Understanding by Design – the synthetic approach to intelligence

Daniel Bisig*+ and Rolf Pfeifer*

* *Dept. of Informatics, University of Zurich, 8050 Zurich, Switzerland*

+*Inst. for Computer Music and Sound Technology, University of the Art Zurich, 8048 Zurich, Switzerland*

Designing and building systems for selected abilities is a core activity of research in artificial intelligence. One of the attractive points of proceeding in this manner is that we not only end up with an actual system, but along the way, we learn a lot about the phenomenon of study, e.g. spontaneous structure formation or rapid locomotion. One of the crucial questions concerns the level at which the system should be designed. In this paper, we introduce the synthetic approach to intelligence. This approach promotes design and development as a research methodology. But in contrast to classical engineering, "Understanding by Design" promotes a bottom-up design stance based on "design for emergence". These ideas are illustrated with a number of case studies from robotics and human-computer interaction.

Conceptual Background and Motivation

"Understanding by Design" is the name of a particular research methodology that is increasingly employed in different kinds of sciences. This methodology forms the main foundation of embodied AI whose research and design concepts are summarized in the text book «Understanding Intelligence» (Pfeifer and Scheier 1999, 2001) and in the more popular science book «How the Body Shapes the Way We Think – A New View of Intelligence» (Pfeifer and Bongard 2007).

Very briefly, "Understanding by Design" is based on the assumption that at least some natural phenomena can best be understood by building artifacts that embody a selected set of these properties. The process of building the artifact and its subsequent exposure to experimentation represents a tangible form of conceptualization that leads to testable hypotheses. In addition, this approach offers the benefit that it often entails engineering innovations that possess general application potentials beyond the particular phenomenon under investigation.

The methodology of "Understanding by Design" combines in itself two types of ideas, the synthetic approach and design for emergence. The synthetic approach introduces engineering practice into scientific research. Correspondingly, a particular phenomenon of interest (e.g. how do we recognize a face in a crowd, or how do we move and walk) is approached from an implementation perspective. Design for emergence tries to minimize designer bias and the pre-definition of the artifact's resultant properties. This represents a novel approach uncommon both to engineering and science.

The synthetic approach originated from within the framework of pre-cybernetic robotics. The term "synthetic method" was employed by psychologist Kenneth Craik to describe the process of testing behavioral theories through machine models [Craik, 1943]. The synthetic approach is not meant as a replacement but rather a complement for the traditional analytical approach. The analytical sciences are very well established and have contributed immensely to the increase of mankind's knowledge about and control over the natural world. Nevertheless, starting from about the second half of the twentieth century, science began to take notice of an increasing number of natural phenomena that seem to notoriously resist clarification. By now it is clear that these phenomena share some fundamental properties that largely defy an analytical approach. These phenomena are typically based on a large number of constituents that operate in parallel and whose interactions with each other and their surroundings must be described by non-linear mathematical relationships. Also, we have become aware of the fact that these phenomena are not anomalies or exceptions, but rather represent the vast majority of systems of interest. Examples abound and can be found in the purely physical world (e.g. climate, star formation, the creation of snow flakes), biological organisms and societies (brains, gene regulation, body movement, swarm behavior, spread of diseases), and man made social, technical and socio-technical systems (cellular automata such as Conway's "Game of Life", the Internet, the stock market, cities, fashion trends). These phenomena are the result of the individual properties of the constituents (the neurons, the "cells" in a cellular automaton, or the humans in a fashion-trend network), and of the complex interaction patterns among these constituents. By building artifacts bottom up from components to compound aggregates to whole systems, the synthetic sciences can study the properties of a whole system and how these properties depend on the interrelationships and behaviors of its components.

Design for emergence emphasizes the relationship between a system's high level and low level properties even stronger than the synthetic methodology does. Emergence comes in different guises (Pfeifer and Bongard, 2007): (i) The behavior of an individual is always emergent in the sense that it is the result of the individual's interaction with the real world. Accordingly, behavior cannot be fully specified by some internal control mechanism but depends on the morphological, material, and environmental conditions as well. (ii) Organisms (or artifacts in engineering) are said to emerge from an evolutionary process. If this process follows principles of natural evolution, then all of the organisms capabilities are emergent properties of a process that is driven solely by the organisms' reproductive success. (iii) A collection of parts or a population of agents can display emergence based on local rules of interaction. For example, proteins assemble into a viral body architecture or a group of birds self-organizes into a swarm. The overall behavior of such systems is not simply the sum of the individual behaviors of its constituent parts but is both quantitatively and qualitatively entirely different. For example insect societies can build sophisticated hive structures despite the limited cognitive capabilities of each single insect. The fact that the human brain can exhibit consciousness is another impressive example. This capability is obviously not present in the brain's individual neurons.

It is impossible to define an analytic causality between local and global properties of an emergent system because they are mutually contingent upon each other: global structure and behaviors are the result of the interaction of the individual components, and the individual components are in turn influenced by the global behavioral patterns. It seems that many of the impressive capabilities of natural systems such as their adaptivity, robustness, and capability for self-repair and reproduction are the result of emergent processes. The import of these insights for science and engineering can hardly be over-estimated. In order for science to study principles of living and intelligent systems, it has to move its focus (at least partially) away from the level at which these phenomena manifest themselves – the global behavior pattern – to lower levels (i.e. material and morphological properties and physical principles) even though they might, at first glance, be unrelated to the topic of interest. Because the analytical treatment of emergent phenomena has proven very hard, scientists have to adopt an empirical position that relies on observation, experimentation and trial-and-error. This means, that in order to understand how changes in emergent system affect its behavior, a scientist has to devise experimental setups and scenarios within which the system can evolve over time. It is within such a setup, that the scientist can test hypotheses about what a system's behavior depends on. For example it can be tested how the shape of swarm depends on the sensory capabilities of the individual agents. Or it can be studied how the quality and diversity of evolutionary adaptations changes in response to differing levels of selection pressure. For these type of experiments, a "trial and error" type of approach is justified by the fact that many of these relationship cannot be predicted but are stumbled upon coincidentally. "Trial and error" can serve as a strategy to sample a vast range of potential interdependencies. In addition, it helps to overcome a scientist's preconception and bias.

With regard to engineering, it is clear that many of the previously mentioned capabilities of natural systems would be desirable for artificial systems as well, such as autonomy, robustness, and the ability to deal with unexpected situations (e.g. robots that interact with humans, vehicles for extraterrestrial exploration). Because of the principles of emergence, it is obviously hard if not impossible to design such artifacts by following a classical engineering methodology (i.e. to move top down from pre-specified high level requirements to low level implementation details). Furthermore, an engineer's expertise and intuition may not only be of very limited utility, but it might prove to be an obstacle to achieve the desired artifact properties. For example, methods from classical control engineering – while extremely efficient in a highly controlled industrial environment – do not work well when applied to building robots that have to deal with the real world. The real world is dynamically changing and only partly predictable and agents therein have to act in the face of uncertainty. For this reason, the design for emergence approach also tries to minimize the human designer's bias and preconceptions. The usefulness and innovative potentials of this approach for engineering has already proven itself. By applying design principles extracted from biology and exploiting ideas from natural evolution, artifacts have been created with minimal human intervention that surpass the capabilities of their human-designed counterparts (Koza et al, 2000).

When designing systems within the context of the synthetic methodology, a scientist usually has to decide between a simulation or robotic's based implementation. Due to the fact that a purely simulation based approach usually benefits from lower costs and time investments as compared to a robotic approach, it might seem hard to justify the development of robots. Nevertheless, robotics forms the main cornerstone of artificial intelligence research. This is based on the fact that behavior as an emergent phenomenon critically depends on physically realistic agent environment interactions. It is very hard to achieve physically realistic behavior in simulation. The indefinite richness of the real world is the main driving force that give rise to the diversity of morphologies and behaviors we witness in biological organisms. No computer simulation is capable of mimicking this richness. A simulation possesses only those properties that have been consciously and deliberately added to it. For this reason, any simulation based experiment is hampered by the absence of properties that have not been implemented and that might turn out to be crucial for an emergent phenomenon the manifest. With that said, it is clear that simulations remain an indispensable tool for scientific research. Experiments that involve a large number of interacting agents can usually be realized in simulation only. The same holds true for evolutionary runs that involve co-adaptations of morphology and control. For this reason, robotic and simulation based experimentation complement each other and will remain the two main methods in synthetic methodology.

Finally, it is important to note that the blending of methodologies from science and engineering has significantly expanded the explanatory power of the sciences. Traditionally, the natural sciences have been pre-occupied with the study of naturally existing systems. The synthetic sciences on the other hand that capitalize on the “Understanding by Design” approach have expanded their focus of inquiry. These sciences no longer develop artifacts as part of their research process in order to imitate natural systems. Quite on the contrary, many of the resulting robots and simulations don’t represent any organisms that exist in nature. For this reason, the synthetic sciences are no longer limited to study natural systems. To paraphrase Christopher Langton, one of the founding fathers of the field of Artificial Life, the synthetic sciences not only study “life as it is” but also “life as it could be”. Accordingly, the “understanding by design” approach allows scientists to gain a more profound understanding about the fundamental aspects of life and intelligence than would be possible if their research focus was limited to natural systems only. The “understanding by design” approach has combined the endeavors of natural science and engineering into a venture that aspires to elucidate the very meaning of being alive and intelligent. As such it has trespassed into the territory of philosophy.

Subsequent to this general and somewhat theoretical introduction of the “Understanding by Design” methodology, the remainder of this article will focus on the presentation of several case studies with the goal to concretize the application of this methodology in the context of embodied intelligence. These examples have been selected in order to highlight the methodology’s impact for a broad range of projects and activities. The examples include projects in basic research, educational activities and art.

Case Studies

Robotic Hand

The robotic hand project is used to investigate the relationship between morphology, intrinsic body dynamics, the generation of information structure through sensori-motor coordinated activity, and learning (Gómez 2007; Gómez et al. 2006; Gómez et al. 2005). The so called Yokoi hand is partially built from elastic, flexible and deformable materials (Figure 1). It’s actuation is based on a muscle tendon system that is inspired by the anatomy of the natural human hand. Furthermore, the robotic hand mimics some of the sensory capabilities of the biological original. Each finger is equipped with sensors for measuring bending, rotation and pressure. Additional pressure sensors are present on the hand's palm and back. In order to control the robotic hand, biologically inspired learning mechanisms have been implemented. These mechanisms enable the hand to explore its own movement capabilities. By correlating the hand's sensory input as a result of its motor outputs, it can also learn to manipulate and grasp objects. The robotic hand project has diversified its initial research focus and is currently moving towards an application as prosthetics device. For prosthetics, it is essential that the utilization of the hand "feels" natural. For grasping, information structure must be induced which implies that there must be rich sensory feedback. Experiments with fMRI show that patients that are provided with even minimal, but correlated, sensory feedback (such as electrical stimulation to the skin or mechanical vibration), integrate their prosthesis into their body schema much faster (Hernandez-Arieta et al., 2008). This project showcases that the synthetic methodology and the notions of morphology and information self-structuring benefit working systems and can lead to spin-offs such as assistive prosthetic devices. In particular, they give indications as to how to augment the sensory-motor “intelligence” of a coupled man-machine system.

Figure 1

Robotic Self-Assembly

The goal of this project is to achieve self-assembly and self-repair in a self-organized robotic system consisting of many modules (Miyashita 2007; Miyashita et al. 2007). Self-assembly is a process through which an organized structure spontaneously forms from simple parts. Despite the fact that this crucial process is ubiquitous in nature, little is known about the mechanisms underlying self-assembly and not much effort has been devoted to abstract higher level design principles. Taking inspiration from biological examples of self-assembly, we designed and built a series of modular robotic systems consisting of cm sized autonomous plastic tiles capable of aggregation on the surface of water. A single module, named Tribolon consists basically of a foamed rubber shape and a small vibrator (Figure 2). Power for the vibrator is provided via an antenna which touches the aluminum ceiling, connected to a power supply. Magnets attached to the modules let them attract or repel each other. The vibration causes the modules to move, and depending on the shape of the modules, different behaviors emerge, such as clustering and rotating. By drawing from our current experience in designing, constructing and controlling macroscopic modular systems, we hope that we will be able to derive conclusions about the level of autonomy that is needed to achieve self-assembly. Our synthetic approach follows the biological principle that components self-construct themselves into organisms in a completely bottom-up fashion.

Figure 2

Educational Robotics

The DREAM (Development of a Robot kit for Education, Art, and More) project aims at developing an educational robot kit consisting of hardware, software and instructional material which aids in communicating conceptual and methodological principles of embodied Artificial Intelligence to people within and outside of the AI community. This kit is intended to promote creative and integrative approaches through a «constructionist» educational approach (Resnick 1996) where students learn by developing and building artifacts. The dream project builds on the AILab's numerous educational activities (Figure 3). For example we are participating in the Roberta network (Petersen et al. 2003, Müllerburg et al. 2002), which aims at developing a robotic education program that specifically targets the interests of girls. The AILab also teaches courses as part of the «bugnplay.ch» competition. This competition deals with art, media and technology. It is organized by the MIGROS Kulturprozent and targets children from 11–20 years. These activities are based on our conviction, that the concepts of embodied AI research are not only relevant to the scientific community within this specific field but have far-reaching implications for scientific research and engineering in general. Furthermore, we have observed that the synthetic methodology is well suited in an education context in that it helps to spawn and maintain a high level of motivation in students and can act as a teaching methodology to communicate even very abstract concepts in a comprehensible and tangible way (Paul et al 2000). Contrary to existing educational robotic kits, the DREAM kit intends to promote principles and methodologies from embodied AI research that are central to the "Understanding by Design" approach. Accordingly, the kit will emphasize bottom up engineering and design for emergence. We believe that these concepts encourage the search for new types of problem solving strategies that are useful and inspiring for anybody who is involved in problem solving and decision making in our highly dynamic, complex and only partially predictable world.

Figure 3

Robotic Art Installation

The project HairMotion is an ongoing collaboration between one of the authors and artist Valerie Bugmann. The project aims to realize an interactive robotic installation that serves as an experimental environment for non-verbal communication. The robots will be stationary and their only means of expression is via the breathing movements of their artificial lungs and their pneumatically actuated hair (Figure 4). The robots are setup in a circular arrangement and face towards an inner space that can be entered by a visitor. Within this inner space, the robots are capable of perceiving the presence and movements of the visitor via a vision based tracking system. The interaction between the robots and the visitor progresses through various stages. Throughout these stages, the robots' relation to the visitor varies, their degree of autonomy and reactivity changes and correspondingly the movements of the visitor and robots transition through periods of synchronization, correlation and independency. The interactive scenario on the installation encourages visitors to explore the relation between their own behavior and those of the robots. The visitors engage the robots into a dialogue of movements in order to identify recurring elements of interaction. Accordingly, the robotic installation creates an experimentation space that allows visitors to create and test syntactic and semantic features of a movement based non-verbal language.

Figure 4

References

Craik, K. J. W. (1943). The Nature of Explanation. Cambridge University Press, Cambridge, 1943, England.

Gómez, G. (2007). Adaptive Learning Mechanisms for Autonomous Robots. PhD Dissertation, Faculty of Mathematics and Science, University of Zurich, Switzerland.

Gómez, G., Hernandez, A. and Eggenberger H. P. (2006). An adaptive neural controller for a tendon driven robotic hand. Proceedings of the 9th International Conference on Intelligent Autonomous Systems. IOS Press, Tokyo, Japan, pp. 298–307.

Gómez, G., Lungarella, M. and Tarapore, D. (2005). Information-theoretic approach to embodied category learning. Proceedings of the 10th International Symposium on Artificial Life and Robotics (AROB 10), Beppu, Japan, pp. 332–337.

Hernandez-Arieta, A., Dermitzakis, K., Damian, D., Lungarella, M. And Pfeifer, R. (in press). Sensory-motor coupling in rehabilitation robotics. To appear in Service Robotics.

Koza, J.R., Keane, M. A. , Yu, J., Bennett, F. H., and Mydlowec, W. (2000) Automatic Creation of Human-Competitive Programs and Controllers by Means of Genetic Programming, Genetic Programming and Evolvable Machines, Kluwer Academic Publishers, Hingham, USA.

Miyashita, S., Hadorn, M., and Eggenberger H. P. (2007) Self-Assemble of Water Floating Active Units. IEEE International Conference on Mechatronics, Harbin, China.

Miyashita, S., Kessler, M., and Lungarella, M. (2008) How Morphology Affects Self-Assembly in a Stochastic Modular Robot. IEEE International Conference on Robotics and Automation (ICRA08), Pasadena, USA.

Paul, C., Hafner, V., Bongard, J. C. (2000) Teaching New Artificial Intelligence using Constructionist Edutainment Robots. GMD Workshop on Edutainment Robots, Sankt Augustin, Germany.

Pfeifer, R. and Bongard, J. (2007) How the Body Shapes the Way We Think – A New View of Intelligence. MIT Press, Cambridge MA.

Pfeifer, R. and Iida, F. (2005) Morphological Computation: Connecting Body, Brain and Environment. Japanese Scientific Monthly, 58(2), pp. 48–54.

Pfeifer, R. and Scheier, F. (1999, 2001) Understanding Intelligence. MIT Press, Cambridge MA.

Resnick, M. (1996) Distributed Constructionism. Proceedings of the International Conference of the Learning Sciences, Northwestern University, pp. 280–284.

Ziegler, M., Iida, F., Pfeifer, R. (2005) «Cheap» Underwater Locomotion: Morphological Properties and Behavioral Diversity. IROS05 Workshop on Morphology, Control and Passive Dynamics.

Ziegler, M., Iida, F., Pfeifer, R. (2006) «Cheap» Underwater Locomotion: Roles of Morphological Properties and Behavioural Diversity. Proceedings of Climbing and Walking Robots (CLAWAR 06).

Figures

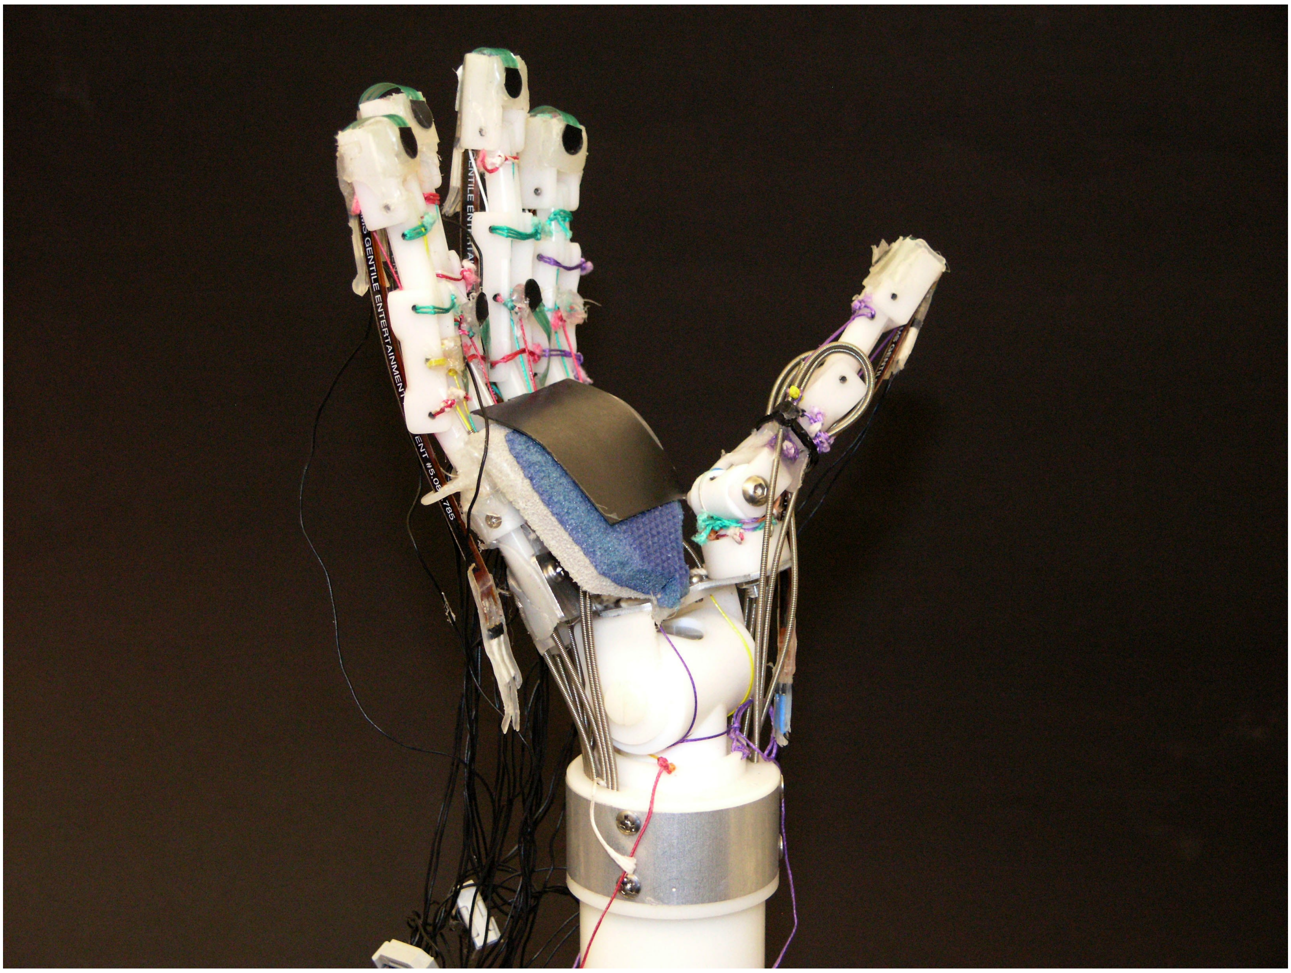

Figure 1: Yokoi Robotic Hand

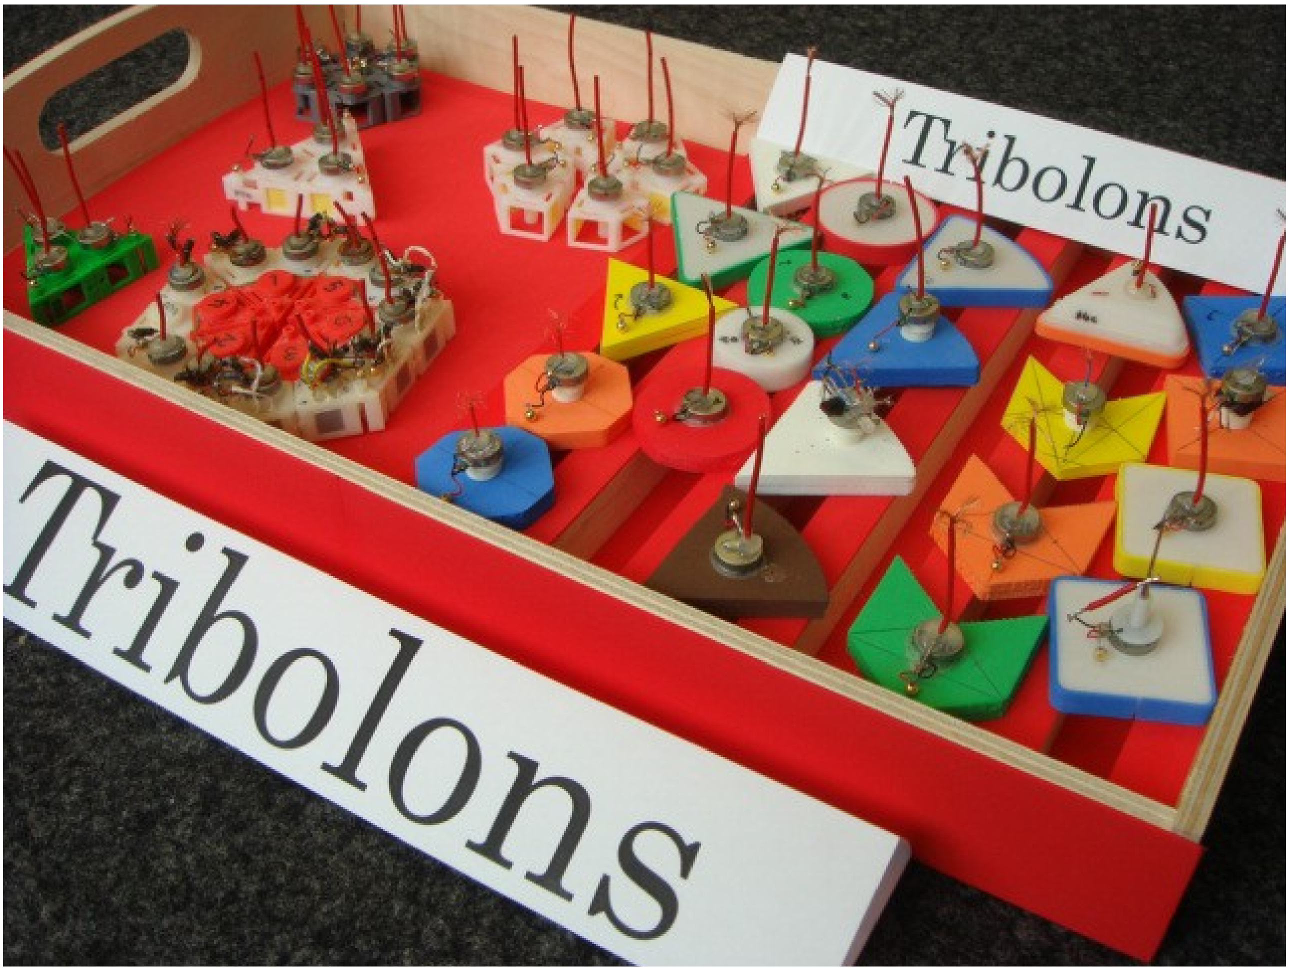

Figure 2: Examples of Tribolon Modules for Self-Assembly

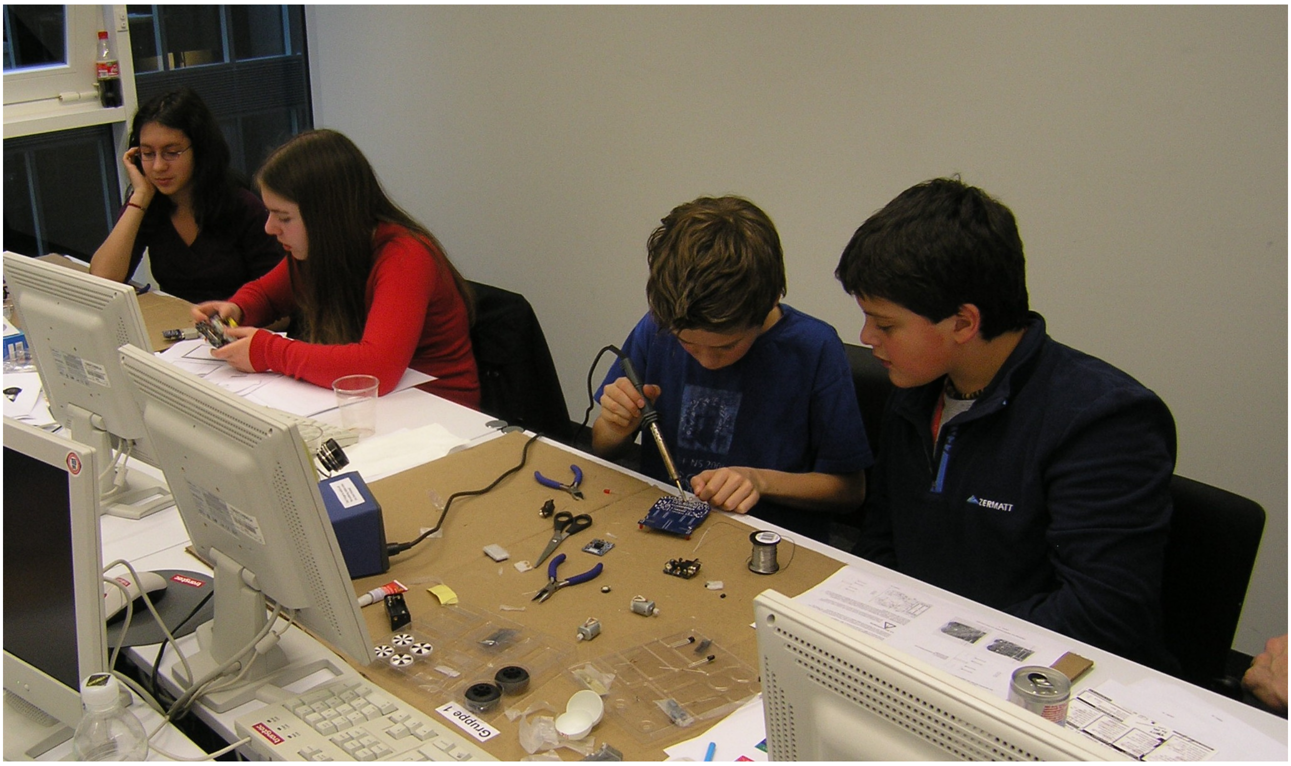

Figure 3: AILab Open Day Robot Building Workshop.

Figure 4: Rendering of a HairMotion Robot.