

Ada: Constructing a Synthetic Organism

Kynan Eng, Andreas Bähler, Ulysses Bernardet, Mark Blanchard, Adam Briska, Jörg Conradt, Marcio Costa, Tobi Delbrück, Rodney J Douglas, Klaus Hepp, David Klein, Jonatas Manzolli¹, Matti Mintz², Thomas Netter, Fabian Roth, Ueli Rutishauser, Klaus Wassermann, Adrian M Whatley, Aaron Wittmann, Reto Wyss, Paul F M J Verschure

Institute of Neuroinformatics, University/ETH Zurich, Switzerland, www.ini.unizh.ch/~expo

¹*Nucleo Interdisciplinar de Comunicacao Sonora, Campinas, Brazil*

²*Psychobiology Research Unit, Tel-Aviv University, Tel-Aviv, Israel*

Abstract

Despite immense progress in neuroscience, we remain restricted in our ability to construct autonomous, behaving robots that match the competence of even simple animals. The barriers to the realisation of this goal include lack of knowledge of system integration issues, engineering limitations and the organisational constraints common to many research laboratories. In this paper we describe our approach to addressing these issues by constructing an artificial organism within the framework of the Ada project – a large-scale public exhibit for the Swiss Expo.02 national exhibition.

1. Introduction

In recent decades an enormous amount of data about the functions of neural systems has accumulated, spanning the range from sub-cellular processes to behavioural observations. The search for underlying principles hidden in these data has yielded a number of models of neural function, for example the famous Hodgkin-Huxley model. Canonical forms of many different neural networks have been analysed (Hopfield networks are a well-known example), and evidence for natural analogues of some of these networks has been found. Knowledge of the anatomical and functional structure of the brain is constantly increasing, in part due to the use of relatively new techniques such as neuronal track tracing and functional magnetic resonance imaging.

Despite the wealth of data on brain structure and function, little is known about the overall computational principles, let alone how they can support robust control of real-world behaving systems. Many examples of systems exist that implement a particular small subset of brain functions in a robust and biologically plausible way, and some have been interfaced to robots. However, most of these robots tend to be demonstration platforms and cannot be viewed as artificial organisms in their own right. Given that humans are able to construct complicated artefacts such as large aircraft, spacecraft and microprocessors, why has so little comparable success been achieved in robotics? The

answer may lie in three related factors: knowledge of system integration issues, engineering limitations and organisational constraints. We attempt to address these issues simultaneously within the framework of a large-scale public exhibition project called Ada.

2. The Ada Project

The public exhibit *Ada: the intelligent space* is the product of interdisciplinary research into brain function at the Institute of Neuroinformatics (INI). One of the goals of Ada is to gain an understanding of large-scale behavioural integration issues in biological organisms, through the experience of constructing an artificial organism. For the general public, Ada is intended to stimulate discussion of brain-like technologies and the social implications of their usage. “She” is an interactive entertainment space (~400 m² including auxiliary areas, see Figure 2) at the Swiss national exhibition Expo.02 in the town of Neuchâtel. Ada will function continuously for 10½ hours a day over 5 months from 15 May to 20 October 2002. It is expected that 0.5 million people will visit Ada during this time.

Conceptually, Ada can be seen as an inside-out robot with visual, audio and tactile input, and non-contact light and sound effectors. Visitors to Ada are immersed in an environment where their only significant sensory stimulation comes from Ada herself (and other visitors). Like many animals, Ada’s output is designed to have a certain level of coherence and convey an impression of a basic unitary sentience to all of her visitors. At the same time, she can communicate on an individual basis with visitors. This is achieved by using a combination of global lighting and background music to communicate overall conditions, with local light and sound effects for individual interactions.

The development of Ada commenced in early 1999 and ramped up to a final team size of over 20 people, drawn from a wide range of disciplines ranging from biological sciences through engineering to musical composition. In addition, specialists in logistics, architecture, public

relations and scenography were brought in as needed. A team this large working on a single project is still unusual in the neurosciences outside private industry, yet it is small in comparison with teams working in other fields such as high-energy physics. We predict that large-scale neuroscience projects of this sort will become more common in the future.

Creating a space rather than a conventional robot allows us to investigate whether biologically based models of control can be generalised to other task domains. In this way we satisfy the definition of general intelligence as given by Newell and Simon [11], where anything can become a task. Other advantages of using a space as opposed to a conventional robot relate to safety and public interaction. With only light and sound as effectors, Ada cannot injure her visitors and she is able to entertain much larger numbers of people simultaneously that would be possible with a conventional robot. Making Ada big, immobile and non-contact also allows for simpler engineering – all but one of the major components of Ada are made from standard off-the-shelf equipment. The problems of miniaturisation, mobility and maximising power density are effectively bypassed, leaving more resources available to concentrate on Ada's functional aspects.

3. Related Projects

There are several research projects dealing with issues related to home automation and "intelligent rooms", as well as many companies offering home automation systems. A typical example of the latter category is the GE Smart series from GE Industrial Systems [8]. These systems offer a substrate for connecting electrical devices and home network services with a common software interface. The control system software is based on rule sets or driven directly by end users, either within the building or via remote links. In this sort of system, the design emphasis is on ease of end-user installation, operation and customisation, rather than advanced behavioural functionality.

More advanced control systems exist in projects such as the Intelligent Room at MIT [10]. This aim of the Intelligent Room project is to develop systems that support human activities in a seamless, flexible way. To date, work has been done on components such as context-aware speech and gesture recognition, flexible resource allocation [7] and an agent-based extension to Java called MetaGlue. Work has also been done on biologically inspired sensory integration using two overlaid topographical maps to relate speech input with visual tracking input [3]. Ada has a similar set of functionalities, but with three main differences. Firstly, Ada is a completed product and is much larger than the Intelligent Room, in terms of physical size, number of components and degree of behavioural integration.

Secondly, the design of the user interaction with the space is immersive rather than invisible – the building does not serve its users' needs in the background, but is an active participant in their experiences. Finally, and most importantly, the space has its own goals which it actively tries to achieve by engaging with its users.

A similar project, also named the Intelligent Space, is being pursued within the Hashimoto lab at the University of Tokyo [1]. The technologies used focus on visual processing and have similarities with Ada and the MIT Intelligent Room, but with a subtle difference to both. The Hashimoto Intelligent Space is designed to be a platform to facilitate communication between the entities that inhabit it – whether they be humans, robots, or components of the space itself. The concept of a Distributed Intelligent Network Device (DIND) is proposed for connecting devices in the space. Each DIND has sensors, processing and communications components; a space is made up of two or more DINDs. In this way the space is seen not as an explicit entity like Ada, but as a common medium that enables interaction between components in a physical area.

There are also robots that, like Ada, seek to emulate the functions of organisms and interact with humans. A well-known example is the humanoid torso *Cog* [2]. The Cog project has so far dealt more with individual competencies rather than overall behaviours, such as visual-motor processing, human interaction with robot facial expressions based on an emotional model, and neural models of arm motor control. The emotional model in Cog runs on a head-only subsystem of Cog called Kismet, and has some similarities to Ada's emotional model – both contain a set of drives (goal functions) and a set of emotional states. While individual components of Cog have achieved impressive degrees of functionality, the individual behaviours have not yet been integrated into a cohesive whole.

A more similar animal-like analogue to Ada is the *Mutant* dog robot [6] and its commercially available successor *Aibo* made by Sony. Despite Aibo's small size (over 10,000 Aibos would fit inside Ada), it can be seen as Ada's closest existing relative. They are both complete systems designed to interact with the general public, and both integrate visual, audio and tactile information to produce behaviour. They both even have an internal emotional model and layered system architectures: Aibo's architecture is agent-based, while Ada's architecture goes further in terms of biological realism by using a hybrid of simulated neural networks and agent-based software components. Sony has formalised its system architecture in the OPENR model for building extensible, customisable robots [5]. The main differences between Aibo and Ada are the obvious ones of appearance and size. By looking like a dog, Aibo (and Cog, by looking like a human torso) has an inherent advantage over Ada for human interactions. A

decision made in designing Ada was to explore the limits of human interactions that could be supported without the use of pre-existing metaphors. Ada also has the possibility for individual human interaction that would require dozens of Aibos. On the engineering front, Aibo has the dual challenges of miniaturisation and minimising power consumption, whereas Ada faces power consumption constraints on a much larger scale.

4. Sensors, Effectors and Core Services

In total (including auxiliary exhibition areas), Ada has 15 video inputs, 360 tactile inputs, 9 audio input channels, 46 mechanical degrees of freedom, 16 output audio channels, 1080 floor tile light outputs (3 per tile), 30 ambient light outputs and 20 full-screen video outputs. All of these inputs and outputs can be addressed independently, giving a rich array of sensory modalities and output possibilities. Ada's multi-modal sensory inputs mimic some of the capabilities of organisms: vision, hearing and touch:

- **Vision:** Pan-tilt cameras called *gazers* are available to Ada for attentional, focused interactions with specific visitors. The cameras have on-board zoom and digital filtering capabilities.
- **Hearing:** There are clusters of three fixed microphones each in the ceiling plane, with which Ada is able to localise sound sources by triangulation. Some basic forms of sound and word recognition are available.
- **Touch:** Ada has a "skin" of 0.66 m wide hexagonal pressure-sensitive floor tiles [4] that can detect the presence of visitors by their weight. This is the only component of Ada that required a significant ground-up engineering effort. Each contains a microcontroller and sits on a serial bus running an industrial automation protocol called Interbus.

As well as sensing, Ada can also express herself and act upon her environment in the following ways:

- **Visual:** Ada uses a 360° ring of 12 LCD projectors to express her internal states visually to visitors. These projectors can show hardware accelerated 3D objects covering multiple screens, and live motion video windows that can move with smooth transitions between screens. There is also a ring of ambient lights for setting the overall visual emotional tone of the space. Local visual effects can be created using the red, green and blue coloured neon lights in each floor tile in Ada's skin.
- **Audio:** Ada is able to generate a wide range of sound effects. She expresses herself using sound and music composed in real-time on the basis of her internal states and sensory input. She can also change the pitch of her output depending on what

she hears from her visitors. The composition is generated using a system called Roboser [13].

- **Touch:** Ada has twenty 16-bit pan-tilt *light fingers* for pointing at visitors or indicating different locations in the space. They are standard theatre lights on a serial bus called DMX, which is also used to control the ambient lights and the gazers.

The core services of Ada support her higher-level behavioural functions. These services include a *tracking system* that uses information from the floor tile pressure sensors to determine the location, speed and direction of visitors. The limited resolution of the floor tiles means that it is not possible to distinguish individual paths in all situations, so in some cases Ada will only know about the presence of groups of people at particular locations. To obtain more information about individual visitors, a *vision system* deploys gazers to collect images of people who have been localised on the floor. The *audio system* localises and recognises basic sounds (such as the word "Ada") to help in identifying salient individuals. On the output side, the *Roboser* audio system composes real-time music and sound effects, a *video server* supports the visualisation of saved and live images, and a *DMX server* controls the light fingers, gazers and ambient lights.

5. Behaviours and Interactions

The degree of success with which visitors can be convinced that Ada is an artificial organism depends strongly on the nature of their interactions. The operation of the space needs to be coherent, real-time, and reliable enough to work for extended periods. As well as this, it must be understandable to visitors (there are several demo modules at the entrance, which play an important role in initial visitor priming) and sufficiently rich in the depth of interactions so that visitors feel the presence of a basic unitary intelligence

To provide for a natural progression in visitor interaction, Ada incorporates at least four basic behavioural functions. First, she can *track* individual visitors or groups of visitors, possibly (but not necessarily) giving them an indication that they are being tracked. At the same time, she can *identify* those visitors who are more "interesting" than others because of their responsiveness to simple cues that Ada uses to probe their reactions. These people are encouraged to form a *group* in part of the space through the use of various light and sound cues. When the conditions are appropriate, Ada rewards a group of visitors by *playing* one of a number of games with them. She continuously evaluates the results of her actions and expresses emotional states accordingly, and tries to regulate the distribution and flow of visitors. These four behavioural functions are decomposed into smaller behaviours that call on the core services as needed.

Table 1: Descriptions of types of software, data storage and learning found at different levels in Ada

Level	Functionality	Software	Learned/stored data	Comments
4: Behavioural modulation	Goal function evaluation Behaviour mode selection Emotional model	Simulated neurons (IQR421*)		Most of the computation at this level has direct biological analogues
3: Behavioural modules	Coordinated high-level interactions	Simulated neurons Software agents	Synapse weights Agent action rules Fuzzy logic rule sets	Agents could be conceived as “meta-neurons” that autonomously collect process their inputs to produce outputs
2: Sensorimotor processes	Filtering of raw input data	Procedural or object-oriented code	Raw data adaptation Output data smoothing parameters Object-based data for complex information	Some of the input data adaptation is similar to neural behaviour, but the bulk of the code at this level is procedural
1: Device I/O drivers	Interface to hardware	Procedural or object-oriented code	Input data parameters Output data parameters Physical configuration	
0: Hardware devices	Motor control Sound production Sensor reading Light setting	On-device logic	I/O tuning parameters, typically set during calibration	Some devices have sophisticated abilities, but these are largely opaque to system developers

* IQR421 is a neural simulation software package developed at the Institute of Neuroinformatics

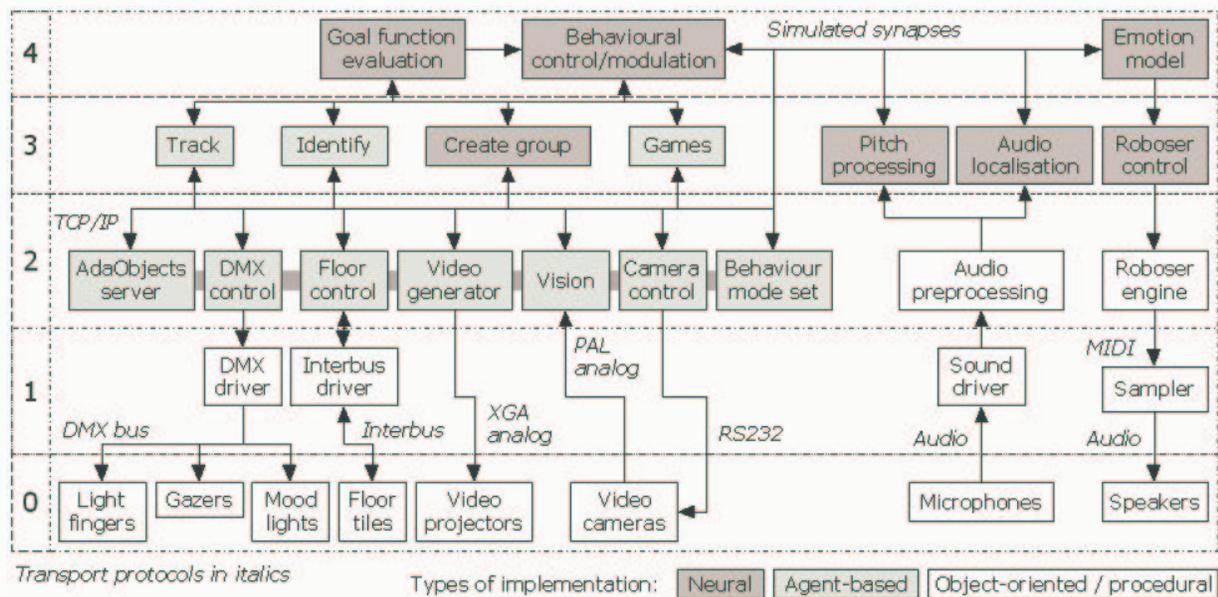


Figure 1: Overview of Ada system architecture, organised into conceptual layers.

6. System Architecture

Ada is not intended to be a brain-like system at the level of modelling every individual neural element. While this might be a desirable long-term goal to allow ultra-

low power implementations using technologies such as analog VLSI chips, the technology and design tools are not sufficiently mature at the moment. Instead, the aim is to provide a hybrid architecture that satisfies all project-related constraints, while allowing maximum

flexibility of implementation and the investigation of system-level integration issues [12].

The architecture of Ada can be roughly sketched out as a series of levels (see Figure 1), with a gradient of decreasing biological plausibility as the proportion of traditional procedural code increases. The types of data stored or “learned” at each level also follows a similar gradient of biological plausibility. Each level contains modules that communicate with other modules in the same layer, as well as with modules in adjacent layers. The metaphor being used here is that of distributed brain-like computation, characterised by tight coupling within individual modules, and loose coupling between modules. The underlying software is a mixture of simulated neural networks, agent-based systems and conventional procedural or object-oriented software. The types of computations and data storage at each of the different levels are summarised in Table 1.

Different communication protocols are used to connect the components of the system. The simulated neural networks use a specialised socket-based protocol, while an asynchronous message-based middleware is used to exchange data between agents. The middleware automatically distributes messages to agents on a subscription basis. Data exchanged in this way includes floor data (system input), visitor tracking data (internal), behavioural states (internal) and DMX device control (system output).

7. System Goals and Action Selection

Conceptually, Ada exists as an artificial organism that tries to maximise its own goal functions, which we interpret as her “happiness”. This means that the system as a whole must implicitly or explicitly compute its own level of happiness, which can then be used to determine if certain actions contribute to this goal. As a first approximation we can write:

$$H = f(g_s, g_r, g_i)$$

H = overall goal or “happiness”

g_s = survival

g_r = recognition

g_i = interaction

Survival is a measure of how well Ada satisfies her basic requirements, which are to maintain a certain flow of visitors over time and to keep these people moving with a certain average speed. *Recognition* quantifies how well Ada has been able to track and collect data about individual people or groups of people, as a precondition for more advanced interactions. This process can be seen as Ada “carving” objects out of the world of her sensory data, which is implemented as a progressive filtering of the sensory data and the creation of objects in an internal database once certain criteria of persistence and coherence have been satisfied.

Interaction measures the number of successful human interactions that Ada has been involved in, with more complex interactions such as games being weighted more highly.

As a system, Ada has the goal of maximising the value of H . There are multiple strategies for achieving this: for example, Ada could encourage high visitor throughput, but in doing so have very few possibilities for recognition and interaction (g_s high, g_r and g_i low). Alternatively, Ada could also achieve an equivalent value of H with only a few visitors in the space, but with high recognition and interaction with each visitor (g_s low, g_r and g_i high). The actual computation of H occurs over multiple levels: an explicit top-level calculation is done using simulated neurons, and in parallel individual behaviours also calculate their own contributions to the parameters for H .

The results of the H calculation are combined with other high-level inputs to select the most appropriate behavioural state for Ada at any point in time. Behaviour selection occurs at multiple levels – for example, the floor tiles display colours that depend on the local effects in use as well as the overall state of the space. At the top level, the neural modulation scheme is used to activate and inhibit subgroups of the underlying behaviours. This modulation can take a variety of forms, including a “hard” winner-take-all (WTA) scheme, a “softer” multiple-winner WTA, or a scheme where the behaviours run completely freely. The extent to which the behaviour modulation needs to be “hard” depends on the subjective evaluation of how the behaviours interact and/or interfere with each other. Because of the constraints imposed by a high visitor flow rate, the behavioural control is normally run in a “hard” mode when the exhibit is very busy.

8. Computational Infrastructure

Ada runs on a 100 Mbit network of about 30 PCs (AMD Athlon XP, 0.5-1.0 Gb RAM) running Linux, including peripheral parts of the exhibit such as the entry corridor displays. Because of the highly bursty network traffic, the cluster is partitioned into several subnets. Specialised driver cards are used for DMX and Interbus communications. In addition, about 24 frame grabbers and 4 sound cards are installed. Four laptops on a wireless LAN enable system testing and tuning to occur while walking around in the main space.

9. Integration & Testing

Since 1998 a number of increasingly large public tests were run to evaluate the feasibility and scalability of the underlying technologies, gauge visitor impressions, and test different interaction scenarios. One of these tests was at the Zurich Festival of Science (Zurich main

station, May 2001, 100,000 visitors), where a system called Gulliver was run for three days in collaboration with the Remote Sensing Laboratories in the Department of Geography of the University of Zurich. Gulliver contained light fingers that could follow visitors, and a floor used as a large collective joystick to control a simulated flight over Switzerland. Zürifäscht, the triennial Zurich city festival in July 2001, provided the setting for the next system test. This time a more Ada-like system called Gulliver II was deployed over three days, including a raised area from where spectators could observe the space, and some displays showing the internal operations of Gulliver. Some basic Ada functionalities were tested, such as visitor tracking with floor tiles and light fingers, displaying visitor trajectories on the floor, sound localisation of handclap noises and a group football game. The two key issues that stood out from the results of the tests were the need for effective visitor flow control, and the importance of communicating Ada's intentions clearly through the use of effective cues and visitor pre-conditioning sequences.



Figure 2: A typical live user interaction scene within Ada. Visible are floor tiles, a light finger highlighting a visitor (centre left), a dynamic 3D visualisation (top) and a live gazer video on the screens (top left).

10. Outlook

Ada is one of the first real-world systems to attempt to replicate brain-like functions on a large scale, in terms of physical size, number of sensors and effectors, and animal-like behavioural integration. She is a convergence of multiple interests from many different parts of society, most of whom are not neuroscience experts. Nevertheless, they have an equally legitimate interest in the project since they are collectively paying for it, and the technologies Ada is based on will directly affect their future. Ada is intended to be a public statement of the state of the art in real-world autonomous system development, and a stepping stone on the way to more effective systems. It is hoped that

she will be used as a benchmark against which future systems can be compared. This began as soon as the exhibit opened, when real-world feedback became available and ongoing system upgrades commenced.

11. Acknowledgments

Ada is supported by organisations including the ETH/University of Zurich, Expo.02, Manor AG, the Velux Stiftung and the Gebert Rüt Stiftung.

12. References

- [1] Appenzeller, G, Lee, J-H and Hashimoto, H, "Building Topological Maps by Looking at People: An Example of Cooperation between Intelligent Spaces and Robots", *Proceedings of the International Conference on Intelligent Robots and Systems (IROS 1997)*
- [2] Brooks, R A, Breazeal, C, Marjanovic, M, Scassellati, B and Williamson, M M, "The Cog Project: Building a Humanoid Robot" (1999), *Computation for Metaphors, Analogy, and Agents*, LNCS 1562, pp 52-87, Springer-Verlag
- [3] Coen, M, "Multimodal Integration – A Biological View" (2001), *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2001)*, Seattle, Washington
- [4] Delbrück, T, Douglas, R J, Marchal, P, Verschure, P and Whitley, A M, "A device for controlling a physical system", EU patent application 99120136.9-2215 (1999)
- [5] Fujita, M and Kageyama, K, "An Open Architecture for Robot Entertainment" (1997), *Proceedings of the First International Conference on Autonomous Agents*, pp 435-442, ACM Press
- [6] Fujita, M and Kitano, H, "Development of an Autonomous Quadruped Robot for Robot Entertainment" (1998), *Autonomous Robots* 5: 7-20
- [7] Gajos, K, Weisman, L and Shrobe, H, "Design Principles For Resource Management For Intelligent Spaces", *Proceedings of the Second International Workshop on Self-Adaptive Software (IWSAS'01)*, Budapest, Hungary
- [8] GE Industrial Systems, GE Smart web page, www.ge-smart.com
- [9] Lee, J-H & Hashimoto, H, "Intelligent Space – Its concept and contents" (2002), *Advanced Robotics* Vol. 16 No. 3: 265-280
- [10] MIT Intelligent Room project web page, www.ai.mit.edu/projects/iroom
- [11] Newell, A, *Unified Theories of Cognition* (1990), Harvard University Press
- [12] Voegtlin, T and Verschure, P, "What Can Robots Tell Us About Brains? A Synthetic Approach Towards the Study of Learning and Problem Solving" (1998), *Reviews in the Neurosciences* 10: 3-4, 291-310
- [13] Wasserman, K C, Blanchard, M, Bernadet, U, Manzolli, J M and Verschure, P F M J (2000), "Roboser: An Autonomous Interactive Composition System", In: I Zannos (Ed.), *Proceedings of the International Computer Music Conference (ICMC)*, pp. 531-534. San Francisco, CA, USA: The International Computer Music Association.