

Summary: “The AI Lectures from Tokyo”, 2 February 2004

Rolf Pfeifer, Britta Glatzeder, and Yasuo Kuniyoshi

This summary is to provide a synopsis of the most important aspects of embodiment. As such, it serves the goals of the lectures set out initially. The first goal was to provide a comprehensive introduction to the concept of embodiment and its implications for our understanding of intelligence. The second goal was to demonstrate that matters can always be seen differently, even if one feels that this is impossible. The third goal was to give some insights into recent research in the field and to familiarize the audience with achievements and challenges in the field.

Let us first look at how our perception of intelligence has changed during the lectures, if at all. Let us start with a concrete question that I am sure many have been asking themselves. While in classical artificial intelligence, much time is devoted to so-called “high-level” activities such as theorem proving, logical reasoning, problem solving, and chess, throughout the lectures we have been looking relatively often at sensory-motor tasks such as moving around, obstacle avoidance, wall-following, block-pushing, running dogs, and walking, and we have been studying a lot of robotic experiments. Now what does walking have to do with what we normally think of as intelligent behavior? Or more generally speaking, how do so-called “low level” sensory-motor processes relate to more “high-level” cognitive ones? This answer is, of course, at the heart of any theory of embodiment.

One of the important goals of artificial intelligence research is, in addition to understanding natural forms of intelligence, the development of an abstract theory. We have made a start in this direction by proposing a set of design principles that capture our new – embodied – understanding. The principles have been separated into two parts, the design procedure principles or “meta principles” and the actual agent design principles (for an overview, see Table 1).

Table 1: Overview of the design principles.

Label	Name	Description
		<i>Methodological principles</i>
M-Princ. 1	Synthetic methodology	Understanding by building
M-Princ 2	Emergence	Systems should be designed for emergence (for increased adaptivity)
M-Princ 3	Diversity-compliance	Tradeoff between exploiting the givens and generating diversity solved in interesting ways
M-Princ 4	Time perspectives	Three perspectives required: “Here and now”, ontogenetic, phylogenetic
M-Princ 5	Frame-of-reference	Three aspects must be distinguished: perspective, behavior vs. mechanisms, complexity
		<i>Agent design principles</i>
A-Princ 1	Three constituents	ecological niche (environment), tasks, and agent must always be taken into account
A-Princ 2	Complete agent	Embodied, autonomous, self-sufficient, situated agents are of interest
A-Princ 3	Parallel, loosely coupled processes	Parallel, asynchronous, partly autonomous processes, largely coupled through interaction with environment
A-Princ 4	Sensory-motor coordination	Behavior sensory-motor coordinated with respect to target; self-generated sensory stimulation
A-Princ 5	Cheap design	Exploitation of niche and interaction; parsimony
A-Princ 6	Redundancy	Partial overlap of functionality based on different physical processes
A-Princ 7	Ecological balance	Balance in complexity of sensory, motor, and neural systems: task distribution between morphology, materials, and control
A-Princ 8	Value	Driving forces; developmental mechanisms; self-organization

In what follows we will use some of the principles to find an answer to the above question, i.e. how sensory-motor tasks relate to what we intuitively term as “thinking”. The term “thinking” has proved controversial as it is often taken to imply conscious thought, and, as pointed out earlier, we prefer to use the broader term “cognition”. But this is not essential, here.

Let us now take the example of grasping to demonstrate how the body and cognition, or more generally, neural processing interact. Natural arms are built of muscles, tendons, ligaments, and bones, materials that are non-rigid to varying degrees. All of these materials have their own intrinsic properties such as mass, stiffness, elasticity, viscosity, temporal characteristics, damping, and contraction ratio to mention but a few. These properties are all exploited in interesting ways in natural systems. For

example, there is a natural position for a human arm which is determined by its anatomy and by these properties. Reaching for and grasping an object like a cup with the right hand is normally done with the palm facing left, but could also be done – with considerable additional effort – the other way around. Assume now that the palm of your right hand is facing right and you let go. Your arm will immediately turn back into its natural position. This is not achieved by neural control but by the properties of the muscle-tendon system: in a sense the system acts like a spring – the more you stretch it, the more force you have to apply and if you let go the spring returns to its resting position. Also, the human arm exhibits intrinsic damping. Normally reaching equilibrium position and damping is conceived of in terms of electronic (or neural) control, whereas in this case, this is achieved (mostly) through the material properties. Or put differently, the morphology (the anatomy), and the materials provide physical constraints that make the control problem much easier – at least for the standard kinds of movements. The main task of the brain, if you like, is to set the material properties of the muscles, the spring constants. Once these constraints are given, the control task is much simpler.

This represents an illustration of the principle of ecological balance, as there is a kind of task distribution between the morphology (the anatomy), the neural system (the control), the materials, and the environment (in this case the environment plays a role as, for example, gravity acts on the arm which also provides some control in terms of “moving” the arm back into its natural position). Exploiting material properties, morphology and interaction with the environment are aspects of the principle of cheap design. Note that even highly complex systems exploit physical and environmental properties, even though they incorporate a lot of redundancy. Materials in themselves have, so-to-speak computational properties or control properties, as they “take over” some of the tasks of the neural system.

But let us take this example one step further to get towards the core of our question. Let us, for the sake of the argument, assume that a baby receives random neural stimulation. The resulting movements of the arms will be far from random: typically the arm will move towards the front of the body. If we further assume that there is a grasp reflex in the hands, then, if the hand by chance encounters an object, it will grasp it (which does not require much control because of the particular morphology of the hand), and as a consequence this object will be brought into the range of the visual field. The implications are as follows. Through the grasping, sensory stimulation will be generated at the finger tips of the baby where humans have very dense and highly sensitive touch sensors. Because the object is brought into the

visual field, there is now correlated sensory stimulation concerning the object in both the visual and the haptic (touch) channel. And this sensory stimulation is correlated and at least temporarily stable. In other words, this is “good” raw material for the neural system to process and also forms the basis for cross-modal learning. Cross-modal learning, i.e. the association of sensory stimulation from different sensory channels (e.g. the visual and the haptic) is at the foundation of concept formation and categorization. Categorization is probably the most fundamental cognitive ability: you have to be able to make distinctions in the real world, otherwise you are not going to survive for very long: distinctions between food and non-food, conspecifics and members of other species, dangerous and benign objects, etc., etc. must be made. Moreover, babies have a tendency to stick objects into their mouths which on the one hand generates additional sensory stimulation from the lips and the tongue (where we also have dense and highly sensitive sensors, and we have the taste buds in addition) and on the other the objects thus positioned always have the same physical distance from the eyes. This way, correlated sensory stimulation is generated in various sensory channels. Note that this generation of data is achieved through the physical interaction with the environment: this is a physical process, not a computational one! And this physical process has strong implication for the neural processing required. So, we are beginning to see the true implications of embodiment. The sensory-motor coordination, the physical interaction with the environment, leads to correlated sensory stimulation which is the basis for learning and categorization. And categorization is the basis of cognition and concept formation, which in turn is at the foundation of our language abilities. In other words, the principle of sensory-motor coordination provides, in some sense, a “link” between the embodied interaction with the environment and higher-level cognitive processes.

Coming back to our question of the relation between walking and thinking we can now imagine that through motion, sensory stimulation is induced, i.e. actively generated, and the specific patterns of sensory stimulation are determined by our body and the specific interactions with the environment. These patterns of sensory stimulation in turn form the basis or the “raw material” for the neural system to form cross-modal associations and categories.

Remember now our discussions of symbol grounding early on in the lectures. Although this is clearly not the final solution yet, you can now begin to understand how symbol grounding might be achieved, i.e. how the body might provide a foundation for category formation and later concept formation, and it is precisely the categories and concepts for which language provides the symbols. The details of how this

actually occurs in humans during ontogenetic development are still unknown, but this seems a highly plausible and we could imagine that we can devise neural network-based mechanisms to achieve this task. George Lakoff and Rafael Nunez in their book “Where mathematics comes from” (Lakoff and Nunez, 2000), argue along similar lines that even highly abstract mathematical concepts such as the imaginary number “i” or the notion of a “power set” may have their origin in the very specific way of how the human body is built.

We are also beginning to see that the distinction between high-level cognitive and low-level sensory-motor starts blurring: categorization is highly sensory-motor, but is also fundamental to cognition, even higher level cognition. So, perhaps, walking and thinking are much more closely related than one would think at first sight. This is further corroborated by the idea from the famous neuroscientists Daniel Wolpert that the evolutionary pressure on the development of the brain has come from need to move and orient. “Plants don’t have brains”, so to speak.

To conclude our illustrations of the design principles, let us briefly look at a few examples that most people would associate with higher levels of cognition. One is conversation. Recall the “social robot” Kismet. It is equipped with a number of basic reflexes for turning towards loud noises, tracking slowly moving objects, habituation, and avoidance (if something comes too close). This leads, as we have seen, to behavior that awfully looks like true social behavior. The question that immediately comes up is: What about human social behavior? Could we imagine that it might also be much more driven by simple reflexes rather than rational thought? As we know from David MacFarland, the famous Oxford University ethologist (animal behavior), rational behavior need not be based on rational thought. This kind of behavior reminds us of the principle of parallel, loosely coupled processes in which the processes run asynchronously, in parallel, and are often only coordinated through the interaction with the environment, in this case, the social interaction.

Conversation is another activity that most of us associate with very high-level cognitive processes. Simon Garrod, a psychologist at the University of Glasgow who has been studying conversation for many years, asks the question why most of us find it difficult to give a speech (monologue), but find conversation (dialogue) very easy, even though, intuitively one would think the opposite to be the case. He proposes that there might be some rather reflex-like processes of “alignment”. He terms this the “interactive alignment account of dialogue processing.” (Garrod, 2003). John Bargh, a psychologist at Yale University, in a recent paper on the foundation of human behavior, argues in a paper entitled “The unbearable automaticity of being”

(Bargh and Chartrand, 1999) that much more that we would like to believe might be driven by simple reflexes and not be under conscious control. Again, we see the distinction between high-level and low-level processes become increasingly blurred. By the way, Garrod and Bargh, both seem to propose something like a “cheap design” principle for social behavior.

So, some distinctions that seem very clear at first sight become blurred which implies that the distinctions are perhaps not as obvious as they seem if we look at the underlying processes. Questioning whether distinctions, taxonomies, and conceptual hierarchies that people propose really hold, is necessary but also adds uncertainty. But this uncertainty can result from a deeper understanding of intelligence. As the famous physicist and Nobel Laureate Wolfgang Pauli said: “Whenever I have understood something, I am confused – but on a higher level.” But we not only add to the confusion, we do have something instead, we have a set of design principles that can help us tremendously understand natural systems, but also build artificial ones.

There are many phenomena that we have not (yet!) explained, such as abstract thinking (e.g. chess), natural language, creativity, and consciousness. But we with the synthetic approach we have enormously powerful approach that, in a bottom-up way, may eventually lead the way to a deeper understanding of these issues. And at the theoretical level we have made a good start with the set of design principles.

It has been our experience that many people find these insights rather disconcerting because they go against our traditional way of thinking in which we like to view ourselves as having our actions firmly under conscious control. Closer inspection, however, yields that there are enormous benefits to being – largely – controlled by reflexes: the reflexes are there and they function by themselves and we don’t have to worry about them. And it they seem to work very well in regulating social behavior. A great relief, indeed!

Before we close this summary, let us briefly review some of the presentations we have experienced in the sessions “the latest from”. Yasuo Kuniyoshi, Toshihiko Nakamura, and Masafumi Okada of the University of Tokyo pursue a dynamical systems approach to intelligence where they adopt a complete agent perspective and study the dynamics of the system-environment interaction. They try to connect this continuous dynamics to the higher-level symbolic activities bi-directionally – bottom-up and top-down – by exploiting the intrinsic transition behavior of the dynamics. The discrete attractor states in a dynamical system can naturally be associated with symbols and memory. This way, it is possible to “ground” the higher-level symbolic

activities in this dynamics. Masayuki Inaba, Hirochika Inoue of Tokyo University, and Minoru Asada of Osaka University, have a strong interest in humanoid robotics where, ultimately, they are interested in human assistance and social interaction. Inoue and Inaba have the vision of assisting the elderly in order to give them as much independence as possible for as long as possible by integrating real task performance (i.e. the robot performing a task for the human) and human interaction functionalities (where the robots supports the human while interacting with him or her). Inaba has also been interested in human-like anatomy and has designed and built the robot “Kenta” (tendon boy) that is equipped with a flexible spine – modeled after the human spine – and many artificial tendons; a very challenging control problems. Shigeru Hirose of Tokyo Institute of Technology is interested in applying biological principles to the construction of robots, and a similar goal is pursued by Rolf Pfeifer and his research group in Zurich. It is important to note that the goal is not to produce a copy of nature, but to understand its principles, find abstractions, and then applying these principles to the design of technology. This is precisely how we proceed when applying the synthetic methodology. Isao Shimoyama has also been inspired by natural systems and has built especially small systems mimicking particular morphologies of insect eyes. He has also made great inroads into MEMS technology (Micro-Electric Mechanical Systems), technology for building very small devices that often combine actuation and sensing, an idea akin to the principle of sensory-motor coordination. Friedrich Pfeiffer, Technical University of Munich, and Rolf Pfeifer both have the goal to understand rapid locomotion and to build pertinent devices. While Friedrich Pfeiffer is focusing on the control, Rolf Pfeifer is applying principles of embodiment, in particular cheap design and ecological balance to design walking and running robots. Hiroshi Yokoi of Hokkaido University has been building artificial hands, closely modeled according to principles known about the function of human hands. His goal is to develop hands that can be attached to humans who have lost a hand, controlled by means of EMG signals (or potentially other means of neural interfacing). He has explicitly been applying, for example, the principles of ecological balance to his research. Rodney Brooks, of MIT CSAIL (Computer Science and Artificial Intelligence Laboratory) has recently started investigating the “foundations of life”, research that is still at the very beginning, but promises exciting results in the future. In his robotics research his current focus is on vision and manipulation, two areas in which current robots have not been very successful yet. For this reason, Hong-bin Zha of Beijing University has been investigating robot vision for many years and developed many pertinent applications, e.g. in the area of security. Visionary Satoshi Murata of Tokyo Institute of Technology

is looking at self-assembly and self-repair, mechanisms that are fundamental for any kind of living systems. This falls under the general topic area of self-organization, a frequent research topic in artificial life that has more recently been increasingly applied in the field of artificial intelligence as well. One intention is to ultimately replace much of what we now do with “control”, by “self-organization”. Albrecht Schmidt of the University of Munich introduced intelligent environments and argued convincingly that we not only want to put sensors into environments but also motor components, another examples that reminds us of the principle of sensory-motor coordination. He also demonstrated Finally, Luc Steels of the Free University of Brussels and Sony Computer Science Laboratory in Paris, pointed out the importance of social interaction aspects in emergent behavior, an issue that a number of speakers had also been stressing as well and to which definitely a lot of research efforts have to be devoted.

The authors

Rolf Pfeifer is a physicist, psychologist, and computer scientist

He is professor of information technology and director of the Artificial Intelligence Laboratory of the University of Zurich
Lecturer: The AI Lectures from Tokyo

Britta Glatzeder is a philosopher

She is a senior research associate at the AI Lab of the University of Zurich and a member of the Parmenides Foundation for the Study of Thinking, Elba
Project manager: The AI Lectures from Tokyo

Yasuo Kuniyoshi is an engineer

He is professor of information science and technology at the University of Tokyo and co-director of the Laboratory for Intelligent Systems and Informatics
Host and lecturer: The AI Lectures from Tokyo