Like the breathability of air

Embodied embedded communication

Pim Haselager
Donders Institute for Brain, Cognition and Behaviour, Radboud University Njmegen

I present experimental and computational research, inspired by the perspective of Embodied Embedded Cognition, concerning various aspects of language as supporting Everett's interactionist view of language. Based on earlier and ongoing work, I briefly illustrate the contribution of the environment to the systematicity displayed in linguistic performance, the importance of joint attention for the development of a shared vocabulary, the role of (limited) traveling for language diversification, the function of perspective taking in social communication, and the bodily nature of understanding of meaning.

Keywords: dialect formation, embodied embedded cognition, embodied meaning, joint attention, language, recipient design, systematicity

1. Introduction

In Everett's wide-ranging and profoundly stimulating book, language is studied at the intersection of culture, cognition and communication, and some of its properties are likened to "the breathability of air" (p. 151). In this paper I would like to present some of my work, inspired by the general perspective of Embodied Embedded Cognition (EEC), as being very much in line with, and in support of, Everett's interactionist approach, although the material presented will be of a computational and experimental kind, instead of anthropological. Hopefully, these different forms of studying language can strengthen the interactionist perspective.

I welcome Everett's emphasis on focusing at where 'everything comes together'. To my mind, in the past cognitive science has suffered from allowing its reliance on the standard, and itself methodologically sound explanatory strategy of 'divide and conquer' to color its portrayals of the phenomena in question. Obviously, to study a phenomenon experimentally, as in cognitive neuroscience and psychology, one has to take the phenomenon out of its natural context and investigate it

Pragmatics & Cognition 20:2 (2012), **263–274.** DOI 10.1075/pc.20.2.04has
ISSN 0929–0907 / E-ISSN 1569–9943

under ideally completely controlled conditions, probing it in various ways while observing its behavior as closely as possible.[1] This is sound scientific practice and the diminished ecological validity of the observations is often a reasonable price to pay for the detailed and replicable observations obtained. But a sound scientific practice has unfortunate consequences, if one starts forgetting that the phenomenon thus studied is not 'the real thing', but an artificial construction. One can learn a lot about humans by systematically manipulating and observing their shadows, but such an approach will never do justice to the full richness of human behavior and cognition. In my view, this is one of the most general and important messages of the Embodied Embedded Cognition approach (see, e.g., Thelen and Smith 1994; Clark 1997; Haselager et al. 2008; van Dijk et al. 2008) that gained momentum in the 90's. The title of Hutchinson's (1995) classic, *Cognition in the Wild*, best captures this idea.

2. Systematicity and the breathability of air

A case in point, and of great influence regarding the study of language, is Chomsky's competence — performance distinction. As is well known, Chomsky (1965: 4) proposed to make "a fundamental distinction between the competence (the speaker-hearer's knowledge of his language) and performance (the actual use of language in concrete situations)". The competence outlines a space of possible behaviors of which the actual behavior of an agent provides only a sample (Fodor 1968: 130–131; Haselager 1997: 107–109). Thus, competence becomes an idealized version of the phenomenon in question, "unadulterated" by the influence of extraneous factors (Katz 1974: 232). As in the case of experimentation, specifying an idealized version of a phenomenon can be good scientific practice, as long as one does not get blinded by its unadulterated beauty. As Simon (1981: 23) noted: "There is a continuing danger that focus upon an ideal competence that resides in some kind of Platonic heaven (or a Cartesian one) will impose normative constraints on the study of actual language behavior".

The debate about productivity and systematicity provides an interesting illustration of how an idealized competence may bias the perception of what needs to be explained, as well as how. One may attempt to explain the productivity and systematicity displayed by linguistic behavior (and thought) by postulating a representational system with constituent structure and compositional syntax and semantics (Fodor and Pylyshyn 1988; Fodor and McLaughlin 1990), such as Fodor's (1975) language of thought. But, as Everett notes (p. 94), at least one of the questions to ask is "Is there a simpler explanation?" This has led to many attempts to model the systematicity of language on the basis of a representational system without constituent structure (e.g., distributed representations utilized by connectionist networks; van Gelder 1990), that have, however, encountered significant difficulties (e.g., Hadley 1994; Haselager and van Rappard 1998). One recent attempt, that I think is congenial to Everett's perspective (e.g., pp. 131–135, 138–139), is to see systematicity of language as capitalizing on the structure inherent in an organism's environment rather than its internal representational system (Frank, van Rooij, and Haselager 2009).

This is not the place for going into the details of the computational model presented in that paper, so I here merely wish to point out its underlying idea. Systematicity as displayed in the *actual practice* of human cognition and language may depend on only 'weak' representational resources *combined* with a largely systematic world under an appropriately wide variety of circumstances (Frank et al. 2009: 374). Likewise, the displayed systematicity of the connectionist model developed by Frank results from the interaction between the architecture and its environment. That is, the representational system need not have full-blown constituent structure, nor does the solution need to lie in strong architectural constraints and quite specific training or learning processes, typical for many early connectionist modeling attempts. From this perspective, a representational system capable of *reflecting* the systematicity in the environment might suffice for displaying a psychologically plausible degree of systematicity. "Under this view, similarities among languages, such as they are, would be like resemblances between bows and arrows across the world", Everett says (p. 87). Likewise, we argue that it is the organism's embeddedness in the world, rather than its internal structure, from which the systematicity of its linguistic behavior derives. I think that Everett's remark (p. 151), made in relation to the learnability of language, about the breathability of air not necessarily being a feature of air but a feature of life as it evolved on this planet, also nicely captures the intent of our work. Systematicity need not be a feature of a representational system (e.g., in virtue of it having a constituent structure), but a feature of the world in which cognitive systems evolved.

3. Social interaction and sharing a vocabulary

Similarly, an organism's embeddedness, specifically related to the interaction with other agents, might help to explain the development of shared vocabularies. In Kwisthout et al. (2008) we took up Tomasello's (1999) suggestion that the ability to engage in various forms of joint attention provides an essential mechanism for, and may have co-evolved with, cultural learning and language. In checking attention, participants (e.g., a child and parent) consider what the other is focusing his/her attention on (establishing what the other is looking at, for instance). In following attention, one is being led by the other to shift one's attention to focus on a new object or event. In directing, one does the opposite, namely actively guiding another's attention to a specific object or event. A computational approach to the study of how agents can develop a common lexicon (shared word-meaning mappings), is to have artificial agents (robots, software agents) play language games, i.e., agents exchanging utterances that stand for features (e.g., shape, color) of objects present in a shared environment (Steels 2001; Steels and Kaplan 2002). The task for these agents is to end up with a common lexicon, a shared set of symbols, so they have to solve the *social symbol grounding problem* (Cangelosi 2006; Vogt and Divina 2007). Everett provides many wonderful illustrations of this problem when describing his attempts to communicate with the Pirahã and other native tribes.

The three forms of joint attention allow individuals to reduce the number of possible meanings when learning a new word. As is well known, Quine (1960) pointed out that each unfamiliar word that we learn could in principle mean an infinite number of things (*gavagai* might mean 'rabbit', 'undetached rabbit parts', 'running furry animal', 'dinner', 'it will rain', and so on). So, in order to learn the meaning of a word in actual practice, one must be able to substantially reduce the number of possible hypotheses. Our focus was on the role of the three joint attention skills for reducing the uncertainty with which the meanings of words can be inferred, by examining their effects on achieved accuracy and on how much they speed up the language game.

Checking attention turned out to be most effective. In our computer simulations we found that in the absence of the capacity to check attention, communicative accuracy simply failed to reach 100%, even if the two other attention skills were present. However, when the capacity for checking attention was present, all simulations converged to 100% communicative accuracy. Still, following and directing attention played a major role in increasing the speed with which a 100% communicative accuracy was achieved, going from around 66,000 language games for checking attention alone to about 18,000 games for directing attention and 2500 games for following attention. Our results (only briefly and incompletely summarized here) indicate that the ability to *check* attention is more crucial than the ability to *follow* attention, which in turn appeared to be more crucial than the ability to *direct* attention. The large differences in the effectivity of these skills for the language game may further be used to suggest hypotheses about their evolutionary ordering.

Assuming that the most effective mechanism evolved first, if our simulations are correct in indicating that checking attention is of primary importance for creating a common lexicon during language games, then it would have been beneficial for early hominids to have these capabilities before more advanced language usage could emerge. So checking attention may have evolved first, following attention second, and directing attention last in human evolution. Of course such general hypotheses would need more investigation, but they might serve as a rough indication of how computational models can lead to evolutionary hypotheses that are in line with Everett's general interactive perspective on communication. Moreover, including the child's growing participation in more complex social interactions in the dynamics of the computational model provides a concrete illustration of the embeddedness of cognition.

4. Technologies don't remain the same: Vowel change

Everett reminds us that language is a cognitive technology (p. 46; cf. Dascal 2002), and that as a tool it is both shaping us and being shaped by us (p. 218). Language changes all the time, and, as far as we can establish, has become more varied and diverse over time. "Why did English stop being German?" Everett asks (p. 325), and answers: "For the exact reasons that you talk more like the friends you grew up with than your own parents; for the reason that you talk more like members of your economic class than another. The basic rule of language change, again, is 'You talk like who you talk with'". As another example of how computational modeling might assist in our understanding of language, and of how much it is connected with the embodied embeddedness of our cognition, I would like to briefly introduce a recent computational model of vowel change as developed by du Pau et al. (in prep.).

Du Pau created a language simulation tool (entitled DEViL for 'Dialect Emergence Virtual Lab') that was used to implement a computational model of vowel change, based on a theoretical framework consisting of three hypotheses coming from Croft's (2000) general theory of language change. The three hypotheses concerning language change are: (1) the *formation hypothesis*, according to which reinforcement of communicative success and punishment of disuse drive the formation of a shared lexicon; (2) the *change hypothesis*, stating that alignment to variation in perceived signals (vowels in our model) drives linguistic change; and (3) the *divergence hypothesis*, claiming that social fragmentation leads to linguistic divergence. In our model, it's where the agents are, and who they interact with, that drives language change; so, during runs of the computational model, agents could do little else but walk and talk, moving to (or staying in) specific regions and verbally interacting with the 'locals'. The communicative behaviour of the agents could be influenced in the model through various factors such as entrenchment of signals (due to reinforcement after communicative success), alignment (the adjustment of an individual's pronunciation to their partner's pronunciation), memory decay (the forgetfulness of an agent concerning the meaning or pronunciation of a signal), and noise.

The verbal interactions consisted of the basic language game (Steels 2001) described in the previous section. In one of the simulations we ran, we tested the hypothesis that geographical constraints are a potential cause of vowel divergence. The distance between languages depends on the interactions between members of the groups that speak them. We simulated agents of two neighbouring speech communities that could change their vowel systems simultaneously for 250,000 time steps, while we manipulated the amount of travelling to and fro. We found that physically separated communities that start with the same language gradually develop their own language systems when simulated geographical or other constraints reduce travelling between them. The speed and degree of divergence depended roughly proportionally on the severity of the constraints imposed on the travel opportunities. Under conditions with a lot of travelling, the languages were found to remain shared. Little travelling led to dialects and eventually non mutually comprehensible languages.

5. Social communication and recipient design

How do humans tune their communicative behaviors to different types of agents, e.g., people that differ (to widely varying extents) in their language and culture? In Amsterdam, where I live, I get asked the way to the Rijksmuseum, van Gogh museum, red light district, concert hall, or just a particular street repeatedly, sometimes by locals, more often by Dutch people from out of town, and very often by people from other countries. Obviously in such cases, we share a lot of common ground, probably much more than found during Everett's fascinating encounters (even if my questioners are from very different countries and cultures, we share at least some knowledge about what a city, a museum, a bus or a taxi are). Still, I can notice myself addressing all these different (types of) questioners differently, e.g., by speaking more slowly and/or loudly, making clearer gestures, etc., often adapting my behavior in response to their facial and bodily expressions of understanding or bafflement. The tuning of my communicative signals to the person in front of me is known as recipient design (Sacks, Scheglof, and Jefferson 1974), discussed by Hockett (1960) under the label of 'feedback'. Cases can be even more complex when I also (have to) take into account the group of overhearers of which the addressee may be part, monitoring their responses in addition to the one who asked me the question, and adapting my communicative actions to them simultaneously (audience design). How do I (and everyone else) do it?

One way to understand my behavior is that I am forming (subconsciously or consciously) hypotheses about what the questioner might want, believe and know, and use this type of perspective taking to improve my communicative actions (Clark and Carlson 1982; Grice 1974, 1989; Levelt 1989). However, one might suggest that involving all this cognitive machinery of hypotheses generation, testing, and updating is computationally very expensive, especially considering the generally easy and rapid communicative performance displayed in everyday circumstances. An alternative, recently attracting adherents, would be that recipient design is done through applying simple heuristics or rules-of-thumb triggered by the presence or absence of certain cues (Galati and Brennan 2010; Epley et al. 2004; Shintel and Keysar 2009). One may think of the example of a dog increasing his pawing when he is being ignored in his entreaties for a treat as a case of the dog following the inferentially relatively simple strategy of recipient design based on heuristics (e.g., 'when insufficient response, repeat more often').

TCG is a communicative task where two players, a sender and a receiver, play a game on a 3 x 3 grid board. The sender knows both her own goal state and that of the receiver, and has to communicate, through game movements only (no verbal or gestural communication is possible during the experiments), the goal state to the receiver, and reach her own goal state as well. Basically, the sender communicates about the game, using game movements only, while playing the game at the same time. This setup allows to study different communication strategies, and their adaptation, in an abstract form. Based on experiments and formal analysis (that, again, do not allow rapid recounting here) we suggest that, even under such relatively impoverished experimental conditions, communication is better explained by the suggestion that the sender engages in a form of perspective taking. Errors on the part of the receiver help to clarify how the receiver is misinterpreting the speakers' communicative intentions, helping the speaker to hypothesize about the 'why' of misinterpretation, which can then form the basis for adjustment of the new communicative signals. In communication, one is really trying to understand the other, in order to be better understood oneself. Even during such brief meetings as with a tourist asking the way in Amsterdam, culture, cognition, and communication come together.

6. Telling stories: Embodied meaning

Finally, I'd like to focus on an example of experimental work in relation to meaning. Meaning is obviously crucial to Everett's studies. He rightly points out (p. 47) that questions about communication are preceded by the question about *why* people want to communicate. This question leads him to the investigation of culture and how communication made it possible to share values (p. 48–49). This seems entirely sensible to me. However, although Everett discusses the importance of the body for culture very often, and especially how culture influences or shapes the body (e.g., p. 28) he does not refer to work in cognitive neuroscience that shows how the body, i.e.,our sensorimotor capacities, is crucial to our understanding the meaning of language. Given his perspective, this is understandable, and I hope here to provide some information, based on a paper by Kerkhofs and Haselager (2006), regarding an experiment of Glenberg and Kaschak (2003), that may complement Everett's interactive position.

The issue is how representations acquire meaning for the system that has them. Harnad (1990: 335) formulated the basic question as follows: "How can the semantic interpretation of a formal symbol system be made intrinsic to the system, rather than just parasitic on the meanings in our heads? (…) The problem is analogous to trying to learn Chinese from a Chinese/Chinese dictionary alone". The study of the *embodiment* of cognition can be helpful here. The body is more than a mere transducer of information between the organism and the environment: It actively shapes the form cognitive tasks can take and also presents possibilities for solving them (Clark 1997; Chiel and Beer 1997; Lakoff and Johnson 1999; Haselager 2004). Therefore one might expect to find traces of an organism's sensorimotor interactions with its environment in the way organisms understand and respond to meaning. The basic idea is that meaning depends on an individual's history of bodily interactions with the world. People recreate or simulate those experiences in response to linguistic input, and use them to produce meaningful behavioral (including but not limited to) linguistic output. From this perspective, perceptual and motor processes are not peripheral to but form the *core* of mental content.

In Glenberg and Kaschak's (2003) experiment participants had to decide whether certain sentences were sensible (e.g., Andy handed you the pizza) or nonsense (e.g., Leonard drank the sun). Three kinds of transfer sentence were used: Imperatives like (1), concrete transfer sentences like (2), and abstract transfer sentences like (3). Note that all sentences have a form where the transfer or movement is towards the "you" person, and an 'opposite' form where the transfer or movement is away from the "you" person.

(1) Open the drawer / Close the drawer

(2) You handed Andy the pizza / Andy handed you the pizza

(3) You told Jim the story / Jim told you the story

The participants had to make their yes/no judgments with a button-box with three vertically aligned buttons. After pressing the middle button the sentence appeared on the screen. The position of the "yes" button was either above or below the middle button. This means that the motor response that the participant has to make is either in line or in conflict with the direction of the motion that is described in the test sentences. Glenberg and Kaschak found that participants are faster to accept a sentence (to respond "yes") when the motor response and the described action *matched*. Importantly, this also applied to the more metaphorical direction involved in telling someone a story versus being told a story by someone. An embodied sense of direction exists even in understanding meaning in that case.

7. Conclusion

There is much more in this rich and readable book that invites links to ongoing research in experimental psychology, cognitive neuroscience, and artificial intelligence. Traditionally these sciences may have taken a more single-minded, less inclusive, less interactive perspective on language, cognition, and culture. Over the last few decades, however, there is a clear tendency to turn back to 'cognition in the wild' and, though still using experimental and computational methods, this tendency at least prevents the methodology from obscuring the phenomena to be studied. A recent plea for a more inclusive approach in cognitive neuroscience is presented by Hasson et al. (2012: 114): "Cognition materializes in an interpersonal space (…) Despite the central role of other individuals in shaping one's mind, most cognitive studies focus on processes that occur within a single individual. We call for a shift from a single-brain to a multi-brain frame of reference". While I applaud the move to a multi-brain frame of reference, I think Everett is doing us a great service in clearly demonstrating that this in itself is not enough. Language "is the cognitive fire of human life" says Everett (p. 327), and to understand the cognitive wildfires of natural organisms, we need to look at much more than just brains, even taken together. The fuel of cognitive fire consists of embodied practices, embedded ways of living and speaking that mesh imperceptibly with one another, then diverge or even split at times, coming back together at other moments. Everett's book shows us how warming it can be to consider this fire in its natural splendor.

References

Blokpoel, M., Van Kesteren, M., Stolk, A., Haselager, P., Toni, I., and Rooij, I. 2012. Recipient design in human communication: Simple heuristics or perspective taking? *Frontiers in Human Neuroscience* 5(52): 1–18. doi: 10.3389/fnhum.2012. 00253.

Cangelosi, A. 2006. The grounding and sharing of symbols. *Pragmatics and Cognition* 14(2): 275–285.

Chiel, H.J. and Beer, R.D. 1997. The brain has a body: Adaptive behavior emerges from interactions of nervous system, body, and environment. *Trends in Neurosciences* 20: 553–557.

Clark, A. 1997. *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: The MIT Press.

Clark, H. and Carlson, T. 1982. Hearers and speech acts. *Language* 58(2): 332–373.

Croft, W. 2000. *Explaining Language Change: An evolutionary Approach*. Harlow: Pearson Education.

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

Dascal, M. 2002. Language as a cognitive technology. *International Journal of Cognition and Technology* 1(1): 35–89.

de Ruiter, J., Noordzij, M., Newman-Norlund, S., Newman-Norlund, R., Hagoort, P., and Levinson, S., 2010. Exploring the cognitive infrastructure of communication. *Interaction Studies* 11(1): 51–77.

de Ruiter, J., Noordzij, M., Newman-Norlund, S., Hagoort, P., and Toni, I. 2007. On the origin of intentions. In Y. Rosseti, P. Haggard, & M. Kawato (eds.), *Sensorimotor Foundations of Higher Cognition*. Oxford: Oxford University Press, 601–618.

du Pau, B., Dijkstra, T., Haselager, W.F.G., and van Heuven, W. In prep. A model for simulating vowel change and its underlying mechanisms.

Epley, N., Keysar, B., Van Boven, L., and Gilovich, T. 2004. Perspective taking as egocentric anchoring and adjustment. *Journal of Personality and Social Psychology* 87(3): 327–339.

Fodor, J.A. 1968. *Psychological Explanation: An Introduction to the Philosophy of Psychology*. New York: Random House.

Fodor, J.A. 1975. *The Language of Thought*. Cambridge, MA: Harvard University Press.

Fodor, J.A. and McLaughlin, B. 1990. Connectionism and the problem of systematicity: Why Smolensky's solution does not work. *Cognition* 35: 183–204.

Fodor, J.A. and Pylyshyn, Z.W. 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition* 28: 3–71.

Frank, S., van Rooij, I., and Haselager, W.F.G. 2009. Connectionist semantic systematicity. *Cognition* 110: 358–379.

Galati, A. and Brennan, S. 2010. Attenuating information in spoken communication: For the speaker, or for the addressee? *Journal of Memory and Language* 62(1): 35–51.

Grice, H.P. 1975. Logic and conversation. In P. Cole and J. Morgan (eds), *Syntax and Semantic 3: Speech Acts*. New York: Academic Press, 41–58.

Grice, H.P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Hadley, R. F. (1994). Systematicity in connectionist language learning. *Mind & Language* 9(3): 247–272.

Haselager, W.F.G. 1997. *Cognitive Science and Folk Psychology: The Right Frame of Mind*. London: Sage.

Haselager W.F.G. and van Rappard J.F.H. 1998. Connectionism, systematicity, and the frame problem. *Minds and Machines* 8(2): 161–179.

Haselager, W.F.G., van Dijk, J., and van Rooij, I. 2008. A lazy brain? Embodied embedded cognition and cognitive neuroscience. In F. Calvo Garzon and A. Gomila (eds). *Handbook of Embodied Cognitive Science*. Oxford: Elsevier, 273–290.

Hasson, U., Ghanzafar, A., Galantucci, B., Garrod, S. & Keysers, C. (2012). Brain-to-brain coupling: A mechanism for creating and sharing a social world. *Trends in Cognitive Sciences* 16(2): 114–121.

Hockett, C.F. 1960. Origins of language. *Scientific American* 203: 88–96.

Hutchinson, E. 1995. *Cognition in the Wild*. Cambridge, MA: The MIT Press.
Katz, J.J. 1974. The relevance of linguistics to philosophy. In G. Harman (ed), *On Noam Chomsky: Critical Essays*. New York: Anchor Books, 229–241.
Kerkhofs, R. and Haselager, W.F.G. 2006. The embodiment of meaning. *Manuscrito* 29(2): 753–764.
Kwisthout, J., Vogt, P., Haselager, W.F.G., and Dijkstra, T. 2008. Joint attention and language evolution. *Connection Science* 20(2–3): 155–171.
Lakoff, G. and Johnson, M. 1999. *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. New York: Basic Books.
Levelt, W. 1989. *Speaking: From Intention to Articulation*. Cambridge, MA: The MIT Press.
Noordzij, M., Newman-Norlund, S., De Ruiter, J., Hagoort, P., Levinson, S., and Toni, I. 2009. Brain mechanisms underlying human communication. *Frontiers in Human Neuroscience* 3(14): 14–27.
Quine,W.V.O. 1960. *Word and Object*. Cambridge, MA: The MIT Press.
Sacks, H., Schegloff, E., and Jefferson, G. 1974. A simplest systematics for the organization of turn-taking for conversation. *Language* 50(4): 696–735.
Shinkel, H. and Keysar, B. 2009. Less is more: A minimalist account of joint action in communication. *Topics in Cognitive Science 1*(2): 260-273.
Simon, H.A. 1981. Cognitive science: The newest science of the artificial. In D.A. Norman (ed), *Perspectives on Cognitive Science* . Norwood: Ablex, 13–25.
Steels, L. 2001. Language games for autonomous robots. *IEEE Intelligent Systems* 16: 16–22.
Steels, L. and Kaplan, F. 2002. Bootstrapping grounded word semantics. In T. Briscoe (ed), *Linguistic Evolution Through Language Acquisition: Formal and Computational Models*. Cambridge: Cambridge University Press, 53–74.
Tomasello, M. 1999. *The Cultural Origins of Human Cognition*. Cambridge, MA: Harvard University Press.
Thelen, E. and Smith, L.B. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action.* Cambridge, MA: The MIT Press.
Van Dijk, J., Kerkhofs, R., van Rooij, I., and Haselager, W.F.G. 2008. Can there be such a thing as embodied embedded cognitive neuroscience? *Theory & Psychology* 13(3): 297–316.
Van Gelder, T. 1990. Compositionality: A connectionist variation on a classical theme. *Cognitive Science* 14: 355–384.
Van Rooij, I., Kwisthout, J., Blokpoel, M., Szymanik, J., Wareham, T., and Toni, I. 2011. Intentional communication: Computationally easy or difficult? *Frontiers in Human Neuroscience* 5(52): 1–18.
Vogt, P. and Divina, F. 2007. Social symbol grounding and language evolution. *Interaction Studies* 8: 31–52.

Author's address

Pim Haselager
Donders Institute for Brain, Cognition and Behaviour
Radboud University Nijmegen
B.02.40, Spinozagebouw, Montessorilaan 3
6525 HR Nijmegen, The Netherlands

pimhabra@yahoo.com
http://www.dcc.ru.nl/~haselag

About the author

Pim Haselager is an Associate Professor at the Radboud University Nijmegen, a principal investigator at the Donders Institute for Brain, Cognition and Behaviour, and the head of the educational program of Artificial Intelligence (AI). He has master degrees in Psychology and Philosophy, and a PhD in Theoretical Psychology. He works primarily within the philosophy of mind and of cognitive science, with a focus on AI and Cognitive Neuroscience (CNS). He is particularly interested in the integration of empirical work (experimentation, computational modeling, and robotics) with philosophical issues regarding knowledge, agency, and intelligent behavior. He investigates the ethical and societal implications of research in the fields of AI and CNS, and the ensuing technologies (e.g., Robotics, Brain-Computer Interfacing, Deep Brain Stimulation). He is a member of the Board of Consulting Editors of *Pragmatics & Cognition* and has recently published in *Neuroethics* (2011, 2010), *Philosophical Psychology* (2011), *Social Neuroscience* (2011), *Journal of Neural Engineering* (2009), *Neural Networks* (2009), *Cognition* (2009), *Connection Science* (2008), and *Pragmatics & Cognition* (2005).