Assistive Ecologies
- Bio-mimetic Design Of Ambient Intelligence

Mikkel Holm Sørensen
IT University of Copenhagen
Dept. of Digital Aesthetics and Communication
Email: megel@itu.dk

Abstract

Facing great challenges in designing better assistive technology with ever more complex and pervasive IT systems we need a general change in our understanding of designing IT. Ways to improve the performance of IT without drowning in complexity are of prime importance. This paper presents bio-mimetic design of Ambient Intelligence as a candidate heuristic for the task. It is shown how better assistive functionally goes hand in hand with improved infrastructural capacities by mimicking natural complex systems. On the basis of a presentation of organizational and developmental issues addressed by modern biological models we will suggest preliminary design constraints for Ambient Intelligence ecologies.

1. Dynamic IT: Ambient Intelligence

The IT colonization of human life presents us with great engineering and design challenges. These challenges primarily concern the infrastructural development and organization of IT [14, 19] and support for emerging heterogeneous and fluctuating user practices [11, 16, 17]. Both aspects requires a better long term design strategy for IT.

The challenge has been acknowledged within the IT research community and now the test will be how we respond. How should we design an increasingly complex and pervasive technology for the benefit of the users and not the mere proliferation of technology? This paper argues that by reducing our control over certain aspects of our technology we will actually get better and more adequate technology. Technology which semi-autonomously provides better assistance by following its own 'interests'.

There is broad consensus that to meet the infrastructural and functional requirements for a future technology IT will have to be endowed with adaptively dynamic capabilities – normally referred to as intelligent technology. Taken together with the tendency to embed IT into all parts of our lives this new technology is most properly captured as Ambient Intelligence (AmbI) [22]. This paper argues for a bio-mimetic approach as the only possible way to bring AmbI to realization. 'Bio-mimetic' means a systematic effort to apply natural principles in design of artifacts. The only extant systems with the functional complexity of prospective IT are biological, and biological processes are our best source of information for how such systems might be constructed and organized.

2. Bio-mimetic design of AmbI

To clarify what is meant by bio-mimetic design in this paper allow me to expand on the notion in order to avoid vagueness and misunderstandings.

On the face of it 'bio-mimetic design' appears somewhat incoherent since 'design' normally denotes an intentional and teleological practice whereas the ordering principles of nature are 'blind' and post hoc. On the other hand according to [15] the crisp dichotomy between the teleology of design and the causality of evolution is only apparent. Because the human brain basically operates by variation and selection dynamics thoughts and ideas (e.g. for design) are evolutionary selected rather than deliberately created. The difference between cognitive and other natural evolutionary dynamics is thus merely regional and not intrinsic.

Acknowledging the terminological fuzziness and being at the same time in agreement with [15] with respect to cognition (c.f. [6, 12]) I still find the notion of 'bio-mimetics' helpful in design heuristic terms. There is a difference between human deliberation and merely reaching solutions by due means (e.g. by an autonomous technology). This difference is not only practical: identifying the 'agens' is pivotal for psychological and ethical issues related to technology.

'Bio-mimetics' also needs qualification in another sense. The core principle of nature mimicked by this design approach - evolutionary dynamics - are not exclusively biological. Self-organization by variation and selection is an ubiquitous ordering principle in nature governing basic physical laws as for example crystal growth to sociological processes such as the proliferation of ideas. [6] Biology is only one, albeit very prominent, domain of evolutionary dynamics. [2] However since the design ideas sketched in this paper ultimately aims for devices with adaptive capacities only found in organisms the conventional term ‘bio-mimetics’ seem appropriate.

With this understanding bio-mimetics serves to fill in the blanks that an inadequate human teleology leaves open with the aid of uncontrolled but powerful evolutionary mechanisms while retaining general norms for functionalities. Bio-mimetics is a ‘meta-governing’ approach controlling the overall functional norms and relying on evolutionary processes to provide the functions required. Since the bio-mimetics relevant for AmbI design only mimic certain dynamic principles and not specific structures in nature the inspiration regards design heuristics, not solutions. (For specific motivations for applying bio-mimetics in AmbI design see: [2, 22, 23]).

3. Self-organization in complex systems

To operationalize the self-maintaining robustness and adaptive capabilities found in natural complex systems for the design of AmbI we must focus on the unique way such systems are structurally and functionally organized. This section briefly presents core characteristics of natural complex systems in order to determine guidelines how to obtain both autonomous robustness and flexible functionality in AmbI. The following is in accordance with an interactivist approach,1 a rich framework for complex systems which provides ideas for designing AmbI.

Self-organization in complex systems emerges from the interactive dynamics of the system constituents. These interactions are governed by the constituents themselves following onboard norms. Functional norms rise spontaneously in thermodynamic open systems. Because they are non-equilibrium, such systems depend on a controlled input of energy and information to maintain their existence, i.e. keeping up their functional organization. In contrast steady state systems remain stable if not perturbed by input of energy.

The dependence on input creates an asymmetry of normativity for the system; some things contribute to self-maintenance, others do not. Norms are not explicit let alone conceptual for most systems, but it does make a difference, e.g. for organisms whether they are hungry or not, whether the environment is hostile or not etc. The theory thus explains the rise of norms on the basis of functional coherence.

Norms are onboard since the system navigates by internal anticipatory cues of possibilities for further self-maintaining interactions which the present action provides. Interaction is cyclic and always for the sake of feedback (of information or energy), not the action *per se*. Retention of the outcome of earlier interactions - success or failure - acts as normative cues for subsequent similar situations. Cues range from infrastructurally ‘wired’ over conditioned response patterns to abstract concepts exclusively found in human cognition.

Because norms relates to the self-maintenance anticipatory cues for possible interactive outcomes are constructed by the system itself. There is no such thing as external norms in this regard, since systems differ (immensely) on which interactions support their self-maintenance. Energy in the form of food is of course a universal requirement for organisms, but the kind and amount of food needed varies enormously. When it comes to information of relevance differences are even greater. Being internal, values nevertheless relate to interactive consequences for the system and hence on external feedback.

3.1. Development and evolution

Natural complex systems adapt for the sake of self-maintenance by both short and long term means. In short terms systems adapt by learning. Through an ability to adjust responses on the fly in relation to changing needs and opportunities adaptive systems differ from merely self-maintaining systems by being recursively self-maintaining, i.e. they manage self-repair. In contrast to systems which only contribute to their own self-maintenance in fixed circumstances adaptive systems can actively change or leave the context if counterproductive.

A candlelight provides fuel for its combustion by melting candle, keeps an above threshold combustion temperature for the flame and attracts oxygen by convectional air turbulence. But it cannot create more candle when burned down, move if oxygen supplies runs critically low or a strong wind starts blowing. A candlelight is self-maintaining in a specific context but it is not self-repairing like organisms which actively honor their inner norms of self-maintenance.

Learning rests on retentiveness and error sensitivity. Retention provide the capability of fixing or associating contingencies related to a successful interaction (memory) and error sensitivity is required in order to adjust interactive strategies if unsuccessful. Error sensitivity and norms are mutually dependent.

In interactivist terms learning is a process towards improving the anticipatory capabilities of the system, i.e. constructing better anticipations for possible outcomes of interactions. Anticipations are contextual and implicit at the lowest cognitive levels and in the beginning of growth in higher level systems, but become increasingly generic

---

[1] The interactivist approach is a trans-disciplinary framework with roots in both evolutionary, pragmatist and phenomenological philosophy, developmental psychology, biology and process physics. Interactivism deals primarily with phenomena such as normativity, functions and meaning as rising from interactive processes in complex systems. [3, 4, 5, 6, 7, 8, 9, 10]

and universal as you ascend the hierarchy of intelligence or in the growth of higher level systems.

Besides ontogenetic adaptive mechanisms, complex systems (as species) improve their adaptability by procreative means. In reproduction combinatorial (sexual) or/and reconfigurable (mutational) possibilities for variation arise which are often more radical. This can be advantageous if the niche changes drastically or in order to explorer non-crowded places in the fitness landscape. However as adaptivity is mostly obtained over generations, retaining the qualities that made parents successful through a high level of replication fidelity is rewarded evolutionarily. [21]

Both short and long term adaptation rest on the same ordering dynamics namely variation and selection cycles. Actually variation and selection is the ordering principle of self-organization *tout court*. [6] Even though this principle is best known as evolutionary selection, it is also the prime driving force in ontogenesis as trial and error cycles. Evidence in neuroscience suggests that variation and selection dynamics are fundamental in the brain. Multiple neuronal groups offer assistance for a given task and some get selected [12]. Through neurological reinforcement rules synaptic strengths increase in the chosen candidates and by training successful groups become dominant.

4. Architectural ideas for AmbI ecologies

On the basis of this brief presentation of the theoretical fundament we can sketch some preliminary biomimetic design principles for AmbI.

Assistance should be provided by devices (both hardware devices and software applications) inhabiting AmbI ecologies striving to honor inner functional norms by achieving positive feedback from users. Users are the main resource in the environment of devices whom the devices quickly come to 'associate' with maintaining functional coherence. Feedback is earned mostly implicitly, typically as frequency of use, but occasionally explicitly by direct reward or punishment from users. Partly programmed and partly evolved fitness criteria determine the relevant kind and amount of feedback. For example in displays with several functions eye gaze monitoring provides feedback by determining which features the user mostly rely on.

The amount of feedback available in different contexts of use, called functional 'domains', relates to the demand for assistance. More use provides access to more resources. Competition and resources regulates the number and kinds of devices in different domains.

Devices will die only 'softly' and their hardware will be reused implementing software from successful 'parents' in similar domains. The software inherited is likely to be basic and only sufficient to get started. Exceptions to this will occur in more stable domains where users benefit from devices with greater specialization and more elaborate inherited code. Inherited specialization is no hindrance for dynamics though, as e.g. in searching assistants using evolutionary algorithms to execute their function. Learning do take place in such devices but is a matter of fine tuning and does not involve the global organization of the device.

Learning will often be required for 'juvenile' devices to compete with 'grown ups' and so they might be supported by beneficial rules, e.g. greater amounts of feedback or earning all the feedback when cooperating with 'older' devices. Some users may choose to 'raise' their own devices and tolerate initial poor performance.

Devices have no maximum age. 'Life cycles' are solely determined by performance; as long as devices perform satisfactory they stay in the game. Yet wear on hardware and accumulation of bad code will mean the 'death' of devices eventually. If the hardware is still ok and relevant inherited software will be implemented.

Variation and selection cycles takes place both at level of devices and their subroutines. Frequency and amplitude of internal variations are reversely proportional to the success of the device. Devices get 'innovative' to survive.

4.1. Two level architecture for AmbI devices

Devices will basically have a two layer infrastructure which I will refer to as the 'lower' and 'upper' level respectively. This architecture resembles the 'reaction', 'routine' and 'reflection' levels from [19]. However I have chosen cognitively more neutral names and avoid including 'reflection' which controversially indicates the possibilities of strong artificial intelligence.

Lower level: The lower level consists of simple, fast and reliable processes either hardwired or governed by preprogrammed algorithms. Like in Maslow's hierarchy of needs the lower level is related to the basic functional norms of the system and information from this level has much weight in upper level processes. The lower level handles basic issues (battery level, bandwidth etc.) error detection and system warning. Functions at the lower level are mainly general since low level specialization requires very stable domains to function due to reduced ongoing adaptivity. Little learning takes place at this level and variations mainly stem from inheritance variations. However changes do occur at the lower level when for example basic responses are inhibited due to higher level adaptation (see below).

Upper level: The upper level is were learning predominantly occurs and processes are governed by evolutionary computing methods. Minor variations are occasionally put forward to promote ongoing adaptation. Processes are more specifically adapted and context sensitive than at the lower level but also slower due to more complex processing. Adaptation to environmental invariants can overrule lower level signals (e.g. alerts) if they prove irrelevant or counterproductive over time.

Until we get closer to genuine AI the upper level will include some processes normally associated with the highest level cognitive systems, i.e. explicit and elaborate interactive strategies. The upper level monitors the state of the system and acts and reports accordingly. If an error occurs the system will enter a mode of lowered assistance ("anxiousness" in [19]) prompting it to deny requests for help from other devices, notify the user of its state if relevant and allocate resources to fix the problem. This is governance of functional coherence but not reflection.

4.2. Breeding AmbI devices

Variation and selection of devices happen through evolutionary computing methods, both at the evolutionary and the developmental level. However AmbI devices differ from organisms in many respects and not all evolutionary rules in AmbI mimic biological laws. Devices do not have an ontogeny, they do not breed or transfer genetic information directly etc. Yet adaptive dynamics can be achieved otherwise. Besides devices can take advantage of their artificiality in enhancing natural processes.

Because smooth functionality is critical in most contexts constraints have to be applied on variations to provide dynamics that is within an acceptable frame of functional deviance. At the developmental level productive learning models making room for 'experiments' while promoting smoothness are important. So we will probably need to design conservatively and allow only for limited dynamics, e.g. by constraining variations to subroutines of minor importance and by providing an option for returning to the previous setting if better.

Another constructive way to deal with the dilemma of dynamics and smoothness is to apply modularization. Development should be divided into consecutive stages to promote global adaptivity. If development is not horizontally aligned some parts might obtain premature local adaptivity counter to the global adaptivity of the device. [21] Flawless comprehending capacities in a conversation enabled device for example might take all the processing power from the proper assistive functions of the device such as retrieving information.

Developing brand new devices is different from the practice of inheriting software from 'parents' in similar domains mentioned above. Novel types of devices will go through a '*pre hoc'* phases. In the *pre hoc* phase candidates are developed by evolutionary techniques including tools from ALife and evolutionary robotics such as simulation of morphological, structural and functional growth. [23] A 'caretaker' environment 'raises' devices until they reach a tolerable level and are launched as 'juveniles'. From there on development will take place real time and in the relevant domain to obtain specific niche adaptation. This '*ad hoc'* developmental phase [23] is identical to the majority of devices with inherited 'parental' software.

Adjustments which cannot take place on the fly (e.g. morphological or structural changes in the case of non-reconfigurable hardware) happen by occasional modifications. Instructions for modification come from sensory measurements, user feedback and servers hosting relevant information from fellow devices.

Other kinds of modulations which seem to guide biological evolution without ruining dynamics might be applied. A modularization of devices into species (and perhaps even sexes) would help to promote variation and thus adaptivity while avoiding useless mutations. Such modularity decreases the combinatorial space and helps to home in on functional optima when developing devices.

4.3. Functional norms in AmbI devices

Since artifacts are not thermodynamic non-equilibrium systems, normative constraints cannot emerge spontaneously. All though they most likely will depend on electrical power to function, they will not structurally disintegrate by lack of energy. Besides the process is reversible and a 'dead' cell phone can be recharged. Norms will have to be imposed or at least initiated on the individual (type of) devices.

Designing norms for selection in devices is very difficult because fitness criteria is practically impossible to determine in advance and externally. Selection happens in relation to niches, and the specific ways a given technology will be used is unpredictable. In addition selection occurs at many levels and the complexity is immense.

Without suggesting that the task is trivial, ways of developing and evolving value systems in devices will be available eventually as a result of the work in ALife and evolutionary robotics. Promising work in evolutionary robotics [13] suggests that fitness functions on a simple "dynamically stable" neuronal-like architecture tend to give very good overall adaptation. In experiments with evolving control mechanisms for robots, simple fitness criteria used to evaluate the performance of strengths in the neural network gave rise to light-seeking behavior not specified by the fitness criteria. The robots constructed 'associations' of contingencies related to recharging when parked under a light source in a charging zone and obtaining positive feedback when outside the charging zone. Over generations the robots improved the behavior enabling them to reach the charging zone within a couple of time steps before total discharge and leaving the zone again immediately after recharging to maximize feedback.

Much of the success was due to the dynamic stability architecture of the robots which resembles an 'open' organization of non-equilibrium systems. Such architectures support learning as well as routine processes dynamically [13]. The results are in agreement with the interactivist view on the autonomous development of functional norms and a sensitivity for contingencies relevant for self-maintenance in adaptive open systems.

To obtain better interactional dynamics, some fitness criteria probably have to be explicitly determined by users. High level functional parameters in particular cannot be encoded beforehand and will develop faster by directly reflecting the satisfaction of the user. Some automatic responses, such as affective sensors, gaze tracking, body-language sensors, will be used to determine feedback, but the selective mechanisms still have to give room for explicit feedback to achieve sensitivity for nuances of functions wanted.

The very organization of fitness feedback will probably have to be adaptive as well to allow for emerging functionalities and the proper feedback for these. Natural higher level adaptive systems often have a nested organization supporting both learning (adjustment), meta-learning (adjustment of the adjustment) meta-meta-learning (strategies for adjustment of adjustments) etc.

4.3 Performance of AmbI ecologies

Through ongoing variation and selection cycles providing best solutions for the task at hand, a tight functional coupling between users and devices will emerge. Variation and selection dynamics will guide functionality at many different and nested levels. These range from sub-routines with seconds between selections (as in the brain) to major services and hardware with longer 'life-cycles' (as organisms). Thus devices adapt to users while subroutines compete to deliver the services necessary for adaptation.

Besides their individual efforts AmbI devices will provide better assistance by engaging in mutually supportive interactions with other system constituents to honor individual values. This might happen in swarm-like or 'symbiotic' collaborations.

Relevant but for some reason deserted functional niches will quickly become populated because even poor performances will immediately pay off when without competition and selection will reward devices seeking toward this domain. By the same logic dense functional niches will reflect greater need for assistance and more resources.

Due to the general interactive nature of the devices, some level of pro-activity and initiative will be manifest. Yet intruding and inappropriate behavior will get selected against like all other malfunctions. Pro-actively 'aggressive' devices will most likely be tailed by devices with less pro-active strategies ready to take over if aggressiveness proves a bad strategy. Devices will thus continuously adjust to the preferred level of activity and adapt to even subtle forms of interaction. Expectedly bio-mimetic AmbI will not only provide better assistance but also in a more calm and context-sensitive way.

The infrastructure in AmbI will grow evolutionarily. Unused devices become extinct and successful ones proliferate. But adaptation works both ways so that useful mutation also creates new demands. Domains can thus be created both by demand (users) and by supply (devices) since both creates an attractor in the fitness landscape.

Devices will to some extent manage self-repair or alternatively be able to suggest treatment. Perhaps special 'doctoring' devices could be developed to detect and repair by relating their fitness to number of 'bugs eaten'.

5. Non-natural architectural features

Not all parts of AmbI architectures should be governed by bio-mimetic principles, because bio-mimetic design will be most effective for the most dynamic aspects of IT. Very simple or already properly functioning technologies will not benefit from a bio-mimetic re-design (or at least not when compared with the effort going into re-designing).

Among the remaining elements of AmbI some parts should also be conventionally designed, and for different reasons. In some respects bio-mimetics should be constrained to guard simplicity and our ultimate control; in others bio-mimetics will simply benefit from adding non-natural methods.

It might seem strange to argue for a limitation of bio-mimetics for the sake of simplicity when the approach is suggested specifically to deal with growing infrastructural complexity. The reason is that fully bio-mimetic AmbI ecologies could become extremely complex and this is counterproductive when we need to interfere with the systems to overrule, adjust or add functions. In short, limiting complexity will be necessary for exercising control. We may risk not being able to analyze the intricate organization and dynamics of the system if we leave it to itself.

Strategic elements of the systems ought to be conservatively designed to allow for our continued insight. Again modularization is an important measure to guide system dynamics without killing them altogether. This could happen, for example, by building in 'dams of dynamics' (together with developmental stages) constraining variations to a determined window of 'variability'. Structurally modularization (building blocks) will also be an important instrument of AmbI systems to keep control with complexity. [23]

On the positive side bio-mimetic design can also be improved by non-natural mechanisms. Even though evolutionary dynamics is the single most efficient means for adaptivity it is not perfect. Besides it can be very slow.

The selection pressure in assistive ecologies will probably be greater than in nature because criteria for functionality are mostly higher than mere survival. To speed up evolution and promote a smooth functionality by maintaining a low threshold of tolerance we might enhance the evolutionary dynamics of AmbI ecologies by non- or rather super-natural means. For example by deploying horizontal inheritance to promote proliferation of successful functions synchronically as well as diachronically. Or by using learning mechanisms in reproduction to decrease bad mutations. An active heuristic selection would exclude known malfunctioning mutations beforehand and facilitate smoother assistance. Filtering must not be too restrictive though, since unpredictable dynamics is the main strength of bio-mimetics and the reason for applying it as complimentary to our limited designing capabilities.

6. Challenges ahead for bio-mimetic AmbI

Radical changes in the way we frame an issue raises many questions. The same goes for bio-mimetics. Some questions are a matter of empirical investigation and others of a more philosophical nature. I do not wish to pretend that the somewhat futuristic perspective put forward here is without huge difficulties, and I will touch upon some of the hurdles for bio-mimetic AmbI to deal with.

An obvious challenge is to model the proper architectures supporting evolutionary dynamics. We will need to find relevant fitness criteria, ways to execute inheritance, decide on reproduction rules, decide which levels and functions are best governed by evolutionary dynamics and tackle many other issues. And as if that was not labor extensive enough some of this has to be done for a range of different devices. Fortunately bio-mimetic AmbI shares these quests with evolutionary robotics, New AI, ALife and evolutionary computing which guaranties more research dedicated to solving the problems.

Generally we will need to find ways to balance human control with the autonomy of AmbI. Both in order for the evolutionary dynamics to proceed properly without untimely and ignorant human interference and to meet the raison d'être of technology: to serve our needs. Even though the classical myth of the revolting machines is highly unlikely, less dramatic and more realistic developments could be problematic enough. Besides the mere autonomy of IT - however functional - might be discomforting enough for some people to render the overall value of bio-mimetic AmbI negative.

On the pragmatic side AmbI could come to suffer from an unwillingness from major IT manufactures to agree on universal standards critical for AmbI. Commercial interests (e.g. monopoly, branding and business secrets) could get in the way. Time will show if the commercial prospects of supporting standardization are greater than sticking to individual formats.

A related challenge for the design of reliable bio-mimetic AmbI is the fact that what make natural complex networks failure-tolerant and robust also render them vulnerable in specific circumstances. Due to the so called power-law distributed structure of complex networks they are highly dynamic, but unfortunately also vulnerable to epidemic cascades whether manmade or 'naturally' emerged. [1] The swift proliferation and destructiveness of computer viruses have demonstrated the problem with all due clarity. The need for universal standards could become the Achilles heel of bio-mimetic AmbI by providing efficient infrastructures for the spreading of destructive codes. If we design systems with build in 'dams of dynamics' to allow for overview as suggested above this might also prevent destructive viral 'floods'. Another and probably complimentary measure could be to develop artificial 'immune' systems. [14]

This caveat relates to the general problem of balancing teleology with technological autonomy. Much use of IT cannot tolerate the 'death' of functions let alone large scale epidemics as nature does all the time. Whether the dilemma is empirical or essential is still hard to say. If essential we may be forced to choose between a sleek but inflexible and 'dumb' conventional IT and bio-mimetic AmbI which has both the virtues and the vices of nature.

7. Conclusion and perspectives

We live in a time of widespread bio-philia. Biology has replaced physics as the 'celebrity' science, bio-tech has replaced IT as the hot business, naturalistic cognitive science is having a major impact on the older human sciences, evolutionary theories are applied on a wide range of phenomena and everyday items like cars, mobile phones, sunglasses and furniture are being given organic forms as if they had been grown and not assembled. Examples are legion. Whether this tendency is part of a bona fide scientific and cultural progress or just an arbitrary fad is always hard to tell when in the middle of it. The truths is probably - as always - somewhere in between. Whatever the historical causes, there are also good reasons for applying a bio-mimetic design paradigm in the years to come.

Challenges in designing IT rising from increased pervasiveness, complexity and demands for a much more flexible and adaptive technology, calls for a fundamental change in our conception of design. Without letting go of our normative standards for technology we should stop clutching to total control and start embracing the power of bio-mimetic design. Designers would no longer create fait accompli artifacts but rather initiate and moderate, providing constraints and adaptive parameters for design processes. The very use would become an integrated part of design and everybody would become designers for themselves, by themselves. [23] Such a change will surely not happen overnight but probably start having impact on conventional design soon.

This paper argues for the perspectives in applying biomimetic design heuristics for coping with issues concerning the development and organization of AmbI systems. We have seen how optimizing internal stability, reliability and autonomy does not collide with improved assistance from AmbI, but that these two aspects on the contrary are mutually constitutive.

We have discussed future challenges related to a biomimetic design approach, some of which seem more concerning than others. Major obstacles will be balancing evolutionary dynamics with a low fault tolerance and determining the proper balance between system autonomy and our need for ultimate control. I have suggested means to meet these important questions, but they will most certainly need years of careful investigation.

8. References

[1] Barabási, A. L.: Linked. The New Science of Networks. Perseus Publishers, Cambridge 2002.

[2] Bentley, P. J.: "An Introduction to Evolutionary Design by Computers." In Bentley, P. J. Ed.: Evolutionary Design by Computers. 1-79. Morgan Kaufmann San Francisco 1999

[3] Bickhard, M. H.: "Function, Anticipation, Representation." In D. M. Dubois (Ed.) Computing Anticipatory Systems. CASYS 2000 - Fourth International Conference. 459-469. Melville, NY: American Institute of Physics. 2001

[4] Bickhard, M. H. and Richie, D. M.: On the Nature of Representation: A Case Study of James Gibson's Theory of Perception. New York: Praeger Publishers. 1983

[5] Bickhard, M. H. and Terveen, L.: Foundational Issues in Artificial Intelligence and Cognitive Science - Impasse and Solution. Amsterdam: Elsevier Scientific. 1995

[6] Bickhard, M. H. and Campbell, D. T.:□Variations in Variation and Selection: The Ubiquity of the Variation-and-Selective-Retention Ratchet in Emergent Organizational Complexity."□Foundations of Science. In press.□

[7] Christensen, W. D. and Bickhard, M. H.: "The Process Dynamics of Normative Function." Monist, 85 (1), 3-28. 2002

[8] Christensen, W. D. and Hooker, C. A.: "An Interactivist-Constructivist Approach to Intelligence: Self-Directed Anticipative Learning." Philosophical Psychology, 13(1), 5-45. 2000

[9] Christensen, W. D. and Hooker, C. A.: "Autonomy and the emergence of intelligence: Organised interactive construction." Communication and Cognition - Artificial Intelligence vol. 17 no. 3-4, pp. 133-157, 2000

[10] Christensen, W. D. and Hooker, C. A.: "Self-directed agents." In MacIntosh, J. (Ed.) Contemporary Naturalist Theories of Evolution and Intentionality, Canadian Journal of Philosophy, 31, Special Supplementary Volume. In press

[11] Ducatel, K. et al.: Scenarios for Ambient Intelligence in 2010. Report for the European Commission Community Research. IPTS Seville 2001

[12] Edelman, G. M.: Bright Air, Brilliant Fire. In the Matter of the Mind. Basic Books. New York 1992

[13] Floreano, D. and Mondada, F.: "Evolutionary Neurocontrollers for Autonomous Mobile Robots." In Neural Networks, 11(7-8). 1998

[14] Ganek, A. G. and Corbi, T. A.: "The dawning of the autonomic computing era". In IBM Systems Journal, Vol. 42, No. 1, pp. 5-18. 2003

[15] Gatherer, D.: The Memetics of Design. In Bentley, P. J. Ed.: Evolutionary Design by Computers. 1-79. Morgan Kaufmann San Francisco 1999

[16] Harris, J. and Henderson, A.: "Evolution in Action: HCI in a World of Pliant Systems." http://www.pliant.org/Papers-Area.html

[17] Kaasgaard, K.: "Beyond Formalisms. The Art and Science of Designing Pliant Systems." Interview with Jed Harris and Austin Henderson. In Kaasgaard, K.: Genres of Usability. Forthcoming http://www.pliant.org/Papers-Area.html

[18] Nolfi, S. and Floreano, D.: Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines. MIT Press/Bradford Books. Cambridge 2001

[19] Norman, D. A. et al.: "Affect and machine design: Lessons for the development of autonomous machines". In IBM Systems Journal, Vol. 42, No. 1, pp. 38-44. 2003

[20] Oyama, S. et al.: "Introduction: What Is Developmental Systems Theory?". In Oyama, S. et al.: Cycles of Contingencies. Developmental Systems and Evolution. pp. 1-11. MIT Press, Cambridge 2001

[21] Sterelny, K.: "Niche Construction, Developmental Systems, and the Extended Replicator". In Oyama, S. et al.: Cycles of Contingencies. Developmental Systems and Evolution. pp. 333-349. MIT Press, Cambridge 2001

[22] Sørensen, M. H.: "It's A Jungle Out There: Toward Bio-Mimetic Ambient Intelligence Ecologies." Computer, Communication and Control Technologies 2003

[23] Sørensen, M. H.: "Design Symbiosis: The Dynamics of Evolutionary Design." Forthcoming