simulations in which the neural network of the artificial organism is made up of two sub-networks, one mapping perceived entities in the world into movements of the non-linguistic effectors and the other one mapping heard words into movements of the phono-articulatory effectors, with the two sub-networks linked by bi-directional connections. Because of these bi-directional connections, if a heard word co-varies with a non-linguistic action in the organism's experience, the organism will respond to the heard word by executing the non-linguistic action (language understanding). If the perception of a thing co-varies with the production of a given word by the organism's phono-articulatory effectors, the organism will respond to the perceived thing with these phono-articulatory movements (language production). Notice that models such as these can also account for thinking as talking to oneself, which is going from the non-linguistic sub-network to the linguistic sub-network, and vice versa (Mirolli & Parisi, 2006, 2009) for some aspects of language learning such as the transition from babbling (which is mapping heard words into produced words inside the linguistic sub-network) to language understanding/production (which requires the acquisition of the appropriate synaptic weights for the bi-directional connections between the two sub-networks), and perhaps with other phenomena such as the faster acquisition of language comprehension vs language production by children and the greater difficulty for ageing individuals to find the words that refer to things than to understand these same words (if we assume that the linguistic network is significantly larger than the linguistic sub-network).

As indicated in the target article, another open question is the understanding/production of sentences as sequences of words possessing a meaning which can be derived from the meanings of the constituent words. (The target article links this "syntactic" aspect of human language with the issue of what distinguishes words from icons, but the notion of an icon is insufficiently clear.) Some computer models try to simulate nouns as words that co-vary with the objects perceived by the artificial organism and verbs as words that co-vary with which the actions with which the artificial organism responds to these objects, and these simulations show that the artificial organisms are able to respond appropriately to new combinations of nouns and verbs, therefore reproducing what is called the "generativity" of language (Parisi, Cangelosi, & Falcetta, 2002; Cangelosi & Parisi, 2004). However, returning to our notion of meaning as processes that take place in the entire brain and that can go in different directions depending on the context, I think that one should consider that pure generativity is quite rare in the actual use of language and that an element of idiomaticity is present in almost all sentences (Wray, 2005). This can be simulated if one assumes that sentences are not understood or produced by putting together the meanings of the constituent words on the basis of general syntactic rules but by a spreading of continuous neural activation in the entire brain.

REFERENCES

Cangelosi, A., & Parisi, D. (2004). The processing of verbs and nouns in neural networks; insights from synthetic brain imaging. *Brain and Language*, *89*, 401–408. doi:10.1016/S0093-934X(03)00353-5

Mirolli, M., & Parisi, D. (2006). Talking to oneself as a selective pressure for the emergence of language . In Cangelosi, A., Smith, A. D. M., & Smith, K. (Eds.), *The Evolution of Language* (pp. 214–221). Singapore: World Scientific.

Mirolli, M., & Parisi, D. (2009). Towards a Vygotskyan cognitive robotics: The role of language as a cognitive tool. *New Ideas in Psychology.*

Parisi, D., Cangelosi, A., & Falcetta, I. (2002). Verbs, nouns, and simulated language games. *Italian Journal of Linguistics*, *14*, 99–114.

Wray, A. (2005). *Formulaic Language and the Lexicon*. Cambridge, UK: Cmabridge University Press.

THE HARD AND THE EASY GROUNDING PROBLEMS

Vincent C. Müller, Anatolia College/ACT, Germany

'Grounding' is somehow desirable for robotics, but it is not terribly clear why, how, and for what. Historically, Searle had attacked the claims that a computer can literally be said to *understand*, or to *explain* the human ability to understand. His argument was that a computer only does syntactical 'symbol' manipulation and thus a system can not "think, understand and so on *solely* in virtue of being a computer with the right sort of program" (Searle, 1980, p. 368). Harnad turned this in to the challenge that the symbols in a machine with intentional properties (like understanding or meaning) should be 'grounded' via causal connection to the world (Harnad, 1990). Both of these authors were working against the then default thesis that human cognition is computation. If this is assumed, it follows the same computational functions can be carried out on different hardware, e.g. silicon chips instead of neurons.

On this background, I see four symbol grounding problems:

1) How can a purely computational mind acquire meaningful symbols?
2) How can we get a computational robot to show the right linguistic behavior?

These two are misleading: The first one cannot be solved (Müller, 2009) and the second one might not involve grounding since it just requires the right output. We need a problem that does not assume computationalism, but goes beyond 'output' towards grounding of public natural language. In analogy to the distinction between a 'hard' and an 'easy' problem about consciousness (Chalmers, 1995), I would suggest:

3) How can we explain and re-produce the *behavioral ability* and *function* of meaning in artificial computational agents?
4) How does physics give rise to meaning?

We do not even know how to start on the 'hard problem' in 4): from physics to intentional states and phenomenal consciousness; so we should tackle the 'easy problem' in 3): behavior and function.

Cangelosi's three 'sub-problems' are part of this, only that I doubt whether our problem is that of a lonely and fully intelligent agent trying to acquire basic symbols, then to transfer the grounding, and finally to negotiate in communication with other agents. I tend to think that an account of *function* of natural language and language acquisition in humans will have to account for the function of language in intelligence and for the social function of language.

REFERENCES

Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, *2*(3), 200–219.

Harnad, S. (1990). The symbol grounding problem. *Physica D. Nonlinear Phenomena*, *42*, 335–346. doi:10.1016/0167-2789(90)90087-6

Müller, V. C. (2009). Symbol grounding in computational systems: A paradox of intentions. *Minds and Machines*, *19*(4), 529–541. doi:10.1007/s11023-009-9175-1

Searle, J. R. (1980). Minds, brains and programs. *The Behavioral and Brain Sciences*, *3*, 417–457. doi:10.1017/S0140525X00005756

IN THE ROOTS OF THE SYMBOL GROUNDING PROBLEM

Angelo Loula, State University of Feira de Santana, Brazil
João Queiroz, Federal University of Juiz de Fora, Brazil

If we pretend to solve the symbol grounding problem, two fundamental facets have to be placed to frame the discussion. The first precedes the symbol grounding discussion and lays in the roots of a yet open issue in artificial intelligence: how can something become *intrinsic* to an (*autonomous*) artificial agent? The second aspect is important to better define the specific problem of *symbol* grounding: what is a *symbol* and how we can distinguish it from other signs?

The motivation for solving the symbol grounding problem by allowing the agent to *autonomously* build *intrinsic* links between symbols to referents, lies on the early perception by Searle (1980) and Harnad (1990), that symbols cannot be imputed by an external programmer. As Harnad (1990) poses, "the semantic interpretation of a formal symbol system [must] be made intrinsic to the system, rather than just parasitic on the meanings in our heads". So if we expect to solve the symbol grounding problem, we need to find out how can things became *intrinsic* to an artificial issue, and not rely on something imposed or pre-defined by the system designer, thus allowing the agent to become truly *autonomous*. Brooks (1990) pointed out the physical grounding hypothesis as the first step for artificial systems, and the situated and embedded cognition approach has been brought forth as a framework where artificial autonomous agents should be physically connected with its environment, and cognitive competences should be a product of the agents history of interactions. Nevertheless, Ziemke (1999), Ziemke and Sharkey (2001), and Froese and Ziemke (2009) pointed out that this process of making something intrinsic to an artificial agent by grounding a higher-level process in lower-level processes can became rather recursive because the lower-level processes should also be intrinsic to the agent, and therefore should be grounded in the next lower-level processes. One might answer that if the agent is autonomous, its goals are the end of this grounding process, but, as Haselager (2007) points out, that might not be an easy task, particularly if autonomy means the establishment of its own goals. Haselager (2007) proposes that it is fundamentally the body and the 'ongoing attempt to keep its stability' that can allow an agent to be autonomous. But the question of how are this body and its internal functioning intrinsic to the agent itself remains. This can turn into a never ending search for a way of making the ultimate level intrinsic and we could end up saying that nothing can be defined a priori and we can not plan any aspect of the agent. Of course, to say something is intrinsic only if all levels are intrinsic might not help keeping the debate and we think a better question is to say: up to what level is a process intrinsic to the agent? And in discussing symbol grounding, up to what (cognitive/body/environment) level are symbols grounded for the agent?

If we wish to ground symbols, we also need to precisely define what it takes for something to be a symbol and how it is different from other signs. A symbol is a type of sign, and, according to C.S. Peirce, a sign is something that refers to something else (its object), in some respect, for an interpreter, in whom its produces an effect (its interpretant). So for something to be a sign, it should not only be connected to its object, but it should also produce an effect in the interpreter, because if there is no outcome, it is not a sign and therefore cannot be a symbol. A proper understanding of sign processes demands a typology describing the kinds of signs that can be involved in particular semiotic processes (e.g., symbolic). Semiotic processes show a remarkable morphological variety. In an attempt to advance in the understanding of semiotic processes, Peirce proposed several typologies, with different degrees of refinement and several relationships to one another. A basic typology in his framework differentiates between iconic, indexical, and symbolic processes. In his "most fundamental division of signs" (Peirce, 1958, p. 2.275), he characterized icons, indexes, and symbols as matching, respectively, relations of similarity, contiguity, and law between the signs and its object. Icons are signs that stand for their objects by a similarity or resemblance, no matter if they show any spatio-temporal physical correlation with an existent object. In this case, a sign refers to an object in virtue of a certain quality which is shared between them. Indexes are signs which refer to their objects due to a direct physical connection between them. Since (in this case) the sign should be determined by the object (e.g. by means of a causal relationship) both must exist as actual events. Symbols are signs that are related to their object through a determinative relation of law, rule or convention. A symbol becomes a sign of some object merely or mainly by the fact that it is used and understood as such by the interpreter, who establishes this connection. Therefore, for a sign to be a symbol, according to Peirce, there is no need for it to be connected to other symbols; the only requirement is that the interpreter is the one responsible for connecting the symbol to its object (Queiroz & Ribeiro, 2002; Ribeiro et al., 2007). Other properties, such as holding relations with other symbols, give rise to a typology refinement where classes of symbols can be characterized. But, in an indexical relation, on the other hand, the interpreter has nothing to do with the sign-object connection; it only remarks an already established spatial-temporal, physical relation, since the index 'direct the attention to their objects by blind compulsion' (Peirce, 1958, p. 2.305). Thus if an agent after hearing a sign cannot establish on its own the referent for this sign but seeks for the referent in the environment, it interprets the sign as an index, but it is able to determine the referent even in its absence, the sign is a symbol. Indexes, as defined by Peirce, has been neglected by research on symbol grounding, but, inline with Peirce's conception that symbols involve indexes (and indexes involve icons), we have shown how indexes and symbols can emerge and how symbols can be emerge from indexical experiences (Loula et al., 2010, in press). Given the definitions of different sign classes and an indication of its interrelatedness and precedence, we can reframe a question: What it takes to go from indexical processes to symbolic processes? And from iconic processes to indexical processes? And how can we further emerge the different subclasses of icons, indexes and symbols?

REFERENCES

Brooks, R. (1990). Elephants don't play chess. *Robotics and Autonomous Systems*, *6*, 3–15. doi:10.1016/S0921-8890(05)80025-9

Froese, T., & Ziemke, T. (2009). Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence*, *173*, 466–500. doi:10.1016/j.artint.2008.12.001

Harnad, S. (1990). The symbol grounding problem. *Physica D. Nonlinear Phenomena*, *42*, 335–346. doi:10.1016/0167-2789(90)90087-6

Haselager, W. F. G. (2007). Robotics, philosophy and the problems of autonomy . In Dror, I. (Ed.), *Cognitive Technologies and the Pragmatics of Cognition* (pp. 61–77). London: John Benjamins.

Loula, A., Gudwin, R., El-Hani, C., & Queiroz, J. (2010). Emergence of self-organized symbol-based communication in artificial creatures. *Cognitive Systems Research*, *11*(2), 131–147. doi:10.1016/j.cogsys.2008.10.002

Loula, A., Gudwin, R., & Queiroz, J. (in press). On the emergence of indexical and symbolic interpretation in artificial creatures, or What is this I hear? In *Proceedings of XII Alife Conference2010*.

Peirce, C. S. (1958). *The Collected Papers of Charles Sanders Peirce*. Cambridge, MA: Harvard University Press.

Queiroz, J., & Ribeiro, S. (2002). The biological substrate of icons, indexes and symbols in animal communication: a neurosemiotic analysis of Vervet monkey alarm-calls . In Shapiro, M. (Ed.), *The Peirce Seminar Papers – The State of the Art* (*Vol. 5*, pp. 69–78). New York: Berghahn Books.

Ribeiro, S., Loula, A., Araújo, I., Gudwin, R., & Queiroz, J. (2007). Symbols are not uniquely human. *Bio Systems*, *90*, 263–272. doi:10.1016/j.biosystems.2006.09.030

Searle, J. (1980). Minds, brains, and programs. *The Behavioral and Brain Sciences*, *3*(3), 417–457. doi:10.1017/S0140525X00005756

Ziemke, T. (1999). Rethinking grounding . In Riegler, A., Peschl, M., & von Stein, A. (Eds.), *Understanding Representation in the Cognitive Sciences* (pp. 177–199). New York: Plenum Press.

Ziemke, T., & Sharkey, N. (2001). A stroll through the worlds of robots and animals: Applying jakob von uexküll's theory of meaning to adaptive robots and artificial life. *Semiotica*, *134*(1-4), 701–746. doi:10.1515/semi.2001.050

THE ICON GROUNDING PROBLEM

Ricardo Gudwin, UNICAMP, Brazil

In his target article, Cangelosi discusses what are, in his view, the three big challenges in the development of a grounded symbol system: the link of symbols to referents, the symbol-symbol relationship and the development of a shared set of symbols among a society of agents. In this comment, I would like to contribute to such discussion, inserting a new dimension, which in my opinion, is being neglected by the community involved in the symbol grounding problem, which I will be calling here the "icon grounding problem", to make an association to the "symbol grounding problem". What is this problem and how it refers to the "symbol grounding problem itself" ? From my point of view, even though it is possible to ground symbols for things and places in terms of sensors/actuators, this approach is doomed to fail when trying to ground whole phrases (and not just single words) directly on sensor/actuator data. Why ? Because this is not the way we (humans) do while perceiving our world. We do not simply have a memory of a manifold of data, a manifold of properties or attributes. We see "objects" in the world. We build "scenes", or "episodes" with these objects. And our phrases do not just correlate with raw data. Usually, our phrases relate to these scenes and these objects. They "describe" these scenes with substantives which refer to these objects and verbs which refer to actions among objects. But, even though our common-sense gives us such illusion, objects and actions are not really there in the world. Objects and actions are "abstractions" created by our mind to explain the data which flows in our sensors and actuators (e.g. pick the example of seeing shapes and faces in clouds). In other words, these objects and actions are "icons" of sensor/actuator data, created by means of a process of "abstraction". So, what is the "icon grounding problem" and how it relates to the "symbol grounding problem" ? My point is that if we directly try to relate:

world <-> sensor/actuator data <-> symbols

we will not be able to succeed. We need to include a newer dimension in this analysis:

world <-> sensor/actuator data <-> ICONS <-> symbols

In my point of view, symbols should not be grounded directly on sensor/actuator data, but symbols must be grounded in ICONS. And ICONS should be grounded in sensor/actuator data. The biggest problem now, is not to find referents for symbols. The referents for symbols should be icons. The problem now is how do we model these icons and ground them in sensor/actuator data. And the crucial aspect in this problem is how we model "abstraction" in such a task. This is the "icon grounding problem". Once we are able to model "abstraction", we will be able to create icons representing not just objects and actions, but whole scenes. And only when we are able to model whole scenes, we will be able to fully ground complete phrases in a natural language. Only when we have a model for "abstraction", i.e., a procedure for building and modifying icons, we will be able to fully solve the "symbol grounding problem". Someone may object that sensor/actuator data are also icons of reality, and they are right. They are a specific kind of icon we call an "image". But, according to Peirce, there are more evolved kinds of icons, e.g. "diagrams" and "metaphors", which need to be explored. The icon grounding problem includes how to transform among different kinds of icons, from the most basic to the more evolved ones. In order to achieve the grounding of whole phrases, we propose that we need to make a change in the ontological world of icons. We need to make a change from the sensory-motor world to the world of objects. Once we are able to conclude this abstraction, this re-interpretation of icons, maybe we can make it easier the association of whole phrases to scenes in this world of objects, and then properly solve the three sub-problems pointed out by the target work of Cangelosi.

REPLY TO COMMENTARIES

Angelo Cangelosi, University of Plymouth, UK

The commentaries to my target article propose a very interesting set of comments and insights on our current progress and solutions to the symbol grounding problem. In this section I will highlight some of the main issues identified by the esteemed colleagues, and which constitute some of the key open issues and challenges to the symbol grounding problem.

Towards Turing-Scale Robotic Experiments

I would first like to start with the commentary of my symbol grounding "maestro", Stevan Harnad. I welcome his invitation to go beyond using only toy-level solutions to the symbol grounding problem and to aim at Turing-scale experimentation. I believe that one of the most fruitful potential contributions of current cognitive robotics advances is to be able to design experiments where a community of robotic agents is able to socially acquire grounded symbolic systems. This is in line with Steels's invitation to address the symbol grounding issues with practical (robotic) experimentation, and with Belpaeme's encouragement to carry out new experiments with more realistic sensorimotor inputs, such as those with the humanoid robot iCub.

Symbol-Symbol Relationships

As to the commentaries on the current progress and solutions to the three sub-problems identified in my original target article, the issue that appears to be pending remains that of the sub-problem (2) on the transition from indexical representation to syntactic symbol-symbol representations. Steels is right to say that impressive progress has happened on the development of (ungrounded!) symbol-symbol manipulation methods. However, the core problem remains that all these solutions are based on symbolic approaches that use symbols and syntactic relations that are hand-crafted by the researcher, and are therefore not grounded in the agent's own sensorimotor system. At the same time, Belpaeme raises the important point that we should investigate how symbol grounding and symbolic linguistic structures meet to create semantics beyond perceptually grounded language.

Some useful solutions have been suggested to address the sub-problem (2), as well as the grounding problem in general. Madden, Lallée and Dominey propose that the study of embodied simulation capacity in robots could provide groundbreaking benefits to meaning representation and event understanding, thus freeing the robot from "being trapped at the indexical level" through the integration of these representations into more articulated, syntactic-like representation systems. Cowley also suggests a solution through the focus on the "problem of coactions" in child development. In babies, dynamics and symbols –how we speak and move – gradually become resources that could be used in assessing and managing situations (which is similar to Madden et al.'s mental simulation capability). Weng points out that incremental and recursive handling of internal (neural) representations can explain how the brain links symbols between themselves.

Other authors have commented on other issues related to the importance of syntactic relationships and rules in symbol systems.

Greco underlines the fact that syntactic elements, such as function words, are procedural patterns about how to glue properly the symbols. Giolito seems to imply the necessity of some kind of pre-linguistic knowledge of syntactic structure that helps the subsequent learning and acquisition of symbol-symbol relationships. However, in both cases, the issues remains of how these cognitive capabilities have been developed both from and evolutionary and a developmental perspective.

Symbols and the Brain

Weng and Parisi both raise the important issue that there seems to be no symbols in the brain. Weng explains that what the brain does is to link non-symbolic neural representations of perceptual entities with an attended part of high-dimensional sensorimotor representations. Parisi also stresses the fact that internal representations of symbols are actually neural patters on activation that have evolved, or have been learned during interaction with the world. Both authors support the development of new neuro-computational models of symbol learning. For example, Weng proposes the 5-chunk brain-mind model that has the potential to explain the developmental emergence, through experience, of symbolic capabilities. I agree with both colleagues that there are no symbols in the brain, but rather internal representation of concepts based on dynamic neural activation patterns. However, the real difference between the neural representation of symbols, such as words and numbers, and the representations of other sensorimotor concepts is the fact that symbolic representations have the property of being discrete, thus allowing recursive combinations (compositions) of symbols to express other higher-order symbolic entities.

Grounding of Abstract Symbols

One of the main challenges in grounding theory remains the explanation of the sensorimotor grounding of abstract concepts and words such as "freedom" and "beauty". Parisi mentions this as one of the main challenges in the field. Greco invites us at looking beyond a pure search for the sensorimotor bases of all words. And he also discusses the importance of a priori constraints determined by our cognitive system and its functional architecture, that could be involved in the grounding of such complex concepts. My suggestion here is that the sensorimotor basis of language grounding, including that of abstract words, should be considered in a broader sense which includes various embodiment levels, such as those of basic instincts and drives, emotions and social instincts.

Autonomy and the Hard Symbol Grounding Problem

Some commentaries focus on the importance of autonomy in the symbol grounding. Muller proposes the distinction between a "hard" and an "easy" version of the symbol grounding problem. The hard version focuses on the essential property of the agent's own intentionality to be able to create and use symbols. The "easy" version aims at the explanation and re-production of the behavioral ability and function of meaning in artificial computational agents and robots. Muller suggests that the easy symbol grounding problem can be successfully investigated through practical experimentations on developmental robotics, whilst the hard problem remains a very big challenge. Loula and Queiroz also discuss the concepts of autonomy and intrinsic links in the grounding problem to understand how can a symbol become intrinsic to an autonomous agent. In particular they are interested in understanding up to what level is a process intrinsic to the agents. They indicate that the agent's body, and its ongoing need and attempt to keep its stability, can explain the concepts of autonomy and intrinsicity in cognitive and symbolic processing.

Semiotic Considerations

Loula and Queiroz stress the importance of the need to have a precise definition of a symbol, within Pearce's semiotic framework that differentiates between icons, indices and symbols. In particular they stress the importance of understanding the (cognitive) transitions from indexical processes to symbolic ones, and also the transition from iconic processes to indexical one. Gudwin proposes a new focus of the symbol grounding problem shifting the attention towards the intermediate representations that link the sensorimotor input with the symbolic level. I believe that his suggestion to require the presence of an intermediate representation stage, that he calls "icon", share similar properties with Harnad's original proposal to have categorical representations as the intermediate representations at the core of the symbolic capability.

Conclusion

The computational approaches to the modelling of linguistic and symbolic capabilities, such as developmental robotics experiments on language learning, offer a great methodological and epistemological opportunity to investigate the fine mechanisms at the basis of the symbol grounding problem. Through these models, ranging from cognitive agents models to Turing-like robotics experimentations, we have a great challenge, and opportunity, to demonstrate some of the great mysteries of human development: how babies are able to acquire autonomously, through interaction with the physical world, own body and their social environment, this shared symbolic system that we call language.

Stevan Harnad was born in Budapest, Hungary. He did his undergraduate work at McGill University and his graduate work at Princeton University's Department of Psychology. He is currently Canada Research Chair in Cognitive Science at Université du Québec à Montréal (UQAM) and Professor of Cognitive Science at the University of Southampton. He is also an External Member of the Hungarian Academy of Sciences. Founder and Editor of Behavioral and Brain Sciences (a paper journal published by Cambridge University Press), Psycoloquy (an electronic journal sponsored by the American Psychological Association) and the CogPrints Electronic Preprint Archive in the Cognitive Sciences. His research is on categorisation, communication, cognition and consciousness.

Luc Steels is professor of Computer Science at the Free University of Brussels (VUB), founder and director (from 1983) of the VUB Artificial Intelligence Laboratory and co-founder and chairman (from 1990 until 1995) of the VUB Computer Science Department (Faculty of Sciences). He is also the director of Sony Computer Science Laboratory in Paris. His scientific research interests cover the whole field of artificial intelligence, including natural language, vision, robot behavior, learning, cognitive architecture, and knowledge representation. At the moment his focus is on dialogs for humanoid robots and fundamental research into the origins of language and meaning. Current work focuses on developing the foundations of semiotic dynamics and on fluid construction grammars.

Tony Belpaeme is an Associate Professor (Reader) in Intelligent Systems in the Centre for Robotics and Neural Systems at the University of Plymouth. He works on developmental robotics and human-robot interaction and has a keen interest in how language impacts on cognition and how this can be used to construct intelligent systems. Relevant to this contribution, he has edited a book on symbol grounding with Stephen J. Cowley and Karl F. MacDorman and published on how cultural learning can support symbol grounding.

Carol Madden grew up outside Philadelphia and received an undergraduate degree in Cognitive Psychology from the University of Notre Dame in 1998. She received her graduate training from the Florida State University, obtaining her Ph.D. in 2004 in Cognitive and Behavioral Sciences. She subsequently did postdoctoral research at the Erasmus University Rotterdam, the Netherlands, where she was awarded a fellowship for her own research project until spring 2010. Carol recently began a permanent position with the CNRS at the INSERM Stem Cell and Brain Research Institute in Lyon France. Her research investigates how humans represent described events through embodied and modality-specific language representations. This research is extended to create artificial systems that can use grammatical cues to activate appropriate meanings that are liked to real perceptual-motor experience.

Stephen Cowley is a Senior Lecturer in Developmental Psychology at the University of Hertfordshire (UK). His empirical work has dealt with prosody in conversations, mother-infant interaction, and coaction by children and adults. He has also investigated how robots and mannekins contribute to classroom activity and medical simulation. Aiming to link linguistics, biology and social practice, he and Nigel Love co-founded the Distributed Language Group (DLG). Instead of modelling verbal systems or language-use, the DLG community stress linguistic heterogeneity. Taking a view from everywhere, they trace language to biological and cultural systems that evolved together with human coaction. Pursuing this view, Stephen has published over 50 research papers and edited various collections of papers. These include a Special Issue of Language Sciences (2007) on "Cognitive dynamics and distributed language" and a Special Issue of Pragmatics & Cognition (2009) on "Distributed language."

Juyang (John) Weng received his BS degree from Fudan University, and MS and PhD degrees from University of Illinois, Urbana-Champaign, all in Computer Science. He is now a professor at the Department of Computer Science and Engineering, Michigan State University, East Lansing, Michigan. He is also a faculty member of the Cognitive Science Program and the Neuroscience Program at Michigan State University. Since the work of Cresceptron (ICCV 1993), he expanded his research interests in biologically inspired systems, especially the autonomous development of a variety of mental capabilities by robots and animals, including perception, cognition, behaviors, motivation, abstract reasoning skills. He has published over 200 research articles on related subjects. He is an editor-in-chief of International Journal of Humanoid Robotics and an associate editor of the new IEEE Transactions on Autonomous Mental Development, a member of the Executive Board of International Neural Network Society, and a fellow of IEEE.

Alberto Greco is full professor of General Psychology, chief of Lab of Psychology and Cognitive Sciences and coordinator of the PhD course in Psychology, Anthropology and Cognitive Sciences at the University of Genoa, Italy. He is vice-president of the Italian Association for Cognitive Sciences (AISC). His main research fields concern epistemological, historical, and methodological aspects of psychology and cognitive science, and cognitive models. His main contributions are about concept definitions, integration between models and perspectives in cognitive science, foundation and use of simulation models, experimental and applied study of categorical processes and of diagrammatic reasoning. Recent extensions concern cognitive robotics in ambient assisted living.

Barbara Giolito, four-year degree in Philosophy of Science at the University of Turin - Vercelli, grade 110/110 cum laude ("The use of the Incompleteness Theorem of Gödel in evaluating the Artificial Intelligence methodology in the study of the human mind") - 1998; Diplôme d'études supérieures in Philosophie et Histoire de la Logique at the Geneva University ("A study on the possibility to refuse the excluded middle principle in classical logic") - 1999; PhD degree in Philosophy of Language at the Eastern Piedmont University "A. Avogadro" at Vercelli ("The referential competence in human language: a neural network approach") - 2005; since 2007 Member of CRESA Centre of Research in Experimental and Applied Epistemology, University Vita-Salute San Raffaele at Milan; since 2008 Lecturer of Epistemology of human sciences, Epistemology and elements of logic, Artificial Intelligence and Logic of scientific discovery and innovation at the University Vita-Salute S.Raffaele at Milan.

Domenico Parisi is a researcher at the Institute of Cognitive Sciences and Technologies of the National Research Council in Rome, where he is part of the research group Laboratory of Artificial Life and Robotics. He is the Director of the review 'Sistemi Intelligenti'. His main research interests are the development of simulation models of individual and social behaviors through neural networks, genetic algorithms and models of artificial life; their philosophical and cultural implications (with regards to the impact of technology on the human mind and the cultural and political evolution of societies) and their applications in the fields of education, entertainment and divulgation.

Vincent C. Müller is an Associate Professor of Philosophy, Anatolia College/ACT, Thessaloniki, Greece (Dept. chair 2000-2004). Education: Philosophy with minors in modern history and linguistics at Phillips University Marburg and at Hamburg University. MA in philosophy at King's College London. PhD in philosophy with a thesis on "Realism and Reference" at Hamburg university, with one year as Visiting Research Student at Oxford University. Publications & Activities: Edition and translation of a collection of papers by Hilary Putnam, a number of articles on the philosophy of language, mind and computing, as well as on ethical problems (in Mind and Language, Philosophy & Phenomenological Research, Minds & Machines and other places), co-edition of several bibliographies and lexica, philosophical translations English-German. Visiting Research Fellow at Princeton University in 2005/6. Currently preparing a book on "Artificial Intelligence: Basic Problems". Coordinator of "EUCogII" - 2nd European Network for the Advancement of Artificial Cognitive Systems, Interaction and Robotics (www.eucognition.org).

Angelo Loula is a Professor at the State University of Feira de Santana, Bahia, Brazil. He is a PhD candidate at the Department of Computer Engineering and Industrial Automation, FEEC/UNICAMP. He received his M.Sc. in Electrical Engineering from FEEC, his B.Sc. in Electrical Engineering from the Federal University of Bahia and a Technologist degree in Computing from Ruy Barbosa Faculty. He is the leader of Group for Research on Artificial Cognition. He is the co-editor of 'Artificial Cognition Systems' (IGI, 2006) and 'Computação Cognição e Semiose' (EDUFBA, 2007), and Advances in Modeling Adaptive and Cognitive Systems (UEFS, 2010). He is the Editor-in-Chief of the International Journal of Signs and Semiotic Systems. His research interests include artificial cognition, semiotic and meaning processes, communication, language evolution and acquisition, artificial life, adaptive behavior, complex systems and self-organization, embodied and situated cognition.

João Queiroz is a professor at the Institute of Arts and Design and at the Graduate Studies Program in Communication, Federal University of Juiz de Fora, Minas Gerais, Brazil. He earned a Ph.D. in Communication and Semiotics from the Catholic University of Sao Paulo (PUC-SP), and a post-doc fellowship in Intelligent Systems at the School of Electrical and Computer Engineering (FEEC-DCA), State University of Campinas (Unicamp, Brazil), and in Philosophy of Biology at the Institute of Biology, Federal University of Bahia (UFBA). He is the Editor-in-Chief of the International Journal of Signs and Semiotic Systems. He is author and co-editor of several books and special issues as well as numerous articles on the emergence and evolution of semiotic complexity in biological and artificial systems. He is the director of the Group for Research in Cognitive Science and Semiotics (UFJF), Group for Research on Artificial Cognition (UEFS), an associate researcher of the Research Group on History, Philosophy and Biology Teaching (UFBA), and member of the Institute for Research on Anticipatory Systems (The University of Texas at Dallas). His research interests include evolution of semiotic complexity; emergence of symbolic based communication; artifacts and cognitive technologies; complexity studies; C.S. Peirce's semiotic and pragmatism.

Ricardo Gudwin is an Associate Professor at the Faculty of Electrical and Computer Engineering, State University of Campinas - Brazil. He was born in Campinas-SP, Brazil, in 1967, and received the B.S. degree in Electrical Engineering in 1989, the M.S. degree in Electrical Engineering in 1992 and the Ph.D. in Electrical Engineering in 1996, all of them from the Faculty of Electrical and Computer Engineering, State University of Campinas - Brazil. His earlier research interests include fuzzy systems, neural networks and evolutionary systems. His current research interests include the study of intelligence and intelligent systems, intelligent agents, semiotics, computational semiotics and artificial cognition. He published two edited books with the tittles "Artificial Cognition Systems" (in conjunction with Angelo Loula and João Queiroz - July 2006 - Idea Group Publishing - ISBN 1-59904-111-1), and "Semiotics and Intelligent Systems Development" (in conjunction with João Queiroz - December 2006). Prof. Gudwin is the head of the "Computational Semiotics Group", and Scientific Member/Director of the Group for Research on Artificial Cognition within the DCA/FEEC/UNICAMP, a member of the board of governors of the SEE - Semiotics-Evolution-Energy Virtual Institute in Toronto, Canada, and a member of the editorial board of the "On Line Journal for Semiotics, Evolution, Energy Development" - ISSN 1492-3157, published by the SEE Virtual Institute.

Angelo Cangelosi is Professor of Artificial Intelligence and Cognition at the University of Plymouth (UK), where he leads the Centre for Robotics and Neural Systems. He obtained a PhD in psychology and computational modeling in 1997 from the University of Genoa, whilst also working as visiting scholar at the National Research Council (Rome), the University of California San Diego (USA) and the University of Southampton (UK). He has produced more than 140 scientific publications, and has been awarded numerous research grants from UK and international funding agencies. He is coordinator of the EU FP7 project "ITALK" (italkproject.org) and the "RobotDoC" Marie Curie doctoral network in developmental robotics (robotdoc.org).

International Journal of Signs and Semiotic Systems

An official publication of the Information Resources Management Association

Mission

The mission of the *International Journal of Signs and Semiotic Systems* (IJSSS) is to provide a forum for comparative, evolutionary, theoretical, and empirical perspectives on computational modeling of semiotic systems and processes, and to present formal and synthetic simulation, implementation, and realization of computer machines and robots involved in these systems and processes. This journal offers a collaborative approach to innovative research geared toward a deeper understanding of semiotic processes and an integrated vision of synthesis in artificial systems. As an integral reference for community researchers in the field, this publication explores recent developments in neurocognitive science, cognitive ethology, second order cybernetics, cognitive semantics, pragmatic linguistics, theoretical semiotics, biosemiotics, and evolutionary biology.

Subscription Information

IJSSS is published semi-annually: January - June; July - December by IGI Global. Full subscription information may be found at www.igi-global.com/ijsss. The journal is available in print and electronic formats.

Institutions may also purchase a site license providing access to the full IGI Global journal collection featuring more than 100 topical journals in information/computer science and technology applied to business & public administration, engineering, education, medical & healthcare, and social science. For information visit www.infosci-journals.com or contact IGI at eresources@igi-global.com.

Copies of Articles

Copies of individual articles are available for purchase at **InfoSci-On-Demand** (www.infosci-on-demand.com), an aggregated database consisting of more than 31,000 research articles written by prominent experts and scholars in the field of information/computer science and technology applied to business & public administration, engineering, education, medical & healthcare, and social science.

Copyright

Correspondence and questions:

Editorial: Angelo Loula & Joao Queiroz
Editors-in-Chief
IJSSS
E-mail: ijsss@artificial.eng.br

Subscriber Info: IGI Global
Customer Service
701 E Chocolate Avenue
HersheyPA17033-1240,USA
Tel: 717/533-8845 x100
E-mail: cust@igi-global.com

The *International Journal of Signs and Semiotic Systems* is currently listed or indexed in: MediaFinder; The Standard Periodical Directory; Ulrich's Periodicals Directory