Thematic Group 5

Intelligent and Cognitive Systems

Report for Public Consultation

Thematic Group Coordinators and Editors
Rolf Pfeifer & Alois Knoll

Version 3.1

January 31, 2006

Contents

1 INTRODUCTION & MOTIVATION

The overall goal proposed here is to construct physically instantiated systems that can perceive, understand, and interact with their environment – but also *evolve* in order to achieve human-like performance in activities requiring context-specific knowledge.

This is far beyond the current state of the art and will remain so for many years to come. Therefore, we propose to focus research efforts on a small set of strategic challenges required to make headway towards this vision. The research dedicated to achieving solutions in this area will result in significant scientific and technological advances. The strategic challenges are motivated by recent research in robotics, neuroscience and cognitive sciences, and can be listed as follows: (i) the role of *growth and development* in cognitive systems, (ii) conceptual understanding of the *action/perception coupling of embodied agents* – with the environment, with humans and with their peers, (iii) the role of *morphology* – how can the body perform processes that are essential to the agent's successful operation, (iv) *research into materials* and what they can do for an agent, and (v) development of a *theoretical framework*.

The purpose of this document is to delineate the research challenges that flesh out certain aspects of this grand vision and map them into more concrete tasks, given the present and emerging technological advances. This builds on (and continues) research pursued in earlier national and EU-level programs, such as Neuro-Informatics for "Living Artifacts", Beyond Robotics, Cognitive Systems, Bio-ICT, and Neuro-IT.

It is interesting to observe that although – according to Moore's law – the computing power has been doubling roughly every one to two years, the intelligence of artifacts in the real world has not even remotely paralleled this development. The reason for this discrepancy is to be found in the fact that the traditional equation "intelligence = computation" is obviously wrong. Traditionally, in neuroscience, psychology, artificial intelligence and robotics, intelligence has been conceptualized as the product of sophisticated control processes within a cognitive or neural architecture. As a consequence, attempts to create intelligent systems have concentrated on the design of computational models of such architectures, whose internal processes solve hard perceptual or cognitive problems and encapsulate knowledge about the world. While this computational approach has been successful at modeling the function of some specific cognitive modules or brain regions, most of the models address only one aspect of the problem of intelligence, the part that "is going on inside the head." Although this is changing, however, in the neurosciences there has been much less of an interest in viewing intelligence as an emergent property of the neural and behavioral activity of the entire organism.

The research topics proposed here are extremely challenging and will require large efforts by many research groups in order to achieve significant progress. In line with recent conceptual and technological developments the research challenges also include a theoretical framework which, rather, than capitalizing on the idea of computation only, could be based on the notions of **embodiment**, the dynamical systems metaphor, complete agents rather than individual components, self-**reconfiguration and self-repair**, **morphology and materials**, and **growth and development**. Progress in the theoretical underpinnings of embodied intelligence will undoubtedly have strong technological implications in areas including robotics, actuator technology, materials, self-assembling systems.

2 VISION AND GRAND CHALLENGES

We propose to tackle the challenge by developing so-called *complete* agents, i.e., agents that are

- embodied & self-sufficient – they can sustain themselves over extended periods of time;
- situated – they can acquire information about the environment through their own multi-modal sensory systems and act accordingly; and
- autonomous – they function independently of external control.

The development of such an agent would constitute a major breakthrough and would, given the current state of the art, obviously require many years to develop. In this framework, intelligence and cognition are properties that emerge as an agent interacts with its environment and encompass many individual skills (problem solving, decision making, reasoning, dialogue with humans, etc.).

By the term **embodiment** we refer to the fact that intelligence requires a body, and all the implications of this assumption. One of the core research topics is how intelligent behavior emerges from the interaction of brain, body, and environment. In recent years, a large community of researchers has begun to realize the importance of a body (and brain-body interaction) for understanding intelligence and its central role in a wide range of processes including perception, object manipulation, movement, and locomotion.

By the term **environment** we refer to the physical and social environment of the agent. It is dynamic and *includes humans and other agents* that can be embodied or reside in ambient networks.

We start by briefly outlining the set of research themes. These are the most salient ones we identified, although they do not form an exhaustive set of research topics into intelligent and cognitive systems:

- *Mind-body co-development and co-evolution:* In order to maximally exploit the design power of evolution and development, controllers and robot morphologies have to evolve simultaneously. The permanent interaction of the body of an agent with the environment during growth enables its “mind” to develop. We need to understand and model the developmental processes in nature and build embodied artificial systems that are inspired by them; embed them into evolutionary processes, and study interaction of physical and information processes during this development. These also require physical growth (see the following point).
- *Systems and materials that can physically grow:* Recent research strongly suggests that physical instantiation and materials play an essential role in behavior. Also through growth, biological organisms can form highly complex morphological structures. In this respect there are promising starting points (e.g. self-assembling materials, modular robotics). Although it is not expected for growable materials that match biological capacities to become available any time soon, some aspects of physical growth can be studied through advanced principles of autonomous modular systems that optimally adapt to a task and the environment.

- *Morphological computation:* Shape and the materials (e.g. muscles) perform important functions for an agent in real time, an idea that that bypasses the concept of classical Turing computation architecture. Morphological computation here refers to processes based on shape (e.g. molecules/DNA, modules of a modular robot) and material properties (e.g. of the muscle-tendon system). The challenge will be to explicitly apply morphological concepts in theoretical and practical explorations of embodied artificial systems.
- *Design for emergence:* Behavior is always the result of the interaction of an agent with the environment, i.e. behavior is emergent, meaning that it cannot be understood (and designed) on the basis of the internal control program (or “brain”) only. The question then is: how can we design purposive (goal-directed) agents without destroying the emergent nature of their behavior?

In each of these topic areas, real physical embodiment plays an essential role. However, given the current state of the art, simulation work (“embodied agent simulations”) will form an important part of the endeavour.

The challenges outlined above will contribute to the development of a unified framework of intelligent and cognitive systems, i.e. a shared theoretical understanding or *epistemology*, within which the developments take place and which provides the context for its evaluation. Given that there is currently no such universally shared framework in the field of intelligent and cognitive systems, it needs elaboration and is hence part of the research program proposed: What is the epistemology that is needed to cope with future developments and challenges in the area of intelligent and cognitive systems? What kind of framework is considered productive within which to conduct research and technological development? The results can be viewed as the first steps toward a theory of intelligence and cognitive systems. It is important that this theory be grounded in experimental work and the other way around, in an evolving “dialectic” relationship.

3 PROPOSALS FOR RESEARCH THEMES

In the sequel we describe the research themes that will help to move towards the vision outlined in the previous section:

3.1 Mind-body co-development and co-evolution

Rationale. In order to achieve the strategic goal of autonomous agents we need concepts, tools, and methods to design true complexity. One of the reasons why we do not have cognitive systems with a high level of intelligence concerns the fact with direct design at the “here and now” time scale, we are encountering a complexity barrier. What will be required are automated or semi-automated design methods. True complexity must emerge from development in the interaction with a physical and social world. Methods of artificial evolution that incorporate models of ontogenetic development are highly promising, as they reflect principles of biological evolution and adaptation. The standard approach in evolutionary robotics has been and still is to use an existing robot platform and evolve the (neural) controller for it. This does not exploit the potential of embodiment for intelligent behavior: mind and body must *co-evolve*. Given the state of technology, it makes sense to differentiate for embodied artifacts between (i) the development of cognitive skills in a fixed body (which is what is currently done; no real breakthroughs are expected from this approach), (ii) the development of artifacts with some morphological adaptation capacity (e.g. “growable sensors/actuators”); and (iii) artifacts with real mind-body co-evolution.

Bottlenecks today. There is a deficit in knowledge (from neuroscience and biology) about mechanisms of ontogenetic development. In particular, we currently lack powerful models and tools for simulating genetic regulatory networks and developmental plasticity, i.e., how the development of organisms is influenced by the interaction with the real world. In addition, growth in artificial systems is not well understood (see 3.2 below) and it is currently not clear how an artificial evolutionary process can be connected to the real world, although there have been some preliminary attempts. This can possibly be achieved through the instantiation of a particular interaction dynamics with the environment or with another agent. A major bottleneck concerns the lack of good dynamic models and physically realistic simulation environments for testing phenotypes. On the hardware side, we lack growing materials, adaptable structures, sophisticated sensory and motor systems, in particular artificial skin, as well as dexterous and flexible and robust actuators (including artificial muscles, deformable tissues, etc.).

Expected benefits (impact). In addition to contributing to the design of such intelligent agents, this research and development will deepen our insights into the mechanisms of phylogenetic and ontogenetic development and the interaction of the two, and how to exploit them for design.

Main research objectives.

- understanding and modelling of ontogenetic processes
- embedding ontogenesis into an evolutionary process (“evo-devo”)
- interaction of physical and information processes during ontogenetic development

Research focus. While the long-term goal is to have true mind-body co-evolution, in the intermediate term, systems with various degrees of “coupling” should be targeted, with a focus on how to extend these systems: mimicking development by initially constraining abilities of a sophisticated robot, and including ideas from modular, self-assembling robotics, computational and behavioral neuroscience, and embodied agent simulation models.

Demonstrator. An embodied agent simulation coupled to the real world: a process of morphogenesis based on genetic and/or neural regulatory networks embedded into an evolutionary scheme with some aspects of the process taking place in the real world. The simulation for testing the phenotypes in a physically realistic environment should be partly realized through an interaction with the real physical world.

Why now? In the areas of artificial life, developmental/evolutionary robotics and behavioral neuroscience the importance of embodiment has been widely recognized. Moreover, developmental biology has achieved a significant level of understanding of genetic regulatory networks, and a number of models have been developed. Also, in the material sciences, growing materials in the sense of self-assembly have been demonstrated. Given this “landscape”, mind-body co-evolution promises to create scientific and technological synergies.

3.2 Systems and materials that can physically grow

Rationale. In order to achieve the strategic goal of autonomous agents we need concepts, tools, and methods to design true complexity. There are a number of reasons why this goal has to date not been achieved. One reason is that design is typically done at the "here and now" scale, whereas biological systems are "designed" at the evolutionary and developmental scales. It is only through growth and adaptation that biological organisms can form highly complex morphological structures starting from a single cell, and it is the permanent dynamical interaction of their body with the physical environment during this growth process which enables the different levels of their "minds" to develop. Harnessing the power of biological development and evolution requires the ability to grow materials in the real world. Although we are not expected to have growable materials that match biological capacities available any time soon, physical growth can be studied through advanced principles of modular systems that optimally adapt to a task and the environment.

Bottlenecks today. The bottlenecks are both conceptual and technological. At the conceptual level, it is the lack of a sufficient understanding of evolution and how it interacts with ontogenetic development at a level that could be exploited for design. While in recent years there has been enormous progress in biology and the neurosciences in understanding developmental processes, in particular on genetic regulatory networks, pertinent models are still largely missing (although there have been some preliminary promising attempts). At the technological level – and this is the most serious limitation that has prevented growth from being exploited for design – it is simply the lack of growable materials that could be employed to mimic developmental processes.

Expected benefits (impact). Artifacts that co-evolve their (possibly distributed) brains (control systems) and their body in permanent interaction with the environment over an extended period of their lifetime (embodied artificial ontogenesis) would exhibit a completely new level of intelligence, adaptivity, and intractability. Moreover, growable materials can be harnessed for systems capable of self-repair which would then enhance the potential of deploying agents in remote and hazardous sites. At the conceptual level, understanding growth and finding ways to handle the accompanying uncertainty will fundamentally change the way we envision potential applications, as growable materials have not yet entered into people's thinking in research and development.

Demonstrator. A demonstration of a robot whose material properties adapt, depending on external influences from the environment. This process should not be merely passive (such as the deformation of tissue on the finger tips in grasping); instead, there should be a local feedback loop regulating the microscopic properties of the materials. An example from nature is the growth of the eye, which is regulated through neural and biochemical circuits based on sensory input while interacting with the environment. This means that adaptation processes could be electronically or chemically controlled. The next level of complexity would be to have a growing system based on self-assembly, i.e., where the pre-manufactured building blocks can be taken from a repertoire. Concretely, this could be a kind of "bone" tissue that grows thicker in places where there is larger stress. First prototypes of such materials are already available.

Main research objectives. The main objective is, ultimately, the development of growable materials. While there is a long way to go there are a number of promising starting points: (i) simulation of artificial morphogenesis using genetic regulatory networks (non-trivial genotype-phenotype mapping); (ii) models of morphogenesis embedded into artificial evolution; (iii) modular robots (macro, micro, and nano-scale self-assembly); (iv) self-assembled materials; (v) developmental/epigenetic robotics; and (vi) electronically controlled microfluidic arrays to control chemical reactions very precisely, e.g. for activation and deactivation of real neural networks' activities.

Research focus. From today's point of view, we see four essential threads of technology research (as opposed to the indispensable conceptual lines of work mentioned above) that should form the basis for an integrated research plan: molecular robotics, distributed growable sensors, growable distributed information processing, and growable motor entities and spatially distributed actuators. If successful, these efforts might help overcome one of the most obvious and most difficult challenges: the limitations of current sensing and actuator technologies. Ideally, it will be possible to formulate – at an appropriate level of abstraction – principles which govern the growth processes in the artifact, such as the mechanisms underlying genetic regulatory networks or the physical and chemical processes underlying self-assembly and how they can be controlled. In parallel, the development of convincing application scenarios should be advocated. This not only pertains to useful deployment on the factory floor, in private homes, outdoor support, etc., but also involves the transfer of parts of the technology to applications that could profit from, say, micro scale machinery with integrated sensing and information processing abilities for medical use. In the context of this "Intelligent and cognitive systems" initiative, it is important that growth not be viewed as an isolated research topic, but always as part of a complete artifact having to operate in the real world. Once growth technology is in place, it will be possible to have mind-body co-evolution in the real world. We do not expect to have mature "growth technology" fully available by the end of FP7, but some preliminary achievements (including simulation) are certainly realistic.

Why now? Recent developments in the emerging fields of self-assembled materials and nanotechnology, developmental robotics, and artificial evolution and morphogenesis are encouraging enough to justify substantial investment into forming a new discipline by combining key research in these fields to work towards completely new types of growing artifacts. In the field of growing materials, there are interesting trends: the number of publications that have the term "self-assembly" in their title, after an initial stagnation at a low level, has shown a significant increase in recent years (Figure 1). This might be an indication that growth in terms of self-assembling materials might soon be a feasible option in the design of intelligent and cognitive systems.

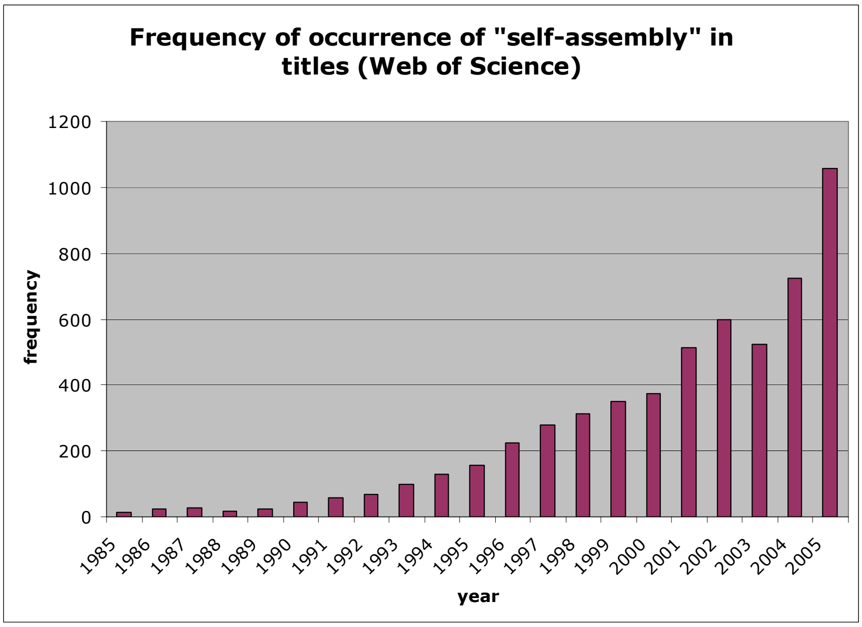

Figure 1: Frequency of occurrence of “self-assembly” in titles (Web of Science).

Thus, even though the technology is only at the beginning, the time seems ripe for a focused research program.

3.3 Morphological Computation

Rationale. Classical artificial intelligence and cognitive science, including cognitive psychology, are founded on the notion of Turing Computation or abstract symbol manipulation. Embodied agents, by their very nature as complex dynamical systems, are fundamentally different from computers: their specific morphologies perform processes essential for the agent’s successful operation. In recent years, computational and behavioral neuroscience have recognized the action/perception coupling as the crucial aspect of human development.

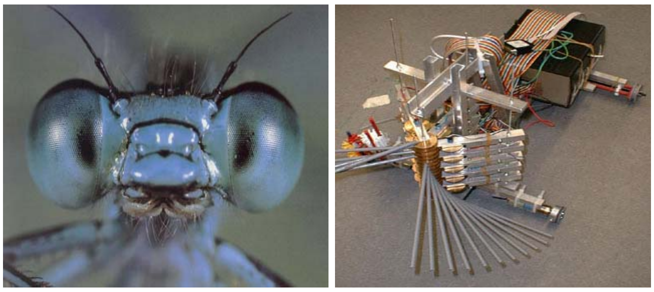

Figure 2: Facet distribution of insect eye and robot system “eyebot” with an artificial compound eye, which can autonomously modify the angular positions of the individual facets.

The term “morphological computation” is used to designate the fact that control not only resides in the brain or a microprocessor, but is also distributed throughout the organism. While there is a certain “trading space” – some functions can potentially be performed by traditional computation or by the morphology – it is not “free trade”, i.e., the different functions are not equivalent. Such processes are performed by morphological and material properties of a system in order to facilitate or to support control tasks. Examples are facet distributions in insect eyes, stiffness properties of muscle-tendon systems, and deformability of tissue on finger tips. Figure 2 shows an insect and a robotic “eye” performing morphological computation. The term, as employed here, encompasses a much broader class of phenomena than classical Turing computation, as it includes aspects of morphology, materials, and system-environment interaction. An essential consideration in “morphological computation” is time. Many processes in the real world, e.g., the reaction to the foot’s impacting the ground in running, requires reaction times that are beyond the computational capacities of the neural system – and also of electronic systems. In other words, the agent has no choice but to delegate some of the control to the materials and the morphology. Morphological computation is always required whenever we are dealing with real physical agents interacting with the real world. A proper understanding of these issues is a prerequisite for the development of truly adaptable artifacts and thus forms an important objective

Bottlenecks today. Because the notion of Turing computation has been and still is so dominant, it is hard to imagine that there might be other kinds. Areas such as DNA or more generally molecular computation have been around for many years, but progress has been slow. Moreover, often in research studying different kinds of computation, there is an attempt to use different substrate and to interpret the result as Turing computation. Another bottleneck concerns the lack of quantitative measures capable of relating morphology, materials, and interaction with the environment. Such measures could help, for example, in answering the question of how much processing – or "computation", so to speak – is taken over by a sensorimotor coordination. A third bottleneck results from the fact that even in robotics, the robot hardware (morphology and materials) and its control are typically designed separately.

Expected benefits (impact). It is clear that Moore's law does not hold for systems that have to interact with the real world, i.e., for intelligent and cognitive systems, because intelligence is not only about computation, at least in the traditional sense. Extending the scope from the virtual world of Turing machines into the real one will open up entirely new design possibilities, conceptually and practically. Because it focuses on all aspects of agent design, this perspective is expected to provide impulses for completely novel types of *morphological* computational architectures, and will thus lead to innovative designs of intelligent and cognitive systems. Moreover, it will define a new research field that will attract scientists from many other areas such as biomechanics, neuroscience, chemistry, physics, material science, and dynamical systems. A concrete example of the potential benefits is recent research on evolvable artificial cells, which will be programmable in the sense that they can, for example, open up to deposit a chemical in a particular environment, which has enormous application potential in medicine. Also, cells could be programmed to form arbitrary structures, which in turn would enable them to perform certain functions. Self-assembly is another area where morphological computation is essential: if all the "work" has to be done by traditional computational methods, as it is done today in most robotics, self-assembly and self-reconfiguration projects, the processes will always require large amounts of computation and they will not scale down to very small units because microprocessors will be needed.

Demonstrator. An artifact demonstrating how the trade-off mechanism for morphology, materials, and interaction with the environment can be applied to the real world. The demonstrator should illustrate the "trading" function and show how the performance of this artifact is improved as compared to the traditional perspective. The notion of "morphological computation" should be illustrated in quantitative terms.

Main research objectives. The notion of "morphological computation", though intuitively plausible, still lacks basic understanding, theoretical rigor, and quantitative underpinning. The main objectives are the elaboration of this concept at the theoretical level, and the development of many case studies illustrating the concept. It is important that explicit consideration is given to morphological computation in the design process.

Research focus. The challenge here is to explicitly apply the idea of morphological computation in theoretical and practical explorations. A proper understanding of these issues is a prerequisite on the way towards truly adaptable artifacts and thus forms an important objective. Concretely, the following aspects might provide specific research questions in this program:

- shape (the shape of the insect eye, the anatomy of the hand-arm-shoulder system, the ear, the whiskers of rodents, complete anatomy for exploiting passive dynamics, etc.); at different scales, the shape of cells and molecules provides functionality to the organism;
- materials (stiffness and elasticity of muscle-tendon system, slack in the joints, ground friction, whiskers, etc.);
- physical nature and distribution of sensors on organism (haptics – skin sensors, visual, laser, acoustic, smell, etc.);

- brain structure (connectivity, plasticity, physical arrangement, length of connections);
- signaling mechanisms in brain and body (hormones, regulatory networks); and
- growth processes (see 2.2) and, more generally, changes in morphology which lead to different functionality of the system; change can be shaped by experience.

Why now? In the design of artifacts having to interact with the real world, researchers have reached critical limitations that can no longer be dealt with using the classical notion of Turing computation only. An illustration is the research on humanoid robots where progress has been stifled by an adherence to traditional computational concepts and it has become clear that novel design principles are required. Also, in a number of research areas – robotics, biomechanics, chemistry, dynamical systems – people are beginning to think along similar lines.

3.4 Design for emergence – purposive behavior

Rationale. The notion of purpose is probably one of the most difficult to deal with, because if we really want to learn something about "purpose", we do not simply want to program it into the system. If we program purpose into the system, we may not learn very much about purpose itself, we just get out of the system what we program into it. Moreover, if we do not allow for emergence (and its intrinsic uncertainty), we will not see any interesting behavior of the system evolving. It will also be difficult to specify purpose – whose categories should this specification be based on: the human designer's/observer's, or the artifact's?

As a consequence, one way or other, we would want to achieve purpose as an emergent phenomenon, i.e., employ mechanisms of emergence to achieve a certain behavior, function, task-orientedness, or ability to interact with a user. We would also like to be able to predict the most probable system architecture that will result from a given initial state. This, of course, means that systems may be forced to take "detours" on the way towards achieving a certain functionality – but at the same time it means that the solution ultimately achieved will most likely be more interesting, and that more can be learned from it.

Design in the real world is always emergent. This means that the behavior of an agent cannot be entirely pre-programmed, but is always the result of an interaction with the real world. Even in classical control theory, if a there is a heavy obstacle in its way, the arm will be stopped. Thus, design must always be "for emergence", it is logically not possible for it to be otherwise. Three kinds of emergence can be distinguished: in an individual agent (at the "here-and-now" time scale), global behavioral patterns emerging from the local interaction of many agents, and emergence from time scales (the ontogenetic or the "here-and-now" emerging from an evolutionary process, or the "here-and-now" emerging from an ontogenetic process). Because design for emergence explicitly takes into account the interaction with the environment (including other agents), the resulting agents tend to be more robust. Therefore, on the way to truly intelligent agents, design for emergence must be understood. The fundamental research question that needs to be answered is how purposive behavior can emerge be achieved without designing the goal into it. It is important to note that emergence is not an all-or-none phenomenon but gradual: certain behaviors are emergent to a higher or lesser extent.

Bottlenecks today. Because of the dominance of classical control theory in robotics and because Cartesian thinking is still widespread in science but also in the population at large, emergence is still considered a mystery.

The conflict that has been holding back progress is that systems are designed for a particular purpose and because emergence is currently still poorly understand, scientists and engineers have failed to date to "design for emergence". The standard way to proceed is still in a top-down manner. The major obstacle is a conceptual one, namely that it is unclear how a particular design goal can be achieved through emergence. While from a perspective of basic research it is sufficient to design a system and observe what kinds of behaviors emerge, from an engineering perspective this is insufficient.

Expected benefits (impact). There will be conceptual and practical benefits. At the conceptual level, understanding how purposive behavior emerges from more basic processes will represent a major theoretical breakthrough and will lead to entirely novel ways of viewing intelligent behavior. In neuroscience it might elucidate how motor signals translate into behavior, and how processes of self-regulation occur. At the practical level it will be possible to break the "complexity barrier" because adaptation to the environment will be taken into account via emergence.

Demonstrator. A methodology that can be quickly applied to a broad class of design problems. This could be instantiated as a set of tools for artificial evolution and morphogenesis. This also requires appropriate application scenarios.

Main research objectives. In this work program there are two major research questions. First, how can goal-directed or purposive behavior come about in biological agents without explicitly postulating goal-hierarchies (which are always highly arbitrary)? And second, given a particular set of tasks or desired purposive behaviors of an artificial agent, how can it be designed without directly programming the agents with goals or purpose (which was and still is the classical perspective)? The reasons for this is that the classical approach requires enormous amounts of computation and has only had limited success in specific types of environments. Moreover, true complexity can only emerge, and direct programming does not yield additional insights.

Research focus. All aspects of emergence should be investigated, in an individual agent ("here-and-now"), emergence from local interactions in groups, and emergence from the different time scales. The latter relates, of course, directly to the issue of mind-body co-evolution. Research in this program should focus on how to achieve goal-directed purposive behavior in agents without explicitly designing the goals into the agents. Specifically, given a set of tasks or desired behaviors the agent should achieve, what are the kinds of design decisions that have to made, and at what level or time scale should they be made? How do the environments have to be structured so that the desired behaviors emerge?

The study of natural evolution and ontogenetic development can provide inspiration. "Evo-devo", i.e., methods of artificial evolution with a non-trivial genotype-phenotype mapping, is promising, as such methods reflect principles of biological evolution. An interesting research topic is open-ended evolution, where the challenge is to achieve desired behaviors while keeping the evolutionary process open-ended. Coupling aspects of the evolutionary process to the real world bears great potential and presents a big challenge (this is part of a different work program).

Design methods are not restricted to evolutionary and developmental ones, i.e., they are not limited to biologically inspired ones, but could, for example, capitalize on mechanisms of morphological computation. For instance, they could exploit mechanisms of self-assembly and re-configuration.

Why now? While the field of adaptive mobile robotics is definitely thriving, there is, among insiders, a definite understanding that current robots do not meet the ultimate goals of intelligent and cognitive systems in most respects. But current robots are mostly designed “by hand” at the “here-and-now” level. So, there is a true need for a novel design methodology, and, at least in some parts of the research community, there is an increasing level of awareness that design for emergence might be a key issue.

3.5 Theoretical framework

Developing a theoretical framework does not constitute a separate research theme; rather, within each of the four themes proposed above, contributions to the theoretical framework should be part of the research. Thus, we expect that as the individual research topics are tackled, the framework will be continuously augmented. This is important because the theoretical reflection implies a certain level of abstraction and thus transferability to other domains. For the time being we would base the theoretical framework on the following assumptions:

- *Systems* that we intuitively consider intelligent always comply with and exploit their ecological niche, and they display a high behavioral diversity. For understanding and designing intelligent systems, three time scales can be distinguished based on the underlying mechanisms: (a) “here and now”, (b) ontogenetic, and (c) phylogenetic. The synthetic methodology, “understanding by building” has proven highly productive in the past because we are interested not only in understanding artifacts, but also in designing and building them. Design can be performed at the “here and now” scale (“hand design”), at the ontogenetic scale (“by learning and development”), and at the phylogenetic scale (“evolutionary methods”).

- *Agents as dynamical systems:* Intelligence always requires a (physical) body and can thus be viewed as a dynamical system. One crucial implication is that through physical interaction, agents create structured sensory stimulation; particularly interesting kinds of interactions, typical of intelligent and cognitive systems, are sensorimotor couplings. Also, as dynamical systems, agents have attractor states, preferred behavioral patterns corresponding to energy efficient motion. Because of their nature as non-linear dynamical systems, real world environments are predictable only to a limited extent. Moreover, given the sensory and time limitations, situated physical agents can only acquire very limited information about their environment. Thus, they have to be designed to function properly under these conditions. One way of achieving this is to “design for emergence”: behavior is always emergent and cannot be directly programmed into the system: (a) from an agent-environment interaction, (b) in terms of time scales, and (c) from local rules in groups of agents. The latter also implies the development of the agents’ low-level communication skills, i.e., the emergence of complex social interaction from simple sensorimotor capabilities.

- *Cognition and intelligence:* Various definitions for the terms cognitive and intelligent exist in the literature. The most widespread use is as a descriptive term for the large class of so-called higher-level processes, that is, processes not directly driven by the sensory and motor systems. In the literature, cognition is typically associated with information processing. Looking at the underlying processing, it turns out that what is called cognition, e.g., the ability for categorization (to make distinctions in the real world), can no longer be clearly separated from sensorimotor processes. In this sense, even walking has certain cognitive qualities. Thus, it is probably best to take the term cognition to designate certain classes of behavior that we consider particularly interesting and intuitively intelligent, e.g., categorization, abstract problem solving, reasoning, and natural language behavior. Cognition is clearly a continuum and arises in embodied systems as they develop from “infant” state into an adult organism and, for example, develop gestures out of elementary motions. Modeling this process, as in developmental robotics represents a grand challenge. Since cognition and intelligence are often associated with consciousness, contributions to the latter should also be part of the program. A similar point holds for artificial emotions.