

# Cui Bono Robo Sapiens?

Alois Knoll

Ever since Czech writer Karel Capek's well-known play "R.U.R. – Rossum's universal robots" [Capek], first published in 1920 and performed in 1921, man-made robots of human-like shape have inspired fiction writers to envision worlds with artificial creatures far superior to human beings – either in friendly coexistence or taking over by eliminating their creators and leaving them behind as an ephemeral step in the evolution of life on earth. Roboticists and AI researchers, by contrast, have come to realise that there is still a long way to go if only parts of such visions are to come true. Almost every aspect of research on humanoid robotics that has been touched on across scientific communities has taken researchers to the edge of current technology. Moreover, it has also become obvious how limited our knowledge about ourselves is when it comes to implanting those skills into a mechanical body that are necessary to enable a robot to mimic basic aspects of human intelligence. However, due to recent developments in enabling technologies [Brooks] (processing power, mechatronics, walking machines, articulated vision heads and more) and also due to findings and developments in other fields (e.g. studies of the human brain, linguistics, psychology), we currently observe a shift in the view of what artificial intelligence is and how it can be put to work in operational autonomous systems. This sets the stage for putting perceptive, cognitive, communicative and manipulatory abilities together to create truly autonomous humanoid robots.

Undoubtedly, there are good reasons for embarking on this demanding research journey: there is no platform other than the adult-sized humanoid that is better suited to study many details of our own "being there" in a dynamically changing man-made environment, e.g. through experimentally validating Maturana's understanding [Maturana] of enactive cognition through *structural coupling* both with the environment (i.e. our semi-structured world), with other humanoids and with (a society of) humans. Hence there is also no better platform to study all the different aspects of artificial embodied minds and their development through the interplay of evolving cognitive and motor skills. Most importantly, multidisciplinary research focusing on real autonomy for humanoids, i.e. their capacity to establish and maintain their own identity through self-control and self-guidance, may pave the way to robot systems (not necessarily of human shape) that not only adapt to dynamically changing environments (such as insect-level agents) but also to situations in their interactive discourse with humans for which their designers did not implement explicit rules a priori. In other words: at the end of this development there may well be robots which not only change and enlarge their initial internal set of states (and break their programmed rules) as a result of continuous learning

and behaviour-based plastic adaptation to the environment, but which are also capable of evaluating the incoming multimodal information stream and explaining their emergent internal states to us humans – in our language.

**Potential benefits of humanoid robot research and development.** The main purpose of the first series of fully functional humanoids, Waseda university's "Wabots", was to undertake practical feasibility studies in mechanical engineering. These developments started in the mid-seventies, when the state of computing technology (but also sensors and vision, energy supply, etc.) was still far from what is needed even for a basic notion of "autonomy". Even though at that time one could not even dream of implementing higher-level cognitive abilities as integral functions of these bodies, there were impressive achievements of the emulation of human motor skills (walking, grasping, even piano-playing). Throughout the 1980's there was also a wave of engineering efforts into human inspired limbs, particularly multi-fingered hands, but it faded away when it became clear that there are very few, if any, immediate industrial applications. From today's perspective, however, as complete humanoids are within reach, we can draw on the experience from all these designs and attempts. We observe a growing enthusiasm about the usefulness of humanoids research because we can now see more clearly what potential benefits, i.e. direct applications but also spin-offs, might be:

- *Service:* Unlike autonomous service robots that perform a more or less limited range of special tasks with or without human supervision, a humanoid robot can in principle use the same tools and appliances as humans and may hence become as flexible in adapting to new tasks as a human being. On condition that it is close enough to human shape and size, it may also operate in totally unchanged man-made environments. Moreover, if it is capable of receiving its tasks by carrying on a dialogue with human instructors involving speech, gestures and facial expressions, then it will provide a functionality that surpasses by far anything that today's service robots have to offer. Completely new markets may evolve. Of all the items in this list, this is undeniably the one that requires most multidisciplinary research efforts before we can even think of building prototypes.
- *Prosthetics:* If we think of the humanoid robot as a collection of prostheses for limbs but to some extent also for sensors, then it becomes clear that prosthetics and humanoid research may very fruitfully profit from each other. While there is still little evidence that "cyborgs" may ever be realised or the human mind be transferred to these machines, prostheses that afford some autonomy of their own may become an alternative to current designs, at least until it is possible to "re-grow" human organs. From this point of view it makes sense to subsume humanoid robotics under "biomedical robotics" or even "medical robotics" in the wide sense. The latter is especially true if we

also take into account the potential of humanoids to become an invaluable help for those elderly people who need permanent home care.

- *Education:* Basically there are two different uses for humanoids in education. (i) Students *build* humanoids to learn in a practical exercise about their mechanical construction and the complex software modules that control it. There are quite a few universities in Japan where this is common practice. To spread such education opportunities around the world, it would be highly desirable to have a standard kit available that contains the essential building blocks in modular form. (ii) Students *use* humanoids to experiment with and enhance their skills. The aim should be to make them very easy to use, to clearly specify their interface so as to enable non-roboticists and even students of non-engineering faculties to quickly become familiar with the robot.
- *Entertainment:* robots of human shape used for animation and advertisements at exhibitions and fun-fairs do not depend on a highly developed set of skills. It is usually rather their bodily appearance that attracts people because they discover human traits in these machines. To maintain a certain “surprise-factor” over time, however, it will be necessary to constantly improve their skills. Depending upon the target application, this may even include grasping and sophisticated navigation, e.g. for showing visitors around, manipulating and explaining the objects on display in a natural way. At the other end of the spectrum, small-sized humanoids may well play the role of toys for children that exhibit a greater wealth of behaviours and may hence be much more interesting for children to interact with than artificial animals.

While products resulting from development in these areas may well succeed in traditional markets and create new ones, the greatest challenge from a research perspective is the use of humanoids as subjects in *Cognitive sciences*. Without doubt there is no other machine conceivable on which we may simulate more realistically the development of cognitive processes in developmental psychology, linguistics, etc. – emulating perception and action in the same world in which human beings grow up. The humanoid robot's body, if equipped with a rich set of human-like sensors, generates a stream of multimodal and multidimensional information about the environment that very closely resembles the input to the human perception system. The “motor side” also requires the control of actuators in an extremely high dimensional workspace to act in the real three dimensional world in real time, similar to what the human nervous system has to control. Research need not be limited to study individual development; one can also imagine the study of “inter-humanoid” relations in humanoid societies or the evolution of collective intelligence in swarms of humanoids.

Interestingly enough, direct military uses are currently hard to imagine: not only because fine-grained delicate manipulation and flexible adaptation to three-dimensional environments are

normally not at the heart of military operations, but in particular because the increasing specialisation of military tasks does not necessitate a general purpose robot modelled on the human soldier (at least not in the battlefield). This does not mean, however, that there is no transfer of technology conceivable: think of the current efforts of building force-amplifying exoskeletons that will require energy sources which humanoids design may well capitalise on.

**Scientific challenges.** Some aspects of humanoids design and control, such as various key components, basic walking, gross body motion, and active vision have been studied extensively and successfully. Methods, architectures and software are available that enable humanoids to perform basic actions, which easily make unbiased people think that these machines are indeed not far from living creatures. However, the central question is true autonomy, i.e. the capacity of self-guidance and self-control. This, in turn, implies necessarily that the designer of the humanoid accepts that the robot evolves over time, increasing the “distance” between states, patterns (of sensor input, of behaviour and of motion) and representations that were originally programmed, and add new ones – at some point in future time it may even have the capacity to change its body shape. From our point of view there are four categories of interdependent research categories that must be pursued in the quest for full autonomy:

**(a) Brains and mind:** The central shortcoming of traditional robot systems, adaptive and capable of learning or not, is their closed set of structurally different behaviours and skills. Their parameters may be adjustable and their behaviours may have the potential to be combined and blended, but they will nevertheless always remain slaves to their designers’ foresight. This will only change once we understand how the humanoid can construct its own individual epistemology. Firstly, this construction process calls for the definition of basic instincts (and primary goals) which, in combination with correspondingly selected behaviours and sensory abilities, allow for this individual development, e.g. through active exploration. Secondly, an appropriate set of constraints and a choice of limiting conditions must keep the internal state of the robot within consistent boundaries during the self-directed process of development. Thirdly, and most importantly, an architecture is needed that embeds and networks these initial skills through the definition of an “information flow” and channeling through selection. This is much like the structure of living creatures at birth: though we all start from a heap of cells, the shape and function of our nervous system is formed before our birth, i.e. before we first come in contact with the world in which we develop our individual skills by training our brain (regardless of the interdependence of this structuring process and prenatal stimuli). From cognitive science and brain research we have a clear indication that the adult human brain implements two dif-

ferent categories of information processing [Christaller]: fast parallel “wired” processing and slower, more adaptive serial processing on different time-scales. The line between the two cannot be drawn sharply and it is also possible that skills deliberately and consciously acquired through learning in “serial mode”, if trained over and over again, may fall below the level of consciousness – driving a car and swiftly switching gears without being aware of the underlying complex motion sequence is one of the prominent examples, playing chess and after a while being able to almost instantly assess the situation on the chess board is a less well-known example [Dreyfus]. Humanoids’ control architectures that are capable of evolving based on only very few built-in dispositions will almost certainly have to be modelled on such distinctions. Furthermore, for controlling the various reflexes but also the high-level skills necessary for social interaction and learning, i.e. the coupling of sensing, behaviour and reasoning (situated perception-action cycle), these architectures will have to contain ingredients from the behaviour based approach. There will be a point in the development, however, when we shall see the return of “old AI” techniques, particularly the notion of representations and the reasoning about them along with some kind of introspection and “explainability”, e.g. in the sense of the ability of the robot to explain why it took what action. The control of the co-evolution of cognitive and manipulative skills by continuous evaluation of sensorimotor feedback stimuli is another challenge. Here, meditating entities, i.e. representations for categories and concepts, must be carved out through continuous interaction with the environment (and other members of a humanoids “society” or human beings). Again, the problem will be the self-directed coordination between action-selection, corresponding sensor interpretation and temporary/persistent representation generation (in a predefined format or not). While there have been some attempts to generate categories autonomously based on sensory stimuli, the subsequent step of successfully self-building a tailored non-trivial reasoning and introspection system on top of these has not yet been demonstrated. We tend to hypothesize, however, that the development of higher cognitive skills, like planning and the evolution of a common ontology as well as a language (and gestures) between members of a society, becomes pretty easy once the essence of categorisation (and grounding of patterns in these categories) is cleared up.<sup>1</sup> Of course, some scientists (e.g., Chomsky) believe that language is genetically coded. Finally, the problems of introspection, reasoning about the robot’s own state in terms of lower-level representations for obtaining some kind of consciousness and anticipation of the behaviour of other members of the society along with the organisation of “bootstrapping effects” in mind development as well as

---

<sup>1</sup> Note the analogy between this hypothesis and R. Brooks’ statement about the „essence of being there“ as well as the reasons he gave for that in his 1987 paper about „Intelligence without reason“. Arguably, the time it takes infants to categorise objects is much longer relative to the time it takes them from there to start speaking – then things speed up considerably.

fundamentally new learning techniques (learning by discovery, by insight, by producing goal-directed random ideas at a conceptual level) will pose questions to researchers far beyond the present scope of humanoids research.

**(b) Interaction with Humans and the Environment.** For humanoid robots to continuously extend their range of skills and transcend their initial state of mind and body, it is necessary to implement both sides of structural coupling in as sophisticated a manner as that found in humans. The ability to interact with humans is highly dependent on the expressive power communicable by the body (face expressions, body language) and the ability of the mind to control and interpret multimodal output and input. We note that traditional systems for “teaching by demonstration” or skill transfer have not met with much success. We identify three main reasons for this failure: (i) Instruction input is monomodal, mostly through a fixed camera. This precludes the system from constructing cross-modal associations by evaluating clues from more than one modality. It also prevents the instructor from giving additional explanations in “natural” modalities, e.g. teaching movements of the hand supplemented by instructive speech statements. (ii) Partly due to monomodality the instruction is not in the form of a dialogue between the instructor and the robot. Dialogic interaction may be the source of additional information in “normal” instruction mode, but it becomes indispensable in the case of error conditions. (iii) Due to input from only one location (i.e. the aforementioned fixed camera) there are no redundant multi-dimensional views available that could be used to emulate to some degree the body and head motion that humans use mostly unconsciously to deal with ambiguous situations or occlusions. Moreover, there seems to be no adequate approach to evaluate the instruction stream in terms of its usefulness for learning. In other words: which parts of the input stream (e.g. hand motions) are innovative with respect to the abilities acquired up to a certain point? Redundancy would help to differentiate between elements in the instruction flow that contribute to the learning goal and those that do not. Humanoid robots, if mobile and equipped with a full set of sensors for all modalities needed for being taught a certain class of tasks (e.g. vision, speech and touch), are perfectly suited to learning through demonstration but to also demonstrating during the learning process what and how well they have learned it, which in turn enables the instructor to repeat critical phases – exactly in the way we would teach a human child. The highly abstract but powerful metaphor for this kind of intelligence development is *imitation*; it encompasses both motor and cognitive aspects [Schaal]. This metaphor subsumes a whole set of individual learning approaches. It holds the potential to acquire structurally new sensor/action sequences, which have in no way been pre-programmed. Implementing them on a humanoid in such a way that the human counterpart takes the humanoid seriously requires that a plethora of individual techniques be available to the robot controller (in the form of predefined

and/or evolving skills): (i) recognition, production and integration of natural language/speech, gestures, facial expressions, (ii) expression of desires, intentions and emotions, (iii) keeping the focus of attention across modalities, (iv) learning to re-recognise specific situations, (v) fault tolerance obtained by multisensor perception and compensation, and (vi) adaptation to individual human habits, culturally different customs and idiosyncrasies.

**(c) Structure and Purpose of Body and Limb Movement.** As argued above, the motion of the humanoid in many dimensions is an indispensable ingredient of learning and the emergence of even primitive forms of “intelligence”. Not only is this plasticity the precondition for continuous structural coupling with the world around us, but it is also the basis for social and physical interaction with human beings or with other humanoids for cooperative tasks (with or without parallel explicit instruction/communication over “channels”). This challenge is the point where humanoid design requires a departure from the methodology used for obtaining insect-level intelligence: not only are motion sequences much more complex but they are also deliberate in the sense that they serve different purposes under different tasks. Identical motion patterns may be triggered by reflexes but also willfully. There is ample evidence that there is a strong link between human motor skill and cognitive development (e.g. [Lakoff]). Our abilities of emulation [Clark], mental modelling and planning of motion are central to human intelligence and, by the way, a precondition for anticipation, but they also critically depend on the experience we have with our own body dynamics as we plastically adapt our body’s shape to the environment. Apart from this direct link to/from the humanoid mind, there are many practical research fields that have seen dramatic advances in the recent past (e.g. bipedal walking with anthropomorphic legs) but that still require detailed study before we come even close to the performance observed in biological systems: (i) whole body gross and fine motion (and its learning by observation and demonstration), (ii) body-centred behaviours (limb coordination) and posture stability, (iii) learning of arm motion, grasping and sensorimotor control for delicate manipulation tasks including adaptation/generalisation to new tasks, e.g. starting from new (unknown) situations of body and/or object configuration, (iv) navigation, planning of movement and collision avoidance under dynamically changing boundary conditions (v), locomotion, gait and foot-placement, and (vi) simulation of body and limb dynamics and the use of such simulators not only for the designer but also as a “mental tool” for the humanoid.

**(d) Architecture and System Design.** The recent discussion about the overall nature of cognitive systems (dynamical systems vs. symbol interpreters [Gelder]) notwithstanding, the architecture of the control system, i.e. the collection of centralised and decentralised control structures for component coordination, must host the abstract functions that en-

able development. While the physical implementation (organic vs. anorganic) and the representation of information flowing through the humanoid may well change over time, the architecture will have to afford mechanisms to structure the “primordial soup” of programmed initial behaviours, skills and instincts. Most likely, the architecture will contain low-level elements for high throughput processing whose adaptation ends after some time with only a minimal degree of plasticity remaining and other elements that warrant long-term adaptivity – the former without explicit representations and the latter relying on them. The main question will be how they can be tied together in such a way that a reasonable convergence is obtained after some time of development. This intriguing problem of architecture layout is inextricably interwoven with the achievements in mechanical and sensor design. Progress achieved here is immediately communicable even to non-experts: (i) mechatronics of hands, feet, legs, arms, heads; new types of actuators including materials for actuators and skin, (ii) sensor design (articulated vision, tactile, directional ears, new sensing principles), and (iii) evolutionary hardware.

**Scanning the issue.** We have a collection of papers that spans the whole gamut of humanoid robot research. One part of this set is in a special issue of “Robotics and Autonomous Systems” (RAS) published by Elsevier Science, the other part is published in this special issue of “Autonomous Robots” (AR), ordered by subject category. This includes the robot structure and physical capabilities, whose properties are studied both through experimental prototypes, as well as simulation systems. Hashimoto et al. (AR) present two fully operational humanoids developed at Waseda university, where much of the pioneering research in the field was done. “Hadaly”, which is a torso on wheels, was developed with human-humanoid interaction in mind; it integrates vision and voice recognition with gesture generation. “Wabian”, by contrast, can walk on legs and was built as a prototype of a humanoid that can carry objects in a household setting. Furuta et al. (RAS) describe the design and construction of compact humanoid robots. Bipedal systems are developed with walking control in which energy efficiency, hardware load balancing, and real-time gait generation are all handled. The paper by Nakamura et al. (RAS) demonstrates a humanoid robot simulator that can emulate the dynamics of motion of various robot structures based on virtual links. Their system can handle bipedal motion, standing, and collision avoidance. The virtual humanoid robot platform can be the virtual counterpart for a robot platform. Kuffner et al. (AR) address similar issues: they develop a new algorithm for automatically generating motions from full-body posture goals under the boundary conditions of dynamical stability and collision avoidance. The algorithm was implemented on a dynamical model of the H5 humanoid robot; results of sample runs are presented in the paper. Kagami et al. (AR) introduce a trajectory generation method for humanoids using the standard ZMP approach. Their fast algorithm is



based on a simplified robot model, it allows for the generation of trajectories even if several limbs are in contact with the environment, and it has been tested on the H5 humanoid.

The following papers consider various approaches to transferring human behaviors and skills to humanoid robots. Ude et al. (RAS) describe a real-time visual system that can learn and interact with humans. The system uses shape and color in conjunction with a probabilistic tracker. This allows the system to avoid being brittle and it is reliable in a complex environment. Stoica (RAS) also proposes a method to transfer skills to the robot by close interaction with a human. This requires anthropomorphic robots, and involves the visual tracking of arm motions and their mapping onto the robot structures. This approach is arguably more efficient and sets the stage for the study of developmental robotics. Fod et al. (AR) propose a method for extracting and classifying the movements of a human arm for imitation learning. They start with a repertoire of movement primitives that an agent can execute, observe the movements of a human arm, cluster the observed data using principal components analysis to associate them with elementary movements taken from the repertoire and can then closely reproduce the original sequence of movements. Storing and retrieving motion patterns is also in the focus of the paper by Riley et al. (AR). They describe a dynamical systems approach for realizing ball catching on the Sarcos humanoid robot, which predicts the ball-hand impact and generates human-like motion trajectories for catching. Billard et al. (RAS) develop a model of human imitation, in this case of two-arm motions. Their approach is biologically inspired, and is based on a hierarchy of artificial neurons. They have a 37-DOF humanoid robot, and human data is used. Both human learning and robot learning data are compared. Vijayakumar et al. (AR) demonstrate how learning in high-dimensional parameter spaces can be performed in real-time based on an infinite stream of input data. The goal is to learn the inverse kinematics of a 7-DOF anthropomorphic robot arm for a specific task. As an example task the arm is to draw an "eight" in space, and using their locally weighted projection regression algorithm, it takes only one minute to learn the local kinematics model. In the paper by Cheng et al. (RAS), human-robot interaction is studied in the rich context of stereo vision, auditory systems and proprioception. This is coupled with high-performance motor control. The system features parallel sensory channel perception analysis and integrated system response.

The remaining papers consider the cognitive development of humanoid robots from very different standpoints. Scasselatti (AR) takes a relatively abstract view. He proposes the development of a "theory of mind" for humanoid robots based on the fundamental social skills for humans: the attribution of beliefs, goals, and desires to other people. His paper briefly summarises some of the theories of the development of theory of mind in human children and discusses the potential application of these theories to building robots with similar capabilities. Asada et al. (RAS) propose the study of cognitive developmental robotics as a way to

emphasize the interaction between the embedded structure in the robot and the environment. This includes learning and development of the robot brain, and a key insight is the necessity to create social environments conducive to robot learning. Coelho et al. (RAS) propose advanced cognitive development through modification of the robot's active response due to perception-action experience. Their insightful proposal to reduce the tremendous complexity in the learning space of a humanoid robot, is to use nature's physical response to robot actions in order to filter behavioral tactics and strategies. Breazeal et al. (AR) explore the potential of using human speech for communicating with humanoid robots and for teaching them. They present an approach for recognizing four prosodic patterns that communicate praise, prohibition, attention, and comfort. This perceptual ability is integrated into their robot "Kismet", which makes it possible for a human instructor to directly manipulate the robot's affective state. The paper by Giszter et al. (RAS) proposes the basis for biomimetic control architectures. They study the mammalian system in terms of its ability to achieve skill formation and motor learning, and they have developed interfaces with motor centers in the brain to directly examine cortical motor learning using a neurobiotic interface to rats and (partially to) monkeys.

**Conclusions and the future.** Being aware of the many forecasts about the future development of "intelligent" machines starting with the early computer pioneers, we shall not try to predict what humanoids might be in ten or twenty years time, exponential growth of mankind's knowledge admitted or not. However, it may be useful to recall some limitations that will have to be overcome before humanoid robotics is accepted as an important field within the community and researchers acknowledge that robotics as a whole, possibly along with other disciplines, really profit from investments of time and capital into humanoids. (i) A scenario, or "benchmark", is needed that makes research results comparable with respect to all relevant skills of the humanoid. It should also show that the task classes involved in it cannot be handled in a more effective and efficient manner by special purpose robots. Ideally, some "competitive component" should be included. As proven by the extremely successful RoboCup, this would spur new results by constantly re-focusing to a (preferably simple) goal. There have been suggestions for "Humanocup" but we think that scenarios involving complex three-dimensional environments are more convincing. Examples would be cooking in a standard kitchen (including an understanding of the recipe), (sub)tasks of home care, e.g. changing bed sheets, or participation in disaster recovery/rescue, etc. (ii) From an engineering point of view a reference or standard platform (at an affordable price) would be highly desirable to have (not only for student education as mentioned above). Based on such a platform an international infrastructure for humanoids research may build up (with shared use of simulation software, of components or full robots, etc.). Only when this is the case will we

see a growing community become interested in humanoids. This also presupposes that more attention is paid to critical components with an integral powerful energy source being the most important. Inexpensive components will lead to a wealth and diversity of implementations. A distributed reliable power supply modelled on the chemical reactions found in the bodies of living creatures with their extremely high efficiency is definitely a long time away – it would also require a complete redesign of actuators. A cable connection of the humanoid with the wall plug is not an option. A mid-term alternative might be fuel cells or small high-speed electrical generators powered by hydrogen. Here we may profit from progress made in the automotive industry. (iii) For obvious reasons humanoids will never be an exact replica of human intelligence. It therefore remains to be investigated systematically what kind of intelligence we may expect. Can we prove that it is doomed to remain on the level of, say, whales, i.e. animals to which we attribute intelligence and whose utterances we can associate with our own emotions, although they would probably not be able to tell us anything of interest even if we could communicate with them? Or will humanoids eventually talk to us on the grounds of an identity of their own? To explore this with some plausibility, a concerted effort by philosophers, cognitive scientists and roboticists is needed, which along the way may also contribute new insight into the different theories of mind-body relations that were developed over the last 2000 years, i.e. the philosophical implications of artificial embodied minds. For the past lack of this dialogue (with few exceptions) a definitive answer remains yet to be given. This is also true for the somewhat esoteric issues of self-repair and self-reproduction together with “artificial embryology”.

Are humanoids the ultimate goal in robotics and will they put an end to its development? Most certainly not. Nevertheless, whatever our individual view of the promise of this field: it can be considered to be beyond all doubt that humanoid robotics is a fascinating multidisciplinary area of research which will be intensely studied for the foreseeable future.

## References

[Brooks] R.A. Brooks and L.A. Stein. Building Brains for Bodies. *Autonomous Robots*, Vol. 1 (1), Nov. 1994, pp. 7-25

[Capek] Karel Capek: RUR - Rossum's Universal Robots : kolektivní drama o vstupní komedii a 3 dejstvích. Praha : Aventinum, 90 pages, 1920

[Christaller] Th. Christaller and C. Kolo and E. Pöppel. Bioinformation. Problemlösungen für die Wissensgesellschaft. In: *Technik, Wirtschaft und Politik*, Bd. 38, Oct. 1999, Heidelberg, Physica-Verlag

[Clark] A. Clark and R. Grush: Towards a Cognitive Robotics. *Adaptive Behavior*, Vol. 7(1)1999, pp. 5-16

[Dreyfus] H. L. Dreyfus. The Current Relevance of Merleau-Ponty's Phenomenology of Embodiment. <http://hci.stanford.edu/cs378/reading/dreyfus-embodiment.htm>

[Gelder] T. van Gelder. The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* 21(5), 1998

[Lakoff] G. Lakoff. *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. University of Chicago Press, 1990

[Maturana] H. R. Maturana and F. J. Varela. *The Tree of Knowledge: The Biological Roots of Human Understanding*. Shambhala Publications, 1992

[Schaal] S. Schaal. Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 3(6)1999, pp.233-242.