

This article was downloaded by: [University of Plymouth]

On: 25 November 2008

Access details: Access Details: [subscription number 789285667]

Publisher Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Connection Science

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713411269>

Beyond the individual: new insights on language, cognition and robots

Luís Seabra Lopes^a; Tony Belpaeme^b

^a Departamento de Electrónica, Telecomunicações e Informática, Universidade de Aveiro, Portugal ^b School of Computing, Communications and Electronics, University of Plymouth, United Kingdom

Online Publication Date: 01 December 2008

To cite this Article Lopes, Luís Seabra and Belpaeme, Tony(2008)'Beyond the individual: new insights on language, cognition and robots',Connection Science,20:4,231 — 237

To link to this Article: DOI: 10.1080/09540090802518661

URL: <http://dx.doi.org/10.1080/09540090802518661>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

FOREWORD

Beyond the individual: new insights on language, cognition and robots

Luís Seabra Lopes^{a*} and Tony Belpaeme^b

^a*Departamento de Electrónica, Telecomunicações e Informática, Universidade de Aveiro, Portugal;*

^b*School of Computing, Communications and Electronics, University of Plymouth, United Kingdom*

Keywords: language; robots; cognition; artificial intelligence; distributed language

1. Introduction

Language sets humans apart from other animals. In the context of robotics, language is a fast and surprisingly efficient channel to communicate thoughts; it is used to instruct and teach and is arguably the most important conduit for transferring and acquiring knowledge. Ever since the early days of artificial intelligence, language has drawn a lot of attention as an intuitive and easy interface to machines, and large efforts are focused on designing speech recognition and natural language interpretation and production. The Holy Grail is of course a system which can interpret and produce language like humans do. This ties in with the historical trend of seeking to make machine interfaces increasingly intuitive and thus user friendly. This has recently culminated in the field of social robotics, in which the fundamental issues and implementations are studied in order to make robots capable of social interaction with people and other robots.

While there have been a number of increasingly successful implementations of non-linguistic behaviour in robots, progress in implementing linguistic behaviour has been slower. Natural language processing (NLP) has already received decades of attention in the artificial intelligence community. The NLP community has a number of well established beliefs and methods, resulting for example in NLP systems using knowledge representation languages with pre-defined syntax and hard-coded semantics. While this works for constrained and static domains, such as automated reply systems or natural language interfaces to automated booking systems, a more complex robotic system will need flexible linguistic and semantic structures (Roy and Reiter 2005). Moreover, robots need to address the symbol grounding problem (Harnad 1990), i.e. how to ground their knowledge in perception and action.

Following Belpaeme, Cowley and MacDorman (2007), who have edited a special issue attempting to reconceptualise symbol grounding from a distributed/cultural perspective, we now set out to look at language and robotics along similar lines. The present special issue collects a number of contributions which each argue for a flexible approach to language in robots, but more

*Corresponding author. Email: lsl@ua.pt

importantly, each sees language as being not restricted to the individual, not as a process located inside a language faculty, but instead as a dynamic and distributed cognitive process rooted in perception, action and culture (Belpaeme & Cowley 2007; Cowley 2007a, b; MacDorman 2007).

2. Robots with language

Looking at recent industrial and service robots and recent research prototypes, one finds many different kinds of robot interface. While science fiction has often influenced and even successfully predicted research and development, contemporary robots are still a long way from the machine interfaces of fiction literature and cinema. For almost a century robots have featured in science fiction and many of them use natural language. Famous examples include C-3PO, an android in the *Star Wars* series which could assist in translation between forms of communication in different cultures. In Clarke's *Space Odyssey* saga, HAL9000 was a computer-based agent capable of spoken language dialogue, face detection, reasoning and emotion processing (Stork 1997).

The reality of language-using robots in research labs and real-world applications is far from what science fiction portrays. However, language is increasingly used as a powerful tool in developing social robots (Seabra Lopes and Connell 2001; Fong, Nourbakhsh and Dautenhahn 2003). Such research developments can be traced back to SHRDLU, a natural language understanding system for a simulated robotic arm manipulating blocks (Winograd 1972). In the last decade, interest in robot assistants and humanoid robotics has led research groups around the world to develop language-based interfaces for robots. Even today, most use approaches not dissimilar to Winograd's SHRDLU system. Many of these projects implement complex integrated robotic systems to carry out real-world tasks. HERMES, for instance, is a humanoid robot that carried out guided museum tours for an extended period of time (Bischoff and Graefe 2002). Similarly, MARVIN gives tours around the lab where it was developed (Koch, Jung, Wettach, Nemeth and Berns 2008). JIJO-2 is an office robot that can guide visitors, deliver messages and arrange meetings (Asoh *et al.* 2001). ARMAR, another robot with a partially humanoid configuration, performs kitchen-related tasks (Stiefelhagen *et al.* 2004). ARTOS centralises the interface to an assisted living environment, an option that the developers justify by arguing that acceptance of an ambient system with a robot as communication partner is higher than of an ambient system where no communication partner is visible (Koch *et al.* 2008).

In the current state of the art, these robots tend to use a limited subset of natural languages such as English or Japanese. The syntax is usually hard-coded in the robot's software system, so that the user's utterances can be appropriately processed. The semantics of the utterances is often also hard-coded. An interesting corpus-based approach has been used for the BIRON robot (Hüwel, Wrede and Sagerer 2006). A large set of so-called 'situated semantic units' (SSU) was designed based on a corpus of situated conversational data. The SSUs, which encode networks of semantic relations representing world and discourse knowledge, form the base of semantic parsing. For the natural instruction-based robot described in Lauria, Bugmann, Kyriacou, Bos and Klein (2001), a so-called 'functional grammar' was also extracted from a corpus of route descriptions. The functional grammar is a list of primitive sentences, words and procedures users refer to when giving instructions to the robot. This vocabulary is grounded via hard-coding of each primitive procedure.

Given the current state of speech recognition technology (sensitivity to noise, to the speaker, to accents) and the spontaneous nature of speech (discontinuous and/or incomplete utterances), robustness of semantic interpretation is a major concern. In the CARL robot, a combination of deep and shallow parsing is used to cope with non-grammatical sentences (Seabra Lopes, Teixeira, Quinderé and Rodrigues 2008). In BIRON, interpretations are rated by semantic coherence,

which may in part be evaluated in the situated context, rather than by grammatical correctness (Hüwel *et al.* 2006).

A recent approach to disambiguate natural language is to use multimodal cues. Several of these robots use visual feedback to activate dialogs and interpret user utterances based on the situated context. The HRP-2 humanoid robot, for example, performs face detection and even portrait drawing (Ido, Matsumoto, Ogasawara and Nisimura 2006). Face detection and recognition is carried out by JIJO-2 (Asoh *et al.* 2001). Pointing gestures of the human interlocutor, as well as gaze direction, are used to constrain interpretations by BIRON (Toptsis, Haasch, Hüwel, Fritsch and Fink 2005), ARMAR (Stiefelhagen *et al.* 2004) and the robot of Hanafiah, Yamazaki, Nakamura and Kuno (2004). Some researchers are using incremental language processing algorithms, at syntactic and semantic levels, to increase robustness and improve disambiguation in situated contexts where typically there is visual feedback (cf. Brick and Scheutz 2007; Kruijff, Lison, Benjamin, Jacobsson and Hawes 2007). The focus is therefore shifting from detailed analysis of formal features to multimodal, directed and dialogical aspects of language.

Social robots learn from interacting with other members of a social group. The role of communication in robot learning is well recognised (Klingspor, Demiris and Kaiser 1997). Instructo-SOAR was one of the first systems to demonstrate learning based on interactive text-based teaching. This system acquired procedural knowledge to be used by a simulated robotic arm (Huffman and Laird 1993). HERMES and JIJO-2 learn maps of the environment based on human guidance. BIRON learns object names, when the user points to and names them. Robots that reason and learn need some sort of semantic memory. Kiatisevi, Ampornaramveth and Ueno (2006) describe SPAK, a frame-based knowledge management platform for interactive robots. In the CARL robot, a knowledge acquisition and management system, based on semantic network notions, enables the accumulation of semantic information obtained from human interlocutors. Question answering uses deductive and inductive inference and takes into account the speech recognition confidence associated to each piece of information (Seabra Lopes *et al.* 2008).

Other, perhaps atypical, approaches to making robots sensitive to language exist as well. Billard and Dautenhahn (1999) describe a robot that learns a synthetic proto-language. Roy and Pentland (2002), Steels and Kaplan (2002) and Seabra Lopes and Chauhan (2007) describe robotic agents that learn names of real-world objects based on interactions with a human caregiver. Ripley is a robotic arm that manipulates objects on a table (Roy, Hsiao and Mavridis 2004). By coupling active vision with physical simulation, Ripley is able to respond to spoken commands such as 'Hand me the blue thing on your left'.

All the above mentioned robots use natural language to some extent in their interactions with humans. However, despite all efforts, robots remain primitive at interpreting and producing natural language. Perhaps new developments in the language sciences can offer suggestions or new insights into what language is and how it can be implemented on machines.

3. The language sciences and robotics

The origins, evolution and acquisition of language and its role in human societies have long been studied by philosophers, linguists, psychologists, neuroscientists and cognitive scientists. In recent years, a distributed view of cognition and language has emerged (Hutchins 1995; Cowley 2007a, b). Control of embodied action becomes an emergent property of a distributed system composed of brain, body and environment. Language ceases to be identified with a formal system. It is both dynamical and symbolic and, for this reason, does not need to be contained within the individual. Instead, language is a heterogeneous set of distributed processes that are embedded in a culture which unites a group of individuals. As a representation, language is seen as a cultural product, perpetually open-ended and incomplete, and partly ambiguous (Love 2007; Linell 2007).

As such, language acquisition and evolution involve not only internal, but also cultural, social and affective processes.

In this context, new research questions open up: How does language transform human cognitive processes? How is language grounded in perception and action? In what ways does human phenomenology depend on linguistic experience? Can a distributed perspective on language clarify the nature of silent rehearsal (internal thought processes)? How can an individual or robot acquire language when language is distributed? How is language used to achieve joint experience? What is the embodied basis of social semiosis? How can the language sciences and the technological sciences aid each other? How can the language sciences contribute to the theoretical and practical study of how humans interact with robots?

While the language sciences have, until now, focused on language in human societies, the robotics and artificial intelligence communities are becoming increasingly active in developing user-friendly robots, that is, robots that are flexible, adaptable and easy to command and instruct. These artificial agents need to cognitively interpret perception and action, accumulate and manipulate semantic information for decision-making and interact with human subjects using natural language. To enable language spread, robots will need to use human action and perception (Cowley, this issue).

There are two main contact areas between robotics and the language sciences. First, robots can be used as simulation models for the empirical study of language origins, evolution and acquisition. This is an extension of the computational modelling approach to language. Second, current knowledge about language as a cultural product can be used to design and develop robots for practical applications.

For example, in semiotics it has been convincingly argued that in order for a word or linguistic expression to mean something, that word or expression needs to be grounded in perception (Harnad 1990). This is usually considered as an individual activity. It is assumed that an agent arrives at semantic representation through interaction with the environment. Semantic representations are assumed to be identical because, it is asserted, this is needed for *communication*. Recently, through experimental work, it has been suggested that linguistic representations and semantics are more flexible and that they adapt on the fly during dialogue (Pickering and Garrod 2004). While this process is active during dialogue, computer simulations and robotic experiments suggest that a similar process is active during language and concept acquisition (Steels and Belpaeme 2005). Individual agents coordinate their linguistic and semantics representations through repeated linguistic interactions. The process differs in timescale: coordinating representations can take hundreds or thousands of interactions. Interestingly, it illustrates how individual representations are shaped through language. This language is not maintained within a single individual, but is incomplete and distributed within a population of language users. Semantic representations do not arise solely through interacting with the environment. Rather, the semantics are agreed and maintained in a population of language users. In a nutshell, the meaning of words is distributed. This idea has been used to explain the typology of colour words and colour categories in humans (Belpaeme and Bleys 2005), but also suggests a novel method of acquiring conceptual information in computers and robots.

New developments in the language sciences will surely influence how language is approached in computer science and robotics. We have mentioned the distributed view, where language and all it affects is no longer contained in a single brain, but rather maintained in a population of language users. This suggests that for robots to be language-enabled, they will need to be sensitive to multi-modal cues from a group of language users, a suggestion which has been welcomed in recent years in the robotics community. Another development in the language sciences, of relevance to robotics, is the evidence for Whorfian effects, in which language impacts on cognition (Boroditsky 2001; Gilbert, Regier, Kay and Ivry 2006). This might show that the language we speak not only affects how we perceive the world, but that language impacts on knowledge acquisition

and shaping of knowledge. Another, perhaps less recent, development is the view that language achieves optimality on all its different fronts, for example speech (Liljencrants and Lindblom 1972), semantics (Hendriks and de Hoop 2001), and syntax and politics (Lazear 1999). This optimality of language, both in representation and transmission, can be found back in statistical natural language processing, but has the potential to inform the design and implementation of language-sensitive robots (see Hofe and Moore, this issue).

4. The papers in this issue

With this special issue, we aim to showcase synergies between robotics and the language sciences, identifying potential for collaboration. We are pleased to realise that the selected papers indeed highlight several synergies and areas of collaboration. The papers of **Hsiao *et al.***, **Seabra Lopes and Chauhan**, and **Doshi and Roy** present artificial agents and robots with the capacity to ground and/or use particular features of language. They have direct relevance for further developments in human-robot interaction but they can also be the subject of study and source of inspiration for work in the language sciences. The paper of **Anderson** highlights how the development of architectures for language-using robots may find inspiration in the organisation of the human brain. By contrast, the papers of **Hofe and Moore** and **Steels and Spranger** show how robotics can be a fundamental tool for the language sciences, by providing models that can be used to test contrasting theories.

Modularity is a widely applied engineering principle. The design of intelligent robots is often organised into modules responsible for such functions as natural language processing, planning, vision or motor control. Does the human brain follow a similar design? **Anderson** presents a discussion of modularity in the brain. The evidence he gathers suggests that the brain, far from being modular, works by redeploying resources. The author invites designers and implementers of language-using systems to design biologically inspired, non-modular and deeply integrated cognitive architectures.

Classical agent architectures in artificial intelligence were based on the sense-model-plan-act (SMPA) approach. Behaviour-based architectures rejected central models and explicit representations, focusing on responsiveness. Some hybrid architectures couple an SMPA layer with a behaviour-based control layer. In this issue, **Hsiao *et al.*** also contribute to the discussion on architectures. In their approach, reactive processes are the basic building block, as in a behaviour-based architecture. However, these processes are incorporated into object schemas and plan hierarchies. These representations can then be used to ground vocabulary for objects and actions. The approach is implemented on Trisk, a robotic arm that manipulates objects on a table. The used structures or representations are currently hard-coded to be consistent with a specific environment. In the longer term, the authors plan to develop learning mechanisms that enable the agent to dynamically construct representations (object schemas, etc.) that fit the setting.

Learning for language grounding is the focus of **Seabra Lopes and Chauhan**. They present an approach for long-term and open-ended category learning that relies on shared attention and teaching actions from a human caregiver to collect training instances. Such a learning approach can support vertical language spread, i.e. from a pre-existing group to a new robotic member. The approach includes a metacognitive component that takes advantage of the on-line nature of the language acquisition process to dynamically reconfigure the categorisation system and optimise memory usage. The paper also emphasises scalability and evaluation. A simple physical agent with visual feedback about objects on a table is used to test the developed learning system. The agent can learn reasonably large numbers of categories of real world objects. Different aspects of the meta-cognitive processing component are evaluated through systematic experiments.

Learning is also the focus of **Doshi and Roy**. They present a decision-theoretic approach for spoken dialogue management based on Partially Observable Markov Decision Processes (POMDP). The approach aims to achieve robustness with respect to noisy or ambiguous speech input as well as lack of knowledge on user preferences. To that end, the authors extend the standard POMDP model that allows a robot to learn user-specific dialogue models online, including new vocabulary and word choice preferences. The approach is demonstrated and evaluated for both a simulated domain and a real robotic wheelchair.

Hofe and Moore present an animatronic model of the human tongue and vocal tract for research on speech production. The goal of their project is to quantify energy consumption during human speech production and better understand how speech energy is optimised with regards to information throughput. One premise is that speech should not be considered as a unidirectional transmission of information, but as a dynamic process governed by both the speaker and the receiver. Hofe and Moore justify the choice for a robotic approach for the studies at hand based on a careful evaluation of alternatives, particularly computational methods based on formant, concatenative and articulatory synthesis. The paper specifies the construction of the robotic model and presents results on tongue positioning.

Steels and Spranger question how robots can build and maintain models of their bodies, such that they can control bodily movements, execute gestural actions and communicate about body parts and actions. They describe a range of setups under which a robot can acquire a body image and show how a group of robots can agree on a language to describe their body images and actions. An interesting conclusion is that a mirror system, while not strictly necessary, speeds up the acquisition of a body image. Their implementations and results might have interesting correlations in real languages, as recent results show that body categorisation is sensitive to culture (Majid, Enfield and van Staden 2006).

This special issue resulted from the *Symposium on Language and Robots*, an event that was organised by the Distributed Language Group (DLG) and took place at the University of Aveiro, Portugal, on 10-12 December 2007. The symposium, which followed the *External Symbol Grounding Workshop* (Plymouth, UK, 2006), brought together scholars who each step away from established approaches in natural language processing, and explore the distributed and extended nature of language. They demonstrate how this view can be used to build language-enabled robots and how constructing new robots can serve to generate and support new insight in the language sciences. Given the large potential for collaboration, a second edition of the *Symposium on Language and Robots* is planned for the near future.

References

- Asoh, H., Motomura, Y., Asano, F., Hayamizu, S., Itou, K., Kurita, T., Matsui, T., Vlassis, N., Bunschoten, R., and Kröse, B. (2001), 'Jijo-2: An Office Robot that Communicates and Learns,' *IEEE Intelligent Systems*, 16(5), 46–55.
- Belpaeme, T., and Bleys, J. (2005), 'Explaining Universal Colour Categories Through a Constrained Acquisition Process,' *Adaptive Behavior*, 13(4), 293–310.
- Belpaeme, T., and Cowley, S.J. (2007), 'Extended Symbol Grounding,' *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, 8(1), 1–6.
- Belpaeme, T., Cowley, S.J., and MacDorman, K.F., eds. (2007), *Symbol Grounding*, Special Issue of *Interaction studies: Social behaviour and communication in biological and artificial systems*, 8(1).
- Billard, A., and Dautenhahn, K. (1999), 'Experiments in Learning by Imitation: Grounding and Use of Communication in Robotic Agents,' *Adaptive Behavior*, 7(3–4), 415–438.
- Bischoff, R., and Graefe, V. (2002), 'Dependable Multimodal Communication and Interaction with Robotic Assistants,' *Proc. 11th IEEE International Workshop on Robot and Human Interactive Communication*, pp. 300–305.
- Boroditsky, L. (2001), 'Does Language Shape Thought? English and Mandarin Speakers' Conceptions of Time,' *Cognitive Psychology*, 43, 1–22.
- Brick, T., and Scheutz, M. (2007), 'Incremental Natural Language Processing for HRI,' *Proc. ACM/IEEE International Conference on Human-Robot Interaction (HRI'07)*, pp. 263–270.
- Cowley, S.J. (2007a), 'The Cognitive Dynamics of Distributed Language,' *Language Sciences*, 29(5), 575–583.

- Cowley, S.J. (2007b), 'Distributed Language: Biomechanics, Functions and the Origins of Talk,' In *Emergence of Communication and Language*, C. Lyon, C. Nehaniv and A. Cangelosi, Heidelberg, London: Springer, pp. 105–127.
- Fong, T., Nourbakhsh, I., and Dautenhahn, K. (2003), 'A Survey of Socially Interactive Robots: Concepts, Design, and Applications,' *Robotics and Autonomous Systems*, 42, 143–166.
- Gilbert, A., Regier, T., Kay, P., and Ivry, R. (2006), 'Whorf Hypothesis is Supported in the Right Visual Field but not the Left,' *Proceedings of the National Academy of Sciences*, 103(2), 489–494.
- Hanafiah, Z.M., Yamazaki, C., Nakamura, A., and Kuno, Y. (2004), 'Understanding Inexplicit Utterances using Vision for Helper Robots,' *ICPR 2004. Proceedings of the 17th International Conference on Pattern Recognition*, pp. 925–928.
- Harnad, S. (1990), 'The Symbol Grounding Problem,' *Physica D*, 42, 335–346.
- Hendriks, P., and de Hoop, H. (2001), 'Optimality Theoretic Semantics,' *Linguistics and Philosophy*, 24(1), 1–32.
- Huffman, S.B., and Laird, J.E. (1993), Learning Procedures from Interactive Natural Language Instructions, *Proc. 10th International Conference on Machine Learning*, Amherst, MA, pp. 143–150.
- Hutchins, E. (1995), *Cognition in the Wild*, MIT Press, Cambridge MA.
- Hüwel, S., Wrede, B., and Sagerer, G. (2006), 'Robust Speech Understanding for Multi-Modal Human-Robot Communication,' *Proc. 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN06)*, Hatfield, UK, September 6–8, 2006, pp. 45–50.
- Ido, J., Matsumoto, Y., Ogasawara, T., and Nisimura, R. (2006), 'Humanoid with Interaction Ability Using Vision and Speech Information,' *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1316–1321.
- Kiatissevi, P., Ampornaramveth, V., and Ueno, H. (2006), A Frame-based Knowledge Software Tool for Developing Interactive Robots, *International Journal of Artificial Life and Robotics*, 10, pp. 18–28.
- Klingspor V., Demiris J., and Kaiser, M. (1997), 'Human-Robot-Communication and Machine Learning,' *Applied Artificial Intelligence Journal*, vol. 11, pp. 719–746.
- Koch, J., Jung, H., Wettach, J., Nemeth, G., and Berns, K. (2008), 'Dynamic Speech Interaction for Robotic Agents,' *Recent Progress in Robotics: Viable Robotic Service to Human*, Lecture Notes in Control and Information Sciences, vol. 370/2008, Springer Berlin/Heidelberg, pp. 303–315.
- Kruijff, G.-J. M., Lison, P., Benjamin, T., Jacobsson, H., and Hawes, N. (2007), 'Incremental, Multi-Level Processing for Comprehending Situated Dialogue in Human-Robot Interaction,' *Language and Robots: Proceedings of the Symposium*. 10–12 December 2007, Aveiro, Portugal, Seabra Lopes, L., T. Belpaeme, S.J. Cowley (eds.), Universidade de Aveiro, 55–64.
- Lauria, S., Bugmann, G., Kyriacou, T., Bos, J., and Klein, E. (2001), 'Training Personal Robots Using Natural Language Instruction,' *IEEE Intelligent Systems*, 16(5), 38–45.
- Lazear, E. P. (1999), 'Culture and language,' *Journal of Political Economy*, 107(S6), 95–126.
- Liljencrants, L., and Lindblom, B. (1972), 'Numerical Simulations of Vowel Quality Systems,' *Language*, 48, 839–862.
- Linell, P. (2007), 'Dialogicality in Languages, Minds and Brains: is There a Convergence Between Dialogism and Neuro-Biology?' *Language Sciences*, 29(5), 605–620.
- Love, N. (2007), 'Are languages digital codes?' *Language Sciences*, 29(5), 690–709.
- MacDorman, K.F. (2007), 'Life after the Symbol System Metaphor,' *Interaction Studies*, 8(1), 143–158.
- Majid, A., Enfield, N. J., and van Staden, M. (2006), 'Parts of the Body: Cross-Linguistic Categorisation,' *Language Sciences*, 28(2–3), 137–147.
- Pickering, M.J., and Garrod, S. (2004), 'Towards a Mechanistic Psychology of Dialogue,' *Behavioral and Brain Sciences*, 27, 169–226.
- Roy, D., and Pentland, A. (2002), 'Learning Words From Sights and Sounds: A Computational Model,' *Cognitive Science*, 26, 113–146.
- Roy, D., and Reiter, E. (2005), 'Connecting Language to the World,' *Artificial Intelligence*, 167, 1–12.
- Roy, D., Hsiao, K.-Y., and Mavridis, N. (2004), 'Mental Imagery for a Conversational Robot,' *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 34(3): 1374–1383.
- Seabra Lopes, L., and Chauhan, A. (2007), 'How Many Words Can My Robot Learn? An Approach and Experiments with One-Class Learning,' *Interaction Studies*, 8(1), pp. 53–81.
- Seabra Lopes, L., and Connell, J.H. (2001), 'Semisentient Robots: Routes to Integrated Intelligence,' *IEEE Intelligent Systems*, 16(5), 10–14.
- Seabra Lopes, L., Teixeira, A.J.S., Quinderé, M., and Rodrigues, M. (2008), 'Integrated Capabilities for Knowledge Acquisition through Spoken Language Interaction in a Mobile Robot,' *Autonomous Robot Systems and Competitions: Proceedings of the Eighth Conference*. 2 April 2008, Aveiro, Portugal, Universidade de Aveiro, pp. 79–86.
- 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.
- Steels, L., and Kaplan, F. (2002), 'AIBO's first words: The Social Learning of Language and Meaning,' *Evolution of Communication*, 4(1), 3–32.
- Stiefelhagen, R., Fügen, C., Gieselmann, P., Holzapfel, H., Nickel, K., and Waibel, A. (2004), 'Natural Human-Robot Interaction using Speech, Head Pose and Gestures,' *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2422–2427.
- Stork, D. (1997), *HAL's Legacy: 2001's Computer as Dream or Reality*, MIT Press, Cambridge, MA.
- Toptsis, I., Haasch, A., Hüwel, S., Fritsch, J., and Fink, G.A. (2005), 'Modality Integration and Dialog Management for a Robotic Assistant,' *Proceedings of the INTERSPEECH 2005*, Lisbon, Portugal, pp. 837–840.
- Winograd, T. (1972), *Understanding Natural Language*, Academic Press, San Diego, California.