Organic Architectures for Large-Scale Environment-Aware Sensor Networks

Paul Lukowicz, Institute for Computer Systems and Networks, UMIT Hall i. Tirol, Austria, paul.lukowicz@umit.at
Erhardt Barth, Institute for Neuro- and Bioinformatics, University of Lübeck, Germany, barth@inb.uni-luebeck.de
Jan Kim, Institute for Neuro- and Bioinformatics, University of Lübeck, Germany, kim@inb.uni-luebeck.de

Abstract

This paper examines how methods inspired by biological processes can be applied to the design of large-scale environment-aware sensor networks. Our ultimate goal are systems containing thousands of sensors spanning large building complexes or even cities that can cooperate to detect and analyse complex situations. We discuss the problems involved in implementing such systems using conventional engineering approaches, sketch a system architecture that can circumvent these problems by evolving a system optimised for a given application domain from a rough generic template, and discuss initial ideas for using biological analogies to perform the required adaptations. These include low-level attention, brain plasticity, and genetic regulatory networks.

1 Introduction

The classical engineering approach to system design requires an exact specification of system response for every possible situation that the system might encounter. This can be achieved by providing an analytical response function, by enumerating all relevant possibilities, or a combination of both. While highly successful as the basis for most of today's technical devices and systems, this approach is increasingly reaching its limits as systems get more and more complex and have to deal with the dynamics of the real world. In such domains, the possible system states and the corresponding system responses are usually analytically intractable. At the same time, the enumeration strategy is not feasible due to combinatorial state space explosion.

To deal with the above problem, the organic computing concept takes inspiration from the way biological organisms are able to adapt to a complex dynamic environment. Rather than to exactly specify every single state and response the idea is to: (i) sketch a rough outline of the system containing plausible possibilities for key states and responses, and (ii) provide a set of rules that allow the system to autonomously adapt and evolve the appropriate states and responses over the course of its lifetime. As obvious as the above might seem, in particular since it has been functioning for biological systems for millions of years, today it is far from clear *where, how, and if at all* this idea can be applied to technical systems. This paper presents the results of an initial, conceptual architecture study that aims to answer the above question. To this end, we look at a specific class of systems that has recently generated much interest: large-scale environment-aware sensor networks and discuss (i) the problems involved in implementing such systems using conventional engineering approaches, (ii) a system architecture that can circumvent these problems by evolving a system optimised for a given application domain from a rough generic template, and (ii) specific initial ideas for using biological analogies to perform the required adaptation.

Self-organisation in large scale sensor networks is a very active research topic [10] that includes some large projects [22, 11, 9]) in which self-organisation is mostly used as a means of optimising routing with respect to resource consumption and service discovery. By contrast, we investigate the use of bioinspired adaptation methods as means of improving the ability of the system to extract meaning from the environment and satisfy high-level goals. Resource consumption related system-state concerns drive the adaptation. In terms of bioinspired analogies other sensor network project have mostly looked at areas like swarm intelligence and genetic algorithms. Brain inspired methods have so far been given little attention, partially because of the focus on routing and resource consumption mentioned above.

In the general area of situation awareness current work has had some success in modelling and detecting predefined situations and user activities (e.g. [12, 34, 33, 35, 23]) and using this information in so-called context-sensitive applications (e.g. [29, 5, 16]). However, the success is still mostly limited to simple situations in a constrained environment. Also, reliability is still a major issue.

In addition to looking at systems order of magnitude larger and more complex than the state of the art, an important aspect of our work is that we look beyond the use of one general solution for all situations. Instead we focus on the ability to adapt to the situation using multiple methods and pre-given patterns.

Considering the general area of organic systems, concepts are just beginning to emerge. Initial ideas that were presented at [1] include, among others, an adaptive approach to the control of traffic lights described by Rochner and Mueller-Schloer, organic middleware concepts presented by Trumler, Bagci, Petzold and Ungerer, concepts for computer crash management described by Haase and Eschmann, and gaze-based human-machine interaction [8]. Bioinspired methods as used in computer vision and genetic algorithms will be discussed in more detail in Section 4.

2 Problem Description and Example Scenario

We deal with sensor networks spanning large building complexes or even cities and analyse how such systems can be made to function as a **single coherent system** that uses a complex, dynamically evolving world model to extract high-level meaning from the events in the environment. Our approach is based on appropriate integration of classical proof-based and bioinspired, self organising methods. As a consequence,we start with a static, generic system template devised by using conventional engineering techniques and apply bioinspired self organisation methods to evolve it into an application-specific sensor network that is able to assign high-level meaning to events in a specific, changing environment.

2.1 Example Scenario

To illustrate our vision of large-scale environment-aware sensor networks and to outline the associated problems, we begin with a relevant example scenario that is both very rich and scalable: an intelligent hospital.

2.1.1 Sensing Modalities

A single hospital room might be equipped with a subset of the following information sources

- Sensors in the bed (bed configuration and patient posture) and other furniture.
- Identification of all equipment and personnel in the room
- RFID readers in the bed and table able to identify medicine vials, food and drink containers etc.
- Location for selected equipment and personnel (ultrasonic or through carpet) in the room
- Ambient sound sensors in the room
- Personal microphones
- Motion sensors in the hospital gowns of the medical personnel for activity recognition
- Information from equipment about settings, activity, and measurement results
- Cameras in the room.

In addition, different systems for tracking personnel around the halls (e.g. personal inertial navigation, RFID, WLAN location, ultrasonic) might be available. RFID readers and switches integrated in appropriate cabinets could be able to track who has taken what medicine out of which cabinet. Something similar could be done for simple and yet important tools like syringes or scalpels.

2.1.2 World Model and Extraction of Meaning

In principle, the above sensors provide enough information to fairly exactly characterise the environment. Examples of the type of information that we would like to be extracted include:

1. Primitive events and actions such as taking a particular measurement on the patient, having a sip of water, standing up, taking a particular vial out of a cabinet etc.
2. Compound activities such as administering medication, eating lunch, or performing a particular examination.
3. Localised processes and situations such as a doctor making his rounds, a personnel meeting, or an operation taking place.
4. Global processes and trends such as a shift change, equipment shortage, long term changes in response times to emergencies etc.

2.1.3 System Goals

Properly used, the above information could increase the reliability and improve the efficiency of a variety of processes in the hospital. With respect to a particular patient, this includes enforcing timely and correct medication or more accurate tracking of the patient's condition by combining information about activity, eating and sleeping patterns with the medical parameters. For the personnel, such a system could provide context sensitive information and notification, help optimise daily schedules, and assist in personal performance analysis. On the organisational level, duty rosters of the personnel, equipment usage, and patient assignments could be optimised.

2.2 Issues to Consider

In a large hospital complex with hundreds of rooms and people, we are looking at a system with thousands or even tens of thousands of sensors. The sensors need to recognise hundreds of different isolated event types, combine them, include additional information and evaluate it with respect to the high-level system goals. Similar numbers of sensors can be found in many other application domains such as factories and large office buildings. By contrast, the current state of the art of commercial systems is the use of sensor networks to monitor selected physical parameters such as room temperatures, tire pressures etc. in simple control loops. Research in ambient intelligence has had some success in going a step further to model and detect predefined situations and user activities and uses this information in so-called context sensitive applications. However, the success is still mostly limited to simple situations in a constrained environment. In addition, the recognition often lacks the robustness needed for real-life applications. In summary, the state of the art is still far from the above vision of extracting high-level meaning in large sensor networks such as the one described above. The reasons for this are in part well known problems facing any attempt aiming at creating an environment-aware system and partly issues specific to the ambient sensor network approach. They can be summarised under three headings: complexity, dynamic development, and resource constraints.

2.2.1 Complexity

There are three main sources of complexity that the system needs to deal with:

- Model Complexity. Activities and situation in the environment can rarely be viewed in isolation. Instead, the different levels and parts of the envisioned model hierarchy are interlinked in a complex manner. Considering the scale at which the envisioned networks are supposed to operate, any attempt to exactly specify all possible variants of such interactions is infeasible.
- Environment Variability. Human actions and interactions tend to be highly variable. Situations that on an abstract level belong to the same class and should be recognised by the system as identical can be quite different on a more detailed level. Such differences can emerge spontaneously as a reaction to specific events, or be a matter of personal style and preferences. Again, an attempt to pre-engineer a system to account for all possibilities faces combinatorial state explosion problems.
- Input Space Complexity. The size of the envisioned sensor networks implies that it is able to collect a large amount of information. This is of course necessary to be able to monitor complex situations in a large intelligent space. On the other hand, it implies that the dimension of the input space is so large that conventional recognition and

clustering algorithms have problems dealing with the data unless an appropriate partitioning is found. This means that it is not feasible to collect data from all the nodes and then process it in a central control system.

2.2.2 Dynamic Development

The only constant about the real world is continuous change. Changes occur on different time-scales, ranging from spontaneously occurring events to long-term trends. Since classically engineered systems tend to have problems dealing with a changing environment, whereas biological systems are good at mastering it, the dynamics are a key reason to look at bioinspired alternatives. In our case the relevant sources of change are:

- World Model Dynamics. The environment includes both abrupt, spontaneous events (e.g. a doctor is called to an emergency during patient examination) and medium to long term evolution. The latter encompasses trends such as systematic changes in the way certain procedures are conducted by certain persons as they get more experienced. In a hospital setting they might also include changes to procedures due to new equipment or regulations.
- Stimuli Dynamics. The quality of information about a given activity or situation contained in a particular sensor signal can greatly vary over time. The variation can be spontaneous (e.g. a light ultrasound receiver being occluded), periodic (e.g. a loud machine being switched on at a particular time of a day making sound recognition difficult), or long-term developments (e.g. a bed making more and more noises as it gets older).
- System Dynamics. The sensor network itself is subject to a number of different developments. On a short-term time scale parts of the system are continuously moving between locations, since we assume sensors to be also part of equipment and the outfit of the people. On the long-term time scale, we must assume sensors to drop out or be added as equipment is updated.

2.2.3 Resource Restrictions

The vision behind ambient sensor networks assumes that the sensor nodes are an integral, permanent part of the environment. This implies two things. First, depending on the specific location and role of the sensor node, there are stringent constraints on its size, weight, and price. Second, many nodes will not be connected to a fixed power supply. Neither will it be possible to regularly change a battery. Thus the nodes will either have to live off energy extracted from the environment or use so little power that a single battery will last for the entire life time of the device. In any case, power consumption is a primary concern limiting the performance and functionality of the nodes. As a consequence, resource constraints and management must be taken seriously. Interestingly, it is widely believed that energy consumption plays a key role in the performance limits of the human brain. Configuration problems of similar type can be found in many distributed and parallel systems. In previous work [2], we have for example discussed how the architecture of context-sensitive, wearable systems can be optimised according to similar criteria. Nonetheless, the problem considered in this paper is special for three reasons: (1) the sheer size of the optimisation space, (2) the need for dynamic reconfiguration at runtime with minimal resources available for the reconfiguration task and (3) the fact that the recognition problem for which the system is to be optimised is known only in very broad terms at system design time and changes as the network develops and modifies its world models.

3 Organic Approach Overview

Our approach aims to structure the sensor network in a hierarchical, parallel way that matches the structure of the world model. This is similar to the way different parts of the brain are structured to deal with complexity and resource restrictions. Based on such a structure, we define a generic system template that will also include feature specifications, pre-trained classifiers and initial guesses at the probability distributions involved in the modelling of high-level processes and goals. In a biological analogy, this would mimic the way some concepts are believed to be genetically pre-wired into the brain. Due to the complexity and variability of the environment in a network spanning a large space, such a model will be incomplete and very inaccurate. To counter this, methods inspired by cognitive processes (e.g. attention), brain plasticity and genetics (regulatory networks), as well as population evolution, will be used to allow the system to adapt to the specific environment in which it is deployed and to keep adapting as the environment changes. The adaptation will include structural changes (including the emergence of new structure corresponding to new concepts), feature and probability distribution optimisations, and strategies for resource management.

3.1 System Structure

In terms of system structure, the key concepts for dealing with complexity while observing the resource restrictions are parallelism and hierarchy. As a consequence, the envisioned architecture is based on the following assumptions:

1. The world model is hierarchic, roughly following the isolated event, compound action, local process, global process-and-trend categories structure presented in the example scenario.
2. As shown in Figure 1, the system consists of sensor nodes that contain a set of sensors, a simple processing unit, a communication interface, and a power supply. The specific sensors, the computing power of the processing unit, and the specifications of the communication interface can vary from node to node. So does the available power supply that can range from a fixed connection to a fully autonomous node scavenging power from the environment. What all nodes have in common is the fact that they each have their own limit regarding which sensors and algorithms can be used at the same time. This means that depending on the required recognition task some nodes might need to operate in different modes, which in turn means that not all types of recognition can be accomplished with optimal accuracy at the same time.
3. In every relevant location a group of sensor nodes is assigned to every primitive event/action that the system needs to recognise at this location. Thus for example in every patient room the sensors in the bed plus possi-

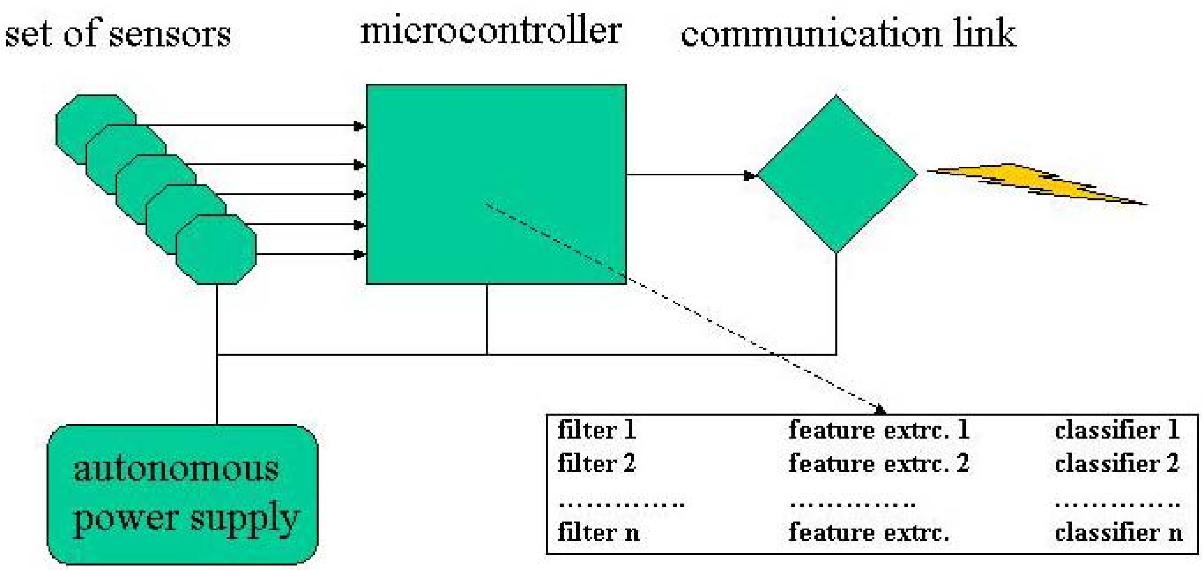

Figure 1 Structure of a sensor node

bly motion sensors in the gown and location sensor in the room would be in charge of detecting the patient getting out of bed.

4. The local event/action sensor groups are hierarchically organised following the hierarchy of the world model (see Figure 2). Thus every compound action, or local/global process is mapped onto a hierarchy of sensor node groups. In the hospital network the groups responsible for the detection of taking out a medicine vial from the cabinet, the group responsible for tracking the nurse along the corridor and the group responsible for recognising the action of taking a pill are interconnected into a recognition system for the process of medication delivery. For the compound actions and local processes there will be parallel groups corresponding to different locations and/or persons. For the global processes, single system wide groups will exists.

3.2 System Template

The system template provides two things. First, it defines an outline of the world model that the structure of the network will follow. This includes estimates of all probability distributions involved in the construction of compound activities and processes out of the primitive events/actions. They can be specified using such conventional approaches as Bayesian networks or probabilistic grammars and estimated using supervised learning techniques. Second, it provides an initial mechanism for the recognition of all primitive events/actions. The specification includes initial guesses at the best sensor nodes, features and node configurations to be used and classifiers pre-trained using a supervised approach. A key property of the template is that it is generic and not optimised to the specific environment in which it will be deployed. This means that one would have for example a generic hospital template applicable to different hospitals. It also means that for each primitive event only a single generic specification is provided. This specification has to be applied to all locations in which this event can then take place (see next paragraph). The same is true for the higher level compound actions and processes. The generic, single instance specification and the fact that we assume a manageable number of primitive events/actions to be the basis for all extraction of meaning allows the template to avoid combinatorial complexity explosion even for large spaces and complex environments. On the other hand it means that the initial performance of the system will be very poor. Also we cannot expect the template to be complete in terms of both the primitives and the compound actions/processes.

3.3 Adaptation

From the system point of view, four kinds of adaptation at different time scales are needed:

1. **Short-term resource management** Most sensors will be

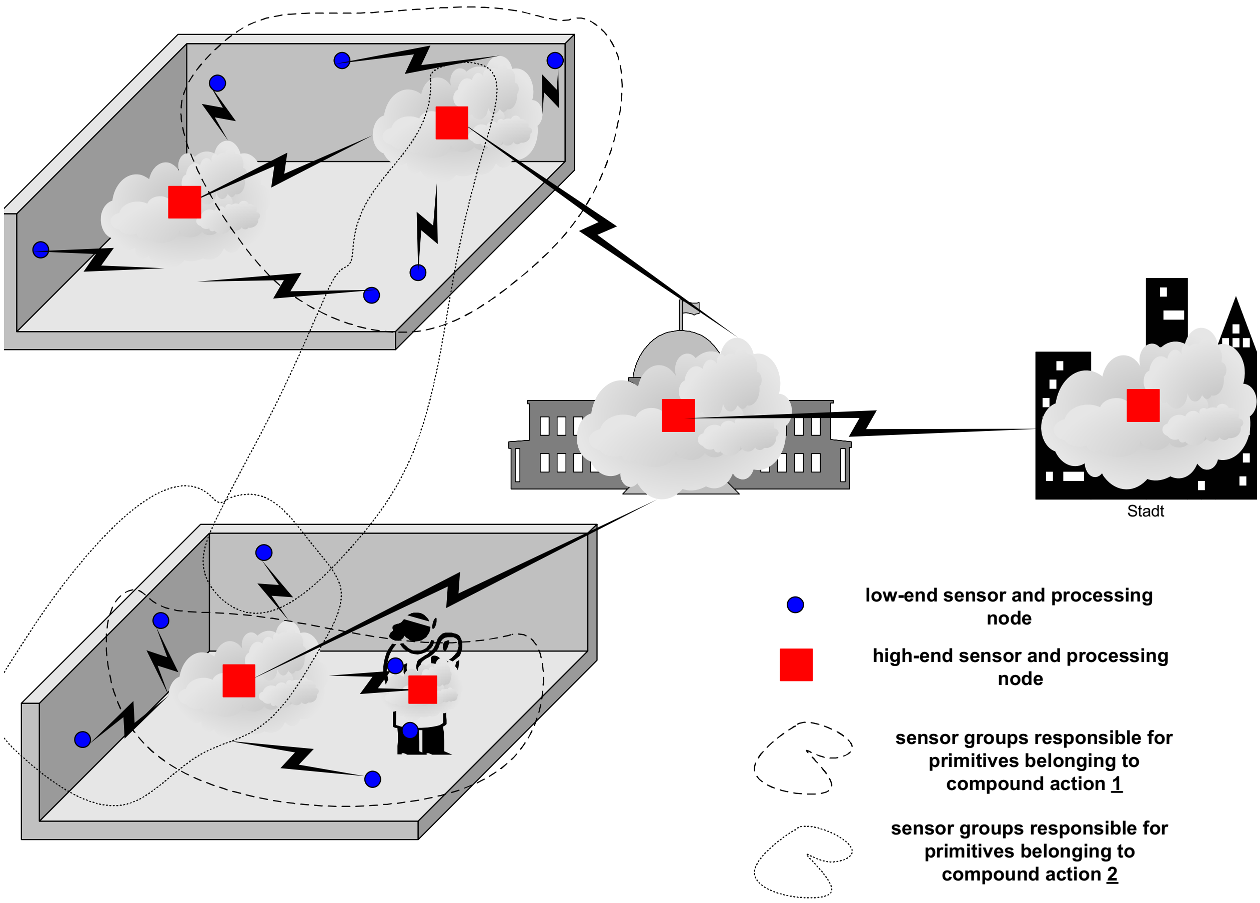

Figure 2 Structure of the envisioned sensor system and an illustration of a possible assignment of sensor node groups to primitives and compound actions.

part of a number of different groups dedicated to different events. In general, different modes of operation will be required by each group (resolution, computed feature, measurement mode). As a consequence it will not be possible to detect all possible events at the same time (at least not with the maximum accuracy). Instead, a mechanism will be needed to decide, which are the most relevant/likely events and focus the system resources on those.

2. **Medium-term, location-specific adaptation** As described above, the same events/actions and processes will be found in different locations and with the participation of different persons. However they are not likely to take place in an identical way. Such variations are major problems for current context-aware systems that often have to be trained specifically for each location and person. The development of methods that will allow a system trained for one location to adapt to another is therefore necessary. The adaptation shall involve the selection/weighting of sensors/features as well as the modification of the probability distributions underlying the classifiers.
3. **Medium-term structural adjustment** As described before, due to environment complexity and variability the system template will in general neither contain all primitive events nor will it accurately reflect all relations between primitive events in compound actions or processes. Thus the system must be able to evolve structures such as new event types, new connections and with it new higher level compound actions and processes. Some of those will be purely internal, needed to better represent higher levels and the goals. Others might actually be presented to human supervisors for evaluation.
4. **Long term system-level evolutionary adjustment** While the variations between locations are a major source of problems, they can also be used to the system's advantage. The different variants of a recognition subsystem for a given event/action developed at different locations can be seen as a population in which a kind of population-wide evolution can take place. While some of the differences will be location-specific, some will simply be location-independent improvements. They can appear accidentally or due to the fact that a particular event has been particularly common at this location. Overall system improvement analogous to evolutionary development of a species population can be achieved by propagating such general improvements to all locations.

4 Bioinspired Adaptation Strategies

For the four types of adaptation needed by the system, we consider the following analogies from biological systems.

4.1 Short-term resource management

This type of adaptation is analogous to the way humans tend to focus their attention on a particular subset of stimuli and ignore others given current goals. In animal vision, the gaze is directed to the objects of interest, which are then processed with high resolution and with colour, whereas the peripheral visual field is processed with decreasing resolution and no colour. Vision is an active and highly selective process, a strategy that makes biological systems very effective in acting in their specific environment. The active vision paradigm has meanwhile become an established area of machine vision: a computer, as well as a human observer, can understand its environment more effectively and efficiently if the sensor interacts with that environment [6]. Low-level attention is usually supposed to be driven by the sensory input data and characterised as a bottom-up process (while top-down indicates cognitive processes which depend on context, the world model etc.). A few models of low-level attention exist, see [14] for a review, that all attempt to find those image features that form a so-called saliency map and are more likely to attract attention. A low-level attention module should be able to decide which kinds of changes in the sensory input are more significant than others. Based on the concept of intrinsic dimension [36, 4], algorithms have been developed that can extract dynamic saliency from a video stream and predict human eye movements, e.g. [7, 8].

Thus, the problem of resource management can be handled by competition introduced by the attention-based control subsystem of the architecture. The process will involve a competition between the magnitude of a novel or unexpected sensory state against the goal-directed focusing of attention on other parts of the state space due to current ongoing operations that try to satisfy a higher level goal.

4.2 Medium-term, location-specific adaptation

It is one of the most stunning abilities of the brain that it can easily recognise objects despite large variances in their appearance. In general, such perceptual similarities cannot be described by a metric in input space [27]. Nevertheless, recent brain-inspired computer-vision systems are able to automatically learn object classes [28] and are beginning to reach acceptable levels of performance. However, current brain-inspired machine learning techniques need to be further developed to be used for location-specific adaptation.

As a complementary approach, techniques inspired by phenotypic adaptation can be used. The capability of a given genome to produce dramatically different phenotypes in response to differences in environmental conditions is another impressing example of the adaptive potential of biological systems. It is well known that, for example, human skin reacts to exposure to sunlight by producing dark pigments, and that plants can strongly vary in shape, leaf thickness, and many other properties in response to their local growth conditions.

The basis of such adaptation is the fact that expression, or activity, of a gene can be regulated. The molecular agents that mediate gene regulation (called transcription factors), are themselves gene products. Thus, the genes in a genome are organised in a directed graph of regulatory interactions called a regulatory network. By combining theoretical research on regulatory networks [25, 15] with recent advances in molecular biology, which enable identification of such networks in model organisms [21], new insights into the organisational principles of regulatory networks are currently beginning to emerge [3]. Furthermore, regulatory networks lend themselves well to various algorithmic abstractions [15, 19, 26]. These factors provide strong motivation for exploring algorithms for adaptation modelled after regulatory networks.

4.3 Medium-term structural adjustment

Recent experiments indicate that the plasticity of the brain can be spectacular. Sharma et. al. [32], for example, have shown that young ferrets can see in the sound zone, i.e. the auditory cortex is largely rewired if driven by visual instead of auditory input and can replace the function of the visual cortex to a large extent. Traditional machine learning does not involve drastic structural changes. Accordingly, neural-network models assume the changing of weights in the neural network as the basis of learning. Although there are a number of approaches for self-organising neural networks that can adapt to changing inputs, e.g. [20, 24], methods that allow for substantial changes in the network structure are still missing. Therefore, new kinds of machine-learning techniques which do not traditionally operate in a given feature space (like a neural network that adjusts its weights) but can continuously change that feature space to optimise a given goal are needed. This implies that feature spaces and classification boundaries therein can both change such as to optimise a criterion. This would be a departure from viewing learning as the process that estimates an unknown (classifier) function, which will then determine the forthcoming decisions. Instead, rather few and changing regions of the feature space would be used to make a decision that will be only optimal at that moment and for the then relevant goals.

4.4 Long-term system-level evolutionary adjustment

Biological evolution is capable of generating information with adaptive value by random variation (mutation, recombination etc.) and to accumulate this information by Darwinian selection. This process has been conceptualised by evolutionary algorithms [13, 30], which have successfully been applied in various optimisation and machine learning tasks. In most evolutionary algorithms, the objective is adaptation to a given set of conditions, i.e. a given, static fitness function. In contrast, natural evolution is capable of retaining genetic information which has temporarily lost its adaptive value as a result of changes in the environment. If the original conditions return on a regular basis, regulatory mechanisms which activate a given piece of information (e.g. a gene) specifically under the conditions where this confers a selective advantage can evolve. It is clear that regulatory networks (see above) are a result of continued evolution of such regulatory mechanisms.

Clearly, activation of genes which provide adaptation to the current conditions ultimately depends on processing the sensor input by which the current conditions are recognised. Regulatory networks are dynamical systems capable of processing such sensor data (which may be represented by pheromones as a biological example, or by input from a sensor network as a technical example). Increase of complexity of regulatory networks during evolution thus enables a living system to adapt to an increasingly wide amount of variations in environmental conditions. Thus, integrating regulatory networks into evolutionary algorithms, as explored e.g. in [18], and combining this with evolutionary mechanisms for continuously maintaining a population with complex diversity [31], e.g. by bioinspired mechanisms for mutation rate control [17] is a promising approach towards realizing large-scale, distributed sensor networks that are capable of long-term, autonomous adaptation to a complex and changing environment.

5 Conclusion

We have proposed an organic architecture for very large-scale ambient sensor networks. The overall goal is to show that with the help of biological analogies large sensor networks spanning building complexes or even cities can be made to function as a single coherent system that uses a complex, dynamically evolving world model to extract high-level meaning from the events in the environment. Towards these goals, we propose to (i) structure the sensor network in a hierarchical, parallel way that matches the structure of the world model, (ii) start with a generic system template that includes feature specifications, pre-trained classifiers and initial guesses of the probability distributions involved in the modelling of high-level processes and goals, and (iii) use methods inspired by visual information processing (salient features), cognition (attention, concept formation, abstraction and schemata learning), brain plasticity, genetics (regulatory networks), and population evolution to allow the system to adapt to the specific environment in which it is deployed, and to keep adapting as the environment changes.

Besides a continuous optimisation of parameters (as in traditional machine-learning), the envisioned adaptation includes structural changes (like the emergence of new structure corresponding to new concepts), optimisation of feature and probability distributions, and strategies for resource management.

The ideas described in this paper are the results of an initial conceptual exploration stage of our work. Next, they need to be converted into specific methods and algorithms and tested in real-life experiments. From there work towards the actual implementation of real-life sensor networks can begin. In addition, an appropriate simulation environment will be needed to be able to flexibly investigate really large scale networks.

6 References

[1] GI-Workshop on Organic Computing, in connection with 34. Jahrestagung der Gesellschaft für Informatik, Ulm, September 2004.

[2] Urs Anliker et al. A systematic approach to the design of distributed wearable systems. *IEEE Transactions on Computers*, 53(3):1017–1033, August 2004.

[3] Albert-László Barabási and Zóltan N. Oltvai. Network biology: Understanding the cell's functional organization. *Nature Reviews Genetics*, 5:101–113, 2004.

[4] Erhardt Barth, Ingo Stuke, and Cicero Mota. Analysis of motion and curvature in image sequences. In *Proc. IEEE Southwest Symp. Image Analysis and Interpretation*, pages 206–10, Santa Fe, NM, April 7–9, 2002. IEEE Computer Press.

[5] Michael Beigl, Hans-W. Gellersen, and Albrecht Schmidt. Mediacups: experience with design and use of computer-augmented everyday artefacts. *Computer Networks*, 35(4):401–409, 2001.

[6] A. Blake and A. Yuille. *Active vision.* MIT Press, Cambridge, MA, 1993.

[7] Martin Böhme, Christopher Krause, Erhardt Barth, and Thomas Martinetz. Eye movement predictions enhanced by saccade detection. In *Brain Inspired Cognitive Systems 2004*, 2004.

[8] Martin Böhme, Christopher Krause, Thomas Martinetz, and Erhardt Barth. Saliency extraction for gaze-contingent displays. In *Proceedings of the 34th GI-Jahrestagung*, volume 2, pages 646–650, 2004. Workshop on Organic Computing.

[9] L. Doherty, B.A. Warneke, B.E. Boser, and K.S.J. Pister. Energy and performance considerations for smart dust. *Int'l Journal of Parallel and Distributed Systems and Networks*, 4(3):121–133, December 2001.

[10] Deborah Estrin, David Culler, Kris Pister, and Gaurav Sukhatme. Connecting the physical world with pervasive networks. *IEEE Pervasive Computing*, 1(1):59–69, 2002.

[11] Hans-W. Gellersen, Michael Beigl, and Albrecht Schmidt. Environment-mediated mobile computing. In *Proceedings of the 1999 ACM symposium on Applied computing*, pages 416–418. ACM Press, 1999.

[12] Hans W. Gellersen, Albercht Schmidt, and Michael Beigl. Multi-sensor context-awareness in mobile devices and smart artifacts. *Mob. Netw. Appl.*, 7(5):341–351, 2002.

[13] D.E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning.* Addison-Wesley, New York, 1989.

[14] Laurent Itti and Christof Koch. Computational modelling of visual attention. *Nature Reviews Neuroscience*, 2(3):194–203, 2001.

[15] Stuart A. Kauffman. Developmental logic and its evolution. *BioEssays*, 6:82–87, 1987.

[16] N. Kern, B. Schiele, H. Junker, P. Lukowicz, and G. Tröster. Wearable sensing to annotate meeting recordings. In *Proceedings Sixth International Symposium on Wearable Computers ISWC 2002*, 2002.

[17] Jan T. Kim. Energy dependent adaptation of mutation rates in computer models of evolution. In Christoph Adami, Richard K. Belew, Hiroaki Kitano, and Charles Taylor, editors, *Artificial Life VI*, pages 248–255, Cambridge, MA, 1998. MIT Press.

[18] Jan T. Kim. LindEvol: Artificial models for natural plant evolution. *Künstliche Intelligenz*, 1/2000:26–32, 2000.

[19] Jan T. Kim. `transsys`: A generic formalism for modelling regulatory networks in morphogenesis. In Jozef Kelemen and Petr Sosík, editors, *Advances in Artificial Life (Proceedings of the 6th European Conference on Artificial Life)*, volume 2159 of *Lecture Notes in Artificial Intelligence*, pages 242–251, Berlin Heidelberg, 2001.

Springer Verlag.

[20] T. Kohonen. *Self-Organization and Associative Memory*. 1989.

[21] Tong Ihn Lee, Nicola J. Rinaldi, François Robert, Duncan T. Odom, Ziv Bar-Joseph, Georg K. Gerber, Nancy M. Hannett, Christopher T. Harbison, Craig M. Thompson, Itamar Simon, Julia Zeitlinger, Ezra G. Jennings, Heather L. Murray, D. Benjamin Gordon, Bing Ren, John J. Wyrick, Jean-Bosco Tagne, Thomas L. Volkert, Ernest Fraenkel, David K. Gifford, and Richard A. Young. Transcriptional regulatory networks in *Saccharomyces cerevisiae*. *Science*, 298:799–804, 2002.

[22] P.J.M. Havinga L.F.W. van Hoesel, L. Dal Pont. Design of an autonomous decentralized mac protocol for wireless sensor networks. In *6th International Symposium on Autonomous Decentralized Systems*, April 2003.

[23] Paul Lukowicz et al. Recognizing workshop activity using body worn microphones and accelerometers. In *Pervasive Computing: Proc. of the 2nd Int'l Conf.*, pages 18–22, 2004.

[24] T. Martinetz and K. Schulten. Topology representing networks. *Neural Networks*, 7(3):507–522, 1994.

[25] Hans Meinhardt. *Models of Biological Pattern Formation*. Academic Press, London, 1982.

[26] Luis Mendoza and Elena R. Alvarez-Buylla. Dynamics of the genetic regulatory network for *Arabidopsis thaliana* flower morphogenesis. *J.theor.Biol.*, 193:307–319, 1998.

[27] David Mumford. Mathematical theories of shape: Do they model perception. *SPIE Geometric Methods in Computer Vision*, 1570:2–10, 1991.

[28] P Perona, M Weber, and M Welling. Unsupervised learning of object classes from natural scenes. *Journal of Vision*, 1(3), 2001.

[29] Cliff Randell and Henk L. Muller. The well mannered wearable computer. *Personal Ubiquitous Comput.*, 6(1):31–36, 2002.

[30] Ingo Rechenberg. *Evolutionsstrategie '94*. Fromann-Holzboog, Stuttgart-Bad Cannstadt, 1994.

[31] Henning Schwöbbermeyer and Jan T. Kim. A comparative analysis of biodiversity measures. In Dario Floreano, Jean-Daniel Nicoud, and Francesco Mondada, editors, *Advances in Artificial Life*, Lecture Notes in Artificial Intelligence, pages 119–128, Berlin, Heidelberg, 1999. Springer Verlag.

[32] J. Sharma, A. Angelucci, and M. Sur. Induction of visual orientation modules in auditory cortex. *Nature*, 404:841–7, 2000.

[33] Thad Starner, Bernt Schiele, and Alex Pentland. Visual contextual awareness in wearable computing. In *ISWC'98*, pages 50–57, October 1998.

[34] Emmanual M. Tapia, Stephen S. Intille, and Kent Larson. Activity recognition in the home using simple and ubiquitous sensors. In *Pervasive Computing: Proc. of the 2nd Int'l Conf.*, pages 158–175, April 2004.

[35] K. Van Laerhoven and O. Cakmakci. What shall we teach our pants? In *ISWC'00*, pages 77–83, October 2000.

[36] C Zetzsche and E Barth. Fundamental limits of linear filters in the visual processing of two-dimensional signals. *Vision Research*, 30:1111–7, 1990.