Extending symbol grounding

Tony Belpaeme and Stephen J. Cowley

Faculty of Technology, University of Plymouth / School of Psychology, University of Hertfordshire

The papers collected in this special issue emerged from an international workshop on symbol grounding organised at the University of Plymouth on 3 and 4 July 2006 by the *Distributed Language Group*. Our goal was to extend the classical view of symbol grounding by recognising that language and cognitive dynamics are mutually constitutive. Specifically, we aimed to do so by bringing researchers who study human signalling together with others who focus on simulating intelligence and language. In the original call for papers, we set out these objectives as follows:

> Specifically, we wish to invite contributions viewing language and cognition as linking what goes on in the head with causal processes that are intersubjective, multimodal, affect-laden, and organised by historically rooted customs and artefacts.… The purpose of the workshop is not so much to present completed work as to find new ways of tackling a complex issue and to launch collaboration among participants to that end.… Since the workshop focuses on how symbol grounding can be reconsidered when language is viewed as a dynamical process rooted in both culture and biology, research related to robotic or computer modelling of symbol grounding, psychological and linguistic viewpoints on cognitive development and semiotic dynamics are of great interest.

Of the papers given at the conference, six have come through our process of selection. Each focuses on ways in which events internal to an agent are grounded in external causal processes. All propose ways in which inner representations (symbols) can interact with the richness of the language-saturated world.

By invoking symbol grounding, Harnad (1990) recognised that any computational model of mind must explain how an agent's representations (or symbols) connect with the external world (cf. Taddeo & Floridi, 2005). While contrasting in their view of computation, each paper in this issue accepts the classical view that (inner) symbols must connect with objects and linguistic signals. At the same time, there is also a striking contrast. While Harnad used ‘symbol’ to denote both the agent’s (inner) representations and linguistic signals, today this systematic ambiguity is widely rejected. Most concur that explaining the evolution and emergence of linguistic signals is quite unlike grounding (inner) symbols into objects beyond the skin. To understand this shift in perspective, one has to revisit the historical context of symbol grounding. Harnad (1990) formulated the symbol grounding problem in response to Searle’s (1980) famous (or notorious) *Minds, Brains, and Programs* article. In this, readers are invited to imagine a Chinese Room where a person applies rules to input that consists of what are seen as senseless squiggles. As is the case for a computer program, this gives rise to output that makes sense *outside* the room. For many this suggests that computers cannot attain human-like intelligence. Like a parable, the text is cleverly constructed. Playing the devil’s advocate, Searle first entertains the view that, because a brain is like a computer executing a program, a computer could pass the Turing Test. Second, he prompts the intuition that syntactic processing is not intentional. Implicitly, any device that passes the test is of merely technical interest. Without intentionality, no computer could attach meaning to its inner symbols. Harnad (1990, 2003) reinterpreted this as a challenge that showed it was important to establish how symbols (linguistic or otherwise) *can* be connected to the world beyond the body. Without a person to interpret them, squiggles on paper or spoken words indeed lack meaning or, for Harnad, are ungrounded. Once inwardly represented, however, they connect with referent-indexed symbols and, in so doing, are *grounded*. The challenge is to show how agents can causally connect symbols to the external world. First, they need senses to identify referents which, eventually, can connect with squiggles or spoken words. To add meaning to the resulting symbols, they need sensorimotor interaction with the world. To solve the problem of symbol grounding, a machine must be a robot rather than a computer.

What kinds of *symbol* do robots need? For Harnad (1990) and computational theorists (Newell & Simon, 1976; Newell, 1980; Pylyshyn, 1984; Fodor 1987), this is a matter of definition. Symbols are arbitrary physical token or syntactic combinations of such tokens that can be manipulated systematically and generatively and assigned a meaning. An agent’s symbols, therefore, stand in for both physical properties (e.g., what is common to cups) and repeatable linguistic signals (e.g., spoken words). By implication, grounding symbols in cups is broadly similar to grounding them in the speech stream.

It is striking that each of the papers presented below rejects this approach. Generally, they propose two ways of extending the classical view of symbol grounding. First, those who take a pragmatic view draw on insights from implementing symbol grounding in simulation or on real-world robots. This leads them, first, to reject appeal to already-combined forms. Instead, symbols are viewed as arising from connections between a referent, its sensorimotor interpretation and, in some cases, a communicative signal. This is a *semiotic* interpretation of symbol (Maturana and Varela, 1992; Clancey, 1997; Clowes, 2007; Vogt, 2002, 2007). For the pragmatists, a symbol is ultimately a set of relations which connect a meaning, signal and referent. As in Peirce's semiotic model, this is constituted by a triad of representamen, interpretant and object. Symbols thus unite a system's sensorimotor capacities with its interpretative ones. To distinguish this from seeing the symbols as meaningless tokens, Vogt (2002) calls this the *physical* symbol grounding problem. Alongside this, a second culturalist group also rejects formal definitions of symbols. They object to the view that brains use similar forms of symbols to index both objects and patterns in the speech stream. Underplaying cultural practices, this view exaggerates the importance of notations. By assuming that verbal patterns are grounded in notational ways, the classical approach forgets that words are ultimately physical signals that contribute to social activity. Unlike cups, they do not draw on sensorimotor invariants. For the brain, indeed, spoken words may entirely lack symbolic representation. In Viger (2007), therefore, a capacity to identify referents is the outcome of a process of acquisition. For Worgan and Damper (2007), the classical approach must be extended by appeal to *signal grounding*. Finally, for Cowley (2007), categorisation of spoken words is marginal during the early stages of learning to talk. The infant-caregiver dyad focuses, not on forms or referents, but coordinating needs around social strategies. To simulate extended symbol grounding, robots too will need (inner) symbols that have richer content than is posited in the classical model.

Language thus ceases to be a system of formal symbols. Rather, it is viewed as an external cognitive resource that allows us to acquire categories and concepts. It functions as agents coordinate behaviour and, in so doing, gain some grasp of how others perceive the world. During the process of grounding symbols in language, therefore, agents play social games that always involve (at least) two brains, computers, robots, or some combination of these systems (cf. Seabra Lopes & Chauhan, 2007). In coordinating, of course, they typically interpret each signal in different ways. Because of individual differences and the dynamic nature of cognition, neither flesh nor silicon will concur on signal meaning. Although this poses no major problems for an individual agent, it ensures that cooperation or communication must take on strategic functions. Since referents are not guaranteed, cultural processes must stabilise interaction. We use language to coordinate what we do and thereby coordinate with what signals represent (Steels and Belpaeme, 2005).

Language thus influences symbol grounding in several ways. It facilitates the acquisition of meaning (Cangelosi and Harnad, 2001) and serves to coordinate referents between individuals (Steels and Belpaeme, 2005). Far from being a monolithic system of determinate forms, language is a dynamic process that is shaped and propagated by users who exploit social and physical constraints. Given that language also influences symbol grounding, the process too must be dynamic. Further, while important in human cognition, our cognitive states (or symbols) depend on much more than verbal patterns. In extending classical views of symbol grounding, therefore, we emphasise emotion and, equally, how categories sensitise to culturally based customs and artefacts. Human interaction and, thus, human symbols, depend on more than a syntactic analysis of words. Beyond this, however, each paper points in a different direction. First, breaking with Fodor, **Viger** argues that the emergence of reference be seen as the culmination of a learning process. Taking the opposite tack, **Cowley** focuses on infancy to stress that the early stages of learning to talk are more dependent on affective and multimodal processes than spoken words. In an attempt to construct a general framework, **Clowes** suggests that language is gradually internalised. Reporting ambitious work, **Vogt and Divina** describe a model that aims at simulating the cultural processes. In building robots that learn words **Seabra Lopes and Chauhan** fall in line in finding that teaching – not categorisation – is a simpler basis for language. Finally, focusing on the evolution of phonetic forms, **Worgan and Damper** stress that many properties of human languages are likely to depend on constraints based in how articulation and audition favour learning. Such processes, of course, use the physics of speech and, thus, underpin symbol grounding.

The six papers extend traditional views of symbol grounding in both applied and theoretical ways. At the engineering end, **Worgan and Damper** criticise the lack of attention paid to the physical aspects of the symbol that is grounded. Pursuing this, they present a computational model in which a shared speech code emerges between agents, and show how constraints on production, perception and acoustics influence speech signals. While Worgan and Damper are concerned with physical signals, **Vogt and Divina** focus on simulations of their semantic counterparts. In their paper, they present a large scale simulation in which agents socially transmit and acquire behavioural skills. As the agents communicate, their signals can have ecological relevance only if they are grounded. Vogt and Divina present the technical challenges involved and put forward a number of learning methods, such as cross-situational learning, inspired by psycholinguistic observations. The contribution of **Seabra Lopes and Chauhan** takes the study of symbol grounding into the real world by presenting an experiment in which an adaptive learning system acquires the meaning of objects through linguistic instruction. Besides demonstrating the characteristics of their learning method, their pragmatic approach clearly shows that implementing symbol grounding in a real-world system is a formidable challenge in which many issues – such as perception, concept representation and concept learning – remain far from solved. **Cowley** argues that too much attention is spent on language as the facilitating medium for symbol grounding, and he observes that infants, from an early age, bootstrap signal grounding by relying on what agents *do together*. The infant uses affect-laden interactions to acquire culturally evolved patterns that enable it to participate in a grounding process. **Clowes** starts out by contrasting the view of symbols as formal tokens with the view that they emerge from a semiotic agreement. In his model, he seeks to reconcile strengths of these approaches. In so doing, he starts from Vygotsky's views on language internalisation and sets out a model of symbol internalisation. In Clowes's model this process is dynamic and self-regulatory. In the most theoretical paper, therefore, **Viger** takes a philosophical perspective on symbol grounding. He focuses on how abstract symbols, for example, *unicorn*, can become intrinsic to thinking. While concurring with Fodor that we need a language of thought, he challenges tradition by suggesting that referents emerge from a long acquisition process. Finally, we leave the last word to **MacDorman** who, looking back over the workshop, proposes that we abandon the symbol system metaphor to ask how humans—and perhaps robot bodies—can construct themselves into persons (MacDorman, 2007).

The symbol grounding problem has shown more life in it than Searle's thought experiment. Indeed, our papers strongly imply that the Chinese Room parable fails because of a simplistic assumption that, as in programs, spoken words can be formally represented. In fact, as all contributors stress, semiosis involves more than form-based computation. While there are many differences between contributors, we are optimistic about future work on both extended symbol grounding and how machines can simulate the development of persons. Specifically, while studies on the evolution, emergence and origins of language have been shown to matter for both theory and engineering, the Distributed Language Group now aims to proceed by reversing this perspective. Given how verbal-patterns can be grounded, we can ask what engineers who design robots to adapt to human users are able to tell us (and want to know) about linguistic signals. This theme will be addressed in a follow-up workshop.

References

Blackmore, S. (1999). The meme machine. Oxford: Oxford University Press.

Cangelosi, A. & Harnad, S. (2001). The adaptive advantage of symbolic theft over sensorimotor toil: Grounding language in perceptual categories. *Evolution of Communication, 4*(1), 117-142.

Clowes, R. (2007). Semiotic symbols and the missing theory of thinking. *Interaction Studies, 8*(1) this issue.

Cowley, S. J. (2007). How human infants deal with symbol grounding. *Interaction Studies, 8*(1), this issue.

Fodor, J. A. (1987). Modules, frames, fridgeons, sleeping dogs and the music of the spheres. In Pylyshyn, Z. W. (Ed.) *The robot's dilemma: The frame problem in artificial intelligence* (Chapter 8). Norwood, NJ: Ablex..

Harnad, S. (1990). The Symbol grounding problem. *Physica D, 42,* 335-346.

Harnad, S. (2003). The Symbol grounding problem. *Encyclopedia of Cognitive Sciences.* London: Macmillan/Nature.

MacDorman, K. F. (2007). Life after the symbol system metaphor. *Interaction Studies, 8*(1), this issue.

Newell, A. & Simon, H.A. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the ACM, 19*(3), 113-126.

Newell, A. (1980). Physical symbol systems. *Cognitive Science, 4*(4), 135-183.

Pylyshyn, Z.W. (1984). *Computation and cognition.* Cambridge, MA: The MIT Press.

Seabra Lopes, L. & Chauhan, A. (2007). How many words can my robot learn? An approach and experiments with one-class learning. *Interaction Studies, 8*(1), this issue.

Searle, J. (1980). Minds, brains and programs. *Behavioral and Brain Sciences, 3*(3), 417-457.

Steels, L. and Belpaeme, T. (2005). Coordinating perceptually grounded categories through language. A case study for colour. *Behavioral and Brain Sciences, 24*(8), 469–529.

Taddeo, M. & Floridi, L. (2005). The symbol grounding problem: A critical review of fifteen years of research. *Journal of Experimental and Theoretical Artificial Intelligence, 17*(4), 419-445.

Viger, C. (2007). The acquired language of thought hypothesis: A theory of symbol grounding. *Interaction Studies, 8*(1), this issue.

Vogt, P. (2002). The physical symbol grounding problem. *Cognitive Systems Research, 3*(3), 429-457.

Vogt, P. & Divina, F. (2007). Social symbol grounding and language evolution. *Interaction Studies, 8*(1), this issue.

Worgan, S. F. & Damper, R. I. (2007). Grounding symbols in the physics of speech communication. *Interaction Studies, (8)*1, this issue.