Creating Brain-like Intelligence

Bernhard Sendhoff, Edgar Körner, Olaf Sporns

2008

Preprint:

This is an accepted article published in Creating Brain-like Intelligence. The final authenticated version is available online at: https://doi.org/[DOI not available]

Creating Brain-like Intelligence

Bernhard Sendhoff[1], Edgar Körner[1], and Olaf Sporns[2]

[1] Honda Research Institute Europe GmbH
Carl-Legien-Str. 30
63073 Offenbach, Germany
{bs,edgar.koerner}@honda-ri.de
[2] Department of Psychological and Brain Sciences
Indiana University, Bloomington
Indiana 47405, USA
osporns@indiana.edu

Abstract. In this chapter, we discuss the new research field brain-like intelligence and introduce and relate the contributions to this volume to each other.

1 Brain-like Intelligence

This volume brings together contributions from researchers across a broad range of disciplines who attempt to define promising avenues towards the creation of brain-like intelligence. What is Brain-like Intelligence? Although it seems necessary to have a good understanding of what one wants to create before one starts, there is no crisp and clear definition. As is often[3] the case, we have to be content by first identifying a list of ingredients and with drawing up brain-like intelligence and other types of "intelligence" that have been put forward in the past.

Over 50 years ago, the field of artificial intelligence was founded during the now famous Dartmouth conference [1]. After an enthusiastic start it quickly became clear that the original goals would be much harder to achieve than anticipated. Since then artificial intelligence has proceeded in many different directions with a number of research "spin-offs" but without realizing the goal of achieving general-purpose intelligent computing applications. Maybe the most successful applications of artificial intelligence have been in the areas of search engines and of logical reasoning systems leading eventually to expert systems. Expert systems succeed when a problem is sufficiently well structured and its complexity can be controlled. Some problems have these properties, however, natural environments do not. This is one of the reasons why classical AI has not been successful in building systems that can autonomously operate in natural environments. Shakey was the first mobile robot that appeared to "reason" about its actions [2]. While it was a technological masterpiece, it was plagued by slow response times, resulting in sluggish interactions with its environment. Even if

[3] There is no unambiguous definition of "life".

we acknowledge the tremendous increase of computational capacity during the last 50 years (roughly a factor of between 10^7 and 10^{11}), there are several factors that still prevent a stricly rule-based approach, for example the richness of the natural environment, its inherent unpredictability and ambiguity, and the hidden combinatorial explosion. Nevertheless, as Sloman argues [3], lessons learned from and within classical artificial intelligence remain relevant to the research program of creating brain-like intelligence and the reaction against everything related to classical AI may even have held up progress.

Computational intelligence is a research spin-off of AI that currently serves as an umbrella for several disciplines with three major directions represented by neural networks, evolutionary computation and fuzzy systems. While these directions started off as separate endeavors, they all addressed weaknesses of AI. AI's lack of flexibility is addressed by fuzzy systems, while AI's lack of adaptivity and learning is addressed by neural networks and evolutionary computation. These research areas have all been and still are successful in many disciplines, however, their contribution to a better understanding of the brain and hence to achieve brain-like intelligence has been limited (although recently Zadeh [4] has argued that fuzzy systems are a prerequisite to achieve human-level machine intelligence). Nevertheless, they continue to serve at least as starting points for our quest for brain-like intelligence. In this volume, we find a range of approaches based on neural networks and evolutionary computation, see the chapters by Suzuki et al. [5], by Elfwing et al. [6] and also the one by Sporns [7].

So where does this leave brain-like intelligence? As the name suggests, we would like to achieve intelligence as demonstrated by brains preferably those of highly evolved creatures. We could call it *pragmatic intelligence* or in an evolutionary sense, *selected intelligence*. This form of intelligence is required for the expression of certain behaviours which are in turn required to guarantee survival and reproduction (constituting the most basic values). This type of intelligence is not an end in itself; it is a means and it is expensive (in humans the brain consumes 25% of energy at 2% of body weight and it also requires the involvement of a very large proportion of genetic information, with about 70% of all genes expressed in the brain). At the same time, this type of intelligence must be available under all circumstances, it must be versatile and flexible, and since the environment constantly changes it must also be evolvable. It must allow the completely autonomous control of a body. It adapts and learns and *if necessary* – higher up on the phylogenetic ladder – it reasons. In a way, this line of thought also demystifies the brain in a similar way as Charles Jennings [8] wrote in his review of the book Accidental Mind [9]: "*... the evolutionary history of the brain is a series of design compromises, culminating in a jerry-built assemblage of redundant systems that make us who we are today.*".

Brain-like intelligence is a system property, it is the result of a well orchestrated interaction between control, body and environment; this is demonstrated by the experiments described by Sporns [7]. All three elements of the system grow and develop together; even if the environment in itself does not change, the perception of the environment changes in the course of development. We emphasize the role of the environment because it is the driving force behind it all. In an abstract sense, the only source of information during evolution and development is the environment, a point made also in the comments by Sloman [3].

In part, this information has been genetically stored, thus influencing development and then contributing to learning and adaptation. However, learning should not only be thought of as confined to task specific machine learning but also has to be newly interpreted from a systems perspective just as intelligence[4].

Brain-like intelligence maintains a representation of the environment including the system itself. It has to cope with a continuous influx of an immense amount of mostly unspecific and for the current system state irrelevant information. Therefore, brain-like intelligence has to relate *what it perceives* to *what it knows*. There is abundant evidence that the brain actively creates its own perception, i.e., it predicts and manipulates the sensory input by feedback from its representation of the world. This seems to be achieved by hierarchically and dynamically organizing and controlling the interplay between different processing streams and areas in a highly organized structure or architecture (e.g. columnar structure). More detailed hypotheses of these processes have been put forward, e.g. [12].

Brain-like intelligence cannot be identified with a singular functionality, it is its versatility, its robustness and its plasticity which makes it the object of our quest. Superior functionality in a single problem domain can often already be achieved by specialized technological systems (think about a simple calculator or a chess programme), however, no technological system can robustly cope with the plethora of tasks encountered by any animal during its life time.

Brain-like intelligence is not the only extension of artificial and computational intelligence that has been proposed in recent years: autonomous mental development [13], cognitive architectures [14] and brain-based devices [15] rate among them. As these designations already suggest, autonomy based on a developmental program to acquire knowledge during interaction with humans lies at the core of Weng et al. [13]. Krichmar and Edelman [15] formulate a list of six basic principles and properties of an intelligent machine which includes active and adaptive behavior, developing from a set of basic innate rules (values). While Krichmar and Edelman place emphasis on the relation to the neural details, Weng et al. do not make this a prerequisite for their approach. Krichmar and Edelman clearly state that the analogy to the natural brain must be taken seriously. Indeed their last criterion suggests to use experimental data from neuroscience to measure performance; we will come back to this point again in Section 7. The review article by Vernon et al. [14] nicely summarizes the trend that the two traditionally opposing camps of relying on physical symbol systems and emergent development systems merge into hybrid systems (e.g. shown in

[4] This "systems perspective" [10] is currently advocated "all over the life-sciences" e.g. in the emerging area of Systems Biology and in pharmaceutical research, where there is a growing trend to build data based models of whole organs and even organisms [11] to predict the effect of newly developed drugs.

Fig. 1 in [16]). Vernon et al. promote the thinking in terms of phylogenetic and ontogenetic development of environmentally embedded systems as also outlined in Section 6 of this chapter. They relate autonomous mental development as has been put forward by Weng et al. to cognitive architectures. At the same time, development is entangled with the environment, and the right sequence of tasks [17] is important for an iterative refinement of skills. All of these approaches are closely related to our understanding of brain-like intelligence, with most of the differences found in the details of the structural and developmental constraints of the model and in the ways by which phylogenetic and ontogenetic development and learning can and must be integrated. By rightly focusing on the interaction of the systems with their environments (low-level perceptual and high-level social/cultural), we should not forget that the system is also internally driven by different layers of microscopic and macroscopic dynamics. It is generally agreed that brain-like systems will have to continuously consolidate and re-arrange their internal state. However, on a microscopic level there is a second continuous internal dynamics including cellular signalling and gene regulatory processes to guarantee continuous operability at the macroscopic level.

The accumulation of ever more detailed biological, cognitive and psychological data cannot substitute for general principles that underlie the emergence of intelligence. It is our belief that we have to more intensively pursue research approaches that aim at a holistic and embedded view of intelligence from many different disciplines and viewpoints. We must augment "locality" by a global control architecture and functional isolation by environmentally integrated systems that are capable of autonomous (phylogenetic and ontogenetic) development and self-organisation consistent with brain evolution. The goal is to create systems that are vertically complete spanning all relevant levels and yet horizontally simplified, i.e. *ignoring something in everything without ignoring everything about something*. This line of thought is mirrored by the diverse contributions which we will briefly relate and cluster in the next sections.

2 The Theoretical Brain: Structure, Dynamics and Information

The aim of theoretical neuroscience is to understand the general principles behind the organization and operation of nervous systems. Therefore, insights from theoretical neuroscience would be the ideal starting point for creating brain-like intelligence. However, in many ways the foci are different. Brain-like intelligence is concerned with nervous systems inside organisms inside environments. This perspective yields theoretical as well as practical consequences. It is necessary to study the transition from microscopic to macroscopic levels of structural and temporal organization and it is necessary to embrace the interaction with the environment in the formulation of the mathematics of brain-like processing. Not surprisingly, extensions of information theory remain the best candidate for approaching the later problem. There have been several early attempts, e.g. [18, 19], and in [7] Sporns describes an intuitive way how the information exchange between system and environment can be formalized.

Information theory is intrinsically connected to entropy and thermodynamics, which with its statistical formulation has been one of the most successful approaches towards connecting a microscopic with a macroscopic level of description. In their contribution [20], Deco and Rolls aim at bridging the gap between cognitive psychology and neuroscience by formulating microscopic models at the level of populations of neurons to explain macroscopic behavior, i.e. reaching a decision. Secondly, they argue that statistical fluctuations due to the spiking activity of the network of neurons are responsible for information processing with respect to Weber's psycho-physical law. Therefore, the probabilistic settling into one attractor is related to finite size noise effects, which would not be observable in a mean field or rate simulation model. The operation of networks in the brain is inherently noisy, which e.g. facilitates symmetry breaking. Thus, stochastic aspects of neural processing in the brain seem to be important to understand global brain function and if this is the case, we can expect that it will equally be important to create brain-like intelligence.

The role of structure and dynamics for brain-like intelligence is also the focus of the chapter by Jost [21]. This collection of mathematical tools for the analysis of brain-like systems highlights the different sources for rich dynamics (network structure, various update functions, e.g. logistic map) as well as the importance of temporal synchronization of flexible neuron groups in a network. Many different organizational levels of dynamics make up the overall information processing patterns in brains. Neural dynamics is self-contained as well as stimulus modulated in an active way (the brain moves) while being organized on an equally dynamical structure that is changing on different time scales during learning, development and evolution. According to Jost all this must be seen in the light of information theory, which is also a driving force behind the experiments described in Sporns' chapter [7] .

Sporns chooses an abstraction level that is closer to the actual biological system. However, the triad of structure, dynamics and environment is also the central focus of Sporns' chapter. The brain is a complex system because it consists of numerous elements that are organized into structural and functional networks which in turn are embedded in a behaving and adapting organism. Brain anatomy has long attempted to chart the connection patterns of complex nervous systems, but only recently, with the arrival of modern network analysis tools, have we been able to discern principles of organization within structural brain networks. One of the overarching structural motifs points to the existence of segregated communities (modules) of brain regions that are functionally similar within each module and less similar between modules. These modules are interconnected by a variety of hub regions that serve to functionally integrate the overall architecture. This duality of segregation and integration can be assessed with information theoretic tools and is at the root of brain complexity.

3 The Embodied Brain

Embodied Cognition advocates the view that the mind is shaped (or, to put it even more strongly, defined) by the body. Although the basic idea of embodiment has been proposed some time ago (see, for example, the writings of the French Philosopher Maurice Merleau-Ponty), it was not prominently pursued in classical artificial intelligence, which instead focused on abstract and deliberately disembodied symbols and rules. Today, Rolf Pfeifer is an ardent advocate of embodied cognition, a view that is reflected in his contribution in this volume [22]. Pfeifer and Gómez outline a number of different cases where complex control problems are fundamentally simplified by appropriate morphology (e.g. sensory morphology and material properties). They argue that information from the environment is structured by the physical characteristics and morphology of both the sensory and the motor systems (information self-structuring), a point also made in the chapter by Sporns. At the same time, Pfeifer and Gómez point towards a trade-off or balance between the maximal exploitation of the dynamics (requiring less sophisticated neural control) and maximal flexibility during operation. In other words, morphological computation is efficient and necessary, however, it is also restrictive with respect to the level of adaptation that can be achieved on a short time-scale, i.e., the individual life-time. It would be interesting to explore how this trade-off is realized by different species in biology.

Embodied cognition is strongly related to how our control of the environment shapes what we experience and how we develop. Possibly the most important means to both control and experience our environment are our hands. Ritter et al. argue in their contribution [23] that manual control and grasping are strongly related to the development of language which is usually seen as an instantiation of "pure" cognition (i.e. symbolic and independent of sensorimotor processes). The term manual intelligence is used by Ritter to denote a paradigm shift towards an understanding of the environment as a source of contact and interaction that is embraced and not avoided by robots. As a result language has developed as an extension of the ability to physically manipulate objects to the skill to mentally re-arrange and assemble ideas and descriptions of objects.

Although we will discuss embodiment in the light of evolution in Section 6, the embodied evolution framework put forward in the chapter by Elfwing et al. [6] complements our picture on embodied cognition. The population aspect inherent in any evolutionary interpretation adds a new perspective to embodied cognition which is intrinsically linked to the aspect of communication put forward by Ritter et al. and the notion of social and imitation learning discussed in [24–26] and summarized in Section 4.

4 The Social Brain

The brain resides in a body which lives in an environment. However, the environment is more than just a complex assembly of physical objects to explore and manipulate. Especially in more highly evolved species, the physical environment of an individual organism consists of other conspecifics which are often in the role of social partners. The importance of this type of interaction often termed social interaction has been studied since a long time in particular in animal behaviour [27]. It is known that the development of brain structure is altered by social deprivation.

In developmental robotics, social interaction has been studied e.g. in the context of imitation learning, which is the focus of the chapter by Grimes and Rao [24] and which is also described in the chapter by Vijayakumar et al. [25]. In kinematic imitation learning the observed behaviour of the teacher has to be mapped into the kinematic space of the observer. Grimes and Rao build a dynamic Bayesian network model of the imitation learning process including sensory-motor learning and implement it on the humanoid robot HOAP. The probabilistic structure of their model allows to deal with uncertainties which are inherent to imitation learning and to incorporate prior-knowledge in a very natural way. HOAP learns from two sources of information: demonstrative and explorative and under two types of probabilistic constraints: matching (observing the teacher's state) and egocentric (constraints of the learners state).

Imitation being basically uni-directional from the teacher to the observer can be seen as a first step towards social interaction. Wrede et al. argue in their contribution [26] that truly bi-directional interaction and communication is the basis of successful infant learning. They describe the necessity of joint attention (a kind of mental focus) for learning of artificial systems in interaction. In this way, the attention of each interaction partner can be directed to a common focus. The top-down process of joint attention can be triggered by bottom-up saliency strategies. The synchrony of information of different modalities supports to achieve this focus on a subset of the available information. Wrede et al. suggest that interaction strategies derived from verbal and non-verbal interaction from which turn-taking and feedback strategies are derived are required in order to build systems that can engage successfully in social interaction.

5 The World inside the Brain – The Brain inside the World

Systems displaying brain-like intelligence need an appropriate structure or architecture of processing layers. In Section 1, we have mentioned the research field of cognitive architectures being deeply related to brain-like intelligence. Such architectures allow the system to operate in the world and to represent the world. It is the amalgamation of these traditionally separate views that have been referred to as hybrid systems and which will be a prerequisite to success. There is a spectrum of approaches ranging from the more operational to the more representational philosophy.

Vijayakumar et al. suggest in their contribution [25] an adaptive control and planning system for robot motion within a statistical machine learning framework that can be coupled to a number of different information sources. They argue that the probabilistic and statistical level of description allows to abstract from the detailed underlying neural organization while being able to represent, process and fuse information in a functionally similar way. In particular, such system architectures can cope well with missing information and are suitable for statistical learning and inference mechanisms. Vijayakumar et al. compare the identification and learning of random latent, i.e. not directly observable, variables to the development of "internal" representations in cognitive architectures. They discuss and extend classical robot control schemes and apply the proposed probabilistic framework to imitation learning in humanoid robotics.

In the chapter by Goerick [28] an architecture called PISA for an autonomously behaving system that learns and develops in interaction with the environment is outlined. The acronym PISA stands for **P**ractical **I**ntelligence **S**ystems **A**rchitecture and consists of a large number of different elements that are described in the chapter. Goerick puts emphasis on the role of internal needs and motivations in PISA and on the approach to incrementally realize such a complex system architecture embedded in the environment. He proceeds with the description of interactive systems that allow motion control, online learning and interaction in different contexts on the humanoid robot ASIMO and that have been developed within the proposed framework. Furthermore, he suggests a notation called "Systematica" that is especially designed to describe incremental hierarchical control architectures.

Suggesting an architecture in between the cognitivist and the emergent view, Eggert and Wersing outline an approach toward a cognitive vision system in their chapter [29]. They focus on the conceptual role of control processes in the visual system to keep the combinatorial complexity of natural visual scenes under control. Questions concerning the high-level representational framework, the low-level sensory processes, the mediating structures of the control and the optimization criteria under which the control processes operate are discussed in the chapter. Eggert and Wersing proceed to highlight a few visual processing "subsystems" that are relevant for the general architecture. Examples are image segmentation, multicue tracking, and object online learning for classification and categorization. The self-referential build-up of a visual knowledge representation is an important element in Eggert and Wersing's chapter. At the same time, they emphasize that visual scene representations are sparse and volatile and therefore only store what is needed and what was accessible under given resource constraints.

We started out this section with a more "brain inside the world" focus represented by the work of Vijayakumar et al. and Goerick, then proceeded discreetly to a slightly (note that we are being very careful here) more "world inside the brain" view in the cognitive vision system discussed by Eggert and Wersing, arriving now at Sloman's chapter [3] where he argues that embodiment is not the solution to the intelligence problem it is merely a facet of it, arguably an important one.

Sloman moves further along the brain-world axis and puts forward that early AI has failed to put sufficient emphasis on the embodiment aspect but that does not mean that all of the earlier work is meaningless in our quest for brain-like intelligence. In his chapter, Sloman addresses several aspects of cognition including the development of children where it is evident that the explanatory power of primarily sensory driven systems with a relatively straightforward dynamics connecting the sensor with the motor side is not sufficient. Instead he makes a case for a multi-level dynamical system where the majority of processing happens decoupled from the direct environmental I/O. However, for Sloman that does not mean more or less but different emphasis on the environment. He suggests to study the features of the environment relevant for animal and robot competences and the different ways biological (and we would add artificial) evolution has (and would) respond to them.

Tsujino et al. [30] outline two models of the basal ganglia for autonomous behavior learning. The system-level model uses a reinforcement learning framework whereas the neuron-level model employs a spiking neural network. The two different levels of abstraction allows the authors to address different questions which are related to the function of the basal ganglia. While the issues of reward setting and input selection are the central focus for the system-level model, the spiking neural network is used to investigate e.g. mechanisms of timing.

6 The Evolved Brain

Brain-like intelligence is *selected intelligence*. Therefore, it is inherently put in an evolutionary context. But what are the consequences? There is an iterative development of functionalities that relates to the phylogenetic development of the architecture that is genetically represented and environmentally adapted. Although Haeckels orginal statement "ontogeny recapitulates phylogeny" is false in its literal interpretation, it is true that phylogenetically older structures generally occur earlier during ontogenetic development. Since in a cascade of hierarchically organized processes that are executed during development, it is easier to implement change at a later stage of the process than at an earlier stage, this is a natural result of an evolving system. However, this puts an immense strain on the richness of the architectural and organizational (in the dynamical system sense) primitives which evolution could manipulate. Evolvability of brains heavily constraints its processing principles. Robustness is a consequence. Were the processing principles brittle any evolutionary change would result in system failure. The genetic representation in itself is a complex information processing structure. Gene regulatory networks build cascading nonlinear dynamical systems that encode information indirectly. High structural and temporal precision is expensive (energy) and cannot be achieved globally. The huge complexity of the brain can – in general – only be represented by the genetic apparatus in a coarse way, the result is an inherent requirement for flexibility. If there is no means to specify each neuron location and neural connection precisely, a system has to evolve that is flexible. In a sense the shortcomings of the evolutionary process is – to a certain degree – responsible for the desirable properties of brain-like intelligence. Evolution is situated design, i.e. system development and system operation are not spatially decoupled like in traditional system design. Therefore, the embodiment discussion is meaningless from an evolutionary perspective (Sloman [3] calls it a tautology); only both together constitute an individual which is subject to selection. However, current discussion on development generally assumes a developing control system inside a fixed (chosen) body; from an evolutionary perspective this makes little sense and it remains to be seen to which degree this separation can be upheld. Note, that this goes beyond the co-evolution of body and brain that has been demonstrated by Lipson and Pollack [31].

Besides the principal relation between brain and evolution, there is also a more pragmatic one. Evolutionary computation offers a powerful approach to the optimization of complex structures on non-differentiable, noisy and multi-modal quality landscapes. In particular in combination with faster more local search techniques (reinforcement learning, gradient descent, BFGS) evolutionary algorithms have proven to be very successful for the adaptation of systems. The field of evolutionary robots [32] demonstrates this. In their chapter [6], Elfing et al. successfully integrate both methods for adaptation in their cyber rodent project. They present a framework for embodied evolution consisting of both a simulation environment and a few hardware robots using an elaborate mating scheme without explicit fitness assignment. The genotypes of the cyber rodents contain information on the neural top-layer controller and on the learning parameters. Reinforcement learning is used for lifetime adaptation. The two-layered control architecture selects learning modules dependent on behavior, environment and internal energy. Suzuki et al. co-evolve active vision and feature selection in a neural architecture in their chapter [5]. Active vision is the process of selecting and analyzing parts of a visual scene. Although the degree of freedom of the neural system is limited (the structure is fixed), the authors nicely demonstrate the selective advantage of active vision in their evolutionary set-up.

7 The Benchmarked Brain

The objective measurement of sucess and progress is a vital element in brain-like intelligence as it is in any other fields of science and engineering [33, 34]. Although varying between science and engineering, the different aspects of system verification and validation are well established. However, in brain-like intelligence we face additional difficulties. Firstly, we are not clear yet, whether we shall position ourselves more within the scope of science or technology. In the first case, success has to be judged by neurophysiological or psychological experimental data as is the case in computational neuroscience. In the second case, the target is to provide evidence that the realized system accomplishes its intended requirements. Of course, in this case we have the initial burden to clearly define what the intended requirements are against which we want to judge our progress. The fact that there is a continuous transition between both extreme stand-points makes the measurement process even more ambiguous. In [35], Herrmann and Ohl clearly follow the "science path" by suggesting "cognitive adequacy" as a measure. They argue that if it is possible to observe the same or similar behaviour in artificial systems as the real brain demonstrates, then it is likely to work like the real brain. They proceed to identify a number of anchor-points suitable for the comparison between system and the real-thing: reaction times: differences/ratios; error rates: cognitively adequate algorithms should make errors under the same circumstances as humans would; perception measures: show similar illusory or ambiguous percepts as those in human perception. This view is similar to the one put forward by Krichmar et al. which we discussed in Section 1.

If we follow the technological stand-point, we have to start by stating our system requirements or by laying out the rules for competition. There are small-scale (benchmarks for machine learning or image processing) and large-scale competitions (RoboCup, Darpa Urban Challenge), however, from the perspective of "brain-like intelligence" they always leave a feeling of dissatisfaction behind. The reason is that functionality can be achieved in many different ways and a functional approach to judging the level of "brain-like intelligence" inside a system or machine would end up in the endeavour to define brain-like intelligence. The results are typically lists with various items and this is where our dissatisfaction comes from, we know that lists are only of temporary validity. As we argued above, we would like to judge the complete system, but against what? In a recent BBC interview [36], Dharmendra Modha, manager of Cognitive Computing at IBM, proposed a radical solution to the problem: "We are attempting a 180 degree shift in perspective: seeking an algorithm first, problems second." Although at first sight an intriguing and perhaps even plausible statement, it would by definition preclude an objective and unambiguous measurement of progress: this is a dangerous path for technology to choose. At the same time, the path which is currently pursued is not much preferable, as more and more publications solve a more or more specific problems that a particular research group has committed itself to. This often does not allow comparison or objective measurement of quality. It does not even guarrantee reproducibility, because the complexity of most systems is too high and the provided detail of implementation information is too low.

This problem receives less attention in the scientific community than it should because it lies at the core of an overtly successful approach towards creating brain-like intelligence. In particular for those who invest in this field of research (research agencies, industry), a solution to this question is vital. In our opinion, the truth will be somewhere between the scientific and the technological stand-point, where function can be combined with experimental observation to lead the way through the jungle of brain-like system architectures.

8 Summary and Conclusion

We are facing a puzzle where we believe that we have identified a couple of pieces to be important (those are the focus of our research) but the overall picture remains fuzzy; we cannot be sure about the importance or centrality of the pieces and we have not yet figured out which pieces will connect to one another.

If we summarize our situation it is a bit like this: we aim for a fuzzy target using mostly relatively brittle approaches while having difficulties to measure progress. So where are the good news? The good news is: We are beginning to move in the right direction, even beyond the progress that is driven by the increase in computing power. The systems we build now are more open, more flexible and more adaptive than those of the past. We have understood that we do not build systems to operate in the environment but with the environment and because of the environment. This collection of papers is evidence of this progress.

It is interesting to note that there seems to be a certain reservation to bridge the gap from neuroscience to advances and new developments in intelligence research, e.g. brain-like intelligence. Judging by the conceptual proximity the number of chapters in this volume that relate to neuroscience is relatively small (mainly the ones by Deco and Rolls, by Sporns and one of the models by Tsujino et al.). This does not seem to be an observation that is restricted to our effort in brain-like intelligence. Indeed the recent collection of papers [37] from researchers in artificial intelligence dedicated to the 50th anniversary of AI has a similarly limited number of neuroscience related contribution. What could be the reason? Of course we can only speculate. However, on the one hand neuroscience is too much (from the view point of brain-like intelligence) focused on the details of neural processing instead of on the large-scale processing principles. On the other hand, research in brain-like intelligence must take care to emancipate – to a certain degree – from the grasp of technology.

There are a number of interesting and important questions that we have not addressed in this chapter and which are also not addressed by any of the contributions to this book. High on this list of omissions rates the question concerning the computational substrate for systems exhibiting brain-like intelligence. It seems that over the last five decades, research in intelligent systems has proceeded by incorporating more and more biological principles into the blueprint for our approach towards intelligence. In this quest, can we ignore whether we compute with cells in an organic systems or with gates in a silicon system? In the brain, we cannot distinguish between hardware and software. The architecture, structure and algorithms have evolved together and it is impossible to say where one ends and the other starts. It would be an evolutionary accident if we could extract some principles out of the context of the remaining ones and still expect this one to perform well.

The community has learned to scale down expectations over the last 50 years. So where can we go in the next fifty? We will build machines that support us robustly and autonomously both in the real and the virtual world. Will they challenge us cognitively? We do not think so. However, every reader is invited to speculate about the future after meeting the present on the next 344 pages. Finally, we cannot phrase it better than Turing [38]: *We can only see a short distance ahead, but we can see plenty there that needs to be done.*

References

1. McCarthy, J., Minsky, M.L., Rochester, N., Shannon, C.E.: A proposal for the Dartmouth summer research project on artificial intelligence. http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html (1955)
2. Nilson, J.: Shakey the robot. Technical Report 323, SRI International AI Center (1984)
3. Sloman, A.: Some requirements for human-like robots: Why the recent emphasis on embodiment has held up progress. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 254–284 (this volume).
4. Zadeh, L.A.: Toward human level machine intelligence – is it achievable? The need for a paradigm shift. IEEE Computational Intelligence Magazine **3**(3) (2008) 11–22
5. Suzuki, M., Gritti, T., Floreano, D.: Active vision for goal-oriented humanoid robot walking. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 311–321 (this volume).
6. Elfwing, S., Uchibe, E., Doya, K.: Co-evolution of rewards and meta-parameters in embodied evolution. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 285–310 (this volume).
7. Sporns, O.: From complex networks to intelligent systems. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 16–32 (this volume).
8. Jennings, C.: The most wonderful organ – not! Nature Neuroscience **10**(11) (2007) 1339
9. Linden, D.J.: The Accidental Mind. Harvard University Press (2007)
10. de Schutter, E.: Why are computational neuroscience and systems biology so separate. PLoS Computational Biology **4**(5) (2008)
11. Zheng, Y., Kreuwel, H.T., Young, D.L., Shoda, L.K., Ramanujan, S., Gadkar, K.G., Atkinson, M., Whiting, C.C.: The virtual NOD mouse: Applying predictive biosimulation to type 1 diabetes research. Annals of the New York Academy of Sciences **1103** (2007) 45–62
12. Körner, E., Matsumoto, G.: Cortical architecture and self-referential control for brain-like computation. IEEE Engineering in Medicine and Biology **21**(5) (2002) 121–133
13. Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., Thelen, E.: Autonomous mental development by robots and animals. Science **291**(5504) (2001) 599–600
14. Vernon, D., Metta, G., Sandini, G.: A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. IEEE Transactions on Evolutionary Computation **11**(2) (2007) 151–180
15. Krichmar, J.L., Edelman, G.M.: Brain-based devices for the study of nervous systems and the development of intelligent machines. Artificial Life **11** (2005) 63–77
16. Duch, W., Oentaryo, R.J., Pasquier, M.: Cognitive architectures: Where do we go from here? In Wang, P., Goertzel, B., Franklin, S., eds.: Frontiers in Artificial Intelligence and Applications, IOS Press (2008) 122–136
17. Sendhoff, B., Kreutz, M.: A model for the dynamic interaction between evolution and learning. Neural Processing Letters **10**(3) (1999) 181–193

18. De Vree, J.: A note on information, order, stability and adaptability. BioSystems **38** (1996) 221–227
19. Sendhoff, B., Pötter, C., von Seelen, W.: The role of information in simulated evolution. In Bar-Yam, Y., ed.: Unifying themes in complex systems – Proceedings of the International Conference on Complex Systems 1997. (2000) 453–470
20. Deco, G., Rolls, E.T.: Stochastic dynamics in the brain and probabilistic decision making. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 33–53 (this volume).
21. Jost, J.: Formal tools for the analysis of brain-like structures and dynamics. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 54–68 (this volume).
22. Pfeifer, R., Gómez, G.: Morphological computation - connecting brain, body and environment. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 69–86 (this volume).
23. Ritter, H., Haschke, R., Steil, J.J.: Trying to grasp a sketch of a brain for grasping. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 87–104 (this volume).
24. Grimes, D.B., Rao, R.P.N.: Learning actions through imitation and exploration: Towards humanoid robots that learn from humans. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 105–141 (this volume).
25. Vijayakumar, S., Toussaint, M., Petkos, G., Howard, M.: A statistical machine learning approach to planning and moving in dynamic environments. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 155–196 (this volume).
26. Wrede, B., Rohlfing, K.J., Hanheide, M., Sagerer, G.: Towards learning by interacting. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 142–154 (this volume).
27. Harlow, H., Harlow, M.: Social deprivation in monkeys. Scientific American **205**(5) (1962) 136–146
28. Goerick, C.: Towards cognitive robotics. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 197–219 (this volume).
29. Eggert, J., Wersing, H.: Approaches and challenges for cognitive vision system. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 220–253 (this volume).
30. Tsujino, H., Takeuchi, J., Shouno, O.: Basal ganglia models for autonomous behavior learning. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 336–359 (this volume).
31. Lipson, H., Pollack, J.B.: Automatic design and manufacture of robotic lifeforms. Nature **406** (2000) 974–978
32. Nolfi, S., Floreano, D.: Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines. MIT Press (2000)
33. Oreskes, N., Shrader-Frechette, K., Belitz, K.: Verification, validation and confirmation of numerical models in earth science. Science **263**(5147) (1994) 641–646

34. Oberkampf, W.L., Trucano, T.G., Hirsch, C.: Verification, validation and predictive capability in computational engineering and physics. Applied Mechanics Review **57**(5) (2004) 345–384
35. Herrmann, C.S., Ohl, F.W.: Cognitive adequacy in brain-like intelligence. In Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K., eds.: Creating Brain-like Intelligence. Volume 5436 of LNCS., Springer Verlag (2009) 322–335 (this volume).
36. Palmer, J.: IBM plans brain-like computers. BBC News (November 2008)
37. Lungarella, M., Iida, F., Bongard, J., Pfeifer, R., eds.: 50 years of artificial intelligence, Springer Verlag (2008)
38. Turing, A.M.: Computing machinery and intelligence. Mind **59** (1950) 433–460