


50th Anniversary Summit of Artificial Intelligence

Summit Proceedings

Monte Verita, Ascona, Switzerland
July 9-14, 2006



50th Anniversary Summit of Artificial Intelligence

Monte Verita, Ascona, Switzerland

July 9-14, 2006

Summit Proceedings

Conference Chair

Rolf Pfeifer, University of Zurich, Switzerland

Program Co-Chairs

Max Lungarella, University of Tokyo, Japan

Fumiya Iida, University of Zurich, Switzerland

Josh Bongard, Cornell University, USA

Program Committee

Rodney Brooks, Massachusetts Institute of Technology, USA

Yasuo Kuniyoshi, University of Tokyo, Japan

Linda Smith, Indiana University, USA

Olaf Sporns, Indiana University, USA

Luc Steels, Sony Computer Science Laboratory Paris, France

Acknowledgment

This event has only been possible thanks to the generous sponsorship of a number of companies and institutions.

SIEMENS

Main Sponsor



MIGROS
kulturprozent



Advisor for PR and sponsor relations

The 50th Anniversary Summit of Artificial Intelligence

– Artificial Intelligence in the 21st century –

Rolf Pfeifer¹, Max Lungarella², Fumiya Iida¹, Josh Bongard³

¹Artificial Intelligence Laboratory, University of Zurich

²Laboratory for Intelligent Systems and Informatics, University of Tokyo

³Computational Synthesis Lab, Cornell University

The discipline of Artificial Intelligence (AI) was “born” in the summer of 1956 in Dartmouth in Hanover, New Hampshire. Half a century has passed and AI – coming a long way since its inception – has turned into an important field whose influence on our daily lives can hardly be overestimated. Many specialized AI systems exist that are at work in our cars, in our laptop computers, and in our personal and commercial technologies. Undoubtedly, in the future the impact of AI on our lives is poised to increase, blurring even more the already fuzzy boundary between man and machine.

Past, present, and future: The first important goal of this workshop is to celebrate AI’s 50th anniversary by reflecting on its history and development, assessing the state of the art, and speculating about the future of the field. Despite the significant advances of AI in the last 50 years – recall IBM’s Deep Blue computer beating the chess world champion Garry Kasparov in a landmark victory in 1997 at a tournament in New York – it is clear that the original challenges set by the first generation of AI visionaries have yet to be met. Not only is natural intelligence far from being understood and artificial forms of intelligence still so primitive compared to natural ones, but seemingly simple tasks like categorization, recognition, and manipulation of objects – which are “easy” for a 3-year-old – remain to be realized artificially. A look at the current research landscapes of psychology and neuroscience reveals the paucity of knowledge about how biological brains achieve their remarkable functionalities, how these functionalities develop in the child, and how they have arisen in the course of evolution. Also, we do not have a good understanding of the cultural and social processes that have helped shape human intelligence. Moreover, because basic theories of natural intelligence are lacking and – despite impressive advances – the required technologies for building sophisticated artificial systems are still not available, the capabilities of current robots fall far short of the intelligence of even very simple animals. At the workshop we will discuss the potential reasons for this unsatisfactory situation. One hypothesis that many researchers have been pursuing in recent years is that this might be due to the strict adherence to the computational paradigm – “cognition as

computation” – and the neglect of embodiment and the interaction with the physical and social world. Other hypotheses concern the lack of computational power and the insufficient incorporation of formal methods. However that may be, an analysis of the discipline reveals that in a large part of the research community a clear paradigm shift is under way – from a purely computational view to one of *embodiment*. In this view, intelligent behavior is not merely the result of computational processes, but emerges from the interaction of brain processes, morphology, material properties, and interaction with the environment.

Interdisciplinarity and cross-fertilization: Embodiment implies that we need to consider all aspects of an organism – brain, body, system-environment interaction – which in turn means that researchers from many disciplines need to participate in the adventure of unraveling the processes underlying intelligent behavior. The second main objective of the workshop is therefore to bring together not only computer scientists, linguists, and psychologists, but also biologists, neuroscientists, engineers, roboticists, material scientists, as well as researchers of dynamical systems and biomechanics. It is our conviction that breakthroughs can only be achieved through a strong cross-fertilization among these fields and by initiating and fostering cooperation between groups from different disciplines.

Goals of AI research: There are three kinds of goals associated with modern AI research – understanding biological systems, abstracting key principles of intelligent behavior, and developing practical applications. The first goal, understanding animals and humans, can be tackled using the synthetic methodology, which can be characterized by the slogan “understanding by building.” Since its early days, this has been the standard approach of AI: You are interested in some phenomenon, say, how humans recognize a face in a crowd or how ants find their way back to the nest after a foraging trip, and you try to understand how this comes about by building an artifact – a robot or a computer program – that mimics certain aspects of this phenomenon. This method has proved extremely powerful. Of course, this step requires the close interdisciplinary cooperation of biologists, neuroscientists, and engineers. Next, principles need to be abstracted so that the insights can also be applied to artificial systems – the question here is what has been learned. This can be viewed as the first traces of a “theory of intelligence.” Finally, the insights can be applied to the design of useful applications. In traditional applications, human intelligence has frequently supplied the motivation for the research, but then the problem was often solved without trying to mimic the biological system, as for example in most approaches to machine learning. In order to be successful an application often does not need to slavishly copy from nature, as best illustrated by IBM’s Deep Blue victory. However, if we are interested in adaptive systems in the real world – and a lot of recent

research in AI is going in this direction – nature can be a great source of inspiration. Again, interdisciplinary cooperation will be a key to technological progress and scientific breakthroughs.

Broad impact: With the highly varied background of the participants and the grand challenges and issues AI addresses, we hope that the impact of the workshop and of the field in general will go far beyond the scientific and engineering discipline of AI proper. There is no doubt that the concepts and fascination developed in AI have already reached society at large, including business, art, entertainment, and the media. This is the reason why we have researchers as well as people from other walks of life who have provided valuable input to the development of the field, such as businesspeople, artists, and journalists. We are convinced that the outcome will not only help us sketch out the future directions of AI research, but also understand how deep the paradigm change represented by embodiment in fact goes. The pertinent ideas will be incorporated into a publication emerging from the workshop (proceedings, book, or handbook), which will form a comprehensive collection of opinions and views.

Program: The various goals are reflected in the structure of the conference program. Because the discussion of core issues is an essential components there will be, in addition to formal lectures and poster sessions, a series of panel discussions on a number of topics. The first panel entitled the “The future of AI – classical or embodied” raises the topic of a paradigm shift, of an “embodied turn”, so to speak. The panel on “Advertising AI to the public and to companies – strategies and methods” will discuss strategies of how to commercialize ideas from AI research. The session on “The new landscape of artificial intelligence – the impact of other research areas” focuses on the new landscape of AI and its relations to other disciplines normally not directly associated with AI such as neuroscience, bionics, biomechanics, and material science, but that do make significant contributions to the field. “Modern AI: beyond ‘cognition as computation’?” asks the question of whether an extension of the notion of computation will be required in order to understand embodied forms of intelligence. And finally, at the level of policy and science management, the subject of funding is of course crucial to the development of AI. The final panel, “Funding AI research”, will discuss pertinent policies and strategies of different funding agencies for AI research.

The first day of the Summit is dedicated to the first objective, to reflecting on history and prospects and to assessing the impact of the field on society at large. The second day will focus on cross-fertilization, that is, on the relation between the various scientific and non-scientific disciplines, e.g. engineering, psychology, neuroscience and neural interfacing, bionics,

biomechanics, and art. The third day, entitled “The insider view of AI” features prominent AI researchers presenting their perspective on the state-of-the art. The fourth day is dedicated to the presentation of important research projects, demonstrating the breadth of today’s field of AI research. Finally, the goal of the fifth and last day is to come up with a synthesis of the ideas generated during the meeting.

In order to reach not only scientists, but also the public, schools, and children, we plan an event where we will explain the basic ideas of AI in non-technical language on the 12th of July. This will include a demonstration of some of the most advanced robots in the world.

Acknowledgements: This event has only been possible thanks to the generous sponsorship of a number of companies and institutions: Siemens (main sponsor), Centro Stefano Franscini, Neuronics, Migros Kulturprozent, Swisscom, and the AI Lab (Department of Informatics, University of Zurich). We would also like to thank matek for their competent advice on PR and sponsor relations.

We hope you will enjoy the conference!

Contents

Conference Program

Speaker Abstracts

Poster Abstracts

Robot Demonstrations

50th Anniversary Summit of Artificial Intelligence



Conference Program

Conference Program

Sunday, July 9

Opening

- 18:00 Welcome reception and welcome address
20:00 – 22:30 World Cup Final

Monday, July 10

- 07:30 Breakfast

History and Prospects of AI

- 08:30 – 09:40 Nils Nilsson – “Routes to the AI Summit”
09:40 – 10:50 Rodney Brooks – “AI: Where from and where to?”
Break (10 min.)
11:00 – 12:10 Panel discussion “The future of AI – Classical or embodied or both?”
(Rolf Pfeifer (M), Seth Bullock, Inman Harvey, Owen Holland, Nils Nilsson, Luc Steels)
12:15 Lunch

Society at Large

- 14:05 – 15:05 Hod Lipson – “Curious and creative machines”
15:05 – 16:05 Cynthia Breazeal – “Socially intelligent robots”
Break (15 min.)
16:20 – 17:20 Adrienne Wortzel – “Archipelago.ch: The dynamic diorama”
17:20 – 18:40 Panel discussion “Advertizing AI to the public and to companies – strategies and methods”
(Rodney Brooks (M), Rudiger Dillmann, Luca Gambardella, Simon Grand, Takanori Shibata)
18:40 – 19:30 Robot demonstration
(Fumiya Iida and Peter Revai)
19:35 Aperio Ticinese

Tuesday, July 11

07:30 Breakfast

Modern AI: Cross-fertilization at Work – Part 1

08:30 – 09:40 Rodney Douglas – “Computations performed by collections of recurrently connected neurons, and their implementation in hybrid VLSI electronic systems”

09:40 – 10:50 Andrew Schwartz – “Useful signals from motor cortex”

Break (10 min.)

11:00 – 12:10 Peter Fromherz – “Biophysical studies on brain-computer interfacing”

12:15 Lunch

Modern AI: Cross-fertilization at Work – Part 2

13:50 – 15:00 Rudolf Bannasch – “Morphological intelligence in bionic applications”

15:00 – 16:10 Metin Sitti – “Biologically inspired miniature robots”

Break (20 min.)

16:30 – 17:40 Michael Dickinson – “The neural control of aerodynamics in fruit flies”

17:40 – 18:50 Kevin O'Regan – “Consciousness”

Break (10 min.)

19:00 – 20:00 Panel discussion “The new landscape of artificial intelligence – the impact of other research areas”

(Olaf Sporns (M), Rudolf Bannasch, Rodney Douglas, Lukas Lichtensteiger, Steve Potter)

20:00 Conference banquet

Wednesday, July 12

07:30 Breakfast

The Insider View of AI – Part 1

08:30 – 09:40 Alex Waibel – “Intelligent systems as mediators in human communication”

09:40 – 10:50 Sebastian Thrun – “The future of driving”

Break (10 min.)

11:00 – 12:10 Luc Steels – “Artificial intelligence and the origins of symbolic culture”

12:15 Lunch

The Insider View of AI – Part 2

13:30 – 14:40 Relaxed poster session 1 (accompanied by coffee, tea, and dessert)

14:40 – 15:40 Inman Harvey – “Get a life”

Break (20 min.)

16:00 – 17:00 Dario Floreano – “Design principles for emergent cooperation and communication in robotic swarms”

17:00 – 18:10 Panel discussion “Modern AI: Beyond 'cognition as computation'?”
(Rolf Pfeifer (M), Ezequiel di Paolo, Frederic Kaplan, Daniel Polani, Sebastian Thrun)

18:10 – 20:00 Local event + Robot demonstration
(Dario Floreano)

20:00 Dinner

Thursday, July 13

07:30 Breakfast

Project Day – Part 1

08:30 – 09:30 Giulio Sandini, Giorgio Metta, David Vernon – “RobotCub - Sharing a body for the advancement of AI”

09:30 – 10:00 Claes von Hofsten – “The development of gaze control in human infants”

10:00 – 10:30 Kerstin Dautenhahn – “Why social intelligence matters in the design and development of intelligent robots”

Break (10 min.)

10:40 – 11:10 Aude Billard – “From HI to AI: Transmitting human skills and knowledge to robots”

11:10 – 11:40 Chrystopher Nehaniv – “Sensorimotor experience, information and development”

11:40 – 12:10 Auke Ijspeert – “Adaptive locomotion in robots with multiple degrees of freedom”

12:15 Lunch

Project Day – Part 2

13:30 – 14:40 Relaxed poster session 2 (accompanied by coffee, tea, and dessert)

14:40 – 15:15 Norman Packard – “Programmable artificial cells”

15:15 – 15:50 Roland Siegwart – “The rise of robots – machines sharing the environment with natural creatures”

Break (10 min.)

16:00 – 16:35 Akio Ishiguro – “Mobiligence: The emergence of adaptive motor function through interaction among the body, brain, and environment”

16:35 – 17:10 Rudiger Dillmann – “Emergent cognitive capabilities for humanoids: Robots learning senso-motor skills and task knowledge from multimodal observation of humans”

Break (10 min.)

17:20 – 17:55 Raja Chatila – “Cogniron: The cognitive robot companion”

17:55 – 18:30 Koh Hosoda – “Synergistic intelligence: A cognitive developmental approach towards emergence of communication”

Break (10 min.)

18:40 – 19:50 Panel discussion “Funding AI research: Large international projects? Alternatives?”

(Giulio Sandini (M), Raja Chatila, Colette Maloney, Rolf Pfeifer, Norman Packard, Alex Waibel; accompanied by drinks, smooth transition to dinner)

20:00 Dinner

Friday, July 14

07:30 Breakfast and packing

Grand Finale – Synthesis and Future Work

08:50 – 10:50 Open Discussion – Synthesis, future work and perspectives

12:15 Lunch

Farewell... see you in 50 years!

50th Anniversary Summit of Artificial Intelligence

A stylized, layered mountain range graphic in shades of gray, positioned to the right of the main title.

Speaker Abstracts

Speaker Abstracts

Synergistic Intelligence Project: A Cognitive Developmental Approach Towards Emergence of Communication

Minoru Asada, Koh Hosoda, Yasuo Kuniyoshi, Hiroshi Ishiguro, Toshio Inui

From HI to AI: Transmitting Human Skills and Knowledge to Robots

Aude Billard

Socially Intelligent Robots

Cynthia Breazeal

AI: Where From and Where To?

Rodney Brooks

Cogniron: The Cognitive Robot Companion

Raja Chatila

Why Social Intelligence Matters in the Design and Development of Intelligent Robots

Kerstin Dautenhahn

The Neural Control of Aerodynamics in Fruit Flies

Michael H. Dickinson

Emergent Cognitive Capabilities for Humanoids: Robots Learning Senso-Motor Skills and Task Knowledge from Multimodal Observation of Humans

Rüdiger Dillmann

Computations Performed by Collections of Recurrently Connected Neurons, and Their Implementation in Hybrid VLSI Electronic Systems

Rodney Douglas

Biophysical Studies on Brain-Computer Interfacing

Peter Fromherz

Get A Life

Inman Harvey

Mobiligence Project: Emergence of Adaptive Motor Function Through Interaction Among the Body, Brain, and Environment

Koh Hosoda, Akio Ishiguro, Hajime Asama

Adaptive Locomotion in Robots with Multiple Degrees of Freedom

Auke Ijspeert

Curious and Creative Machines

Hod Lipson

Sensorimotor Experience, Information, and Development

Chrystopher L. Nehaniv

Routes to the Summit
Nils J. Nilsson

Consciousness
J. Kevin O'Regan

Artificial Cell Engineering as a Computational Problem
Norman Packard

The RobotCub Project: An Open Humanoid Framework for Research on Embodied Cognition
Giulio Sandini, Giorgio Metta, David Vernon

Useful Signals from Motor Cortex
Andrew Schwartz

The Rise of Robots – Machines Sharing the Environment with Natural Creatures
Roland Siegwart

Biologically Inspired Miniature Robots
Metin Sitti

Artificial Intelligence and the Origins of Symbolic Culture
Luc Steels

The Future of Driving
Sebastian Thrun

The Development of Gaze Control in Human Infants
Claes von Hofsten, Kerstin Rosander

Intelligent Systems as Mediators in Human Communication
Alex Waibel

Archipelago.ch: The Dynamic Diorama
Adrienne Wortzel

Synergistic Intelligence Project: A Cognitive Developmental Approach Towards Emergence of Communication

Minoru Asada*, Koh Hosoda*, Yasuo Kuniyoshi[†], Hiroshi Ishiguro*, Toshio Inui^{††}

JST ERATO ASADA Synergistic Intelligence Project

*Osaka University, [†]University of Tokyo, ^{††}Kyoto University

[asada | hosoda | kuniyoshi | ishiguro | inui]@jeap.org <http://www.jeap.org>

The recent advances in information and robot technology (in short, IRT) enable us to build humanoid robots that have a large number of DoFs and various kinds of many sensors. Human-like motions are realized on these robots, and the shape and appearance have become closer and closer to us.

However, the current robotics lacks the faculties of language communication with ordinary people and of intelligent behavior generation in various situations such as at home. The fundamental relationship between humans and robots would become more important since these robots would be introduced into our lives in near future, and therefore the mechanisms of adaptation and development of both humans and robots should be taken into account in order to find the correct direction of the technology in future.

"Synergistic Intelligence" (hereafter, SI), the title of JST (Japan Science and Technology Agency) ERATO (Exploratory Research for Advanced Technology) Asada project, emerges intelligent behaviours through the interaction with environment including humans. Synergistic effects with brain science, neuroscience, cognitive science, and developmental psychology are expected. SI is one approach to a new discipline called "Humanoid Science" that aims at providing a new way of understanding ourselves and a new design theory of humanoids through mutual feedback between the design of human-like robots and human-related science.



"Humanoid Science" under which a variety of researchers from robotics, AI, brain science, cognitive science, psychology and so on are seeking for new understanding of ourselves by constructivist approaches, that is expected to produce many applications.

SI adopts a methodology called "Cognitive Developmental Robotics" (hereafter, CDR)¹ that consists of the design of self-developing structures inside the robot's brain, and the environmental design: how to set up the environment so that the robots embedded therein can gradually adapt themselves to more complex tasks in more dynamic situations. Unstructured terrains are opponents for adaptive walkers to negotiate with in order to generate dynamic motions. The caregiver's behaviour to a

robot is one environmental design issue since parents, teachers, and other adults adapt themselves to the needs of children according to each child's level of maturity and the particular relationship they have developed with that child.

One of the most formidable issues in SI is "Nature vs. Nurture": to what extent should we embed the structure, and to what extent should we expect the development triggered by the environment? A symbolic issue is "Language Acquisition." How can robots emerge the symbol in the social context? What is the essential element in this process?

This paper presents an introduction of SI and its preliminary studies on emergence of communication. The project aims at building cognitive developmental artificial agents (humanoids), understanding natural agents (humans), and their mutual feedbacks (see Figure). The project consists of four groups: (1) Physio-SI: whole body dynamic motions such as walking, running, and jumping, (2) Perso-SI: cognitive developmental robotics including body image, imitation, and language communication, (3) Socio-SI: emergence of communication and society by androids, and (4) SI-mechanism: neuroscientific supports for Physio, Perso, and Socio-SIs.

References

1. M. Asada, K. F. MacDorman, H. Ishiguro, and Y. Kuniyoshi. Cognitive developmental robotics as a new paradigm for the design of humanoid robots. *Robotics and Autonomous System*, 37:185–193, 2001.

From HI to AI: Transmitting human skills and knowledge to robots

Aude Billard

Learning Algorithms and Systems Laboratory
École Polytechnique Fédérale de Lausanne, Switzerland
aude.billard@epfl.ch

In this talk, I will report on our recent progresses to developing natural means of transmitting human knowledge about tasks and skills to robots. This work exploits various means of human-machine interactions, and, in particular, the ability to imitate.

Up to now, providing robots with the ability to imitate even simple gestures was sufficiently complex to keep the community concentrated on developing purely reflexive imitation capabilities. Very recently, the community has started to realize the importance to provide robots with the ability to interpret the user's intention and predict the user's actions.

In our work, we have progressively added complexity to our algorithms for imitation learning, taking inspiration in the various stages of imitation learning in children: starting from reflexive imitation of body motions and building up the way to informed and selective imitation of goal-directed tasks. A core issue of our recent work stresses the idea of endowing robots with the ability to be selective during the imitation.

Socially Intelligent Robots

Cynthia Breazeal
The Media Lab
MIT, Cambridge, MA, USA
cynthiab@media.mit.edu

What is a socially intelligent robot? For me, the ultimate vision of a socially intelligent robot is one that is able to communicate and interact with us, understand and even relate to us, in a personal way. It is able to engage people as a full-fledge partner and participates in fundamental forms of human interaction such as collaborative teamwork, social support, and social learning or teaching.

Human social intelligence is certainly the most advanced example, but the level of social intelligence exhibited by companion animals would also find many pragmatic uses for robots [Fong et.al. 2003]. Unlike the original goal of Artificial Intelligence, which is to create a technological system with human equivalent intelligence, my goal is to create robots that are human-synergistic and human-compatible. Specifically, robots should bring value to us and are valued by us because they are different from us in ways that enhance and complement our strengths and abilities. I argue that the goal of creating robots that can engage us as full-fledges partners is as challenging and deep as the original goal of Artificial Intelligence because it requires scientists of the natural and artificial to deeply understand human intelligence, behavior, and nature across multiple dimensions (i.e., cognitive, affective, physical, and social) in order to design synthetic systems that support and complement people and our goals. It also requires scientists to understand the dynamics of the flesh-and-blood-human-with-robot system. Theories about disembodied minds or even embodied minds operating in isolation fall far short of this goal.

Machines today simply do not understand “people as people.” For instance, by and large, robots treat us either as other objects in the environment, or at best they interact with us in a manner characteristic of socially impaired people. For instance, robots are not really aware of our goals and intentions. As a result, they don’t know how to appropriately adjust their behavior to help us as our goals and needs change. They generally do not flexibly draw their attention to what we currently find of interest so that their behavior can be coordinated and information can be focused about the same thing. They do not realize that perceiving a given situation from different perspectives impacts what we know and believe to be true about it. Consequently, they do not bring important information to our attention that is not easily accessible to us when we need it. They are not deeply aware of our emotions, feelings, or attitudes. As a result they cannot prioritize what is the most important to do for us according to what pleases us or to what we find to be most urgent, relevant, or significant. They do not readily learn new skills and abilities from interacting with or observing people. As a result, they cannot take advantage of the tremendously rich learning environment of humans. These shortcomings must be addressed for robots to achieve their full beneficial potential for us in human society (and vice versa).

Promisingly, there have been initial and ongoing strides in all of these areas [e.g., Breazeal, 2002; Scassellati, 2000 and for reviews see Picard, 1997; Fong et al, 2003; Schaal, 2000]. In particular, in my own group we have been steadily working to endow

robots with socio-cognitive skills (and evaluating their impact in human subject studies) on a wide range of human-robot interactions such as collaborative teamwork and social learning. We have been developing an architecture based on embodied cognition theories from psychology [e.g., Barsalou, 2003; Sebanz et. al., 2006] to give our humanoid robot visual and mental perspective taking abilities using a simulation theoretic framework. Specifically, *Simulation Theory* holds that certain parts of the brain have dual use; they are used to not only generate our own behavior and mental states, but also to predict and infer the same in others. To understand another person's mental process, we use our own similar brain structure to simulate the introceptive states of the other person. This is the process by which the robot infers its human partner's goals, attention, beliefs, and affect from observable behavior.

Within a teamwork task, the robot is able to compare and reason about how these human internal states relate to its own in order to provide the person with informational support and instrumental support [Gray et. al, 2005]. For example, in the case of informational support, the robot can relate its own beliefs about the state of the shared workspace to those of the human based on the visual perspective of each. If a visual occlusion prevents the human from knowing important information about that region of the workspace, the robot knows to direct the human's attention to bring that information into common ground. Furthermore, based on principles of Joint Intention Theory, the robot uses a versatile range of non-verbal behaviors to coordinate teamwork and establish and maintain mutual beliefs about progress in the task [Hoffman & Breazeal, 2004]. In the case of instrumental support, the robot can infer the human's intent (e.g. a desired effect on the workspace) from observing their behavior. If the human fails to achieve their intent, the robot can reason about how it might best help the human achieve their goal either by achieving that goal for them or by providing mutual support that helps the human achieve his or her goal. The representations by which the robot reasons and plans is inspired by embodied cognition theories [e.g, Barsalou, 2003; Sebanz et. al., 2006].

Within a social learning context, the robot uses its perspective taking abilities to interpret the intent behind the human's demonstrations [Breazeal et. al., 2006]. Imagine a scenario where the demonstrations are provided a person who does not have expertise in the learning algorithms used by the robot. As a result, the teacher may provide sensible demonstrations from a human's perspective; however, these same demonstrations may be insufficient, incomplete, ambiguous, or otherwise "flawed" from the perspective of providing a correct and sufficiently complete training set needed by the learning algorithm to generalize properly. We have tackled this issue by designing the robot to be a socially cognitive learner in a tutelage-based scenario. As the robot observes the human's demonstrations, it internally simulates "what might I be trying to achieve were I performing these demonstrations in their context?" The robot therefore interprets and hypothesizes the intended concept being taught not only from its own perspective, but from the human teacher's visual perspective as well. Through this process, the robot successfully identifies ambiguous demonstrations given by the human instructor, and clarifies the human's intent behind these confusing demonstrations. Once these problematic demonstrations are disambiguated, the robot correctly learns the intended task. In sum, I believe that maintaining mutual beliefs and common ground in human-robot teaching-learning scenarios will make robots more efficient and understandable learners, as well as more robust to the miscommunications or misunderstandings that inevitably arise even in human-human tutelage.

References

1. Barsalou, L.W., Niedenthal, P.M., Barbey, A., & Ruppert, J. (2003). Social embodiment. In B. Ross (Ed.), *The Psychology of Learning and Motivation*, Vol. 43 (pp. 43-92). San Diego: Academic Press.
2. Breazeal, C. (2002) *Designing Sociable Robots*. MIT Press. Cambridge MA.
3. Breazeal, C., Kidd, A. L. Thomaz, G. Hoffman, M. Berlin (2005) "Effects of Nonverbal Communication on Efficiency and Robustness in Human-Robot Teamwork." *Proceedings of IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS 2005)*.
4. Breazeal, C., Berlin, M., Brooks, A., Gray, J., Thomaz, A.L. (2006) "Using perspective taking to learn from ambiguous demonstrations" *Robotics and Autonomous Systems: Special Issue on Learning by Demonstration*. To appear.
5. Fong, T., Nourbakshsh, I., Dautenhahn, K. (2003). A Survey of Social Robots. *Robotics and Autonomous Systems*, **42**, pp. 143 – 166.
6. Gray, J., Breazeal, C., Berlin, M., Brooks, A. and Lieberman, J. (2005), "Action Parsing and Goal Inference using Self as Simulator." *Proceedings of Fourteenth IEEE Workshop on Robot and Human Interactive Communication (Ro-Man05)*, Nashville, TN.
7. Hoffman, G., and Breazeal, C. (2004). "Collaboration in Human-Robot Teams." *Proceedings of AIAA First Intelligent Systems Technical Conference*, Chicago, IL.
8. Meltzoff, A., and Decety, J. (2003) What imitation tells us about social cognition: a rapprochement between developmental psychology and cognitive neuroscience. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **358**, 491-500.
9. Picard, R. W. (1997) *Affective Computing*. MIT Press. Cambridge, MA.
10. Sebanz, N., Bekkering, H., and Knoblich, G. (2006) Joint action: bodies and minds moving together. *Trends in Cognitive Science* **10:2**, 70-76.
11. Scassellati, B. (2000) Investigating Models of Social Development Using a Humanoid Robot, in Barbara Webb and Thomas Consi, eds., *Biorobotics*, MIT Press. Cambridge MA.
12. Schaal, S. (1999) Is imitation learning the route to humanoid robots? *Trends in Cognitive Science* **3(6)**, 233-242.

AI: Where from and where to?

Rodney Brooks
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology, USA
brooks@csail.mit.edu

When Artificial Intelligence research started fifty years ago the computers that were available were tiny by even the standards of throwaway computer we would today put in a birthday card. But the aspirations and hopes for what we could achieve back then have not changed much until today. We have still not provided the salvation expressed in the hopes of the dozen people who met at Dartmouth College in the summer of 1956.

We have done much however, we have built systems and changed the world in ways that probably were not imagined in 1956. Artificial Intelligence has not done it alone – there has been an interplay of AI and other disciplines from the first days feeding of each other in tangled webs that will be very hard to pick apart should one insist on distinguishing the AI contributions from others.

All technologies of information and communication, from hardware to search engines, are built with and employ techniques that were once the province of Artificial Intelligence. Indeed, in the United States, at least, the three leading computer science departments (Stanford, MIT and CMU) that grew rapidly through the sixties and seventies had Artificial Intelligence Laboratories or AI subgroups and key and intellectually dominating faculty members from the Dartmouth conference. And other departments that were always ranked in the top ten in the country had significant AI components--indeed many of the top ten academic CS departments in the US are now chaired by faculty who have been part of AI labs or groups. At MIT the largest single lab on campus, accounting for close to 10% of all research is the Computer Science and Artificial Intelligence Laboratory, and there are other labs on campus where both disciplines are actively pursued together. Computer science and artificial intelligence have merged in the United States.

But what of the goals of AI? Scratch an AI person and deep down they want to create an artificial system which has human level equivalence in its abilities in the world. But what are the components of such a system?

In the early days there was everything else, and then there was thinking. Thinking was the grail of early AI. The ability to think explicitly, to solve puzzles, to play games, were the abilities that had set the early AI researchers apart from their families and friends. They were better at it than others around them and they were lauded for their intelligence. Surely then, those capabilities were the key to intelligence. And this is exactly where early AI researchers concentrated their work. And they made progress. On search, heuristics, pattern matching, and even statistical machine learning. The results fed into other technologies and became part of their fabric.

Only later was it realized that the things that are easiest for a two year old child are the hardest part of our intelligence to reconstruct. Perception and action are much harder than we first thought. Some retreat to application areas where perception and action are

not needed. And good work does follow from those decisions. But to get to the true dream of AI we must not shy away from these very difficult areas.

The real work is yet ahead of us. We have unbelievably vast computational resources at our fingertips today. We have not yet learned how to harvest them as well as those available to a 2,000 neuron polyclad flatworm. Part of the reason is that real "intelligence" is just part of a bigger system and indistinguishable from it, just as CS and AI have merged. For real intelligence the body, the physics, the context, and the environment are all part of what makes things tick. AI researchers will be struggling for the next 50 years in trying to untie this new gordian knot. They will be forced to develop new methods of study beyond the classical reductionist methods on which they have been trained. They will not do this alone, but in conjunction with those in other disciplines such as biology, social science, and who knows, maybe even climatologists.

AI and its tentacles have transformed the world over the last fifty years. It will do it again over the next fifty years. It's going to be quite a ride.

Short Biography

Rodney Brooks is the Panasonic Professor of Robotics and the Director of the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL). He is also cofounder and Chief Technical Officer of iRobot Corporation. He received two degrees in mathematics from the Flinders University of South Australia and the Ph.D. in computer science from Stanford University in 1981. He held research position at Carnegie Mellon University and MIT, and a faculty position at Stanford, before joining the faculty at MIT in 1984. His research at MIT has been centered on behavior-based robotics, first with mobile robots, then humanoid robots, and more recently concentrating on building a new class of robots to cohabit the world and assist people in all aspects of their lives. He is a Fellow of the American Association for Artificial Intelligence (AAAI), the American Association for the Advancement of Science (AAAS), the Association for Computing Machinery (ACM), and is a member of the National Academy of Engineering (NAE).

Cogniron: The Cognitive Robot Companion

Raja Chatila
LAAS - CNRS

7, Avenue Colonel Roche - 31077 Toulouse, France

“Can machines think?” was the founding question for AI posed by Alan M. Turing in 1950. AI was “officially” born as a discipline in 1956 to address this problem. By the late sixties, the Shakey project by Nils Nilsson and colleagues was a first tentative to answer the question by including perception and action to the thinking part, until Rodney Brooks shifted the focus in the mid eighties away from thinking to sensing and acting.

Since then many AI and robotics research projects are still trying to provide an answer. New directions and new concepts have emerged, many of which inspired by results and advances in understanding natural intelligent systems, namely animals and humans, and others built on the conclusions drawn from the results of implementing and experimenting with well engineered robots.

There is no real theoretical foundation today for answering this question which stirred up controversy since the beginning. However, what became clear is that there are several components to intelligence, and that they must be investigated jointly, in an integrative perspective.

The Cogniron project (“The Cognitive Robot Companion”) funded by the European commission’s FP6-IST-FET “Beyond Robotics” program tries to address the problem, shifting from intelligence to cognition¹. Four cognitive capacities are studied in the project (in reference to natural cognition in humans and animals):

- Understanding of space, objects and situations
- Decision-making
- Learning
- Communication and interaction

These capacities cannot be considered separately. The project considers them in the framework of three “key experiments”: a robot home tour in which the interactive and spatial cognitive capacities are emphasized, a pro-active and curious robot which focuses on the decision-making and interactive capacities, and the task learner which aims at experimented learning by demonstration and imitation.

The cognitive capabilities have to operate concurrently and continuously. It is therefore of utmost importance to understand how they are interleaved in the robot’s cognitive architecture, which organizes the components of the system at several levels of granularity. Finally, we are interested in evaluating how much has been accomplished. Even if every individual function composing the system can be benchmarked, their concurrent operation requires an overall performance evaluation

¹ The partners of the project are LAAS, EPFL, Fhg IPA, KTH, University of Bielefeld, University of Amsterdam, Hertfordshire University, University of Karlsruhe

taking into account the integrated nature of the system. This evaluation is based on metrics related to the complexity of the task, the complexity of the environment and the amount of information the robot has on its environment.

Why social intelligence matters in the design and development of intelligent robots

Kerstin Dautenhahn
Adaptive Systems Research Group
University of Hertfordshire, UK
K.Dautenhahn@herts.ac.uk

Since its origin 50 years ago Artificial Intelligence research has been strongly inspired and motivated by human intelligence: human thinking and problem-solving dominated until the late 1980's, whereby chess-playing, theorem-proving, planning skills etc. were taken as benchmarks of human and artificial intelligence. In contrast to these skills that are of great interest in particular to adult members of western societies, more recently sensorimotor skills emphasizing the embodied nature of human intelligence (including locomotion, object manipulation etc.) are considered to be the more fundamental but certainly more biologically and developmentally plausible milestones that researchers are aiming at, highlighting the close relationships between mind, body and environment. In such a “nouvelle AI” viewpoint a robot is more than a “computer on wheels” as it has been considered in AI for decades. A 'nouvelle AI' robot is embodied, situated, surrounded by/responding to/interacting with its environment. A 'nouvelle AI' robot takes its inspirations not necessarily from humans: insects, slugs or salamanders can be equally worthwhile behavioural or 'cognitive' models depending on the particular skills that are under investigation.

Despite impressive examples of sensorimotor skills in present day robots, reaching human-like intelligence remains a big challenge. In my talk I will argue that one particular aspect of human intelligence, namely social intelligence, might bring us closer to the goal of making robots smarter (in the sense of more human-like and believable in behaviour): the social environment cannot be subsumed under “general environmental factors”: humans interact differently with each other than with a chair, or a stone.

Despite this change in viewpoint from the so-called 'classical' to the 'nouvelle' direction of AI, social intelligence has not yet been recognized as a key ingredient of artificial intelligence while it has been widely investigated in fields where researchers study human development and intelligence. Acknowledging the social nature of human intelligence and its implications for artificial intelligence is an exciting challenge that requires truly interdisciplinary viewpoints.

My talk will argue that social intelligence is a key ingredient of human intelligence, and as such a candidate prerequisite for any artificially intelligent robot. In order to illustrate studies into social robots I will give concrete examples of current research that I have been involved in. This will cover a developmental perspective in the context of the RobotCub project. Here, interaction games and interaction histories allow a robot to extend its temporal horizon. In this way situatedness goes beyond the here and now and includes meaningful experiences (from the perspective of the robot) during its 'life-time', ultimately leading to robots with autobiographic memory of meaningful experiences. Other research emphasizes the importance of interaction games investigating interaction kinesics and means of a robot to use body (facial and other)

expressions to regulate social interactions. Further examples will also cover research into the design of robot companions, cf. Cogniron project. A robot companion in a home-environment needs to 'do the right things', i.e. it has to be useful and perform tasks around the house, but it also has to 'do the things right', i.e. in a manner that is believable and acceptable to humans. Such design-oriented work is also relevant for the Robotcub perspective of studying development in a child-sized humanoid robot that researchers as well as subjects in experimental trials need to interact with.

Human-robot interaction is a highly challenging area that requires interdisciplinary collaboration between AI researchers, computer scientists, engineers and psychologists where new methods and methodologies need to be created in order to develop, study and evaluate interactions with a social robot.

Humans are above all social animals. For artificially intelligent robots, can it be otherwise?

The Neural Control of Aerodynamics in Fruit Flies

Michael H. Dickinson,
Caltech, MC 138-78, Pasadena CA, 91125
flyman@caltech.edu

A central feature in the natural history and behavior of fruit flies is the ability to fly through the air in search of food, mates, and oviposition sites. What neural processes enable these tiny flies to take-off and efficiently explore their local environment? Like all forms of locomotion, flight behavior results from a complex set of interactions, not just within circuits in the brain, but among neurons, muscles, skeletal elements, and physical process within the external world. To control flight, the fly's nervous system must generate a code of motor information that plays out through a small but complicated set of power and steering muscles. These muscles induce microscopic oscillations in an external skeleton that drive the wings back and forth 200 times each second producing a time-variant pattern of aerodynamic forces that the fly modulates to steer and maneuver through the air. The animal's motion through space alters the stream of information that runs through an array of visual, chemical, and mechanical sensors, which collectively provide feedback to stabilize flight and orient the animal towards specific targets. The goal of the research in my laboratory is to 'reverse engineer' this flight control system, and thus determine the means by which the nervous system controls the animal's trajectory through space.

Drosophila, like many flies, search and explore their environment using a series of straight flight segments interspersed with stereotyped changes in heading termed *saccades*. Each saccade is a rapid maneuver in which the fly turns 90 degs in less than 50 ms. Using a combination of tethered and free flight methods, we have investigated both the sensory signals that trigger these rapid turns as well as the aerodynamic means by which the animals produce the required torque. The results suggest that the saccades represent a collision avoidance reflex initiated by the visual system. In tethered flight simulators, flies vigorously turn away from poles of visual expansion. The reflex is so strong that under closed-loop conditions flies actively orient toward poles of visual contraction – a result that is counterintuitive considering the visual flow flies would expect to encounter during forward flight. The frequency and spatial distribution of saccades are altered by the presence of attractive odors – which helps to explain why flies hover over fruit bowls and rotting bananas.

Once triggered, hard wired sensory-motor circuitry executes a rapid all-or-none program that directs a saccade either to the left or to the right. Although angular acceleration during the turns approach $20,000 \text{ degs s}^{-1}$, both the changes in motor output and the resultant alterations in wing motion required to produce saccades are quite subtle. Further, a high speed analysis of saccades indicates that flies must generate torque to start the turn, and counter-torque to stop. This result suggests that, despite their small size, fruit fly flight body dynamics are dominated by inertia and not friction during the brief saccades. In addition, free flight saccades are much shorter than fictive saccades in tethered flight, underscoring the importance of sensory feedback in regulating saccade duration. Evidence suggests that whereas the visual system triggers the saccade, the signal to initiate the counter-turn that terminates the maneuver arises from the mechanosensory halteres, which are more sensitive than the eyes to rapid

rotation. This research illustrates how processes within the physical world function with neural and mechanical features of an organism's design function to generate a complex behavior.

Short Biography

Michael Dickinson was born in Seaford, Delaware in 1963, but spent most of his youth in Baltimore and Philadelphia. He attended college at Brown University, originally with the intent of majoring in Visual Arts, but eventually switched to Neurobiology, driven by a fascination for the mechanisms that underlie animal behavior. While in college, he studied the roles of neurons and neurotransmitters in the control of leech feeding behavior. He received a Ph. D. in Zoology at the University of Washington in Seattle in 1991. His dissertation project focused on the physiology of sensory cells on the wings of flies. It was this study of wing sensors that led to an interest in insect aerodynamics and flight control circuitry. Michael worked briefly at the Max Planck Institute for Biological Cybernetics in Tübingen, Germany, and served as an Assistant Professor in the Dept. of Anatomy at the University of Chicago in 1991. He moved to University of California, Berkeley in 1996 and was appointed as the Williams Professor in the Department of Integrative Biology in 2000. Dickinson moved to Caltech in July, 2002 and is currently the Abe and Esther Zarem Professor of Bioengineering and Biology.

Dickinson's research interests broadly concern the mechanistic basis of animal behavior. Specifically, he has studied the flight behavior of insects simultaneously at several levels of analysis, in an attempt to integrate cellular physiology, biomechanics, aerodynamics, and behavior. His awards include the Larry Sandler Award from the Genetics Society of America, the Bartholemew Award for Comparative Physiology from the American Society of Zoologists, a Packard Foundation Fellowship in Science and Engineering, and the Quantrell award for Excellence in Undergraduate Teaching at the University of Chicago. In 2001, he was awarded a MacArthur Foundation Fellowship.

Emergent Cognitive Capabilities for Humanoids: Robots Learning Senso-Motor Skills and Task Knowledge from Multimodal Observation of Humans

Rüdiger Dillmann
Universität Karlsruhe (TH), Germany
dillmann@ira.uka.de

Humanoid robot systems are designed to interact with humans in terms of conversation about the task to be done and how to do the task and finally how to execute it. In addition humanoids act goal-oriented and are capable to react on disturbances or unexpected events in a competent manner. Such robot systems operate in dynamic human centered scenarios which require capabilities such as adaptivity, perception, categorisation, action and learning. Examples for such systems are service robots or humanoids that interact in the immediate environment of humans in a context dependant goal-oriented manner and cooperate with humans. The behavior of such robots is characterized by active sensing processes, fusion of sensor data, perception, interpretation and the selection of appropriate actions as well as their superimposed control systems. Thereby learning and shaping of sensory and motoric abilities as well as the active observation and interpretation of situations and actions are of major interest.

A probate approach for learning knowledge about actions and sensomotoric abilities is to acquire knowledge about human actions by sensorial observation, trying to imitate, to understand and to transfer these abilities in the sense of learning by demonstration to the robot. This requires human motion capture, observation of interaction, object state transitions and observation of spatial and physical relations between objects. By doing this, it is possible to acquire so-called "skills", situative knowledge as well as task knowledge, and can be introduced to new and unknown tasks. New terms, new objects and situations, even new types of motion can be learned with the help of a human tutor or be corrected interactively via multimodal channels. The term multimodality describes communication channels which are intuitive for humans, such as language, gesture and haptics (physical human-robot contact). These are to be used for commanding and instructing the robot system.

The field of programming by demonstration has been evolved strongly as a response to the needs of generating flexible programs for service robots and is largely driven by attempts of modeling human behaviour and to map it onto virtual Androids or humanoid robots. It comprises a broad set of observation techniques processing large sets of data from high speed camera systems, laser, data-gloves and even exoskeleton devices. Some systems operate with precise a-priori models other use statistical approaches to approximate human behaviour. In any case observation is done to identify motion in space and time, interaction with the environment and its effects, useful regularities or structures and its interpretation in a given context. With this goal, systems have been developed which combine active sensing, computational learning techniques, multimodal dialogues to enrich the semantic system level, memorisation techniques as well as mapping strategies to make use of the learned knowledge to control a real robot. One important paradigm is that objects and action representations cannot be separated and form the building blocks for cognitive robot systems. Thus, so called object-action

complexes -OACS- can be derived to unify different sensor, actuator and object representations including language and allow the robot to understand its environment (EU-6.framework IP PACO-PLUS).

In the context of the SFB 588 "Humanoid Robots" a mobile two-arm system with five-finger hands, a flexible torso as well as a head with visual and acoustic sensors is being developed, which appears to behave like a human. The locomotion system and the behaviour of the robot as well as the multimodal interaction are tailored to humans. For the support of human-robot cooperation it is important for the robot to acquire the aim of humans, to remember already corporately accomplished actions and to apply this knowledge in each individual case in the correct way. The status of the Humanoid Robot Project is outlined and the achieved results are discussed.

Short Biography

Prof. Dr.-Ing. Rüdiger Dillmann received his Ph. D. in Electrical Engineering at the Universität Karlsruhe (TH) in 1980. Since 1987 he is Professor at the Institute of Computer Science and Technical Engineering (CSE) (<http://www.iain.ira.uka.de/>) and since 2001 director of the research group, Industrial Applications of Informatics and Microsystems (IAIM) at the Universität Karlsruhe (TH). Since 2002 he is also president of the Research Centre for Information Technologies in Karlsruhe (FZI).

As leader of these two institutes Prof. Dillmann supervises several fundamental and industrial research in the areas of robotics with special interest on intelligent robot systems, computer assisted surgery and interactive simulation systems in the context of learning. Prof. Dillmann is author or co-author of more than 200 scientific publications and several books.

He is involved in many synergistic activities like: Director of the German collaborative research centre "Humanoid Robots", IEEE-RAS chairman of the German chapter, Chairman of the German Society of Information Science (GI), section 4.3/1.4 "Robotic Systems", Chairman of the German Association of Engineers (VDI-GMA), member of the IEEE-RAS AdCom, Editor in chief for the journal "Robotics and Autonomous Systems", Elsevier.

Computations performed by collections of recurrently connected neurons, and their implementation in hybrid VLSI electronic systems

Rodney Douglas

Institute of Neuroinformatics, ETH and University of Zurich, Switzerland

rjd@ini.phys.ethz.ch

Animal brains are dramatically more effective in dealing with real-world tasks than even the most advanced computers. In mammals, the neocortex is very likely the subsystem most relevant for intelligent and effective interaction with the world, and it is one region where we can hope to understand the relationship between neuronal architecture and the computation that it supports. Fortunately, the evidence accumulating since the fundamental work of Gilbert and Wiesel indicates that the basic architecture and operation of cortex might be explicable in terms of the relationship between relatively few types of excitatory and inhibitory neurons^{4,5}. In our quantitative studies of the static connections weights between neurons in cat visual cortex, we have found that one prominent feature of the neocortical circuit of cat visual cortex is the high degree of connectivity between pyramidal cells in the superficial cortical layers, suggesting that the fundamental computational process of cortex depends on direct recurrence between these pyramids².

Populations of such recurrently connected neurons can implement 'soft Winner-Take-All' (sWTA) circuits that have interesting computational properties, that are quite different to conventional computing circuits^{6,7}. For example, analog amplification and digital multistability are generally seen as incompatible functions and are separated into two classes of electronic technology. But, in the neocortical circuits multistability can coexist with analog responses. The sWTA circuits exhibit population coding, gain modulation, focal attention (signal selection), and spatiotemporal pattern generation, all of which are characteristics of neocortical computation^{3,1,11,10}, and they can be fabricated in custom Very Large Scale Integrated electronic circuits composed of either rate- or spiking-neurons^{7,9,8}. Although the properties of sWTA's are now well understood, and there are some examples of how they can be used in particular neuronal and technological applications, there is little understanding of how one could build a general processor with these circuits. In this talk I will describe our steps towards building a general relational processor that, probably like the neocortex, uses the sWTA as its principle element.

Short Biography

Rodney Douglas is Professor of Neuroinformatics, and Co-Director at the Institute of Neuroinformatics of the Swiss Federal Institute and the the University of Zurich. He graduated in Science and Medicine at the University of Cape Town. After obtaining a Doctorate in Neuroscience, he moved to the Anatomical Neuropharmacology Unit in Oxford, where he continued his research on the anatomy and biophysics of the microcircuitry of cerebral together with Kevan Martin. As Visiting Associate, and then



Visiting Professor at Caltech, he extended his research interests in neuronal computation to the modeling of cortical circuits using digital methods (together with Christof Koch), and also by the fabrication of analog VLSI circuits (together with Misha Mahowald). In 1996 he and Kevan Martin moved to Zurich to establish the Institute of Neuroinformatics. In 2000, Douglas was awarded the Körber Foundation prize for European Science. His current research interests include; experimental anatomy and physiology of visual cerebral cortex; theoretical analysis and simulation of cortical circuits; design and fabrication of neuromorphic systems that exploit analog Very Large Scale Integration methods to construct electronic circuits that perform analogous signal processing and computational functions to biological neuronal networks.

References

1. Ben-Yishai B. and Bar-Or R.L and Sompolinsky H., Theory of orientation tuning in visual cortex, *Proceedings National Academy of Science USA*, 92(9)3844-8, 1995.
2. Binzegger, T. and Douglas, R. and Martin, K. A quantitative map of the circuit of cat primary visual cortex, *Journal of Neuroscience*, 24:(39) 8441-53, 2004
3. Douglas, R. J. and Koch, C. K. and Mahowald, M. and Martin, K. A. C. and Suarez, H. H, Recurrent Excitation in Neocortical Circuits, *Science*, 269:981–985, 1995.
4. Douglas, R. J. and Martin K. A. C. and Witteridge D., A canonical microcircuit for neocortex, *Neural Computation*, 1:480–488, 1989.
5. Douglas, R. and Martin, K. Neuronal Circuits of the Neocortex, *Annual Review of Neuroscience*, 27: 419-51, 2004.
6. Hahnloser R., and Douglas, R., and Mahowald, M. and K. Hepp, Feedback interactions between neuronal pointers and maps for attentional processing, *Nature Neuroscience*, 2:746-752, 1999.
7. Hahnloser, R. and Sarpeshkar, R. and Mahowald, M. and Douglas, R.J. and Seung, S., Digital selection and analog amplification co-exist in an electronic circuit inspired by neocortex, *Nature*, 405:947–951, 2000.
8. Indiveri, G. and Chicca, E. and Douglas, R. A VLSI array of low-power spiking neurons and bistable synapses with spike-timing dependent plasticity, *IEEE Transactions on Neural Networks*, 2005, (In Press)
9. Liu, S.-C. and Kramer, J. and Indiveri, G. and Delbruck, T. and Douglas R., Analog VLSI: Circuits and Principles, *MIT Press*, 2002.
10. Poggio, T. and Bizzi, E., Generalization in vision and motor control, *Nature*, 43:768-774, 2004.
11. Salinas E. and Abbott L.F, A model of multiplicative neural responses in parietal cortex, *Proceedings National Academy of Science USA*, 93(21):11956-61, 1996.

Biophysical Studies on Brain-Computer Interfacing

Peter Fromherz

Department of Membrane and Neurophysics
Max Planck Institute for Biochemistry
Martinsried/München

The term Artificial Intelligence was introduced for the emulation of “intelligent” functions of animal brain by digital electronic computers. We may ask whether a technology of such intelligent functions is not better implemented with the same material as it is used in animal brain, i.e. with nerve cells in electrolyte. A complete information processing system would be a hybrid of a “thinking” neuronal system combined with a computing electronic system. We must solve two fundamental problems to achieve that goal: On one hand, we have to “understand” the brain, i.e. to know the “trick” how the slow dynamics of nerve cells is able to perform such astonishing tasks as fast pattern recognition and fast control of motor dynamics. On the other hand, we have to interface neuronal and electronic systems on a microscopic level in both directions in order to take full technological advantage of “thinking” in the hybrid system. It is the latter problem that we consider in our work. We assemble simple hybrid devices of nerve cells and semiconductor chips and study the basic biophysics of interfacing. By enhancing the complexity of the hybrids step by step, we look how far we come to implement novel functions. A side aspect of the approach is that it may also help to solve the first problem by developing novel neurophysiological techniques.

On the side of the neuronal system, we consider the three levels of ion channels, of individual nerve cells and of brain tissue. On the side of microelectronics, we study the basis of two-way interfacing using transistors and capacitors of simple silicon chips and then transfer the method to more involved chips that are fabricated by industrial CMOS technology.

The structure and the electrical properties of cell-silicon contacts are studied with luminescent dyes that are embedded in the cell membrane. The reflecting surface of silicon gives rise to a change of fluorescence due to interference effects such that the distance of cell and chip can be determined. Alternate voltages applied to silicon induce a spectral shift of fluorescence that is used to determine the electrical resistance of the cell-chip contact.

The mechanism of electrical interfacing is studied with recombinant channels for Na^+ and K^+ ion in the membrane of cells that are cultured on capacitors and transistors. When voltage ramps are applied to a chip, capacitive current gives rise to a voltage across the cell membrane and opens the ion channels. When the channels are open, ionic current flows along the cell-chip contact and gives rise to a gate-voltage that modulates the source-drain current of a transistor.

We implemented the two-way interfacing of nerve cells from invertebrates (snails) and mammals (rats). Interfacing of invertebrate neurons is more advanced, because these cells are larger such that the strength of coupling to capacitors and transistors is higher. By improving the quality of capacitors (high-k insulators) and of transistors

(low-noise design), significant progress was recently achieved with individual mammalian neurons, too.

Elementary hybrid circuits were assembled with two neurons: (i) One neuron is stimulated from a capacitor, its activity is coupled to a second neuron through a synapse and the excitation of the second neuron is recorded with a transistor. (ii) An excited neuron is coupled to a transistor, its signal is shaped on the chip and used to stimulate a second neuron with a capacitor. More complicated networks were cultured on the chip with defined geometry using controlled outgrowth by chemical and topographical patterns. The yield of the resulting hybrid systems, however, was still rather low.

Neuronal networks from brain were interfaced to silicon chips by culturing slices from rat hippocampus. In that case, the two-way interfacing with capacitors and transistors was implemented through local populations of neurons. It was possible to induce learning effects (LTP, LTD) from the chips.

CMOS chips with 16000 recording transistors on one squaremillimeter were developed with an inert surface of titaniumdioxide, such that the mechanism of interfacing was the same as with simple silicon chips. One-way multi-site interfacing was achieved with small networks of snail neurons as well as with cultured slices from rat brain. Experiments with more involved CMOS chips for two-way multisite interfacing are in progress.

Reference

Fromherz, P. (2005). The Neuron-Semiconductor Interface. In: Bioelectronics – from theory to applications. Eds. I. Willner & E. Katz (Wiley-VCH, Weinheim) pp 339-394.

Short Biography

Peter Fromherz is a director at the Max Planck Institute for Biochemistry in Martinsried/Munich and professor for Experimental Biophysics in the Physical Faculty of the Technical University Munich. He completed his PhD in Physical Chemistry in 1969 at the University Marburg. Subsequently, he led a research group at the Max Planck Institute for Biophysical Chemistry in Göttingen. In 1981 he became a full professor for Experimental Physics at the University Ulm. Since 1994 he is a scientific member of the Max Planck Society. His present interests are the interfacing of semiconductor chips with neuronal systems and the development of voltage-sensitive dyes for brain research.

Get A Life

Inman Harvey
CCNR/EASy, University of Sussex
inmanh@susx.ac.uk

To create artificial analogues of adaptive, intelligent cognitive life-forms requires both technical ability and an appropriate level of understanding. Over the last 50 years the former, in terms of computing and materials, has developed out of all recognition, but the latter has been rebuffed at a number of dead-ends. I shall sketch the failure of the computer metaphor for the brain, the failure of the first wave of Artificial Neural Networks, the failure of Neuroscience, the unanswered questions of behavior-based robotics – where is the Juice?

I shall nevertheless advocate a realistic optimism, with a focus on issues of metabolism, homeostasis and dirt. Autopoiesis is the ultimate form of homeostasis, and is what grounds *meaning* for an agent – in terms ultimately of what, in its world, is important for its continued survival. So artificial agents lack the juice in so far as their survival is not actually under threat. Dirt is important as the underlying explanation for why strong Alife in a computer simulation is a non-starter.

We seek to build artificial agents that have meanings and values in their worlds. But also in doing so, the characteristics of the models we build betray our own meanings and values. Does the evidence suggest that the typical practitioner of AI is naive, unsophisticated, uncultured and heartless – and quite possibly a militaristic control freak to boot?

Short Biography

Inman Harvey's first degrees from Cambridge were in Mathematics, Moral Science, and Social Anthropology. After many years running his own businesses, primarily import-export out of Afghanistan, he came to the University of Sussex in 1988 to research in Evolutionary Robotics and artificial evolution. He is a founder member of the Evolutionary and Adaptive Systems (EASy) group there. His current research interests also include Gaia theory, Neural Networks, a Dynamical Systems approach to minimal cognition studies, the Maximum Entropy Production principle.

Mobiligence Project: Emergence of Adaptive Motor Function through Interaction among the Body, Brain, and Environment

Koh Hosoda[†], Akio Ishiguro[‡] and Hajime Asama*

[†]Department of Adaptive Machine Systems, Osaka University

[‡]Department of Electrical and Communication Engineering, Tohoku University

*Research into Artifacts, Center for Engineering, the University of Tokyo

hosoda@ams.eng.osaka-u.ac.jp, ishiguro@ecei.tohoku.ac.jp,

asama@race.u-tokyo.ac.jp

http://www.arai.pe.u-tokyo.ac.jp/mobiligence/index_e.html

Mobiligence Project Abstract

Animals behave adaptively in diverse environments. Adaptive behavior, which is one of the intelligent sensory-motor functions, is disturbed in patients with neurological disorders. However, the mechanisms for generating intelligent adaptive behavior are not thoroughly understood. Such an adaptive function is considered to emerge from the interaction of the body, brain, and environment, which requires movement of the subject. We call such adaptive intelligent function *mobiligence*.

The project is designed to investigate the mechanisms of mobiligence by collaborative research in biology and engineering. In the course of this collaborative project, the following steps will be carried out:

1. biological and physiological examinations of animals;
2. modeling of biological systems;
3. construction and experiments on artificial systems by utilizing robotic technologies;
4. creation of a hypothesis and its verification.

The goal of this project is to establish the common principle underlying the emergence of mobiligence.

Research Approach to the Mobiligence Project

In this project, the mobiligence mechanism is elucidated by the constructive and systematic approaches through the collaboration of biologists and engineering scientists who developed biological models by integrating physiological data and kinetic modeling technologies (see Figure 1). In other words, the Mobiligence Project is pursued by integrating biology and engineering, i.e., physiological analysis (biology), modeling and experiments on artificial systems (engineering), verification of models (biology), and discovery and application of principles (engineering).

In the project, we focus on three adaptive mechanisms:

1. Mechanism whereby animals adapt to recognize environmental changes;
2. Mechanism whereby animals adapt physically to environmental changes; and
3. Mechanism whereby animals adapt to society.

Research groups for each of the categories listed above are organized. The three groups conduct their respective research and clarify the universal, common principle

underlying the mechanism of mobiligence. The Planned Research Team studies the following specific subjects: analysis of the environmental cognition and the adaptive mechanism in reaching movements; analysis of the physical adaptive mechanism in walking; and analysis of the adaptive mechanism observed in the social behaviors of insects. In addition, the Planned Research Team clarifies the common principle underlying mobiligence from a dynamic viewpoint. Furthermore, we study adaptive mechanisms relating to various objects by publicly inviting proposed topics and clarify the universal, common principle therein.

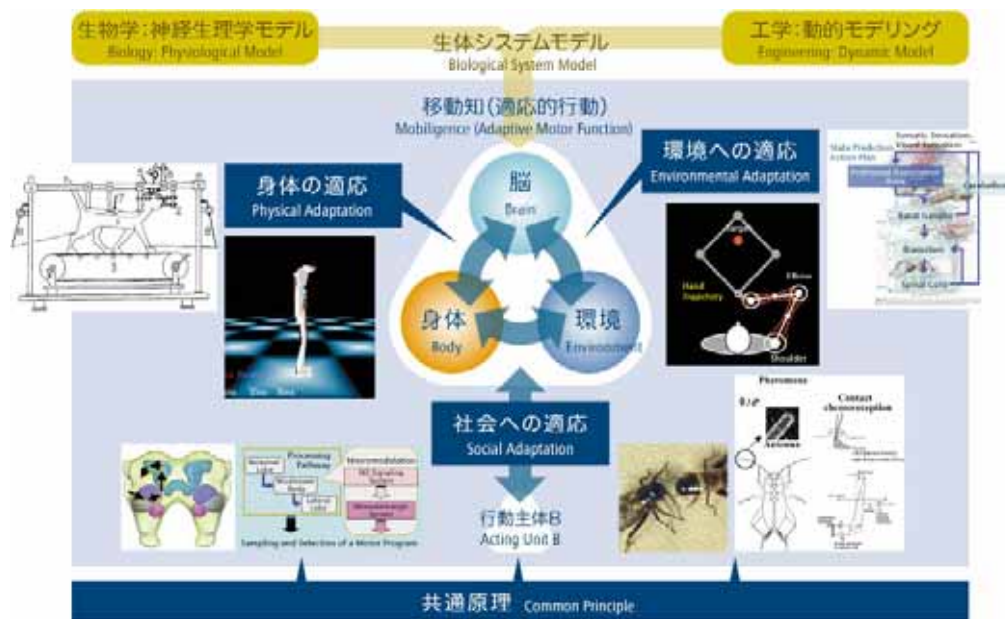


Figure 1 Overview of the Mobiligence Project

Objectives of the Mobiligence Project

Various adaptive motor function mechanisms used by animals will be studied. In the medical field, the results of our research will contribute to the discovery of a method to improve motor impairment and develop rehabilitation systems. In addition, in the engineering field, the results of our research will contribute to the determination of the constructive principles of artificial intelligence systems. Furthermore, we will explore the new research field, mobiligence, establish a research organization that integrates biology and engineering, and implement programs to train engineering scientists and biologists to conduct biological and engineering research, respectively.

For more information on the project, visit the project web site:

http://www.arai.pe.u-tokyo.ac.jp/mobiligence/index_e.html

Adaptive locomotion in robots with multiple degrees of freedom

Auke Ijspeert
Biologically Inspired Robotics Group
École Polytechnique Fédérale de Lausanne, Switzerland
auke.ijspeert@epfl.ch

The agility and efficiency of animal locomotion tend to fascinate engineers. The skills to coordinate multiple degrees of freedom, using compliant actuators (muscles and tendons), and massively parallel control (the central nervous system), give animals an agility and energy efficiency rarely replicated in man-made robots. In vertebrate animals, the problem of coordinating and modulating the multiple rhythmic signals necessary for locomotion are generated by central pattern generators (CPGs) in the spinal cord. CPGs are networks that need only simple control signals to initiate and modulate complex locomotion patterns.

In this talk, I will present how the concept of CPGs can be very useful to control robots with multiple degrees of freedom. Results of several projects ranging from snake-like to humanoid robots will be presented. In particular, I will present the first results in controlling the crawling and walking of the simulated iCub robot, the humanoid robot developed in the framework of the Robotcub project. The controllers are constructed out of coupled nonlinear oscillators. Different interesting properties of the system will be presented such as the possibility to learn arbitrary rhythmic signals, limit cycle behavior (i.e. stable rhythm generation), integration of sensory feedback, and modulation of speed and direction of locomotion. Interactions with other components of the Robotcub project will also be discussed, in particular recordings of infant crawling and sitting behaviors in infants carried out by Kerstin Rosander at Uppsala University, and more general sensorimotor coordination.

Curious and Creative Machines

Hod Lipson
Computational Synthesis Laboratory
Cornell University, Ithaca NY, USA
hod.lipson@cornell.edu

One of the hallmarks of human intelligence is the ability to design: To synthesize a set of elementary building blocks in order to achieve some novel, high-level functionality. Imagine a Lego set at your disposal: Bricks, rods, wheels, motors, sensors and logic components are your “atomic” building blocks, and you must find a way to put them together to achieve a given high-level functionality: A machine that can moveⁱ, say. You know the physics of the individual components' behaviors; you know the repertoire of pieces available, and you know how they are allowed to connect. But how do you determine the combination that gives you the desired functionality? This is the problem of *Synthesis*.

The Second Half of AI

In the last two centuries, engineering sciences have made remarkable progress in their ability to analyze and predict physical phenomena. We understand the governing equations of thermodynamics, electromagnetics, and fluid flow, to name but a few. Numerical methods such as finite elements allow us to solve these constitutive equations with good approximation for many practical situation. We can use these methods to investigate and explain observations, as well as to predict the behavior of products and systems long before they are ever physically realized.

But progress in systematic *synthesis* has been frustratingly slow. Robert Willis, a professor of natural and experimental philosophy at Cambridge, wrote back in 1841:

*[A rational approach is needed] to obtain, by direct and certain methods, all the forms and arrangements that are applicable to the desired purpose. At present, questions of this kind can only be solved by that species of intuition that which long familiarity with the subject usually confers upon experienced persons, but which they are totally unable to communicate to others. When the mind of a mechanic is occupied with the contrivance of a machine, he must wait until, in the midst of his meditations, some happy combination presents itself to his mind which may answer his purpose.”*ⁱⁱ

Almost two centuries later, a rational method for general open-ended synthesis is still not at hand. Engineering design is still taught today largely through apprenticeship: Engineering students learn about existing solutions and techniques for well-defined, relatively simple problems, and then – through practice – are expected to improve and combine these to create larger, more complex systems. How is this synthesis process done? We do not know, but we cloak it with the term “creativity”. Even fields far from engineering, such as humanities and arts, share the same conundrum: You can appreciate good poetry, music, and sculpture, but how do you systematically create it?

The field of Artificial Intelligence has not escaped this inevitable course either. Over the last fifty years, AI – and its modern incarnation as machine learning in particular – has been primarily occupied with modeling and prediction, but not synthesis of *new things*. Learning from examples, combining logical facts, and propagating constraints, leave us interpolating inside the convex hull of our existing knowledge. I am not

claiming that this is either easy or that it is not useful, nor that it has been fully mastered. But it is a fundamentally different direction than the quest for open-ended synthesis, where the results are unbounded in their complexity and performance.

The Evolutionary Nature of AI

Artificial Intelligence is almost an oxymoron: Whenever breakthroughs are achieved – from Deep Blue’s mastery of chess to Stanley’s autonomous traversal of the Mojave desert – some are quick to point out that something is still missing. If it was manually designed, can be *truly* intelligent? Perhaps it is the ability to create *new things*, that would ultimately convince the skeptics. While computers can compute – and now analyze – almost anything, open ended creativity is still the unconquered holy grail still seen as distinctively human.

Human intelligence is ultimately a natural biological phenomenon, and like any other biological phenomenon, it is a product of evolution. Many theses have been written about the evolutionary origin of intelligenceⁱⁱⁱ, and one argument is that intelligence was driven by the need to create and use new tools. Not blindly execute an innate recipe for building a nest, a dam, or a hive – but a true *adaptive* ability to construct new things that exploit current resources, strengths and weaknesses of others.

Indeed the two standing examples of systematic synthesis we have to inspire us are both evolutionary: One is natural evolution, governed by Darwinian natural selection and variation. The other example is engineering design – not by the mythical maverick designer, but by a slow evolutionary progress, accumulating successive small variations and recombination of exiting technologies made by millions of ordinary designers, subject to the natural selection of the market^{iv}. These evolutionary processes are admittedly slow, inefficient, and provide no guarantees of optimality or even success, but perhaps there are fundamental limits on the conversion of energy into new information – a kind of thermodynamic law^v.

On Curiosity and Creativity

Perhaps the most fascinating form of intelligence is the one that combines open-ended synthesis with open-ended analysis. *Curiosity* is the pursuit of new knowledge: Not only passively searching for patterns in data, but actively probing and perturbing the world to extract new information – like a child asking questions. Asking the right question is again an open-ended synthesis problem, involving creation of new predictive hypotheses and generation of actions to best test their consequence^{vi}. As any parent knows – curiosity and creativity are hallmarks of intelligence. Can we make such curious machines? Will we relinquish some control over what they discover and create?^{vii} Are we ready to give up on our human-centric claim to curiosity and creativity?

Conclusion

I am not alone in this quest for a new AI that can creatively generate new things^{viii} and ask the right questions, nor am I unique in my view that nature’s evolutionary processes provide the key; but *open-ended* evolutionary computation and active learning^{ix} have existed on the periphery of mainstream AI for decades. In this fiftieth anniversary of AI, I seek a new thrust – from analysis to synthesis, and from learning machines, to curious and creative machines.

Short Biography

Hod Lipson is an Assistant Professor at the departments of Mechanical & Aerospace Engineering and the faculty of Computing & Information Science of Cornell University in Ithaca, NY. Prior to this appointment, he was a postdoctoral researcher at Brandeis University's Computer Science Department and a Lecturer at MIT's Mechanical Engineering Department. He received his Ph.D. from the Technion – Israel Institute of Technology in 1998. Lipson's research focuses on new methods for autonomous adaptation in behavior and morphology of robotic systems, with broader impacts to design automation and manufacturing technologies. His work uses primarily biologically-inspired approaches, as they bring new ideas to engineering and new engineering insights into biology.

ⁱ Lipson, H., Pollack J. B., (2000), "Automatic Design and Manufacture of Artificial Lifeforms", *Nature* 406, pp. 974-978

ⁱⁱ Robert Willis (1841), *Principles of Mechanism*

ⁱⁱⁱ Geary D., (2004) *Origin of Mind: Evolution of Brain, Cognition, and General Intelligence*. American Psychological Association

^{iv} Ziman J. (Ed.), (2000) *Technological Innovation as an Evolutionary Process*, Cambridge University Press

^v Kauffman, S., (1993) *The Origins of Order : Self-Organization and Selection in Evolution*, Oxford University Press

^{vi} Bongard J., Lipson H. (2005) "Nonlinear system identification using coevolution of models and tests", *IEEE Transactions on Evolutionary Computation*, 9(4): 361-384

^{vii} Kelly K, (1995) *Out of Control: The New Biology of Machines, Social Systems and the Economic World*, Perseus Books Group

^{viii} John Koza (2005), *Genetic Programming: Routine Human-Competitive Machine Intelligence*, Springer

^{ix} Angluin D. (1987) Learning regular sets from queries and counterexamples. *Information and Computation*, 75:87–106, 1987.

metric to experiential episodes of various temporal horizons, it is possible to impose geometric order on a robot's temporally extended sensorimotor experiences, at various temporal scales.⁵ The structure of these dynamic spaces of experiences provides an agent-centric enactive representation of interaction histories with the environment, grounded in the temporally extended sensorimotor experience.^{2,3,5,6} Potentially an agent can act using this dynamically growing, developing space of memories to return to familiar experiences, predict the effect of continuing on a current behavioural trajectory, and explore at the boundary of what is already mastered (cf. Vygotsky's zone of proximal development). By using temporally extended experiences to guide action and interaction, we will have the beginnings of post-reactive robotics and episodic intelligence in artificially constructed enactive agents that grow, develop, and adapt their cognitive structures with a broader temporally horizon.⁴

Selected References

1. K. Dautenhahn & T. Christaller, Remembering, Rehearsal and Empathy - Towards a Social and Embodied Cognitive Psychology for Artifacts. In: *Two Sciences of the Mind. Readings in Cognitive Science and Consciousness*, S. O'Neill, P. Mc Kevitt, & E. Mac Aogain, Eds. John Benjamins, pp. 257-282, 1996.
2. N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, & R. te Boekhorst, "Using Sensory-Motor Phase-Plots to Characterise Robot-Environment Interactions", *IEEE Computational Intelligence in Robotics & Automata (IEEE CIRA'05)*, special session on Ontogenetic Robotics, pp. 581-586, 2005.
3. N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, & R. te Boekhorst, "Interaction Histories: From Experience to Action and Back Again", *Proc. Fifth International Conference on Development and Learning (May 31-June 3, 2006 - Bloomington, Indiana, USA)*, IEEE Computer Science Press, (in press).
4. C. L. Nehaniv, K. Dautenhahn, & M. J. Loomes, Constructive Biology and Approaches to Temporal Grounding in Post-Reactive Robotics, In: G. T. McKee & P. Schenker, Eds., *Sensor Fusion and Decentralized Control in Robotics Systems II* (September 19-20, 1999, Boston, Massachusetts), Proceedings of SPIE Vol. 3839, pp. 156-167, 1999.
5. C. L. Nehaniv, "Sensorimotor Experience and Its Metrics: Informational Geometry and the Temporal Horizon", *IEEE CEC 2005 (2005 Congress on Evolutionary Computation, 2-5 September 2005, Edinburgh, Scotland)*, Proceedings Vol. 1, pp. 142-149, 2005.
6. C. L. Nehaniv, N. A. Mirza, K. Dautenhahn, & R. te Boekhorst, "Extending the Temporal Horizon of Autonomous Robots", *Proc. 3rd International Symposium on Autonomous Minirobots for Research and Edutainment (AMiRE 2005, Awara-Spa, Fukui, JAPAN, 20-22 September 2005)*, Editors: K. Murase, K. Sekiyama, N. Kubota, T. Naniwa, & J. Sitte, Springer Verlag, pp. 389-395, 2006.
7. L. Olsson, C. L. Nehaniv, & D. Polani, "Sensory Channel Group and Structure from Uninterpreted Sensor Data", *IEEE NASA/DoD Conference on Evolvable Hardware 2004*, IEEE Computer Society, pp. 153-160, 2004.
8. L. Olsson, C. L. Nehaniv, & D. Polani, From Unknown Sensors and Actuators to Actions Grounded in Sensorimotor Perceptions, *Connection Science*, special issue on Developmental Robotics (D. Bank & L. Meeden, guest eds.), volume 18(2), 2006 (in press).
9. L. Olsson, C. L. Nehaniv, & D. Polani, "Measuring Informational Distances between Sensors and Sensor Integration", *Artificial Life X*, MIT Press, 2006 (in press).

Routes to the Summit

Nils J. Nilsson
Artificial Intelligence Laboratory
Department of Computer Science
Stanford University
nilsson@cs.stanford.edu

Achieving human-level artificial intelligence has turned out to be an extremely difficult task—much like climbing an unconquered peak. I exploit this mountain-climbing analogy to describe some of AI's history, achievements, detours, and prospects.

Many people have dreamed of intelligent artifacts. That's a good thing, because most expeditions begin with and are sustained by dreams. In mythology, we have the "Golden Maids" who served the god Hephaestus. And there is Galatea, a statue sculpted by Pygmalion and brought to life by Venus. Leonardo Da Vinci made sketches for a robot knight. And, of course, we have Karel Capek's "Rossum's Universal Robots" and Isaac Asimov's science fiction stories. But along with dreams, serious climbers must have some clues about how to conquer the mountain. Early clues about what would be needed to build intelligent artifacts can be found in Aristotle's syllogisms, Leibniz's ratiocinator, Boole's propositional algebra, and Frege's *Begriffsschrift* (concept writing).

Life itself provides several additional important clues. McCulloch and Pitts showed that networks of simple models of a biological neuron could compute all computable functions. George Miller and Noam Chomsky were among those who helped launch the field of cognitive science, providing hints about some of the higher-level functions implemented in our brains. And, because intelligent life forms evolved, perhaps the very processes and history of evolution can be employed to show us the way toward intelligent artifacts.

From engineering come additional important clues. Among these are feedback mechanisms, automatic machinery such as the Jacquard loom, and, most importantly, the computer. Turing and Shannon are among the first to offer detailed suggestions about how computer programs could play chess—a game that requires intelligence.

Armed with these ideas, several "base camps" were laid out. Frank Rosenblatt established camp "Neural Nets." Allen Newell and Herb Simon established camp "Cognitive Processes." John McCarthy established camps "Logical Methods" and "Commonsense." Marvin Minsky established camps "Heuristic Programming" and "Frames." Lotfi Zadeh established camp "Fuzzy," and Ed Feigenbaum established camp "Knowledge Engineering." There were many other camps as well, and all of them overlapped and shared ideas and resources. Many used the same "climbing gear," such as LISP and graph-searching methods.

At first the going was rather easy. Several preliminary pitches were scaled without too much difficulty. Climbers hailed the successes of the Perceptron, GPS, Dendral, Mycin, Shrdlu and the robots Shakey and Freddy. Ebullience reigned amongst the climbers. Optimistic predictions were made about summiting soon.

It wasn't long, however, before commentators pointed to difficulties ahead—ones such as "the combinatorial explosion," "Godel's barrier," "brittleness," and "the mesa phenomenon." They said that all of the early successes were achieved in "limited"

terrain. About the same time, some of the climbers abandoned the goal of reaching the summit, saying it was nobler and more important to use their gear to help people along the routes. Others claimed that none of the proposed routes would “go” and that alternative routes exploiting “emergence,” “subsumption,” or something entirely different would be required.

So, the climbers re-grouped. They improved much of their climbing gear and developed some new gear and techniques—Lisp machines, Bayes networks, sophisticated search strategies, Monte Carlo methods, Walksat, default logics, POMDPs, hidden Markov models, reinforcement learning, genetic programming, and support-vector machines, among others. Indeed, many of these methods were so powerful and useful that even more climbers abandoned the climb and detoured into green valleys to use their expertise on problems in biology, business, and defense—problems that didn’t have very much to do with summiting.

Now, some fifty years after setting out for the peak, the remaining climbers are re-focusing on their original goal. Most of them agree that a combination of routes and techniques will be needed to reach the summit and that doing so might even take another fifty years or longer. More information about how the brain works is helping to inspire novel computational techniques. Much faster and less expensive computers allow more climbers to try out new strategies more quickly. Some really hard pitches have already been completed, such as automatically driving a Volkswagen 132 miles through desert terrain. My talk will explore some of these new developments and conclude with some personal guesses about promising routes ahead.

However long it takes, the summit is still there!

Short Biography

Nils J. Nilsson, Kumagai Professor of Engineering (emeritus) in the Department of Computer Science at Stanford University, received his PhD degree in electrical engineering from Stanford in 1958. He spent twenty-three years at the Artificial Intelligence Center of SRI International working on statistical and neural-network approaches to pattern recognition, co-inventing the A* heuristic search algorithm and the STRIPS automatic planning system, directing work on the integrated mobile robot, SHAKY, and collaborating in the development of the PROSPECTOR expert system. He has published five textbooks on artificial intelligence. Professor Nilsson returned to Stanford in 1985 as the Chairman of the Department of Computer Science, a position he held until August 1990. Besides teaching courses on artificial intelligence and on machine learning, he has conducted research on flexible robots that are able to react to dynamic worlds, plan courses of action, and learn from experience. Professor Nilsson served on the editorial boards of the journal *Artificial Intelligence* and of the *Journal of Artificial Intelligence Research*. He was an Area Editor for the *Journal of the Association for Computing Machinery*. He is a past-president and Fellow of the American Association for Artificial Intelligence and is also a Fellow of the American Association for the Advancement of Science. He was a co-founder of Morgan Kaufmann Publishers, Inc. Professor Nilsson is a foreign member of the Royal Swedish Academy of Engineering Sciences and is a recipient of the IEEE “Neural-Network Pioneer” award, the IJCAI “Research Excellence” award, and the AAAI “Distinguished Service” award.

Consciousness

J. Kevin O'Regan
Laboratoire Psychologie de la Perception
CNRS - Université Paris 5

Consider a robot programmed so that it *acts* in every way as though it is conscious. For example when injured, it screams and shows avoidance behavior, imitating in all respects what a human would do when in pain. The robot is able to talk about its pain, and it reasons and acts like it has the pain. The philosopher Ned Block would say that the robot has Access Consciousness of the pain.

However all this would not guarantee that to the robot, there was actually *something it was like* to have the pain. Something extra might be required for the robot to *actually experience* the pain, and that extra thing is what Ned Block calls Phenomenal Consciousness.

How can we actually implement the two kinds of consciousness in a robot? Implementing Access Consciousness would involve programming complicated intentional, decisional, and linguistic behaviors into the robot. But, theoretically at least, since all these abilities are forms of behavior, even if very complicated ones, there is no logical objection to this being possible. It may be very difficult, it may require new types of algorithms which we have not managed to build into present-day computing systems, but there is no *logical* impossibility. Indeed, artificial intelligence researchers are busy working towards this goal.

On the other hand the situation is quite different when it comes to Phenomenal Consciousness. If Phenomenal Consciousness is really what is left to explain after all possible behaviorally describable mechanisms have been invoked, then it might even be the case that this very definition prevents Phenomenal Consciousness from being studied in a scientific way, and makes it impossible to implement in a robot.

I shall put forward a way of looking at Phenomenal Consciousness which provides a solution to these problems. The new "sensorimotor" approach is a different way of conceiving what we mean by Phenomenal Consciousness. It is based on the idea that sensory experiences corresponding to Phenomenal Consciousness are not caused by the activation of internal representations of outside events, but rather, sensory experiences are *skills that we exercise*. The fact that we have physical bodies with sensors that we can actively move around, plays an essential role in the theory.

When developed the theory explains many of the important questions about phenomenal consciousness: Why there is something it is like rather than nothing it is like to have a sensory experience. Why different sense modalities like vision and hearing differ in the feels that they produce. Why feels are so hard to define and why we cannot be sure other people feel the same feels as we do.

The theory also successfully predicts certain surprising phenomena in human perception that I shall describe: change blindness, sensory substitution, illusions of ownership and certain phenomena in color psychophysics.

Ultimately I will suggest how robots with bodies can have sensory experiences and Phenomenal Consciousness like humans.

Short Biography

After studying theoretical physics at Sussex and Cambridge Universities, I moved to Paris in 1975 to work in experimental psychology at the Centre National de Recherche Scientifique where I worked on eye movements in reading. My interest in the problem of the perceived stability of the visual world led me to question established notions of the nature of visual perception, and to discover, with collaborators, the phenomenon of "change blindness". My current work involves exploring the empirical consequences of a new "sensorimotor" approach to vision and sensation in general (see my "Sensorimotor Manifesto": <http://lpp.psych.univ-paris5.fr/tikiwiki>). I am interested in applying this work to robotics. I am currently director of the Laboratoire Psychologie de la Perception, CNRS, Université Paris 5.

Artificial Cell Engineering as a Computational Problem

Norman Packard
ProtoLife Srl
Venice, IT

The European integrated project "Programmable Artificial Cell Evolution" will be reviewed with the view of analyzing the computational issues on three levels: (i) what are the computational problems inherent in the design of artificial cells? (ii) once artificial cells exist, how may they be programmed? (iii) what is the computational potential of artificial cells? Special attention will be given to the role of self assembly in the context of cells with very primitive informational (genetic) chemistry. Computation implemented by such systems will be cast as a microscopic form of morphological computation, a type of non-Turing, non-digital, computation taking place at the molecular level. Evolutionary computation techniques used to solve the engineering problems facing artificial cell design will be discussed in the broader context of burgeoning materials discovery industry.

Short Biography

Norman Packard has worked in the areas of chaos, learning algorithms, predictive modeling of complex time series, statistical analysis of evolution, artificial life, and complex adaptive systems. Packard holds a B.A. from Reed College (1976) and Ph.D. in Physics from University of California at Santa Cruz (1983). After post-docs at IHES (Bures-sur-Yvette) and IAS (Princeton), he joined the physics department at the University of Illinois, Urbana-Champaign in 1987, where he became an associate professor before leaving to become a co-founder of Prediction Company in 1991. Prediction Company's business is based on building predictive models for financial markets. It has a long-standing exclusive relationship with Union Bank of Switzerland for the implementation of a technology-based trading system based on the predictive models. Packard served as CEO of the company from 1997 to 2003, when he left his management position to become chairman of the board of directors for the company. Packard is currently working in a new scientific and business direction based on development of evolutionary chemistry in programmable microfluidic technology. Long-range applications of this technology include the fabrication artificial cells from non-living material, and their programming for useful functionality. Packard is co-founder of a new company, ProtoLife S.r.l., which aims to develop these ideas in the private sector. Packard has had a long-standing involvement with the Santa Fe Institute, currently serving on its external faculty.

The RobotCub Project: An Open Humanoid Framework for Research on Embodied Cognition

Giulio Sandini, Giorgio Metta and David Vernon
LIRA-Lab, DIST
University of Genoa, Italy

RobotCub is an international research project dedicated to the investigation of embodied cognitive systems through the ontogenic development of a humanoid robot. This robot, equipped with a rich set of innate action and perception capabilities, should develop over time an increasing range of cognitive abilities by recruiting ever more complex actions and thereby achieving an increasing degree of prospection (and, hence, adaptability and robustness) in dealing with the world around it.

To facilitate this research, we are creating an open-system 53 degree-of-freedom humanoid robot: iCub, which is the same size as a 2 year-old child (approx. 90 cm tall), will be able to sit and crawl, and engage in dexterous two-handed manipulation. The upper body has 41 degrees of freedom: 7 for each arm, 9 for each hand, 6 for the head, 3 for the torso and spine. In addition, each leg will have 6 degrees of freedom. The head comprises a 3 degree of freedom serial neck, with 2 degrees of freedom for independent vergence, and 1 degree for freedom for a common eye tilt. The hands feature underactuation of the fingers, with absolute position sensors on finger joints, tension sensors on finger tendons, and tactile sensors on each finger. The sensory system will include a binocular vision system, touch, audition, and inertial sensors.

Our research position is that cognition is created in a developmental agent-centred manner through embodied physical interaction involving exploration, manipulation, imitation, and communication. Thus, ontogenesis cannot be short-circuited: the cognitive system initially deals with immediate events and increasingly acquires a predictive capability. The process of cognitive development involves several stages, from coordination of eye-gaze, head attitude, and hand placement when reaching, through to more complex and revealing exploratory use of action. This is typically achieved by dexterous manipulation of the environment to learn the affordances of objects in the context of one's own developing capabilities. The ontogenetic process is driven by both exploratory and social motives. Exploratory drives encompass the discovery of the potential of the system's own actions and the discovery of novel regularities in the system's perception-action space. Social drives involve the identification of mutually-constructed patterns of behaviour through interaction between agents.

A key objective of the RobotCub project is to make the iCub the platform of choice for research in embodied cognition. Therefore, both the iCub humanoid robot and all embedded cognitive software will be a freely-available open system to be released under a GNU General Public Licence. Thus, the scientific community can use it, copy it, and alter it, provided that all alterations to the humanoid design and the embedded software are also made available under the RobotCub open licence. We are actively encouraging the community to use the iCub in this manner and, within a year, we will launch a competitive call for research proposals based on the iCub. Following a review, those who submit a successful proposal will be awarded a financial grant to build and

exploit a complete iCub. In addition, a Research and Training Site (RTS) is being set up as a reference site for open system integration, upgrades, releases, licensing. It will also offer training facilities and several iCubs will be available for 3rd party research.

RobotCub is funded by the European Commission, Cognition Unit, under the 6th Framework Programme, Project 004370. Further details can be found at <http://www.RobotCub.org> and <http://www.iCub.org>.

Useful signals from motor cortex

Andrew Schwartz

Motor Laboratory, Department of Neurobiology

University of Pittsburgh, USA

abs21@pitt.edu

Large portions of the brain are active during behavioral tasks. Neuronal impulses are the medium of information transmission in the CNS. The basic tenet of systems neurophysiology is to draw relationships between this widespread activity and behavioral parameters. Such principles, not yet established, could lead to a basic understanding of brain function.

In the last 25 years, it has become clear that neuronal discharge in the precentral cerebral cortex is modulated with the direction of arm movement. Each neuron fires maximally in a “preferred direction” of movement, fires minimally in the opposite direction and has cosine-graded activity in between. The firing rate-movement direction tuning is robust in terms of reliability and in our ability to find this relation in many other brain structures. When the activity patterns of multiple neurons are combined with population algorithms, movement intention can readily be extracted from these signals. The hand’s trajectory is represented isomorphically ahead of the actual hand movement in the population. This shows that these volitional arm movements are specified incrementally and continuously. The neural representation of the trajectory contains the psychophysical invariants found in natural movements related to the stereotypic slowing for curved parts of the movement (2/3 power law, law of animate movement, curvature-based segmentation). The ability to extract this isomorphic representation gives us a foundation for studying otherwise covert operations of the brain.

The lag between the extracted 'neural trajectory' and hand movement was found to be directly proportional to the local curvature of the trajectory. This lag increased (the neural signal was further ahead of the corresponding portion of the arm trajectory) as the radius of curvature decreased. The lag increase corresponded to a decrease in hand speed. When eye position was tracked during the same tasks, it was found that the eyes would saccade to the curved portion of the trajectory so that the moving hand would cross the foveal fields as it went around the high curvature part of the trajectory. This all suggests that more 'control' (information transmission) is exerted in the curved parts of the movement. If directional information has a fixed channel capacity (rate limited) through the system, then it will take longer to transmit larger quantities (bits) of information. In portions of the task requiring more information, this will be reflected in a longer latency between direction specification in motor cortex and movement execution, a general slowing of movement and an increased capacity for correction (visuomotor).

Visual information is thought to traverse a number of cortical structures before reaching the motor cortex. The ventral premotor cortex is a structure just preceding the motor cortex in this chain. We designed a task to dissociate vision of the moving hand from the actual movement of the hand. In this motor illusion task, subjects moved their hands in a virtual reality environment along a visualized oval. As repeated circuits of the object were made, the gain between the hand movement and visual cursor was gradually increased in the horizontal dimension. Subjects unconsciously adjusted the movement of their arm movement to a circle to maintain the visualized oval shaped

cursor movement. The visualized, perceived movement was of an oval, while the actual movement of the cursor was circular. Neural trajectories extracted from the motor cortex were circular, but those from the ventral premotor cortex were oval-shaped. The motor cortex extractions matched those of the actual movement. The premotor cortex signals reflected the perceived movement.

The signals we extract from the motor cortex contain a high fidelity representation of arm movement. This is an ideal control source for prosthetic arm movement. The major technical hurdle to overcome in this regard, is the necessity of recording from a population of neuronal units simultaneously. Technological advances in electrode arrays, electronic signal conditioning and computational performance have now made it possible to demonstrate the feasibility of this approach.

Using a virtual reality display, monkeys were trained to move their hands from the center of a cube to its corners in separate movements. The 3D position of the hand was tracked in real-time as populations of single units were recorded. Tuning functions relating movement direction to discharge rate were calculated for each unit. The monkeys then performed the task in the absence of any movement using their recorded brain signals that were processed in 30 ms intervals to generate an extracted population signal reflecting the animal's intention to move the cursor. Monkeys could perform this task very well. The brain-controlled cursor moved almost as fast as when hand-controlled. The animal was able to move to novel targets immediately, control the speed of the cursor and to move it in arbitrary directions. The control was robust, with the animal typically performing for more than an hour with about a 90% success rate for several weeks. When carrying out these experiments, we found that the directionality (preferred directions) of the recorded units would change across the transition from hand-control to brain-control. Although the reason for this is unknown, the way these changes took place in individual units was consistent across days.

The apparent change in preferred direction was also a major finding in another study. Using the 3D VR brain-control paradigm, we first had monkeys acquire a high performance level. Then, we took a subset of the recorded units and assigned new preferred directions to them by rotating their original preferred directions in the extraction algorithm. This led to an altered cursor movement when the animal initially used brain control to reach the displayed targets as the movements were rotated by the reassignment. The animals quickly recovered from this perturbation and were moving the cursor to the middle of the targets within a few minutes. When the preferred directions of the recorded units were measured throughout the task they were seen to rotate to compensate for the errant movement. Interestingly, this rotation was more evident in the units that had reassigned preferred directions, even though the movement was generated from the total population. Somehow the units responsible for the perturbation were more likely to generate the recovery to that perturbation. This shows that learning is taking place as the monkey sees the altered behavior generated from the extraction process and changes the way the recorded units fire during the task to compensate for the imposed error.

As a final demonstration of prosthetic control, the brain-control signal was used to control a physical device. This prosthetic arm consisted of 4 DOF arranged anthropomorphically with a binary gripper at the end. The shoulder of the device was placed near the monkey's own shoulder and the animal was able to operate the arm in 3D space to reach out, grasp a piece of food and bring it to its mouth. This shows how the basic directional properties of cortical units can be used to better understand neural movement generation principles and how they can be used to control practical devices.

Short Biography

Andrew Schwartz received his PhD in physiology from the University of Minnesota in 1984, followed by a fellowship with Apostolos Georgopoulos in 1987, where they studied motor cortical representations of reaching. Currently he is a professor at the University of Pittsburgh where he continues to investigate cortical control of arm movement which has extended to the development cortical prosthetics.

The Rise of Robots – Machines Sharing the Environment with Natural Creatures

Roland Siegwart

Institute of Robotics and Intelligent Systems, Autonomous Systems Lab

ETH Zurich, Switzerland

rsiegwart@ethz.ch <http://www.asl.ethz.ch/>



Robots are rapidly evolving from factory “workhorses” that are physically bound to their work-cell to machines evolving in our environment and interacting with natural creatures. This development makes systems obviously much more complex and raises important new questions that are meanwhile addressed by a very interdisciplinary research community. The fundamental questions linked with the development of such robots and systems span from locomotion and control, to natural interaction and social intelligent up to functional / semantic representations.

In this talk various projects of the speaker’s research group will be presented and discussed. It starts with tiny and simple robots building mixed societies with cockroaches. Based on simple interaction rules and behaviors, the robots are able to significantly influence the collective decision process of the mixed society in test environment. However, these promising results can unfortunately not be simply scaled up to more complex environments and interactions. The fundamental problem lies in the fact, that collective intelligent is essentially limited by the competences and performance of the individuals of the society. Future robots able to integrate in human society and offer useful services require first of all the perception and representation capacities that can cope with complex settings. This has driven our research in the context of the COGNIRON project towards functional-based environment representations, which we consider as fundamental for higher cognitive functions. We suggest in a first step a hierarchical probabilistic representation of space that is based on objects arranged in topological object graph. In our most recent EU project BACS (Bayesian Approaches to Cognitive Systems) we try in a consortium of ten partners to further enhance and apply the Bayesian approach for solving complex cognitive tasks. Our

research and future directions are strongly influenced by the long-term robot experience we made at the Swiss National Exhibition in 2002 and the more recent work on a theater robot and intelligent cars.

Short Biography

Roland Siegwart is Full Professor for Autonomous Systems at ETH Zurich. He received his M.Sc. ME in 1983 and his Doctoral degree in 1989 at ETH Zurich. After his Ph.D. studies he spent one year as a postdoc at Stanford University where he was involved in

micro-robots and tactile gripping. From 1991 to 1996 he worked part time as R&D director at MECOS Traxler AG and as lecturer and deputy head at the Institute of Robotics, ETH. From 1996-2006 he was a full professor for autonomous systems and robots at the Ecole Polytechnique Fédérale de Lausanne (EPFL), and 2002-2006 also vice-dean of the School of Engineering. He was the funding chairman of the Space Center at EPFL and deputy director of the national center of competence in research on Interactive Multimodal Information Management. Roland Siegwart is strongly involved in EU projects spanning from cognitive science to intelligent cars. He leads a research group of around 35 people working in the field of robotics, mechatronics and product design. He is an active member of various scientific committees and co-founder of several spin-off companies. He was/is the general Chair of IROS 2002 and AIM 2007, and Vice President for Technical Activities (2004/05) and distinguished lecturer (2006/06) of the IEEE Robotics and Automation Society.

Biologically Inspired Miniature Robots

Metin Sitti

NanoRobotics Laboratory, Department of Mechanical Engineering
and Robotics Institute

Carnegie Mellon University, Pittsburgh, PA 15213, USA

sitti@cmu.edu

Inspiration by nature has enabled alternative design ideas for new materials, structures, systems, and control and communication methods since the beginning of the human history. Biological systems have evolved to find just-good-enough solutions to survive. By understanding and adapting the underlying principles of these solutions to engineering systems, new miniature robots which can operate in unstructured environments robustly and efficiently are investigated in this work. Especially considering the unknowns and complexity of the physics and dynamics at the micro and nanoscales, biological inspiration could have a significant inspiration source to develop new miniature robot designs. On the other hand, the developed bio-inspired micro/nano-robots would be used to understand nature better at the small scales.

First inspiration by nature in this study is repeatable attachment mechanisms to develop robust and agile miniature climbing and crawling robots in complex environments such as inside human intestines and smooth and rough wide range of surfaces on earth and in space. Mechanical interlocking used by plant hooks, insect and lizard claws and human fingers, suction-cups used by octopus under water, micro-structured or smooth foot pads with a secreted oil used by ants, cockroaches, crickets, and tree frogs, and dry micro and nano-structured foot-hairs by geckos, anoles and spiders are possible repeatable attachment mechanisms in nature. Each mechanism works robustly and efficiently for different surface roughness (from centimeter scale down to atomic scale), types (hardness and hydrophobicity) and conditions (dirty, wet or dry). Using the mechanical interlocking and dry micro/nano-fibers, geckos are very agile and robust climbers on wide range of surface materials, roughness and even dirt. Understanding the principle of gecko foot-hair adhesion on smooth and micro/nano-scale rough surfaces using interatomic surface forces such as van der Waals forces and hierarchical compliance, synthetic gecko adhesives are proposed to be analyzed, designed and fabricated. Prototype polymer micro-fibrillar adhesives are fabricated using micro-molding and optical lithography techniques. Fabricated polyurethane and silicone rubber micro-fibers can enhance adhesion and adapt to micron scale surface roughness. Using these micro-patterned dry adhesives, tank, tri-foot wheels, and legged type climbing robots are designed, built, and demonstrated for climbing on smooth and micron scale rough surfaces. Furthermore, by coating these micro-fibers with silicone oil similar to beetle foot-hairs, an endoscopic capsule robot inside intestines is designed and demonstrated. This miniature biomedical robot with an on-board camera can open its legs with the adhesive foot-pads to attach to the intestinal wall when desired for minimally invasive imaging, biopsy, local drug delivery, and surgery in intestines.

As the second topic, miniature robots with legged locomotion on water are proposed inspired by water striders and basilisk lizards. Water striders with 10s of milligram weight can stay on water surface using surface tension based lift forces due to their very hydrophobic hairy supporting legs and can move on water up to 1.5 m/s speeds by

rowing two side legs. Modeling, optimal design, manufacturing, and control results and issues of a Water-Walker robot inspired by these insects are presented. Current robot prototype with optimized leg shape can lift weights up to 9.5 gr and move fast on water surface using Teflon coated stainless micro-wire legs. On the other hand, basilisk lizard with 10s or 100s of gram weight uses very fast rotation of its two legs with a specific elliptic trajectory at 5-10 Hz frequencies. By slapping and stroking their feet into the water, the lizard affects a momentum transfer which provides both forward thrust and lift. The design of a bio-inspired Water-Runner robot utilizing similar principles is discussed, modeled, and prototyped. Functionally, the robot uses a pair of identical four bar mechanisms, with a 180° phase shift to achieve bipedal locomotion on the water's surface. Computational and experimental results are presented and reviewed with the focus being a maximization of the lift to power ratio. After optimization, two legged models can experimentally provide 12-15 g/W of lift while four legged models can provide 50 g/W of lift. This work opens the door for biped and quadruped robots to become ambulatory over both land and water, and represents a first step toward studies in amphibious stride patterns; step motions equally conducive to propulsion on water and land.

Finally, swimming micro-robots are proposed inspired by *E. Coli* and *S. Marcescens* bacteria. These bacteria dominantly use their flagella to propel themselves efficiently at very low Reynold's numbers such as 10^{-4} . Modeling the micro-fluidics of these around 20 nanometer diameter and around 10 micrometer long helical flagella, bio-inspired synthetic helical flagella are proposed. Using a scaled up flagellar propulsion test setup and a scaled up swimming bacterial robot, effect of flagella geometry, number and rotation speed, and torque are demonstrated. Currently, micro-scale swimming robots are being designed and built by external magnetic field based actuation and hybrid integration of biological bacteria to an inorganic robot body. The latter new hybrid approach would enable a steerable micro-robot with a diameter down to 10s of micrometers in the future.

All of the above miniature biologically inspired robots with their many open research challenges can directly access to small spaces and scales and can be very agile, distributed, massively parallel, light weight, and inexpensive. Tens or hundreds of these miniature robots would enable distributed system and swarm robotic system platforms in the future. They could revolutionize health-care, environmental monitoring, space exploration, search and rescue, and entertainment applications in the future.

Short Biography

Metin Sitti received the BSc and MSc degrees in electrical and electronics engineering from Bogazici University, Istanbul, Turkey, in 1992 and 1994, respectively, and the PhD degree in electrical engineering from the University of Tokyo, Tokyo, Japan, in 1999. He was a research scientist and lecturer in the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley during 1999-2002. He is currently an assistant professor and the director of the NanoRobotics Laboratory in Department of Mechanical Engineering and Robotics Institute at the Carnegie Mellon University. His research interests include micro/nano-robotics, biologically inspired micro/nano-systems, MEMS/NEMS, and biomedical micro/nanotechnology. He received the National Science Foundation CAREER award

in 2005 and the Struminger award for his micro/nanoengineering teaching activities. He is elected as the Distinguished Lecturer of the IEEE Robotics and Automation Society for 2006-2007. He received the best paper award in the IEEE Robotics and Biomimetics Conference (2004), the best paper award in the IEEE/RSJ International Conference on Intelligent Robots and Systems (1998), and the best video award (2002), the best student paper nomination (2001), the best manipulation paper nomination (2004), and the best paper nomination (2000) in the IEEE Robotics and Automation Conference. He is the chair of the IEEE Nanotechnology Council, Nanorobotics and Nanomanufacturing Technical Committee and the IEEE Robotics and Automation Society, Rapid Prototyping in Robotics and Automation Technical Committee. He is an editorial board member of Journal of Micromechatronics, Journal of Nanoscale Science and Engineering, and International Journal of Control, Automation, and Systems.

Artificial Intelligence and the Origins of Symbolic Culture

Luc Steels

VUB AI Lab and Sony Computer Science Lab, Paris

AI has made incredible progress the past decades by studying intelligence from an information processing view, that is by asking the question what kind of information processes can achieve planning, language communication, vision, problem solving, memory, etc. AI has not simply used computer science as the foundation for this investigation but considerably expanded computer science with new programming languages, new models of computation, new forms of memory organisations, new ways of distributing computation, etc. It has not only produced a host of ideas on the information processing underlying intelligence but translated this into workable technologies so that we can now build applications, like robots using model-based vision to navigate in real time in the real world based on visual sensing, which were unthinkable even ten years ago. The impact on other disciplines has been dramatic. The information processing view has aided neuroscience to move away from a pure biochemical/physiological investigation of the brain to develop a 'computational' neuroscience that studies how the brain is able to do the kind of massive complex information processing that AI models predict should be there. It has aided the mind sciences (linguistics, psychology, anthropology) to move away from mere descriptive research to information processing models of intelligence. So the information processing level has emerged as the appropriate way to bridge the gap between mind and brain and thus to link together mind science and brain science.

One could say that the most rapid advances in AI were achieved in its first three decades because so little existed before. In the nineties, the occasionally excessive focus on information processing was corrected by bringing back the body and the ecological environment as a constraining and enabling force of intelligence, and the exclusive focus on the single individual gave way to multi-agent systems, so that the social shaping of intelligence became another enabling force.

In this talk I focus on the kind of research that I have been pushing through with my team more recently, building up on decades of earlier AI achievements. We have focused on the question of the origins and evolution of symbolic culture, which is both the foundation and the unique province of human intelligence. I will survey some of our recent experiments in which groups of situated and embodied agents autonomously bootstrap communication systems which have similar properties as those found in human languages, particularly grammar and perspective reversal. This extraordinarily challenge has not only yielded a new view on (symbolic) intelligence, which emphasizes dynamics and becoming rather than static end-states of competence, but also many new ideas on the architecture of intelligent systems and how fluidity and constant adaptation can be achieved.

The Future of Driving

Sebastian Thrun
Artificial Intelligence Laboratory
Stanford University, USA
thrun@stanford.edu

Every year, we lose more than a million people worldwide to traffic accidents; 43,000 in the U.S. alone. Cars also make us lose time: the average U.S. worker spends more than an hour per day in her car. Self-driving cars carry the promise of making highways safer and increasing people's productivity. They will also increase the throughput of highways, and enable older adults to sustain their independence longer. But most importantly, making cars drive themselves is an ultimate robotic challenge, one that requires real advances in real-world AI.

This presentation will report about a recent research on autonomous cars. The speaker will start with the DARPA Grand Challenge, a robot competition created by the U.S. Government to spur innovation in autonomous driving. The goal was to build a robot that could drive 131 miles through unrehearsed desert terrain in ten hours or less, based on a coarse GPS-referenced description of the course. While in 2004, none of the competing vehicles mastered more than 5% of the total distance, five vehicles finished in 2005, with Stanford's Stanley robot coming in first. To many, this was a watershed moment for AI: within a year, the AI community achieved what many experts thought to be impossible within the next decade: building competent self-driving cars!

The winning vehicle, Stanley, was an AI robot. It extensively relied on some of the best AI has to offer: probabilistic reasoning, machine learning, computer vision, and real-time planning and control. However, none of the winning robots are quite ready yet for deployment. They still have to learn how to drive safely in traffic, how to drive fast, and how to handle environments as diverse as cities, suburban neighborhoods, and highways. This creates an enormous challenge for the field of AI, one that will ultimately impact all aspects of society.

The speaker's research group has recently begun working on these problems. He will discuss research on precision localization, high-speed navigation, and probabilistic modeling of traffic. He will discuss his work in the perspective of the overall objective of advancing society through self-driving cars, and point out new research directions for AI. Let's get AI into the driver's seat!

Short Biography

Prof. Dr. Sebastian Thrun is the Vance D. and Arlene C. Coffman Associate Professor of Computer Science and Electrical Engineering at Stanford University, where he also serves as the Director of the Stanford AI Lab. His research focuses on robotics and artificial intelligence. Thrun has delivered numerous invited plenary presentations at leading conferences and symposia. He also served as the inaugural general conference chair of Robotics Science and Systems (RSS) 2005; general conference chair of Neural

Information Processing Systems (NIPS) 2003 and Conference on Automated Learning and Discovery (CONALD) 1998; co-chair of the International Symposium of Robotics Research (ISRR) 2005; and program chair of NIPS 2002. He is a founding member of RSS and Vice President of the NIPS Foundation. Thrun has been included in the "Brilliant Ten" by Popular Science. One of his robots has been named the "Best Robot Ever" by Wired Magazine.

The development of gaze control in human infants

Claes von Hofsten and Kerstin Rosander
Uppsala University, Sweden

Gaze control serves two functions: switching gaze between visual targets and stabilizing gaze on a particular target. Switching gaze is crucial for looking because visual acuity is so much higher in the centre than in the periphery of the visual field. Stabilizing gaze is equally important. The slightest movement of the image on the retina deteriorates acuity. Stabilizing gaze involves both head and eye movements and is guided by at least three kinds of information: visual, vestibular and proprioceptive. It is a complicated dynamic problem. Both eye and head movements are involved in maintaining gaze on the object attended to and the movements of these body parts must be timed and scaled to each other. The head also moves for other reasons and these movements must be compensated for. Because of the information processing lags in the nervous system and the mechanical lags of the effectors, all these adjustments must be predictive to avoid the gaze to lag the target.

Infants develop such gaze stability during the first half year of life. Two relatively independent control mechanisms are involved. One of them has the task of maintaining a stable gaze in space while moving around and it is primarily controlled by vestibular information. This system matures during the first month of life. The task of the other one is to stabilize gaze on a moving objects and it is primarily controlled by visual information. This system matures during the 2nd and 3rd months of life. The vestibular and the visual systems for stabilizing gaze are not coordinated to begin with and their developments differ. Because of this, problems arise when both the subject and the object moves and when the subject begins to track object with both head and eyes. There are two kinds of problems that need to be solved. First, body movements that are not part of the tracking need to be compensated but not the movements that are part of the tracking effort. Secondly, the movements of the head and eyes must be coordinated.

Intelligent Systems as Mediators in Human Communication

Alex Waibel

Department of Computer Science
Carnegie Mellon University and Universität Karlsruhe
ahw@cs.cmu.edu

Introduction

Effective communication is perhaps the most visible and important ability for intelligent agents to operate in a physical world, to form relationships with and learn from other agents and to form societies that magnify and empower the abilities of the individual. As the importance of computer communication has been recognized, we have seen considerable interest in this area of AI, advancing the state of the art from basic command & control speech devices all the way to sophisticated modern multimodal perceptual interfaces. These interfaces can no longer be seen as mere input devices, that attempt to replace the keyboard. Rather they emerge as intelligent participants and mediators in human-human as well as multi-agent interaction. This broader view of computer interfaces is proactive, social, and engaged and offers tremendous opportunities for numerous practical applications.

In this presentation, I will discuss two grand challenge problems, that instantiate this view and turn computer devices into mediators and proactive assistants for human-human interaction. The first is cross-lingual communication assisted/mediated by machine. The latter, Computers in the Human Interaction Loop (CHIL), i.e., computer that provide services based on the perceived observed implicit needs of humans, rather than explicit interaction.

Cross-Lingual Communication

One of the human-human communication grand challenges is given by a growing need for cross-lingual communication devices that provide translation of or information access to speech, text, images (road-signs), to cross the linguistic divide. Computer systems should also attempt to do this in a human-centered fashion so as to turn the linguistic separation transparent and enable free unincumbered human-human interaction. With the globalization of society, and the opportunities and changes that go along with it, there is an ever more pressing need for computer assistance in this area.

Delivering effective cross-lingual support, however, is a formidable technical challenge for Computer Science and AI. The pure component technologies, Speech, Machine Translation, Natural Language Understanding, are all still only partially solved and remain active areas of research. Performance still falls far short of human performance in both error rates and robustness. Small deviations in application domain, variability or noise conditions have dramatic effects on overall system performance. Still, research on integrated systems that include such components is moving forward and expanding at a rapid pace. Speech translation, in particular, represents the combination of several difficult problems (recognition, translation, synthesis) and has grown from curiosity and feasibility studies into the now best funded and most intensely pursued research goal in human language technologies. Early systems were restricted to small domains, limited vocabularies, speaker dependent, and could accept only well-formed, syntactically well-formed (read) speech. Current efforts

are beginning to see the emergence of first domain-unlimited, spontaneous speech translation systems that could provide translation assistance in complex domains, such as lectures, political speeches, and broadcast news. Work on suitable human-centered interfaces has also progressed to deliver output from such translation function in various forms, including targeted audio speakers, translation goggles, video screens, and PDAs.

Despite considerable advances, a number of research problems remain to be addressed. These problems all affect performance but also cost of building and maintaining cross-lingual systems: 1.) Effective handling of tonality, morphology, orthography and segmentation, 2.) Foreign accents and foreign words/names introduce modeling difficulties, 3.) the sheer number of languages (~6000 by most estimates) makes traditional porting approaches via training on large corpora or rule based programming impractical and prohibitively expensive, 4.) Scaling multilingual and cross-lingual technologies such as Machine Translation and Cross-lingual Retrieval to all possible language pairs (N^2) and domains, dialects, languages and applications.

In my talk I will describe and demonstrate some of the emerging applications and systems, and will discuss scientific problems and approaches on the road to more effective, more robust and more portable multilingual language assistance and services.

Computers in the Human Interaction Loop (CHIL)

After first building computers that paid no intention to the communication with humans, we have in the past decades developed ever more sophisticated interfaces that provided human-machine interaction by putting the "human in the loop" of computers. These interfaces have improved usability by providing more appealing output (graphics, animations), more easy to use input methods (mouse, pointing, clicking, dragging) and more natural interaction modes (speech, vision, gesture, etc.). Yet the productivity gains that have been expected have largely not been seen and human-machine interaction still remains a partially frustrating and tedious experience, requiring excessive attention to the direct interaction with the machine and to dealing with technical artifacts. We must therefore consider a third paradigm of computer use, in which people increasingly interact with people, and move the machine into the background; where machines observe human activities and provide services implicitly, and -to the extent possible- without explicit request.

We call this second grand challenge problem "Computers in the Human Interaction Loop (CHIL), i.e., a class of computer services that attempt to provide assistance to humans *implicitly* and *proactively* based on a full perceptual description and understanding of human-human events, interactions, activities and ultimately needs. Each of these services relies on a number of sophisticated perceptual technologies (speech, vision, etc.) that are only now becoming possible. Putting the "Computer in the Human Interaction Loop" (CHIL), instead of the other way round, frees humans to perform other tasks. It's realization though is another formidable technical challenges. The machine must now *always* observe and understand humans, model their activities, their interaction with other humans, the human state as well as the state of the space they are in, and finally, infer intentions and needs. >From a perceptual user interface point of view, we must process signals from sensors that are always on, frequently inappropriately positioned, and subject to much greater variability. We must also not only recognize WHAT was seen or said in a given space, but also a broad range of additional communicative information, such as the WHO, WHERE, HOW, TO

WHOM, WHY, WHEN of human interaction and engagement. To achieve suitable performance, these and other technologies are currently being developed. They are being advanced and accelerated by international benchmarking and evaluation campaigns, that emerge in the community. Finally, integrated services are implemented and demonstrated at various research labs around the world. In this talk, I will describe a variety of multimodal interface technologies that we have developed to answer these questions and demonstrate several CHIL type services that take advantage of such perceptual interfaces.

Archipelago.ch: The Dynamic Diorama

Adrianne Wortzel
Division of Design and Technology
New York City College of Technology
City University of New York, USA
sphinx@camouflagetown.tv

This talk puts forward the function of narrative as a sub-field of creative robotics, and attempts to show that the creation of narrative scenarios in parallel with developing research in artificial intelligence is a useful tool when embedded in the methodologies of scientific and engineering research. Dramatic narrative is dynamic and, when built as a ventilated system, can be grown alongside, and integrated with, the myriad and changing paths of research. A dramatic scenario developed parallel with research both provides and receives a layering of context and meaning which amplifies both the science and the art.

As an example I will present an ongoing project entitled "archipelago.ch" which emanates from robotic research developed at the Artificial Intelligence Laboratory of the Department of Informatics, University of Zurich, Switzerland. By working with original bottom-up developing robotic systems, such as those originated in the AILab, artists can abandon sculptural or choreographic concerns to develop a dramatic scenario which is true to the capabilities of a particular robot or robotic system. Such scenarios, hand-in-hand with being artistic expressions, can have the potential to re-enter and inform the science from which they emerge.

The above text is paraphrased from a paper, "Narrative in Robotics Scenarios for Art Works," written collaboratively by Dr. Daniel Bisig, Senior Researcher, AILab and myself, and presented and published at the Proceedings of the Symposium on Robotics, Mechatronics and Animatronics in the Creative and Entertainment Industries and Arts, AISB2005 at the SSAISB Convention, University of Hertfordshire, Hatfield, UK.

Short Biography

Adrianne Wortzel's art works explore historical and cultural perspectives using new technologies to create interactive robotic and telerobotic installations and performance productions in both physical and virtual networked environments.

An in-progress *Archipleago.ch*, is a film depicting a "galapagos" where the indigenous creatures are the robots created by resesarchers at the Artificial Intelligence Laboratory, University of Zurich. This residency was made possible by a 2004 Artists-in-Labs;

Recent works include *Eliza Redux* (<http://elizaredux.org>), a physical robot offering virtual psychoanalytic sessions emulating Joseph Weizenbaum's ELIZA program, created in collaboration with Michael Schneider (Artist, Physical Computing-NYU) and Robert Schneider (Professor of Computer Science-Lehman College). *The Veils of Transference*, a video of a pre-scripting psychoanalytic session between a human and a

robot, *Camouflage Town*, a telerobotic installation exhibited in *Data Dynamics* at The Whitney Museum of American Art (Spring 2001). *Sayonara Diorama*, (1998) a play with human and robotic actors; the tale of a fictive second Voyage of the Beagle by Darwin thirty years after the first; *The Ship's Detective*, in Cooper Union's *Technoseduction* exhibition (1997); *The Hidden Archivists at the Anchorage* at *Creative Time's Art in the Anchorage* (1997); *NoMad is An Island* at *Ars Electronica97* (Linz, Austria); and *Tableaux Vivant Dan Une Monde Parfait* in *Areale 99* (Baitz, Germany). Her work is documented at <http://artnetweb.com/wortzel>.

Her recent projects have been made possible by support from the Swiss Artists-in-Labs Program, the Artificial Intelligence Laboratory at the University of Zurich, the Franklin Furnace Fund for Performance Art Award, the PSC-CUNY Research Foundation, the National Science Foundation, and the Greenwall Foundation.

She is a Professor of Communication Arts at New York City College of Technology, CUNY, on the Doctoral faculty of the Interactive Technology and Pedagogy Certificate Program of the CUNY Graduate Center, and an Adjunct Professor of Mechanical Engineering at the Cooper Union for the Advancement of Science and Art where she is also the Director of *StudioBlue*, a telerobotic arena for performance productions. She is currently building another StudioBlue at Citytech.

50th Anniversary Summit of Artificial Intelligence



Poster Abstracts

Poster Abstracts

Chemistry of Irrationality: Towards the Science of Doubt and Confusion

Andrew Adamatzky

Morphodynamics and Perception: A Conceptual Stance and Some Clues from Evolutionary Robotics

Fernando Almeida e Costa, Ian Macinnes, Inman Harvey

Enculturation Without Language

Sepand Ansari, Caro Lucas

The Man-Machine Interaction: FMRI Study for an EMG Prosthetic Hand with Biofeedback

Alejandro Hernandez Arieta, Hiroshi Yokoi

Extending the Robot's Operational Time and Space

Lijin Aryananda

Cognition Network Technology: A Novel Knowledge Based Approach for Analyzing and Processing Complex Data. Method and Application Samples in Image Analysis

Martin Baatz, Arno Schäpe, Günter Schmidt, Maria Athelougou, Gerd Binnig

Communication With and Within Artificial Societies

Fatmah A. Baothman

Humanoid Robots — A Tool to Understand Human Intelligence

Sven Behnke

Perspective Taking: An Organizing Principle for Learning in Human-Robot Interaction

Matt Berlin, Jesse Gray, Andrea L. Thomaz, Cynthia Breazeal

A Human-like Robot Torso ZAR5 with Fluidic Muscles: Full Controlled by Using Data Suit and Gloves

Ivo Boblan, Rudolf Bannasch, Andreas Schulz, Hartmut Schwenk

Engineering Modular Interfaces in Neural System Design

Martin Boerlin, Paul F. M. J. Verschure, Tobi Delbruck, Kynan Eng

A Robot That Builds a Simulator of Itself: A Counterrevolutionary Tale

Josh Bongard, Victor Zykov

Ants, Bees, Humans and Other Social Natural and Artificial Intelligent Autonomous Agents

Fabio P. Bonsignorio

Can “Intentional” Behaviors Emerge From Random Associations?

Simon Bovet

Social Human-Robot Interaction

Cynthia Breazeal

Whither Embodiment?

Andrew Brooks

Adaptive Dynamical Systems: A promising Tool for Embodied Artificial Intelligence

Jonas Buchli, Ludovic Righetti, Auke Jan Ijspeert

Teaching Cognitive Robotics using Multiple Online Remote Robots

Catalin Buiu, Nicolae Moanta

Mysterious Tools

Seth Bullock

Who Made the Decision?! An Investigation into Distribution of Memory, Actions and Decision Makings in a Multi Agent System

Roozbeh Daneshvar, Caro Lucas

AI Re-Emerging as Research in Complex Systems

Kemal A. Delic, Umeshwar Dayal

Anthropocentricity and the Social Robot: Artistic and Aesthetic Investigations into Machine Behaviours

Louis-Philippe Demers

The Humanoid Robot Project – Holistic Movement

Raja Dravid, Gabriel Gomez, Peter Eggenberger Hotz

AI Approaches For Next Generation Telecommunication Networks

Gianni A. Di Caro, Frederick Ducatelle, Luca M. Gambardella

Evolutionary Humanoid Robotics: First Steps Towards Autonomy?

Malachy Eaton

A Robotic Model for Rat Tactile Sensing

Miriam Fend, Rolf Pfeifer

Applying Data Fusion in a Rational Decision Making Architecture of a Believable Agent

Benjamin Fonooni, Behzad Moshiri, Caro Lucas

Beyond Information Transfer: The Emergence of Embodied Communication

Tom Froese

Dynamic Meta-Learning

Matteo Gagliolo

Ant Colony Optimization for Advanced Logistics Systems

Luca Maria Gambardella

Investigating the Interplay Between Morphology and Behavior Using the Same, but Adaptive Neural Controller

Gabriel Gomez, Peter Eggenberger Hotz

The Communication and Commercialization of Innovation: Observations and Implications

Simon Grand

Complex Behaviour Recognition in Interacting Agents

Verena V. Hafner

Adaptive Multi-Modal Sensors

Kyle I. Harrington, Hava T. Siegelmann

From Artificial Intelligence to Artificial Life in the First 50 Years; To Homeostasis and Entropy Production in the Next 50?

Inman Harvey

Acting Lessons for Artificial Intelligence

Guy Hoffman

Right Body, Right Mind

Owen Holland

Analog-Digital Model of Cortical Memory Consolidation

Lars E. Holzman, Hava T. Siegelmann

Designing Physical Impedance for Engineering Adaptive Intelligence

Koh Hosoda

In the Search of Principles Underlying Cognitive Phenomena

Martin Hülse, Keyan Zahedi, Steffen Wischmann, Frank Pasemann

From Locomotion to Cognition: Cheap Design Approach to Adaptive Behavior

Fumiya Iida, Gabriel Gomez, Rolf Pfeifer

Interactive Intent Imitation for Humanoid Robots Based on Dynamic Attention Prediction and Control

Tetsunari Inamura, Naoki Kojo, Kazuyuki Sakamoto, Masayuki Inaba

A Situated, Embodied and Dynamical Systems Approach to Understanding Learning and Memory

Eduardo Izquierdo-Torres, Inman Harvey

The Sincere Gatekeeper

Terence Jardine

The Progress Drive Hypothesis

Frédéric Kaplan, Pierre-Yves Oudeyer

Acquirement of Body Schema Using Intelligent EMG Prosthetic Hand

Ryu Kato, Alejandro Hernandez Arieta, Kojiro Matsushita, Hiroshi Yokoi

What is a Perception-Action System? From Agency to Animality

Fred Keijzer, Daan Franken, Marc van Duijn

Implications of Embodiment and Situatedness on the Structure of Simulated Fish Schools

Hanspeter Kunz, Charlotte K. Hemelrijk

High Performance Knowledge Systems: from PUFF to ICE

John Kunz

How Morphology Affects Learning of a Controller for Movement

Nathan Labhart, Shuhei Miyashita

A Formal Definition of Intelligence for Artificial Systems

Shane Legg, Marcus Hutter

Towards Programmable Self-Assembly of Microrobot Ensembles

Lukas Lichtensteiger

Information Processing as Embodied Interaction

Max Lungarella, Olaf Sporns

The Promises and Challenges of Socially Assistive Robotics

Maja J. Matarić, Adriana Tapus

Designing Dynamical Stable Locomotors Across the Reality Gap – Introduction of Edutainment as Heuristic-Diversity Design

Kojiro Matsushita, Hiroshi Yokoi, Tamio Arai

Technical Proposal for an Active Floating Element

Shuhei Miyashita, Maik Hadorn, Peter Eggenberger Hotz

Does a Synesthetic Mechanism Aid Robot's Language Learning?

Yukie Nagai

Tactile Manipulation of Unknown Objects

Lorenzo Natale, Eduardo Torres-Jara

Fugue – an AI Artwork for the 21st Century

Gordana Novakovic, Peter Bentley, Anthony Ruto, Rainer Linz

Computational Auditory Scene Analysis – Towards Listening to Several Things at Once

Hiroshi G. Okuno

Artificial Death for Attaining System Longevity

Megan Olsen, Hava Siegelmann

CAVIAR – Building Blocks for a Large-Scale Spiking Vision System

Matthias Oster, Raphael Berner, Patrick Lichtsteiner, Rafael Paz-Vicente, Rafael Serrano-Gotarredona, Alejandro Linares-Barranco, Tobi Delbrück, Philipp Häfliger, Rodney Douglas, Anton Civit, Bernabe Linares-Barranco, Shih-Chii Liu

Simple Linear Spring Is Not Sufficient? Stabilizing Passive Dynamic Running by Exploiting the Intrinsic Body Dynamics

Dai Owaki, Akio Ishiguro

Formalizing Emergence to Accelerate the Accretion of Embodied Intelligence

Frank Pasemann

The Anticipatory Nature of Representations

Giovanni Pezzulo, Gianluca Baldassarre, Rino Falcone, Cristiano Castelfranchi

Intelligence from Information Structure

Daniel Polani

The Role of Neurobiology in Artificial Intelligence: To Make It Less Artificial

Steve M. Potter

An Extension to Subsumption Architecture Inspired by Human Nervous System

Yoosef Ramezani, Behzad Moshiri, Caro Lucas

Design Methodologies for Central Pattern Generators: Toward “Intelligent” Locomotion in Robots

Ludovic Righetti, Auke Jan Ijspeert

Evolutionary Robotics and Perceptual Supplementation: Dialogue Between Two Minimalist Approaches

Marieke Rohde, Ezequiel Di Paolo

Adaptive Mechanics: Compliant Legged Locomotion

Juergen Rummel, Andre Seyfarth, Fumiya Iida, Elmar Dittrich

Self-Organization of Complexity in Embodied Agents

Alexander Schmitz, Fumiya Iida

Applying OWA Operator in a Multi-Agent Architecture of an Emotional Robot

Mohamed Shayganfar, Behzad Moshiri, Caro Lucas

A Philosophy of Designing Human Interactive Robot for Psychological Enrichment and Robot Therapy

Takanori Shibata

How Should the Body and the Brain be Coupled? A Robotic Case Study with a Modular Robot

Masahiro Shimizu, Takafumi Mori, Toshihiro Kawakatsu, Akio Ishiguro

Active Vision for Neural Development and Landmark Navigation

Mototaka Suzuki, Dario Floreano

Non-Linear Models for the Dynamics of Motivation

Rene te Boekhorst

Towards an Experience-Based Artificial Intelligence

Ricardo A. Téllez, Oscar Vilarroya

An Ultimate Goal of AI and Tasks for Four Periods

Jiabei Wu

“Cheap” Underwater Locomotion: Roles of Morphological Properties in Underwater Locomotion

Marc Ziegler, Fumiya Iida, Rolf Pfeifer

Bioinspired Indoor Microflyers

Jean-Christophe Zufferey, Dario Floreano

Chemistry of Irrationality: Towards the Science of Doubt and Confusion

Andrew Adamatzky

Faculty of Computing, Engineering and Mathematical Sciences
University of the West of England, Bristol, United Kingdom
andrew.adamatzky@uwe.ac.uk

Brains do not compute. They hardly ever rationalize. Being predominantly far from mental equilibrium, brains are apt for distorting causes and effects. Brains are the place where "...derationalized by passion, deactualized by memory, ideas and purposes are reborn as irrational beliefs and symbols"¹.

Adopting theoretical constructs from reaction-diffusion computing² and massive-parallel models of collective intelligence³ we design computational model of mentality depicted as a large-scale pool of irrationally behaving entities.

We construct models of quasi-chemical interactions between happiness, sadness, anger, fear and confusion, and then characterize all modes of integral and spatio-temporal dynamics in affective mixtures. We show how distortions of conventional forms and structures in crowd-minds lead to drastic changes in reasoning and inference, and divert doxastic collectives from conventional routes of knowledge accumulation to complex non-linear dynamics of mixtures of belief, delusion, ignorance, knowledge and doubt.

In computational experiments with constrained doxastic mixtures we prove that norms can barely improve, and hardly control, behavior of crowd-minds. Non-trivially behaving combinatorial, pre-logical, systems are derived therefore to provide a formal foundation of the irrationality of a crowd-mind.

We envisage that our results in dynamics of affective mixtures, where collective pre-emotions develop to sadness and anxiety, may form a basis for emerging science of confusion.

Findings in behavior of doxastic mixtures — evolving towards doubt and ignorance — characterize a collective pre-knowledge, a subject of a science of doubt. Evolving of a pre-structure to a structure, governed by irrationality, may provide a starting point for studies in disordered mentality.

Models and paradigms developed might be applied to mathematical studies of affective collective intelligence, computational models of minds near the state of mental disorder, design of massively parallel prototypes of artificial consciousness, software implementations of affective cognition and design of hardware prototypes of emotional controllers for robots.

References

1. Moscovici, S. (1985). The discovery of masses. In: *Changing Conceptions of Crowd Mind and Behaviour*, eds. C. F. Graumann and S. Moscovici (Springer, Berlin), 5–25.
2. Adamatzky, A., De Lacy Costello, B. & Asai, T. *Reaction-Diffusion Computers* (Elsevier, 2005).
3. Adamatzky, A. *Dynamics of Crowd-Minds* (World Scientific, 2005).

Morphodynamics and Perception: a conceptual stance and some clues from evolutionary robotics

Fernando Almeida e Costa, Ian Macinnes and Inman Harvey
Centre for Computational Neuroscience and Robotics
University of Sussex, Brighton, UK
F.AlmeidaCosta@sussex.ac.uk, ian.macinnes@gmail.com,
I.R.Harvey@sussex.ac.uk

One of the deepest conceptual insights brought about by New A.I. was the realization that cognition results from the opportunistic exploitation of all morphodynamical properties present in the agent's body and environment, allowing the minimization of control at the algorithmic level. These properties structure the agent's perceptual world. The importance of body morphologies has been stressed in a significant body of work in A.I. and cognitive science. Here, we are aiming at a larger scope of morphodynamical properties, which includes both the morphological organisation of the agent's body and the morphological structuring of the environment "as perceived" by the agent. Gestalt theory in psychology of perception and Jakob von Uexkull's functional circle hypothesis, in biology/ethology, provide two general frameworks for the understanding of active perception as a morphologically-based ability. It will be shown that these frameworks must be seen within the broader scope of a *morphological turn* that occurred in the second half of the 20th century in various areas of research. In all these cases a morphologically-based – as opposed to information-theoretic based – way of conceiving cognitive activity was put forward. Some authors go as far as defending that developments in the theory of dynamical systems opened the possibility of a non-reductionist physics of perception and meaning (see Petitot, 2000). This possibility will be discussed.

A dynamicist approach regards organisms as being perturbed by and responding to some cues they were evolutionarily selected to respond to within their environment, rather than mirroring or extracting information from the outside world. The morphological structuring of the perceived environment – the salencies to which an agent is able to respond – is highly constrained by the particular morphologies of its body, and by the dynamics of those morphologies. To exploit this aspect for engineering and investigative purposes, a particular method in evolutionary robotics is proposed, based on the functional circle hypothesis by Jakob von Uexkull (Macinnes and Di Paolo, 2005a). A particular experiment based on that method is presented.

A functional circle is an abstract structure that describes the functional relationship between an organism, its "perceived world", and its environment. According to the functional circle hypothesis, a perceptual sign of an object give rise to a perceptual cue, the *subjective experience* of that object in the organism's *Umwelt**. This leads to an effector cue which drives the animal to perform some action, changing the organism's relationship to the object. After the action is performed the functional circle self extinguishes.

The proposed method consists of changing the mutational operators to evolve functional circles instead of directly evolving sensorimotor loops. The example experiment applies a model of functional circles to Karl Sim's block agents as a means of co-evolving the agent's morphodynamics and its perceptual world, and it is argued that this technique enables a closer coupling between body, controller, and environment than directly evolving sensorimotor loops. This is supported with

statistical evidence. Each simulated robot is constructed from blocks joined by powered hinge joints and attributes of both the morphology and neural network controller are genetically specified. The robots are evolved to move across a flat surface without using wheels or rolling parts. They must adopt a strategy of producing from their controllers and bodies a pattern of states that result in coherent forward movement for their individual morphologies.

The evolving functional circle hypothesis predicts that adding multiple perceptual cues will produce robots that are more adapted to their environment than they would be otherwise. A comparative analysis of the results of experiments in simulated robots suggests that this is the case.

An explanation will be put forward. The specific positions of the sensors using mutable locations together with body morphology define spatial and temporal relationships with the environment. Co-evolving the agent's morphology, locations of its sensors, and controllers, evolve these relationships as well which implies that we are evolving perceptual cues, and therefore evolving perceptual worlds. Models of the perceptual worlds of the robots and organisms are often represented as being a part of the controller or nervous system. The experiment shows that this is misleading: allowing the morphology to evolve permits the agent to select and respond to appropriate cues from its body and environment, and therefore also evolves its modeled perceptual world. It also demonstrates that the agent exists as part of the environment it has adapted to, as it requires and expects the sequence of environmental cues that it has evolved to respond to in order to function in a coherent manner.

* The word *Umwelt* was used by von Uexküll to describe the “own world” of an organism, constructed through its particular (as a species, morphodynamically organised in a certain way) interaction with the environment.

References

- Beer, R. 1995. A dynamical systems perspective on agent-environment interaction. In *Artificial Intelligence*, 72:173-215.
- Maccines, I. and Di Paolo, E. 2005a. The advantages of evolving perceptual cues. *Adaptive Behaviour*, 14(2), forthcoming
- Petitot, J. 2000 (1st ed.1992). *Physique du Sens – de la théorie des singularités aux structures sémio-narratives*, Paris, CNRS Éditions.
- Pfeifer, R. (2000) On the role of morphology and materials in adaptive behavior, *From animals to Animats. Proceedings of the 6th Int. Conference on the Simulation of Adaptive Behavior*. Cambridge, Mass.: MIT Press
- Von Uexküll, J. 1934. A Stroll through the worlds of animals and men: A picture book of invisible worlds. In Schiller, C., editor, *Instinctive Behaviour: The development of a modern concept*. International Universities Press, Inc., New York, 1957 edition, English Translation.

Enculturation without language

Sepand Ansari*, Caro Lucas*

*Control and Intelligent Processing Center of Excellence
University of Tehran, Tehran, Iran
sepans@ece.ut.ac.ir, lucas@ipm.ir

Language is mostly described as the means of communication or a system for manipulating symbols. Language is the key requirement of emergence of culture which is the universal capacity to classify, codify and communicate the experiences symbolically. This communication and ability to share experience is extremely desirable for Artificial Intelligence researchers but achieving a dynamic language in multi agent systems is a demanding task. The early attempts to achieve that were based on the internal mental representation of the agent. When agents have central mental representation, sharing this representation with each other is only passing communication symbols that can be easily translated to the mental representation symbols, to each other. But long before the appearance of Artificial intelligence, philosophers such as Heidegger and Wittgenstein, had shown that the whole Cartesian mental representation way of thinking about the mind's relation to the world is philosophically wrong. Wittgenstein described language as a set of *language-games* within which the words of our language function and receive their meaning. This view of *meaning as use* represents a break from the classical view - as presented by himself in his early thoughts - of *meaning as representation*. From the phenomenological point of view to Artificial Intelligence, proposed by Hubert Dreyfus [1], symbolic mental representation leads Artificial Intelligence to some sort of egocentric predicament. The egocentric predicament in philosophy is the idea that all our knowledge of the world must take the form of mental representations within our own minds (sensations, images, ideas, and so on), which the mind then operates upon in various ways. Thus we can never have any direct contact with reality outside our minds, and it becomes impossible to move justifiably from our own experimental data to the existence of an external world.

Therefore because of complexity, vagueness and nondeterministic nature of language, moving toward achieving an emergent protocol and language between autonomous agents, especially when eliminating mental representation, is vastly tied with philosophical concepts such as consciousness. Thus some researchers inspired language-game concept for language learning [2].

On the other hand, Dreyfus' critics are now widely influencing AI researchers. New agent architectures are proposed that eradicate the concept of central representation [3]. Dreyfus argues that human-like intelligence would require having a human-like being in the world, which would require them to be embodied in the environment, and have social acculturation more or less like ours.

This embodiment, distribution and communication can be seen in Stigmergy. In stigmergy individual parts of the system communicate with one another by modifying their environment and a social behavior emerges from their interaction. In stigmergy, the intelligence is the result of *being in the world*, and memory is distributed between the agents and even the environment. Stigmergy and more particularly, ant colony algorithms are vastly used in Artificial Intelligence [4]. Data mining and path finding in graphs are good examples of swarm intelligence. But all these usages have a characteristic in common which is static and deterministic environment. In the static environments, the hard coded rules that shape the behavior of the society are

sufficient for solving problems. But in more realistic, complex and dynamic environments, these simple rules are no longer adequate and the collective behavior should be evolved and so new social behaviors emerge through the evolution. Therefore, the social behavior is evaluated in this approach instead of the behavior of the individuals.

This kind of stigmergy should be based on open functionality. Unlike most evolutionary algorithms that a fixed desired functionality is coded in the fitness function, in an open functionality evolutionary system there is no *a priori* fitness function [5]. In this approach, just like A-life, no fixed fitness function is imposed from outside by the designer.

This stigmergy provides the communication needed for agents to be socialized, and since there is a behavioral competence between agents and societies, the social behavior evolves as the environment changes. Here the main questions arise. How far the stigmergy can go in the process of socialization? How powerful is the implicit language emerged in stigmergy in comparison to explicit languages?

References

1. H. L. Dreyfus *What Computers Still Can't Do: A Critique of Artificial Reason* (MIT Press 1992)
2. L. Steels, *Language games for autonomous robots*, (IEEE Intelligent Systems 16 (5) 2001)
3. R. Brooks. *Intelligence without representation*. (Artificial Intelligence, 47, 1991)
4. R. Daneshvar and C. Lucas, *Is Environment a Common Memory?* (Cooperative Multi Agent Learning workshop in The 16th European Conference on Machine Learning (ECML), Porto, Portugal, 2005)
5. L. Steels, *Emergent functionality through on-line evolution*. (Proceedings of Alife IV. Cambridge, MIT Press).

The Man-Machine Interaction: FMRI Study for an EMG Prosthetic Hand with Biofeedback

Alejandro Hernandez Arieta, Hiroshi Yokoi
Precision Engineering Department
The University of Tokyo, Tokyo, Japan
alex@prince.pe.u-tokyo.ac.jp

The mutual adaptation between man-machine opens new possibilities in the development of better user friendly interfaces that not only adapt to the user's characteristics, but also, permits the adaptation of the user to the machine. In this aspect, the artificial intelligence has open several doors to improve the man-machine interaction.

The artificial intelligence allows the machines to learn by themselves and adapt to different situations. Using this, the development of intelligent machines becomes possible. Now, if the intelligent machine has a way to interact with us, the adaptation can be done both ways. There are several examples of the use of feedback to improve the man-machine interface. One example is the use of sound to acquire cues in the interaction with the machine [1]. These studies show the improvement in the interaction when we increase the number of communication channels between the man and the machine. The problem with sound cues is that need the conscious effort to be recognized. Hunter et al. [2] shows another example of the importance of increasing the communication channels. In this study they show how the multiple sensory stimuli contribute to the conscious awareness of the body, and how it can be used to change the abnormal body awareness that occurs after limb amputation. This effect is also known as cortical reorganization where the brain after losing the stimuli from the amputated limb, due to the cross-modality, received input signals from the adjacent neurons, resulting in what is called "phantom limb". Our body is a multimodal system that uses several channels to obtain the current status of our bodies, if one channel fails; there are still others that help to provide the missing data. The user of a prosthetic hand needs to overcome the lack of tactile and proprioceptive data with visual feedback, which causes to fatigue faster because of the increment of conscious effort to control the hand [3]. These mechanisms need the implementation of a feedback source that enables the user to develop extended physiological proprioception [4]. We find some examples in the application of "tactile feedback" using vibration [5] or electrical stimulation [6]. On the hand, in the man-machine interfaces studies we find haptic interfaces that provide tactile feedback. Regrettably, those cannot be applied to prosthetic devices where the user presents partial or complete loss of the arm, which are our interest in this study. Therefore, we need to find a different way to provide with sensorial information to the human body. It is been demonstrated that the brain works with correlative information, therefore when provided with simultaneous stimuli, the brain can associate the stimuli into a unique event [8] Using this knowledge, we can force the brain into produce new sensations, provided that the stimulus is simultaneous, so the person using a prosthetic hand can have sensorial feedback besides the visual.

With the proper stimuli combination, the body image can be change, allowing for the human body to adapt to external devices. In order to test this hypothesis we proposed the used of an adaptable EMG controlled prosthetic hand with tactile feedback using electrical stimulation. Lotze et al [9] showed the positive effects in the use of myoelectrical prosthesis to revert cortical reorganization. This is possible due

to the combination of the intentionality from the amputee to move the absent limb, which results in muscular movement, and the visual feedback provided by the myoelectric device. We think that if we include even more sensory channels, in this case, tactile feedback, the adaptation to the prosthetic device can be enhanced, resulting in subconscious control of the device. With this we close the control loop having the highly plastic human brain adapt to the prosthetic hand, and the prosthetic hand adapt to the users particular characteristics. We performed an fMRI study in order to measure the adaptation of the human brain when using an adaptable EMG prosthetic hand. Our results show that the brain can interact better with the prosthetic hand when receiving the tactile feedback. Also, we confirm the cross modality of the brain, because the amputee's brain was able to identify the action of grasping an object and the electrical stimulation as one event, resulting in the activation of the somatosensory area of the right hand, even though the stimulation was performed on the left arm, and the person does not have a right hand.




Figure 1 shows the image resulting after the fMRI scanning for the amputee grabbing a cylinder using the robot hand with electrical stimulation functioning as tactile feedback. The image on the right is the zoom out of the motor and somatosensory area related to the hand and arm. The upper part shows the activation of the motor area in charge of the right hand movement, the lower part shows the reaction from the somatosensory area related to the hand. It is important to notice that the subject does not have the right arm to touch any object.

References

- [1] Rauterberg M. Different Effects of Auditory feedback in Man-Machine Interfaces. Human Factors in Auditory Warnings. ISBN 0-291-39849-9 (1999)
- [2] Hunter J, Katz J, Davis K. the effect of tactile and visual inputs on phantom limb awareness. Brain, Vol 126, p. 579-589 (2003)
- [3] Weir, R.F. ff., and Childress, D.S. and Licameli, J.N., "Motor Prostheses and the Lack of Physiologically Appropriate Feedback: Their function and connections with Manipulators, Robots and other Human-Machine Systems". Human and Machine Haptics, Section 3.21. Eds: M. Cutkosky & Srinivasan, (MIT Press); March 98
- [4] Simpson, D. C., (1974), "The Choice of Control System for the Multimovement Prosthesis: Extended Physiological Proprioception (e.p.p.)" in The Control of Upper-Extremity Prostheses and Orthoses (Herberts, P., Kadefors, R., Magnusson, R., and Petersen, I., Eds.), Springfield, Charles Thomas.
- [5] Alvaro Rios Poveda. Myoelectric prosthesis with sensorial feedback. University of New Brunswick's MyoElectric Controls/Powered Prosthetics Symposium. 2002.
- [6] M. Yoshida, Y. Sasaki, N Nakayama, "Sensory feedback for biomimetic prosthetic hand", BPES 2002, In Japanese "The 17th living body and physiology engineering symposium".
- [7] Makoto Shimojo, Takafumi Suzuki, Akio NAMIKI, Takashi Saito, Masanari Kunimoto, Ryota Makino, Hironori Ogawa, Masatoshi ISHIKAWA, and Kunihiro Mabuchi : Development of a System for Experiencing Tactile Sensation from a Robot Hand by Electrically Stimulating Sensory Nerve Fiber, 2003 IEEE International Conference on Robotics and Automation (Taipei, Taiwan, 2003.9.16) pp.1264-1270
- [8] Carrie Armel K. and Ramachandran V. S. Projecting sensations to external objects: evidence from skin conductance response Proc. R. Soc. B. 270, p. 1499 – 1506, 2003
- [9] Lotze M, Grodd W., Birbaumer N., Erb M., Huse E., and Flor H. Does use of a myoelectric prosthesis prevent cortical reorganization and phantom limb pain? . nature neuroscience, volume 2 no 6, june 1999

Extending the Robot's Operational Time and Space

Lijin Aryananda

Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology, Cambridge, MA, US
lijin@csail.mit.edu

Our motivation is to develop a robotic creature, Mertz, that ‘lives’ among us on a regular basis while learning from and about people through its daily social experience (see Figure 1). In particular, we aim for the robot to incrementally learn to recognize individuals, acquire simple words, form associations among multi-modal percepts and robot’s own actions, and adapt based on regularities observed in the environment. Our approach is inspired by the role of experience in the world and social interaction, both of which have been shown to be crucial to the infant’s learning process. Developmental approaches in robotics and the importance of social interaction for robot learning have been widely explored^{1,2}. We propose that in order to learn about people through social experience, the robot must be situated in our physical and social environment, i.e. experiencing the world and interacting with people everyday in different human spaces.

In this paper, we present the lessons and challenges that we encountered during the consequential process of extending the space and time in which the robot operates. Although it is still difficult for humanoid robots to operate robustly in noisy environment, the issue of robustness has not received adequate attention in most research projects³. Since robots will ultimately have to operate beyond the scope of short video clips and end-of-project demonstrations, we believe that a better understanding of these challenges is valuable for motivating further work in various areas contributing to this interdisciplinary endeavor.



Fig. 1. Mertz, an active vision head robot with 13 degrees of freedom (left) and a picture of the robot interacting with passersby during an experiment session (right).

Previous and Current Work

We have designed the robot toward days of continuous operation at various public spaces⁴. We are currently working on several projects, including face and speaker recognition, multi-modal attention system, spatiotemporal association of perceptual and action events, etc. During this entire development process, we have encountered a number of expected and unexpected challenges in extending the range of the robot’s operational time and space. We will briefly describe some of these findings below.

Lessons and Challenges

Perception has been blamed to be one of the biggest hurdles in robotics and certainly has posed many difficulties in our case. We generally found that many existing vision and speech technology are not suitable for our setting and constraints. Vision algorithms for static cameras are unusable because both cameras pan independently. The desktop microphone required for natural interaction with multiple people generates decreased performance compared to the headset typically used for speech

recognition. Drastic lighting changes inside the building and conducting experiments in different locations have forced us to go through many iterations of the robot's perceptual systems. Something that works in the morning at the laboratory may no longer work in the evening or at another location. Many automatic adaptive mechanisms, such as for the camera's internal parameters to deal with lighting changes throughout the day, are now necessary.

For a robotic creature that continuously learns while living in its environment, there is no separation between the learning and testing periods. The two are blurred together and often occurring in parallel. Mertz has to continually locate learning targets and carefully observe to learn about them. These two tasks are conflicting in many ways. The perceptual system is thus divided between fast but less precise processes for the first task and slower but more accurate algorithms for the latter. Similarly, the attention system has to balance between being reactive to new salient stimuli and persistent to observe current learning target. This dichotomy is interestingly reflected in the *what* and *where* pathways of our visual cortex, as well as the endogenous and exogenous control of visual attention.

Humans' tendency to anthropomorphize generally makes the robot's task to socially interact simpler. However, requiring the robot to interact with multiple people for an extended duration has called for a more sophisticated social interface. One can imagine that a friendly robot that makes eye contact and mimics your speech can be quite entertaining, but not for too long. While the premise that social interaction can guide robot learning is promising, it also suffers from the "chicken and egg" problem in a long-term setting, i.e. in order to sustain an extended interaction, the robot also needs to be able to learn and retain memory of past events.

In all engineering disciplines, we tend to focus on maximizing task performance. Whenever people are present, Mertz's task is to detect and engage them in interaction. We learned that when the robot is on all the time, in addition to its tasks, the robot also has to deal with down time, when no one is around. All of a sudden the environment's background and false positive detection errors become a big issue. During an experiment session, the robot kept falsely detecting a face on the ceiling, stared at it all day, and ignored everything else. Lastly, as the software complexity grows, the harder it becomes to keep the entire system running for many hours. Memory leaks and infrequent bugs emerging from subsystem interactions are very difficult to track. Moreover, a robot that runs for many hours per day and learns from its experiences can easily generate hundreds of gigabytes of data. While having a lot of data is undeniably useful, figuring out how to automatically filter, store, and retrieve them in real time is an engineering feat.

References

1. M. Lungarella, G. Metta, R. Pfeifer, G. Sandini. *Developmental Robotics: A Survey*. Connection Science, 15(4):151-190, 2003.
2. T. Fong, I. Nourbakhsh, and K. Dautenhahn. *A Survey of Socially Interactive Robots*. Robotics and Autonomous Systems, 42:143-166, 2003.
3. R. Bischoff, V. Graefe. *Design Principles for Dependable Robotic Assistants*. Int. Journal of Humanoid Robotics, vol. 1 no. 1: 95-124, 2004.
4. L. Aryananda, J. Weber. Mertz: *A Quest for a Robust and Scalable Active Vision Humanoid Head Robot*. IEEE-RAS/RSJ Int. Conference on Humanoid Robots, 2004.

Cognition Network Technology: A Novel Knowledge-Based Approach for Analyzing and Processing Complex Data. Method and Application Samples in Image Analysis

Martin Baatz, Arno Schäpe, Günter Schmidt, Maria Athelougou, Gerd Binnig
Definiens AG, Munich

Data analysis in general and image analysis in particular require multi-scale and semantic approaches when dealing with complex problems. The clusters of interest will most likely be embedded in or will themselves represent hierarchical structures. Even in non-hierarchical situations it may not be possible to extract the relevant objects (segmentation) in a straight forward manner. Instead, a stepwise generation of the intermediate object on different scales may lead to the desired result. A high level computer language designed for modelling such complex cognition processes must offer a limited number of generic building blocks which can be combined into a program. Because of the multi-scale aspect, those building blocks have to be reusable for very different procedures as well as for those on very different scales.

Cognition Networks were developed for representation, knowledge extraction and simulation of complex systems and data. This new dynamic object model combines methods from many other well known approaches for handling complexity like semantic networks, Bayesian networks, cellular automata, neuronal networks, expert systems and programming languages together with new aspects like self-similarity and local adaptive computing. A *Cognition Network* is built by *objects*. All objects may carry various kinds of data and may be linked by (link) objects. Links are objects themselves, and therefore carry data *and* further links. In addition to this, any object may carry semantic meanings and procedural attachments. For the purpose of complexity reduction and for the purpose of an appropriate representation of the structure and the semantics of the original input data, *Cognition Networks* are self similar in the following aspects: (1) basic properties and data structures are similar for all objects; (2) the network has a hierarchical structure, i.e. an object can be linked to a sub network in order to represent structure on different scales; (3) object links themselves can be linked and (4) procedures and methods are applied in the same manner to all objects, explicitly to objects on different scales. Points (2) and (3) produce the fractal topology of the *Cognition Network* (Fig. 1).

A Cognition Network is able to store, represent and extract knowledge from a complex input like images or texts. The knowledge stored in a Cognition Network is represented by the network structure of all objects and the contained data. A large and valuable part of that information is contained by link objects and by sub networks. Data analysis within a Cognition Network is a dynamic process controlled by a specific sub network or knowledge base that was created using a high level programming language called the Cognition Network Language. This knowledge base contains the procedural knowledge and the descriptive knowledge for a given analysis task in a structured form. From the input data of the analysis problem an *instance* Cognition Network is constructed by the procedural objects from the rule base. Objects of the instance network are linked into the rule base by classification procedures. Descriptors used for classification are intrinsic information of the object and/or information provided by the network structure of the object, addressing criteria such as composition, embedding or distance. An object or cluster composed by a number of basic data units such as pixels can carry far more and more relevant information than a single data point. This is supported by appropriate clustering and segmentation procedures [2]. The hierarchical structure of clusters in the Cognition Network allows the simultaneous representation of structures in images on different scales. In such a




Fig.1 Cognition Network: note the fractal topology of the hierarchical object network with link objects and procedural attachments.

structure, each object is directly linked to all neighbouring objects, sub objects and super objects. When operating over this network a considerable amount of structural and relational information can be accessed. [3]. The additional information accessible and the improvement of the signal to noise ratio concerning the information provided by objects result in more detailed and more robust classification results [4]. Based on the labelling of objects, procedural objects can operate on specific semantic sub networks and/or to the networked environment of individual objects. This enables local adaptive network procedures to change the object's classification and the network structure itself. Starting with the initial input data - for example pixel in an image - the constantly alternating application between local evaluation of the current state of the network and locally adaptive procedures results in a self-creation and self-organisation of the network. The final state of the network then represents the information extracted from the data like the structures of interests in images. The knowledge base can be adapted to solve specific problems by interactively pointing to sample regions in an image and than transferring this implicit knowledge into explicit procedural and descriptive knowledge.





Fig 2: Computation of the proliferation index of crypts in the small intestine of mice. Example from a high-throughput screening. The procedure has extracted a crypt with a longitudinal cross-cut section and the contained mitotic and non-mitotic nuclei from PCNA stained tissue. The procedure works despite the strong structural variety in which crypts occur. Image data courtesy Novartis, Basel





Fig. 3: Extraction of nuclei and mitochondria in electron microscopic images. Although spectral information exists only in one dimension (black/white) the procedure finds objects using the spectral and the relational information in the image.

Image data courtesy ICF LMU Munich

Cognition Network Technology has proven robustness and reliability through numerous fully automated high throughput image analysis applications in areas like histopathology (Fig 2) or electron microscopy (Fig 3), for instance. It specifically finds structures of interest even in challenging cases such as low signal to noise ratio images, heterogeneous or variable structures of interest or tasks which include a complex semantic [1].

References

- [1] G. Binnig, M. Baatz, J. Klenk, et. al.: "Will Machines start to think like humans ?" *Europhysics News* Vol.33 No.2, 2002.
- [2] M. Baatz, A. Schäpe: "Multiresolution Segmentation – an optimization approach for high quality multi-scale image segmentation. *Angewandte Geographische Informationsverarbeitung XII, Beiträge zum AGIT-Symposium*, Salzburg, 2000.
- [3] M. Baatz, A. Schäpe, G. Schmidt: „Verfahren zur Verarbeitung von Datenstrukturen“, *German Patent Application* Nr. DE 19960372.3 of 14 Sept. 1999
- [4] M. Baatz, A. Schäpe: "Object-Oriented and Multi-Scale Image Analysis in Semantic Networks". *Proc. of the 2nd International Symposium on Operationalization of Remote Sensing*, Enschede, ITC, 1999.

Communication With and Within Artificial Societies

Fatmah A. Baothman

Departments of Computer Science

King Abdul Aziz University, Jeddah, Saudi Arabia

fbaothman@kau.edu.sa and al_batuul@yahoo.com

Technology advancements resulting in the creation of artificial intelligent societies have raised many challenges for researchers¹. One prime concern, with these societies, has been the means to establish natural social interaction in dynamic, cross-cultural environments. The main complexity faced is that needed to enable agents existing in **World of Three Dimensions (W3Ds)** – real, virtual and digital – to evolve their knowledge via a collaborative learning mechanism using natural language to engage in efficient communication acts. This work discusses the need for enabling autonomous agents' survival in W3Ds with **Hyper Communication Interactions (HCI)** capabilities focusing on natural speech. We now have recognizers with lexicons as large as those of well-educated humans and avatars such as Baldy and Badr². The best of these have simulated vocal tracts whose articulators can be inspected as they speak³. An important aspect of this kind of speech communication is the representation of speech sounds. In this note, the focus is on speech encoding units called the sub-segmental primes (see fig.1). These elements link and symbolize the intentional knowledge at the skeletal tier with the acoustic and phonetic output associated with speech articulation and perception. An important aspect of sub-segmental elements is that they can be recognized in isolation or in combination⁴. Moreover, these elements form cognitive entities which, constrained by language parameters⁵. Elements populate timing slots are arrayed in the 'melody tier' in a 'syllabic' constituent structure. These primes proved to be alternative to Chomsky & Halle representations^{6,7}. A linguistic agent using a version of **Government Phonology (GP)** elements was designed in previous work⁸ and its sub-segmental representations were used to model all attested phonological processes of Arabic⁹(see fig. 2).

Technology advancements resulting in the creation of artificial intelligent societies have raised many challenges for researchers¹. One prime concern, with these societies, has been the means to establish natural social interaction in dynamic, cross-cultural environments. The main complexity faced is that needed to enable agents existing in **World of Three Dimensions (W3Ds)** – real, virtual and digital – to evolve their knowledge via a collaborative learning mechanism using natural language to engage in efficient communication acts. This work discusses the need for enabling autonomous agents' survival in W3Ds with **Hyper Communication Interactions (HCI)** capabilities focusing on natural speech. We now have recognizers with lexicons as large as those of well-educated humans and avatars such as Baldy and Badr². The best of these have simulated vocal tracts whose articulators can be inspected as they speak³. An important aspect of this kind of speech communication is the representation of speech sounds. In this note, the focus is on speech encoding units called the sub-segmental primes (see fig.1). These elements link and symbolize the intentional knowledge at the skeletal tier with the acoustic and phonetic output associated with speech articulation and perception. An important aspect of sub-segmental elements is that they can be recognized in isolation or in combination⁴. Moreover, these elements form cognitive entities which, constrained by language parameters⁵. Elements populate timing slots are arrayed in the 'melody tier' in a 'syllabic' constituent structure. These primes proved to be alternative to Chomsky & Halle representations^{6,7}. A linguistic agent using a version of **Government Phonology (GP)** elements was designed in previous work⁸ and its sub-segmental representations were used to model all attested phonological processes of Arabic⁹(see fig. 2).




Fig. 1: The link between Constituent Structure and Sub-segmental or Element tier




Fig. 2: A Linguistic Melody Geometry Agent

The present work surveys the majority of autonomous agent communication toolkits. It addresses the problems faced in implementing sub-segmental elements in these toolkits. The goal is to design a **Hyper Communication Agent Architecture (HCAA)**, not limited to speech, suitable for human-agent¹⁰ and multi- agent interaction between W3Ds. A practical start of agent interaction has been made in Alife systems¹¹. This work also provides empirical comparative experiments performed on agents incorporating evolutionary neurocontroller techniques¹².

The neurocontrollers design accepts stimuli from three types of input nodes: type-1 indicating the W3D environment; type-2 identifying the source of communication (human, artificial, or other remote agent-creature); and type-3 specifying the sensor-data used (speech, other: auditory, haptic, and visual). Experimental controllers were devised with three to a maximum of nine hidden layers, and trained to fix internal parameters of standard feed-forward connectivity. Output nodes were designed to represent actions – specified by command words. Each output was associated with a phonological element, and the number of outputs units was chosen to reflect the elements needed to represent words in the intended command lexicon. It was observed as the numbers of hidden layers were changed from 3 to maximum 9, the fitness of the feed- forward networks almost doubled. The conducted experiments indicate that neurocontrollers with enough internal memory promise significant achievements for future communication applications with and among artificial societies.

References

1. Pfeifer, Rolf. & Scheier, Christian. *Understanding Intelligence* (MIT Press, Boston, 1999).
2. Ouni, S., Cohen, M. M., & Massaro, D. W. Training Baldi to be multilingual: A case study for an Arabic Badr. *Speech Communication*, 45(2), pp.115-137, 2005.
3. Ouni, S., Massaro, D.W., Cohen, M.M., Young, K. & Jesse, A. Internationalization of a Talking Head. 15th International Congress of Phonetic Sciences, August 3-9, Barcelona, 2003.
4. Harris, John. *English Sound Structure* (Blackwell press, Oxford, 1994).
5. Kaye, Jonathan. *Phonology: A Cognitive View* (Lawrence Erlbaum Associates, New Jersey, 1989).
6. Charette, Monik. *Conditions on Phonological Government* (Cambridge University Press, Cambridge, 1991).
7. Brockhaus, Wiebke & Michael Ingleby. Unary Primes and Constituent Structure. In Language Engineering. *Bulletin de la Communication Parlée* 5, pp. 109-29, 2000.
8. Baothman, Fatmah, Phonology Based ASR for Arabic, PhD thesis, School of Computing and Engineering, University of Huddersfield, 2003.
9. Baothman, Fatmah & Ingleby, Michael. *Representing Coarticulation Processes in Arabic Speech*, Journal of Arabic Linguistics Society, in Perspectives on Arabic Linguistics XVI, Boudelaa, Sami (ed.), pp. 95–102, 2006.
10. Dickinson, Ian. Human-Agent Communication. Hewlett-Packard Laboratories, 1998. Available at [<http://www.hpl.hp.com>].
11. Kaplan, F. & Schemata, Semiotic. Selection Units for Linguistic Cultural Evolution. Proceedings of Artificial Life VII, pp. 372-381, Cambridge, 2000.
12. Mukesh J. Patel, Vasant Honavar & Karthik Balakrishnan. *Advances in the Evolutionary Synthesis of Intelligent Agents* (MIT Press, Massachusetts, 2001).

Humanoid Robots — A Tool to Understand Human Intelligence

Sven Behnke

Humanoid Robots Group, Computer Science Institute

University of Freiburg, Georges-Köhler-Allee 52, 79110 Freiburg, Germany

behnke@informatik.uni-freiburg.de

Introduction

Humanoid robots, robots with an anthropomorphic body plan and human-like senses, are enjoying increasing popularity as research tool. More and more groups worldwide work on issues like bipedal locomotion, dexterous manipulation, audio-visual perception, human-robot interaction, adaptive control, and learning, targeted for the application in humanoid robots. These efforts are motivated by the vision to create a new kind of tool: robots that work in close cooperation with humans in the same environment that we designed to suit our needs. While highly specialized industrial robots are successfully employed in mass production, these new applications require a different approach: general purpose humanoid robots. The human body is well suited for acting in our everyday environments. Stairs, door handles, tools, and so on are designed to be used by humans. The new applications will require social interaction between humans and robots. If a robot is able to analyze and synthesize speech, eye movements, mimics, gestures, and body language, it will be capable of intuitive communication with humans. A human-like action repertoire also facilitates the programming of the robots by demonstration and the learning of new skills by imitation. Last, but not least, humanoid robots are used as a tool to understand human intelligence¹. This is a consequence of the constructive approach to AI. In contrast to programming symbol manipulation systems, building robots avoids the symbol grounding problem. Embodiment and situatedness require dealing with the problems of the real-world: perception and action. Various biomimetic robots have been built to test theories about animal perception and behavior. Likewise, the construction of humanoid robots is one promising approach to understand human intelligence.

Addressing all of the above aspects simultaneously exceeds the current state of the art. Today's humanoid robots display their capabilities in tasks requiring a limited subset of skills. My group develops humanoid robots for two applications: playing soccer and guiding people through a museum.

Humanoid Soccer Robots

Kitano and Asada proposed the RoboCup humanoid challenge² as the millennium challenge for AI and advanced robotics. It is to construct by 2050 a team of fully autonomous humanoid robot soccer players able to win a soccer game against the winner of the most recent FIFA World Cup. To facilitate research towards this long-term goal, the RoboCup Federation

organizes since 1997 annual soccer competitions in five leagues, focusing on different




Fig. 1: Left: Two KidSize robots of NimbleRo 2005 playing soccer. Right: 2006 TeenSize robot Robotinho.

research aspects. While team play and learning are major topics in the simulation league, the real-robot leagues developed solutions for robust real-time perception, omnidirectional locomotion, and ball handling. The RoboCup championships grew to the most important robotic competition worldwide³. Fig. 1 shows the robots Jupp and Sepp (19DOF, 60cm, 2.3kg), which competed for my team NimbRo at RoboCup 2005. The robots are controlled by a Pocket PC, which interprets the images of a wide-angle camera. They are able to walk omnidirectionally, can kick the ball hard and get up from the ground reliably⁴. Jupp and Sepp came in second in the overall Best Humanoid ranking. In the 2 vs. 2 soccer games, played for the first time, they reached the final, which was won 2:1 by Team Osaka. Their larger sibling Max won the Penalty Kick competition in the TeenSize class. Robotinho, which we constructed for the 2006 TeenSize class is also shown in Fig. 1. It has 21DOF, is 100cm tall, and weighs only 5kg.

Humanoid Robots for Intuitive Multimodal Communication

Guiding people through a museum requires a different subset of the skills mentioned above. Here, the intuitive multimodal communication with visitors is important. Fig. 2 shows our humanoid robot Fritz explaining a smaller robot to a human. Fritz is 120cm tall. It has 16DOFs in the legs and the arms and 16DOFs in the head. Its eyes are movable cameras. Fritz detects faces in the captured images and remembers the persons around it⁵. While it focuses its attention to one person, it keeps eye contact with the others to involve them into the conversation. Speech recognition and speech synthesis generate some small talk and allow the visitors to select exhibits that the robot explains. Fritz uses pointing gestures with its eyes, head, and arms to draw the attention of the visitors to the exhibits. The animated mouth, the eye brows, and its voice express emotions.

Conclusion

Playing soccer and guiding people through a museum cover relevant subsets of the skills needed to build capable humanoid robots. My group made progress in both domains. I will bring at least one humanoid robot to the summit for demonstrations. See www.NimbRo.net for images and videos of the robots.

References

1. B. Adams, C. Breazeal, R.A. Brooks, and B. Scassellati: Humanoid Robots: A New Kind of Tool. *IEEE Intelligent Systems*, 15(4), pp. 25-31, 2000.
2. H. Kitano and M. Asada: "The RoboCup humanoid challenge as the millennium challenge for advanced robotics", *Advanced Robotics*, 13(8), pp. 723-737, 2000.
3. S. Behnke: "Playing Soccer with Humanoid Robots", to appear in *KI - Zeitschrift Künstliche Intelligenz*, 2006.
4. S. Behnke, M. Schreiber, M. Bennewitz, J. Stückler, H. Strasdat, and J. Schwenk: "Designing a team of soccer-playing humanoid robots", In *Proc. of ISR/Robotik*, Munich, 5/2006.
5. M. Bennewitz, F. Faber, D. Joho, M. Schreiber, and S. Behnke: "Towards a Humanoid Museum Guide Robot that Interacts with Multiple Persons", In *Proc. of Int. Conf. on Humanoid Robots (Humanoids)*, Tsukuba, 12/2005.




Fig. 2: Our humanoid Fritz explaining a smaller robot.

Perspective Taking: An Organizing Principle for Learning in Human-Robot Interaction

Matt Berlin, Jesse Gray, Andrea L. Thomaz, Cynthia Breazeal
The Media Laboratory
MIT, Cambridge, MA, U.S.A.
[mattb, jg, alocker, d, cynthiab]@media.mit.edu

The ability to interpret demonstrations from the perspective of the teacher plays a critical role in human learning. Robotic systems that aim to learn effectively from human teachers must similarly be able to engage in perspective taking. We present an integrated architecture wherein the robot's cognitive functionality is organized around the ability to understand the environment from the perspective of a social partner as well as its own. The performance of this architecture on a set of learning tasks is evaluated against human data derived from a novel study examining the importance of perspective taking in human learning. Perspective taking, both in humans and in our architecture, focuses the agent's attention on the subset of the problem space that is important to the teacher. This constrained attention allows the agent to overcome ambiguity and incompleteness that can often be present in human demonstrations and thus learn what the teacher intends to teach.

Perspective Taking Architecture

We believe that socially situated robots will need to be designed as socially cognitive learners that can infer the intention of a human's instruction, even if the teacher's demonstrations are less than perfect for the robot. Our approach to endowing machines with socially-cognitive learning abilities is inspired by leading psychological theories and recent neuroscientific evidence for how human brains might infer the mental states of others. Specifically, *Simulation Theory* holds that certain parts of the brain have dual use; they are used to not only generate behavior and mental states, but also to predict and infer the same in others.^{1,2}



Fig. 1. The Leonardo robot and simulator.

We present an integrated architecture which runs on a 65 degree of freedom humanoid robot and its graphical simulator. Our architecture incorporates simulation-theoretic mechanisms as a foundational and organizational principle to support collaborative forms of human-robot interaction, such as tutelage-based learning. Our implementation enables a humanoid robot to monitor an adjacent human teacher by simulating his or her behavior within the robot's own generative mechanisms on the motor, goal-directed action, and perceptual-belief levels.

One important feature of this perspective taking process is that it focuses the agent's attention on the subset of the problem space that is important to the teacher.




Fig. 2. Input domains consistent with the perspective taking (PT) vs. non-perspective taking (NPT) hypotheses. The student's attention is focused on just the blocks that the teacher is aware of/attending to.

Focusing on a subset of the input/problem space directly affects the set of hypotheses entertained by the learning algorithm, and thus directly affects the skill transferred to the agent via the interaction with the teacher. This constrained attention allows the agent to overcome ambiguity and incompleteness that can often be present in human demonstrations.

Human Subjects Study and Evaluation

We conducted a novel human subjects study that highlights the important role that perspective taking plays in learning within a socially situated context. Participants engaged in four different learning tasks involving foam blocks. 20 participants observed demonstrations provided by a human teacher sitting opposite them (the social condition), while 21 participants were shown static images of the same demonstrations, with the teacher absent from the scene (the nonsocial condition). The specific skills acquired by the participants through these demonstrations were inferred from their subsequent behavior. For every task, differences in the skills acquired between the social and nonsocial conditions were highly significant.

The tasks from our study were used to create a benchmark suite for our architecture. In our simulation environment, the robot was presented with the same task demonstrations as were provided to the study participants. For every task and condition, the rule learned by the robot matched the most popular rule selected by the humans. This strongly suggests that the robot's perspective taking mechanisms focus its attention on a region of the input space similar to that attended to by study participants in the presence of a human teacher. Thus, our humanoid robot can apply perspective taking to draw the same conclusions as humans under conditions of high ambiguity.

Task	Condition	PT Rule	NPT Rule	Other	<i>p</i>
Task 1	social	6	1	13	***
	nonsocial	1	12	8	
Task 2	social	16	0	4	***
	nonsocial	7	12	2	
Task 3	social	12	8	-	***
	nonsocial	0	21	-	
Task 4	social	14	6	-	***
	nonsocial	0	21	-	

Fig. 3. Differential rule acquisition for study participants in social vs. nonsocial conditions. ***: $p < 0.001$.

Others have looked at the use of visual perspective taking in collaborative settings.³ This is the first work to examine the role of perspective taking for introceptive states (e.g., beliefs and goals) in a human-robot learning task. The result of our work is a novel, integrated approach where perspective taking is used as an organizing principle for learning in human-robot interaction.

References

1. Davies, M. & Stone, T. *Introduction*. In M. Davies & Stone, T., eds., *Folk Psychology: The Theory of Mind Debate* (Blackwell, Cambridge, 1995).
2. Barsalou, L.W., Niedenthal, P.M., Barbey, A., & Ruppert, J. *Social embodiment*. In B. Ross, ed., *The Psychology of Learning and Motivation* 43 (Academic Press, San Diego, 2003).
3. Trafton, J. G., Cassimatis, N. L., Bugajska, M. D., Brock, D. P., Mintz, F. E., & Schultz, A. C. Enabling effective human-robot interaction using perspective-taking in robots. *IEEE Transactions on Systems, Man, and Cybernetics* 35(4):460–470 (2005).

A Human-like Robot Torso ZAR5 with Fluidic Muscles: Full Controlled by Using Data Suit and Gloves

Ivo Boblan¹, Rudolf Bannasch², Andreas Schulz², and Hartmut Schwenk¹

¹ Technische Universität Berlin, FG Bionik und Evolutionstechnik,

² EvoLogics GmbH, F&E Labor Bionik,

Ackerstr. 76, 13355 Berlin, Germany, boblan@bionik.tu-berlin.de, www.zar-x.de

“Without embodiment is artificial intelligence nothing.” Our anthropomorphic robot ZAR5 (in German *Zwei-Arm-Roboter* in the version 5) is the first biological inspired and complete fluidic muscle driven robot torso which can full controlled by a data suit and two five finger data gloves.

Our robot is a human-like torso with two arms and five-finger hands which are strict developed according to bionical considerations. The combination of biology and robots leads to smoother and compliant movement which is more pleasant for us as people. Biologically inspired robots embody non-rigid movement which are made possible by special joints and actuators which give way and can both actively and passively adapt stiffness in different situations. The more the technical realization corresponds with the biological role model the successful is the reflection of the true reality. If we want learn more about the control architecture and there functionality in the human being, we have to build an exact copy of the natural role model as much as possible to reproduce our conceivability's of artificial intelligence.

Biological inspired is not only the morphology - size, proportions and load-bearing inner structures - but also the physiology – moving mechanical parts and muscle tendon systems – and also parts of the all driven control architecture. The better the morphology is understood and transferred to the artificial body the better can act the physiological parts and finally the controlled software. Morphology, physiology and control are an entity and have to always consider together.

All joint angles of the data suit wearing man are read and transferred to the corresponding robot joints. The angle transmitters of the two-arm suit are pots hold by a cushy exoskeleton. The angular values of the two data gloves come from strain gauges which sit above of each finger joint. All angular data are read at every 20ms and transferred to the CAN bus connected microcontrollers of the robot body. Each main body part of the whole robot: right and left hand and right and left side of the body is controlled by a system of two microcontrollers. One microcontroller organized the control loops of the connecting joints of this body part and the other is responsible for the generation of the PWM signals for the fast switching valves. The main PC located in the base of the robot gives the data of the operator or a pre-configured batch file directly to the controlling microcontroller.

We decide between different function modes to control the movements of the joints. It depends on the one or antagonistic muscle driven setups of the joints and on the made mission of the movements. The stand-alone working PWM controller drives the activated valves and finally the connected muscles of the joints to the desired positions. Before these tasks are finished they are normally overwritten by the next datagram's from the suit or gloves and lead to complex movement trajectories.

The strict decentralized control architecture enabled a parallel, failure tolerant, cheap designed and robust piloting of the robot joints.



First workshop photo with handling over



Holding an object with both hands



Demonstration at the Hannover Messe 2006



Holding objects to demonstrate the functionality



Shake hands with German Federal Chancellor



ZAR5 and Indian prime minister at the HM 2006

Engineering Modular Interfaces in Neural System Design

Martin Boerlin*, Paul F M J Verschure*[†], Tobi Delbruck* and Kynan Eng*

*Institute of Neuroinformatics, University/ETH Zurich, Switzerland

[†]Technology Department, ICREA and Universitat Pompeu Fabra, Barcelona, Spain
kynan@ini.phys.ethz.ch

Introduction

The problem of connecting separately developed computing modules in robotic systems is generally solved by shared interface definitions at the physical (connectors, signals), logical (data representations) and semantic (high-level data structures) levels. While physical and logical interfaces are relatively easily defined, semantic interfaces can quickly become complex to the point where both sides need to use a common software base or cumbersome software architectures (e.g. the Common Object Request Broker Architecture).

Future robotic computing systems that more closely resemble the nervous systems of natural organisms will face problems in transmitting semantics across noisy (spiking) neural interfaces. One way of solving this problem is to have a receiver learn to decode its input, reducing the need for shared semantics (Fig 1). This approach would enable the development of more modular, interchangeable neural systems – a kind of “neural plug-and-play”. In the nearer term, these systems would also be better able to deal with long-term on-line adaptation to partial system damage, sensor drift, mechanical component misalignment, etc. Some useful types of semantics which can be learned by a receiver across a neural interface include:

1. Topological (neighbourhood) relationships between sender-side elements
2. Automatic input categorization and identification, possibly to drive changes in receiver-side processing and actions
3. Combining (1) and (2) to make deductions concerning internal sender-side data flow and/or data processing. Decoding this type of semantics could be seen as a simple kind of “code sharing” across the interface.

Much work has already been done on problem (2), e.g. in spike sorting algorithms and unsupervised learning based on input statistics. Here we describe our approach to the somewhat less-studied problem of (1) – learning topological relationships between sender-side elements – as a pre-requisite for attacking problem (3).

Learning Topology from Data

A fundamental semantic relationship in neural computing systems is spatial and/or topological adjacency. Preserving topology across an interface can be achieved by physical-level spatial encoding, at the expense of requiring one transmission conduit per spatial location. An alternative is to only transmit multiplexed data for important events, using unique identifiers for spatial locations at the logical level. This is commonly referred to as an address-event representation (AER), a logical layer often used in neuromorphic systems¹ such as analog VLSI vision sensors² and tactile sensor networks³. The topographical relationships between adjacent elements in an AER sender can be learned in a weight matrix on the receiver side by using the temporal relationships between the transmitted events. This is achieved with a spike timing




Fig. 1. Computational module interfacing using (A) shared design-time semantics, and (B) independent sender and receiver encoding/decoding.

dependent plasticity (STDP) learning rule⁴ combined with predictive Hebbian learning:

- If an event occurred at some location at time t , then a subsequent event at a different location at time $t + \delta t$ probably occurred near the first location (STDP).
- If, using our current estimation of the connectivity, we can correctly predict the label of the next incoming event, then we can be more confident of our prediction (predictive Hebbian learning).

The neighbours can then be estimated using a soft winner-take-all operation, combined with a threshold to determine edge and corner elements. Weight normalization ensures that they remain bounded and that the process is responsive to changes in the input topology. We tested our topology learning algorithm on real-world data from a tactile sensory floor³ and a neuromorphic vision chip². In both cases the algorithm was able to reconstruct the sender-side topology within a relatively short time (<1 h for the tactile floor, a few min. for the vision chip). Errors in the learned topology were found in areas where little data was received, i.e. around the edges of the tactile floor and in regions of the vision chip image with little motion input. The algorithm was able to adapt to simulated input “damage” and adjust the learned topology accordingly.




Fig. 2. Development of learned topology representations from sensory data for a 360-tile tactile sensory floor (hexagonal tiling, top) and a 64x64 pixel silicon retina (square tiling, bottom). Above each image is the number of events processed. Learning weeds out potential connections.

Discussion

We have shown how a basic type of semantic information in a neural processing system – topological organization – can be learned from incoming data, offering design-time advantages in terms of enhanced system modularity and runtime advantages including simple system configuration/upgrading and adaptability to damage or drift. In future work we will extend this work to deal with other semantic information such as multi-modal data representations and stimulus tracking, and we will investigate the effects on input resolution and signal degradation on learned receiver side representations.

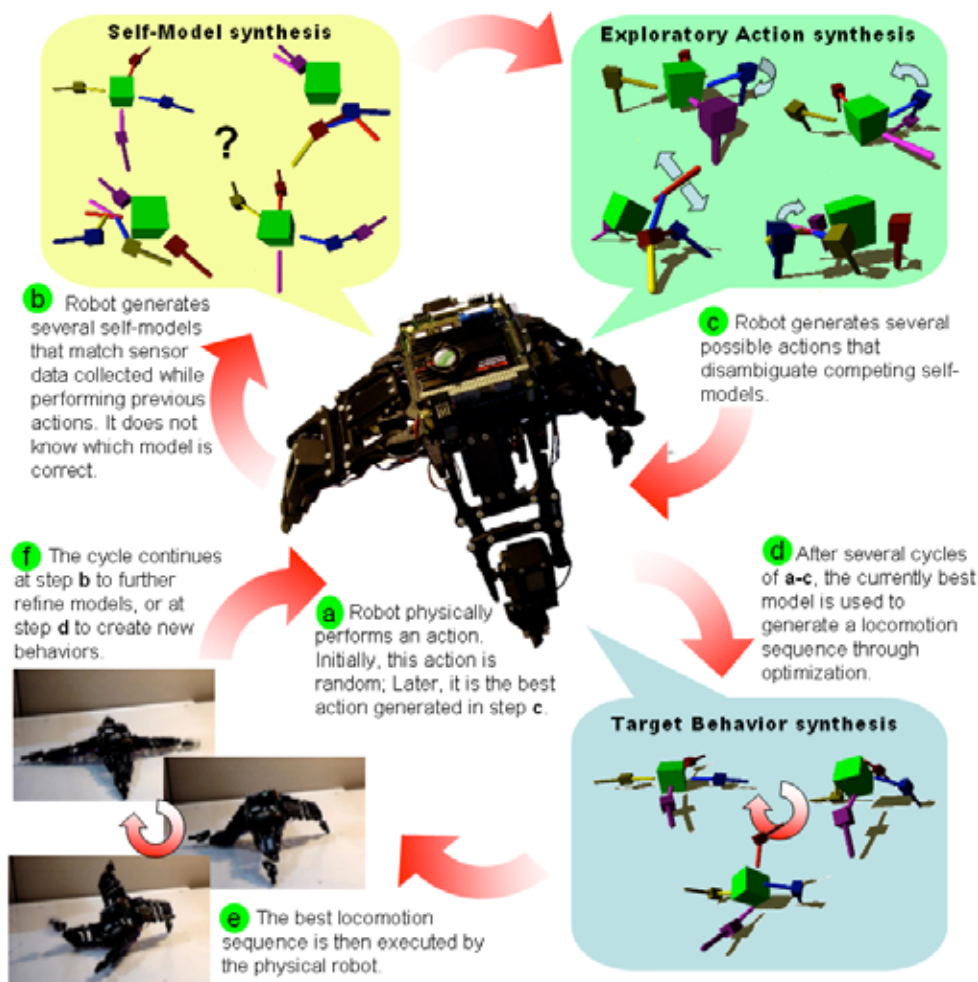
References

1. Serrano-Gotarredona, R. et al. AER Building Blocks for Multi-Layer Multi-Chip Neuromorphic Vision Systems. *Advances in Neural Information Processing Systems* **18**, 1217-1224 (2005).
2. Lichtsteiner, P., Posch, C. & Delbruck, T. A 128x128 120dB 30mW Asynchronous Vision Sensor that Responds to Relative Intensity Change. *International Solid State Circuits Conference (ISSCC 2006)*, 508-509 (2006).
3. Delbrück, T., Douglas, R. J., Marchal, P., Verschure, P. F. M. J. & Whatley, A. M. A device for controlling a physical system. in *US 6,603,082; EP 1091334; JP2001180810* (ETH Zurich, US, EU, JP, 1999).
4. Song, S., Miller, K. D. & Abbott, L. F. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience* **3**, 919-926 (2000).

A Robot That Builds a Simulator of Itself: A Counterrevolutionary Tale

Josh Bongard, Victor Zykov
Sibley School of Mechanical and Aerospace Engineering
Cornell University
josh.bongard@cornell.edu

Higher animals use some form of an internal model of themselves for planning complex actions and predicting their consequence^{i,ii,iii,iv,v}, but it is not clear if and how these self-models are acquired or what form they take. Despite this, Brooks' AI revolution in the 1980s^{vi} was a response to the slow progress of model-based robotics since the 1950s, and ushered us into the current AI epoch in which model-free robotics is fashionable. Although simple and robust behaviors can be achieved without a model at all^{vi,vii}, Brooks himself anticipated the "tantalizing possibility"^{viii} that a robot could one day autonomously create models of itself along with new behaviors, but cited "deep issues" preventing it from becoming viable.



A robot that can build a simulator of itself. The robot performs a random action (a), then synthesizes candidate models of its own morphology (b). Those models are then used to synthesize a new action (c) that induces maximal disagreement among the models. This cycle repeats until a sufficiently accurate model is achieved. At that point, the best model is used to synthesize a locomotory behavior (d), which is then executed by the physical robot (e).

Here I report a successful physical realization of this concept, and further demonstrate the reward mechanism for directing the self-modeling phase: disagreement among competing internal models.

More specifically, I will show how low-level sensor-motor correlations can give rise to an internal predictive self-model, which in turn can be used to develop new behaviors (see figure). I will demonstrate both computationally and experimentally how a legged robot automatically synthesizes a predictive model of its own topology (where and how its body parts are connected) through limited yet self-directed interaction with its environment, and then uses this model to synthesize successful new behaviors: behaviors synthesized in simulation that transfer well to the physical robot. Importantly, we have found that approximate models are sufficient for synthesizing behavior in simulation and then executing them on the physical robot.

In the experiments I will present, the physical robot learns how to move forward based on only 16 brief self-directed interactions with its environment. These interactions were unrelated to the task of locomotion, and are driven only by the objective of disambiguating competing internal models. This method is superior to the three current approaches to automated behavior generation for robotics (evolving behaviors all on a physical device, which requires at least hundreds of evaluations on a physical robot^{ix}; evolving in a robot simulator, which requires the manual construction of the simulator first^x; or adapting a previous behavior online, which assumes the existence of a previous behavior^{xi}) because it does not require extensive physical trials, and does not assume a simulator or previous behavior.

This finding may help develop more robust machines, as well as shed light on the relation between curiosity^{xii} and cognition^{xiii} in animals and humans: creating models through exploratory action, and using them to create new behaviors through introspection.

ⁱ Wolpert, D. M., Ghahramani, Z. & Jordan, M. I. An internal model for sensorimotor integration. *Science* **269**, 1880–1882 (1995).

ⁱⁱ Imamizu, H., Miyauchi, S., Tamada, T., Sasaki, Y., Takino, R., Pütz, B., Yoshioka, T., Kawato, M. Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* **403**, 192–195 (2000).

ⁱⁱⁱ Wolpert, D., Miall, R. C. & Kawato, M. Internal models of the cerebellum. *Trends in Cognitive Science* **2**, 2381–2395 (1998).

^{iv} Llinas, R. R. *The i of the Vortex* (MIT Press, Cambridge, MA, 2001)

^v Maravita, A., Spence, C. & Driver, J. Multisensory integration and the body schema: Close to hand and within reach. *Current Biology* **13**(13), R531–R539 (2003)

^{vi} Brooks R. Intelligence without representation. *Artificial Intelligence* **47**, 139–159 (1991).

^{vii} Saranli, U., Buehler, M. & Koditschek, D. E. RHex: A simple and highly mobile hexapod robot. *Intl. J. of Robotics Research* **20**(7), 616–631 (2001).

^{viii} Brooks, R. A. in Proc. 1st Eur. Conf. on Artificial Life (eds Varela, F. J. and Bourgine, P.) 3–10 (1992).

^{ix} Nolfi, S. & Floreano, D. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines* (MIT Press, Boston, 2000)

^x Jakobi, N. Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adaptive Behavior* **6**(2), 325–368 (1998)

^{xi} Collins, S. H., Ruina, A., Tedrake, R. & Wisse, M. Efficient bipedal robots based on passive-dynamic walkers. *Science* **307**, 1082–1085 (2005).

^{xii} Kanwisher, N. & Wojciulik, E. Visual attention: Insights from brain imaging. *Nature Reviews Neuroscience* **1**, 91–100 (2000).

^{xiii} Crick, F. & Koch, C. A framework for consciousness. *Nature Neuroscience* **6**, 119–126 (2003).

Ants, Bees, Humans and Other Social Natural and Artificial Intelligent Autonomous Agents

Fabio P. Bonsignorio
Heron s.r.l. and Dept. of Mechanics, University of Genova
V. Ceccardi 1/18, Genova, Italy
fabio.bonsignorio@gmail.com

The paper exposes and discusses the concept of 'network embodied cognition', based on natural embodied neural networks, with some considerations on the nature of natural 'collective' intelligence and cognition, and with reference to natural biological examples, evolution theory, neural network science and technology results, network robotics.

It shows that this could be the method of cognitive adaptation to the environment most widely used by living systems and most fit to the deployment of artificial robotic networks. In nature there are many kinds of loosely coupled networks of intelligent agents, largely varying in terms of quantity of agents and cognitive and adaptive capacity (i.e. of computational needs) of each agent. On our planet cooperating biological neural networks seem to be an ubiquitous solution to the adaptation of natural intelligent systems to the environment, much like as DNA information coding is the basis of life. In a distributed agent system a task is performed by a varying number of cooperating units. The intelligence, computation and cognition capabilities of each autonomous agent of the swarm, herd, flock, and group determine, intuitively, the reach and complexity of the tasks that the collective entity performs.

Very different examples of that schema of adaptation to the environment are ant colonies and human societies. At the 'lower' extreme the intelligent unit of adaptation to the environment is the whole ant colony and the collective behaviour emerges from the cooperation schemes of the ants, at the 'higher' extreme – human societies – the collective behaviour is decisive, but the individual intelligent units are much more intelligent, in terms of ability in cognitive and environment interaction tasks of a whole ant colony. The reach of effectiveness and efficiency of a human society is much wider. In the artificial domain the single robot intelligence can spring from artificial neural networks, from symbolic processing, various adaptive methods, from non-linearity decoupling schemes, to polynomial identification. The computation can be based on single Von Neumann machines, on multiprocessor facility or on computer networks. In the natural domain the most widely used method of 'intelligence', computation and 'cognition' are 'embodied' biological neural networks.

A very simplified model of a biological neural network – not considering, usually, the 'plasticity' of natural examples – is constituted by artificial neural networks, a schematic model of natural neural systems. In an artificial neural network the computation is based on




Figure 1

the triggering of an output signal when a threshold of a sum of weighted connection values is reached (the original model is given by Rosenblatt's 'perceptron' , proposed in 1958). There are several important results concerning artificial neural networks which suggest some general remarks. While a single layered perceptron have some environment mapping limits, the 'multilayered perceptron', see Fig. 1, mapping capabilities are remarkable. Hornik et al., and Funahashi (1989) have demonstrated that an artificial neural network with an input layer, an output layer and a number of 'hidden' layers (a 'multilayered perceptron') can approximate within a given error boundary any real valued vector function, by means of an appropriate learning procedure applied to the weights on the connections.

This has led to many successful applications in automatic learning in AI and robotics. If we consider the process of cognitive adaptation to the environment as a learning process in the sense of AI learning theory, whose final result is the 'fit' behaviour of the (living) agent, we can draw some interesting conclusions. We can suppose that it is always possible to build a neural network approximating (in the sense of probabilistic approximate learning) any environment of given complexity, for instance measured by its Vapnik-Chervonenkis dimension. This can be interpreted saying that a learning system for a physical (embodied) agent based on a neural network can be taught to interact effectively in (almost) any environment, or 'to know' (almost) any environment. Several learning procedure for artificial neural networks have been demonstrated. Particle swarm optimization (Kennedy and Eberhart) allow to tune the weights of connections by means of a swarm of agents in the environment: if we assume that any agent has an (almost) identical neural network and (almost) perfect communication between the agents this allow a collective learning (tuning of the weights) of the multi agent system.

Other approaches are genetic algorithms (imitating genetic natural evolution), evolutionary programming, reinforcement learning. Ant algorithms mimic the ant colonies learning process based on external storage of information through pheromones paths (mathematically modeled by Millonas). The importance of 'embodiment' is well shown by the MIT biped passive walker and by theoretical investigation by many people, for example by Pfeifer and Iida, which make clear that part of the 'computation' needed by control, intelligence and cognition are in fact performed by the physical morphology of the agent and by its physical relations within the environment.

A general schema seems to emerge. The unit of (intelligent/cognitive/computational) adaptation to the environment is constituted by loosely coupled groups of neural networks embedded, or more properly 'embodied', into physical agents sensing and acting cooperatively in the physical environment. The weights of the connection are determined in part by biological inheritance (modeled by genetic algorithms optimization), in part through social cooperative exploration (modeled particle swarm optimization) and individual tuning (modeled by reinforcement learning). The information can be maintained in part inside the neural networks of the individual of the group (communicated from agent to agent for instance by means of bees' waggle-dance or human language) in part externally (ant pheromone paths and human libraries). In part in the morphology of the agent body itself. Assuming a common (simplified) measure for the complexity of the environment, e.g. VC dimension, it should be possible to estimate roughly the dimensions of the neural networks of the individual agent of a large natural multiagent system, on the basis of the agent number, cognitive capacity and modalities of collective information maintenance and intra-communication modalities.

These ideas are discussed with reference mainly to the philosophical views of Piaget (to know is to know how), Merleau-Ponty (the role of 'imitation' in understanding), Bateson (the importance of 'relation') , Marx (collective learning of the 'masses' through the 'praxis') and some result of human brain studies ('mirror neurons').

Can “Intentional” Behaviors Emerge From Random Associations?

Simon Bovet
Artificial Intelligence Laboratory
University of Zurich, Switzerland
bovet@ifi.unizh.ch

Understanding natural behaviors

One of the main goals of modern Artificial Intelligence and Cognitive Science is to understand the principles underlying intelligent behaviors. Also, most “intelligent” systems that are known — and from which inspiration is drawn — are in fact natural agents. It is therefore legitimate to ask the following question: how much can we understand about intelligent systems which are products of natural evolution, i.e. which were not *intentionally* designed for any specific task, by investigating artificial agents precisely constructed, at least partly, to perform particular tasks? Consequently, we propose here an alternative approach consisting of building a system a priori not designed for any particular task, and observing what behaviors can emerge from its interaction with the environment.

Model and experiments

In the proposed control architecture, each sensor or motor the robot is equipped with is represented by a population of artificial neurons. All these populations are systematically coupled together by artificial synapses with Hebbian-like plasticity. Note that no distinction is made between sensors and motors, and that the homogeneous network presents no specific structure. In the following experiments, we observe the behaviors that are produced by externally exciting or inhibiting one arbitrary neuron in the network.

In a first experiment¹, we use a robot equipped with two wheels, an omnidirectional camera and a tactile whisker sensor. When exciting the neuron corresponding to tactile stimulation (i.e. making the robot “feel” as if its whisker was touching an object), propagation of neural activity across the network produces the following, seemingly “intentional” behavior: the robot moves toward and follows any object which is in the center of its field of view, until its whisker gets in contact with the object.

In a second experiment², the same control architecture is used on a wheeled robot equipped with camera, temperature and ambient light sensors. The behaviors observed when modifying the activity of either the temperature or light sensor correspond exactly to two navigation strategies observed in insects: the robot returns to its initial position using either path integration or visual landmark homing, respectively.

In a third experiment³, a robot is engaged in a more complex task: in a T-shaped maze, a tactile cue indicates the arm of the maze where at its end a reward is delivered to the robot. After several runs (the side of the cue and the reward being each time randomly reassigned), we observe that the robot learns to turn into the correct arm using the tactile stimulation. A closer inspection reveals that the robot is able to learn this delayed reward task *without* keeping any information about its action and the cue — an assumption otherwise present in all current reinforcement-learning models.

Conclusions

These experiments demonstrate the potential of our novel approach, which consists in observing behaviors that can emerge from a system not specifically designed for any particular task. We show that seemingly intentional behaviors do not necessarily require specific assumptions on the control architecture, but can emerge from the interaction with the environment of an agent using a homogeneous neural network learning only cross-modal correlations. Moreover, this method allows some forms of natural behaviors to be explored from a new perspective: for instance, how two different insect navigation strategies can rely on one single mechanism; or how memory can be at least partly off-loaded into the environment and exploited for a delayed reward learning task.

References

1. Bovet, S. & Pfeifer, R. *Emergence of Coherent Behaviors from Homogenous Sensorimotor Coupling* (Proc. Int. Conf. on Advanced Robotics, Seattle, 2005).
2. Bovet, S. *Emergence of Insect Navigation Strategies from Homogeneous Sensorimotor Coupling* (Proc. Int. Conf. on Intelligent Autonomous Systems, Tokyo, 2006)
3. Bovet, S. & Pfeifer, R. *Emergence of Delayed Reward Learning from Sensorimotor Coordination* (Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Edmonton, 2005)

Social Human-Robot Interaction

Cynthia Breazeal
The Media Lab
MIT, Cambridge, MA, USA
cynthiab@media.mit.edu

What is your vision of robots in the future? I envision socially and emotionally intelligent robots that communicate and interact with us, understand and even relate to us, in a personal way. They are able to engage us as full-fledge partners and participate in fundamental forms of human interaction such as collaborative teamwork, social support, and social learning or teaching.

Robots today, however, simply do not understand “people as people.” For instance, by and large, robots treat us either as other objects in the environment, or at best they interact with us in a manner characteristic of socially impaired people. For instance, robots are not really aware of our goals and intentions. As a result, they don’t know how to appropriately adjust their behavior to help us as our goals and needs change. They are not deeply aware of our emotions, feelings, or attitudes. As a result they cannot prioritize what is the most important to do for us according to what pleases us or to what we find to be most urgent, relevant, or significant. They do not readily learn new skills and abilities from interacting with or observing people. As a result, they cannot take advantage of the tremendously rich learning environment of humans. These shortcomings and others in socio-emotional intelligence must be addressed for robots to achieve their full beneficial potential for us in human society (and vice versa).

Promisingly, there have been initial and ongoing strides in all of these areas [e.g., Breazeal, 2002; Scassellati, 2000 and for reviews see Picard, 1997; Fong et al, 2003; Schaal, 2000]. In particular, in my own group we have been steadily working to endow robots with socio-emotive-cognitive skills (and evaluating their impact in human subject studies) on a wide range of human-robot interactions such as collaborative teamwork and social learning [see <http://robotic.media.mit.edu> for our group publications]. We have been developing an architecture based on embodied cognition theories from psychology [e.g., Barsalou, 2003; Sebanz et. al., 2006] to give our humanoid robot visual and mental perspective taking abilities using a simulation theoretic framework. Specifically, *Simulation Theory* holds that certain parts of the brain have dual use; they are used to not only generate our own behavior and mental states, but also to predict and infer the same in others. To understand another person's mental process, we use our own similar brain structure to simulate the introceptive states of the other person. This is the process by which the robot infers its human partner's goals, attention, beliefs, and affect from observable behavior.

Within a teamwork task, the robot is able to compare and reason about how these human internal states relate to its own in order to provide the person with informational support and instrumental support. For example, in the case of informational support, the robot can relate its own beliefs about the state of the shared workspace to those of the human based on the visual perspective of each. If a visual occlusion prevents the human from knowing important information about that region of the workspace, the robot knows to direct the human's attention to bring that information into common ground. Furthermore, based on principles of Joint Intention Theory, the robot uses a versatile range of non-verbal behaviors to coordinate teamwork and establish and maintain mutual beliefs about progress in the task. In the case of instrumental support, the robot can infer the human's intent (e.g, a desired effect on the workspace) from observing their behavior. If the human fails to achieve their intent, the robot can reason about how it might best help the human achieve their goal either by achieving that goal for them or by providing mutual support that helps the human achieve his or her goal. The representations by which the robot reasons and plans is inspired by embodied cognition theories [e.g, Barsalou, 2003; Sebanz et. al., 2006].

Within a social learning context, the robot uses its perspective taking abilities to interpret the intent behind the human's demonstrations. Imagine a scenario where the demonstrations are provided a person who does not have expertise in the learning algorithms used by the robot. As a result, the teacher may provide sensible demonstrations from a human's perspective; however, these same demonstrations may be insufficient, incomplete, ambiguous, or otherwise "flawed" from the perspective of providing a correct and sufficiently complete training set needed by the learning algorithm to generalize properly. We have tackled this issue by designing the robot to be a socially cognitive learner in a tutelage-based scenario. As the robot observes the human's demonstrations, it internally simulates "what might I be trying to achieve were I performing these demonstrations in their context?" The robot therefore interprets and hypothesizes the intended concept being taught not only from its own perspective, but from the human teacher's visual perspective as well. Through this process, the robot successfully identifies ambiguous demonstrations given by the human instructor, and clarifies the human's intent behind these confusing demonstrations. Once these problematic demonstrations are disambiguated, the robot correctly learns the intended task. In sum, I believe that maintaining mutual beliefs and common ground in human-robot teaching-learning scenarios will make robots more efficient and understandable learners, as well as more robust to the miscommunications or misunderstandings that inevitably arise even in human-human tutelage.

Unlike the original goal of Artificial Intelligence, which is to create a technological system with human equivalent intelligence, my goal is to create robots that are human-synergistic and human-compatible. Specifically, robots should bring value to us and are valued by us because they are different from us in ways that enhance and complement our strengths and abilities. I argue that the goal of creating robots that can engage us as full-fledged partners is as challenging and deep as the original goal of Artificial Intelligence because it requires scientists of the natural and artificial to deeply understand human intelligence, behavior, and nature across multiple dimensions (i.e., cognitive, affective, physical, and social) in order to design synthetic systems that support and complement people and our goals. It also requires scientists to understand the dynamics of the flesh-and-blood-human-with-robot system. Theories about disembodied minds or even embodied minds operating in isolation fall far short of this goal.

References

1. Barsalou, L.W., Niedenthal, P.M., Barbey, A., & Ruppert, J. (2003). Social embodiment. In B. Ross (Ed.), *The Psychology of Learning and Motivation*, Vol. 43 (pp. 43-92). San Diego: Academic Press.
2. Breazeal, C. (2002) *Designing Sociable Robots*. MIT Press. Cambridge MA.
3. Please see <http://robotic.media.mit.edu> for our group publications
4. Fong, T., Nourbakshsh, I., Dautenhahn, K. (2003). A Survey of Social Robots. *Robotics and Autonomous Systems*, **42**, pp. 143 – 166.
5. Picard, R. W. (1997) *Affective Computing*. MIT Press. Cambridge, MA.
6. Sebanz, N., Bekkering, H., and Knoblich, G. (2006) Joint action: bodies and minds moving together. *Trends in Cognitive Science* **10:2**, 70-76.
7. Scassellati, B. (2000) Investigating Models of Social Development Using a Humanoid Robot, in Barbara Webb and Thomas Consi, eds., *Biorobotics*, MIT Press. Cambridge MA.
8. Schaal, S. (1999) Is imitation learning the route to humanoid robots? *Trends in Cognitive Science* **3(6)**, 233-242.

Whither Embodiment?

Andrew Brooks
The Media Lab
MIT, Cambridge, MA, USA
zoz@media.mit.edu

In his introduction to Blake and Yuille's book on Active Vision, Rodney Brooks described with forceful eloquence classical AI's reductionist approach to making difficult problems computationally tractable through simplistic simulations: "an intellectual trap had been sprung" [11]. As intelligent robotics research has accelerated, the value of embodiment is now widely (though not universally) accepted. But are there similar traps in this direction as well?

Arguments in favor of embodiment have been strongly presented over the last fifteen years. In Brooks' active vision example, the practical ability to affect the world and perceive the results gives an agent increased analytical powers compared with a passive viewer. Moreover, being physically situated in the real world not only grounds the reasoning and behavior of the agent but provides a hard metric of its performance: it has been argued that an autonomous agent must be capable of coping with the uncertainties associated with operating in the real world in order to be considered truly intelligent (e.g., [10]). The physical participation in the world experienced by intelligent biological systems has also been used to support embodied cognition theories stating that the embodiment of these systems is inseparable from their resultant reasoning capabilities (e.g. [2]), and that this may be true also of man-made intelligent machines.

Valid practical arguments have also been raised against the short-term use of embodiment, however, specifically concerning the accepted work model involved in university robotics research. Perhaps the most inflammatory rhetoric on the subject has come from AI luminary Marvin Minsky, who infamously described intelligent robotics research as a "fad" that wastes years of graduate students' lives [1]. Hardware development and maintenance can indeed be a significant distraction from core efforts. Here we see the kernel of the reverse intellectual trap: not oversimplification but overcomplexity. Sadly, many robotics laboratories demonstrate that the insistence on embodiment for problems that are clearly too narrow for situated intelligence to even begin to develop is responsible for a similar reduction to absurd simplicity. Devotion to physical mimicry of biological systems via "all on-board sensing", for example, often results in a lab full of colored markers — little more than a physical simulation of a blocks world. In our work we use a distributed suite of on-board and environmental sensors [4, 5]. Embodiment should be a means for concentrating on problems that necessitate physical presence, not for arbitrarily introducing new stumbling blocks.

One such important problem is human-robot communication. While all sorts of tricks can be performed to give robots more accurate models of the physical world, it is the behavior of their human counterparts that is in many ways the most unpredictable. Understanding that robots will not operate in a social vacuum is important, and significant work has been devoted to various aspects of this area (e.g. [14, 3]). Embodiment is crucial both to understanding human behavior in terms of human physicality, and to learning how one's own embodiment can contribute to social cooperation with humans. Humanoid robot development is often justified by the fact that human environments are physically designed for humans; but we must remember that these environments are also populated with humans whom the robot must learn from, teach, respond intelligently to, and otherwise relate to.

A major component of the above is body understanding, a necessary subcomponent of higher-level socially intelligent embodied behaviors such as imitation (e.g. [20]). Imitation has the potential to be a powerful tool for teaching new skills to robots. Moreover, there is strong evidence that imitation is an important factor underlying mental simulation [e.g. 18]. This mechanism of making assumptions and predictions about other agents [e.g. 12] requires that individuals be enough "like" one another in body and psychology, and to be able to mentally "map" each other's bodies onto their own. Biological evidence has been found for "mirror neurons" that perform this mental body mapping in primates [16]. We have been developing similar mechanisms for humans to bodily communicate their own physical correspondence to humanoid robots [8]. We are currently working on mechanisms for bodily attention that allow the

robot to more appropriately determine when and how to refine its body correspondence, and to allow the human to apply deliberate supportive “scaffolding”.

Once body correspondence has been achieved, an area that is in need of more exhaustive research is the use of the body for communication. Humans convey a significant amount of information to one another through gesture and “body language”, and while gestural interfaces in general have received extensive attention (see [19] for an overview) the nuanced ways in which bodily expression underpins a great deal of copresent human communication are in need of more attention from robotics researchers. One such example is deictic reference, in which minimal gesture and speech are combined to achieve common ground between agents and objects; the process is more complex than it may appear, and we therefore developed a system that formalizes this activity into a robust “deictic grammar” that incorporates bodily activity as well as spoken phrases [9]. In general, however, “body language” is not a true language with discrete rules and grammars, but it does convey coded messages that humans can interpret [17]. Many of these messages are also culturally dependent, and robots that work with humans need to know how to make use of this communication channel in a socially appropriate manner. The bulk of robotics research into bodily expression has concentrated on facial representation of emotional content, believing this to be culturally stable [13]. However more recent theories have challenged the accepted notion that facial expressions are unconscious and culturally invariant representations of internal emotional state, arguing instead that they are quite deliberate communications that are heavily influenced by the context in which they are expressed [15]. We have similarly demonstrated the value of subtle bodily expression to task performance by human-robot teams [6], and developed new algorithms for full-body communication in humanoid robots [7].

The debate over the value of embodiment in general intelligent machine research cannot be resolved with a general pronouncement. Intelligent disembodied systems, by some definition, may be feasible. But an understanding of embodiment in the setting of social engagement and intelligent interchange with humans is irreplaceable. We must attend to how this fundamental aspect of human cognition can best be exploited in our goal of developing intelligent systems that operate in concert with us.

References

- [1] Baard, M. *AI Founder Blasts Modern Research*. Wired News, May 13 2003.
- [2] Barsalou, L.W., Niedenthal, P.M., Barbey, A.K., and Ruppert, J.A. *Social embodiment*. In Ross, B.H., editor, *The Psychology of Learning and Motivation*, volume 43, pages 43–92. Academic Press, 2003.
- [3] Breazeal, C. *Towards sociable robots*. *Robotics and Autonomous Systems*, 42(3–4):167–175, 2003.
- [4] Breazeal, C., Brooks, A.G., Gray, J., Hancher, M., McBean, J., Stiehl, W.D., and Strickon, J. *Interactive robot theatre*. In Proc. Int. Conf. on Intelligent Robots and Systems, Las Vegas, Nevada, October 2003.
- [5] Breazeal, C., Brooks, A.G., Gray, J., Hoffman, G., Kidd, C., Lee, H., Lieberman, J., Lockerd, A., and Chilongo, D. *Tutelage and collaboration for humanoid robots*. *Int. J. Humanoid Robots*, 1(2):315–348, 2004.
- [6] Breazeal, C., Kidd, C.D., Lockerd Thomaz, A., Hoffman, G., and Berlin, M. *Effects of nonverbal communication on efficiency and robustness in human-robot teamwork*. IROS 2004.
- [7] Brooks, A.G. and Arkin, R.C. *Unpublished technical report*. Sony Intelligence Dynamics Laboratories, August 2005.
- [8] Brooks, A.G., Berlin, M., Gray, J., and Breazeal, C. *Untethered robotic play for repetitive physical tasks*. In Proc. ACM International Conference on Advances in Computer Entertainment, Valencia, Spain, June 2005.
- [9] Brooks, A.G. and Breazeal, C. *Working with robots and objects: Revisiting deictic reference for achieving spatial common ground*. In Proc. First Annual Conference on Human-Robot Interaction (HRI’06), Salt Lake City, UT, March 2006.
- [10] Brooks, R.A. *Intelligence without reason*. Proc. Int. Joint Conf. on Artificial Intelligence, pages 569–595, 1991.
- [11] Brooks, R.A. *Introduction*. In Blake, A. and Yuille, A.L., editors, *Active Vision*. MIT Press, 1992.
- [12] Davies, M. and Stone, T. *Mental Simulation*. Blackwell Publishers, Oxford, UK, 1995.
- [13] Ekman, P. and Davidson, R.J. *The Nature of Emotion*. Oxford University Press, 1994.
- [14] Fong, T., Nourbakhsh, I., and Dautenhahn, K. *A survey of socially interactive robots*. *Robotics and Autonomous Systems*, 42:143–166, 2003.
- [15] Fridlund, A. *Human Facial Expression: An Evolutionary View*. Academic Press, San Diego, CA, 1994.
- [16] Iacoboni, M., Woods, R.P., Brass, M., Bekkering, H., Mazziotta, J.C., and Rizzolatti, G. *Cortical mechanisms of human imitation*. *Science*, 286:2526–2528, 1999.
- [17] Machotka, P. and Spiegel, J. *The Articulate Body*. Irvington, 1982.
- [18] Meltzoff, A. and Decety, J. *What imitation tells us about social cognition: A rapprochement between developmental psychology and cognitive neuroscience*. *Phil. Trans. R. Soc. London B*, 358:491–500, 2003.
- [19] Pavlovic, V.I., Sharma, R., and Huang, T.S. *Visual interpretation of hand gestures for human-computer interaction: A review*. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):677–695, July 1997.
- [20] Schaal, S. *Is imitation learning the route to humanoid robots?* *Trends In Cognitive Sciences*, 3(6):233–242, 1999.

Adaptive Dynamical Systems

A promising tool for embodied artificial intelligence

Jonas Buchli, Ludovic Righetti, and Auke Jan Ijspeert
Biologically Inspired Robotics Group
Ecole Polytechnique Fédérale de Lausanne, Switzerland
jonas@buchli.org, ludovic.righetti@a3.epfl.ch, auke.ijspeert@epfl.ch

Introduction

Nonlinear Dynamical Systems (NDS) are an interesting tool to devise locomotion controllers for mDOF robots, e.g. CPGs for quadrupeds [1]. Furthermore, this approach has a strong foundation based on the investigation of biological coordination tasks [7]. One of the difficulties that limits the usability of this approach to robots is the problem of how to design a suitable NDS to control a given robot and, if the NDS can be found, how to tune its parameters. Furthermore, one would like to have a certain flexibility and adaptivity in the controllers in order for them to adapt to changing body properties and non stationary, complex environments.

In a “proof of principle” implementation in a simulation [2] we showed that we can extend a simple dynamical system (i.e. an oscillator) with an additional state variable and the corresponding evolution law (i.e. differential equation) in order to make it adaptive to a mechanical structure. The mechanical structure (body) and the adaptive frequency oscillator (controller) make up a simple adaptive locomotion system. This locomotion system is capable of adapting to changing body properties or an addition of external load. In this contribution we show that we can successfully implement these concepts in the real world. As briefly discussed in [2] we consider the therein presented system as an example of a much broader class of systems. Namely, we propose that the extension of dynamical systems by state variables with different time scales could be a useful tool for many applications in robotics, machine learning and other areas where dynamic adaptive behavior is desirable. We call such dynamical systems which are extended with additional state variables, which evolve on different (normally slower) time scales than the rest of the state variables, Multiscale Dynamical Systems. In contrast to conventional adaptive control, these systems and the feedback loops are usually strongly non-linear. This allows for interesting pattern formation capabilities.

Adaptive Dynamical Systems

We implement an adaptive dynamical system by using a plastic NDS, as described in the following. The controller is described by the dynamical system which has four additive contributions, i.e. the *Effective flow* $\mathbf{F}_{\Sigma}(\mathbf{q}, t)$ is the sum of all the parts described in the following and is the effective dynamics: (1) The *Intrinsic flow*: $\mathbf{F}_0(\mathbf{q})$ describes the part of the system that we consider fixed (i.e. constant over time). (2) The *Plastic flow* $\mathbf{F}_p(\mathbf{q}, t)$ describes the plastic part of the system, e.g. the adaptation in the neural network (learning) of a subject, metabolic adaptation, etc. (3) *External influences*: $\mathbf{F}_e(t)$ this term describes disturbances from the environment, sensory input, feedback, etc. It can also be used to impose a training signal on the system. (4) Finally the *Noise* term $\mathbf{F}_{\xi}(t)$ includes thermal and other noise sources as well as not modeled dynamics. The *adaptation of the dynamics* is modeled by the plastic flow. The idea of the plastic dynamical system is the following: The history of the phase point leaves traces in the system, i.e. history of the systems behavior *and* external influences shape future dynamics. There are two main contributions to the change of the dynamics: (1) The “memory trace function” \mathbf{T} and the (2) forgetting dynamical system \mathbf{R} :

$$\dot{\mathbf{F}}_p = \underbrace{\mathbf{T}(\mathbf{F}_{\Sigma}, \mathbf{q}, t)}_{\text{trace}} + \underbrace{\mathbf{R}(\mathbf{F}_{\Sigma}, \mathbf{q}, t)}_{\text{forgetting}}$$

Example: Controller adapting to body dynamics

As a concrete example of the above ideas we will show the implementation of a controller which is able to autonomously adapt to the body dynamics of a quadruped underactuated robot [5]. The robot is actuated only at his hip joints, and has springs in the knee joints. Such a robot thus has a very pronounced body dynamics in terms of resonant frequencies. From former studies on the robot (and animals) it is known that the resonant frequencies are a good choice for efficient locomotion. We present, based on the above general concepts, a very simple online adaptive controller based on adaptive frequency oscillators [2, 6]. The controller needs no complicated signal processing, no algorithmic description and there is no separation between learning substrate and learning algorithm, which makes the whole system treatable in a unified way.

By the interaction of the controller and the body-environment systems, a distinct locomotion pattern emerges, results from simulations furthermore suggest that the gait pattern and the energy consumption distribution is in line with observations in mammals [4, 3].

Outlook

Despite its application to robotic locomotion, we are convinced that the presented approach will be useful in at least two ways to bring forward our understanding of intelligent processes: (1) as a mathematical framework for learning and adaptive systems which allows to formulate problems in the unified, rigorous language of dynamical systems, and (2) more specifically on the road from locomotion to cognition, it will allow us to gradually implement intelligent high-level controllers which are grounded in the real world.

References

- [1] J. Buchli and A.J. Ijspeert. Distributed central pattern generator model for robotics application based on phase sensitivity analysis. In *Proceedings BioADIT 2004*, volume 3141 of *Lecture Notes in Computer Science*, pages 333–349. Springer Verlag Berlin Heidelberg, 2004.
- [2] J. Buchli and A.J. Ijspeert. A simple, adaptive locomotion toy-system. In *Proceedings SAB'04*, pages 153–162. MIT Press, 2004.
- [3] J. Buchli, L. Righetti, and A.J. Ijspeert. Adaptive dynamical systems for movement control. In *Proceedings of 3rd International Symposium on Adaptive Motion in Animals and Machines – AMAM 2005*, page 7. Verlag ISLE, Ilmenau, 2005. Abstract.
- [4] J. Buchli, L. Righetti, and A.J. Ijspeert. A dynamical systems approach to learning: a frequency-adaptive hopper robot. In *Proceedings of the VIIIth European Conference on Artificial Life ECAL 2005*, *Lecture Notes in Artificial Intelligence*, pages 210–220. Springer Verlag, 2005.
- [5] F. Iida, G.J. Gomez, and R. Pfeifer. Exploiting body dynamics for controlling a running quadruped robot. In *Proceedings of the 12th Int. Conf. on Advanced Robotics (ICAR05)*, pages 229–235, 2005.
- [6] L. Righetti, J. Buchli, and A.J. Ijspeert. Dynamic hebbian learning in adaptive frequency oscillators. *Physica D*, 216(2):269–281, 2006..
- [7] G. Schöner and J.A.S. Kelso. Dynamic pattern generation in behavioral and neural systems. *Science*, 239(4847):1513–1520, 1988.

Teaching Cognitive Robotics using Multiple Online Remote Robots

Catalin Buiu, Nicolae Moanta

Department of Automatic Control and Systems Engineering
“Politehnica” University, Bucharest, Romania
cbuiu@ics.pub.ro

Introduction

In teaching cognitive robotics using mobile robots for experiments there are many situations when experiments on local robots are difficult or even impossible because of the low availability of robots (which are often expensive pieces of equipment) and of the large number of students that have to take the experiments. This is why more and more teachers are using the advantages offered by telepresence, educational technologies and software agents^{1,2} to offer the possibility to take experiments on online robots that can be accessed from anywhere at any time (Fig. 1).

This paper presents some related research efforts carried on at our Laboratory of Autonomous Robotics in order to use multiple online robots for teaching cognitive science and artificial intelligence concepts.

By using a telepresence client³ (see Fig. 2), the students can remotely control 1 or 2 robots that are connected to the same computer. The control algorithms can be very simple, such as Braitenberg vehicles, simple obstacle avoidance and wall following applications or more complex control algorithms, such as fuzzy logic and neural networks based controllers.

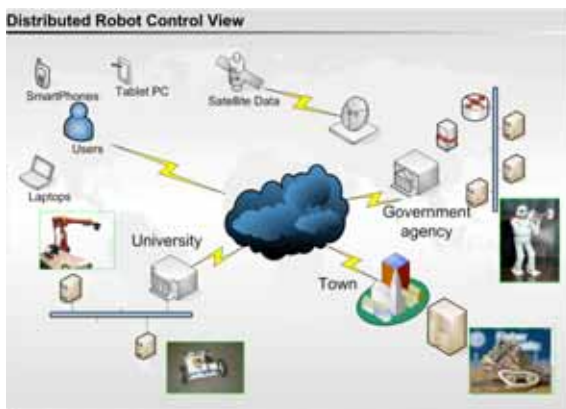


Fig. 1. Distributed robot control.




Fig. 2. Simple telepresence client.

Recently, a new architecture³ (see Fig. 3) has been developed where there is one online robot that can be remotely controlled using a Web interface and Web services.

This paper proposes a new distributed architecture (Fig. 4) that is currently under development at our Laboratory and which involves multiple online robots that are located in different laboratories. By using this architecture and the related Web interface the student can control the multiple robots that are able to communicate to each other. Additional robots can be added at any time to this network of online robots.




Fig. 3. Web control application

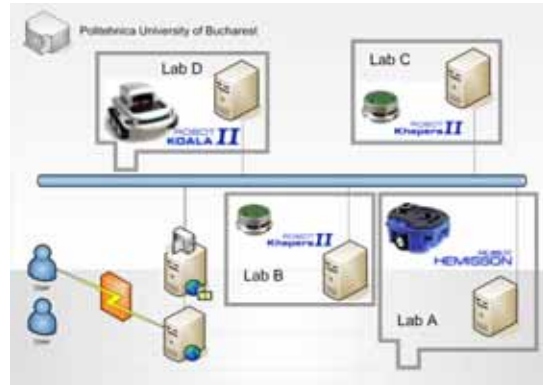


Fig. 4. Network of collaborative online robots.

Technical overview

The application is based on the recent set of technologies for building and running connected systems named Windows Communication Foundation. The communications infrastructure between robots is built around the Web services architecture. The service-oriented programming model simplified the development of connected systems. Advanced Web services support provides secure, reliable, and transacted messaging along with interoperability.

Conclusions

By using the advantages of new educational and telecommunication technologies, we proposed several architectures for using online remote robots for teaching basic concepts of cognitive robotics and artificial intelligence. While in the beginning the robots were connected to the same computer, now the remote user can access a network of multiple online robots that can communicate to each other.

References

1. Buiu, C., I. Dumitrache, F. Mihai & D. Gillet. *Experiments in Remote Collective Robotics. New Pedagogical Scenarios*. In: *Proceedings of ED-MEDIA 2004*, World Conference on Educational Multimedia, Hypermedia and telecommunications, Lugano, June 21-26, 2004.
2. Buiu, C. *Hybrid Educational Strategies for a Laboratory Course on Cognitive Robotics*, IEEE Transactions on Education, submitted.
3. Buiu, C. & Moanta N. *Distributed Application for Mobile Robot Remote Motion Control*, 7th IFAC Symposium on Advances in Control Education, Madrid, June 2006.

Mysterious Tools

Seth Bullock

School of Electronics and Computer Science
University of Southampton, Southampton, SO17 2BJ, UK
sgb@ecs.soton.ac.uk

This paper explores the role of *explanatory opacity* in the development and use of models and tools derived from AI and Alife systems. This issue is shown to have been significant for the very first instances of simulation modelling via machine intelligence. Moreover, it is argued that the same issue is critically important today for the prospects of Alife-style simulation models as tools with which to explore and possibly manage or control complex adaptive systems.

The artificial sciences have twin mysteries as their explanatory targets: the nature of life and mind. Increasingly, we can describe (perhaps somewhat informally) what it is to *be* living or to *be* minded. The system-level phenomena associated with each are ready to hand: metabolising, reproducing, recognising, anticipating, etc. Moreover, via research in neuroscience and organismic biology, we have also gained an increased understanding of the material components that constitute natural minds and living creatures and their organisation. But these advances have not pierced the key mystery which seems to remain as deep as ever: how do those material components come to bring about the systemic properties of life and mind by virtue of their organisation?

By contrast with the natural sciences, artificial intelligence (AI) and artificial life (Alife) have tended to pursue this question by attempting to *engineer* systems that demonstrate the required system-level behaviours. In this pursuit, AI and Alife researchers are sometimes inspired and directed by what we know of the component properties or organisational structures of natural life and intelligence. As such, it is perhaps understandable that the history of both AI and Alife is strewn with artificial systems that have appeared (at least initially) to share some of the mystery associated with life and mind: Turing's deliberately enigmatic intelligence test, Eliza's suspiciously ersatz conversation, the internal representations of artificial neural networks distributed across a "hidden" layer, the surprisingly sophisticated behaviour of Braitenberg's vehicles, the purposeful organisation of artificial swarms, the creative potential of coevolutionary arms-races, the "emergent complexity" of an artificial life world.

For each of these synthetic systems, global behaviour tends to be (a posteriori) predictable—we can say what the systems do or achieve, and describe patterns and trends in their behaviour. We also have a rather complete knowledge of the systems' low-level components and their interactions, since we engineered them. Again, the mystery lies in understanding how the systemic behaviour arises from system components and their interactions. This mystery is sometimes amplified for rhetorical or dramatic effect (Eliza) or to draw attention to our preconceptions (Braitenberg's *Vehicles*), but it is also often real and challenging. The notion of distributed representation within feed-forward artificial neural networks, for example, was certainly challenging when first articulated.

However, the issue to be considered here involved the situation we face when objects of AI and Alife enquiry (neural networks, genetic algorithms, etc.) are appropriated as scientific or engineering *tools*. Often this transition occurs before we have a strong grasp of what it is that, say, evolutionary algorithms or agent-based

models are good for. The explosion of artificial neural network models within cognitive science in the late 80s, and the current appetite for agent-based simulation models within, e.g., social science, are examples of the kinds of “mysterious tools” that AI and Alife are capable of generating.

Interestingly, the first example of machine intelligence generating such a tool highlights the issue. Charles Babbage, inventor of the first automated computing devices, chose to employ his machinery to run a simple simulation model of geological processes. His aim was to demonstrate that a single process unfolding mechanically without intervention could nevertheless generate behavior that exhibited discontinuities (analogous to the geological record containing evidence of supposed “miracles”). For present purposes, what is interesting here is that this aim could only be met by exploiting what was for the audience a mysterious relationship between the machine’s mechanisms and the machine’s output. Babbage deliberately reinforced this mystery during *in situ* demonstrations of the machine’s behaviour in his Marylebone residence, where he manipulated the audience’s expectations in order to maximise the surprise elicited by the machine’s behavioural jumps. For a modern audience, it may be trivial to understand that a computer can at first do one thing and then, automatically, start to do something else, but for Babbage’s audience the implications of Babbage’s demonstration were more profound.

However, this use of mysterious machinery was not without its detractors. Reverend William Whewell, for instance, denied “the mechanical philosophers and mathematicians of recent times any authority with regard to their views of the administration of the universe” (W. Whewell, *Astronomy and General Physics Considered with Reference to Natural Theology*, Pickering, London, 1834). For Whewell, the reason for this denial lay not in any limits on machine *performance*, but in the loss of explanatory clarity that accompanied their use:

“Whewell brutally denied that mechanised analytical calculation was proper to the formation of the academic and clerical elite. In classical geometry ‘we tread the ground ourselves at every step feeling ourselves firm’, but in machine analysis ‘we are carried along as in a rail-road carriage, entering it at one station, and coming out of it at another ... It is plain that the latter is not a mode of exercising our own locomotive powers ... It may be the best way for men of business to travel, but it cannot fitly be made a part of the gymnastics of education’” (S. Schaffer, “Babbage’s intelligence: Calculating engines and the factory system”, *Critical Inquiry*, 21(1), 203-227, 1994).

Here, the context of the disagreement between Babbage and Whewell will be spelled out, and its relevance to the current use of “artificial worlds” as useful (predictive) models will be discussed. It is concluded that where machines automate adaptive, autonomic, intelligent or complex behaviour (rather than mere rote procedures), Whewell’s worries are legitimate, and as such should be a pressing concern for the sciences of the artificial.

Who Made the Decision?! An Investigation into Distribution of Memory, Actions and Decision Makings in a Multi Agent System¹

Roozbeh Daneshvar, Caro Lucas
School of Electrical and Computer Engineering
University of Tehran, Tehran, Iran
roozbeh@daneshvar.ir, lucas@ipm.ir

In this paper we have investigated a group of multi agent systems (MAS) in which the agents change their environment and this change triggers behaviors in other agents of the group (in another time or another position in the environment). The structure makes a super organism in the group such that new behaviors is observed from the whole group. This distribution exists in many aspects like a super memory (or even a super brain) that exists in the environment and is not limited to memories of individuals. As another instance, the actions performed by agents are transformed via a chain of interactions and the actions observed by the group are not necessarily created by individuals (if the environment is capable of transforming). Another general instance is distributed decision making that is done by the group of agents which is in a higher level consisting both individual and group decision makings, and can be viewed as emergent rather than consciously planned.

Introduction

When the agents in a MAS communicate with each other, the communication media is not limited to facilities for direct communication as there are also approaches for indirect communication between agents. One of these media for indirect communication is the environment the agents are located in. The agents change their common environment and this change triggers behaviors of other agents. When an agent fires behaviors of another agent (even if this process is done indirectly and implicitly), a kind of communication between two agents has formed (while this communication might happen in different positions and different times). Two instances of these systems are stigmergy and swarm intelligence.

The surrounding space which agents are located in, is a physical common part that performs the role of a media for agent communications. This is only one aspect of the environment while the environment is not limited to physical surrounding space of agents. We can find other aspects (not necessarily physical) that play the role of a common environment. For instance, as a part of the environment, we can name the agents themselves. The agents of a MAS are parts of their common environment as they keep records of past events (which are able to change their state respectively). When the state of an agent triggers behaviors of another agent, a kind of communication has formed between two agents (not necessarily including the media agent). An example of this role of individuals as the environment is rumor in societies of people.

We have investigated the elements that make a MAS intelligent and we have considered how we can enhance a system without necessarily enhancing the individuals. We have considered the use of agents with bounded rationality that make simple decisions according to their perceived state of the environment and these simple decision makings lead to higher level decision makings in the system layer.

¹ Any kind of military uses from the content and approaches of this article is against the intent of the authors

The agents are unconscious about their non-intentional behaviors. These systems are inspired by the natural organic systems where there is a greater brain making decisions that does not exist in individuals. The meta structure existing in groups of natural creatures or cells has the ability of containing a super memory that does not exist in individuals (like ants and their pheromone trails in an ant colony) and has abilities of data transformations (like reactions of ants to pheromone trails in the environment) and meta decision makings (like the performance of an ant colony comparing the abilities of ants). As an example, in [1] we have a group of RoboCup soccer players with emergent behaviors and some virtual springs that connect them together, each spring representing decision concern. The environment (potential field) acts as a superbrain guiding the action of each player without the awareness of how the different concerns have been fused into that action.

Distributed Memory, Actions and Decision Makings

An implicit and non declarative memory exists in a MAS in which the agents communicate via their environment. When an agent changes the state of the common environment and this change fires a behavior in another agent at a later time, the changed part of environment performs the role of a memory element that remains when time passes and is used by another member of the group. The environment is a general implicit memory that keeps track of agent-agent and agent-environment interactions and hence the state of environment is changed. This kind of memory is not limited to a simple recorder of changes as it affects the process of decision making in agents and performs the role of a distributed decision maker also (this is the reason we call it a distributed brain also). As an observer, we can quantify the memory existing in the environment and show it as a quantity in the system (while it is not necessarily this much specified for the agents). The environment also plays the role of a media for transformation of actions. When an agent performs an action (provided there is proper structure for this transformation), this action is capable of movement among members of the team. For instance, in a group of flocking birds, the change in position made by one of the individuals affects the neighbors while their change affect others respectively. The chain effects of actions can be evaluated as quantities. In [1] when an agent makes a movement, this action affects other members of the team also, as they are connected together and the spring forces make them move. A group of environment elements act as decision maker parts. These elements might be passive and they still affect the process of decision making (their state fires behaviors of agents and affects their decision making and hence a part of the group decision making is done by environment elements). Whether the environment elements take part in decision making process or not depends on the structure of the environment and the interactions between agents and environment. In [1] the position of other agents (as a part of the environment) leads to actions of agents and hence the decision making is done in the system level while the agents are not aware of.

References

1. M. A. Sharbafi, R. Daneshvar & C. Lucas, *New Implicit Approach in Coordination between Soccer Simulation footballers* (Fifth Conference on Computer Science and Information Technologies, 19-23 September, Yerevan, Armenia, 2005)

AI Re-Emerging as Research in Complex Systems

Kemal A. Delic, Umeshwar Dayal
Hewlett-Packard Co.
HP Technology Services Group – HP Labs
[kemal.delic | umeshwar.dayal]@hp.com

The history and the future of Artificial Intelligence could be summarized into three distinctive phases: embryonic, embedded and embodied. We briefly describe early efforts in AI aiming to mimic intelligent behavior, evolving later into a set of the useful, embedded and practical technologies. We project the possible future of embodied intelligent systems, able to model and understand the environment and learn from interactions, while learning and evolving in constantly changing circumstances. We conclude with the (heretical) thought that in the future, AI should re-emerge as research in complex systems. One particular embodiment of a complex system is the Intelligent Enterprise.

The Early Life of AI Research

Research in AI started with the noble objective of creating computer programs that exhibit human intelligence, promising some great achievements in a projected timeline. It has passed through several alternating cycles of brave promises and grave disappointments. Some embryonic applications in search, game playing, language understanding, expert systems, vision, robotics, and automatic programming were interesting, but stayed typically at the level of (un-scalable) prototypes.



Fig. 1. AI Research as Medieval Alchemy

Research in AI started with the noble objective of creating computer programs that exhibit human intelligence, promising some great achievements in a projected timeline. It has passed through several alternating cycles of brave promises and grave disappointments. Some embryonic applications in search, game playing, language understanding, expert systems, vision, robotics, and automatic programming were interesting, but stayed typically at the level of (un-scalable) prototypes.

This early period spawned the development of some key technologies, much like medieval alchemist who invented important chemical compounds and processes while trying to turn lead into gold (see Fig. 1), early AI researchers developed several key technologies while chasing their field's overly ambitious goals. That period saw the development of important conceptual models, which later served as the core of AI-inspired technology developments. Several technological domains such as knowledge representation, information extraction, semantic inferencing, machine learning, probabilistic reasoning and data mining were born out of AI.

Entering the Mature Age of AI Technologies

Paradoxically, AI has achieved success through invisible but working technologies, embedded into solutions, while disappearing from the glitzy public media and avoiding prolonged debates about the nature of intelligence. Parallel advances in IT enabled the realization of the early plans of AI communities, which were split now into several (competitive) schools of thought. It is important to mention that AI researchers also imported techniques from applied mathematics, especially from probability and statistics. In the second, mature age, toy prototypes evolved into serious technologies which evolved into real-world business systems.

Chess, go and backgammon machines able to compete against champions, global search engines capable of searching billions of web pages, embedded car and consumer device technologies, automated robotic production lines, financial screening and automated trading engines, decision support systems, space exploration voyagers, multi-player Internet games, and mobile intelligent agents are only a few notable examples of AI-inspired technologies that have attained serious deployment in consumer, business, and scientific systems. The

emerging self* (management, healing, configuration) or autonomic technologies promises to spawn yet another wave of technology advances based on the early AI ideas and concepts.

On the Re-Emerging Future of AI

At present, we seem to be right now in the phase of omnipresent needs for embedded AI systems emulating intelligent behavior(s) within business and/or consumer systems. The emergence of global, large-scale systems has brought radical technology improvements for creating, transferring and processing torrents of computer-generated data. We have already seen AI deliver on some of its early promises during its middle age (after 50 years), for example by replacing human labor with robots. Replacing human intellectual feats with machines may take another 50 years and might require a very different AI architecture -- for instance a hybrid one with paired silicon and wet chips bringing artificial and living matter together with corresponding/corroborative 'computation' and 'cognition' activities (see Fig. 2). We may yet witness the emergence of embodied intelligence realized as intelligent (omnipresent, dependable) systems.

By comparing the history of AI research with the history of research in complex systems (which, coincidentally, also started in the early 1950's) one can conclude that both fields are exploring similar avenues.

Turing and Simon are considered to be pioneers by both communities. Simply speaking, the design and

architecture of man-made artificial objects and systems are very often inspired by the biological, nature-born systems, as Leonardo da Vinci has observed early in 15th century.

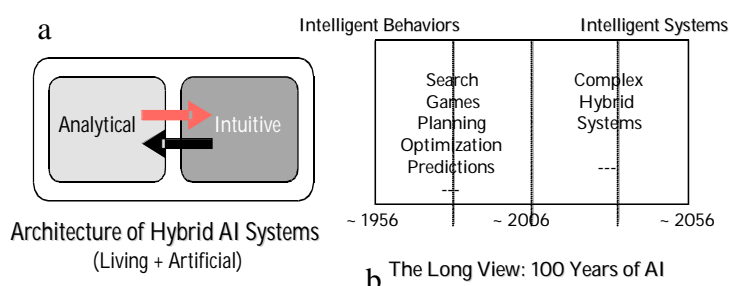


Fig. 2. a. Hybrid AI Systems – Embodied

b. The Long View of AI Research

Complexity, as a phenomenon, emerges typically from close interactions between the living and artificial worlds, resulting in non-linear behaviors. Such systems evolve, adapt and exhibit learning features. Thus, it would be natural and beneficial to join forces around the Science of Complex Systems. This will have (at minimum) the following benefits: enlarging and widening research, increasing the chances of creating a very large set of valuable technologies and walking away from ambitious but controversial term of “Artificial Intelligence”.

We are surrounded by the host of omnipresent, complex systems (cells, markets, companies, supply networks, etc) for which elucidation of closed-loop control patterned after natural, biological systems combined with knowledge representation, learning and analytic techniques may lead the creation of large-scale embodied systems. As is Intelligent Enterprise, for example, system of the high, practical value and important scientific relevance [see The Rise of The Intelligent Enterprise³]. We believe that the future of AI will be far more successful as research in Complex Systems over the next 50 years.

References

1. Turing A.M. Computing Machinery and Intelligence (MIND, 1950).
2. Simon H. The Sciences of the Artificial (MIT Press, Boston, 1981).
3. http://acm.org/ubiquity/views/k_delic_4.pdf

Anthropocentricity and the Social Robot: Artistic and Aesthetic Investigations into Machine Behaviours.

Louis-Philippe Demers
Z-Node Research Fellow
Hochschule für Gestaltung und Kunst Zürich
e-mail: lpd@processing-plant.com

This paper discusses the notion of anthropomorphism and perceived behaviours in the social robots from the biased view of several artistic robotic installations and performances. The author presents these observations as a source of inspiration for embedding behaviours in robots. Investigating anthropocentricism, these works mix machines from the very abstract geometric to the very representative zoomorphic shapes. The robot was exploited as the medium in atypical human analogies and situations. In *La Cour des Miracles*, we staged the misery of the machine. In *L'Assemblée*, 48 robotic arms gather in an arena to create crowd behaviours. In *Armageddon*, robots were Angels and God's messengers while in *Devolution*, they were part of a biological metaphor with dancers. As we attribute intent to outside agents that act upon the physical world¹, one might question the level of anthropomorphism needed in social robots² and also reflect if this projection is an inevitable reflex or not.³ Social robots have mainly embraced the humanoids with friendly behaviours as the mode of intercommunication⁴, should we further ask, if this alley channels the potential of the robot intelligence.

Weak anthropomorphism.

Kinetic art, usually mechanomorphic, feeds on continuous transformation and participation of the viewer. The movement (or perceptible change of state) of an object can be seen in part as its objective nature, while its perception can be its subjective counterpart. Consequently, a rather abstract inert shape can become fluid, organic and eventually anthropomorphic, by the sole means of contextualization and movement. In figure 1, a simple motor mounted on springs creates a rich range of chaotic movement, staging this object in a cage anthropomorphises its essence resulting with the viewers perceiving it as an untamed miserable entity in *La Cour des Miracles*. Without an immense degree of computation, the behaviour is carried out by a juxtaposition of this social mis-en-scène and the inherent complex dynamic characteristics of the structure. Equally, shapes of figure 2 were created by a set of discrete manipulators⁵ where these geometries are asked to perform to an audience. Beyond the aesthetic of the hypnotic organic movements of these machines, audiences readily address the intent. This uncanny manifestation does not push the viewer to retract from the dialogue but rather induces a fascination to understand and further interact with the object. The weak anthropomorphism is here an advantage as it frees the “sign from the signified”. It enables a multiplicity of readings from a simple starting shape: an array of cubes.



Fig. 1. Untamed machine.



Fig. 2. Organic cubes.

Anthropomorphism through acting methods for robotic characters.

To explore the acceptance of artificial behaviours we will look at the theatre and the art; both providing fictitious environments to stimulate a suspension of disbelief. Stage performers share similarities with the social robots in that they both utilize gesture, body and physical action to incarnate behaviours. Acting methods may call for psycho-physical unity where behaviour is inherently physically grounded;⁶ the walking table of figure 3 manages to navigate even under a deliberate poor gait. The behaviour is a collaboration of the unstable equilibrium of the construction and the staging. The introduction of a latent failure in the gait not only creates a poetic moment but also gives a supplementary spark of life to the object, as it is similarly proposed for social robots.⁷ Acting methods propose opposite stances be taken by actors: presence or absence. The presence calls upon the performer's experience to dwell into his/her experience to deliver the character, absence requires an abnegation of the self to produce a pure rendering of the directors' directives and scripts. The beggar of figure 3 had no experience of misery neither of being poor. Its shape was a square box (symbol of a chest) that could rock over a hinge (body language of imploring). The beggar performer lean towards absence while the table is rooted more in presence via the physicality of its shape.




Fig. 3. Characters: beggar (left), walking table (right).

Anthropocentricity and the Fake.

We could associate Baudrillard's symbolic orders⁸ with the degree of anthropomorphisation of the machine: it is the reflection of a basic reality, it masks and perverts a basic reality, it marks the absence of a basic reality and finally, it bares no relation to reality whatsoever. The first three call upon anthropomorphic incarnations of the robot while the last is pure simulacra. These artistic explorations fuel themselves at the growing blurred division between the man and the machine and demonstrate the paradox of artificial life. Stuck between the real and the artificial, the flesh and the metal, the sign and the signified, the anthropomorphisation of the robot suffers from Multiple Ontologies Disorder, a high-level manifestation of human-robot schizophrenia.⁸ Since the principal of artificial reproduction favours the human body and the human existence as construct, is anthropocentricity at the centre of this disorder?

References

1. Dennett, D. *The Intentional Stance*. (The MIT Press, Boston, 1987)
2. Duffy, B.R., "Anthropomorphism and The Social Robot", *Robot as Partner: An Exploration of Social Robots*, 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems, September 30 - October 4, EPFL, Switzerland, 2002.
3. Kennedy, J.S. *The New Anthropomorphism* (Cambridge Press, 1992)
4. Fong, T. et al. *A Survey of socially interactive robots* (Robotics and Autonomous Systems 42, 2003, pp. 143-166)
5. Suthakorn, J. & Chirikjian, G.S., *Design and Implementation of a New Discretely-Actuated Manipulator*, Proc. of ISER 2000, Hawaii, December 2000.
6. Hoffmann, G. *HRI: Four Lessons from Acting Method* (MIT Media Lab)
7. B.R. Duffy, G. Joue, "The Paradox of Social Robotics: A Discussion", AAAI Fall 2005 Symposium on Machine Ethics, November 3-6, 2005, Hyatt Regency Crystal City, Arlington, Virginia
8. Baudrillard, Jean. *SIMULACRA AND SIMULATION*. Tr. Sheila Faria Glaser. Ann Arbor: University of Michigan. 1994. Originally published in French by Editions Galilee, 1981. 164 pages.
9. Demers, L.P. & Vorn, B. *Schizoid Ontologies of Cybernetic Lures* to appear in "Flesh Eating Technologies", edited by Sara Diamond, Sylvère Lotringer, M.A. Moser, The Banff Centre, Feb. 1998.

The Humanoid Robot Project Holistic Movement

Raja Dravid^{1 2}, Gabriel Gomez¹, and Peter Eggenberger Hotz¹

¹Artificial Intelligence Laboratory

University of Zurich, Andreasstrasse 15, Zurich, CH-8050

²ddd forschen + entwickeln, Hardturmstrasse 68, Zurich, CH-8005

[dravid, gomez, eggen]@ifi.unizh.ch

One of the key aspects of understanding human intelligence is to investigate how humans interact with their environment. Performing articulated movement and manipulation tasks, (e.g., walking, hitting a flying object, using tools made for humans) in a constantly changing environment, have proved more difficult than expected, although many impressive humanoid robots have been built to date. We have thus sought to re-investigate the underlying concepts of natural movement in accordance to human morphology.




Fig. 1. Prototyping the thorax and insertion points for the tendons.
(a) Fiberglass thorax. (b) Insertion points on the thorax.

Humans are composed of mostly soft, compliant materials. Muscles can act as springs, supportive structures, shock absorbers and as actuators, depending on the situation. This has consequences to how we move. Although we are capable of isolating movement to single joints, natural movement is generally distributed in different parts of the body. In this sense the body can be looked upon as a chain with different links all interacting with each other (see Dravid et al., 2002a and b). The impetus of motion can be generated in one part of the body, say the hips and travel through the shoulder complex, resulting in a swinging motion of the arm, not unlike that of a whip. We coin this form of movement “holistic movement”.

In order to investigate holistic movement we propose to first concentrate on the shoulder and arm complex. To this end we plan to build a tendon driven robotic shoulder and arm, resembling human morphology (Fig. 1). This design plan will be followed down to the ulna and radius configuration of the lower arm. In previous work we have investigated the mechanical properties of such arrangements (Fig. 2).

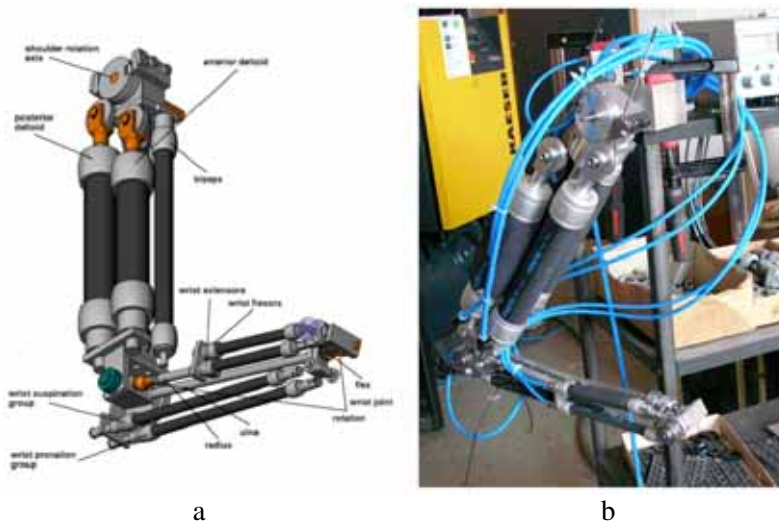


Fig. 2. Pneumatic arm. (a) CAD design. (b) First prototype.

Unlike conventional approaches we will construct the shoulder and arm complex around a fiberglass thorax serving as a sliding surface for the shoulder blade (Fig. 2a). The humerus, ulna and radius will be made out of plastic skeletal bones wrapped in fiberglass. A surgical shoulder and elbow joint will then be grafted to the bones. Ligaments made out of dyneema, a new fiber composite, will serve as connective tissue (i.e., tendons) between the anthropomorphic insertion points lying on the thorax and the moving bones (Fig. 2b). We will simulate the 9 prime movers of the shoulder complex and 10 muscles of the arm with new generation pneumatic actuators made by Festo Inc. Feedback is provided through pressure, force and positional information.

A previously built robotic head with stereo color vision and an anthropomorphic tendon driven hand will also be available (See Fig. 3 and Gomez et al., 2006). Results may then be used in robotics, tele-operation and rehabilitation applications.




Fig. 3. Existing robotic components. (a) Robotic head with stereo vision. (b) Tendon driven robotic hand.

References

- Dravid, R., Iida, F. and Paul, C. (2002a). *Control of lateral bounding for a pendulum driven hopping robot*, Proc. of 5th International Conference on Climbing and Waling Robots (CLAWAR 2002), pp. 333-340.
- Dravid, R., Iida, F., and Paul, C. (2002b). *Design and control of a pendulum driven hopping robot*, Proceedings of International Conference on Intelligent Robots and Systems 2002 (IROS 02), pp. 2141-2146
- Gomez, G., Hernandez, A. and Eggenberger Hotz, P. (2006). *An adaptive neural controller for a tendon driven robotic hand*. In Proceedings of the 9th International Conference on Intelligent Autonomous Systems (IAS-9). T. Arai et al. (Eds.), IOS Press, Tokyo, Japan, pp. 298-307.

AI approaches for next generation telecommunication networks

Gianni A. Di Caro, Frederick Ducatelle and Luca M. Gambardella
Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA)
USI/SUPSI, Lugano, Switzerland
[frederick | gianni | luca]@idsia.ch

Computer networks are an important example of distributed dynamic systems which are omnipresent in our daily life. The strategic importance and intrinsic constraints of such systems imply the need for distributed control, especially at the routing level, to make the network behavior adaptive to changes in topology, data traffic, services, etc.. Therefore, the control and machine learning communities have always been interested in the field of computer communications. Already in the 1950's, Bellman and Ford applied dynamic programming to the problem of routing optimization in networks.^{1,2} While the Bellman-Ford routing algorithm implements distributed control, it offers limited adaptivity, so that its performance degrades considerably in rapidly changing scenarios. Over the years, people have explored more adaptive versions of the original Bellman-Ford algorithm, encountering, however, several difficulties. More recently, researchers have investigated new routing algorithms which provide better adaptivity, building on advances in machine learning. In 1987, Nedzelnitsky and Narendra developed an approach based on stochastic learning automata.³ In 1994, Boyan and Littman proposed Q-routing,⁴ an adaptation of the Bellman-Ford algorithm which uses ideas from the Q-learning algorithm developed in the context of reinforcement learning. In 1998, Di Caro and Dorigo proposed AntNet,⁵ which was derived from the Ant Colony Optimization (ACO) metaheuristic,⁶ and implements a distributed Monte Carlo sampling system to learn routing decisions.

Despite the very good performance shown by these and many other adaptive routing algorithms, current network technology still relies on mainly static algorithms. Internet routing algorithms such as RIP and BGP are derived from the basic Bellman-Ford algorithm. They have capabilities to deal with infrequent topological changes (e.g. caused by network failures), but do not provide traffic adaptivity. The main strategy to deal with traffic fluctuations and provide guaranteed quality of service is overprovisioning of network resources, making fully adaptive routing algorithms unnecessary in practice.

However, this status-quo is now changing rapidly. Advances in wireless technology, such as Bluetooth and WiFi, allow for more freedom when setting up or changing a data network (e.g., users can move freely across the network and create local ad hoc multi-hop networks), while the introduction of new communication models and user services, such as peer-to-peer networking and voice-over-IP, leads to new and changing demands in terms of data traffic. Networks are becoming increasingly dynamic and heterogeneous. And since these new networks are more user-centered, their characteristics are determined by the users, rather than by a central authority, such that overprovisioning is no longer an effective option. This evolution is expected to accelerate, increasing the need for new, dynamic control algorithms. These algorithms should learn about the current network and user context, adapt decision policies to it, and even self-tune internal parameters. This is the approach advocated in the view of *autonomic communications*.⁷

While the arrival of new generation networks and autonomic communications renews the case for adaptive routing, it poses a challenge which goes beyond what

earlier developed learning routing algorithms can deal with. The mobility of network nodes and the changes in data traffic patterns due to the appearance of new services leads to different network modes, defined by characteristics such as bandwidth, connectivity, etc.. The network mode can evolve over time, or different network modes can coexist in the same heterogeneous network. Moreover, constraints imposed by the network technologies add further complexity. We believe one approach to deal with these challenges is to integrate several learning and behavior paradigms to create a fully adaptive, multi-modal controller. We followed this approach in the area of mobile ad hoc networks (MANETs), which are networks consisting of wireless, mobile hosts communicating in multi-hop fashion without the support of any infrastructure. MANETs are a paradigmatic example of the dynamic new generation networks. The algorithm we proposed, AntHocNet,⁸ is an ACO inspired routing algorithm. However, it contains several elements not present in other ACO routing algorithms such as the earlier mentioned AntNet. Specifically, it combines ACO based Monte Carlo sampling and learning with an information bootstrapping process, which is typical for dynamic programming and some reinforcement learning approaches. Operating the two learning mechanisms at different speeds allows to obtain an adaptivity, robustness and efficiency which neither of the subcomponents could offer alone. Moreover, the balance between the use of proactive and reactive behaviors allows both to anticipate and to respond in timely fashion to disruptive events. AntHocNet's innovative design, which sets it apart from other MANET routing algorithms, has been shown to give superior performance over a wide range of simulation scenarios with different characteristics in terms of mobility, data traffic, etc..

We believe that the good performance of AntHocNet for MANETs is an indication that such an integrated approach can be the way to go to provide adaptivity in dynamic multi-modal networks. However, a more fundamental challenge is to bring these challenging environments back to machine learning research, where it offers an opportunity to study difficult but real distributed dynamic environments and to support the implementation of the new algorithms needed to drive the progress of network development.

References

1. Bellman, R. On a routing problem. *Quarterly of Applied Mathematics*, 16(1) (1958).
2. Ford, L. R. and Fulkerson, D. R., *Flows in networks*, Princeton Univ. Press (1962).
3. Nedzelnitsky, O. V. and Narendra, K. S. Nonstationary models of learning automata routing in data communication networks. *IEEE Trans. on Systems, Man, and Cybernetics*, SMC-17 (1987).
4. Boyan, J. A. and Littman M. L., Packet routing in dynamically changing networks: A reinforcement learning approach. *NIPS* 6 (1994).
5. Di Caro, G. A. and Dorigo, M., AntNet: Distributed stigmergic control for communications networks. *J. of Artificial Intelligence Research (JAIR)*. 9 (1998).
6. Dorigo, M., Di Caro, G. and Gambardella L. M., Ant algorithms for discrete optimization, *Artificial Life*, 5(2) (1999).
7. Kephart, J., and Chess, D., The vision of autonomic computing, *IEEE Computer magazine*. 36(1) (2003).
8. Ducatelle F., Di Caro G. A., and Gambardella L. M., Using ant agents to combine reactive and proactive strategies for routing in mobile ad hoc networks, *International Journal of Computational Intelligence and Applications*, 5(2) (2005).

Evolutionary humanoid robotics: first steps towards autonomy?

Malachy Eaton

Department of Computer Science and Information Systems

University of Limerick, Ireland

malachy.eaton@ul.ie

Introduction. Evolutionary humanoid robotics is a branch of evolutionary robotics¹ dealing with the application of evolutionary principles to the design of humanoid robots. In the experiments outlined here we use a genetic algorithm to choose the joint values for a simulated humanoid robot with a total of 20 degrees of freedom (elbows, ankles, knees, etc.) for specific time intervals (keyframes) together with maximum joint ranges, in order to evolve bipedal locomotion. An existing interpolation function fills in the values between keyframes; once a cycle of 4 keyframes is completed it repeats until the end of the run, or until the robot falls over. The humanoid robot is simulated using the Webots mobile robot simulation package² and is broadly modeled on the Sony QRIO humanoid robot³ (see Fig. 1). In order to get the robot to walk a simple function based on the product of the length of time the robot remains standing by the total distance traveled by the robot was devised. This was later modified to reward walking in a forward (rather than backward) direction and to promote walking in a more upright position, by taking the robots final height into account⁴.

Joint range restriction. In previous experiments it was found that the range of movement allowed to the joints by the evolutionary algorithm, that is the proportion of the maximum range of movement allowed to the robot for each joint was an important factor in evolving successful walks⁴.

Initial experiments placed no restriction on the range of movement allowed and walks did not evolve unless the robot was restricted to a stooped posture and a symmetrical gait; even then results were not impressive. By restricting possible movement to different fractions of the maximum range walks did evolve, however as this was seen as a critical factor in the evolutionary process it was decided in the current work to include a value specifying the fraction of the total range allowed in the humanoid robots genome.

Experimental details. The genome length is 328 bits comprising 4 bits determining the position of the 20 motors for each of 4 keyframes; 80 strings are used per generation. 8 bits define the fraction of the maximum movement range allowed. The maximum range allowed for a particular genome is the value specified in the field corresponding to each motor divided by the number of bits set in this 8 bit field plus 1. 8 bits was chosen as reasonable walking patterns were seen to evolve when the range was restricted by a factor of 4 or thereabouts in previous experiments. The genetic algorithm uses roulette wheel selection with elitism; the top string being guaranteed safe passage to the next generation, together with standard crossover and mutation. Maximum fitness values may rise as well as fall because of the realistic nature of the Webots simulation. Two-point crossover is applied with a probability of 0.5 and the probability of a bit being mutated is 0.04. These values were arrived at after some experimentation.



Fig. 1 The humanoid robot walking

Experimental results. We ran three trials of the evolutionary algorithm on a population size of 80 controllers for 700 generations of simulated robots, taking approximately 2.5 weeks simulation time on a reasonably fast computer, corresponding to approximately 7.5 weeks of “real time” experimentation. A fitness value over about 100 corresponds to the robot at least staying standing for some period of time; over 500 corresponds to a walk of some description. The results obtained were interesting; walks developed in all three runs, on average after about 30 generations, with fine walking gaits after about 300 generations. This is about half the time on average that walking developed with a fixed joint range. We can see from Fig. 2 that the joint range associated with the individual with maximum fitness fluctuates in early generations; typically low values (high movement ranges) initially predominate as the robot moves in a “thrashing” fashion. Then the movement range becomes restricted for the highest performing individuals, as a smaller range of movement increases the likelihood that the robot will at least remain standing for a period, while hopefully moving a little. Then in later generations typically the movement range gradually becomes relaxed again, as a greater range of movement facilitates more rapid walking, once the robot has “learnt” how to remain upright.




Fig 2. Fitness and joint range, averaged over 3 runs

Discussion. In this work we have demonstrated one of the first applications of evolutionary algorithms to the development of complex movement patterns in a many-degree-of-freedom humanoid robot. Perhaps the time has now arrived for a more serious and detailed discussion on the possible ethical ramifications of the evolution of human-like robots. Such robots may be able to take our place in the workforce, or in other fields, and there may well also be other significant

social consequences. Other, more technical, issues arise – while Asimov’s three laws of robotics may appear a little dated, it could be important to avoid the appearance in the home or workplace of unexpected evolved “side effects” which may escape a rigorous testing regime. For example one of the walking behaviours evolved in our work involved the robot walking (staggering) in a manner amusingly reminiscent of an intoxicated person. While this gait proved surprisingly effective, not many people would relish the prospect of being shown around a new house for sale by a seemingly drunken robot! In conclusion, if indeed we are now beginning to see the first tentative “steps” towards autonomous humanoid robots, perhaps this is an appropriate forum to now look forward to the harnessing of this technology for the benefit of mankind.

References

1. Nolfi S, Floreano D. “Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines”. (MIT press, 2000).
2. Michel O. “Webots: Professional Mobile Robot Simulation”, International Journal of Advanced Robotic Systems, Vol. 1, No. 1, pages 39-42.
3. Mojon S (2003) “Realization of a Physic Simulation for a Biped Robot” Semester project, School of Computer and Communication Sciences EPFL, Switzerland 2003.
4. Eaton M, Davitt T.J. “Automatic evolution of bipedal locomotion in a simulated humanoid robot with many degrees of freedom” in *Proceedings of the 11th International Symposium on Artificial Life and Robotics*, Jan 23-25, 2006, Japan, pp. 448-451. [Best paper award]

A Robotic Model for Rat Tactile Sensing

Miriam Fend, Rolf Pfeifer
Department of Informatics
University of Zurich, 8050 Zurich, Switzerland
[fend | pfeifer]@ifi.unizh.ch

The rat whisker system has served as model for biological research on tactile sensing for many years. Its advantages are manifold: rats are almost as sensitive with their whiskers as we are with our fingertips, but stimuli can be applied in a very controlled manner by moving only one or several whiskers. Furthermore, the morphology of the whiskerpad is preserved in the brain stem, thalamus and the primary sensory cortex, where it is called the barrel cortex.




Fig. 1. a) Amouse robot (b) For locomotion behaviors, the relation of the sensory vs. the physical space of the agent has to match.

Recently, the robotics community started to develop artificial whiskers inspired by the biological example. The synthetic approach [1] of building an artefact can enhance our understanding of biological systems because parameters such as morphology and material can be changed that are not accessible in the biological agent.

Morphology of Artificial Whiskers

We have equipped a robot with active artificial whiskers and an omnidirectional camera (figure 1a) to study the relation between morphology and different tasks as well as the tactile capacity of the sensor for texture discrimination. Whiskers extend the sensing range (sensory space in figure 1b) beyond the rigid body of the robot (physical space in figure 1b). The relation of physical and sensory space strongly influences the performance of the robot, and it is dependent on the specific task. We have shown that the natural morphology resembles more closely to a morphology well suited for wall following but not for obstacle avoidance. This suggests that whiskers on natural agents are more optimized to tasks such as wall following [2].

Tactile Texture Discrimination

In a second line of research, we have investigated the potential of our artificial whiskers for texture discrimination. While tactile categorization of surfaces is challenging, especially with a mobile robot, we have shown that it is possible to discriminate different textures. Furthermore, the discriminatory capacity is enhanced by sensorimotor feedback.

While whiskers are novel sensor for robots, we also show how research on the robot and biological investigations can complement each other.

References

1. Pfeifer, R. & Scheier, C. *Understanding Intelligence* (MIT Press, Boston, 1999).
2. Fend, M., Bovet, S. & Pfeifer, R. *On the influence of morphology of tactile sensors for behavior and control*, Robotics and Autonomous Systems, in press

Applying Data Fusion in a Rational Decision Making Architecture of a Believable Agent

Benjamin Fonooni*, Behzad Moshiri[†], Caro Lucas[†]

*Departments of Engineering

Islamic Azad University Sciences and Research Branch, Tehran, Iran

bfonooni@hotmail.com

[†]Control and Intelligent Processing, Center of Excellence, School of Electrical and Computer Engineering

University of Tehran, Tehran, Iran

moshiri@ut.ac.ir, lucas@ipm.ir

1. Introduction

This research project is heading for creating a believable agent which able to decide rationally and use emotions to regulate decisions that has been made. As we look inside the general architecture¹, it is obvious that agents which are equipped with emotions will deliberate processes that assumed to be their best choices. The rationality means acting appropriately in the various situations. However, it enables applications to have more believable interactions between man and machine which is the most important consideration of these agents. Combination of emotions, rationality and personality will yield to believable agents. The fundamental objective of this research work is the improvement of the decision making algorithm using Data Fusion mainly in the domains applied to humans in which both rationality and emotions have effective roles in decision making process.

2. Agent's Architecture

As Figure 1 depicts, architecture has three main components. Input

component responsible for gathering environmental information and delegating processed results to other components. Rational component has the main function of keeping and updating the rational state. Emotional component is responsible for updating emotional state of an agent according to the stimuli created by rational component and feeds action selector to choose one action among all which are rationally equivalent. The heart of the rational component architecture constitutes from a production system with forward chaining approach. Usually, the system contains a series of rules to make conditions-actions and uses the sensory input of the agent to percept conditions and verifies in relations of the rules, then chooses which actions to use towards the outside world or of other agents. Goal management system has a role of managing and selecting goals of an agent, perhaps, the goal selection is the most important process for any goal based agent. Making a proper decision

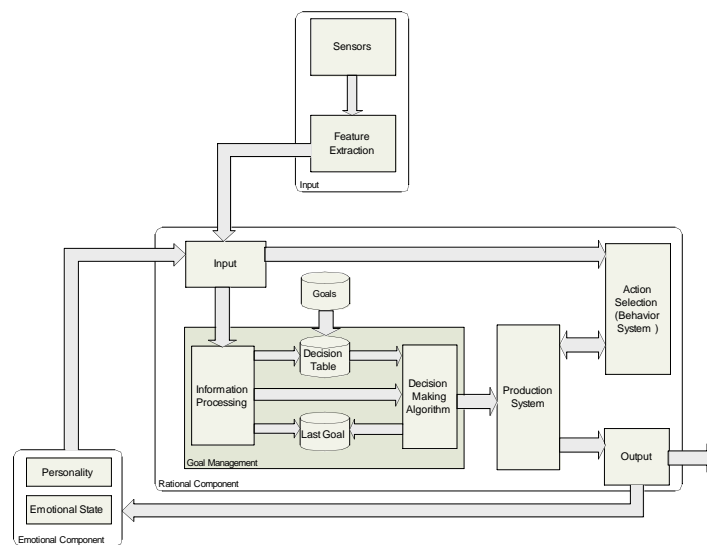


Fig. 1. Architecture for Rational-Emotional Agent

according to the agent's rational state which is obtained from production system and personality needs an algorithm which has better results compared to previous ones². In the following section we will explain the role of Data Fusion and aggregated operators in decision making process. The role of emotional component in action selection mechanism has not been covered in this article.

3. Applying Data Fusion in multi-attribute decision making algorithm

Our algorithm needs to specify a set of attributes for each goal which has to be assigned by a user. Set of attributes includes priority, importance and easiness which are considered as collection of aggregated objects in the unit interval, a_i , which has to be ordered and stored in decision table. Agent's personality is used as coefficient which determines its degree of pessimism. In case of using OWA operator^{3,5} in our decision making algorithm, a weighting vector W should be defined and initialized. The main question would be obtaining the weights associated with OWA, because it models process of aggregation used on data set. We used a back propagation method to learn from agent's observations⁴. Suggested algorithm is described below:

1) Each aggregated value will be calculated by classic Hurwicz's multi-attribute method and is considered as our desired value:

$$\rho \text{Max}_i a_i + (1 - \rho) \text{Min}_i a_i = d \quad \rho : \text{Agent's personality}$$

2) Following learning algorithm should be applied to estimate the corresponding weights:

$$\lambda_i(l+1) = \lambda_i(l) - \beta w_i (b_{ki} - \hat{d}_k)(\hat{d}_k - d_k) \quad \beta : \text{Learning rate}$$

$$w_i = \frac{e^{\lambda_i(l)}}{\sum_{j=1}^n e^{\lambda_j(l)}}$$

$$\hat{d}_k = b_{k1} w_1 + b_{k2} w_2 + \dots + b_{kn} w_n \quad \hat{d}_k : \text{Current estimation of } d_k$$

Parameters λ_i determine the OWA weights and are updated with back propagation of the error $(\hat{d}_k - d_k)$.

Finally, after 100 iterations, the best \hat{d}_k with maximum value will be selected as current agent's goal and will be delegated to the production system. Also priority of each goal will be decreased and checked out with a threshold so that it would fade out after a while.

References

1. Camurri, A. & Caglio, A. *An Architecture for Emotional Agents* (IEEE Multimedia Journal, October-December 1998).
2. Camurri, A. & Volpe G. *A Goal Directed Rational Component for Emotional Agent* (Systems, Man, and Cybernetics, IEEE SMC '99 Conference Proceedings, 1999).
3. Fullér, R. & Majlender, P. *On obtaining minimal variability OWA operator weights* (Fuzzy Sets and Systems 136(2): 203-215, 2003)
4. Filev, D. P. & Yager, R. R. *Learning OWA Operator Weights From Data* (Proceedings of the Third IEEE International Conference on Fuzzy Systems, Orlando, 468-473, 1994)
5. Beliakov, G. *How to Build Aggregation Operators from Data* (INTERNATIONAL JOURNAL OF INTELLIGENT SYSTEMS, VOL. 18, 903-923, 2003)

Beyond Information Transfer: The Emergence of Embodied Communication

Tom Froese
C.C.N.R., Departments of Informatics
University of Sussex, Brighton, UK
t.froese@sussex.ac.uk

Introduction

The “Communication Is Information Transfer” metaphor is one among the most deeply entrenched constitutive assumptions of traditional AI research. However, there is a growing consensus in the cognitive sciences that the functional definition of communication as an exchange of information with selective advantage is inadequate because it only pertains to the descriptive domain of the observer. What is needed instead is a shift of focus to the underlying dynamical mechanisms and morphological structures which causally enable communication to emerge as a coordination of behavior between embodied, situated, and autonomous systems. In what follows I will first briefly outline the progress that has been made in artificially evolving behavioral coordination in multi-agent systems, and then describe in a few words why these developments are particularly relevant for the future of embodied cognitive science as a whole.

Traditional AI: communication as information transfer

Almost since the beginning of AI there have been attempts at using synthetic models to gain insights into the origins and nature of communication. Within this broad framework the two main paradigms of cognitive science can be clearly distinguished: the computational approach characterizes communication in functional terms as a transfer of information between sender and receiver over some kind of medium, while the embodied approach defines communication in operational terms as a form of behavioral coordination taking place between two or more structure-determined dynamical systems¹. Recently, the computational approach has come under criticism; for while there are examples of communicative behavior between animals which can be said to describe a certain state of affairs, these descriptions require the existence of previous consensual agreement which can only be achieved on the grounds of pre-existing communicative abilities². This shortcoming of the computational approach is reflected in simulations in which a communication channel is explicitly made available to the agents as part of the experimental setup^{3, 4}. In this manner the problems regarding the biological origins of communication are ignored: communicative behavior does not first have to originate from non-communicative behavior but merely has to be fine-tuned by artificial evolution.

Embodied AI: communication as behavioral coordination

In contrast, the use of evolutionary robotics methodology⁵ in combination with a dynamical systems approach to cognition⁶ provides a promising framework for investigating the origin of biologically grounded communicative behavior via synthetic means. Researchers have successfully used this kind of approach to evolve embodied behavioral coordination without dedicated communication channels. For example, movement coordination through role allocation has been evolved between two simulated autonomous agents equipped only with wheels and proximity sensors⁷, and a similar task has been successfully implemented with actual robots⁸. These

developments are important for the new paradigm because they present us with concrete examples where non-trivial, situated, and embodied communication between autonomous agents has evolved from non-communicative behavior. Embodied AI's focus on the dynamics of behavior, as well as on the morphology of the organism and its environment, permit an operational analysis of how these three theoretical entities are related, and how they enable non-trivial communicative behavior to emerge.

Final remarks

Despite of all the talk of an immanent paradigm shift in the cognitive sciences, what we actually find is an empirical and philosophical stalemate. Indeed, there is a real danger that this will remain so in the foreseeable future. While embodied AI is increasingly successful in explaining the foundations of immediate action-in-the-world, traditional AI nevertheless remains more successful in accounting for theoretical and abstract cognition. However, there is a strong possibility that artificial life research into communication could enable embodied cognitive science to finally move beyond the computational paradigm. It presents us with one example of how to escape the current restriction to low-level sensorimotor accounts of cognitive behavior while still retaining the same methodological commitments. This is because an embodied form of communication allows an agent to successfully interact in the higher-level cognitive domain of social phenomena by making use of the same primary cognitive capabilities which have been evolved for intelligent action-in-the-world. Finally, in this way embodied AI could potentially provide us with a new path towards the holy grail of AI: human like intelligence.

References

1. Maturana, H.R., & Varela, F.J. *The Tree of Knowledge: The Biological Roots of Human Understanding* (Shambhala Publications, Boston, 1987).
2. Di Paolo, E.A. "An investigation into the evolution of communication" (*Adaptive Behavior*, **6**(2), pp. 285-324, 1997).
3. Werner, G.M. & Dyer, M.G. "Evolution of Communication in Artificial Organisms" (*Artificial Life II*, Langton, C.G., et al. (Eds.), Addison-Wesley, pp. 659-687, 1991).
4. MacLennan, B.J., & Burghardt, G.M. "Synthetic ethology and the evolution of cooperative communication" (*Adaptive Behavior*, **2**(2), pp. 161-187, 1993).
5. Harvey, I., Di Paolo, E., Wood, R., Quinn, M. and Tuci, E. A. "Evolutionary Robotics: A new scientific tool for studying cognition" (*Artificial Life*, **11**(1-2), pp. 79-98, 2005).
6. Van Gelder, T. "The dynamical hypothesis in cognitive science" (*Behavioral and Brain Sciences*, **21**, pp. 615-665, 1998).
7. Quinn, M. "Evolving communication without dedicated communication channels" (*Proc. of the 6th Euro. Conf. on Artificial Life*, Kelemen, J. and Sosik, P. (eds.), Springer, pp. 357-366, 2001).
8. Quinn, M., Smith, L., Mayley, G., & Husbands, P. "Evolving Formation Movement for a Homogeneous Multi-Robot System: Teamwork and Role-Allocation with Real Robots" (*Cognitive Science Research Paper*, **515**, University of Sussex, 2002).

Dynamic Meta-Learning

Matteo Gagliolo

IDSIA, Galleria 2, 6928 Manno (Lugano), Switzerland

University of Lugano, Faculty of Informatics,

Via Buffi 13, 6904 Lugano, Switzerland

matteo@idsia.ch

Most solvable AI problems can be addressed by more than one algorithm; most AI algorithms feature a number of parameters that have to be set. Both choices can dramatically affect the quality of the obtained solution, and the time spent obtaining it.

Algorithm Selection, or *Meta-Learning*, techniques [1,2] typically address these questions by solving a large number of problems with each of the available algorithms, in order to learn a mapping from $(problem, algorithm)$ pairs to expected performance. The obtained mapping is later used to select and run, for each new problem, only the algorithm that is expected to give the best results. This approach, though being preferable to the far more popular "trial and error", poses a number of problems. It presumes that such a mapping can be learned at all, i.e., that the actual algorithm performance on a given problem will be predictable with enough precision before even starting the algorithm - often not the case with stochastic algorithms, whose performance can exhibit large fluctuations among different runs (see, e.g., [3]). It also assumes problem instances met during the training phase to be statistically representative of successive ones. For these reasons, there usually is no way to detect a relevant discrepancy between expected and actual performance of the chosen algorithm. Finally, it neglects computational complexity issues: ranking between algorithms is often based solely on the expected *quality* of the performance; and the time spent during the training phase is not even considered, although it can be large enough to cancel any practical advantage of algorithm selection.

The *Algorithm Portfolio* paradigm [4,5] consists in selecting a *subset* of the available algorithms, to be run in parallel, with the same priority, until the fastest one solves the problem. This simple scheme is more robust, as it's less likely that performance estimates will be wrong for all selected algorithms, but it also involves an additional overhead, due to the "brute force" parallel execution of all candidate solvers.

In our view, a crucial weakness of these approaches is that they don't exploit any feedback from the actual execution of the chosen algorithms. We tried to move a step in this direction, introducing *Dynamic Algorithm Portfolios* [6]. Instead of *first* choosing a portfolio *then* running it, we iteratively allocate a time slice, sharing it among all the available algorithms, and update the relative priorities of the algorithms, based on their current state, in order to favor the most promising ones. Instead of basing the priority attribution on performance quality, we fix a target performance, and minimize the time to reach it. To this aim, we search for a mapping from $(problem, algorithm, current\ algorithm\ state)$ triples to *expected time* to reach the desired performance quality. The mapping is obtained training a parametric model of algorithm runtime distribution. To further reduce computational complexity, we focus on *lifelong-learning* techniques that drop the artificial boundary between training and usage, exploiting the mapping during training, and including training time in performance evaluation.

The obtained selection technique is generic, not depending on algorithm-specific properties. We present experiments with genetic algorithms and satisfiability problem solvers.

The target of our work is to obtain a fully dynamic meta-learning agent that learns to use a set of algorithms by solving a set of problems, with minimal a-priori knowledge, and minimal performance overhead, allowed by a continuous cycle of runtime feedback and re-allocation of computational resources.

References

1. J. R. Rice, The algorithm selection problem, in *Advances in computers*, vol. 15 (Academic Press, NY, 1976).
2. R. Vilalta, Y. Drissi, A perspective view and survey of meta-learning, *Artif. Intell. Rev.*, vol. 18, no. 2, pp. 77-95, 2002.
3. C.P. Gomes, B. Selman, N. Crato, H. Kautz, Heavy-tailed phenomena in satisfiability and constraint satisfaction problems, *J. Autom. Reason.*, vol. 24, no. 1-2, pp. 67-100, 2000.
4. B. A. Huberman, R. M. Lukose, and T. Hogg, An economic approach to hard computational problems. *Science*, vol. 275, pp. 51--54, 1997.
5. C. P. Gomes and B. Selman, Algorithm portfolios, *Artificial Intelligence*, vol. 126, no. 1-2, pp. 43--62, 2001.
6. M. Gagliolo and J. Schmidhuber, Dynamic algorithm portfolios, in *Ninth International Symposium on Artificial Intelligence and Mathematics*, Fort Lauderdale, Florida, January 2006.

Ant Colony Optimization for Advanced Logistics Systems

Luca Maria Gambardella, luca@idsia.ch, www.idsia.ch/luca
IDSIA, Istituto Dalle Molle di Studi sull'Intelligenza Artificiale, Galleria 2, 6928 Manno-Lugano,
Switzerland

A traditional business model is articulated in three stages: production, distribution, and sales. Each one of these activities is usually managed by a different company, or by a different branch of the same company. Research has been trying to integrate these activities since the 60s when multi-echelon inventory systems were first investigated (Clark and Scarf, 1960), but, in the late 70s, the discipline which is now widely known as Supply Chain Management was not delivering what was expected, since the integration of data and management procedures was too hard to achieve, given the lack of real integration between the Enterprise Resource Planning (ERP) and the Enterprise Data Processing (EDP) systems (Sodhi, 2001). Only in the early 1990s did ERP vendors start to deploy products able to exploit the pervasive expansion of EDP systems at all levels of the supply chain. The moment was ripe for a new breed of companies to put data to work and start to implement and commercialise Advanced Logistics Systems (ALS), whose aim is to optimise the supply chain seen as a unique process from the start to the end.

The first ALSs were the preserve of big companies, who could afford the investment in research and development required to study their case and to customise the application to interact with the existing EDP systems.

While ALSs were first deployed, researchers in the field of Artificial Intelligence were first investigating new “meta-heuristics”, heuristic methods that can be applied to a wide class of problems, such as Ant Colony Optimisation – ACO (Dorigo *et al.* 1999, Bonabeau *et al.* 2000), Tabu Search (Glover and Laguna, 1997), Iterated Local Search (Stützle and Hoos, 1999), Simulated Annealing (Kirkpatrick *et al.*, 1984). The integration of optimisation algorithms based on innovative meta-heuristics with ALSs for Supply Chain Management opens new perspectives of AI applications in industry. Not only big companies can afford ALSs, but also small and medium enterprises can use state-of-the-art algorithms, which run quickly enough to be adopted for online decision making.

In the following we present three industrial applications where Ant Colony Optimization meta-heuristic has been used to solve different size logistic problems, from a small distribution problems of few vehicles to a large problem that involve more than thousands vehicles. Algorithms based on ACO are multi-agent systems that exploit artificial stigmergy for the solution of combinatorial optimization problems: they draw their inspiration from the behaviour of real ants, which always find the shortest path between their nest and a food source, thanks to local message exchange via the deposition of pheromone trails. The remarkable advantage of ACO based algorithm over traditional optimisation algorithms is the ability to produce a good suboptimal solution in a very short time. Moreover, for some problem instances, ACO algorithms enhanced with local optimisation capabilities, have been proven to be the best overall (Gambardella *et al.*, 1999).

First we present an ALS for planning the sales and distribution process of fuel oil, which includes an advanced ACO-based algorithm for the optimisation of the Vehicle Routing Problem. The system has been adopted by a small size company in Switzerland with the goal to optimize fuel distribution in an area where a lot of delivery constraints are present. In particular there are accessibility restrictions since the distribution it is often executed in urban environments. The system has two optimisation modules: off-line and on-line. Particular care has been given to the integration of these advanced algorithmic modules within the operational company information system. The system displays an interface that lets the human planner manually enter vehicle routes and, on-demand, these routes can be automatically optimised by the algorithm, which returns its results in a few minutes. The planner can then decide whether the computer generated results can be accepted as they are, or they require further refinements and adaptations. This kind of interactive usage has proven very successful in getting the technology adopted by planners in small-medium enterprises, where the obstacles to technology adoption are usually higher.

In this second application the client is one of the major supermarket chains in Switzerland. The problem is to distribute palletized goods to more than 600 stores, all over Switzerland. The stores are the customers of the vehicle routing problem, since they order daily quantities of goods to replenish their local stocks. The stores want the goods to be delivered within time windows, in order to plan in advance the daily availability of their personnel, allocating a fraction of their time to inventory management tasks. The supermarket chain has recently reorganized its logistic process, since it concentrated the distribution process from seven inventories, distributed in various locations, to a central inventory. There are three types of vehicles: trucks (capacity: 17 pallets), trucks with trailers (35

pallets), and tractor unit with semi-trailer (33 pallets). According to the store location, only some vehicles can access it. In some cases the truck with trailer can leave the trailer at a previous store and then continue alone to other less accessible locations. All the routes must be performed in one day, and the client imposes an extra constraint stating that a vehicle must perform its latest delivery as far as possible from the inventory, since it could be used to perform extra services on its way back. The objectives are cost minimization (cost per km) and tour minimization (to limit the number of vehicles). The objective function evaluates each solution (a set of tours) according to the following scheme: (1) first minimize the number of tours, (2) then minimize the costs per kilometer and the cost of violating the soft time windows. In this application the orders are known with sufficient advance to be able to run an off-line optimisation. This problem solved with AntRoute, a modified version of the MACS-VRPTW ant algorithm (Gambardella *et al.* 1999). The algorithm was adapted to the problem in order to handle the choice of the vehicle type, thus, at the start of each tour the ant agent chooses a vehicle. Two ant colonies were used, one minimizing the number of vehicles, and the other one the length of the tours. A waiting cost was introduced in order to prevent vehicles arriving too early at the stores. Local search moves allow to improve the quality of the solutions, exchanging stores between routes or reversing the visit order. AntRoute is able to compute a solution to the problem in less than five minutes in contrast with four hours of the manual planners. The solution outperforms the human planners creating tours which are shorter, use less vehicles, and are less expensive.

In this third application the client is a major logistic operator in Italy. The distribution process involves moving palletized goods from factories to inventory stores, before they are finally distributed to shops. A customer in this vehicle routing problem is either a pick-up or a delivery point. There is no central depot, and approximately a fleet of 1'200 trucks is used. Routes can be performed within the same day, over two days, or over three days, since the Italian peninsula is quite long and there's a strict constraint on the maximum number of hours per day that a driver can travel. All pick-ups of a tour must take place before deliveries. Orders cannot be split among tours. Time windows are associated with each store. There is only one type of truck: tractor with semi trailer. The load is measured in pallets, in kilograms, and in cubic metres. There are capacity constraints on each one of these measurement units, and the first one that saturates implies the violation of the constraint. Vehicles are assumed to be infinite, since they are provided by flexible sub-contractors. Subcontractors are distributed all over Italy, and therefore vehicles can start their routes from the first assigned customer, and no cost is incurred in travelling to the first customer in the route. Loading and unloading times are assumed to be constant. This is a rough approximation imposed by the client, since they have been insofar unable to provide better estimates, accounting for waiting times at the store, which are quite variable and unpredictable. It is a conservative and risk-averse approach. In this application AntRoute has been modified to deal with a unique objective function. The algorithm is still able in few minutes to compute solutions that are better than the solutions produced by human planners. The algorithm adapts its work to different problems in different scenarios and it is very flexible and efficient in producing operative solutions. This is due to the ant algorithm where good problem solution structures emerge during the computation thanks to the dynamic use of the artificial pheromone.

Advanced Logistic Systems ask for high performance artificial intelligence algorithms. In the last years complex industrial applications has been solve in few minutes by meta-heuristics algorithms inspired by natural systems. This is a new successful trend that shows that artificial intelligence algorithms are mature for industrial applications.

References

- Bonabeau, E., M. Dorigo, and G. Theraulaz, Inspiration for optimization from social insect behaviour, *Nature* 406, 39–42, 2000.
- Clark, A.J., and H. Scarf, Optimal policies for a multi-echelon inventory problem, *Management Science*, 6, 475–490, 1960.
- Dorigo M., G. Di Caro and L. M. Gambardella. Ant Algorithms for Discrete Optimization. *Artificial Life*, 5,2, pp. 137-172, 1999.
- Gambardella, L., É. Taillard, G. Agazzi. MACS-VRPTW: a multiple ant colony system for vehicle routing problems with time windows. In: D. Corne et al. (eds.), *New ideas in optimization*, McGraw-Hill, London, 1999.
- Glover, F. and M. Laguna, *Tabu Search*. Kluwer Academic Publishers, Boston, MA, 1997.
- Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, *Science*, 220, 671–680, 1983.
- Sodhi, M. S., Applications and Opportunities for Operations Research in Internet-Enabled Supply Chains and Electronic Marketplaces, *Interfaces*, 31(2), 56–69, 2001.
- Stützle, T., and H. Hoos, Analyzing the Run-time Behaviour of Iterated Local Search for the TSP, In *Proceedings of MIC'99, III Metaheuristics International Conference*, 1999.

Investigating the Interplay between Morphology and Behavior Using the Same, but Adaptive Neural Controller

Gabriel Gomez and Peter Eggenberger Hotz

Artificial Intelligence Laboratory

University of Zurich, Andreasstrasse 15, CH-8050, Zurich, Switzerland

{gomez, eggen}@ifi.unizh.ch

In biological systems, at all stages of development, the nervous system must be able to innervate and adapt functionally to any changes in the size and relative proportions of the body. Until now no engineering methods exist to tackle with unforeseen and concurrent changes in morphology, task-environment, and neural structure. We expect to contribute with our approach to the solution of this problem.

Predefining all possible sensors and movement capabilities of a robot, will certainly reduce its adaptivity. In order to be adaptive, a neural controller must be able to reconfigure itself to cope with environmental and morphological changes (e.g., additional sensors, not only in number, but different types of sensors could be added, sensors could be repositioned over time or became damaged, additional or fewer degrees of freedom could be used, the actuation type and strength could also be varied).

As not all possible changes can be anticipated by the designer, the system should be capable to explore its own movements and coherently adapt its behavior to the new situations. This self exploratory activity has been well studied in infants. The exploration of the infants' own capacities is one of the primary driving forces of development and change in behavior. In babies, spontaneous movement creates both tasks and opportunities for learning. Infants explore, discover, and select among all possible solutions from the exploration space, those that seem more adaptive and efficient ([3]).

Aiming to endow our robots with such adaptivity we present a common basis to investigate the growing of a neural network, value systems and learning mechanisms, the so-called: "ligand-receptor" concept ([1], Fig 1a). This concept is used to teach a robotic hand with 13 degrees of freedom, complex dynamics provided by a tendon driven mechanism of actuation and different types of tactile sensors to grasp objects (Fig. 2). In our implementation, receptors abstract proteins of specific shapes able to recognize specifically their partner molecules. Ligands, on the other hand, are molecules moving around, which also have specific shapes, and are basically used as information carriers for their receptors. The shape of a receptor determines which ligand can stimulate it, much in the same fashion, as notches of jigsaw pieces fit exactly into the molds of other pieces. When a receptor is stimulated by a matching ligand (signaling molecule), the following mechanisms can be elicited on a neuron: connect to a neuron expressing a partner receptor, release a ligand molecule, express a receptor. As a result of the interplay of these processes, the specification of the neural network (i.e., number of neuronal fields, size of each neuronal field, a set of receptors expressed by each neuronal unit, a set of signals that can be released by the sensor neurons) can be obtained and then embedded as a neural controller for the robotic hand ([2], Fig 1b).

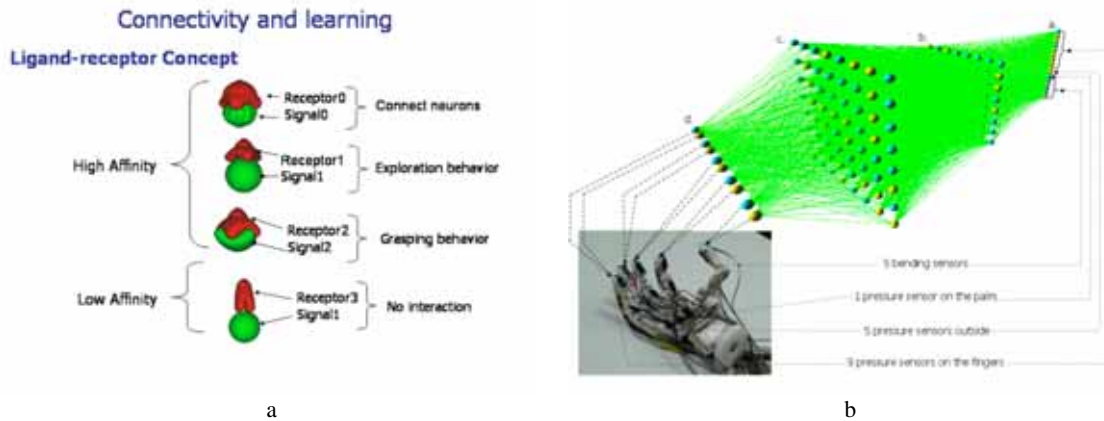


Fig. 1. Ligand-receptor concept. (a) Ligand-receptor interaction based on the affinity between the two entities. (b) Neural network and its connections to the robotic hand.

This biological approach allows us to systematically investigate the interplay between morphology and behavior using the same, but adaptive neural controller. We can make dramatic changes in the morphology and because the neural network is adaptive, it can handle such changes. The basic neuromodulatory system remains the same, but the specific of the arrangement will change. The proposed neural network allows the robotic hand to explore its own movement possibilities to interact with objects of different shape, size and material and learn how to grasp them (Fig 1b).

The experiments are carried out with two different prototypes of the robotic hand, the first prototype was built on aluminum and equipped with standard FSR pressure sensors (Fig. 2a), the second prototype was built on industrial plastic and equipped with pressure sensors based on pressure sensitive conductive rubber (Fig 2b). This made the second prototype approximately half of the weight of the first one. Furthermore, changes in the power of the servo motors and in the length of the tendons made the second prototype not only lighter but stronger.




Fig. 2. Tendon driven robotic hand. (a) First prototype. (b) Second prototype. (c) Final grasp of a tennis ball.

References

- [1] Eggenberger Hotz, P. (2003), *Evolving Morphologies and Neural Controllers Based on the Same Underlying Principle: Specific Ligand-Receptor Interactions*. In Hara, F. & R. Pfeifer, (eds.), *Morpho-functional Machines: The New Species (Designing Embodied Intelligence)* Springer-Verlag, pp. 217-236.
- [2] Gomez, G., Hernandez, A., Eggenberger Hotz, P. (2006). *An adaptive neural controller for a tendon driven robotic hand*. In *Proceedings of the 9th International Conference on Intelligent Autonomous Systems (IAS-9)*. T. Arai et al. (Eds.), IOS Press, Tokyo, Japan, pp. 298-307.
- [3] Smith, L. and Gasser, M. (2005). *The Development of Embodied Cognition: Six Lessons from Babies*. *Artificial Life*. Vol. 11, Issues 1-2, pp. 13 - 30 - Winter-Spring 2005. Special Issue on Embodied and Situated Cognition.

The Communication and Commercialization of Innovation: Observations and Implications

Dr. Simon Grand
RISE Management Research
University of St. Gallen HSG
www.rise.ch

Idea

In the area of Artificial Intelligence, as in other fields of scientific research and technological innovation, researchers and entrepreneurs face a basic challenge: *How is it possible to communicate and commercialize (radically) new ideas, concepts, solutions and products?*

We will highlight a series of paradoxes and challenges inherent in the communication and commercialization of innovation, as they emerge in recent research in the areas of technological innovation and strategic entrepreneurship, as well as in a series of cases from various industries and contexts. In parallel, we will deduce a series of implications for the successful communication and commercialization of innovative ideas and concepts, in general as well as in the fields of Artificial Intelligence and the Cognitive Sciences in particular.

Challenges

In line with Joseph Schumpeter, the founding figure of innovation and entrepreneurship research, we understand by innovation the creation and invention of ideas and concepts, as well as their transformation into new products and solutions, which are successfully realized and accepted on the market and in society. Innovation thus implies the transformation of unconventional, crazy, surprising ideas and concepts (ex ante) into successful, self-evident, unquestioned solutions and products (ex post), through their communication and commercialization. These processes imply a series of challenges:

- Innovations imply most of the time value creation as well as value destruction, through the replacement of existing solutions and established competencies („creative destruction“ in the famous terminology of Joseph Schumpeter).
- Whether the development and commercialization of a new idea or concept is perceived as value creation or rather as value destruction is a question of perspective, situation and context, as well as an issue of particular interests and strategies.
- (Technological) innovation processes imply a diverse series of actors (including universities, companies, entrepreneurs, technologies, research, financial resources, legal regulations, political institutions, public opinion, ...), which together form a „value constellation“.
- Explication, communication, legitimation and justification are thus fundamental for the commercialization and realization of new ideas and concepts, given the systemic, dispersed, collective nature of the existing as well as the intended new value constellation.
- There is an inherent, highly controversial debate in the management community concerning the right measures and criteria for evaluating whether a new solution or product is a „true“ innovation (including degree of newness, systemic nature of the innovation, ...).

- Due to the complex and dynamic nature of innovation processes, it is difficult to explain why a particular idea is more successful than alternative ideas, leading to abbreviated explications of the relevant underlying mechanisms („frame of reference“ problem).

While these challenges are well explored for many important research and technology areas (including Life Sciences, Information Technology, ...), it will be interesting to discuss their relevance for understanding the successful communication and commercialization of new ideas and concepts in the field of Artificial Intelligence and the Cognitive Sciences.

Implications

Management and entrepreneurship under these circumstances is at the same time very important and critical, but also complex and situative. We will argue that our current understanding and conceptualization of the management and entrepreneurship of innovation needs itself substantial reinvention and innovation. The thesis will be, that new insights and concepts from Artificial Intelligence and the Cognitive Sciences might be one of the most promising starting points for such a reinvention, in a sense a way of carrying on the influential contributions of Herbert Simon, which he made based on his view of Artificial Intelligence to economics and management: What would be the contributions of the new concepts in Artificial Intelligence and the Cognitive Sciences to management, both conceptually and practically?

Complex Behaviour Recognition in Interacting Agents

Verena V. Hafner
TU Berlin, Fakultät für Elektrotechnik und Informatik, DAI Labor
Berlin, Germany
vvh@ieee.org

Introduction

Intelligent and complex agent behaviour can sometimes arise from very simple rules of interaction with the environment or other agents. Here, the physical interaction plays an important part [1]. A famous example for such emergent behaviour are the vehicles from Braitenberg's thought experiments [2]. The behaviours of these vehicles were even interpreted as expressions of love or fear.

The other direction is much more complicated: How can an observed or experienced agent behaviour be understood in terms of rules, goals, or even intentions?




Fig. 1 a) AIBO robots interacting with each other. b) Infantoid robot being imitated by the author using two coloured pet baby rabbits for easy visual discrimination. c) Simple 'QuakeII' environment in which the interaction with two human players was performed.

Mapping Behaviours

In Hafner and Kaplan [3], we introduced the concept of interpersonal maps and applied it to a series of experiments with Sony AIBOs (see figure 1a). Interpersonal maps are an extension of somatosensory body maps and involve two agents. They not only represent the morphology (sensors and actuators) of the agents, but also specific characteristics of the interaction for a fixed time period. These body maps are created by using information-theoretic measures between pairs of sensors. Details can be found in [3].




Fig. 2. a) Interpersonal map of a non-interactive player behaviour in Quake II. b) Interpersonal map of a chasing behaviour in Quake II.

The concept is independent from the actual sensors used, thus could be applied to a range of different situations. One example we are currently investigating is understanding human player behaviour in a computer game, in order to create intelligent

artificial game-bots that behave as human players do [4] (see figure 1c). In figure 2, the interpersonal maps for two human players performing different kinds of interactive behaviour are displayed. The player maps on the left are separate for two

non-interacting players, the maps on the right are overlapping and similar sensors are spatially close for two interacting players performing a chasing behaviour. It enables us to roughly discriminate the kind of behaviour, however, it is impossible to reconstruct the behaviour itself by only using the maps.

Another area where these interpersonal maps could advance research in artificial intelligence is the study of imitation behaviour in developmental robotics. A couple of experiments has been performed with the 'Infanoid' robot [5] (see figure 1b) whose movements have been imitated by a human experimenter. Resulting interpersonal maps can show the grade of interaction (e.g. unrelated movements, delayed imitation, immediate imitation) which reflects the imitation behaviour.

Summary

We believe that understanding the interaction behaviour of agents in real world or realistic environments is crucial for understanding intelligence expressed in such a way. One of the challenges is to understand how emergent behaviour is created. The dynamics play an important factor of such intelligent behaviour.

References

1. Pfeifer, R. and Scheier, C. *Understanding Intelligence* (MIT Press, Boston, 1999).
2. Braitenberg, V. *Vehicles. Experiments in Synthetic Psychology* (MIT Press, 1984)
3. Hafner, V. V. and Kaplan, F. *Interpersonal Somatosensory Maps*, Proceedings of the AISB 2005 Third International Symposium on Imitation in Animals and Artifacts, pp. 48-53, Hatfield, UK (2005)
4. Roth, M. and Bauckhage, C. and Hafner, V.V. *Mining Behaviors Using Interpersonal Maps* (submitted to SAB2006 Workshop on Adaptive Approaches for Optimizing Player Satisfaction in Computer and Physical Games)
5. Kozima, H. *Infanoid: A babybot that explores the social environment*, K. Dautenhahn, A. H. Bond, L. Canamero, B. Edmonds (eds.), *Socially Intelligent Agents: Creating Relationships with Computers and Robots*, Amsterdam. (Kluwer Academic Publishers, pp.157-164, 2002)

Adaptive Multi-Modal Sensors

Kyle I. Harrington*, Hava T. Siegelmann[†]

*School of Cognitive Science

Hampshire College, Amherst, USA

[†]Department of Computer Science

University of Massachusetts, Amherst, USA

kharrington@hampshire.edu, hava@cs.umass.edu

Real world applications of robotics and artificial intelligence sometimes require input from unknown environments (Pederson, 2001). For such cases an internal representation of the environment must be created on-the-fly. The use of multi-modal sensors allows general solutions to be used for various classes of unknown environments. By adaptively changing the modality of sensors it is possible to, in many cases, reduce the number of sensors necessary to adequately sense the environment, as well as reduce the amount of bandwidth used by sensors.

Here we present a self-constructive approach to developing a multi-cellular artificial organism capable of adjusting sensor modalities based upon sensory input from a dynamic environment (Fig. 1). The results demonstrate the robustness of this system within an environment of which the organism has no a priori knowledge.

Organism Design

The multi-cellular organism exists in a 2.5D lattice (Fig. 1a.), where the half dimension represents the ability of cells to stack. Four cell types are used, root stem cells, structural, sensor stem, and sensor cells. A root stem cell can differentiate into root stem cells, structural and sensor stem cells, and a sensor stem cell can

differentiate into sensor stem cells and sensor cells. We use chemical diffusion to regulate cell differentiation, division, and migration, in a way similar to Eggenberger's (1997) use of transcription factors, and the use of gases within a GasNet (Husbands et al., 1998).




Fig. 1. Three views of the system. (a) the cellular organism (b) the environment to be sensed (c) the sensor activity within the organism for the exact location on (b)

The organism is initialized as a single root stem cell, then iteratively grows via cell division until it reaches equilibrium.

Final size and steps until equilibrium are determined by both diffusion of growth chemicals and the differentiation criteria. The self-constructive growth allows the organism to self-repair. If cells are removed from the organism; diffusion causes the organism's chemical levels to decrease below equilibrium, causing division to resume until equilibrium is reattained.

	1	2	3	4	5	6	7	8	9	10
S1	R	G	R	G	R	R	G	R	G	R
S2	R	B	B	B	R	B	R	B	B	B
S3	G	G	R	G	G	G	G	G	G	G
S4	R	R	R	R	R	R	R	G	G	G
S5	R	G	R	G	R	G	G	R	G	R

Fig. 2. Sensors tend to focus on two modes, and alternate between the two based upon which is more frequent.

Each sensor is capable of being in one mode at a time. For each active sensing mode, A, there are two chemicals that act as sensory memory, C_{SAS} , short-term and C_{SAL} , long-term memory, the former with a higher diffusion coefficient than the latter. These chemicals are produced by,

$$C_{SAS} = D_{SAL} * S_A \quad C_{SAL} = D_{SAS} * S_A$$

where, S_A is the activation value of the sensor for mode A. The active sensing mode switches if, either $C_{SAS} < C_{SAL}$, or $C_{SAL} < \Theta_{SA}$, where Θ_{SA} determines the minimum amount of activity to stay in mode A. For us each sensor was capable of three sensory modalities, red, green, and blue.

Conclusions

The mean number of sensors in each modality was recorded across a run of 10 organisms. Each run was 60,000 steps. After 40,000 steps 20% of the cells were removed to demonstrate self-repair. The rapid change between modalities in Figure 2 is the feature of our organism that allows it to properly adapt to its environment.




Fig. 3. The top line is the total number of sensors and the bottom are the division of active modalities.

Figure 3 presents results supporting our organism's capability of adapting its sensors to the distribution of color within the environment by maintaining a distribution of sensor modes which is representative of the environment.

Future research will examine the utility of interpolating sensor values from the sensor chemical memory in order to develop a more complete internal representation of the environment.

References

1. Pedersen, L., Wagner, M., Apostolopoulos, D., Whittaker, W.R.. *Autonomous robotic meteorite identification in Antarctica* (Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation, 2001)
2. Eggenberger, P.. *Evolving Morphologies of Simulated 3d Organisms Based on Differential Gene Expression*. (Proceedings of the Fourth European Conference on Artificial Life, Springer-Verlag, 1997)
3. Husbands, P., Smith, T., Jakobi, N., and O'Shea, M.. *Better Living Through Chemistry: Evolving GasNets for Robot Control*. Connection Science, (10)3-4:185-210, 1998.

From Artificial Intelligence to Artificial Life in the first 50 years; To Homeostasis and Entropy Production in the next 50?

Inman Harvey
CCNR/EASy, University of Sussex
inmanh@susx.ac.uk

The spectrum of approaches towards Artificial Intelligence has broadened immensely over the last 50 years, away from logic and abstract disembodied reasoning, and towards a more biological view of intelligence as adaptive embodied behaviour. The logical and abstract end of the spectrum was, from the 1950s, based on the computational metaphor for the brain or the mind; this remains present and thriving, but it is no longer the only game on the block.

Two prominent movements became noticeable during the 1980s: Artificial Neural Networks (ANNs), and Behaviour-based robotics. ANNs were often initially seen as answers (or the guesses of computer scientists) to how a brain might do computations *in vivo*, highly parallel but using simple components. But many of those advocating behaviour-based robotics (or taking a dynamical approach to ANNs) were not proposing an alternative form of computation, but rather rejecting the idea that computation had anything much to do with cognition at all. Some took an Evolutionary Robotics approach (Harvey et al 1997) towards designing dynamical systems for generating adaptive behaviour in artefacts, whilst rejecting the need (or the sense) of interpreting these systems as computational.

Humans can of course perform relatively simple computations, and make machines to do more complex ones. But if one places humans in a biological and evolutionary context, the ability to reason is ‘merely’ a very recent sophisticated behaviour in one species, that is built on top of several billion years of development in adaptive behaviour. Naturally, that species being rather self-centred, it has placed rather too much emphasis on this; but some of us can try and be less biased.

Taking stock now in 2006, this is the context in which much of Artificial Life, Evolutionary Robotics, and a Dynamical Systems approach to Cognition is carried on. Let us now speculate how this spectrum of approaches may broaden still further in the next 50 years; what follows are my personal guesses.

We have already broadened the notion of Intelligence to be one form of adaptive behaviour; a more fundamental analysis of adaptive behaviour is required. For instance, in order to build robots that till now have been lacking ‘the juice’ (as Brooks calls it, Brooks 2002) perhaps we have to build genuinely self-creating, self-repairing, autopoietic (Varela et al 1974) machines; machines with a metabolism, that maintain their organisation through extracting and degrading energy and matter from the environment. We will need to understand the fundamentals of how biological organisms are grounded in the physical world.

A core concept is *homeostasis* (Ross Ashby 1952); indeed autopoiesis is homeostasis of the organisation and identity of an organism, its maintenance in the face of perturbations. Studies (Harvey 2004) coming out of analysis of homeostasis in Daisyworld models suggest that homeostasis in any system may be rather simple to arrange, given some very basic assumptions. Provided a system maintains its identity, and interacts with perturbing forces, then homeostasis will come ‘for free’; and with it, Cognition and Adaptive Behaviour. But why should such living systems have originated in the natural world?

Recent work (Dewar 2003) has given a firmer grounding to a Principle of Maximum Entropy Production in non-equilibrium steady-state systems such as this planet. This makes it reasonable to expect that when conditions *allow* profligate, energy-degrading, entropy producing systems such as living organisms to be viable, then it is overwhelmingly *likely* that they will indeed arise. A planet with organisms generates entropy faster than one without. These very general principles are not specific to the materials used, and are equally applicable to artificial and to carbon-based life-forms.

So based on very recent work, my speculation is that, whilst the middle and the other end (the computational, rational end) of the AI spectrum will continue to see progress, what will be new and thriving in the coming decades will be a move towards more understanding of the *physical basis of living systems*, both real and artificial. Life means enhanced Entropy Production, Homeostasis, Adaptive Behaviour, and Cognition (Stewart 1991). One specialised form of Cognition is abstract reasoning.

References

- Brooks R.A., (2002): *Flesh and Machines: How Robots will change us*. Pantheon Books, NY.
- Dewar R.C. (2003): Information theory explanation of the fluctuation theorem, maximum entropy production and self-organized criticality in non-equilibrium stationary states. *J. Phys. A: Math.Gen.*, **36** 631-641.
- Harvey, I., Husbands, P., Cliff, D., Thompson, A., and Jakobi. N. (1997): Evolutionary Robotics: the Sussex Approach. *Robotics and Autonomous Systems*, **20**, 205-224.
- Harvey, I., (2004). Homeostasis and Rein Control: From Daisyworld to Active Perception. *Proc. 9th Int. Conf. on the Simulation and Synthesis of Living Systems, ALIFE'9*, Pollack, J., Bedau, M., Husbands, P., Ikegami, T., and Watson, R.A. (eds), pp. 309-314. MIT Press, Cambridge MA.
- Ross Ashby, W, (1952): *Design for a Brain*, Wiley, New York.
- Stewart, J., (1991): Life=Cognition: The Epistemological and Ontological Significance of Artificial Life. Proc of ECAL 91, Paris, Bourguin P. & Varela F. eds, MIT Press, pp 475-483.
- Varela, F. J., Maturana, H. R.. & Uribe, R. (1974): Autopoiesis: the organization of living systems, its characterization and a model. *Biosystems* **5** 187–196.

Acting Lessons for Artificial Intelligence

Guy Hoffman
The Media Laboratory
MIT, Cambridge, MA, USA
guy@media.mit.edu

Theater actors have been staging artificial intelligence for centuries. If one shares the view that intelligence manifests in behavior, one must wonder what lessons the AI community can draw from a practice that is historically concerned with the infusion of artificial behavior into such vessels as body and text. Like researchers in AI, actors construct minds by systematic investigation of intentions, actions, and motor processes with the proclaimed goal of artificially recreating human-like behavior. Therefore, acting methodology may hold valuable directives for designers of artificially intelligent systems. Indeed, a review of acting method literature reveals a number of insights that may be of interest to the AI community, three of which are outlined below:

Psycho-physical Unity

Acting teacher Augusto Boal states firmly: “one’s physical and psychic apparatuses are completely inseparable. [...] A bodily movement ‘is’ a thought and a thought expresses itself in corporeal form.”¹ The utter unity of mental and motor expression is stressed throughout acting method, regardless of school, and dating back at least as far as the 19th century.² The physical aspects of psychological motivation become even more central in modern Stanislavskian method,³ which trains actors to both uncover physical manifestations of internal processes, and conversely evoke mental processes *through* physical action. This approach also finds mounting support in recent findings in cognitive psychology, which link cognition to a multi-modal simulation of physical experience.⁴

In recent years the AI community has finally come to recognize the value of embodied cognition. To many of us it seems increasingly futile to develop intelligence without regard to the physical aspects of sensation and action. Yet in the majority of our systems, predominantly semantic “decision” modules are still connected to “sensory” and “motor” subsystems by lines as thin as box-chart arrows or network packet streams. Taking a hint from acting method, it may be time to tear down the barrier between Thinking and Doing, and explore an alternative in which intelligent behavior does not merely communicate with physical behavior, but is part and parcel of the same process – an indivisible whole.

Mutual Responsiveness

Just as thought does not happen in a physical void, useful agents do not act in a social void, and the most apt intelligent machines should interact well with humans and other artificial agents. Much like the shift in AI from single-agent approaches to multi-agent and human-interaction systems, theater practice is increasingly concerned with the relationship between actors as much as it is with each actor’s individual performance. Acting guru Sanford Meisner is most famous for placing much of the content of a scene in the interaction *between* the actors, endorsing a seemingly odd repetition exercise in which actors can only repeat what their scene partner says. He often stated that “what you do doesn’t depend on you, it depends on the other

fellow.”⁵ Others, like Ruth Maleczek, speak of behaviors “bouncing off” the other actor and subsequent actions coming “directly from the response of the other actor.”⁶

In AI, much more emphasis can be placed on the emergence of intelligent behavior from an agent’s pure reaction to other agents, and in particular to humans who often exist in the agent’s environment. We may look to exercises executed by actors when practicing reaction, repetition, and breaking away from mirroring for possible insight into the mechanism of mutual responsiveness that could prove crucial for intelligent behavior.

Continuous Inner Monologue

Perhaps the most significant contribution of the Stanislavski system was the elimination of so-called “representational acting”, a beat-to-beat development of purely external expression. Instead, modern actors work in terms of overarching motives, objectives, obstacles, and intentions, which eventually lead to action selection.³ Moreover, actors are expected to carry out a continuous *inner monologue* throughout their stage presence, leading up to their lines and preventing their text to be uttered as a series of isolated parts.

Maintaining continuity through inner processes is also good advice for artificially intelligent agents, and prescriptive if we are aiming for naturally behaving agents. This should be of particular note for those of us building robots interacting with human counterparts. If we are to steer away from the command-and-response interaction so prevalent in our dealings with artificial agents, action selection should not stem only from the most recent input but grow out of a continuous and multi-layered stream of constantly changing internal parameters. How to reconcile this advice with the requirement for mutual responsiveness laid out above is a worthy challenge for the AI community, one into which the acting discipline – not a stranger to this paradox – might also have relevant insights.

Summary

Well outside of the spotlight of even the most interdisciplinary of AI research, actors have for decades confronted problems that are surprisingly related to the ones the AI community tackles. The guidelines described in this document are examples indicating the potential benefit that the AI community can glean from a closer look at a discipline that has repeatedly concerned itself with the faithful production of artificially intelligent behavior.

References

1. Boal, A. *Games for Actors and Non-Actors* (2nd ed., Routledge, New York, 2002).
2. Stebbins, G. *DelSarte System of Expression* (2nd ed., E.S.Werner, New York, 1887).
3. Moore, S. *Training an Actor: The Stanislavski System in Class* (Viking, New York, 1968).
4. Barsalou, L.W., Niedenthal, P.M., Barbey, A., & Ruppert, J. *Social embodiment*. In B. Ross (Ed.), *The Psychology of Learning and Motivation*, (Academic Press, San Diego, 2003).
5. Meisner, S. and Longwell, D., *Sanford Meisner on Acting* (Vintage, New York, 1987).
6. Sonenberg, J. *The Actor Speaks* (Three Rivers, New York, 1996).

Right body, right mind

Owen Holland

Department of Computer Science, University of Essex, United Kingdom
owen@essex.ac.uk

What AI should be about

Forget about computer vision, or speech recognition, or robot navigation – the right true end of AI is the creation of artificial systems like ourselves. (If you doubt this, go to the movies....) Partial systems, or systems like animals, or systems that solve some of the same problems but in radically different ways – think Big Blue – simply do not add up to success, however useful and encouraging they might seem to us as steps along the way. So what are the necessary ingredients of systems like ourselves? What most differentiates a person from a rock, other than being alive? The answer has two parts: from the outside, we produce intelligent and effective behaviour; and from the inside, we are conscious. If we go back to the 1950s, we find that there is a clear identification of intelligence with human thinking, and that the nature of human thinking was then implicitly conscious. For example, in 1959, Newell and Simon¹ declared the aim of their enterprise to be “...the explanation of complex human behavior” and they judged the efforts of themselves and their peers to have shown that it was “...no longer necessary to talk about the theory of higher mental processes in the future tense”. In the last fifty years – with the honourable exception of the COG project² – progress seems to have consisted in equal measure of sporadic real achievement and continuous lowered expectations. What is needed right now is the restoration of the ambition of the pioneers, and a new accommodation with the expectations of the public who fund us: *we should commit ourselves to building an artificial entity that is both intelligent and conscious.*



Fig 1. Cronos, the anthropomorphic robot, designed and built by Rob Knight

How to go about it

The last fifty years, inside and outside AI, have surely taught us enough to embark on this quest. We know a lot about the structure of the brain, and how its components work. We now appreciate the importance of embodiment in constraining and enabling both reactive and cognitive performance in animals and robots³. The new research area of machine consciousness has a dozen workshops under its belt already. The thesis of the adequacy of ungrounded symbol manipulation was challenged in the 1980s by its antithesis, the equally rigid doctrine of the adequacy of embodied reactive architectures, and we are now seeing a movement towards synthesis. There is a growing realization that control systems engineering has a lot to say about the successful design of autonomous embodied systems, and the first bridges are being built between control theory, psychology, neuroscience, and consciousness⁴. It is difficult to see what is holding us back – other than our own caution. We know enough – enough, at least, to fail in a useful way.

First steps

At the Universities of Essex and Bristol, we have embarked on a project that we hope will serve as the forerunner of other projects that will eventually end in achieving a truly human-like artificial intelligence – one that is conscious. Many lines of evidence within consciousness studies – notably those reviewed and synthesized by Thomas Metzinger⁵ – identify the core of consciousness as being a dynamic internal representation of the body; this virtual body is in constant interaction with an internal representation of the world, achieving intelligence and producing consciousness. In order to produce an artificial agent that has the right kind of internal representation of its body, we have produced a robotic platform that is as functionally similar to that of a human as seems necessary and practicable. It has a skeleton of bone-like elements, joined by freely moving joints (there are 45 degrees of freedom in the torso alone), and driven by paired series-elastic actuators. Such a robot, which we call an anthropomimetic robot⁶ to distinguish it from other merely humanoid robots, is radically different in many ways from conventional implementations. Every movement or imposed load produces a reaction that is transmitted through the elastically interconnected multi-d.o.f structure, and even the simplest action requires the predictive cancellation of unwanted consequences on a whole-body basis, making every movement a whole-body movement. This is not a problem that needs to be engineered out – it is the problem that is solved by the brain, and the internal model of the body is both enabled and constrained by this solution.

The robot, Cronos, is equipped with a visual system modeled closely in both structure and function on that of humans. The single eyeball contains a high resolution colour camera which is mapped to a foveal representation, and which is processed by biomimetic receptive field neurons before being fed to a saliency mapping system for gaze control, and a complex spiking neural network for feature extraction and further processing. The investigative plan is to expose the robot to a complex environment in which it is required to achieve some mission, and to design, learn, or evolve appropriate internal representations of the body and the world capable both of guiding intelligent behaviour, *and of giving rise to phenomenal consciousness*.

The irrelevance of success

No-one expects us to succeed, and they are usually happy to supply plenty of reasons for their view. We don't expect to succeed at our first attempt either, but we remain convinced that both the target and the central method are correct and appropriate for the 21st century, and our secondary objective is to convince others of this.

References

1. Newell, A. and Simon, H.A. (1959) The simulation of human thought. The RAND Corporation, Mathematics Division, Paper P-1734.
2. The Practical Requirements for Making a Conscious Robot. Dennett, D.C. (1994) *Philosophical Transactions of the Royal Society*, A, 349, 1994, 133-46.
3. Pfeifer, R. & Scheier, C. *Understanding Intelligence* (MIT Press, Boston, 1999).
4. Grush, Rick (2004). The emulation theory of representation: motor control, imagery, and perception. *Behavioral and Brain Sciences* 27:377-442.
5. Metzinger, T. (2003) Being No-one: The Self-Model Theory of Subjectivity. MIT-Bradford
6. <http://cswww.essex.ac.uk/staff/owen/research.htm#Anthropomimetic%20robots>

Analog-Digital Model of Cortical Memory Consolidation

Lars E. Holzman, Hava T. Siegelmann
Department of Computer Science
University of Massachusetts, Amherst, USA
[lholzman | hava]@cs.umass.edu

We describe a cortically inspired recurrent neural network that is an extension of work done by Hahnloser, et. al. (2000) on a digital-analog ring circuit. We show analytically and empirically that this network has properties that make it a prime candidate for memory consolidation.

Recent research suggests that fixed time processing occurs on information after exciting the receptive fields of the cornea and before reaching the memory storage apparatus in the Hippocampus (Van Rullen, 2003). Drawing from this, the network we propose operates only for a fixed time on each input. This is in stark contrast to previous models of memory consolidation such as the Hopfield network (Hopfield, 1982) which stop only when the network reaches a stable state. The time necessary to converge is not known in advance, varies with each input, and in fact may not even be deterministic. Furthermore, when models are built to be stored in memory, typically only noisy examples are available. We show that using noisy examples our network creates a condensed form of the concept that maintains the properties of the inputs necessary for their intended use. Finally, the memories created by our network are dynamic and change when the memory is recalled and when more input is provided. We then predict how memories and concepts are changing with monotonic or non-monotonic changes of the input, with changing in more than one coordinate and when different memories collapse to one.

The network

As with the ring-circuit our network possesses digital state (neurons are active or inactive) and analog scaling (active neurons show different levels of gain). However, we extend the circuit to a two-dimensional form to model the early stages of the visual cortex. To do so we lay a set of excitatory neurons out in a grid and update them using:

$$\tau \frac{dE_{(a,b)}}{dt} + E_{(a,b)} = \left[e_{(a,b)} + \sum_{i=-2}^2 \sum_{j=-2}^2 A(\text{dist}((a,b), (i,j))) E_{(i,j)} - \beta I \right]^+$$

where τ is a timing constant, $E_{(a,b)}$ is the activity of the neuron at grid location (a,b) , $e_{(a,b)}$ is the activity of external inputs for location (a,b) , A is a function that converts distance to weighting, dist is the distance between two neurons, β scales the inhibition, I is the activity of the one inhibitory neuron, and $[\]^+$ is one sided positive rectification.

We analyze the network by examining the update rule in two parts. The first part $e_{(a,b)}$ taken on its own is shown to be an unbiased estimator of the parameter of a Bernoulli trial (and thus represent the mean value of inputs to the neuron). The second part, consisting of the summation, inhibition and rectification is shown to perform Gaussian smoothing on the produced parameter estimates. By assuming that inputs arrived on periodic fixed time intervals we see that the hybrid rule consisting of both

parts is capable of building a model that has fewer neurons representing the model and less noise than in parameter estimation but maintains the strength of the signal being studied. Results indicative of this for a simple signal and noise model can be found in Fig. 1.

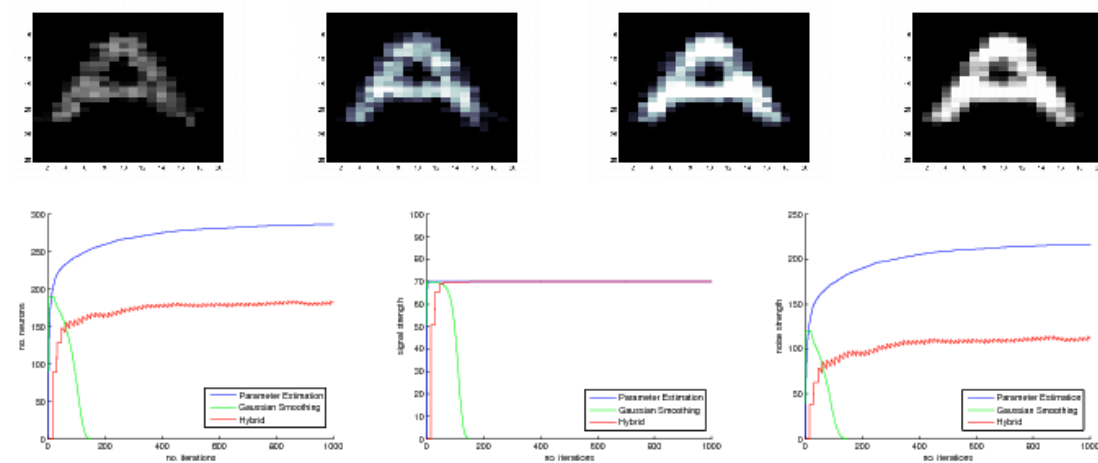


Fig 1. The effect of a fixed-time convergence analog-digital neural network on noisy inputs (top), number of neurons for partial and full update rules (left), strength of signal for partial and full update rules (center), strength of noise for partial and full update rules (right).

To retrieve memories we use a surprise metric between new inputs and existing memories. Itti and Baldi formalize the notion of surprise as being the change in a model with respect to a new input. We develop a variant of their metric using specific properties of our network to construct a “surprise minimization” framework of memory. In this framework models are selected for a new input when the surprise of the input with respect to a particular model is the lowest. The model selected is then updated with the new input. This causes the model to track concepts if they are changing monotonically without losing the compactness or signal strength properties shown for the static case.

This work is an important example of how modern artificial intelligence can create understandable models of psychological phenomenon which can be analytically studied. It uses analysis and simulation to make testable predictions which anchor the work in the neuroscience.

References

1. Hahnloser, R. H., Sarpeshkar, R., Mahowald, M., Douglas, R. J., & Seung, H. S. “Digital selection and analog amplification coexist in a silicon circuit inspired by cortex”. *Nature*, 405, 947--951. (2000)
2. Hopfield, J. J. "Neural networks and physical systems with emergent collective computational abilities", *Proceedings of National Academy of Sciences*, vol. 79 no. 8 pp. 2554–2558, (1982)
3. Itti, L. & Baldi, P. Bayesian Surprise Attracts Human Attention, *Advances in Neural Information Processing Systems*, Vol. 19 (2006)
4. VanRullen, R. & Koch, C. “Is Perception Discrete or Continuous”. *TRENDS in Cognitive Sciences*. 7 (5). (2003)

Designing Physical Impedance for Engineering Adaptive Intelligence

Koh Hosoda

Department of Adaptive Machine Systems, Osaka University

Yamdaoka 2-1, Suita Osaka, 565-0871, Japan

hosoda@ams.eng.osaka-u.ac.jp

<http://www.robot.ams.eng.osaka-u.ac.jp>

Designing impedance is a key issue for engineering adaptive intelligence

Behavior of an autonomous agent emerges from interaction among the physical body, brain, and environment. To deal with the environmental changes, therefore, designing impedance of the body and brain is essential. The role of physical impedance is not so much studied and engineered while brain impedance, that is, adaptive control is studied so far since it can be realized easily by computers. Physical impedance will provide:

1. quick response against the environment which is essential for realizing dynamic motion and
2. rich information about the environment since the interaction between the agent and the environment is tight

To show how the physical impedance contributes to adaptability on these points, two research projects are on going.

Biped robots driven by antagonistic pairs of pneumatic actuators

To show the contribution of the physical impedance to change interaction between the agent and the environment, biped robots are developed and investigated (Figure 1). Their joints are driven by antagonistic pairs of pneumatic actuators that can change the mode to interact the environment. It can bring behavioral change to the robot, walking and running each of which needs different impedance. If we want to build an agent which can perform several modes of interaction (in this case, locomotion), the ability to change the physical impedance is indispensable.



Figure 1: A biped that can walk and run

Anthropomorphic soft fingertip with distributed receptors

Physical impedance of the agent will provide rich sensing ability as well as more modes to interact with the environment. An anthropomorphic soft fingertip is developed which is equipped with many receptors inside (Figure 2). Since it is soft and tightly interacts with the object (environment),



Figure 2: An anthropomorphic fingertip and its application

obtained information about is more than that can be obtained by the traditional rigid sensors.

These studies are more emphasizing on anthropomorphism than the other work on embodied artificial intelligence: antagonistic and elastic drive for realizing dynamic motion and elastic fingertip with randomly distributed receptors. Since the morphology of a human plays a great role for human intelligence, we may be able to infer more than the existing approaches. For contribution to the robotics, we may be able to utilize the power of morphology of anthropomorphism for realizing *intelligent* behaviors.

References

- [1] Takashi Takuma, Koh Hosoda, and Minoru Asada, "Walking Stabilization of Biped with Pneumatic Actuators against Terrain Changes", Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp.2775-2780, Aug. 2005.
- [2] Koh Hosoda, Takashi Takuma, and Masayuki Ishikawa, "Design and Control of a 3D Biped Robot Actuated by Antagonistic Pairs of Pneumatic Muscles", Proceedings of 3rd International Symposium on Adaptive Motion in Animals and Machines, CD-ROM, Sep. 2005.
- [3] Koh Hosoda, Yasunori Tada, and Minoru Asada, "Anthropomorphic robotic soft fingertip with randomly distributed receptors", Robotics and Autonomous Systems, Vol.54, No.2, pp.104-109, 2006.
- [4] Yasunori Tada, Koh Hosoda, and Minoru Asada, "Learn to Grasp Utilizing Anthropomorphic Fingertips together with a Vision Sensor", Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp.486-491, Aug. 2005.

In the search of principles underlying cognitive phenomena

Martin Hülse^{1†}, Keyan Zahedi¹, Steffen Wischmann², Frank Pasemann¹

¹Fraunhofer Institute for Autonomous Intelligent Systems,
Schloss Birlinghoven, Sankt Augustin, Germany

[martin.huelse | keyan.zahedi | frank.pasemann]@ais.fraunhofer.de

²Bernstein Center for Computational Neuroscience,
University Göttingen, Germany
steffen@chaos.gwdg.de

Introduction

The dynamical system approach to natural cognitive systems is formulated as an empirical hypothesis, that can only be validated “if in a long run, the best theories of cognitive processes are expressed in dynamical terms.”¹ From this point of view scientists emphasize the importance of concrete examples of minimal cognitive systems developed or demonstrated with autonomous robot systems². The formal analysis of such minimal systems are the prerequisites for dynamical explanations and “abstracting general principles” of artificial cognitive systems, and might also give us the right tools and experience for the study of complex natural cognition as dynamical systems.^{2,3}

Offering minimal solutions and new questions

Applying the approach of evolutionary robotics⁴ and artificial life⁵ we are able to evolve recurrent neural networks (RNNs) of general type. They serve as neuro-controllers for autonomous robots that act successfully in an open, noisy and changing environment. Due to the stochastic character of the applied evolutionary algorithm⁶ the development of small, but complex RNNs

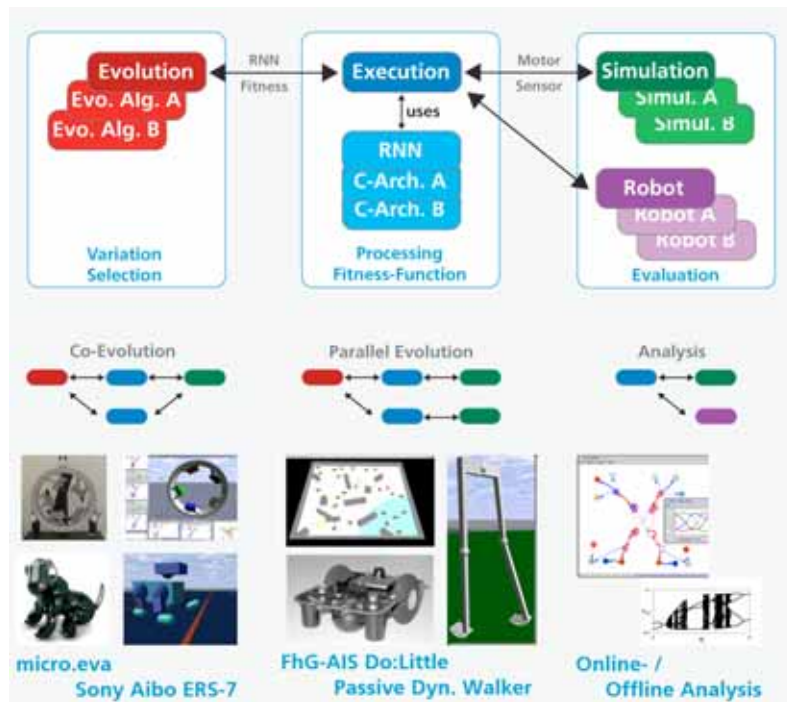


Fig.1. Some applications of the ISEE package

can be forced. The minimality of the resulting RNNs allows us the study of their behavior relevant dynamics in a reasonable depth. Many examples demonstrate the utilization of complex phenomena for behavior control, like bistability⁶, quasi-periodic⁷ and chaotic attractors⁸. But minimal control structures of autonomous systems support also the study of embodiment⁹. Finally, the “embedding” of complex neurodynamics into the sensorimotor-loop can also offer interesting questions in the field of dynamical systems theory. Within this theory neuro-control for autonomous robots can be formally described as parameterized dynamical systems that underlie

specific boundary conditions. These boundary conditions are represented by the characteristics of the system's parameters. Specific parameter characteristics, like the time scale on that a parameter is operating, can evoke interesting transient effects. Such transient effects can, for instance, create adaptive low-pass-filter¹⁰.

ISEE: A structure evolution environment

All the above mentioned examples suggest the strong interdependency of synthesis and analysis of complex behavior control in order to get new insights into the neuro-dynamical effects and phenomena underlying cognitive phenomena. Our integrated structure evolution environment (ISEE) provides a unified framework for the evolution and analysis of RNNs. The evolution is open to diverse robot platforms and experimental setups due to a general interface that enables the connection to each reasonable robot simulation.

The activities of a RNN can be analyzed during the robot-environment interaction, but also off-line to simulate the RNNs as isolated dynamical systems, e.g. computation of bifurcation diagrams and iso-periodic plots, etc. The versatility of ISEE is sketched in the figure, where the general architecture and some sample applications are given, like the evolution of complex control for sensor-driven walking machines, co-evolution of different populations of cooperating controllers, co-evolution of morphology and control of a simulated bipedal walker and finally the evolution of swarm behavior.

Among others, these examples indicate that evolution and analysis of minimal, but complex neural control provide a promising framework for further studies of cognitive phenomena.

References

1. Port, R. F. & Gelder, v. T. *Mind as Motion* (MIT Press, Cambridge 1995).
2. Beer, R. D. A dynamical systems perspective on agent environment interaction. *Artificial Intelligence*, **72**, 173-215.
3. Pasemann, F., Neuromodules: A dynamical systems approach to brain modeling, in: Hermann, H., Pöppel, E. & Wolf, D. (eds.) *Supercomputing in Brain Research – From Tomography to Neural Networks*.
4. Nolfi, S. & Floreano, D. *Evolutionary Robotics* (MIT Press, Cambridge 2000).
5. Langton, C. G. *Artificial Life* (MIT Press, Cambridge 1995).
6. Hülse, M., Wischmann, S. & Pasemann F.: Structure and function of evolved neuro-controllers for autonomous robots. *Connection Science*, **16**, 294–26.
7. Wischmann, S., Hülse, M., Knabe, J. & Pasemann F.: Synchronization of internal neural rhythms in multi-robotic systems. *Adaptive Behavior*, in press.
8. Pasemann, F.: Evolving neurocontrollers for balancing an inverted pendulum, *Network: Computation in Neural Systems*, **9**, 495-511.
9. Pfeifer, R. & Scheier, C. *Understanding Intelligence* (MIT Press, Boston, 1999).
10. Manoonpong, P., Pasemann, F., Fischer, J. & Roth, H. : Neural Processing of Auditory Signals and Modular Neural Control for Sound Tropism of Walking Machines. *International Journal of Advanced Robotic Systems*, **2**, 223-235.

From Locomotion to Cognition: Cheap Design Approach to Adaptive Behavior

Fumiya Iida, Gabriel Gomez, and Rolf Pfeifer
Artificial Intelligence Laboratory, Department of Informatics
University of Zurich, Andreasstrasse 15, CH-8050 Zurich, Switzerland
iida@ifi.unizh.ch

In the traditional approach of robotics and artificial intelligence, one of the underlying implicit assumptions is that feedback control in the precisely modeled environment is the sole generator of intelligent behavior. The main body of research, therefore, has focused on computation, i.e. constructing a world model and making stepwise decisions of motor actions. A conceptual impact has been made during the last few decades, which has led to the paradigm of behavior based robotics, and more recently embodied artificial intelligence (Brooks, 1991; Pfeifer and Scheier, 1999; Iida et al., 2004). This approach deals with adaptive behavior in the context of decentralized control and physical system-environment interactions, and thus the body as a physical entity became a central issue of interest. In order to deal with complex dynamic environment with limited resources of sensors and actuators, adaptive behaviors cannot be reduced to control and computation only. The physical constraints of the system's own body and the environment have to be exploited. From this perspective, the concept of "cheap design" was formulated to systematically explore the relation between adaptive behaviors and morphological properties (Pfeifer and Scheier, 1999).

A number of researchers in biology and robotics have been attracted by the study of legged locomotion, because a legged robot in the real world continuously encounters new demanding situations derived from the physical constraints of their own body and the environment. We have been investigating the issues of legged robot locomotion by mainly focusing on the following two questions: how body dynamics can be used for behavioral diversity, and how body dynamics influences the sensory stimulation, and more generally cognitive processes.

In animals' locomotion a lot of control is not achieved by the brain alone, but through the interaction with the environment and the material properties of the limbs such as elasticity and damping (e.g. Dickinson et al., 2000). Based on these insights, roboticists have started building artificial creatures by exploiting similar principles for their designs (e.g. Collins, et al., 2001). In our investigations of one-, two- and four-legged robots (Figure 1), we found that legged locomotion can be achieved through extremely simple motor control by exploiting morphological properties and body dynamics (e.g. elasticity, rigidity, body weight distribution, and passive joints). An important discovery is that the exploitation of body dynamics can result in significant behavioral variations that are required for autonomous adaptive systems. It was shown that energy efficient self-stabilizing gait patterns can be generated by implementing passive elastic joints in the legged robots (Iida, et al., 2005; Iida et al., 2006).



Figure 1 Representative robotic platforms: (a) Stumpy, (b) Puppy, (c) BioLeg I, (d) J-Walker, and (e) Mini-Dog.

Body dynamics also significantly influences sensory stimulation and neural information processing. Through the interaction with the environment, i.e. through the robot's own dynamics, patterns of sensory stimulation are induced in different sensory channels. If these patterns originate from attractor states of the robot (e.g. from a particular gait pattern), the nature of these attractors will be reflected in the structure of the sensory information (Iida and Pfeifer, 2006). Thus, because these attractor states are the reflection of physically meaningful interactions (e.g. running forward at certain hopping height), by analyzing this sensory information, we as observers, but also the robot itself, can acquire a kind of basic type of symbol, i.e. discretely identifiable entities within a continuous sensory information space.

Although we are still in a nascent stage of our exploration, a further understanding of body dynamics will provide significant insights into adaptive behaviors of embodied systems. By implementing mechanisms for regulating body dynamics and more sophisticated sensory-motor coordination, we will be able to develop a set of interesting physical system-environment interactions on top of which a form of high-level cognition could be developed.

References

1. Brooks, R. A. (1991). Intelligence without reason, Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91), Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, 569-595.
2. Pfeifer, R. & Scheier, C. (1999). *Understanding Intelligence*, MIT Press: Cambridge MA.
3. Iida, F., Pfeifer, R., Steels, L., and Kuniyoshi, Y. (Eds.) (2004). Embodied artificial intelligence. Lecture Notes in Computer Science, Vol. 3139, Springer, ISBN: 3-540-22484-X.
4. Dickinson, M. H., Farley, C. T., Full, R. J., Koehl, M. A. R., Kram, R., and Lehman, S.: How animals move: An integrative view, *Science*, Vol 288, 100-106, 2000.
5. Collins, S. H., Ruina, A., Tedrake, R. L., Wisse, M.: Efficient bipedal robots based on passive-dynamic walkers, *Science* (18 February 2005), V307, p.p.: 1082-1085, 2005.
6. Iida, F., Gomez, G. J., and Pfeifer, R. (2005). Exploiting body dynamics for controlling a running quadruped robot. Proceedings of the 12th Int. Conf. on Advanced Robotics (ICAR05). July 18th-20th, Seattle, U.S.A., 229-235.
7. Iida, F., Minekawa, Y., Rummel, J., and Seyfarth, A. (2006). Toward a human-like biped robot with compliant legs, *Intelligent Autonomous Systems – 9*, Arai, T. et al. (Eds.), IOS Press, 820-827.
8. Iida, F. and Pfeifer, R. (2006). Sensing through body dynamics. *Journal of Robotics and Autonomous Systems*, (in press).

Interactive intent imitation for humanoid robots based on dynamic attention prediction and control

Tetsunari Inamura^{*†}, Naoki Kojo[†], Kazuyuki Sakamoto[†], and Masayuki Inaba[†]

^{*}National Institute of Informatics, Japan

[†]Departments of Mechano-Informatics, The University of Tokyo, Japan

[inamura | kojo | sakamoto | inaba]@jsk.t.u-tokyo.ac.jp

Introduction

Intent imitation is a higher concept than simple imitation, such as copying of motor command. With intent imitation, robots have to recognize a users' intent and modify the original motion patterns to achieve the desired purpose taking into consideration the difference in the physical conditions between humans and humanoid robots. It is difficult to acquire and describe users' intentions by only observations. We focus on a dialogue for the purpose of recognizing tasks or attention points, which is basic and important for acting as most of us do in daily-life. Conventional research has focused on dancing or behavior in a toy-world, which do not need to take into consideration the relationship between a humanoid robot's body and the surrounding environmental objects, because the attention points in the daily-life environment are difficult to deal with. We propose a framework where humanoid robots imitate humans' behaviors not only from the viewpoint of motor commands but also for the purpose of tasks by prediction and control of attention points.

In this paper, we propose an interactive imitation learning mechanism from the viewpoint that a robot should focus on human motions, changes in the surrounding environment that effect its motions, and the attention points. The interactive learning mechanism enables robots to develop purposive behavior with a combination of the taught attention points.

Mutual association model of sensorimotor pattern and attention

We propose a mutual association model in which motor commands, sensory data, and attention points could be mutually recalled as time advances. Using this framework, robots could predict the desired effect of motor actuation and modify the original motion taking into consideration the differences between humans and humanoid robots. To realize the temporal mutual association model, we have adopted a proto-symbol space method^[1] based on the continuously hidden Markov model (CHMM) shown in Figs.1 and 2. With this method, humanoid robots can recognize a human's behavior, abstract the motions into a symbol representation, and generate example motions for imitation.




Fig. 1: Continuous HMM for mutual association model of sensorimotor and attention points.

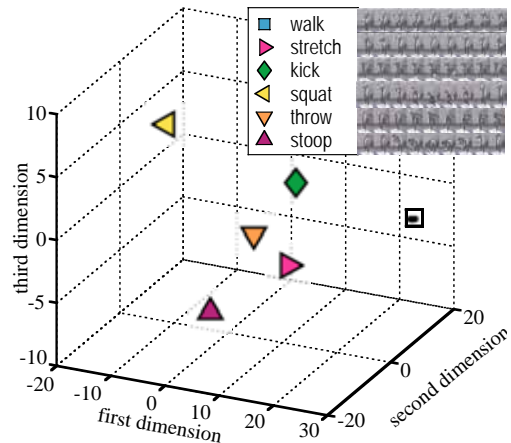


Fig. 2: Proto-symbol space that converts temporal sensorimotor pattern into symbol representation.

Definition and usage of “attention points”

In this paper, the attention points represent the target factors of imitation, in other words, the primitive intent. There are many imitation factors for humanoid robots, such as joint trajectories, relationships between self-body and target objects, and gaze points of cameras. Conventional researches on robotic imitation have treated the trajectories and self-behaviors of motions from simple things like dancing. In contrast, we focus on object handling to achieve daily-life tasks. We adopted some kinematic constraint conditions to be used as the attention points, such as a constraint on the position/posture of the hands and the relative position of both hands (Fig. 3). The attention points are instructed by humans, and represented as an index number to be treated in the same way as other sensory patterns. For the representation, humanoid robots could recall desired attention points by observing a motion pattern.




Fig.3 Attention points (constraint conditions) for objective imitation

Experiments on a humanoid robot

We adopted a HRP2W^[2] as a humanoid robot platform for the interactive motion acquisition and objective behavior imitation. We have practiced teaching and generating the daily life behaviors to confirm the effectiveness of the proposed method.

During the teaching phase, pouring water into a glass, carrying a cup without spilling it, and cleaning a desk were selected and performed by a human. Joint angles for each behavior are measured by the motion capturing system. For the pouring behavior, the robot uses the restriction condition of a horizontal constraint. In the imitation phase, the robot could recognize the irregular pouring motion as taught from the original pouring motion using the symbol space as shown in Fig. 4. The humanoid robot then modified the performed irregular motion to satisfy the constraint condition.

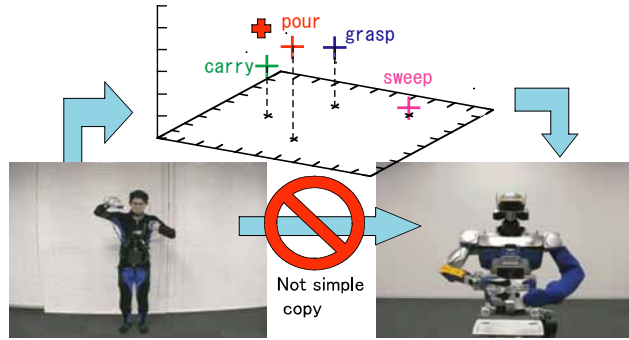


Fig. 4 Intent imitation using symbol space

Conclusion

In this paper, we focused on decision making from the attention points in order for humanoid robots to imitate humans' objective behaviors in daily life. In the current stage, humans have to teach the attention points; however, if humanoid robots can recognize the constraint conditions, users would be free of instructions on the attention points. In addition, such situations can be regarded as a huge step toward the realization of objective imitation for humanoid robots.

References

1. T. Inamura *et al*, *From Stochastic Motion Generation and Recognition to Geometric Symbol Development and Manipulation: Proc. of IEEE-RAS Int'l Conf. on Humanoid Robots* (2003).
2. T. Inamura *et al*, *HRP-2W: A Humanoid Platform for Research on Support Behavior in Daily life Environments: Proc. of 9th International Conference on Intelligent Autonomous Systems* (2006).

A Situated, Embodied and Dynamical Systems Approach to Understanding Learning and Memory

Eduardo Izquierdo-Torres, Inman Harvey
Centre for Computational Neuroscience and Robotics
University of Sussex, Brighton, U.K.
[e.j.izquierdo | inmanh]@sussex.ac.uk

In the days of traditional artificial intelligence cognition was seen as computation, memory was regarded simply as discrete storage space and learning would correspond to the changing of the contents of this space. We are privileged to live in more interesting times, where cognition is now widely understood (but not universally) as arising from the real-time interaction between bodies, brains and environments (Thelen & Smith, 1994; Harvey, 1996; Beer, 1997; Pfeifer & Scheier, 2002).

Learning is a fundamental aspect of cognitive activity because it allows organisms to adapt to the ongoing changes in their environments. Despite much progress in the new artificial intelligence, learning continues to be one of the activities whose mechanisms are understood as taking place ‘inside the brain’, with a particularly strong association to synaptic plasticity in neuronal networks (see for example Floreano & Urzelai, 2001). Although it is the case, in some organisms, where the dynamics of synaptic changes play a role in the adaptive modification of behavior, it is not necessarily the case that synaptic plasticity is either necessary or sufficient for learning behavior (see for example Phattanasri et al., submitted).

As cognitive activity, learning and memorization are behaviors that arise from the interaction between internal dynamics, body and environment. To be more precise, learning corresponds to a *change of behavior* that serves to adapt to the changing conditions of the environment. Accordingly, an act of memorization corresponds to the organisms’ ability to make decisions in its current environment taking into account the history of its previous interactions with the world.

Interestingly enough, in situated, embodied and continuous-time dynamical system agents interacting with an environment, changes in behavior according to past experiences are inevitable (Beer, 1997). This is because the different components that comprise the agent and environment interact in different time-scales; some at very fast time-scales (e.g. neurons) while others at much slower time-scales (e.g. morphodynamics), with the potential for interaction ranging over the continuum. In the extreme case of an agent with completely reactive internal dynamics (i.e. one that does not have internal state), the body’s intrinsic physical properties (e.g. inertia) in addition to its history of interactions with the environment (e.g. position in relation to other objects) allow it to be influenced by past experiences (Izquierdo-Torres & Di Paolo, 2005), and thus, to be capable of learning. In the less extreme cases, where some internal state is possible, the potential for the modification of behavior is simply richer.

According to this view, the question of interest shifts from what is learning and what is not to how the different time-scales of the components throughout the brain, body and environment interact to produce a particular learning behavior. Also of interest is to understand the mechanisms that allow the modification of behavior towards an improved adaptation to the changing environment. In this sense, the main question that our research asks is, how agents use past experiences to influence their future behavior. This we believe corresponds to a fundamental challenge that a

situated, embodied and dynamical systems approach to understanding cognition faces today.

In order to tackle these challenges we employ evolutionary robotics techniques (Harvey et al., 1997). The methodology is guided primarily by an attempt to make only the fewest possible assumptions about the sort of mechanisms an agent ‘requires’ to perform learning and memorization behavior while at the same time allowing it to exploit its embodiment and situatedness as much as possible.

Unfortunately, the study of learning has been impregnated by a computational view, where discrete tokens of presentation, recognition, reward are used in the experimental set ups. But not all studies of learning have been biased in this way, one particularly good example of a more ecological view of learning is the study of the development of social preferences in young animals for their parents or other stimuli ever since Konrad Lorenz’s vivid descriptions of avian imprinting (Lorenz, 1981).

Our current efforts are directed towards evolving agents that can perform imprinting-like learning behaviors (Izquierdo-Torres & Harvey, in press). In particular our research seeks to understand, the broader set of dynamical mechanisms that allow embodied and situated agents to (what an external observer would describe as) ‘record’ a feature from the environment within a continuum and later ‘make-use’ of this ‘stored-information’ to make a decision, with particular emphasis in understanding the role of the agent’s morphodynamics and situatedness in the generation of such behavior.

We are interested in these issues from an evolutionary perspective as well, so questions like: what sort of evolutionary pressure is needed to evolve agents with the capacity to retain a particular memory throughout its lifetime? This corresponds to certain aspects of learning irreversibility and the evolution of agents with critical learning periods, but our research is also concerned with understanding reversible learning.

References

1. Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. In *Artificial Intelligence*, **72**:173-215.
2. Floreano, D. and Urzelai, J. (2001). Evolution of plastic control networks. In *Autonomous Robotics* 11:311-317.
3. Harvey, I. (1996) Untimed and misrepresented: connectionism and the computer metaphor. In *AISB Quarterly*, no. 96, pp. 20—27.
4. Harvey, I., Husbands, P., Cliff, D., Thompson, A. and Jakobi, N. (1997) Evolutionary Robotics: the Sussex Approach. In *Robotics and Autonomous Systems*, **20**:205--224.
5. Izquierdo-Torres, E. and Di Paolo, E. (2005) Is an Embodied System Ever Purely Reactive? In M. Capcarrere et al (Eds.) *Proc. of the 8th European Conf. on Artificial Life*. p252-261. Springer-Verlag.
6. Izquierdo-Torres, E. and Harvey, I. (in press). Learning on a continuum in evolved dynamical node networks. To appear in *Artificial Life X: Proc. of the Tenth Int. Conf. on the Simulation and Synthesis of Living Systems*. MIT Press.
7. Lorenz, K. (1981). *The Foundations of Ethology*. Springer-Verlag.
8. Pfeifer, R. & Scheier, C. (2002) *Understanding Intelligence*. MIT Press, Cambridge, MA.
9. Phattanasri, P., Chiel, H.J. and Beer, R.D. (submitted). The dynamics of associative learning in evolved model circuits. To appear in *Adaptive Behavior*.
10. Thelen, E. & Smith, L.B. (1994). *A Dynamic Systems Approach to the Development of Perception and Action*. Bradford Books/MIT Press, Cambridge, MA.

The Sincere Gatekeeper

Terence Jardine

School of Information Technology & Electrical Engineering
University of Queensland, Brisbane, Australia
s4056847@student.uq.edu.au

As a “gatekeeper” to my own thoughts on artificial intelligence, I can only hope to be sincere if I wish to effectively communicate those ideas using abstractions. Even then, only those with a similar frame of reference shall find my words on the topic to be within their own understanding. “In communication, some gates are closed and others opened because of the assumptions and attitudes held by people – ‘gatekeepers’ – along the way.”¹ If a reader is satisfied with what is conveyed in a piece of writing, it is nearly certain that the piece of writing was read, processed and, perhaps, understood. Similarly, an information appliance, as a gatekeeper, shall find it important to satisfy natural minds (or artificial ones) that it is a sincere gatekeeper. This is especially the case if it were to serve an acceptable and productive purpose.

For an information appliance to be viewed as the personification of sincerity, which is likely to result in a great amount of trust being placed in any news sourced from it, the appliance's processes must be designed to rely on a similar frame of reference to that of the user. This would entail a quantitative analysis of a language, presumably the language in which information is to be conveyed. This might be done in conjunction with an artificial mind or be the sole work of some very ambitious natural ones (e.g. researchers). In any case, a possible step towards an easier transference of these frames of reference is the utilization of tools or mediums that are specific to the task. The identification of relatively fundamental concepts such as colour or heuristics for the categorization of shapes, both as parts of or not parts of a natural language, is essential for a greater appreciation of the abstract notions that might be part of one mind's frame of reference.

One who means what he says is a sincere person.

It is possible to extend this common idea to artificial minds for the purposes of emulating a frame of reference within an artificial mind. If a natural mind observes that an information appliance is mindful but also displays the behaviours of openness and conveyance of truth, it could be labelled sincere and, inherently, be personified. The information appliance must, first of all, communicate its apparently meaningful thoughts in order for this to happen.

Quantitative Analysis

Alternative representations can aid in the confirmation or summary of the quantitative nature of data. Proven examples are the chart and balance sheet. As is the case with the technical analysis of charts, indicators or trends are abstractions that may help productively assess or predict the influences or behaviours of dimensions that affect the value of a characteristic. In the case of a financial balance sheet, one needs to know how to assess the numbers in order to potentially identify a misconception about the state of a business. Knowing how to do this is an equivalent to knowing what to think or, in other words, which heuristics to use.

Representing the Frame of Reference

A statement may be deemed to be part of a message. Further, representations for the

meaning of a statement might include the relevant relationships that form part of the identifiable frame of reference. Where heuristics responsible for an artificial comprehension of meaning correspond with notions that are expressible in a language, an artificial mind could potentially intellectualize that language and also provide natural minds with the opportunity to confirm information gathered from it.

An example representation

<i>Words (language)</i>	<i>Heuristics (data mining technique)</i>	<i>Relationship(s) within frame of reference</i>
The An article ...
canary Instance of bird ...
is To be ...
yellow Instance of colour ...
. End of sentence ...
Comprehension of Language		Intellectualization of Language

Further analysis by the artificial mind

<i>Words (language)</i>	<i>Heuristics (data mining technique)</i>	<i>Relationship(s) within frame of reference</i>
Instance
of
colour
Comprehension of Language		Intellectualization of Language

Artificial Intelligence and User-Centered Design of Information Appliances

An information appliance requires some user-centered design for there to be reliable communication between it and its user. Indicating that the appliance is using an identifiable frame of reference, one that can be intellectualized by the user, is a key factor in whether or not artificial intelligence is suitable for an information appliance that channels a wide range of information.

References

1. Zimmerman, D. E. & Clark, D. G. *Guide to Technical and Scientific Communication* (Random House, Inc., 1987).

The Progress Drive Hypothesis

Frédéric Kaplan and Pierre-Yves Oudeyer
Sony Computer Science Laboratory Paris
kaplan@csl.sony.fr,oudeyer@csl.sony.fr

*« Instead of trying to produce a programme to simulate the adult mind,
why not rather try to produce one which simulates the child's? »
Alan Turing, Computing Machinery and Intelligence, Mind, 1950.*

Children learn to control their bodies, manipulate objects, interact with people, and utter words in just a few months. Their developmental trajectories are remarkably structured. Each new skill is acquired only when associated cognitive and morphological structures are ready. For example, children typically learn first to roll over, then to crawl and sit, and only when these skills are operational, do they begin to learn how to stand. Likewise, sudden transitions occur from stages characterized by seemingly insensitivity to input, to stages of extraordinary sensitivity to new data. Some pieces of information are simply ignored until the child is ready for them. In other words, children naturally become interested in different situations in a specific order of complexity in order to most effectively acquire complex knowledge and skills.

Most existing models in psychology or neuroscience fail to account for the open-ended nature of developmental processes. Development is either reduced to an innately defined maturational process controlled by some sort of internal clock, or, on the contrary, pictured as a passive inductive process in which the child or the animal simply catches statistical regularities in the environment. More generally, epigenetic developmental dynamics as a whole are rarely addressed as an issue as research tends to focus simply on the acquisition of particular isolated skills.

In the recent years, we have explored an alternative view using computational and robotic models. This view hypothesizes that a basic impulse to search, investigate and make sense of the environment and progress in learning lies at the origin of this remarkable force behind the developmental cascade. This driving force shapes environmental exploration in specific ways permitting efficient learning. Infants engage in exploratory activities for their own sake, to drive their own development, not simply to achieve an extrinsic goal. Of course, adults help by scaffolding their environment, but this is just help: eventually, infants decide by themselves what they do, what they are interested in, and what their learning situations are. Far from passive shaping, development has to be viewed as a fundamentally active and autonomous process. We call this view: “the progress drive hypothesis”.

We have designed a working prototype of an intrinsic motivation system permitting to measure and maximize learning progress. This system was shown to work in real-time in continuous spaces, both with a virtual agent set-up and with a real robotic set-up with continuous motor and/or perceptual spaces. Starting from an initially unstructured sensorimotor space, this system is capable of discovering regions characterized by different progress levels and to choose to focus on the ones which lead to maximal learning progress. These situations, neither too predictable nor too difficult to predict, are called “progress niches”. Progress niches are not intrinsic properties of the environment. They result from a relation between a particular environment, a particular embodiment (sensors, actuators, feature detectors and

techniques used by prediction algorithms) and a particular time in the developmental history of the agent. Once discovered, progress niches progressively disappear as they become more predictable.

We conducted several experiments using this system to control the development of an AIBO robot. First results were obtained in the domain of locomotion, discovery of object affordances and prelinguistic communication. What is fundamentally new in these experiments is that learning dynamics, embodiment and environmental factors are becoming controllable variables. One experiment can be conducted with the same learning system, but using a different body placed in a different environment. Likewise, the effect of small changes in the intrinsic motivation systems can be studied while keeping the embodiment and environmental aspects similar. This is a unique feature that no other method in psychology or neuroscience can approach.

At this stage, the existence of “progress drive” remains a hypothesis, but this type of embodied models helps refine our intuitions, suggest novel lines of empirical investigation with humans, and build concepts that shed a different light on children's fantastic learning capacities. Our hope is that it shall help us to recast the nature/nurture debate using a new scientific vocabulary; one that digs deeper than innate and learned.

In addition, this approach might also provide radically new techniques for building intelligent robots. Indeed, as opposed to the work in classical artificial intelligence in which engineers impose pre-defined anthropocentric tasks to robots, the techniques we develop endow the robots with the capacity of deciding by themselves which are the activities that are maximally fitted to their current capabilities. Our developmental robots autonomously and actively choose their learning situations, thus beginning by simple ones and progressively increasing their complexity. More than 50 years, Alan Turing prophetically announced that the child's mind would show us the way to artificial intelligence. It is now time to take this advice seriously.

References

Oudeyer, P-Y and Kaplan, F. (2006) The discovery of communication, *Connection Science*, 18 (2)

Kaplan, F. and Oudeyer, P-Y (in press) The progress drive hypothesis: an interpretation of early imitation, in Dautenhahn, K. and Nehaniv, C. *Models and Mechanisms of Imitation and Social Learning: Behavioural, Social and Communication Dimensions*, Cambridge University Press.

Acquirement of Body Schema using Intelligent EMG Prosthetic Hand

Ryu Kato, Alejandro Hernandez Arieta, Kojiro Matsushita, Hiroshi Yokoi

Department of Precision Engineering

The University of Tokyo, Tokyo, JAPAN

[kato | alex | matsushita | hyokoi]@prince.pe.u-tokyo.ac.jp

Introduction: In the field of cognitive science, one of importance issues is to reveal the mechanism of acquiring body schema and body image ^[1] for the purpose of understanding human intelligence. Among those researches, human hand is often focused as the expression of human intelligence because it is voluntarily controlled by human intention. Therefore, developmental robotics mainly use robots, which consist of head and hands, as platform to implement the mechanism of acquiring body schema as known as one of representative works of embodied artificial intelligence.

As for body schema in conventional physiology and psychology, it has been considered that body image changes little once the image is developed because of the fact that amputees sense severed limb (this phenomena is called as “phantom limb”)^[2]. However, Ramachandran reported that the body image has rapid plasticity ^[3]. Therefore, it is important to understand the process that old body schema changes to new one.

In this paper, we focus on how an amputee acquires a new body schema by using intelligent EMG (electromyographic signal) prosthetic hand. The characteristic of this prosthetic hand is the followings: the mechanism mimics human hand so that the prosthetic hand has similar appearance, size, and weight to a human hand and, moreover, the amputee can attach the prosthetic hand as its natural hand as possible as shown in Fig.1; EMG-to-motion classifier system as shown in Fig.2 controls motion of prosthetic hand by observing surface EMG signal as human intention so that the amputee achieves eight motions on the prosthetic hand.




Fig. 1 Five-fingered robot hand with interference driven finger and wrist based on parallel-wire mechanism (18 joint-DOFs and 13 control-DOFs).

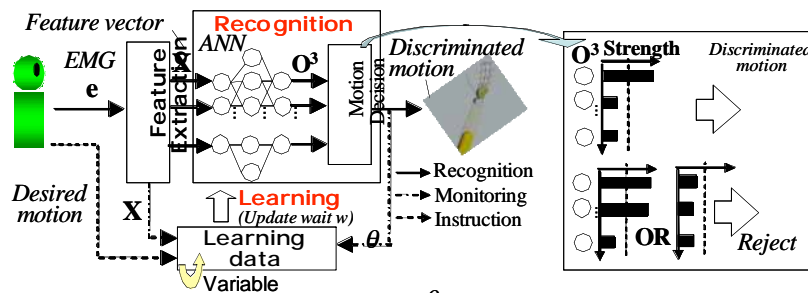


Fig.2 EMG-to-motion classifier system using on-line learning method

For the analysis, we applied f-MRI (functional magnetic resonance imaging) to normal person and amputee in order to record brain activities during use of the prosthetic hand.

Analysis: We conducted some experiments for brain function analysis. Normal person and amputee executed eight right-forearm motions on the prosthetic hand and their brain activities were recorded with f-MRI at certain: the 1st recoding was conducted when the experiment started (1st analysis); the 2nd recoding was conducted when the subjects used the prosthetic hand for three hours (2nd analysis); the 3rd recording was conducted only for the amputee when the subject used for three months (3rd analysis). Furthermore, for the comparison between natural hand and the prosthetic hand, we recorded brain activity when normal person executes three motions with his natural left-hand. Before each recording, sufficient training was done in order to achieve eight motions in best state at the point.

Figure 3 shows results of the f-MRI analysis. M_{max} represents the maximum number of motions with more than 80% of discriminating rate. In Fig.3(a), the f-MRI data shows that activation of primary motor area (M1) and primary somato-sensory area (S1) corresponding to the right forearm were widely growth. Moreover, M_{max} has increased from 3 to 6. Meanwhile, in the case of normal person in Fig.3 (b), it has shown that M1 and S1 were also widely growth. As for the comparison to the natural hand of normal person, it is observed that M1 and S1 were strongly activated and this result was similar to amputee's. These results mean the fact that the more strongly subjects recognize motions on prosthetic hand as their own, the more strongly M1 and S1 activate. Also, it seems that amputees can acquire new body image as original limb by get used to this prosthetic hand.

In addition, the prosthetic hand had no tactile and deep sensory feedbacks so that these results were acquired with only visual feedback. This indicates that human has the redundancy on sensory.




Fig. 3. Brain activities of amputee and normal person during use of prosthetic hand (row indicates training duration; column indicates right-side, top, and left-side views of brain image).

References

- [1] S. Gallagher, "Body Image and Body Schema: A Conceptual Clarification," *Journal of Mind and Behavior*, Vol. 7, pp. 541–554, 1986.
- [2] M. L. Simmel, *The conditions of occurrence of phantom limbs*, *Proceedings of the American Philosophical Society*(Vol. 102).
- [3] V. S. Ramachandran and D. Rogers-Ramachandran, *Synaesthesia in phantom limbs induced with mirrors*, *Proceedings of the Royal Society of London*,(Vol. 263, pp. 377–386, 1996).

What is a Perception-Action System? From Agency to Animality

Fred Keijzer, Daan Franken & Marc van Duijn
Department of Theoretical Philosophy
University of Groningen, Groningen, The Netherlands
[f.a.keijzer] [d.franken] [m.van.duijn] @rug.nl

What is cognition?

Cognitive science may have been around for fifty years now, but we still do not have very articulate ideas about what cognition is. “Often, the class of processes that we regard as cognitive is defined by ostension. Cognitive processes include such processes as perceiving, remembering, thinking, reasoning and ... language” (Rowlands, 2003, p.157). But while we may agree that these are good examples of cognitive processes, it is not clear what exactly makes them so.

The turn to situated and embodied interpretations of cognition, has now tied the notion of cognition more strongly to perception-action relations (Brooks, 1999; Clark, 1997; Hurley, 1998; Keijzer, 2001; Pfeifer & Scheier, 2001). Positive aspects of this change are that perception-action relations seem closer to a definite, physical interpretation, and also stay closer to the biological and evolutionary background of cognition (Allman, 1999).

Nevertheless, to answer what cognition might be, this shift will not be helpful when it suffers from the same problem: What is a perception-action system? Which criteria are we to use and what systems are to be included? Do heat-seeking missiles, dogs, animats, software agents all provide examples of perception-action systems?

Analysis

A possible explanation for this continuing demarcation problem is the following. Concepts like cognition, perception and action are concepts that derive from applying an intentional stance to a particular system. If such a stance is applied, then such notions are part of an interpretation of the system in intentional terms, making the system count as a cognitive one. However, we are free to apply the intentional stance to whatever we want, whether it is a falling stone or a human being. Thus, taking the intentional stance provides a way to interpret systems as being cognitive—or perceiving and acting—independent of any particular physical organization of the so described systems. At the same time, there is an opposing intuition that there must be something about the systems themselves that makes them cognitive or not: Humans and falling stones are very different in this respect. The question concerning what cognition is, is a reflection of this opposing tendency, and aims to delimit the set of cognitive systems on the basis of its physical or dynamical make up.

The result of these opposing tendencies is a double bind when it comes to answering the question what cognition is. The intentional stance provides us with an intuitively plausible cognitive domain, but makes it problematical what kind of physical systems could constitute such a domain. In contrast, one can start out from specific kinds of physical or dynamical systems, but this will be questioned from the start as it would always exclude domains that we currently think of as involving intelligence or cognition.

Moving on: Animality

A possible way to make progress is to bite this bullet, and to focus on particular kinds of systems as a starting point for thinking about cognition. Bearing the lessons from embodied cognition in mind, it is important that perception-action remains central, compared to, for example, neuroscience. Still, this focus on perception-action systems needs to be pushed further by turning to unquestioned natural representatives (Lyon, 2006) of perception-action systems: animals.

The question that needs to be asked is: What is it about animals that turns them into perception-action systems? A way to go then would be to research in much greater detail the physical and dynamical underpinnings of animal behavior (Keijzer, 2005). This would set up a research domain dealing with the set of processes involved for which we propose the label *animality* (Keijzer, 2006). While the notion of agency applies widely, animality is more restrictive: It refers to a particular physical setup exhibited in biological organisms, which is responsible for generating the agentic characteristics exhibited by these systems. What this particular physical setup amounts to remains to be seen, but existing work within embodied and situated cognition as well as neuroethology provides suitable handholds with concepts such as sensorimotor contingencies, subsumption and behavioral and neural pattern generation.

It is plausible that by focusing on, and investigating the details of natural perception-action systems, progress can be made with grounding cognition and perception-action as particular natural phenomenon. Modern biology made huge leaps by going for the details; there are good reasons to hope that the same can happen here.

And Artificial Intelligence?

This move would also be important for AI. A large scale empirical investigation of natural, embodied and situated intelligence will constitute a fact finding enterprise that ought to provide new insights and principles concerning intelligence which will in turn enrich AI research. As the life sciences have been an important source of facts, ideas and inspiration for ALife, the same should apply to the studies of natural intelligence.

References

1. Allman, H. *Evolving brains* (Scientific American Library, New York, 1999).
2. Brooks, R. A. *Cambrian Intelligence* (MIT Press, Boston, 1999).
3. Clark, A. *Being There* (MIT Press, Boston, 1997).
4. Hurley, S.L. *Consciousness in Action* (Harvard, Cambridge, MA, 1998).
5. Keijzer, F.A. *Representation and Behavior* (MIT Press, Boston, 2001).
6. Keijzer, F.A. Theoretical behaviorism meets embodied cognition. *Philosophical psychology*, 18(1), (2005) 123-143.
7. Keijzer, F.A. Differentiating animality from agency: Towards a foundation for cognition. *Proceedings of CogSci2006* (2006).
8. Lyon, P. The biogenic approach to cognition. *Cognitive Processing*, 7(1), (2006) 11-29.
9. Pfeifer, R. & Scheier, C. *Understanding Intelligence* (MIT Press, Boston, 1999).
10. Rowlands, M. *Externalism* (Acumen, Chesham: Acumen, 2003).

Implications of Embodiment and Situatedness on the Structure of Simulated Fish Schools

Hanspeter Kunz*, Charlotte K. Hemelrijk**

*AI Lab, Department of Information Technology, University of Zurich, Switzerland

**Centre for Ecological and Evolutionary Studies,

Univ. of Groningen, The Netherlands

hkunz@ifi.unizh.ch, C.K.Hemelrijk@rug.nl

Artificial Intelligence and collective behavior

Modern Artificial Intelligence seeks to understand intelligent behavior by finding its underlying principles. Inspiration is often drawn from the behavior of natural intelligent systems. As it is difficult to infer the internal mechanisms from observation only, often a synthetic approach, i.e. understanding by building, is chosen. The same also applies to collective behavior: Here, the overt group-level behavior emerges from the behavior at the individual level, which in turn depends on the internal mechanisms of the agents. The goal is to explain the (possibly) complex and typically adaptive behavior at the group-level by a small set of simple internal mechanisms at the individual level. Our intent here is to apply these concepts to schooling behavior in fish and to emphasize effects of embodiment and situatedness in the explanation of collective behavior.

Schooling, embodiment and individual differences

Schooling is one type of collective behavior and individual-based models (or agent-based simulations) have demonstrated that, via processes of self-organization, schooling may emerge in the absence of a leader or external stimuli. Instead, agents follow simple rules and use local information only. Despite the complexity of the collective behavior (schooling), the behavioral rules of the agents are simple: Fish avoid neighbors which are too close, align to those which are at intermediate distances and are attracted to others further away.

Typically, the individuals are modeled as completely identical and point-masses. While such a model is perfectly suited to explore general characteristics of schooling behavior it fails to explain the internal structure found in schools of natural fish, which are often composed of fish of different size, for example.

In a first model we study the effects of individual differences in body size on the structure of the school. We find that no matter the initial arrangement of the individuals, the fish segregate by size (see Fig. 1). Note, that this structure is an automatic consequence of the embodiment; the larger range of avoidance of large fish cause them to avoid others at larger distances than small fish do. Therefore, after some time, they end up at the periphery of the group.




Figure 1: End configuration a simulated fish school of large (black) and small (gray) fish. The passive effect of embodiment (large fish occupy more space and maintain a larger personal space) leads to emergent size segregation.

Schooling and situatedness

In agent-based models (of schooling) situatedness is usually implemented as follows: agents interact with each other within the distance of their sensory range. This overlooks the fact that nearby agents may obstruct the perception of agents further away. Further, because of the strong mutual attraction resulting from the high number of fish that occupy each others sensory ranges, schooling models might predict unrealistically high densities if groups are large (see Fig. 2A, dotted line).

Therefore, in a second model we reflect situatedness more realistically by taking into account that in a group many of the neighbors, although within sensory range, are nevertheless hidden behind others which are closer. We assume that a fish only interacts with the relatively low number of neighbors which are perceived directly, i.e. not occluded by others.

As a consequence the distance over which fish interact and also the number of neighbors perceived depends on the density (in a dense group perception is low). This provides a negative feedback to avoid the highly packed groups we got with our previous poorly situated model (see Fig. 2A, dots).

Another result of representing situatedness and embodiment in more detail is a higher perception distance at the front compared to that at the sides. This is a direct consequence of the elongated body shape of fish, and might have further implications on the behavior of the individuals.

Conclusion

Embodiment and situatedness are important concepts in studies of modern Artificial Intelligence. These studies often provide simpler explanations than pure cognitive approaches. We believe the same applies to agent-based simulations, where these concepts have not yet been applied systematically. We expect new and fruitful insights as regards the influence of embodiment and situatedness on collective behavior.

References

1. Pfeifer, R. & Scheier, C. *Understanding Intelligence* (MIT Press, Boston, 1999).
2. Hemelrijk, C. K. & Kunz, H. *Density distribution and size sorting in fish schools: an individual-based model* (Behavioral Ecology 16(1), 178-187, 2005)

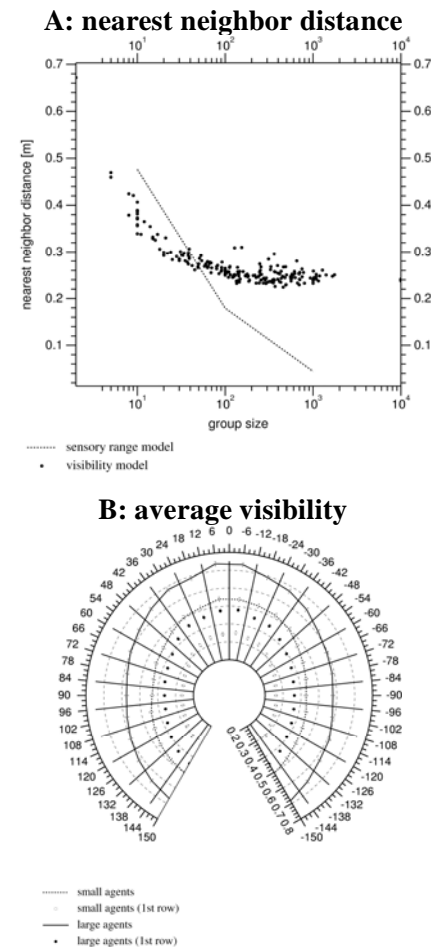


Figure 2: Panel A depicts the average nearest neighbor distance (nnd) vs. group size. In the “sensory range model” the nnd decreases to very low values (corresponding to high density) whereas for the “visibility model” it stabilizes. Panel B shows the average distance of visible agents vs. the direction.

High Performance Knowledge Systems: from PUFF to ICE

John Kunz

Center for Integrated Facility Engineering
Stanford University, Stanford CA, USA

kunz@stanford.edu

AI emerged from the 1956 Dartmouth Conference. Twenty-one years later, my colleagues and I reported the design and initial performance of what we think became the first application of AI to

INTERPRETATION: THE REDUCED TLC AND IVC INDICATE A MILD RESTRICTIVE DEFECT. THE FORCED VITAL CAPACITY, FEV1/FVC RATIO AND MID-EXPIRATORY FLOW ARE REDUCED AND THE AIRWAY RESISTANCE IS INCREASED, SUGGESTING MODERATE AIRWAY OBSTRUCTION. FOLLOWING BRONCHODILATION, THE EXPIRED FLOWS SHOW ONLY SLIGHT IMPROVEMENT. HOWEVER, THE DECREASE IN AIRWAY RESISTANCE INDICATES SOME REVERSIBLE COMPONENT. THE LOW DIFFUSING CAPACITY INDICATES A LOSS OF ALVEOLAR CAPILLARY SURFACE, WHICH IS SEVERE.



Figure 1: The early PUFF system created and explained diagnoses (Left). The current multi-user ICE process embeds models and analyses in a social context (Right).

be used in practice: the PUFF pulmonary function system (1). Today, easily recognizable descendants of that first “expert system” run in medical offices around the world, as do many other AI applications. My research now focuses on Integrated Concurrent Engineering (ICE), a computer and AI-enabled multi-participant engineering design method that is extremely rapid and effective (2). This brief note compares the early PUFF and the current ICE work, identifies symbolic representation and reasoning as some of the methods we used at the time of that first expert system that remain completely relevant today, and identifies some of the findings and lessons of the intervening years, fundamentally the move to model-based multi-discipline, multi-method, multi-agent systems.

Contrasts

Goal: the original PUFF system did diagnostic reasoning, or analysis. Recent ICE work does design, or synthesis, which is intellectually much more challenging.

Knowledge: the PUFF system used heuristic knowledge, coded as production rules, from a single domain. The ICE method uses multiple theoretically founded symbolic models, including the function or design intent, form or design choices and behaviors of integrated Product, Organization and Process models (2).

Reasoning: PUFF used automated production rule interpretation, while ICE is a mixed initiative method that includes manual synthesis and symbolic and numeric analysis of different integrated models.

Performance: A prospective (144 case) study measured PUFF performance at 89-96% agreement (SD 3.8 – 4.7) of the system to independent experts, while the experts had 92% (SD 1.6) mutual agreement. In multiple sessions, ICE reliably achieves a drop in information processing latency in excess of four orders of magnitude (> 2 days, which is high performance in practice, to ≤ 1 minute with $> 4\sigma$ reliability) and two orders of magnitude for design session duration (e.g., > 1 month to two hours; 1 year to four days) while maintaining or improving perceived design quality.

Explanation: PUFF uses text description of its diagnostic reasoning and conclusions. ICE uses graphic, tabular, text and verbal explanation of descriptive design content, predictions and their bases, explanations of prediction and design choice rationale and evaluation of design adequacy given requirements.

What remains the same

Good Knowledge Representation makes Reasoning (relatively) easy. PUFF is the unique domain area in my personal experience in which a pure rule-based knowledge representation was simultaneously appropriate for the domain expert, for the developing “knowledge engineers” and as a programming language for the system content. Scores of applications later, some successful and some not, it is my universal experience that excellent declarative representations are required to make the programming simple enough to both do and to maintain.

It remains difficult and crucial to define appropriate metrics of performance. In both our early and recent work, we spent almost time as much discussing potential performance metrics as we did doing the validation tests that verified the baseline performance of the existing practice and the measured performance of the knowledge system. We find that definition of those performance metrics to be a substantial contribution of the work we do. We now conclude that latency is fundamentally important performance metric of ICE, and underlying ICE design and operating mechanisms must work to reduce both its mean and variance to extremely low levels.

System performance can exceed human performance. Our measured system performance was very high. Astonishingly at the time, the statistical performance of our analysis system exceeded the performance of expert humans in practice because they had distractions, became weary and simply missed important features of the data. Our recent work finds the same result. A well-developed knowledge system has the capacity to perform at a higher level of reliability than expert in their practitioners – because the human experts are normally isolated, busy and readily distracted by disruptions, fatigue and boredom.

What is different

We now routinely depend on many knowledge sources and many computational and social methods. My post-PUFF application efforts always involved creating the new domain knowledge to enable scientifically founded *model*-based reasoning, where the emphasis was on the model, not the reasoning method. The models need to be good enough so that reasoning can be developed, explained, maintained and extended. Data, graphics and social engagement of stakeholders also provide power.

Users expect high performance knowledge processing to be social, involving many participants, and our systems have evolved to engage many stakeholders simultaneously using a variety of models, knowledge and data sources, reasoning and analysis methods. The change from single to multiple-user focus is my most recent and probably most significant change from the early days of AI.

References

1. Aikins, J.A., Kunz, J.C, Shortliffe, E.H., “PUFF: An expert System for Interpretation of Pulmonary Function Data,” *Computers and Biomedical Research* 16, 199-208 (1983).
2. Garcia, A., Kunz, J., Ekstrom, M., and Kiviniemi, A., “Building a Project Ontology with Extreme Collaboration and Virtual Design and Construction,” *Advanced Engineering Informatics*, Volume 18 # 2, pp 71-83, (2004). Also CIFE TR152, <http://cife.stanford.edu/online.publications/TR152.pdf>

How Morphology Affects Learning of a Controller for Movement

Nathan Labhart, Shuhei Miyashita
Artificial Intelligence Laboratory, Department of Informatics
University of Zurich, Switzerland
[labhart | miya]@ifi.unizh.ch

According to the “Principle of Ecological Balance” [1], the control mechanism (neural substrate) of an agent has to match its morphological complexity, and vice versa. Inspired by the centipede, where locomotion is achieved by controlling a number of two-legged body segments, we investigate in this research how the morphology affects the learning of a neural controller for similarly segmented artificial agents.

Like building blocks, simple two-wheeled modules are combined in order to form various morphologies, where the wheels are restricted in movement from to to simulate leg movement and exploit ground friction. A three-layered neural network is employed for learning a control mechanism for each wheel (or pair of wheels) such that the distance traveled within a certain time is maximised (Fig. 1).




Figure 1: 2-dimensional example configuration with controller

In this project, the agents are modeled using a physics simulator. A learning algorithm is implemented in order to train the network for a specific morphology; for example, the two configurations depicted in Figs. 2 and 3 both consist of the same number of modules, but the position of their center of mass is different (in the “hat”-shaped configuration, it is in the middle of its length, whereas in the “baseball cap” configuration it is closer to one end). This asymmetry leads to different friction on the ground, which is reflected in a different neural controller.

By systematically exploring various configurations, the influence of embodiment on performance (i.e. movement) is examined. The findings contribute to a better understanding of the relationship between morphology and control and give insights on how a control mechanism is learned depending on a specific morphology.




Figure 2: Hat-shaped configuration




Figure 3: Cap-shaped configuration

References

- [1] R. Pfeifer, C. Scheier: Understanding Intelligence. MIT Press, Cambridge (1999).
- [2] J. C. Bongard, R. Pfeifer: A Method for Isolating Morphological Effects on Evolved Behaviour. In: B. Hallam, D. Floreano et al (Eds.): Proceedings of the Seventh International Conference on the Simulation of Adaptive Behaviour (SAB2002), MIT Press, pp. 305–311 (2002).
- [3] S. Murata, E. Yoshida, A. Kamimura, H. Kurokawa, K. Tomita, S. Kokaji: M-TRAN: Self-Reconfigurable Modular Robotic System, IEEE/ASME Transactions on Mechatronics, 7 No.4, pp. 431–441 (2002).
- [4] A. Ijspeert, J-M. Cabelguen: Gait Transition from Swimming to Walking: Investigation of Salamander Locomotion Control Using Nonlinear Oscillators. In: H. Kimura, K. Tsuchiya, A. Ishiguro, H. Witte (Eds.): Adaptive Motion of Animals and Machines. Springer, Tokyo, pp. 177–188 (2005).
- [5] Y. Nakamura, T. Mori, S. Ishii: Natural policy gradient reinforcement learning method for a looper-like robot, Proceedings of AROB 11th '06, pp. 51

A Formal Definition of Intelligence for Artificial Systems

Shane Legg, Marcus Hutter
Dalle Molle Institute for Artificial Intelligence
Galleria 2, Manno-Lugano 6928, Switzerland
e-mail: [shane | marcus]@idsia.ch

A fundamental difficulty in artificial intelligence is that nobody really knows what intelligence is, especially for systems with senses, environments, motivations and cognitive capacities which are very different to our own. In our work we take a mainstream informal perspective on intelligence and formalize and generalize this using the reinforcement learning framework and algorithmic complexity theory. The resulting formal definition of intelligence has many interesting properties and has received attention in both the academic [4,5] and popular press [1,2].

Although there is no strict consensus among experts over the definition of intelligence for humans, most definitions share many key features. In all cases, intelligence is a property of an entity, which we will call the *agent*, that interacts with an external problem or situation, which we will call the *environment*. An agent's intelligence is typically related to its ability to succeed with respect to one or more objectives, which we will call the *goal*. The emphasis on learning, adaptation and flexibility common to many definitions implies that the environment is not fully known to the agent. Thus true intelligence requires the ability to deal with a wide range of possibilities, not just a few specific situations. Putting these things together gives us our informal definition: *Intelligence measures an agent's general ability to achieve goals in a wide range of environments.*

To formalise this we combine the extremely flexible reinforcement learning framework with algorithmic complexity theory. In reinforcement learning the agent sends its *actions* to the environment and receives *observations* and *rewards* back. The agent tries to maximise the amount of reward it receives by learning about the structure of the environment and the goals it needs to accomplish in order to receive rewards. To denote symbols being sent we will use the lower case variable names o , r and a for observations, rewards and actions respectively. The process of interaction produces an increasing history of observations, rewards and actions, $o_1 r_1 a_1 o_2 r_2 a_2 o_3 r_3 a_3 \dots$. The agent is simply a function, denoted by π , which is a probability measure over actions conditioned on the current history, for example, $\pi(a_3 | o_1 r_1 a_1 o_2 r_2 a_2 o_3 r_3)$. How the agent generates this distribution is left open, for example, agents are not required to be computable.

The environment, denoted μ , is similarly defined: $\forall k \in \mathbb{N}$ the probability of $o_k r_k$, given the current history is $\mu(o_k r_k | o_1 r_1 a_1 o_2 r_2 a_2 \dots o_{k-1} r_{k-1} a_{k-1})$. As we desire a general definition of intelligence for arbitrary systems, our space of environments should be as large as possible. An obvious choice is the space of all probability measures, however this causes serious problems as we cannot describe some of these measures in a finite way. The solution is to require the measures to be computable. This allows for an infinite space of possible environments with no bound on their complexity. It also permits environments which are non-deterministic as it is only their distributions that need to be computable. Additionally we bound the total reward to be 1 to ensure that the future value $V_{\mu}^{\pi} = E \sum_{i=1}^{\infty} r_i$ is finite. We denote this space E .

We want to compute the general performance of an agent in unknown environments. As there are an infinite number of environments, we cannot simply take an expected value with respect to a uniform distribution. If we consider the agent's perspective on the problem, it is the same as asking: Given several different hypotheses which are consistent with the observations, which hypothesis should be considered the most likely? This is a fundamental problem in inductive inference for which the standard solution is to invoke Occam's razor: *Given multiple hypotheses which are consistent with the data, the simplest should be preferred*. As this is generally considered the most intelligent thing to do, we should test agents in such a way that they are, at least on average, rewarded for correctly applying Occam's razor. This means that our a priori distribution over environments should be weighted towards simpler environments.

As each environment is described by a computable measure, we can measure the complexity of these in the standard way by considering their Kolmogorov complexity. Specifically, if U is a prefix universal Turing machine then the Kolmogorov complexity of an environment μ is the length of the shortest program on U that computes μ . We can now define the *universal intelligence* of an agent π to be its expected performance,

$$Y(\pi) := \sum_{\mu \in E} 2^{-K(\mu)} V_{\mu}^{\pi}$$

By considering V_{μ}^{π} for a number of basic environments, such as small Markov decision processes, and agents with simple but very general optimization strategies, it is clear that Y correctly orders the relative intelligence of these agents in a natural way. A very high value of Y would imply that an agent is able to perform well in many environments. The maximal agent with respect to Y is the theoretical AIXI agent which has been shown to have many strong optimality properties, including being self-optimizing in all environments in which this is at all possible for a general agent [3].

Universal intelligence spans simple agents right up to super intelligent agents, unlike the Turing test. Furthermore, the Turing test cannot be fully formalized as it is based on subjective judgments. An even bigger problem is that the Turing test is highly anthropocentric. Universal intelligence does not have these problems as it is formally specified in terms of the more fundamental concept of complexity.

References

1. C. Fievet. Mesurer l'intelligence d'une machine. In *Le Monde de l'intelligence*, volume 1, pages 42-45, Paris, November 2005. Mondeo publishing.
2. D. Graham-Rowe. Spotting the bots with brains. In *New Scientist magazine*, volume 2512, page 27, 13 August 2005.
3. M. Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, Berlin, 2004. 300 pages, www.idsia.ch/~marcus/ai/uaibook.htm.
4. S. Legg and M. Hutter. A universal measure of intelligence for artificial agents. In *Proc. 21st International Joint Conf. on Artificial Intelligence (IJCAI-2005)*, Edinburgh, 2005.
5. S. Legg and M. Hutter. A formal measure of machine intelligence. In *Proc. Annual machine learning conference of Belgium and The Netherlands (Benelearn-2006)*, Ghent, 2006.

Towards programmable self-assembly of microrobot ensembles

Lukas Lichtensteiger
Department of Chemistry and Chemical Biology
Harvard University, Cambridge, USA
llichtensteiger@gmwgroup.harvard.edu

Intelligence and the need for adaptive bodies. Today's robotic systems still predominantly have fixed (and usually rigid) body structures that are specified by a designer and allow for very little subsequent adaptation. This is in stark contrast to biological systems where bodies can grow and adapt to changes in internal or external conditions, can self-repair and even self-reproduce. Such life-like properties are essential for autonomous survival in open real world environments, and there is increasing evidence that the ability to continuously adapt body morphology in combination with neural control architecture is also an important factor for the emergence of intelligence^{1,2,3}. Recently, several research groups have started to bridge this gap between non-living and living matter by building robots consisting of large ensembles of relatively autonomous, modular building blocks inspired by biological cells or molecules. By simply re-arranging modules these robots can then adopt virtually any desired shape^{4,5,6,7,8}. However, constructing robotic systems from very small modules such as molecules faces the problem that "programming" these modules to assemble into a specific structure and to perform a desired function is very hard. On the other hand, systems that consist of relatively large, conventional robotics modules can easily be programmed but have severe technological limitations in terms of their structural flexibility.

Programmable dynamical self-assembly of microelectro-chemical robotic compound structures. Self-assembly, i.e., using self-organization for assembling larger structures from individual building blocks, is one of the most important techniques used in biology for development of complex functional structures. It is also widely used in chemistry as a bottom-up assembly and manufacturing method⁹. One of its inherent advantages is that the target structures are necessarily thermodynamically stable, i.e., self-assembly tends to produce structures that are relatively defect-free and self-healing. Consequently, in order to address the above challenges for building robots with adaptive bodies, our approach tries to combine the ease of programmability of microelectronic devices with the structure-forming capabilities of self-assembling chemical building blocks. We hope that this approach will allow us to produce multi-module robotic systems featuring certain life-like properties such as the ability to autonomously grow and adapt their body morphology.

References

1. Brooks R. A. (2001). *The relationship between matter and life*. Nature, **409**:409–411.
2. Pfeifer R. and Scheier C. (1999). *Understanding Intelligence*. MIT Press, Cambridge, MA.
3. Lichtensteiger, L. (2004). *Bodies That Think Quickly and Learn Fast: On the Interdependence of Morphology and Control for Intelligent Behavior*. PhD Thesis, University of Zurich, Switzerland.

4. Murata, S. et al. (1999). *Self-reconfigurable mechanical structures*. Proc. SPIE International Symposium on Intelligent Systems and Advanced Manufacturing, Boston, Sept. 1999.
5. Yim, M., Duff, D., Roufas, K. (2000). *PolyBot: a Modular Reconfigurable Robot*. Proc. of IEEE Intl. Conf. on Robotics and Automation (ICRA2000), IEEE Press, NJ.
6. Ichikawa, S. et al. (2003). *Emerging of Group Formation – Morphological Configuration of Multi-Robot System*. In: Hara, F. and Pfeifer, R (eds.), *Morpho-functional Machines: The New Species. Designing Embodied Intelligence*. Springer, Tokyo.
7. Støy, K., Shen, W.-M., and Will, P. (2003). *Implementing Configuration Dependent Gaits in a Self-Reconfigurable Robot*. In: Proceedings of the 2003 IEEE International Conference on Robotics and Automation (ICRA'03), Tai-Pei, Taiwan, September 16-18.
8. Zykov V., Mytilinaios E., Adams B., Lipson H. (2005). *Self-reproducing machines*, *Nature*, **435**:163-164
9. Whitesides, G., and Grzybowski, B. (2002). *Self-Assembly at All Scales*. *Science* **295**:2418-2421.

Information Processing as Embodied Interaction

Max Lungarella¹ and Olaf Sporns²

¹Dept. of Mechano-Informatics, The University of Tokyo, Tokyo, Japan

²Dept. of Psychological and Brain Sciences, Indiana University, Bloomington, USA

maxl@isi.imi.i.u-tokyo.ac.jp, osporns@indiana.edu

The brain-as-a-computer metaphor has been the dominant mode of thought of the past 50 years of research in AI. It is a powerful metaphor, which also thanks to the incredible increase of computing power, has led to many important theoretical advances and applications – ranging from natural language processing to the indexing techniques used in search engines. While it is tempting to think of the brain as some kind of super-efficient “information processor” (at least at some level of abstraction), the question remains if the processing is *exclusively* done by it (as the metaphor suggests). Is information just “out there”, an infinite tape ready to be loaded and processed by the cognitive machinery located between our ears? Could it be that embodied interaction with the real world plays an important (often neglected) role in the computational process itself? Intelligence is physically grounded after all, and emerges and develops in a body situated in a social context! In this view, embodied systems such as humans and robots are not exposed to a massive flow of unstructured sensory information, but because of their behaviors and their morphologies, their sensory inputs have predictable structure and the multimodal data entering their “cognitive” architectures possess statistical regularities. The major challenge posed by the need to process huge amounts of information in real time is thus simplified. But how does embodiment actively support and promote intelligent information processing and the structuring of sensory inputs? How can the dynamics of the physical interaction between an embodied system and its surrounding environment be quantified?

To provide answers to such questions and move towards a theory of “embodied information processing”, it is important to identify some organizing principles aimed at capturing heuristics and design ideas in a concise way. Our view is expressed by a fundamental principle of adaptive behavior: “embodied agents do not passively absorb sensory information from their surrounding environment, but through their actions on the environment; they proactively structure, select, and exploit information.” This principle points to the critical role of the mutual interaction between body, control structure, and environment to induce statistical regularities and information structure at various levels of the control architectures. Its implications for neural information processing are at least three-fold. First off, the structure of the sensory inputs and the one of the neural dynamics are interdependent and cannot be separated. Not only does learning heavily rely on both of them, but also a higher matching between perceived environmental structure and neural connectivity is achieved. The second implication is that sensorimotor activity and embodiment can generate statistical (intra and inter-modal) structure in the sensory inputs and should therefore be exploited, e.g. for design. It can be shown, for instance, that maximizing information structure is highly effective in generating coordinated behavior in a simulated creature subject to behavioral and information theoretic cost functions [4]. Third, part of the “processing” can be taken over by the dynamics of the agent-environment interaction. It follows that only sparse (but well-timed) neural control needs to be exerted if the self-regulating and stabilizing properties of the natural dynamics are exploited – whereby the specific morphology of the body and its natural

dynamics shape the repertoire of preferred movements (without requiring too much control).

In order to quantify the information structure contained in sensory and motor channels of embodied systems, it is important to develop methodologies for their analysis [1,2]. One potential approach is to use tools from information theory to quantify the gain in information (or processing power) that is specifically due to embodiment. We exemplified how this might be achieved by applying such tools to a robot capable of saliency-based attentional behavior [1], a ring-like structure [3], humanoid, wheeled and quadrupedal robots [5] as well as the evolution of sensorimotor behavior in a simple creature [4]. We found that effectively coordinated motor activity can lead to patterns of decreased entropy, and increased mutual information, integration and complexity in the sensory data. We also introduced and discussed methods for identifying, quantifying, and classifying (potentially invariant) local features and global patterns that emerge from sensorimotor dynamics and embodied interaction. These methods included measures of complexity and integration as well as a set of novel tools based on spectral and wavelet analysis. To extract undirected and directed informational exchanges between coupled systems, such as brain, body, and environment, we also studied a set of measures aimed at detecting asymmetric couplings and directional information flow between coupled systems. We applied one of these measures (e.g. transfer entropy), and found that patterns of non-causal as well as causal relations exist which can be mapped between a variety of sensor and motor variables sampled by morphologically different robotic platforms [5].

The conceptual view of *perception as an active process* has gained much support in recent years. Our work not only provides additional evidence for this view, but also suggests a quantitative link between embodiment and information. Perception cannot be treated as a purely computational problem that unfolds entirely *within* a given information processing architecture. Instead, perception is naturally embedded within a physically embodied system, interacting with the real world. It is the *interplay* between physical and information processes that gives rise to perception. Our research aims at paving the way towards a formal and quantitative analysis of the specific contributions of embodiment to perceptual processing through the active generation of structure in sensory stimulation. The idea of inducing information structure through physical interaction with the real world has important consequences for understanding and building intelligent systems, and highlights the fundamental importance of morphology, materials, and dynamics.

References

1. Lungarella, M., Pegors, T., Bulwinkle, D. and Sporns, O. (2005). Methods for quantifying the information structure of sensory and motor data. *Neuroinformatics*, 3(3):243—262.
2. Lungarella, M., Ishiguro, K., Otsu, N. and Kuniyoshi, Y. (2006). Methods for quantifying the causal structure of bivariate time series. *Int. J. of Bifurcation and Chaos* (in press).
3. Pitti, A., Lungarella, M. and Kuniyoshi, Y. (2005). Quantification of emergent behaviors induced by feedback resonance of chaos. *Recent Advances in Artificial Life: Advances in Natural Computation*, 3(15):199—213.
4. Sporns, O. and Lungarella, M. (2006). Evolving coordinated behavior by maximizing information structure. *Proc. of 10th Int. Conf. on Artificial Life* (accepted for publication).
5. Sporns, O., Karnowski, J. and Lungarella, M. (2006). Mapping causal relations in sensorimotor networks. *Proc. of 5th Int. Conf. on Development and Learning* (accepted for publication).

The Promises and Challenges of Socially Assistive Robotics

Maja J Matarić and Adriana Tapus
Interaction Lab, University of Southern California
941 West 37th Place, RTH 407, MC 0781, Los Angeles, CA, USA
{mataric, tapus}@usc.edu

Introduction

Human-Robot Interaction (HRI) for *socially assistive* applications is a growing and increasingly popular research area at the intersection of robotics, health science, psychology, social science, and cognitive science. Assistive robotics has the potential to enhance the quality of life for large populations of users. In response to the rapidly growing elderly population, a great deal of research attention has been dedicated toward the study and development of robot pets and companions aimed at reducing stress and depression [6, 11]. Individuals with physical impairments and those in rehabilitation therapy are also potential beneficiaries of socially assistive technology, both for improved mobility [13] and for improved outcomes in recovery. Finally, individuals with cognitive disabilities and developmental and social disorders (e.g., autism [2, 12]) constitute another growing population that could benefit from assistive robotics in the context of special education, therapy, and training. An effective socially assistive robot must understand and interact with its environment, exhibit social behavior, and focus its attention and communication on the user in order to assist in achieving specific goals.

Socially Assistive Robotics

A defining property of *socially assistive robotics* is its focus on the *social interaction*, rather than the *physical interaction* between the robot and the human user. This is a challenging domain because the robots are interacting with vulnerable users, resulting in ethical issues. Our work addresses a new niche: contact-free social robotic assistance. The physical embodiment of the robot plays a key role in its socially assistive effectiveness. It is well established that people attribute intentions, goals, emotions, and personalities to even the simplest of machines with life-like movement or form [9]. Because of this combination of properties, embodiment constitutes a key means of establishing human-robot interaction, specifically with the goals of having the user respond to the robot and become engaged in a goal-driven interaction with it. Some social robotics research has already been performed [1, 3, 6, 8]. However, social robotics has not yet tackled the complex challenges of assistive tasks, where the overall goal is to achieve measurable progress toward improved health, education, or training. Socially assistive robotics, our field of research focus, presents a new paradox: the goal of retaining user engagement can be in conflict with the health/training/education goals. The robot's physical embodiment, its physical presence, and its shared context with the user, all play fundamental roles in time-extended, sustained, goal-driven interactions in assistive domains.



Figure 1: Our therapist robot

As part of physical presence, the appearance of the robot is one of the important issues in human-robot interaction; it must be appropriately matched to the robot's cognitive and interactive capabilities. The more human-like the robot appears, the higher the expectations of people interacting with it are. In socially assistive robotics, *believability* plays a more important role than realism. Hence, a child-like appearance or anthropomorphic but not highly realistic appearance is typically more suitable for assistive tasks. Our therapist robot, shown in Figure 1, is designed with this philosophy in mind; even more standard mobile robots have already been successfully applied in our work toward therapist robots that assist, encourage and socially interact with users in the context of convalescence, rehabilitation, and education [4, 5, 7, 10]. Our work to date shows that the robot's personality and its social competence, expressed through body language and verbal interaction, are

likely more important than its physical appearance. Our robots are equipped with a basic set of task-oriented and social basic behaviors that explicitly express their desires and intentions physically and verbally [10].

Summary

Our work to date demonstrates the promises of socially assistive robotics, a new research area with large horizons of fascinating and much needed research. Our ongoing efforts are aimed at developing effective embodied assistive systems, and extending our understanding of human social behavior. Hence, even as socially assistive robotic technology is still in its early stages of development, the next decade promises assistive robotic platforms and systems that will be used in hospitals, schools, and homes in therapeutic programs that monitor, encourage, and assist their users. It is therefore important that potential users, well beyond the technical community, become familiar with this growing technology and help shape its development toward its intended positive impact on numerous lives.

References

1. Breazeal, C. *Designing Sociable Robots*, MIT Press, Boston 2002.
2. Davis, M., Dautenhahn, K., Nehaniv, C., Powell, S. *Towards an interactive system eliciting narrative comprehension in children with autism: A longitudinal study*, Proc. 3rd Cambridge Workshop on Universal Access and Assistive Technology (CWUAAT), 6th Cambridge Workshop on Rehabilitation Robotics, "Designing Accessible Technology" Fitzwilliam College, University of Cambridge, 10 - 12 April 2006.
3. DiSalvo, C., Gemperle, F., Forlizzi, J. and Kiesler, S. *All Robots Are Not Created Equal: Design and the Perception of Humanoid Robot Heads*, Proceedings Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques. London, England 2002.
4. Eriksson, J. and Matarić, M., J. *Hands-Off Assistive Robotics for Post-Stroke Arm Rehabilitation*, Proceedings IEEE International Conference on Rehabilitation Robotics (ICORR-05), Chicago, IL, 2005.
5. Feil-Seifer, D., and Matarić, M., J. *Socially Assistive Robotics*, Proceedings, IEEE International Conference on Rehabilitation Robotics (ICORR-05), Chicago, IL, 2005.
6. Fong, T., Nourbakhsh, I. and Dautenhahn, K. *A survey of socially interactive robots*, Robotics and Autonomous Systems 42 (3-4):143-166, 2003.
7. Gockley, R., and Matarić, M., J. *Encouraging Physical Therapy Compliance with a Hands-Off Mobile Robot*, Proceedings of the First International Conference on Human Robot Interaction (HRI-06), Salt Lake City, Utah, 2006.
8. Kiesler, S. and Goetz, J. *Mental Models and Cooperation with Robotic Assistants*, Proceedings, Conference on Human Factors in Computing Systems (CHI), Minneapolis, Minnesota, 2002.
9. Reeves, B. and Nass, C. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press, New York, NY. 1998.
10. Tapus, A., and Matarić, M., J. *Introvert vs. Extrovert Hands-Off Robot for Post-Stroke Rehabilitation Therapy*, Proceedings, 10th International Symposium of Experimental Robotics (ISER-06), Rio de Janeiro, Brazil, 2006.
11. Walton, M., *Meet PARO, The Therapeutic Robot Seal*, CNN, 2003.
12. Werry, I., Dautenhahn, K., *The Emergent Nature of Human-Robot Interaction: An Experimental Study with Children with Autism*, Modeling Biology, MIT Press, 2005.
13. Yanco, H., A. *Wheelesley, A Robotic Wheelchair System: Indoor Navigation and User Interface*, in Lecture Notes in Artificial Intelligence: Assistive Technology and Artificial Intelligence, edited by V.O. Mittal, H.A. Yanco, J. Aronis, and R. Simspon. Springer-Verlag, 1998, pp. 256-268, 1998.

Designing Dynamical Stable Locomotors Across the Reality Gap - Introduction of Edutainment as Heuristic-Diversity Design -

Kojiro Matsushita, Hiroshi Yokoi, Tamio Arai
Departments of Precision Engineering
The University of Tokyo, Tokyo, Japan
[matsushita | hyokoi | arai-tamio]@prince.pe.u-tokyo.ac.jp

Introduction

In embodied artificial intelligence, the inter-dependence between morphology and controller is an important issue because the inter-dependence is considered as one of the most important factors [1] for achieving dynamically stable locomotion. One of the most successful of these applications was the work of Sims [2], in which artificial creatures were automatically designed within a three-dimensional physics simulation. The simulation generated a variety of locomotive creatures with unique morphologies and gaits, some of which have no analogy in the biological world. This suggested that the interdependence between morphology and control plays an important role in the evolution of locomotion. However, evolutionary design is still in its infancy: coupled evolution is conducted only in simulation and how to best represent morphology, controller, environment, and fitness function is not clear yet; differences between virtual and real worlds have not been elucidated so that results from the virtual world are not always transferable to the real world, especially in the case of dynamic systems, although work is focusing on this problem [3]. Therefore, the work of Lipson [4], who demonstrated automated manufacture of evolved simulated robots, constrained his system to static locomotion. Thus, it is important to make clear the constraints for transferring dynamical behavior from simulation to reality.

Coupled Evolution in Virtual World and Rapid Prototyping in Real World

We mainly focus on finding the necessary design components for dynamically stable locomotion and, to that end, applied two methodologies interdependently. One methodology is coupled evolution of morphology and controller in three-dimensional simulation. As the main characteristic, the evolutionary design produces a variety of functions and structures of robots (diversity) because the design knowledge is not limited by human bias. However, prescribing the possible morphologies, environments, and fitness functions for evaluating behavior are the most difficult issues, so that mostly simple robots are designed in this way. Meanwhile, another methodology is rapid prototyping in the real world. In this approach robots are constructed with plastic bottles and RC servo motors by connecting them together with glue and, therefore, robots are built with less technical difficulties such as economizing machining time and easy assembling so that it is relative easy to acquire several functional robots (e.g. legged locomotion) by utilizing human experience (heuristics). Both methodologies have their own advantages and disadvantages, which leads us to propose inter-dependence use of these two methods (fig.1) in order to extract the necessary design components for dynamically stable locomotion.

First of all, we conducted an edutainment course as a practical demonstration of the heuristic-diversity design. Fig.2 shows the results of this course – we have acquired a variety of locomotors such as legged, crawling, and jumping robots. Thus, some of the robots designed by students achieved functional locomotion and, then, these robots were analyzed in simulation for the purpose of clarifying which were the

effective components. As results of analysis, we have found that symmetrical design has much to contribute to dynamical locomotion. We propose a meta-method in which iterating between these two methodologies clarifies more of the design components for dynamically stable locomotion and enables robot designs to more easily cross the reality gap.

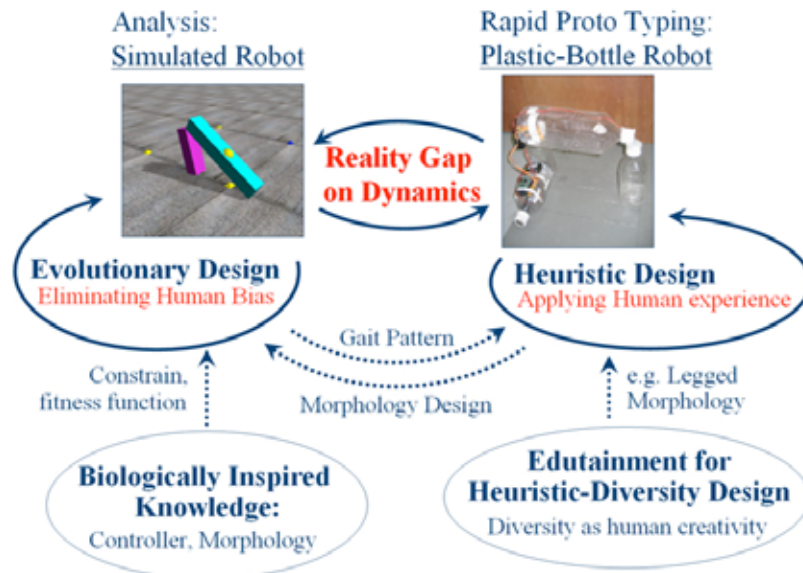


Fig.1 Conceptual design system for investigation of dynamically stable locomotion



Fig.2 Edutainment: Locomotors design by students

References

1. Pfeifer, R. & Scheier, C. *Understanding Intelligence* (MIT Press, Boston, 1999).
2. Sims, K. *Evolving virtual creatures* (Computer Graphics Annual Conference Proceedings, pp.43-50, 1994).
3. Jakobi, N. *Evolutionary robotics and the radical envelope of noise hypothesis*. (Adaptive Behavior, 6(2), pp.325-368, 1998)
4. Lipson, H. & Pollack, J. *Automatic design and manufacture of robotic lifeforms* (Nature 406, pp.974-978, 2000).

Technical Proposal for an Active Floating Element

Shuhe Miyashita, Maik Hadorn, and Peter Eggenberger Hotz
Department of Informatics
University of Zurich, Switzerland
[miya | hadorn | eggen]@ifi.unizh.ch

Introduction and Background

What is the essential difference between physical interaction and informative ordering when an autonomous and distributed system forms the morphology? In this research, we have been trying to realize complex informative ordering, exploiting physical interaction. Recent advances in robotics reveal the importance of autonomous self-construction and embodiment for building intelligent systems. While currently most robot construction and repair is performed manually [1], [2], this will be quite difficult when (a) the complexity of the systems exceeds a certain threshold, and (b) if these systems have to be truly adaptive. With conventional engineering hitting a complexity barrier it seems very useful to draw inspiration from natural systems, such as cells. Through natural evolution they have come up with many interesting solutions for some of the

problems that future robotics will have to deal with, like self-organization and adaptivity to changing environments, fault tolerance and self-repair, self-programming and self-replication, to name but a few. The objective

of this research is to achieve self-assembly and self-repair in a self-organized robotic system consisting of hardware, focusing especially on the shape, which is quite essential as a factor of aggregation but difficult to deal with as finite states. Also towards this end, the size of the individual modules must be reduced significantly (from dm to cm).





Fig. 1. left: experimental setting. **right:** a proposed floating element.

Proposed approach

Figure 1 shows a prototype of the proposed floating element. The total weight of this module, which can float on water is approximately 2.5g and the diameter is 3.5cm. The element consists of a vibrator for the actuator, an antenna which touches a ceiling, and a rod which goes into water. Therefore energy can be supplied constantly from the ceiling to the water via the vibrator. A magnet is attached to the bottom so that elements can attract or repel each other. One of the advantages of this model is that it replaces a mechanical connecting system by magnetic force and the repelling force by vibrator. This enables the system to become small and light in comparison to the model presented in previous researches [3], [4]. One of the important points that we emphasize on here is that this model tries to exploit several interactive mechanisms - in this case fluid dynamics as well as several physical-level-forces, such as magnetic force and mechanical interactions. As a result of several experiments, we noticed that the shape of the element plays an important role in the aggregation behavior of the system. We tried out several types of shapes - square, circle, and rounded square - and

found that the corner of the square shape acts as a potential energy barrier and prevents the elements from aggregation. That is, just by cutting the corner, the aggregation speed is increased. This research has shown that by taking into account the specific physics of an agent-environment interaction, control can be simplified, i.e. the computational requirements can be strongly reduced. This is a necessary prerequisite for studying artificial "organisms" composed of a large number of modules. We are providing a small but significant step towards massively modular self-reconfigurable robots. There are two major implications of this research. One is that it presents an approach towards tackling the "complexity barrier" in engineering, and on the other hand it is of theoretical importance because the concept of "morphological computation" - incorporating morphology and materials - provides a new way of conceptualizing computation.

References

1. T. Fukuda and Y. Kawauchi: Cellular robotic system (CEBOT) as one of the realizations of self-organizing intelligent universal manipulator, Proceedings of the 1990 IEEE Conference on Robotics and Automation, pp. 662-667 (1990).
2. S.Murata, E.Yoshida, A.Kamimura, H.Kurokawa, K.Tomita, S.Kokaji: M-TRAN: Self-Reconfigurable Modular Robotic System, IEEE/ASME Transactions on Mechatronics, **7** No.4, PP.431-441 (2002).
3. M. Shimizu, T. Kawakatsu and A. Ishiguro: "A Modular Robot That Exploits a Spontaneous Connectivity Control Mechanism", In The International Conference on Intelligent Robots and Systems 2005, pp.2658-2663 (2005).
4. P. J. White, K. Kopanski, H. Lipson: Stochastic Self-Reconfigurable Cellular Robotics, IEEE International Conference on Robotics and Automation 2004, pp. 2888-2893 (2004).

Does a synesthetic mechanism aid robot's language learning?

Yukie Nagai

Applied Computer Science, Faculty of Technology, Bielefeld University
P.O. Box 100 131, D-33501 Bielefeld, Germany
yukie@techfak.uni-bielefeld.de

Synesthesia is an involuntary experience that the stimulation of one sensory modality causes a perception in one or more different senses¹. A few adults are reported to have the capability whereas all newborns are synesthetic². They do not keep sensations separate from one another, but rather mix them and respond to the total amount of energy. Ramachandran and Hubbard³ suggest that synesthesia may have evolved language. When people perceive an object in the vision, they evoke an equivalent perception in the auditory sense, which could be a prototype of symbols. Inspired by this idea, I suggest that a synesthetic mechanism helps a robot to learn language. Compared to the existing statistical approach to language learning^{4,5}, a robot that can detect equivalent relationships between different senses will efficiently acquire language.

On the other hand, caregivers' scaffolding, called motherese and motionese⁶, is also an important factor in language learning based on synesthesia. Caregivers are known to modify their speech and actions when interacting with infants so that the infants easily detect the visual and auditory inputs and extract important information from the inputs. I suppose, in addition, when modifying speech and actions they add some sort of equivalent information to the inputs. When teaching the meanings of "large/small," for example, caregivers will pronounce "large" loudly with showing a large gesture while pronounce "small" softly with a small gesture. In the cases of "long/short" and "up/down," they may change the length of the speech and action or the pitch and the moving direction of the action according to the meanings of the words. I thus suggest that the interaction between a synesthetic mechanism of a robot and caregivers' motherese/motionese facilitate language learning by the robot.



Fig. 1. The idea for a language learning model of a robot.

I am currently developing a robotic learning model based on synesthesia and designing psychological experiments to examine the quantitative characteristics of motherese/motionese.

References

1. Cytowic, R. E. Synesthesia: Phenomenology and Neuropsychology. *PSYCHE*, 2(10), 1995.
2. Maurer, D. and Maurer, C. *The World of the Newborn*, New York: Basic Books, 1988.
3. Ramachandran, V. S. and Hubbard, E. W. Hearing Colors, Tasting Shapes. *Scientific American*, pp. 53-59, 2003.
4. Roy, D. K. and Pentland, A. P. Learning words from sight and sounds: a computational model. *Cognitive Science*, 26, pp. 113-146, 2002.
5. Yu, C. and Ballard, D. H. A Unified Model of Early Word Learning: Integrating Statistical and Social Cues. *In Proc. of the 3rd Intl. Conference on Developmental and Learning*, 2004.
6. Brand, R. J. et al. Evidence for 'motionese': modifications in mothers' infant-directed action. *Developmental Science*, 5(1), pp. 72-83, 2002.

Tactile Manipulation of Unknown Objects

Lorenzo Natale^{*†}, Eduardo Torres-Jara^{*}

^{*}Computer Science and Artificial Intelligence Lab
Massachusetts Institute of Technology, Cambridge, MA, USA

[†]LIRA-Lab, DIST

University of Genoa, Genoa, Italy
[lorenzo | etorresj]@csail.mit.edu

An important property of embodied agents is their ability to interact with the environment in which they operate. This is considered of fundamental importance for the emergence of intelligent behavior. Recent work in robotics has shown how simple actions (like poking and prodding) can facilitate perception and learning¹. Grasping is particularly appealing in this context because provides direct access to physical properties of objects (like shape, volume and weight) that are difficult to perceive otherwise. Unfortunately this aspect has rarely been investigated, with a few exceptions^{2,3}. In part this is because current robots have very limited perceptual capabilities. In particular, tactile sensing is often inadequate or inexistent. For this reason most of the research on manipulation has focused on vision and has left haptic sensing overlooked. This paper pushes the idea that sensitive haptic feedback dramatically simplifies manipulation and improves the ability of robots to successfully interact with unknown objects.




Fig. 1. Robot Obrero. Some details of the hand and tactile sensors are visible at the bottom of the picture.

The experiments reported in this paper were carried out on the robot Obrero⁴ (Figure 1). The robot consists of a 2 degrees of freedom (DOF) head a 6 DOF arm and a 5 DOF hand. A monocular camera mounted on top of the head acquires visual information. The arm and hand are equipped with series elastic actuators⁵ (SEA) which provide force feedback at each joint. Tactile feedback is provided by 160 sensors distributed on the fingers and palm. These tactile sensors⁶ are highly sensitive to perpendicular and lateral contact forces and were designed to facilitate contact with the objects and significantly increase friction.

In this paper Obrero exploits its sensing capabilities to grasp a number of objects individually placed on a table. No prior information about the objects is available to the robot. Vision is used at the beginning of the task to direct the attention of the robot and to give a rough estimation of the position of the object. Tactile feedback allows the robot to refine this initial estimation during the task. The robot reaches for the object and explores with the hand the area around it. During exploration the robot exploits tactile feedback to find the actual position of the object and grasp it. The mechanical compliance of the robot and the control facilitate the exploration by allowing a smooth and safe interaction with the object. In figure 2, we observe a

sequence of the robot grasping one of the objects. Preliminary analysis of the data collected in these experiments shows that the haptic feedback originated by the interaction between the objects and the robot carries information useful for learning.



Fig. 2. An example. Sequence of the robot grasping a porcelain cup. Frame 1: the cup is presented to the robot. Frame 2: the robot reaches for the cup. Frames 3 to 6: the robot explores the space and uses tactile feedback to find the object and adjust the position of the hand around it. Frames 7 and 8: the robot grasps and lifts the cup.

References

1. Metta, G. and Fitzpatrick, P. *Early Integration of Vision and Manipulation*. In *Adaptive Behavior special issue on Epigenetic Robotics*. Vol 11(2) (2003) pp. 109-128.
2. Natale, L., Orabona, F., Metta, G. and Sandini, G. *Exploring the world through grasping: a developmental approach*. In *6th CIRA Symposium*, Espoo, Finland, June 27-30, 2005.
3. Okamura, A. M. and Cutkosky, M. *Feature Detection for Haptic Exploration with Robotic Fingers*. *The International Journal of Robotics Research*. Vol 20 (12), 2002, pp. 925-933.
4. Torres-Jara, Eduardo. *Obrero: a platform for sensitive manipulation*. *Humanoid Robots, 2005 5th IEEE-RAS International Conference* on Dec. 5, 2005 Page(s):327 – 332
5. Williamson, M. *Series elastic actuators*. Master's thesis, MIT, Cambridge, MA, 1995.
6. Torres-Jara, E., Vasilescu, I. and Coral R. *A Soft Touch: Compliant tactile sensors for sensitive manipulation*. MIT-CSAIL Technical Report MIT-CSAIL-TR-2006-014, 2006.

Fugue¹ – an AI Artwork for the 21st Century

Gordana Novakovic*, Peter Bentley*, Anthony Ruto*, and Rainer Linz[†]

*Department of Computer Science, University College London, UK

[†]Freelance composer, Melbourne, Australia

[g.novakovic|p.bentley|a.ruto]@cs.ucl.ac.uk; rlinz@aanet.com.au

Fugue is a scientifically informed art project (a product of the emerging discipline of art and science²) based on the functioning of the human immune system. It is an interactive piece, and operates within the framework of an artificial immune system algorithm, evolving in real-time, and expressed through vision and sound. The emergent, evolving nature of the Artificial Immune System algorithm, the use of repetition in the form of a succession of variations of ‘events’, and the



Fig 1. Some of the visual elements representing the components of the immune system in Fugue.

complex structural and functional interrelationships between the individual elements and processes are strongly related to the musical form of counterpoint, which formed one of the inspirations for the artistic concept for Fugue. The sound is presented as a ‘mental soundscape’, a resonance of the function of the immune system in the body. Fugue symbolises the inseparable interconnectedness between all particles and functions of a living body, which is shaped by its inner functions as much as by its interaction with the world.

The AI Connection

The core algorithm, developed by the computer scientist Peter Bentley^{3,4} in the context of artificial intelligence, has now become a form of artificial life, an open system that endlessly changes and evolves, creating emergent interactions at different levels in a state of constant becoming. The Artificial Immune System software creates the dynamics of the virtual immune system drama. It also constructs and implements the architecture of Fugue by providing the functional structure for the communication channels between the visuals and the sound. And finally, its embedding within a fugue-like structure enables it not only to represent the processes involved, but also at the same time to paint a larger picture of the role of the immune system in the functioning of the human body and mind. In the full scale installation, to be exhibited for the first time in Belgrade in July, the human participants will engage the system in a spontaneous non-verbal dialogue, influencing both the unfolding of the immune system drama and the nature of the participants’ experience.

Art and the Embodied Mind

The digital revolution has changed our experience of the physical world and the nature of perception⁵. There is no doubt that true novelty in contemporary art is found

in the concept of interactivity: an active, responsive art work is a process rather than an object, and the audience is now a participant⁶. In an interactive installation, the application of technology creates new forms of non-verbal communication. The whole body of the installation engages in a dialogue with the human body. The sensory system of the installation matches the participant's senses, and the computerised 'nervous system' of the installation matches h/er nervous system. The entire being of the participant is encircled with sounds, images, harmonies/disharmonies, noise/silence, and is electrified by the largely unknown emissive properties of the installation. In truth, the active audience becomes amalgamated with the installation, and the conventional boundaries of the human body (and brain) are called into question.




Fig. 2 The first showing of the Fugue prototype at the New Forms Festival, Vancouver, 2005

It might be thought that AI, as a discipline, would have been a major force in visualising the kind of future it is creating for the human race. However, this seems to have been left largely to futurists such as Virilio⁷ and to artists and philosophers. So what will be the role of art in a future dominated by technology and saturated with artificial intelligence? It might be possible that, like meditation, art will prove to be a means of ending the endless stream of repetitive thinking and bringing attention away from an

awareness of clock-time, and of liberating our body for a full awareness of itself and its place in the world, and also of the broad spectrum of stimuli and meditative forces emitted by an artwork and unfolding in the body of the participant.

References

- [1] www.fugueart.com
- [2] www.cs.ucl.ac.uk/research/tesla/
- [3] Bentley, P.J., Novakovic, G. and Ruto, A. (2005) Fugue: an Interactive Immersive Audiovisualisation and Artwork using an Artificial Immune System. In *Proceedings of the Fourth International Symposium in Artificial Immune Systems (ICARIS)*, LNCS 3267 pp1 - 12 Springer-Verlag GmbH.
- [4] www.cs.ucl.ac.uk/staff/P.Bentley/
- [5] Ascott, R. (2001). Planetary Technoetics: art, technology and consciousness, Conference University Paris 8 (LabArt): Roy Ascott & Leonardo/Olats
- [6] www.gordananovakovic.com/writings/electronic_cruelty.html
- [7] Virilio, P. (2000). *Open Sky*: Verso, London

Acknowledgments

This work is funded by the Arts and Humanities Research Council, Arts Council England, and The Leverhulme Trust.

Computational Auditory Scene Analysis - Towards Listening to Several Things at Once -

Hiroshi G. Okuno

Departments of Intelligent Science and Technology
Graduate School of Informatics, Kyoto University, Sakyo, Kyoto 606-8501, Japan
okuno@i.kyoto-u.ac.jp

Computational Auditory Scene Analysis (CASA)¹

“Listening to several things at once” is a dream of many people and a goal of AI and robot audition. Psychophysical observations reveal that people can listen to at most two things at once. Robot audition is, hence, an essential intelligent function for robots working with humans on a daily basis. Because robots will encounter many different kinds of sounds and noises, robot audition should be able to recognize a mixture of sounds and be noise robust. The three main capabilities of robot audition are sound source localization, separation, and recognition of separated sounds. These functions should require minimum *a priori* information about the robot’s acoustic environments and speakers² so that these robots can be deployed in various kinds of environments. Here, our studies on CASA are briefly summarized.

Sound Source Localization Based on Audio-Visual Integration

Sound source localization is not as accurate as visual localization due to reverberation and noises, while it is more robust against occlusion. Figure 1 depicts real-time multiple speaker tracking based on audio-visual integration³. Sound sources are localized by calculating the interaural intensity difference (IID) and interaural phase difference (IPD) between left and right channels. Because the accuracy of localization is about 10 degrees, the ambiguities are dissolved by integrating visual localization obtained by stereo vision. Thus, the robot can turn to a speaker even if it cannot see him or her. It can also turn to a speaker when he or she starts talking. These behaviors improve social interaction between the robot and people.




Fig. 1. Real-time multiple speaker tracking system with binaural microphones and stereo cameras.

Sound Source Separation Based on Active Direction Pass Filter (ADPF)

This system with IID and IPD was extended to separate sound sources from a mixture of sounds. This system is called the active direction-pass filter (ADPF) because it extracts frequency bins of the same direction. Because the sensitivity depends on the direction, the pass range for the direction varies, that is, it is narrower for central directions, while it is wider for peripheral directions. Separated sounds are recognized with speaker- and direction-dependent acoustic models trained by sounds that are transformed by ADPF with a single sound source. The performance of recognition for three simultaneous word utterances was about 70% to 78%⁴. This system needs a lot of *a priori* information.



Fig. 2. Humanoid SIG asking favorite color to group with three speakers answering at once⁴.

Missing-Feature Theory Based Speech Recognition with Missing Feature Mask

The missing-feature theory is used to reduce the amount of *a priori* information in recognizing separated sounds. It uses a set of missing feature masks to avoid unreliable features in recognition. However, ADPF does not give clues for missing feature mask generation. Therefore, we use geometric source separation (GSS) and a multi-channel post-filter with 8 microphones (Fig. 3). GSS is an adaptive beam former, and the post-filter calculates the leakage between channels. This leakage is used to generate a missing feature mask. For three simultaneous



Fig. 3. Three people giving meal orders at once.

speakers, the performance of recognizing isolated word recognition with a clean acoustic model is improved from 57.4% to 83.9% for the center speaker for an open test. This performance was attained by automatic missing feature mask generation. For two simultaneous speakers, the average performance rate is 89.4% and 94.4% for open and closed tests. Figure 3 shows the resulting system actually recognizing three simultaneous speakers standing about 1.5 m away from the robot, Robovie.

Recognizing Environmental Sound as Extended Onomatopoeia

Because the robot can hear environmental sounds and music besides voiced speech, these sounds are recognized as onomatopoeia, that is, sound-imitation words⁵. Because a literal representation of environmental sounds carries ambiguities, onomatopoeia is extended by using pitch and duration like musical transcription to disambiguate them. This music-like representation of environmental sound in XML is our first step towards signal-to-symbol transformation. Future work should include recognizing non-voiced speech separated by GSS as extended onomatopoeia. We believe such an approach will improve robot-human verbal communication.

References

1. Rosenthal, D. & Okuno, H.G. (eds) *Computational Auditory Scene Analysis* (Lawrence Erlbaum Associates, 1998).
2. Okuno, H.G., *et al.* Challenge Problem for Computational Auditory Scene Analysis: Understanding Three Simultaneous Speeches, *Proc. of IJCAI-97*, pp.30-35.
3. Nakadai, K., *et al.* Real-Time Auditory and Visual Multiple-Object Tracking for Robots, *Proc. of IJCAI-01*, pp.1425-1432.
4. Nakadai, K., *et al.* Improvement of Recognition of Simultaneous Speech Signals Using AV Integration and Scattering Theory for Humanoid Robots, *Speech Communication*, Vol.44, Issues1-4 (2004) pp.97-112.
5. Ishihara, K., *et al.* Automatic Sound-Imitation Word Recognition from Environmental Sounds focusing on Ambiguity Problem in Determining Phonemes, *PRICAI 2004*, LNAI 3157, pp.909-918 (Springer-Verlag, 2004).

Artificial Death for Attaining System Longevity

Megan Olsen, Hava Siegelmann
Department of Computer Science
University of Massachusetts, Amherst, USA
molsen@cs.umass.edu, hava@cs.umass.edu

Self-monitoring systems have the desired property of surviving damages. Many biological systems are self-monitoring to some extent, and can hint at possible ways of attaining this property. Inspired by biology, we propose a self-developing and self-monitoring system in which agents present the possibility of mutating their genes. The agents can repair genes to some extent but if they acquired un-repaired mutations and are no longer aiding the environment they are considered *aberrant* and should therefore die. This gives rise to an artificial creature made of multiple agents in which local death prevents the unhealthy agents from destroying the system. We call this system HADES (Healing and Agent Death Encouraging Stability).

The main application of the study is in distributed systems such as sensor networks, where agents communicate to accomplish the system's goal(s). If an agent becomes injured it should first try to self-repair [1]. If these injuries are not repairable, they may escalate so that it is beneficial that the agent kills itself. This improves on the robustness of previous systems.

HADES may shed light also on tumor formation and suppression. Current understanding of how cancer evolves requires a multi-stage model [3]. Our system also requires multiple steps of mutation in a specific order for a tumor to develop, otherwise the individual agent will either correct its own problems or kill itself to preserve the system's health. Using our system we will address the question of how a system that was in a healthy steady state can lose agents and become weak. We will compare protocols of possible apoptosis, and propose "dynamical-treatments." We will check the new idea of neighboring cells inducing awareness and death.

The System

We assume five gene types inspired by biology. Oncogenes control cell splitting and determine if an agent will replicate. The tumor suppressor genes ensure that the oncogenes do not cause excessive splitting. The repair genes attempt to repair any damage that has occurred to other genes. There are also apoptosis genes that enable death unless they are damaged. Blood supply genes ensure that the agents do not develop more than their nutrition supply allows.

HADES begins replication from a stem agent. The stem agent and blast agents replicate with a high probability whenever they detect a lack of agents in the system. A mature agent replicates with some fixed probability if there is room around it for the daughter to exist. Boundaries stop healthy cells from further replication.

Death of an agent is inspired by cell apoptosis. Apoptosis of an agent mainly occurs whenever an agent detects through self-monitoring that its genes are damaged and can not be repaired. However, if the genes controlling death are damaged the agent may not realize it should die and will continue to replicate, spreading its damaged genes to its daughters. Eventually the surrounding healthy agents will be pushed out of the way by this cluster of troubled agents. By a new mechanism we introduce, the healthy agents that were pushed aside will start sending local signals




Fig. 1: Genes in the order of their necessary mutations

toward the pushing agents to make them aware that they are not functioning properly. Any agent that receives the signal will interpret it as a “kill” command, although it is the agent's own prerogative to decide to kill itself. An agent is convinced to die by sensing the strength of the “kill” signal in its area, implying that multiple agents are sending it.

Results

We differentiated four cases of damage and response. The first case was the benchmark, with mutations occurring in random order. The second case had no functioning repair, the third had no functioning apoptosis, and the last had ordered mutations set to be the worst case so that they are the best to create a tumor. In the worse case, the first mutation is at the repair gene, and the second one is the apoptosis gene; this way we stop the agent's ability to repair and to die. The third mutation must be the tumor suppressor gene, so it would not keep the replication genes from expressing themselves. The fourth gene is then the replication gene, with the blood supply gene being the last one. If these orders are not followed, then a tumor is not formed.

The number of agents in the system during equilibrium in all cases was around 2250 until 25000 iterations. All cases were run for 1705810 iterations. Splitting occurred with probability 0.0025 and natural death not due to apoptosis was set for 0.0024, based on breast cancer literature [2]. The probability for mutation at each split was 0.00099 for each type of gene, based on having 10^6 base pairs per each of 5 genes that we modeled, and 10^9 other base pairs.

We first recorded efforts of cells to die. If the worst case ordering is forced, agents try 1525 times to die. However, if repair never works the agents have 517 attempts to die. If apoptosis never works the agents try 1134676 times to die, but of course every attempt fails. If the mutations are randomly ordered then only 375 attempts to die are made. Each attempt to die was fixed to have a 50% chance of succeeding. Therefore, a system that begins healthy does not need as much death as a system that is unhealthy. Death is an important factor in sustaining health, as it increases with the number of aberrant agents.

Death, however, can have a negative side effect on the health of the system over time. While the equilibrium size was 2250 for long time, around the 25000th iteration the number of cells in the organisms went down in steady state to about 1000. This is probably since the values for the splitting probabilities and the natural death probabilities were chosen to be .0025 and .0024 respectively. When apoptosis is allowed this splitting probability can not match the death rate, so the size decreases. This gives an interesting view of some body areas that have fewer cells later in life. It is interesting to see how the required ratio between these two probabilities must be interpreted to enable stable equilibrium for different life spans.

A tumor can still be created, though with little probability, when the mutations occur in the particular worst case order. After one aberrant cell forms, a tumor will grow exponentially fast since aberrant cells have mutation in their splitting gene which makes them try to split frequently, and since tumor cells are blind to the distance maintained between healthy cells.

We are currently incorporating an additional mechanism to attain longevity where death signals are sent by cells which are being pushed by others. The assumption is that only healthy cells keep respectable distance. Preliminary results demonstrate the strength of this mechanism.

References

1. Bokareva, T., Bulusu, N., and Jha, S. SASHA: Towards a Self-Healing Hybrid Sensor Network Architecture. *Proceedings of the 2nd IEEE International Workshop on Embedded Networked Sensors (EmNetS-II)*, Sydney, Australia, May 2005.
2. Humphreys, R., Krajewska, M., Krnacik, S., Jaeger, R., Weiher, H., Krajewski, S., Reed, J., Rosen, J. Apoptosis in the Terminal Endbud of the Murine Mammary Gland: A Mechanism of Ductal Morphogenesis. *Development* 122, 1996.
3. Ritter, G., Wilson, R., Pompei, F., Burmistrov, D. The Multistage Model of Cancer Development: Some Implications. *Toxicology and Industrial Health*, 2003.

CAVIAR – Building Blocks for a Large-Scale Spiking Vision System

Matthias Oster*, Raphael Berner*[†], Patrick Lichtsteiner*, Rafael Paz-Vicente[†], Rafael Serrano-Gotarredona[§], Alejandro Linares-Barranco[†], Tobi Delbrück*, Philipp Häfliger[†], Rodney Douglas*, Anton Civit[†], Bernabe Linares-Barranco[§] and Shih-Chii Liu*

*Institute of Neuroinformatics, Uni-ETH Zurich, Switzerland

[†]Department of Informatics, University of Oslo, Norway

[§]Inst. of Microelectronics, [†]Univ. of Seville, Spain

[mao | shih]@ini.phys.ethz.ch

Neuronal networks implemented in VLSI consist of building blocks that resemble their biological counterparts as close as possible. For example, neurons perform analog computation in continuous time, and spikes are transmitted asynchronously along virtual synaptic connections that preserve real-time. These building blocks have now evolved to a state where they can be assembled into large-scale artificial systems. We report the current state-of-the-art of the field by describing the CAVIAR project (Convolution Address-Event-Representation Vision Architecture for Real Time), the largest multi-chip system assembled until now (Fig.1)^{1,2}.

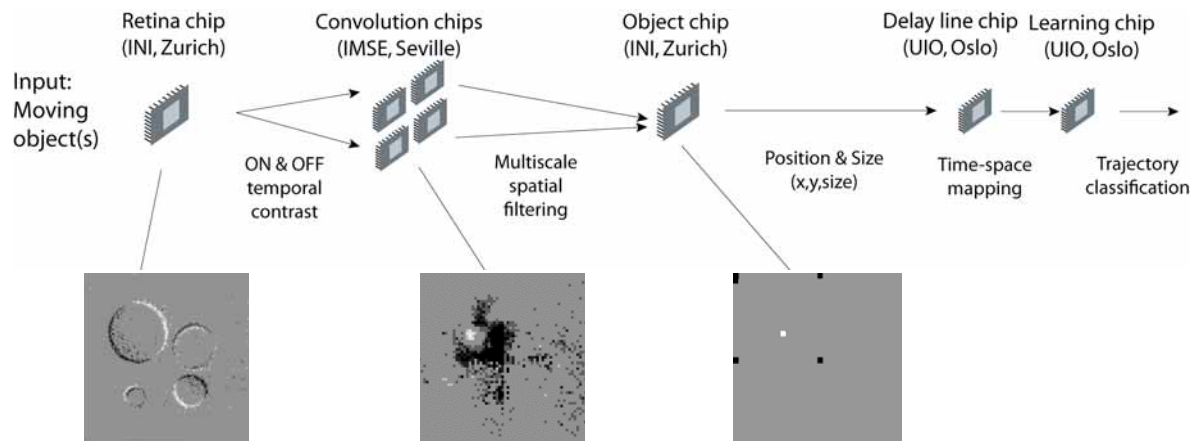


Fig.1: Overview of the CAVIAR system: the artificial retina detect temporal contrast of moving objects, here 4 black disks rotating on a white background, and transforms it into a spiking representation (left inset): black dots represent ‘OFF’ events (in response to negative temporal contrast edges), white dots ‘ON’ ones (response to positive contrast edges). The spikes are transmitted to a stage of convolution chips. Each spike-based convolution chip is programmed to detect the center of a disk of specific size (middle inset): white dots mark spikes that represent positive results of the convolution, black ones negative results. Four convolution chips are tiled to increase the resolution; in other configurations they can be programmed to detect different objects. The spike output of the spatial filter process is cleaned by the ‘object’ chip, which detects the object position (right inset) using a winner-take-all network. The white dot marks the spike output of the object chip; black represents the activity of the inhibitory neurons involved in the computation. Object position and size (in case of the convolution kernels programmed for different ball sizes) are then expanded over time in the delay line chip and the resulting trajectories are classified by the learning chip. Additional modules can be used to monitor the spike trains in the system, to map the synaptic connections, and to inject artificial spike trains into the system.

The main objective of CAVIAR is to develop an infrastructure for constructing a bio-inspired, hierarchically structured multi-chip system for sensing, processing, and actuation. All modules within CAVIAR use the Address-Event-Representation (AER), an asynchronous inter-chip communication protocol. Senders, e.g. pixels or neurons, generate spikes that are represented on the AER bus as source address events. The events can be merged with events from other senders and can be broadcast to multiple receivers. Arbitrary synaptic connections are implemented by remapping the digital addresses. For performance reasons, commonly used connectivity patterns are implemented on-chip, like a convolution operation and a winner-take-all network.

A set of portable interface boards has been developed with transmission rates up to 10^7 spikes/s with a delay of less than $1\mu\text{s}$ between sender and receiver, thus preserving the real-time character of spiking connections between the neurons. Through AER, the architecture and connectivity of the system can be easily changed, e.g. the number and the configuration of convolution chips in the feature extraction stage can be adjusted. The building blocks in CAVIAR are:

Building block	Function	Size (pixels/neurons)
Retina	detects temporal contrast edges	$128^2 \times 2$ (on/off)
Convolution Chip	programmable 32×32 convolution kernel	4×32^2 (tiled / parallel)
'Object' Chip	multi-dimensional spike-based winner-take-all	32^2 or 4×16^2 (4 synapses each)
Delay line	programmable delays	880 elements
Learning chip	associative Hebbian learning	32 (64 synapses each)
Interfaces	connectivity; monitoring and injecting spike trains	up to 2^{16} addresses

One strong advantage of AER systems is that computation is event driven and thus can be very fast. Computation inside the total 37920 neurons and pixels is performed in parallel, opening new ways to explore fast and concurrent processing of sensory data. Each of the building blocks processes spike rates up to 2Mspikes/s, but typical spike rates found in the system are much lower as the information is more sparsely coded along the chain of computation.

Besides further development in terms of size, speed and reliability of building blocks and interfaces, future systems will incorporate recurrent networks and embed learning functions at several stages. By its similarity with the physical substrate of natural intelligence, Neuromorphic engineering provides an ideal test bed for exploring more sophisticated intelligent artificial systems that interact with the real-world.



Fig.2: System assembly (from right to left):

- retina (inlay)
- multiple convolution chips (two boards)
- 'object' chip

References

1. R. Serrano-Gotarredona et al.: **AER Building Blocks for Multi-Layer Multi-Chip Neuromorphic Vision Systems**, *Advances in Neural Information Processing Systems (NIPS)*, 18: 1217-1224, Dec, Y. Weiss and B. Schölkopf and J. Platt (Eds.), MIT Press, 2005
2. CAVIAR project description: <http://www.imse.cnm.es/caviar/>

Simple Linear Spring Is Not Sufficient?

Stabilizing Passive Dynamic Running by Exploiting the Intrinsic Body Dynamics

Dai Owaki* and Akio Ishiguro[†]

*Dept. of Computer Science and Engineering

Nagoya University, Nagoya, Japan

[†]Dept. of Electrical and Communication Engineering

Tohoku University, Sendai, Japan

owaki@cmplx.ecei.tohoku.ac.jp, ishiguro@ecei.tohoku.ac.jp

Introduction

The behavior of a robot emerges through the dynamics stemming from the interaction between the control system, mechanical system, and environment¹. Considering the fact that the control and mechanical systems, which are the targets to be designed, are positioned at the source of this interaction, they should be treated with equal emphasis in the design process. However, as can be seen from the terms of “control system and controlled system” or “controller and controlled object”, traditionally these two systems have been clearly distinguished by their dominant relationship. In other words, system enhancements have been achieved mainly by increasing the complexity of control systems. This, however, causes serious problems, particularly in terms of adaptability and energy efficiency.

Under these circumstances, recently the importance of the following suggestions has been widely recognized: (1) there should be a “well-balanced coupling” between control and mechanical systems; (2) one can expect that quite interesting phenomena, *e.g.* real-time adaptability and high energy efficiency, will emerge under such well-balanced coupling; and (3) the well-balanced coupling between control and mechanical systems should be varied depending on the situation. Since this research field is still in its infancy, it is of great worth to accumulate various case studies at present.

In light of these facts, as an initial step toward this goal, this study intensively discusses

the effect of the intrinsic dynamics of a robot's body on the resulting behavior, in the hope that the mechanical systems appropriately designed will allow us to significantly reduce the complexity of control algorithm required as well as to increase the robustness against the environmental perturbation. To this end, we focus on the property of leg elasticity of a passive dynamic *running* biped, and investigate how this influences the stability of running. The reason why we have employed a passive dynamic running biped as a practical example is that in contrast to a passive dynamic *walking* biped where the behavior is generated through the circulation between kinetic energy and potential energy, a passive dynamic running biped can additionally exploit elastic energy stemming from the leg elasticity. Due to this, a passive dynamic running biped is expected to show different form of the mechanical system's contribution to the resulting behavior.

The Model

Figure 1 illustrates a model of a passive dynamic running biped employed. In the figure, the motion performed during one period of running is shown for clarity. We assume that the state of touch-down is treated as the initial state of running. As illustrated in the figure, running can be described as a periodic alternation between

stance and flight phases.

Therefore, the equations of motion for bipedal passive dynamic running should be described for each phase. In this study, according to the work done by Seyfarth *et al.*², a spring-mass model was employed for describing the stance phase. On the other hand, during the flight phase, the position of the center of mass was simply determined by the gravitational acceleration.




Fig. 1. The model employed.

Simulation Results

In order to investigate the effect of the spring characteristics on the stability of running, we have observed the convergence to the limit cycle. To this end, we have implemented a linear spring and a nonlinear spring into the leg, each of which is expressed as $f(z) = Kz$ and $f(z) = Kz^2$, respectively. It should be noted that the biped with the nonlinear springs converges to the limit cycle rapidly (Fig. 2). These results indicate that the vector field is significantly modified in the case of nonlinear springs, such that the trajectory could converge rapidly to the limit cycle. In sum, the property of leg elasticity significantly influences the structure of the vector field in the state space which specifies the self-stability of the mechanical system. This is an unexpected result. To the best of our knowledge, this has never been explicitly discussed so far. We expect that this significantly not only reduces the temporal complexity of control algorithm but also alleviates constraints on the precision of measurement required for the sensory feedback.

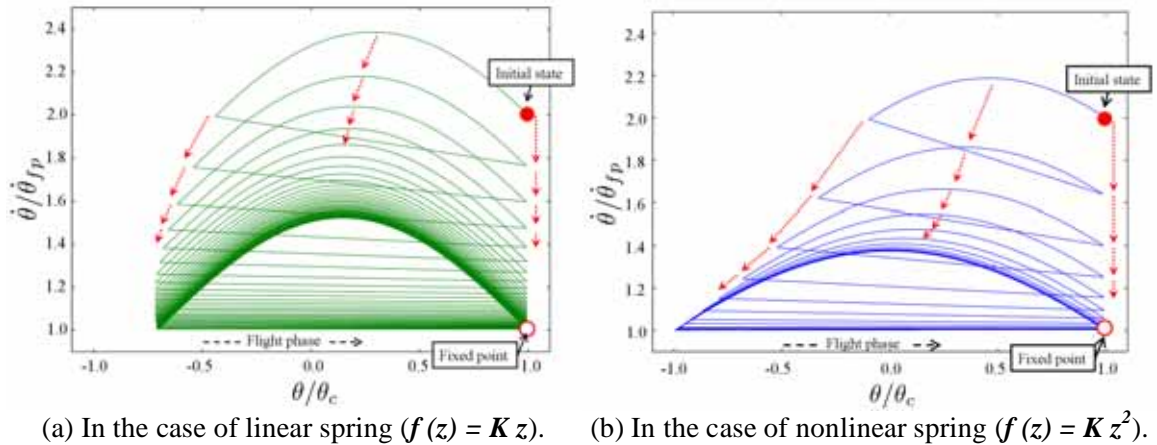


Fig. 2: Comparison of the convergence to the limit cycle between the linear and nonlinear springs. Note that the biped with the nonlinear spring shows rapid convergence to the limit cycle.

References

1. Pfeifer, R. and Scheier, C. *Understanding Intelligence* (MIT Press, Boston, 1999).
2. Seyfarth, H. Geyer, M. Gunther, and R. Blickhan, *A Movement Criterion for Running*, (Journal of Biomechanics, Vol.35, pp.649--655, 2002.)

Formalizing Emergence to Accelerate the Accretion of Embodied Intelligence

Frank Pasemann

Fraunhofer Institute AIS, Sankt Augustin, Germany
frank.pasemann@ais.fraunhofer.de

Introduction

The emergence of spatially and/or temporally coherent structures is a basic phenomenon observed for interacting non-linear systems; and it is recognized that an understanding of emergent phenomena is of fundamental importance in the study of living organisms, and, the point made here, especially for the development of embodied intelligent systems.

Understanding intelligence in the sense of being able to deal with the physical properties of ones eco-niche in a life sustaining way [1], a higher level of intelligent or cognitive behavior can be expected to result from an emergent process, induced by the growing complexity of an agents internal composition, and the way it interacts (morphology, motor system) with its environment and perceives this interactions (sensor system).

Often the attempt to define emergence on the background of, and with reference to a particular theory is often assumed to be counter-productive because of its arbitrariness and limitations. Nonetheless, if one wants to make this concept a productive analytical tool for investigations, a formal definition may be of help, even if one approaches the problem only with regard to the limited applicability in the context of a specific theory.

Modular Neurodynamics and Evolutionary Robotics

Following a modular neurodynamics approach to cognitive systems [4] and applying it to Evolutionary Robotics [2], the realizable reactions and behaviors, as well as the capacity of cognitive abilities, like different types of memory, prediction, and planning, depend crucially on the richness of the attractor structure of the underlying neural control system.

Dynamical systems of this kind can in general not be constructed in terms of a fully connected neural system. It therefore is appropriate to start with specialized neuromodules developed already for specific sensor systems, motor configurations and tasks, and then using evolutionary fusion techniques [5] to generate enfolding structures by adding sensors (*sensor fusion*) and motors together with additional neurons and connections to solve a more comprehensive task. When applying such a fusion processes new qualitative behaviors can appear which are emergent in the sense that these – desired - properties (or solutions for a given task)




Fig. 1 a) A simple obstacle avoidance controller with two inputs and two motor neurons, b) light tropism (4 inputs one neuron) added by fusion (for a Khepera robot)

are neither realizable by the original modules, nor can they be foreseen beforehand. The question then arises, if one can find general (mathematical) conditions for a fusion process, under which emergent properties or behaviors have to be expected. Or, what is almost as effective and perhaps more realistic to be achieved, is to characterize those coupling structures which will suppress emergent dynamical phenomena.

Measuring Behavior Relevant Dynamical Complexity

Concentrating on neural control systems (neuromodules) as parametrized discrete-time dynamical systems, a quantitative notion of emergence has to be based on a convenient behavior oriented measure of dynamical complexity. Like emergence, the term complexity is context dependent, and there are many different attempts to define this concept. But measures based on generalized dimensions, entropies, and Lyapunov exponents give rise to computational difficulties in high-dimensional systems, and other approaches based on local linearization, will probably miss the global character of emergent dynamical properties.

Here measures are discussed which seem to be appropriate for describing the power of a neuromodule **A** with respect to its possible contribution to the “cognitive” abilities of the whole system. They are based on a discrete-time variant of *topological complexity* introduced in [3]. After defining a convenient equivalence

relation on attractor configurations (structural stability) of neuromodules, the measure describes the “distance” of an attractor configuration to a trivial one (globally stable fixed point attractor) by counting the minimal number of bifurcations one has to cross going along all possible paths in parameter space. Taken over all possible attractor configurations of the neuromodule **A** one derives its complexity $\gamma(\mathbf{A})$.




Fig. 2. A self-excitatory neuron **A** (hysteresis element) recurrently coupled to a self-inhibitory neuron **B** (period-2 oscillator) gives a parametrized 2-neuron system (**A,B**) capable of periodic, quasi-periodic and chaotic dynamics. Its dynamical complexity $\gamma(\mathbf{A,B})$ is larger than that of the disjoint system **A** and **B**; i.e. $\gamma(\mathbf{A,B}) > \gamma(\mathbf{A}) + \gamma(\mathbf{B})$.

References

1. Pfeifer, R. & Scheier, C. *Understanding Intelligence* (MIT Press, Boston, 1999).
2. Nolfi, S., & Floreano, D. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines* (MIT Press, Cambridge, 2000).
3. Thom, R. *Structural Stability and Morphogenesis* (Benjamin, Reading, 1975).
4. Pasemann, F., Neuromodules: A dynamical systems approach to brain modelling, in: Herrmann, H., Pöppel, E. & Wolf, D. (eds.) *Supercomputing in Brain Research - From Tomography to Neural Networks* (World Scientific, Singapore, 1995), pp. 331 - 347.
5. Hülse, M., Wischmann, S. & Pasemann, F., Structure and function of evolved neuro-controllers for autonomous robots, *Connection Science*, 16:4, 2004, p.294 - 266.

The Anticipatory Nature of Representations

Giovanni Pezzulo, Gianluca Baldassarre, Rino Falcone, Cristiano Castelfranchi
Institute of Cognitive Science and Technology - CNR
Via S. Martino della Battaglia, 44 - 00185 Roma, Italy
[giovanni.pezzulo|gianluca.baldassarre]@istc.cnr.it

Since the beginning of AI, *intelligence* was conceived as the capacity to solve a problem by working on internal representations of problems, i.e. by acting upon “images” or “mental models” with simulated actions (“reasoning”), before acting in the world. Successively, the concept of “representations” has been attacked in many ways. Recently, many converging evidences in psychology and neurobiology indicate a crucial role of anticipatory representations for many cognitive functionalities such as visual attention¹ and motor control⁹. As suggested by the discovery of mirror neurons⁸, representations are mainly action-oriented and deeply based on the motor apparatus. Barsalou² and Grush⁴ try to provide unitary accounts of these phenomena and anticipatory functionalities now begin to be explored from a computational point of view^{9,5}.

We think that by *conceiving representations as mainly anticipatory* it is possible to *reframe many of the central claims of AI*. In fact, the ability that characterizes and defines a “true mind”, as opposed to a merely adaptive systems, is that of building representations of the non-existent, of what is not currently (yet) “true” or perceivable. A real mental activity begins when the organism is able to endogenously (i.e. not as the consequence of current perceptual stimuli) produce an internal perceptual representation of the world (“simulation” of perception)³. For example, the organism can generate the internal “image” for matching it against perceptual inputs while actively searching for a given object or stimulus while exploring the environment; or it can use it as prediction of the stimulus that will probably arrive, and match its predictions against actual stimuli, and be confirmed, disconfirmed, or surprised. But it can also form mental representations of the current world to work on it, modifying this representations for virtually “exploring” possible actions, events, results: “what will happen if...?”; or maintain concurrent representations, such as motor plans, and select among them. Expectations are not only representations: they can have motivational, axiological, or deontic nature; saying us not only how the world is, was, will be; but how the world *should* be, how the organism would like the world to be. Anticipatory representations can thus be used as *goals* driving the behavior. This is what mind really is: *conceiving and desiring what is not there*: the presupposition for hallucinations, delirium, desires, and utopias.

We aim at providing a *unitary account of the role of anticipation in many cognitive functionalities*, including sensorimotor interaction with the environment, attention, planning and goal selection; and to *integrate them into an unitary architectures*. Anticipatory representations offer two advantages: 1) they make it possible to build up more and more complex functionalities exploiting less complex ones (e.g., off-line planning exploiting on-line planning); 2) even if they are used for detaching from reality (as in visual imagery, or in planning), they are fully grounded: they are acquired in the past experience (e.g. with supervised learning⁹) and can be compared with actual stimuli. Fig. 1 shows our model for an *oculo-motor coordination* system in which many concurrent *perceptual* and *motor schemas* control a camera and a gripper^{6,5}. In the framework of the EU funded project MindRACES (FP6-511931),

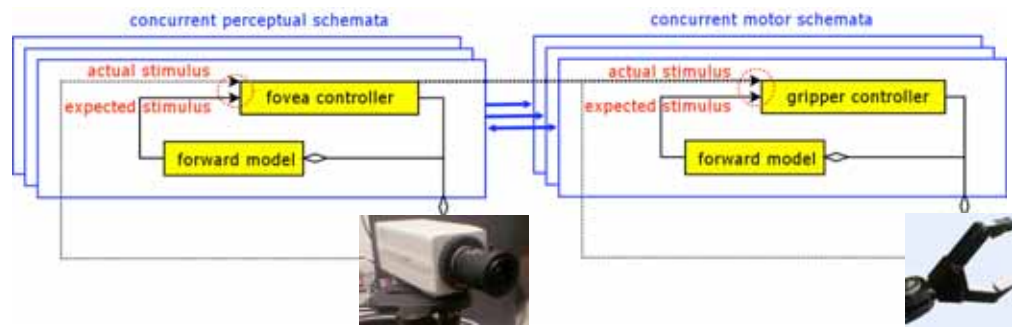


Fig. 1: Coupled perceptual-motor schemas for oculo-motor coordination

this model is being used for realizing a system that has to pick-up with its gripper insects having different sizes, velocity and trajectories on the basis of visual input.

In our systems, anticipation has five main roles: 1) *Action control*: in the case of perceptual schemas, this means orienting the fovea towards relevant inputs (e.g., relevant colors and trajectories); in the case of the motor schemas, this means selecting the most appropriate gripper action (e.g., specialized for quick or slow, big or small insects). Moreover, some perceptual and motor schemas are *coupled*: active perceptual schemas specialized in tracking some trajectories or colors pre-activate motor schemas for picking related insects and vice versa. 2) *Decision*: many competing motor plans are generated and maintained for the same or for different targets, and choice depends on *predictive accuracy*. Schemas predicting better are selected: the rationale is that schemas predicting well are “well attuned” with the current course of events⁹; prediction is an evaluation of schemas efficacy. 3) *Replacing the actual input* if sensors are unavailable or unreliable. 4) *Compensating time delays*. 5) *Erasing the auto-generated input* (e.g., for avoiding to consider as target the own moving gripper).

The anticipatory representations provided by the forward models (e.g. implemented using fuzzy logic or neural networks) offer also a bridge for more complex functionalities such as *offline planning*: possible outcomes of events can be simulated and compared offline by exploiting the same machinery involved in online visual and motor planning, but without sending commands to the effectors. During this operation the expected stimuli replace actual ones and serve as inputs for chaining the schemas. This offers two more advantages: 1) “detaching” the representations from the sensorimotor loop by setting up *hierarchies of schemas* representing abstract concepts; 2) *using goal states*, and not current stimuli, *for the selection of action*.

References

1. Balkenius, C., Hulth N. (1999) Attention as Selection-for-Action: a Scheme for Active Perception *Proceedings of EuRobot-1999*. Zurich.
2. Barsalou, L. W. (1999) Perceptual symbol systems *Behavioral and Brain Sciences* 22 (4): 577-660
3. Castelfranchi, C. (2005) Mind as an Anticipatory Device: For a Theory of Expectations. *BVAI 2005*: 258-276
4. Grush, R. (2004) The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27 (3): 377-396.
5. Jordan, M.I., Rumelhart, D. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16:307—354
6. Ognibene, D. Mannella, F., Pezzulo, G., Baldassarre G. (2006). Integrating Reinforcement-Learning, Accumulator Models, and Motor-Primitives to Study Action Selection and Reaching in Monkeys. *Proceedings of ICCM 2006*
7. Pezzulo, G. Calvi G. (2005). Fuzzy-based Schema Mechanisms in AKIRA. *Proceedings of CIMCA'2005*. Vienna.
8. Rizzolatti, G., Fadiga, L., Gallese, V., Fogassi, L. (1996) Premotor cortex and the recognition of motor actions, *Cognitive Brain Research* 3 131-141.
9. Wolpert, D. M., Kawato, M. (1998) Multiple paired forward and inverse models for motor control. *Neural Networks* 11(7-8):1317-1329

Intelligence from Information Structure

Daniel Polani
Algorithms and Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire

The von Neumann picture of computation and its unbounded possibilities made the algorithmic model a dominant motif of much of past work in artificial intelligence. This algorithmic model implements the “top-down” construction of an intelligent system. In other sciences, e.g. physics or biology the main body of work consists of observation and discovery instead of creation of phenomena (with the obvious exception of the engineering branches of these sciences). Intelligent information processing in biology has evolved without guidance by an “intelligent designer”. This indicates that to understand the concept of intelligence and to harness it in artificial systems we might need to find ways to *evoke* the phenomenon of intelligence rather than to *construct* it.

Recent work in AI, specifically biologically inspired research, such as Neural Networks, Genetic Algorithms or Artificial Immune Systems moves in this direction. Here typically one attempts either a very accurate biological modelling, giving rise to complex computational models or one resorts to simplified models of biological systems. The central question is whether the properties that are responsible for a particular success in the biological counterpart have been appropriately transferred into the synthetic system.

“No Free Lunch”-type considerations have been considered a threat to the AI enterprise: it would be fundamentally impossible to learn efficiently in arbitrary worlds. On the other hand, even if we have not yet been able to reproduce this satisfactorily in artefacts, we *do* have an existence proof for the possibility that intelligence can emerge from simple beginnings: the evolution of biological life and intelligence. In addition, higher levels of intelligence have emerged on different routes. For instance, the evolution of the high intelligence level of the octopus took place independently from that of the mammals.

Thus, it is clear that No-Free-Lunch arguments are probably not relevant for the emergence of intelligence in a real-world scenario: the world is *not* arbitrary but intricately structured. It is constrained by a subtly intertwined set of properties. The simplest among these are symmetries, continuity and smoothness. One hypothesis is that it is these real-world properties that shape the emergence of intelligence in biology (“intelligence through embodiment”), more concretely, the information structure of these properties.

Information-theoretic principles have been suggested as a principle guiding living agents through the structure in the world, since the advent of information theory, e.g. in early cybernetics (Ashby, 1952). Information theory is being used in learning algorithms such as ID3 (Quinlan, 1983) and information-theoretic models have found a broad spectrum of applications, ranging from data mining to the modeling of biological information processing (Lee et al., 2000).

The sophistication and complexity of biological brains is so enormous that it would be of significant importance to identify principles that guide their development and dynamics. It turns out, in fact, that information is a central resource in many respects: biological organisms do often exhaust the available information channels to the limit

(Laughlin et al., 1998) and closely related Bayesian mechanisms are increasingly being established as a basis of biological learning (Kording and Wolpert, 2004).

Applied as a principle, this helps to restrict the arbitrariness of possible information processing architectures, to understand biological information processing and to harness it for artificial systems. Prominent instances are Linsker's infomax principle (Linsker, 1988), as well as new conceptual tools, such as the *information bottleneck method* (Tishby et al., 1999). A measure for brain complexity as well as a set of approaches to characterize the information structure of sensorimotor data have been suggested in (Tononi et al., 1994, Lungarella et al., 2005).

In our work, we extended existing infomax approaches to the full perception-action loop combined with information flow techniques (Wennekers and Ay, 2005). This provides a set of approaches to attain the self-organized structuring of perception-action loops from first principles, based on the information flow structure emerging in a given embodiment (Klyubin et al., 2004); the same framework serves to discover promising avenues for evolution or development (Klyubin et al., 2005). Similarly, we have used information distances to discover relevant structure in sensorimotor maps for hardware robots (Olsson et al., 2006). It should be noted that the success of these approaches, compared to passive data mining or computer vision approaches, or even the one-way Linsker infomax approach, relies centrally on having *active* interaction with the world, i.e. on closing the perception-action channel into a complete cycle.

The importance of the information as resource in biology indicates that artificial implementations of intelligent systems can profit significantly from a wide-range use of information-theoretic techniques. Our results show that, while the use of passively acquired information is common in AI, closing the perception-action loop adds considerable additional structure to information flows and provides hereto unexploited avenues towards the emergence of intelligence in biological and artificial systems.

References

- Ashby, W. R. (1952). *Design for a Brain*. Wiley & Sons, New York.
- Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2004). Organization of the information flow in the perception-action loop of evolved agents. In *Proceedings of 2004 NASA/DoD Conference on Evolvable Hardware*, pages 177–180. IEEE Computer Society.
- Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2005). Empowerment: A universal agent-centric measure of control. In *Proc. IEEE Congress on Evolutionary Computation, 2-5 September 2005, Edinburgh, Scotland (CEC 2005)*, pages 128–135. IEEE.
- Kording, K. P. and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427:244–247.
- Laughlin, S. B., de Ruyter van Steveninck, R. R., and Anderson, J. C. (1998). The metabolic cost of neural information. *Nature Neuroscience*, 1(1):36–41.
- Lee, T.-W., Girolami, M., Bell, A. J., and Sejnowski, T. J. (2000). A unifying information-theoretic framework for independent component analysis. *Computers and Mathematics with Applications*, 39:1–21.
- Linsker, R. (1988). Self-organization in a perceptual network. *Computer*, 21(3):105–117.
- Lungarella, M., Pegors, T., Bulwinkle, D., and Sporns, O. (2005). Methods for quantifying the information structure of sensory and motor data. *Neuroinformatics*, 3(3):243–262.
- Olsson, L., Nehaniv, C. L., and Polani, D. (2006). From unknown sensors and actuators to actions grounded in sensorimotor perceptions. *Connection Science*, 18(2). Special Issue on Developmental Robotics, Douglas Blank and Lisa Meeden, editors. (In Press).
- Quinlan, J. R. (1983). Learning efficient classification procedures and their application to chess end games. In Michalski, R. S., Carbonell, J. G., and Mitchell, T., editors, *Machine Learning - An Artificial Intelligence Approach*, pages 463–482. Tioga, Palo Alto, CA.
- Tishby, N., Pereira, F. C., and Bialek, W. (1999). The information bottleneck method. In *Proc. 37th Annual Allerton Conference on Communication, Control and Computing, Illinois*.
- Tononi, G., Sporns, O., and Edelman, G. M. (1994). A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Proc. Natl. Acad. Sci. USA*, 91:5033–5037.
- Wennekers, T. and Ay, N. (2005). Finite state automata resulting from temporal information maximization. *Neural Computation*, 17(10):2258–2290.

The Role of Neurobiology in Artificial Intelligence: To Make It Less Artificial

Steve M. Potter
Laboratory for Neuroengineering
Department of Biomedical Engineering
Georgia Institute of Technology, Atlanta, USA
steve.potter@bme.gatech.edu
<http://neuro.gatech.edu>

We have, between our ears, a supremely versatile, efficient, capable, intelligent machine that consumes less than 200 Watts of power. It surprises me how little attention in the AI field has been directed toward actual brains. Artificial Neural Networks, arguably the most brain-inspired AI spinoff, are constructed of “units” so simple as to be mere cartoons of the neurons that inspired them. We ought to learn more about how the embodied nervous system accomplishes its feats, and use that knowledge to design AI that is less artificial and more brain-like.

What do we already know about NI (Natural Intelligence) that can inform AI?

Probably the clearest difference, from my neurobiologist’s perspective, between animals and artificial intelligences is the huge number of senses animals have. The continuous flow of information into the brain from the sense organs is enormous. When it gets there, it hits a network whose degree of parallelism is not rivaled by any human-made artifact. There are about 100 billion neurons in our brains, each connected to 1,000-10,000 others with 200,000 km of axons. There seems to be a bias in AI and robotics that “sensors are expensive” and so we make the most of a very few of them. To make AI less artificial, we could strive to incorporate as much sensing power as we dare imagine. We should also note that in the brain, delays are not a problem, but part of the computation. The subtle timing of arrival of action potentials carries information about the dynamics and statistics of the outside world (Gerstner et al., 1997). These are analog quantities; everything in the brain is analog and it computes with timing, not boolean logic. Brain-inspired AI of the future will be massively parallel, have many sensors, and will make use of the dynamics of interactions between analog signals (Maass et al., 2002).

What do we not know about how brains work, but could learn?

To realize this dream of AI that is closer to NI, there are a number of basic questions about how brains work that must be pursued, such as **What is a memory?** and **How do biological networks work?** We know that in brains, unlike in digital computers, the CPU and the memory are one and the same thing. The neurons and glial cells both store and process information in a spatially distributed manner. But we have only a very vague and fuzzy idea of just how they do that. The Blue Brain Project is about to set a giant supercomputer (the son of Deep Blue) to the task of simulating just one cortical minicolumn of a few thousand neurons (Markram, 2006). There is a lot going on at the networks-level that we don’t even have the vocabulary to think about yet. Neurobiologists all believe that memories are stored by changes in the physical structure of brain cells, but we don’t all agree about what those changes might be, let alone how the changes are executed when salient sensory input is received. Neurons have a stunning diversity of morphologies (Mel, 1994), more than any other cell type. There is evidence that some aspects of their shape are altered by experience (Majewska and Sur, 2003; Leuner et al., 2003). But how that relates to a memory being stored is not known.

New Neuroscience Tools

In the Laboratory for Neuroengineering at Georgia Tech, we are developing new research tools to allow neurobiologists to address such fundamental questions. We have created a new type of experimental animal, the *hybrot*. This is a hybrid robot, an artificial embodiment controlled by a network of living neurons and glia cultured on a multi-electrode array (MEA) (DeMarse et al., 2000; DeMarse et al., 2001; Bakkum et al., 2004). We developed the hardware and software necessary to create a real-time loop whereby neural activity is used to control a robot, and its sensory inputs are fed back to the cultured network as patterns of electrical or chemical stimuli (Fig. 1; Potter et al., 2006).




Fig. 1: Hybrot scheme. A living neuronal network is cultured on a multi-electrode array (MEA) where its activity is recorded, processed in real time, and used to control a robotic or simulated embodiment. The robot's sense data is converted to electrical stimuli that are fed back to the neuronal network within milliseconds. The hybrot's brain (MEA culture) can be imaged continuously on the microscope while its body behaves and learns. This may reveal the morphological correlates of memory formation.

These *embodied cultured networks* bring in vitro neuroscience models out of sensory deprivation and into the real world. An MEA culture is amenable to high-resolution optical imaging, while the hybrot is behaving and learning, from milliseconds to months (Potter and DeMarse, 2001; Potter, 2005). A network of even a few thousand living neurons is vastly more complex than any existing artificial neural network. By studying them with these new tools, we may learn some new aspects of network dynamics, memory storage, and sensory processing that could be used to make AI a bit less artificial.

References

- Bakkum, D. J., Shkolnik, A. C., Ben-Ary, G., Gamblen, P., DeMarse, T. B., & Potter, S. M. (2004). Removing some 'A' from AI: Embodied Cultured Networks. In F. Iida, R. Pfeifer, L. Steels, & Y. Kuniyoshi (pp. 130-45). New York: Springer.
- DeMarse, T. B., Wagenaar, D. A., Blau, A. W., & Potter, S. M. (2000). Neurally-controlled computer-simulated animals: a new tool for studying learning and memory in vitro. *Soc. Neurosci. Abstr.*, 26, 467.20.
- DeMarse, T. B., Wagenaar, D. A., Blau, A. W., & Potter, S. M. (2001). The neurally controlled animat: Biological brains acting with simulated bodies. *Autonomous Robots*, 11, 305-10.
- Gerstner, W., Kreiter, A. K., Markram, H., & Herz, A. V. M. (1997). Neural codes: Firing rates and beyond. *Proceedings of the National Academy of Sciences of the United States of America Proc. Natl. Acad. Sci. U. S. A.*, 94, 12740-41.
- Leuner, B., Falduto, J., & Shors, T. J. (2003). Associative Memory Formation Increases the Observation of Dendritic Spines in the Hippocampus. *J. Neurosci.*, 23(2), 659-65.
- Maass, W., Natschlager, T., & Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.*, 14(11), 2531-60.
- Majewska, A., & Sur, M. (2003). Motility of dendritic spines in visual cortex in vivo: Changes during the critical period and effects of visual deprivation. *Proc. Natl. Acad. Sci. U. S. A. Proc. Natl. Acad. Sci. U. S. A.*, 100(26), 16024-29.
- Markram, H. (2006). The blue brain project. *Nat Rev Neurosci.*, 7(2), 153-60.
- Mel, B. W. (1994). Information-processing in dendritic trees. *Neural Computation Neural Comput.*, 6, 1031-85.
- Potter, S. M., & DeMarse, T. B. (2001). A new approach to neural cell culture for long-term studies. *J. Neurosci. Methods*, 110, 17-24.
- Potter, S. M. (2005). Two-photon microscopy for 4D imaging of living neurons. In R. Yuste, & A. Konnerth (pp. 59-70). Cold Spring Harbor Laboratory Press.
- Potter, S. M., Wagenaar, D. A., & DeMarse, T. B. (2006). Closing the Loop: Stimulation Feedback Systems for Embodied MEA Cultures. In M. Taketani, & M. Baudry. New York: Springer.

An Extension to Subsumption Architecture Inspired by Human Nervous System

Yoosef Ramezani, Behzad Moshiri, Caro Lucas
Control and Intelligent Processing Center of Excellence, School of
Electrical and Computer Engineering, University of Tehran, Tehran, Iran
yooseframezani@yahoo.com,moshiri@ut.ac.ir,lucas@ipm .ir

Introduction

Behavior based architectures like subsumption architecture are one of the relatively new architectures try to analyze and model complicated systems with behavior approach and due to its distributed function, they have compatibility to use in various systems. For example, Rodney Brooks has introduced a way of robot control by dividing it to several layers of behavior¹. In subsumption architecture we try to find some simple behavior and model each of simple behaviors. Then find inputs and outputs for each behavior and understand what mechanisms are needed between inputs and outputs of each behavior. We can gain more complex behaviors by repeating this procedure and gain behavior of whole system. When there is a paradox between behavior outputs we make a policy that they do not activate those behaviors simultaneously. Thus it is necessary that we perform a level determination between behaviors. The superiority of behaviors can be in 2 ways: Inhabitation or Suppression. In inhabitation a behavior can inhibit the output of another behavior. In suppression a behavior can inhibit the output of other behavior and send its own output. In nature, there is another system that show similar pattern like modularization and interacting between modules, but it is more powerful. Human nervous system is the most well known system today, which has the ability to work with real world problems. These attributes are unique and can do a good complexion of data processing in different levels in existence of its variations. Thus it can be a good model for behavior based architecture. Brain and neurons function are used several times for other purposes². Now we want to introduce a novel architecture inspired by peripheral nervous system and synaptic function between neurons. We call it SS+. A neuron simply has three parts; cell body, dendrites, and axons. Data receives by dendrites and after some processes in cell body, results send by axons. It is important to note that information sent by axons has not the previous properties and all have changed to some signals. Each neuron has a basis signal. Relation between neurons occurs in synapses. Each synapse may be one of these three types: excitatory, inhibitory or facilitatory. Excitation means increasing in signal frequency or amplitude. Inhibition means decreasing in them (see Fig 1-a). Facilitation causes simpler signal transmitting in synapses. Now we can construct suppression with excitation and inhibition.

Implementation

For example, same problem of mobile robot navigation can be modeled simply by SS+ (see Fig 1-b).

Robot senses environment and goal and obstacles by its sensors and move toward the goal by sending a command to an actuator. At first we suppose three simple behaviors: *Random movement*, *Goal seeking* and *Obstacle avoidance*. Each behavior sends its output to actuator motor separately and excite or inhibit or facilitate other behavior outputs if needed. The model was implemented in an unknown environment with some fix and moving obstacle. Results were good enough and robot acted simply

in presence of goal and obstacle movement. In this model each behavior is very simple because some behavior complexity can be computed in a separated behavior and when needed, this behavior can regulate previous behavior outputs by facilitation, inhibition or excitation. This causes many behaviors can be modeled very simpler. For example, a higher level that aggregates signals come from sensors can understand how many obstacles are there in surrounding environment and if the environment is obstacle-free enough sends appropriate facilitating signal for behaviors which run actuators.

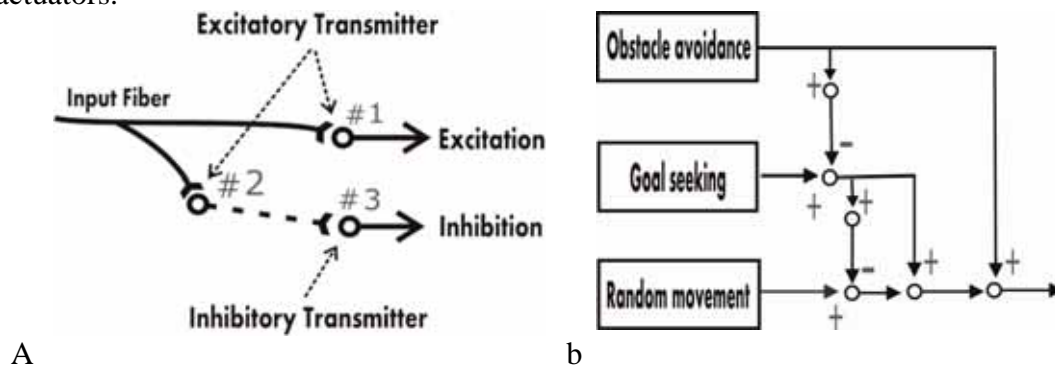


Fig. 1. a. Excitation and inhibition in synapses **b.** robot navigation problem modeled by SS+

Some benefits of SS+ are: -We can build suppression property of subsumption architecture with excitation and inhibition. -We can use excitation property for fusion purposes when there is a similarity between behavior outputs by an algebraic sum. -It is simple to model some task by facilitation property. For example, an additional higher behavior that recognize environment is obstacle free can facilitate behavior which want faster movement. A widespread facilitatory signal to all synapses lead robot to a faster mood and a widespread inhibitory signal to them lead robot to a slower mood. Thus a higher behavior level which recognize robot situation in time can change robot's mood to proper ones. -When behavior outputs are in signal form it can interact with other behavior output simply because it is more near to real world. -When we suppose basis frequency for a behavior and whole system is balanced for a situation, we can perform a change in one or more synapses in two directions, positive or negative. -With three property Excitation, inhibition and facilitation the addition of a new behavior to a designed model is very simple.

References

1. Brooks, R. A., A robust layered control system for a mobile robot, *iee journal of robotic and automation*, ra-2, 1986, pp 14-23.
2. Foss, J., *The percept and vector function theories on the brain*, philosophy of science. 1996.

Design methodologies for Central Pattern Generators: Toward “intelligent” locomotion in robots

Ludovic Righetti and Auke Jan Ijspeert
Biologically Inspired Robotics Group,
Ecole Polytechnique Fédérale de Lausanne, Switzerland
[ludovic.righetti | auke.ijspeert]@epfl.ch

Introduction

The control of locomotion in legged robot and especially in humanoid robots is the first step to embodied cognition and intelligence. This is the step that allows the robot to interact and to discover its environment. However, it is not straightforward how to design good controllers so that the robot can move in unpredictable environment. Unlike animals, robots are not really able to adapt to a changing environment. It is known from biology that the coordination of the limbs during periodic movements is done in the spine of animals [1,2]. The involved neural circuits are called Central Pattern Generators (CPGs). These are self-contained distributed neural networks that can generate all the complex signals that control the coordination of the muscles. Taking inspiration from biology led to very successful locomotion controllers [3-6]. However, the design of such CPGs remains really difficult and very few methodologies are available to construct such systems [4,5], most of the time they need extensive optimization procedures and fine-tuning. The goal of our work is to provide generic design methodologies to construct CPGs, by using the dynamical systems approach. We show through the design of a controller for a crawling baby humanoid robot, which will be used in the RobotCUB project [7], how we can use mathematical tools from the dynamical systems framework to design CPGs.

Crawling Humanoid Robots

As part of the RobotCUB project, our controller is built in order to allow the robot to explore its environment by moving on its arms and legs (i.e. crawling). The CPG is made of originally coupled oscillators. These oscillators are spring-like systems that are bounded in energy and which have a nonlinear spring constant. They exhibit limit cycle behaviour and we can control independently the duration of the ascending and descending phases of the oscillations (i.e. the duration of the swing and stance phases).

By using group-theoretic arguments [8] we can easily infer a minimal network that can generate the desired spatio-temporal pattern. With this methodology, our controller comes with generic properties which are important for robotics. The system is stable against perturbations, which will allow the integration of sensory feedback and a tight coupling with the environment. We can easily modulate the pattern in frequency and amplitude and moreover the duration of both the swing (when the limb lifts off the ground) and stance (when the limb touches the ground) phases can be controlled independently. With this simple controller we show that we are able to generate trajectories that correspond to the ones of real crawling babies and we successfully apply this to the control of the simulated iCub humanoid (which is currently under construction).

Conclusion

Although our design is specific to the control of crawling in a humanoid robot, our approach is quite general and could be successfully applied to the control of other

kinds of legged robots. The oscillator we use in association with the symmetry arguments [8] give a simple and generic method to construct a CPG for any kinds of gait for legged robot. The next step would be to study the coupling of the CPG with the environment in order to make it adaptive to changes in the environment. The structure of our controller should allow such a coupling, as was already shown in [6] and more recently in [9] where they were able to adapt parameters of the oscillators to the body dynamics.

Acknowledgements

This work was made possible thanks to a Young Professorship Award to Auke Ijspeert from the Swiss National Science Foundation (A.I.) and to the support of the European Commission's Cognition Unit, project no. IST-2004-004370: RobotCub (L.R.).

References

- [1] S. Grillner, "Neurobiological bases of rhythmic motor acts in vertebrates", *Science*, vol. 228, no. 4696, pp. 143–149, 1985.
- [2] G. N. Orlovsky, T. G. Deliagina, and S. Grillner, *Neuronal control of locomotion: from mollusc to man*. Oxford University Press, 1999.
- [3] G. Endo, J. Nakanishi, J. Morimoto, and G. Cheng, "Experimental studies of a neural oscillator for biped locomotion with qrio," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, Barcelona, Spain, 2005, pp. 598–604.
- [4] A. Ijspeert, J. Nakanishi, and S. Schaal, "Learning rhythmic movements by demonstration using nonlinear oscillators," in *Proceedings of the IEEE/RSJ Int. Conference on Intelligent Robots and Systems (IROS2002)*, 2002, pp. 958–963.
- [5] L. Righetti and A. Ijspeert, "Programmable central pattern generators: an application to biped locomotion," in *International Conference on Robotics and Automation (ICRA2006)*, 2006, pp. 1585-1590.
- [6] G. Taga, "Emergence of bipedal locomotion through entrainment among the neuro-musculo-skeletal system and the environment," *Physica D: Nonlinear Phenomena*, vol. 75, no. 1–3, pp. 190–208, 1994.
- [7] G. Sandini, G. Metta, and D. Vernon, "Robotcub: an open framework for research in embodied cognition," 2004, *IEEE-RAS/RJS International Conference on Humanoid Robotics*, Santa Monica, CA.
- [8] M. Golubitsky and I. Stewart, *The symmetry perspective: from equilibrium to chaos in phase space and physical space*. Basel, Boston, Berlin: Birkhauser, 2002.
- [9] J. Buchli, L. Righetti, and A.J. Ijspeert. "A dynamical systems approach to learning: a frequency-adaptive hopper robot". In *Proceedings of the VIIIth European Conference on Artificial Life ECAL 2005*, Lecture Notes in Artificial Intelligence. Springer Verlag, 2005

Evolutionary Robotics and Perceptual Supplementation: Dialogue Between Two Minimalist Approaches

Marieke Rohde, Ezequiel Di Paolo

Centre for Computational Neuroscience and Robotics (CCNR)

Department of Informatics, University of Sussex, Brighton, UK

[m.rohde|ezequiel]@sussex.ac.uk

Artificial intelligence (AI) is one of the disciplines constituting the interdisciplinary research field cognitive science. Classically, the metaphor of cognition as information processing has served as a bridge principle between AI and other disciplines, and this bridge is unidirectional, i.e. mental phenomena are reduced to abstract computations, physically implemented in the brain (see Fig. 1, a), thereby rendering AI modelling the intellectual core of cognitive science. We want to introduce an alternative interdisciplinary framework that relies on hermeneutic circular integration of sub-disciplines (see Fig1, b), rather than unidirectional reduction, and does not rely on the information processing metaphor as bridge principle, thereby creating space for alternative views of AI and cognition. The three areas forming the cornerstones of the explanatory triangle are phenomenological enquiry (introspection), empirical research and computational modelling.




Fig. 1. The two interdisciplinary frameworks

Varela (1996) discusses the reciprocal link between phenomenology and the empirical sciences. Here, we want to focus on the elaboration of the double link between the empirical sciences and computational modelling, considering as a specific example evolutionary robotics and perceptual supplementation. This clarification can be seen as one jigsaw piece in a larger endeavour to establish this alternative explanatory framework.

Empirical research in perceptual supplementation (PS) uses devices to “transform stimuli characteristic of one sensory modality (for example, vision) into stimuli of another sensory modality (for example, touch).” (Lenay et al. (2003), p. 2). With this technique, also known as “sensory substitution” (Bach-y-Rita, e.g. (2004)), adaptation to new sensorimotor couplings can be investigated. For instance, in Lenay (2003), the perception of space is analysed, identifying the necessary and sufficient conditions for something to be perceived as distal: Subjects are equipped with just a single photo cell attached to the finger, which controls a tactile stimulator, a minimal PS device. This minimalism, apart from providing methodological advantages, serves to highlight the role of time-extended sensorimotor coordination as the basis of perceptual invariance. The findings are explained through hermeneutic analysis of subjective experience and movement trajectories/performance. The individual consideration of either of the two methods would not have met the profoundness of the explanation thereby gained.

We believe that including computational models into the analysis will lead to an even richer account (as also recognised by the cited group, see Stewart & Gapenne (2004)). Our “house speciality” in the CCNR at the University Of Sussex is Evolutionary Robotics (ER), “a new technique for automatic creation of autonomous robots [...] inspired by the darwinian principle of selective reproduction of the fittest.” (Nolfi & Floreano (2000), preface). Typically, the models generated with this ap-

proach are deliberately minimal (Beer (2003), Harvey et al. (2006)), in order to remain tractable. They serve as tools for grounding and questioning preconceptions about fundamental aspects cognition. *Most other methods in AI do not include this explicit self-critical factor.* Beer (2003) argues that the minimalism of this method allows us to perform the necessary mental gymnastics to deal with real, dynamical, and context-dependent cognitive performance. A concern frequently uttered is: “Will these models scale up in complexity?” We believe that this desire to complicate ER models arises from the misguided ambition to make them approach the complexity of traditional theories about cognition. In fact, the minimalism of ER matches the minimalism of PS (at least as it is practiced by the Compiègne group), as both have emerged from the need for tractability and controlled settings in explaining complex cognitive phenomena. Furthermore, its inherent embodiment and situatedness makes ER a very suitable modelling technique for findings from PS research. Such models can make behavioural strategies and prior assumptions explicit and control the degree of designer intervention, thereby exploring novel principles of AI design. This is the link from PS to ER. At the same time, they can help to derive theories from empirical findings by means of abstraction. New hypotheses can be generated, influencing the design of further experiments. This is the link in the other direction, from ER to PS.

With ER falling naturally into place, our vision of a new interdisciplinary framework is, in its crude structure, complete. It does not rely on reduction of mental states, but on hermeneutic analysis, in which different disciplines inform and constrain each other. It does not need a metaphor like “cognition as information processing”, but recognises cognition as embodied activity. In this framework, the minimalism of ER and PS is not an obstacle, but a merit on the way to explain human-level cognition.

References

1. Bach-y-Rita, P. (2004): Tactile Sensory Substitution Studies. *Ann. Acad. Sci* 1013: 83-91.
2. Beer, R. (2003) The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior* 11(4), 209-243.
3. Harvey, I., Di Paolo, E., Wood, R., Quinn, M. & Tuci, E. A. (2005). Evolutionary Robotics: A new scientific tool for studying cognition. *Artificial Life*, 11(1-2):79-98.
4. Lenay C. (2003) Ignorance et suppléance : la question de l'espace, HDR 2002, Université de Technologie de Compiègne.
5. Lenay C., Gapenne O., Hanneton S., Marque C. & Genouëlle C. (2003) Sensory Substitution, Limits and Perspectives. In: *Touch for Knowing*, John Benjamins Publishers, Amsterdam.
6. Nolfi, S & Floreano, D. (2000). *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge MA: MIT Press.
7. Stewart, J. & Gapenne, O. (2004). Reciprocal modelling of active perception of 2-D forms in a simple tactile-vision substitution system. *Minds and Machines*, 14(3), 309-330.
8. Varela, F. (1996) Neurophenomenology: A methodological remedy for the hard problem. *J. Consc. Studies* 3, 330-350.

Adaptive Mechanics: Compliant Legged Locomotion

Juergen Rummel*, Andre Seyfarth*, Fumiya Iida*[†], Elmar Dittrich*

*Locomotion Laboratory, University of Jena,
Dornburger Str. 23, 07743 Jena, Germany

[†]Artificial Intelligence Laboratory, University of Zurich,
Andreasstrasse 15, 8050 Zurich, Switzerland
[juergen.rummel | andre.seyfarth]@uni-jena.de

Introduction

What are the underlying principles for adaptive locomotion behaviors of humans and animals? Is a complex control strategy of walking and running implemented in the brain and spinal cord to follow pre-defined trajectories, or do we need a specific leg design for stable locomotion? Previous studies have shown that the leg acts like a spring during running^{1,2} and simple control strategies with little sensory feedback may lead to stable and robust locomotion^{3,4}. Based on these theories we built and investigated a series of simple single-legged and bipedal robots. In this paper, we introduce an overview of our robot project and its results.

Single-Legged Robots

From simulations of a simple spring-mass model we know that a leg needs control during flight but not during stance phase. The natural dynamics of the system play an important role for stabilizing periodic movements. To test this strategy we use a symmetric leg that is designed similarly to the spring-mass system (see Fig.1a). It consists basically of a springy telescope leg and a motor at the hip. With this construction and a swing-leg retraction control⁴, we can show that hopping over obstacles can be stabilized after one or two steps.

In a second approach on single-legged applications we investigate asymmetric legs. We started with a two-segmented system and one passive joint (see Fig. 1b). These kinds of robots also contain one motor at the hip but there is no sensory information about kinematics or ground interaction. Here we implement a simple control algorithm that swings the upper segment back and forth in a symmetric manner. We observe that the asymmetric leg has a preferred direction for hopping which changes by using a higher control frequency⁵. Interestingly, the same control strategy can be observed in human walking and running where the biological leg is understood as a compliant system.



Fig. 1. a) Single-legged robot testbed with telescope leg. b) Single-legged Fujubot with an asymmetric leg design. c) The bipedal robot JenaWalker.

Bipedal Robots

The third type of applications, we describe here, are bipedal robots. The leg geometry is more related to human legs with three segments. Both hips are directly driven by motors as described above. The cheap design version JenaWalker shown in Fig. 1c is built with passive knee and ankle joints. The segments are linked with biarticular springs that are comparable with muscle-tendon units in human legs. For driving the hip joint we use the same feed-forward control strategy as mentioned above, namely swinging the upper segment symmetrically back and forth. While this control is very simple and the leg is passively compliant, we find movement patterns similar to human walking⁶. Moreover, system motion patterns are robust against parameter changing. This robot suggests that the compliant leg design becomes more important than a precise feedback control to achieve stable and robust locomotion.

Conclusion

By understanding the principles of legged locomotion, we obtain a number of additional insights into the control mechanisms of animals' adaptive behaviours. In particular, the simple case studies of asymmetric and compliant mechanics showed that legged locomotion would not require complex control strategies. A considerable part of adaptive intelligent behaviours in animal and robot locomotion can be achieved through the physical interaction derived from leg designs, simple actuation and the environment.

Acknowledgments

This research is supported by an Emmy-Noether grant of the German Research Foundation (DFG) to Andre Seyfarth: (SE1042/1).

References

1. Blickhan, R. *The spring-mass model for running and hopping* (J Biomech 22, 1217-1227, 1989).
2. McMahon, T.A., Cheng G.C. *The mechanics of running: How does stiffness couple with speed?* (J Biomech 23, 65-78, 1989).
3. Seyfarth, A., Geyer, H., Guenther, M. and Blickhan, R. *A movement criterion for running* (J Biomech 35, 649-655, 2002).
4. Seyfarth, A., Geyer, H. and Herr, H. *Swing-leg retraction: A simple control model for stable running* (J Exp Biol 206, 2547-2555, 2003)
5. Rummel, J., Iida, F. and Seyfarth, A. *One-legged locomotion with a compliant passive joint* (Intelligent Autonomous Systems 9, Arai et al, eds., IOS Press, 566-573, 2006)
6. Iida, F., Minekawa, Y., Rummel, J. and Seyfarth, A. *Toward a human-like biped robot with compliant legs* (Intelligent Autonomous Systems 9, Arai et al, eds., IOS Press, 820-827, 2006)

Self-Organization of Complexity in Embodied Agents

Alexander Schmitz and Fumiya Iida
Artificial Intelligence Laboratory
Department of Informatics, University of Zurich
Andreasstrasse 15, CH-8050 Zurich, Switzerland
alex@cognitivescience.at, iida@ifi.unizh.ch

Introduction

Increasing the complexity of autonomous agents is one of the most significant challenges in order to understand the nature of intelligent adaptive behaviors. For achieving diverse and adaptive behaviors in many different kinds of uncertain environments, biological systems have a substantially higher level of complexity in anatomical and physiological structures, for example¹. The goal of this project is to explore simple mechanisms for generating and handling such complexity in artificial agents. In particular, we explore three domains of self-organization processes. We demonstrate how simple underlying mechanisms can produce non-trivial behaviors by taking advantage of the interaction between the environment, the control and the body dynamics.

(i) Body Dynamics

Exploiting the physics of the system-environment interaction is an important basis for adaptivity in autonomous agents². We demonstrate through two case studies how passive dynamics and the interaction with environment (i.e. friction) can be used on the one hand for generating behavioral diversity, and on the other for simplifying the control.

We built a wheeled robot that can go to every point in space only by changing the speed of one motor. This robot goes straight when it moves forward but turns when it goes backward because of the rolling friction. The control of this robot becomes simpler than generally conceivable. This principle can also be generalized in a fish robot that can swim to every point in 3D space by using the fluidic friction.

The running dog robot “Puppy” shows another way to accomplish complex behaviors by exploiting morphology in non-trivial ways. It achieves relatively robust rapid legged locomotion without any need of sensors by using the body dynamics induced by the elastic properties of the mechanical system. The control can be extremely simple and be handled by a basic feedforward motor pattern, while the passive dynamics stabilize the system using immediate feedback through the body.³ Morphological properties are also important to increase behavioral complexity: the robot is able to achieve high jumps by extending the length of the hind legs.

(ii) Adaptive control and behavioral diversity

Control without sensors is parsimonious. Although it is sufficient for stable behavior, it is not always optimal, since it is dependent on exploiting a relatively stable environment. In general, the more the control is pre-determined by the human designer, the less adaptive an agent will be. Therefore learning is essential for an autonomous agent to achieve behavioral diversity in many unknown environments.

The control architecture of Puppy makes it easy to generate adaptive behaviors consistent with the body dynamics. A simulation model of Puppy showed that the relation between the speed and a control parameter (i.e. the phase delay between the hind- and forelegs) is nearly monotonous. We used simple hebbian learning to find correlations

between control parameters and the resulting behavior. After a short learning phase the simulated robot was able to run with a desired speed.

A camera and more elaborated hebbian learning in the real-world dog-robot are used to explore more correlations induced by the body dynamics. By now the robot is able to turn slightly to follow objects. We will change the morphology to enable turning within a very small radius. Puppy will learn a predictive model of its environment by using internal simulations of the interaction with the environment. After a short learning period it will be able to run quickly without hitting objects in new environments.

(iii) Self-organization of body structures

By using self-organization processes of small building blocks, we are able to simplify the design processes of more complex agents. Self-assembled hierarchical structures at all scales (protein-folding, cell-membrane, tissue, etc.) are a common phenomenon in living creatures. They enable complex adaptive structures such as muscles and densely packed sensors. Agents become smaller and more reliable through more redundancy and self-repair.

Cell membranes can be easily formed artificially by using self organization.⁴ Currently we are exploring simple methods to make vesicles attach to each other by electrostatic adhesion. This process so far only enables cells to form relatively small clusters, but it is an first attempt toward more complex structures.

Discussion and conclusion

In this project we are exploring only a few case studies, and there are a number of different aspects of self-organization processes which need to be explored in the future. For example, although we explored only one kind of time-scale in each of the case study, it would be very important to understand how self-organization processes in two different time scales (e.g. development processes of self-assembling of microstructures like cells and macroscopic musculoskeletal structure) work together to achieve more sophisticated adaptive behaviors. As a conclusion, although we are still in nascent stage of exploration, it could be said that self-organization at all these scales is necessary for intelligent adaptive behaviors.

References

1. Pfeifer, R. & Scheier, C. *Understanding Intelligence* (MIT Press, Boston, 1999).
2. Dickinson, M. H. et al. *How animals move: an integrative view*. (Science vol 288, April 2000, pp. 100-106, 2000)
3. Iida, F., Gomez, G. J. & Pfeifer, R. *Exploiting body dynamics for controlling a running quadruped robot*. In *International Conference of Advanced Robotics* (ICAR, 2005).
4. Noireaux, V. & Libchaber, A. *A vesicle bioreactor as a step toward an artificial cell assembly* (Proc. Natl Acad. Sci. USA 101 17669–74, 2004)

Applying OWA Operator in a Multi-Agent Architecture of an Emotional Robot

Mohamed Shayganfar*
Department of Engineering,
Islamic Azad University,
Science and Research branch,
Tehran, Iran
mohamad@shayganfar.com

Behzad Moshiri, Caro Lucas
Control and Intelligent Processing Center of Excellence,
School of Electrical and Computer Engineering,
University of Tehran,
Tehran, Iran
moshiri@ut.ac.ir | lucas@ipm.ir

1. Introduction

All Intelligent creatures that we know of have emotions. Humans, in particular, are the most expressive, emotionally complex, and socially sophisticated of all. There have been many computational models of artificial emotions using different techniques to implement the concept of emotional agents through these years. In this paper, as a modification of the applied architecture, we propose applying **Ordered Weighted Averaging (OWA)** operator in the procedure of both internal and environmental assessments of releasers, which all are considered with respect to the robot's wellbeing and its goals according to the architecture of the MIT's sociable robot (Kismet)¹. Each Releaser can be thought of as a simple "cognitive" assessment that combines lower-level perceptual features into behaviorally significant perceptual categories.

2. Proposed Approach

It is common for biologically inspired architectures to be constructed from a network of interacting elements (e.g., subsumption architecture, neural networks, or agent architectures). The agent architecture is implemented where each computational element is conceptualized as a specialist. Hence, each drive, behavior, perceptual releaser, motor and emotion-related process is modeled as a different type of specialist that is specially tailored for its role in the overall system architecture². (see Figure 1)

In this modified architecture each releaser is evaluated based on different obtained factors from other parts of the overall architecture. Particularly, *drives* which are implemented in three distinct processes of social, stimulation, and fatigue styles, will show the un/desirability status of the given stimuli to the robot, *current affective state* of the robot will reduce the misclassification of the releasers' activation, *behavioral state* of the robot also plays an important role in disambiguating certain perceptual conditions, and finally robot's *perceptual state* can contribute to the

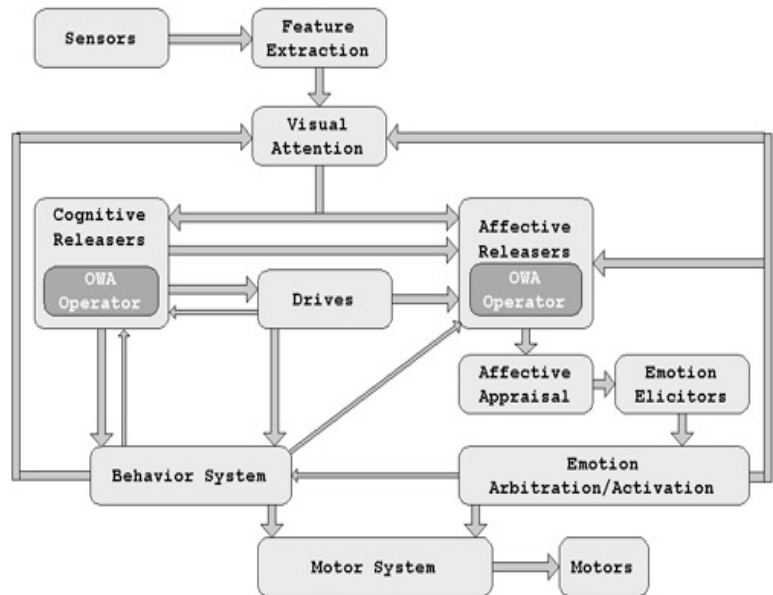


Fig. 1. Modified proposed architecture of the MIT's sociable robot

affective state on their own or in combination with other stimuli. Hence, it will be more biologically inspired if the releasers as a decision-making subsystem in this architecture could fuse the possible diverse decisions from the other decision subsystems.

Yager introduced the OWA operator to provide a family of aggregators having the properties of mean operators ³. One of the key points in the OWA operator is to determine its associated weights. A number of methods have been developed to obtain the OWA weights ⁴.

The main problem is releasers activation in our proposed architecture that consists of four components. The first component is a collection of handcrafted releasers, $X = \{x_1, \dots, x_p\}$ that could be activated. The second component is a collection of 5 criteria relevant in the ranking process including drives, perceptual, behavioral, affective states and the goal of the robot. The third component is a group of 9 ordinary people whose opinions are solicited in ranking the alternatives, which is considered due to the importance of the believability parameter of a sociable robot. The last component consists of 3 experts whose opinions solicited in ranking the same alternatives and can be more reliable as a parameter of decision-making. All these people are asked to provide an evaluation of the alternatives. This evaluation consists of a rating for each alternative on each of the criteria, where the rating are chosen from the scale $\{1,2,3,4,5\}$, where 5 stands for the most relevant, 4 stands for more relevant, 3 stands for a neutral relevancy, 2 stands for less relevant, and 1 stands for the least relevant criterion. Consequently, each person provides a 5-tuple (a_1, \dots, a_5) for each releaser. The next step is to find the overall evaluation for an alternative by a given person. In this stage we aggregate the individual persons' evaluations to obtain an overall value for each releaser. It means that we only have a 9-tuple (b_1, \dots, b_9) which is the result of ordinary people's opinions, and a 3-tuple (b_a, b_b, b_c) which is the result of experts' opinions. In the next stage we aggregate the ordinary people and experts evaluations to obtain an overall value for each releaser.

Yager suggested a method to compute the weights of the OWA operator using linguistic quantifiers. The quantifier is considered as a **Regular Increasing Monotone (RIM)** linguistic quantifier Q :

$$Q_\alpha(r) = r^\alpha, \alpha \geq 0$$

It is supposed that the most criteria quantifier is $Q(r) = r^2$.in our project

Taking into consideration that we have 5 criteria, the weights derived from Q_α , representing the statement **most criteria**, are determined by:

$$w_i = Q\left(\frac{i}{n}\right) - Q\left(\frac{i-1}{n}\right), \quad i = 1, 2, \dots, n$$

References:

1. Cynthia L. Breazeal, Designing Sociable robots, (MIT Press, 2002)
2. Cynthia L. Breazeal, Emotion and Sociable Humanoid Robots, (Int. J. of Human-Computer Studies, Vol. 59, 2003, 119 -155)
3. Ronald R. Yager, On Ordered Weighted Averaging Aggregation Operators in Multicriteria Decision making, (IEEE Transaction on Systems, Man and Cybernetics, 1988)
4. Zeshui Xu, An Overview of Methods for Determining OWA Weights, (Int. J. of Intelligent Systems, Vol. 20, 2005, 843-865)

A Philosophy of Designing Human Interactive Robot for Psychological Enrichment and Robot Therapy

Takanori Shibata

Intelligent Systems Research Institute, AIST
SORST, JST

There are two categories of robots; one is industrial robot and the other is service robot. When we design the industrial robot, we evaluate it in terms of objective measures such as speed, accuracy, and cost. The industrial robot is considered as danger for humans and isolated from them. On the other hand, as the service robot interact with humans, we evaluate it in terms of subjective measures such as interesting, comfortable, and beautiful as well as the objective measures. There are two sub-categories of service robot; one is the physical service and the other is psychological service. The applications of robots for physical service are guidance, physical support for walking, and so on, and they have more weight on the objective measures than the subjective measures. On the other hand, those for psychological service are companion, communication partner, entertainment, mental therapy and so on, and they have more weight on the subjective measures. In the presentation, I will explain a philosophy of designing human interactive robot for psychological enrichment. First, I will explain results of psychological experiment that show importance of evoking human's association in interaction with a robot. Second, I will compare design strategies of human interactive robots such as Paro (Fig. 1) and AIBO. Third, I will explain results of subjective evaluation of the robots by humans. Forth, I will explain applications of robot for mental therapy at hospitals and elderly institutions (Figs. 2 and 3). Finally, I will summarize ways of designing human interactive robot for psychological enrichment and robot therapy.



Fig. 1 Seal Robot, Paro (<http://paro.jp>)



Fig. 2 Robot Therapy at Karolinska Hospital, Sweden



Fig. 3 Robot Therapy at a Nursing Home, Japan

How Should the Body and the Brain be Coupled?

- A Robotic Case Study with a Modular Robot -

Masahiro Shimizu*, Takafumi Mori*, Toshihiro Kawakatsu**, and Akio Ishiguro[†]

*Dept. of Computational Science and Engineering

Nagoya University, Nagoya, Japan

**Dept. of Physics

[†]Dept. of Electrical and Communication Engineering

Tohoku University, Sendai, Japan

shimizu@cmlpx.ecei.tohoku.ac.jp, moritaka@cmlpx.cse.nagoya-u.ac.jp,

kawakatsu@cmpt.phys.tohoku.ac.jp, ishiguro@ecei.tohoku.ac.jp

Abstract - This paper discusses experimental verifications of a two-dimensional modular robot called “Slimebot”, consisting of many identical modules. The Slimebot exhibits adaptive reconfiguration by exploiting a fully decentralized algorithm able to control its morphology according to the environment encountered. One of the significant features of our approach is that we explicitly exploit “emergent phenomena” stemming from the interplay between control and mechanical systems in order to control the morphology in real time. To this end, we particularly focus on a “functional material” and “mutual entrainment” among nonlinear oscillators, the former of which is used as a spontaneous connectivity control mechanism between the modules, and the latter of which acts as the core of the control mechanism for the generation of locomotion. Experimental results indicate that the proposed algorithm can induce locomotion, which allows us to successfully control the morphology of the modular robot in real time according to the situation without losing the coherence of the entire system.

Recently, a modular robot (or reconfigurable robot), consisting of many mechanical units (hereinafter called modules), have been attracting lots of attention. Since the relative positional relationship among the modules can be altered actively according to the situation encountered, a modular robot is expected to show significant abilities, e.g., adaptability, fault tolerance, scalability, and flexibility, compared with a robot on a fixed-morphology basis¹. In order to fully exploit the advantages mentioned above, (1) each module should be controlled in a fully decentralized manner, and (2) the resultant morphology of the entire system should emerge through the module-to-module and module-to-environment interactions.

In light of these facts, this study is intended to deal with an emergent control method which enables a modular robot to change its morphology in real time according to the situation encountered without the use of any global information as well as without losing the coherence of the entire system. Since there still remains much to be understood about how such emergent systems can be created, in this study, we employ the following working hypothesis:

Well-balanced coupling between control and mechanical systems plays an essential role to elicit interesting emergent phenomena, which can be exploited to increase adaptability, scalability, fault tolerance, and so on.

Based on this working hypothesis, we have so far developed a two-dimensional modular robot, called Slimebot¹. In this study, in order to realize an emergent control method, the coupling between the control and mechanical systems of Slimebot has

been carefully designed as follows: we have particularly focused on a functional material, i.e., a genderless Velcro strap, and mutual entrainment among nonlinear oscillators, i.e., van der Pol(VDP) oscillators, the former of which is used as a spontaneous connectivity control mechanism between the modules, and the latter of which acts as the core control mechanism for the generation of locomotion and ensures the scalability. Simulation results indicate that the proposed method can induce amoebic locomotion, which allows us to successfully control the morphology of the modular robot in real time according to the situation without losing the coherence of the entire system (see Fig. 1).

To verify the feasibility of our proposed method, experiments with a real physical Slimebot are also significantly important. In this paper, we explain how we have designed the hardware of Slimebot as an autonomous decentralized system and show a first preliminary experimental result from the view point of real-time adaptive reconfiguration in an environment containing obstacles (see Fig. 2). Here, the Slimebot negotiates its environment without losing the coherence of the entire system while each module changes its connectivity between modules spontaneously. Since the experimental study is still in the initial stage, this paper deals with the locomotion of Slimebot consisting of several real physical modules. The experimental result, however, includes intrinsic emergent property that enables adaptive behavior by coupling between the control and mechanical systems appropriately.




Fig. 1. Representative data of the transition of the morphology in the case of 500 modules (see from left to right). Note that no active control mechanism that precisely specifies connection/disconnection among the modules is implemented.




Fig. 2. Adaptive reconfiguration with 4 modules. See from (a) to (f). We can see a typical example of the spontaneous connectivity control provided by the functional material (see the rear module).

References

1. Fukuda, T. & Kawachi, Y. *Cellular Robotic System (CEBOT) as One of the Realization of Self-Organizing Intelligent Universal Manipulators* (in Proc. of the 1990 IEEE International Conference on Robotics and Automation, pp.662-667, 1990).
2. Ishiguro, A. Shimizu, M. & Kawakatsu, T. *Don't Try to Control Everything!: An Emergent Morphology Control of a Modular Robot* (in Proc. of the International Conference on Intelligent Robots and Systems 2004(IROS2004), pp.981-985, 2004).

Active Vision for Neural Development and Landmark Navigation

Mototaka Suzuki and Dario Floreano

Ecole Polytechnique Fédérale de Lausanne (EPFL)

Laboratory of Intelligent Systems, CH-1015 Lausanne, Switzerland

[Mototaka.Suzuki | Dario.Floreano]@epfl.ch

Introduction

Brains and sensory systems are characterized by limited bandwidth and computational resources. At any point in time, we can focus our attention only to a limited set of features or objects. One of the most remarkable –and often neglected– differences between machine vision and biological vision is that computers are often asked to process an entire image in one shot and produce an immediate answer whereas animals are free to explore the image over time searching for features and dynamically integrating information over time.

Coevolution of Active Vision and Feature Selection¹

We show that the co-evolution of active vision and feature selection can greatly reduce the computational complexity required to produce a given visual performance. *Active vision* is the sequential and interactive process of selecting and analyzing parts of a visual scene. *Feature selection* instead is the development of sensitivity to relevant features in the visual scene to which the system selectively responds. Each of these processes has been investigated and adopted in machine vision. However, the combination of active vision and feature selection is still largely unexplored.

In our experiments behavioral machines equipped with primitive vision systems and direct pathways between visual and motor neurons (Fig. 1) are evolved while they freely interact with their environments. We describe the application of this methodology in three sets of experiments, namely, shape discrimination, car driving, and robot navigation. We show that these systems develop sensitivity to a number of oriented, retinotopic, visual-feature-oriented edges, corners, height, and a behavioral repertoire. This sensitivity is used to locate, bring, and keep these features in particular regions of the vision system, resembling strategies observed in simple insects.

Active Vision and Visual Development^{2,3}

In a further set of experiments we investigate the *ontogenetic development* of receptive fields in an evolutionary mobile robot with active vision. In contrast to the previous work where synaptic weights for both receptive field and behavior were genetically encoded and evolved on the same time scale, here the synaptic weights for receptive fields develop during the life of the individual. In these experiments, behavioral abilities and receptive fields develop on two different temporal scales, phylogenetic and ontogenetic respectively. The evolutionary experiments are carried




Fig. 1. The architecture for active vision and feature selection.

out in physics-based simulation and the evolved controllers are tested on the physical robot in an outdoor environment (Fig. 2).

Such a neural architecture with visual plasticity coupled with a freely moving behavioral system allows us to explore the role of active body movement in the formation of the visual system. More specifically we study the development of visual receptive fields and behavior of robots under active and passive movement conditions. We show that the receptive fields and behavior of robots developed under active condition significantly differ from those developed under passive condition. A set of analyses suggest that the coherence of receptive fields developed in active condition plays an important role in the performance of the robot.



Fig. 2. The Koala mobile robot by K-Team S.A. with a pan/tilt camera in an outdoor environment.

Active Vision and Sequential Landmark Detection⁴

Lastly, active vision may also be useful to perform landmark-based navigation where landmark relationship requires active scanning of the environment. Here we explore this hypothesis by evolving the neural system that controls the vision and behavior of a mobile robot equipped with a pan/tilt camera so that it can discriminate visual patterns and arrive at a predefined goal zone. The experimental setup employed here requires the robot to actively move its gaze direction and integrate information over time in order to accomplish the task.

We show that the evolved robot can detect two separate features in a sequential manner and discriminate the spatial relationships. Since the system can perform active vision and sequentially store the events of visual feature detection, we do not need expensive computational power nor large memory storage capacity which would be required to resort to image memorization and matching. Although there is evidence that insects may indeed adopt such an image memorization and matching strategy, it is tempting to speculate that their tiny brain with restricted memory capacity may favor a more economical strategy such as shown here.

References

1. Floreano, D., Kato, T., Marocco, D. and Sauser, E. (2004). Coevolution of active vision and feature selection. *Biological Cybernetics*, 90(3):218-228.
2. Floreano, D., Suzuki, M. and Mattiussi, C. (2005). Active vision and receptive field development in evolutionary robots. *Evolutionary Computation*, 13(4):527-544.
3. Suzuki, M., Floreano, D. and Di Paolo, E. A. (2005). The contribution of active body movement to visual development in evolutionary robots. *Neural Networks*, 18(5/6):656-665.
4. Suzuki, M. and Floreano, D. (2006). Evolutionary active vision toward three dimensional landmark-navigation. In *Proceedings of the Ninth International Conference on the Simulation of Adaptive Behavior (SAB'06)*, September 25-29, 2006, Roma, Italy, Springer Verlag, in press.

Non-Linear Models for the Dynamics of Motivation

Rene te Boekhorst

Biological Neural Computation Group and Adaptive Systems Research Group
The University of Hertfordshire, School of Computer Science
College Lane, Hatfield, Hertfordshire AL10 9AB, United Kingdom
r.teboekhorst@herts.ac.uk

Based on a review on the effect of ethology on the development of AI, I claim that progress in “biologically-motivated” robotics has been limited because it is insufficiently realized that biologists are more inspired by A.I. and cognitive science than the other way around. By adopting conventional ethological theory in the design of robot architectures, roboticists disappointed by the shortcomings of traditional AI thus unwittingly reintroduce the overly rationalist view they would like to get rid off in the first place.

As an alternative, I propose to start from a simple, general model for studying the dynamics of motivation and suggest its implementation in robots to develop new insights in the study of behaviour. The incentive is that: 1) concepts of classical ethological lack unification, are too much based on black box reasoning and therefore do not link up with physiology as it claimed to be and 2) the alternative rationalist-analytic approach often leads to superfluous and contrived explanations. These shortcomings are due to the habit of seeking separate explanations for each observed phenomenon and the tendency to ascribe behaviour patterns solely to cognitive or genetic qualities of individuals. I will illustrate how dynamical systems models circumvent these drawbacks and generate insights into the generation of behaviour by bringing together ethological and (neuro-) physiological concepts that were hitherto thought to be disconnected. The hypotheses derived in this way are parsimonious in that a multitude of patterns can be traced back to one and the same minimal set of interactive dynamics. As mechanistic implementations of principles discovered in formulae and silica, robots form a critical extension to mathematical models and simulations because they confront us with important real world conditions and physical constraints that are hard to program or would go otherwise unnoticed. The combination of dynamical systems modeling and the implementation of these models in robots (of which the behaviour is then studied as if they were animals!¹) should therefore lead to deeper explanations than the functionalistic top-down approaches of cognitive science and neo-darwinian evolutionary theory.

Background and Scope of the Framework

Bateson² acutely noted that “scientists studying behaviour are faced with dynamical systems that have an awkward way of altering their characteristics when conditions change. The way to study such systems is studying them as *processes*, not by taking snapshots or by abstracting linear causal chains”. To set the stage for this, I will sketch the development of ethology from its simple ideas about reactive behaviour to theories of motivation with ramifications to cybernetics, information theory, statistical analysis and Markov modelling of temporal patterns, computational views (Artificial Intelligence and cognitive science) and the rationalist, neo-Darwinist stance. I will point out the epistemological relationships between these various approaches, the split between mechanistic and rationalist ethological theories and, conform to Bateson’s statement above, identify a dynamical systems approach with the former. Next, I will suggest the first steps towards such formalization. It will form the basis for a general framework which should bring together concepts from as well from ethology as

physiology. Rather than taking systems with known mathematical properties “from-the-shelf” and reword them into appropriate interpretations, the formal framework to be developed here builds up from “first principles”. It takes the form of non-linear differential equations that model the changes in motivational intensity and is based on simple ethological considerations about energy allocation, feedback and the working of external stimuli. The increase in motivational intensity is steered by the fraction of available energy that remains after part of it has been used up by the current level of motivational intensity. This fraction itself is modelled as a Hill function, which gives the system a “physiological flavour”. The resulting cubic differential equation already unites certain features of behaviour that hitherto were subject to separate interpretations. For instance, it acts as a threshold because an unstable intermediate equilibrium arises as a consequence of the systems' own dynamics and thus allows catastrophic switching (cusp catastrophes have been used as a phenomenological description for behavioural switching in models by Zeeman³ and have been investigated subsequently by behavioural scientists⁴ and physiologists⁵). The bifurcation parameter responsible for changing the number of fixed points by shifting up the unstable equilibrium (and hence the threshold) can be linked to the intensity of another motivation or the level of external stimulation. “Competition” among motivations, resulting in a “winner-takes-all” type of action selection, follows as a direct extension of the standard model by letting another motivation tap from the same energy source, leading to a kind of Lotka-Volterra dynamics. By being structurally similar to the Fitzhugh-Nagumo equations for neuronal activity the model has a rich repertoire of dynamics including relaxation oscillations and excitable behaviour. Finally, I will clarify how the framework may incorporate the notion of endorfine/dopamine driven “limits of motivational reward”⁶ and thus links it to theoretical investigations of emotion. Whether this simple framework leads to “interesting” (i.e. counter-intuitive and hence spurring further investigation) real-world behaviour can only be found out by implementing it in a physical robot. I will point out the type of experiments with robots we plan to perform at the Adaptive Systems Research Group (University of Hertfordshire) as part of our contributions to the HUMAINE Network of Excellence on emotion research and the ROBOTCUB Framework VI Integrated project for developmental robotics.

References

1. Boekhorst, I.J.A. (2001). Freeing machines from Cartesian chains. In: “*Cognitive Technology: Instruments of Mind*” (eds. M. Beynon, C.L.Nehaniv & K. Dautenhahn). Lecture Notes in Artificial Intelligence 2117. Springer, Berlin, p. 19-39.
2. Bateson, P. (1991). Are there principles of behavioural development? In: “*The Development and Integration of Behaviour. Essays in Honour of Robert Hinde*” (ed. P. Bateson). Cambridge University Press, Cambridge, p. 19 – 39.
3. Zeeman, E.C. (1976). Catastrophe theory. *Sci Amer.*, 234, 4, 65-83.
4. Koene, P. (2005). Stochastic catastrophe analysis of effects of density on behaviour. An example in laying hens. In: “*Measuring Behavior 2005*”, 5th International Conference on Methods and Techniques in Behavioral Research, Wageningen, the Netherlands
5. Kikuchi, A., Unno, N., Horikoshi, T., Kozuma, S. & Taketani, Y. (2006) Catastrophe theory model for decelerations in fetal heart rate. *Gynecologic and Obstetric Observations*, 61, 72-79.
6. Esch, T. & Stefano, G. B.(2004). The neurobiology of pleasure, reward processes, addiction and their health implications. *Neuro Endocrinol Lett* 4, 25, 235-251.

Towards an Experience-Based Artificial Intelligence

Ricardo A. Téllez*, Oscar Vilarroya*[†]

Departments of *Engineering of the Technical University of Catalonia
and [†]Cognitive Neuroscience Unit of the Autonomous University of Barcelona
Barcelona, Spain
rtellez@lsi.upc.edu | oscar.vilarroya@uab.es

Certain theories of cognitive development, of the evolution of cognition, and of knowledge representation (Nelson 1895, Barsalou 1999, Donald 1991, among others) have indicated that the episode is a central element to understand the first stages of cognitive development, as well as of certain basic cognitive abilities, such as intelligent behavior. Developing these approaches, we introduce the experion theory of cognition (Vilarroya 2002) and we set the future lines of research for applying this theory to the generation of artificial intelligence.

I. The Real Meaning of an Experience

Let's define experion as a slice of the life of a cognitive being, with a limited duration constrained by (bottom-up) neural synchronization and (top-down) attentional dynamics (Vilarroya, in press). An experion would consist of the web of all the states of all sensors, motors, emotions, internal states and motivations (understanding motivation as an internal drive of the being that generates desired states for the being), conscious and unconscious, of the being at a particular moment. The states of all these elements are a function of certain physical constraints of the system, of certain predispositions of the system, of the dynamical interaction between the different elements and of previous experions.

All experions experienced during life are stored in the brain of the being. The cognitive process of a being is then seen as a concatenation of experions. This process of experion accumulation establishes relationships (i.e. similarity) between the present experion and past experions so that the experion is stored modifying its nature and that of relevant past experions according to the type of relationships established.

All the experion storage has as goal to be an action selection mechanism, which should result on a maximal survival of the being. The final motor answer given by a being on a determined situation is then produced by the activation of all the present experions that the being has accumulated during his life, on that particular situation and moment.

II. Towards Experion Based AI

The experion theory is presumed to be underneath all biological cognition, from the most simple up to the most complex. Therefore, it should be possible to generate an artificial being that behaves in that way using the experion theory. In order to do that, we identify the following points to be addressed for an artificial intelligence based on that theory:

1. The AI must be able to acquire/generate experions from its sensors, actuators, internal states and motivations. How the process which gives rise to an experion should be implemented?
2. The AI must be able to generate its concepts about life from its experions. How should those concepts be generated from different experions?
3. One of the most important mechanisms of this theory is the one that finds relations between several experions that are very different from a panceptual state (Vilarroya 2002), but that are very similar from a conceptual point of view (i.e. one is a metaphor of the other, Lakoff & Nuñez, 2002). Until now, most of the artificial

- systems that tried to find such relations were based on symbolic methods, but basing the analysis on experions may provide a new light. How could relations be found between different experions which are indeed related?
4. When generating a new experion, a lot of variables take part in the process, but only a few are really important for the present situation. This is what is called the figure/ground relation. How a figure/ground relation is established from the current experion?
 5. Finally, all the mechanisms explained must have a reason to be there, i.e. they must drive the artificial entity on its environment. How do the relations among all these mechanisms and the action selection mechanism are established in order to help the AI survive in its environment?

These are the lines that will trace our research in the following years.

References

- Barsalou, L.W. *Perceptual symbol systems*. Behavioral and Brain Sciences, 22:577-609, 1999
- Donald, M. *Origins of the Modern Mind: Three stages in the evolution of culture and cognition*. Harvard University Press, 1991.
- Johnson, M. *The body in the mind*. University of Chicago Press, 1987.
- Lakoff, G. & Nuñez, R. *Where mathematics comes from*. Basic Books, 2000.
- Nelson, K. *Making sense. The acquisition of shared meaning*. Academic Press, 1985.
- Nelson, K. *Language in cognitive development. Emergence of the mediated mind*. Cambridge University Press, 1996.
- Schank, R.C. and Abelson, R.P. *Scripts, plans, goals and understanding: an inquiry into human knowledge structures*. Erlbaum, 1977.
- Vilarroya, O. *The dissolution of mind: A fable of how experience gives rise to cognition*. Editions Rodopi, 2002.
- Vilarroya, O. *The experion time constraints*. In press

An Ultimate Goal Of AI And Tasks For Four Periods

Jiabei Wu

Chengdu Textile College, San Wayao Street, Chengdu 610063, P. R. China

Visiting scholar at the Artificial Intelligence Lab, University of Zurich

jiabei@ifi.unizh.ch

1. Introduction

Last year, as a visiting scholar I had the opportunity to study at the Artificial Intelligence Lab of the University of Zurich. In this best academic circle, I have learned about new developments in Artificial Intelligence, and I have begun to consider what a goal of AI could be. I now suggest an ultimate goal of AI and tasks for four periods.

2. The Ultimate Goal of AI

What are the big problems of humanity? In today's world, war and hunger are almost overcome, so the human's enemy is no longer himself. Sometimes there are viruses, hurricanes, earthquakes, and other attacks. But the biggest threat is that our sun, being halfway through its lifespan, will only last for another 4.5 billion years. (<http://en.wikipedia.org>).

What could be an ultimate goal of AI? As there are so many scientists working in this field, it should contribute more to humanity. By the inspiration of nature's organisms, which always want to survive, and inspired by Japanese psychologist Masanao Toda's "Fungus Eaters", I suggest that an ultimate goal of AI is "that humans should survive in the Cosmos".

This is a difficult goal. Human natural intelligence is not enough to achieve this goal. So we have to develop artificial intelligence.

This is a complicated goal. Researchers of only one discipline are not able to achieve this goal. So we have to set up interdisciplinary study researchers. Just like the AI Lab of University of Zurich.

This is an urgent goal, and maybe most people think it is far away, but it depends on the time perspective. Considering the long way research still has to go, you will find it's really urgent.

3. Tasks for the Four Periods of Artificial Intelligence

How to achieve the ultimate goal? In my opinion, the goal is divided into four periods, each of which contains different tasks and topics.

1) The period of the Turing machines. In this period, researchers wanted to use computers to study human cognition, and they focused on information processing. Although computers can be used to simulate virtually any natural process, including brain processes, as a simulation device the computer is not used as a metaphor for intelligence but only as a formal tool. The idea of computation was formalized by Alan Turing, so we call this period the period of the Turing machines. The topics are "computation" and "representation". The main tasks are algorithms and programs.

2) The period of the embodied intelligence. During this period, researchers began to realize that intelligence manifests itself in behavior, and that intelligence must have a body. According to the ideas of Rodney Brooks, it is called "embodied intelligence". (Brooks, R. A., and Breazeal, C.)

Twenty years of research in this field have generated an enormous number of stunning results and insights. Thanks to the "understanding by building" approach, the

theoretical underpinnings of AI have been made, including basic concepts of complete agents, adaptive behavior through neural networks, basic approaches, such as the Subsumption architecture, artificial evolution and artificial life, Dynamical systems; principles of intelligent systems such as design principles of autonomous agents, the principle of parallel, loosely coupled processes, principle of sensory-motor coordination, the principles of cheap design, redundancy, and ecological balance, the value principle, human memory.

The period of the embodied intelligence is not finished, but goes on continuously. We should understand more about the human brain. A Chinese researcher suggested, “it’s time to explore the secret of the brain of human, to get the intelligence form the brain, to build higher intelligence system to serve the human” (Jin Fan 2000).

3) The period of building “AIRE”. Exploring the cosmos is a difficult task; for astronauts it is not only dangerous but they also need to work for a long time in outer space, far away from Earth and their relatives. We should build a highly complex robot for the exploration of the cosmos. Similar to the “Fungus Eater”, this agent should be able to explore an ecological niche by itself. Before a formal name has been established, let me use the name “AIRE” here, its means “Artificial Intelligence Representation on Earth”. In fact, AIRE is an agent for a particular task. But it will work over large distance and long time. So it must be self-sufficient, situated, and use nonsupervised learning. The key is that the AIRE should be able to explore different ecological niches autonomously and, unlike the “Fungus Eater”, decide whether an environment is suitable for humans. I infer that analyzing images and other data, distinguishing and handling ecological niches will be the main topics in this period. Current achievements in fields such as image processing and sensor systems are not yet sufficient for an AIRE, so many advances have to be made before it is possible to build such agents.

As this is a goal for the entire human race, I suggest that the United Nations Educational, Scientific and Cultural Organization support the research necessary for developing AIRE.

4) The period of sending “AIRE” to the cosmos. It is widely known that on December 4, 1996, NASA launched the Mars Pathfinder spacecraft from Kennedy Space Center. It landed on Mars on July 4, 1997. Of course there is still a long way to go until science is advanced enough to send AIRE out of the Solar System to different star systems. To find and to establish new space for humans to live, I infer the main tasks will be how super speed spacecrafts and agents utilize the resources in the cosmos.

4 Conclusions

It is clear from the foregoing discussion what the ultimate goal of Artificial Intelligence should be and what the tasks for the four periods are. If the ultimate goal is realized, humanity’s brilliant culture and age-old history can be eternal in the cosmos -- thanks to AI research.

5 References

- Brooks, Rodney A., and Breazeal, Cynthia (2000). Embodied intelligence. Cambridge, MA: MIT Press.
- Fan, Jin (2000). Principle and Method of Computational Intelligence Basic on Neural, The Southwest JiaoTong University Press, Chengdu, China, P51.
- Pfeifer, Rolf, and Scheier, Christian (1999). “Understanding Intelligence”, The MIT Press Cambridge Massachusetts London, England, Contents, P38-P43, P83, P102, P607
- http://en.wikipedia.org/wiki/Main_sequence

“Cheap” underwater locomotion: Roles of morphological properties in underwater locomotion

Marc Ziegler, Fumiya Iida, Rolf Pfeifer
Artificial Intelligence Laboratory
University of Zurich, Switzerland
[mziegler | iida | pfeifer]@ifi.unizh.zh

Recently, there have been a number of studies on dynamic system-environment interactions in animals and robots to understand the nature of locomotion behaviours. In particular, the role of dynamic morphological properties (e.g. elasticity and rigidity of body structure, and weight distribution) have been regarded as a central issue to achieve real-time adaptability, energy and computational efficiency and rich behavioural diversity [1], [2]. The common ground for legged or swimming locomotion are the environmental conditions forming a frame of external boundaries for the system and the undulatority of motion [3]. The goal is to find intrinsic principles in locomotion which help to understand how nature achieved such a variety of different locomotion and improve future robots dealing with the real world. Our research is based on the thought of "cheap design" [4], [5], where the intrinsic material properties and adequate morphologies take over some of the computation for "free". What follows are the first experimental results of a fish-like swimming robot with only one degree of freedom for actuation and which is therefore very easy to control. Nevertheless, this robot exhibits surprisingly rich behavioural diversity in all three dimensions of the underwater environment.

In general, there are three important design principles required for underwater locomotion. First, a system has to maintain the stability of buoyancy. Second, propulsion force has to be considered. For fish-like swimming, it is of particular importance to consider how vortices can be created and exploited for locomotion. And third, locomotion direction needs to be controlled.

The experimental platform

The morphological design of the, here disassembled, robot is shown in figure 1 (a). It consists of the front (a stiff plate) and, of similar size and form, the flexible tailfin. One of the important morphological features with respect to the first design principle is the weight w and buoyancy b balance. Weight is located close to the vertical body axis and below the horizontal body axis, whereas floating parts are placed away the vertical and above the horizontal body axis. As a consequence two things happen in water. First, when body and tail fin are aligned, the robot remains in an upright position and also comes back after a disturbance. Second, when folded, because of the weight and float distribution, the robot also rolls to one side. This enables not only turning, but also going up. Similar to a corkscrew, it winds itself upwards, as it can be seen in the picture sequence in figure 1 (b).

In the first experimental runs, the motor is controlled by an open loop controller following a sinusoidal curve. The variable parameters are amplitude, frequency and offset ("steering angle"). Basic experiments in forward swimming compared different material properties (flexibility) of the tailfin with respect to different amplitude and frequency combination. Regarding the second principle, it can be shown, that for every material property a set of amplitude-frequency combination can be found, to maximize swimming speed. Comparing the different materials at their optimal control-combination, neither too soft, nor too stiff, but flexible (spring-like) material

shows overall best performance (figure 1 (c)). Experiments on turning/swimming up show how the turning angular velocity is more or less matching to the characteristics of forward velocity. This implies the control parameters being highly dependent on the material property of the tailfin for better performance of forward velocity. This conclusion matches to several observations in biological aquatic locomotion [6].




Fig. 1. (a) components of the fish robot, (b) sequence of upward movement, (c) forward velocity of the three different material properties

For better understanding how the material influences the swimming behaviour and to underline our first visual analysis, a second robot was build. This time, a bending sensor is added in its tailfin, measuring the deflexion during swimming. Comparing the sinusoidal curve of the motor control with the sinusoidal curve of the bending feedback, the peak of deformation does not necessarily match the peak of motor signal for different amplitude-frequency combinations. This observation and the actual deflexion potentially allow conclusions on actual behaviour (e.g. turning) and performance (e.g. velocity) and might be a first step towards the third principle of controlling the direction of locomotion.

Although we have explored only some parts on the possible morphological variations, e.g. elasticity of the tailfin and weight distribution, the experimental results provided significant insights toward a comprehensive understanding of underwater locomotion.

References

1. Dickinson, M. H., Farley, C. T., Full, R. J., Koehl, M. A. R., Kram, R., and Lehman, S.: *How animals move: An integrative view*, Science, (Vol 288, 100-106, 2000).
2. Kubow, T. M., and Full, R. J.: *The role of the mechanical system in control: a hypothesis of self-stabilization in hexapedal runners*, (Phil. Trans. R. Soc. Lond. B, 354, 849-861, 1991).
3. Blickhan, R.: *Axial aquatic and pedal terrestrial locomotion - Form, structure, and movement*, (Verh. Dtsch. Zool. Ges. 86.2, 5-11, 1993).
4. Pfeifer, R. & Scheier, C. *Understanding Intelligence*, (MIT Press, Boston, 1999).
5. Blickhan, R., Wagner, H., and Seyfarth, A. *Brain or muscles?*, (Rec. Res. Devel. Biomechanics, 1, 215-245, 2003).
6. Sfakiotakis, M., Lane, D. M., Davis, J. B. C.: *Review of Fish Swimming Modes for Aquatic Locomotion*, (IEEE JOURNAL OF OCEANIC ENGINEERING, Vol.24, NO 2, APRIL, 237-252, 1991).

Bioinspired Indoor Microflyers

Jean-Christophe Zufferey and Dario Floreano
Ecole Polytechnique Fédérale de Lausanne (EPFL)
Laboratory of Intelligent Systems (LIS), CH-1015 Lausanne, Switzerland
jean-christophe.zufferey@epfl.ch
<http://lis.epfl.ch>

There are not yet fully autonomous flying robots capable of maneuvering in small cluttered environments as insects do. The substantial weight and energy constraints typically encountered in this kind robotic application preclude the use of powerful processors and classical distance sensors (laser range finder, ultrasonic sensors, etc.). Moreover and due to their highly dynamic motion, flying systems require fast sensory-motor mapping despite the very limited processing power available onboard.

Since 2001, we explored bio-inspired approaches to build and control a range of indoor flying robots (Fig. 1). Taking inspiration from flying insects like flies is motivated by the fact that (i) they generally display efficient flight control capability in complex environments in spite of their limited weight and tiny brain, (ii) the sensory modalities they are using for flight control have artificial counterparts (sensors) that fits the limited available payload, and (iii) a large body of literature has been produced by biologists on their anatomy, sensors, processing pathways, and behaviors.




Fig. 1: a) Blimp, 150g, 2002 [1,4], b) F2, 30g, 2004 [2,4,5], c) MC1, 10g, 2006 [3]

The latest prototype we built is named MC1 and has an overall weight of 10 g including visual, inertial, and airflow sensors. It is capable of automatic take-off, speed regulation, and obstacle avoidance in a 7x6-m room equipped with randomly textured walls. To avoid collisions, it computes optic-flow from its onboard CMOS camera and fuses it with rotation rate information provided by a MEMS gyroscope. It has already demonstrated robust operation during several test flights, which lasted up to 10 minutes of autonomous operation.

References

- [1] Zufferey, J.C., Guanella, A., Beyeler, A., Floreano, D. (To appear) *Flying over the Reality Gap: From Simulated to Real Indoor Airships*. Autonomous Robots, Springer US.
- [2] Zufferey, J.C. and Floreano, D. (2006) *Fly-inspired Visual Steering of an Ultralight Indoor Aircraft*. IEEE Transactions on Robotics, 22(1) pp. 137-146
- [3] Beyeler, A., Mattiussi, C., Zufferey, J.C. and Floreano, D. (2006) *Vision-based Altitude and Pitch Estimation for Ultra-light Indoor Aircraft*. Proceedings of the IEEE International Conference on Robotics and Automation, Orlando (ICRA'2006).
- [4] Zufferey, J.C. (2005) *Bio-inspired Vision-based Flying Robots*. PhD Thesis n° 3194, Faculté Sciences et Techniques de l'Ingénieur, EPFL.
- [5] Zufferey, J.C. and Floreano, D. (2005) *Toward 30-gram Autonomous Indoor Aircraft: Vision-based Obstacle Avoidance and Altitude Control*. Proceedings of the IEEE International Conference on Robotics and Automation, Barcelona (ICRA'2005).

50th Anniversary Summit of Artificial Intelligence



Robot Demonstrations

Robot Demonstrations



The Artificial Mouse (AMouse) robot is a multisensory robot equipped with visual sensors and an artificial whisker system in order to perform robot experiments on navigation and learning based on multisensory cues. It serves as a means for validating models of multi-modal and sensorimotor integration gained from neurophysiological experiments and neural modelling. The project investigates the interdependence of sensory morphology (mechanical whisker properties, arrangement within whisker array) and the processing of the sensory signals. Active sensing (body motion and whisking behaviour) is studied to derive insights in the sensorimotor control system and underlying perception.

AMouse

Miriam Fend
Simon Bovet

Artificial Intelligence
Laboratory, University of
Zurich
[www.amousse.de]



In our research, we aim to better understand how legged animals and humans move. We try first to identify general principles of legged locomotion, and then to apply these principles to therapeutical or technical approaches, for instance, in rehabilitation, prosthetics, or robotics.

Bio Leg I and II

Jürgen Rummel
Andre Seyfarth

Locomotion Laboratory,
University of Jena
[www.lauflabor.de]



Here at Essex we are exploring this new field with the CRONOS series of robots, designed and constructed by Rob Knight. The inspirations for the first examples are the human skeleton and its musculature.

We are all familiar with the skeleton, and to some extent with muscles, but we are much less familiar with the ways in which the muscles are connected to the skeleton and act on it to produce movements and to exert force on the external world. CRONOS, a design study, illustrates our approach.

CRONOS

Owen Holland

Essex University
[cswwww.essex.ac.uk]




ECO-ROBOT is a humanoid robot kit which has 17 degrees of freedom and teaching function, and you can enjoy humanoid robot world in a day. You can assemble it within a day only using a plus driver tool and windows available PC even if you are beginner for robots.

ECO-ROBOT

Marc Ziegler
Rolf Pfeifer

Artificial Intelligence
Laboratory, University of
Zurich
iXs Research Corporation
[www.ixs.co.jp/index-e.html]



In biological systems, at all stages of development, the nervous system must be able to innervate and adapt functionally to any changes in the size and relative proportions of the body. Until now no engineering methods exist to tackle with unforeseen and concurrent changes in

morphology, task-environment, and neural structure. We expect to contribute with our approach to the solution of this problem.

Hand

Gabriel J. Gomez
Alejandro H. Arieta
Hiroshi Yokoi
Rolf Pfeifer

Artificial Intelligence
Laboratory, University of
Zurich

[www.ifi.unizh.ch/ailab]

Dept. Precision
Engineering, University
of Tokyo

[www.arai.pe.u-tokyo.ac.jp/dcm/index-j.htm]



With its new base and stylish design, the new Katana 1.2 catches the eye like no other robot. But it's not only the visual features that really make it unique: under its hood, the robot hides highly reliable state-of-the-art technology and a bunch of features that make it one of the most versatile industry-compliant robots available.

Katana Arm

Hansruedi Fröh

Neuronics AG
[www.neuronics.ch]



The project MediaFlies implements an interactive multi-agent system that incorporates flocking and synchronization in order to generate a constantly changing visual and acoustic output. It relies on prerecorded or live video and audio material, which is fragmented and recombined via the agents' activities. Video fragments stem from a

video frame ring buffer which is continuously updated. Audio material is also stored in a ring buffer and reorganized via granular synthesis. Agents engage in synchronization by adapting texture and grain properties in an attempt to recreate the original media material. The success of synchronization depends on the flocks coherence and velocity. Interaction is based on video tracking. It allows users to influence the flocks behavior by attracting or dispersing agents and thereby affects the balance between disturbance and recognizability of the systems audio and video feedback. The project draws its inspiration from the biological phenomena of flocking and synchronization. Simulations of these phenomena form the basis for the generative behavior of MediaFlies.

MediaFlies

Daniel Bisig

Artificial Intelligence
Laboratory, University of
Zurich

[www.ifi.unizh.ch/ailab]



There are two categories of robots; one is industrial robot and the other is service robot. When we design the industrial robot, we evaluate it in terms of objective measures such as speed, accuracy, and cost. The industrial robot is considered as danger for humans and isolated from them. On the

other hand, as the service robot interact with humans, we evaluate it in terms of subjective measures such as interesting, comfortable, and beautiful as well as the objective measures. There are two sub-categories of service robot; one is the physical service and the other is psychological service. The applications of robots for physical service are guidance, physical support for walking, and so on, and they have more weight on the objective measures than the subjective measures. On the other hand, those for psychological service are companion, communication partner, entertainment, mental therapy and so on, and they have more weight on the subjective measures.

Paro

Takanori Shibata

Advanced Industrial
Science and Technology
[www.aist.go.jp]
[paro.jp]



Evolutionary design is still in its infancy: coupled evolution is conducted only in simulation and how to best represent morphology, controller, environment, and fitness function is not clear yet; differences between virtual and real worlds have not been elucidated so that results from the virtual world are not

always transferable to the real world, especially in the case of dynamic systems, although work is focusing on this problem. Therefore, the work of Lipson, who demonstrated automated manufacture of evolved simulated robots, constrained his system to static locomotion. Thus, it is important to make clear the constraints for transferring dynamical behavior from simulation to reality.

PET Bot

Kojiro Matsushita
Hiroshi Yokoi

Dept. Precision
Engineering, University
of Tokyo
[www.koj-m.sakura.ne.jp]




This project is a part of the Biorobotics research activities at the AILab of the University of Zurich. The main objective of this project is to explore the design principles of biologically inspired legged running robots. In particular this project focuses on a minimalistic model of rapid

locomotion of quadruped robots inspired by biomechanics studies. The goal of this project is, therefore, to achieve technology for a form of rapid legged locomotion as well as to obtain our further understanding of locomotion mechanisms in biological systems.

Puppy I

Fumiya Iida
Gabriel J. Gomez
Rolf Pfeifer

Artificial Intelligence
Laboratory, University of
Zurich
[www.ifi.unizh.ch/ailab]



The group's research interests are at the intersection between robotics, computational neuroscience, nonlinear dynamical systems, and adaptive algorithms (optimization and learning algorithms). We develop and apply computational methods with solid mathematical foundations to a variety of

problems related to modeling, optimization, and control in computer science, biology, and robotics. We also take inspiration from biology to produce novel types of robots with adaptive locomotion and sensorimotor coordination abilities, and use the robots to investigate hypotheses of how central nervous systems implement these abilities in animals.

Puppy II

Jonas Buchli
Auke J. Ijspeert
Fumiya Iida

Biologically inspired
Robotics Group, EPFL
[birg.epfl.ch/page27899.html]
Artificial Intelligence
Laboratory, University of
Zurich
[www.ifi.unizh.ch/ailab]



Higher animals use some form of an internal model of themselves for planning complex actions and predicting their consequence, but it is not clear if and how these self-models are acquired or what form they take. Despite this, Brooks' AI revolution in the 1980s was a response to the slow progress of model-based robotics since the 1950s,

and ushered us into the current AI epoch in which model-free robotics is fashionable. Although simple and robust behaviors can be achieved without a model at all, Brooks himself anticipated the "tantalizing possibility" that a robot could one day autonomously create models of itself along with new behaviors, but cited "deep issues" preventing it from becoming viable.

Quad

Josh Bongard
Hod Lipson

Computational Synthesis
Laboratory, Cornell
University
[www.mae.cornell.edu/bongard]



Humanoid robots, robots that have an anthropomorphic body, are in the last years enjoying increasing popularity as a research tool, especially in Japan. This has mainly two reasons: first of all the usefulness of a human-like body for the deployment in an environment which has

been designed for humans and secondly the hope that a humanoid body can facilitate the understanding of human intelligence.

In the junior research group complex humanoid robots are being developed. This includes work in the areas: mechanics, electronics, perception, behavior control and learning. Our research focuses mainly on dynamic bipedal walking and intuitive multimodal communication with humans.

Robocup Humanoid

Sven Behnke

University of Freiburg
[www.informatik.uni-freiburg.de]



The common grounds for legged or swimming locomotion are the environmental conditions forming a frame of external boundaries for the system and the undulatory of the motion. The goal is to find the intrinsic principles in locomotion which

help to understand how nature achieved such a variety of different locomotion and improve future robots dealing with the real world.

Robot Fish

Marc Ziegler
Fumiya Iida
Rolf Pfeifer

Artificial Intelligence
Laboratory, University of
Zurich
[www.ifi.unizh.ch/ailab]



Roomba is an intelligent and effective vacuuming robot. All Roomba Vacuuming Robots feature iRobot's unique AWARE™ Robot Intelligence Systems. AWARE uses dozens of sensors to monitor Roomba's environment, and adjusts Roomba's behavior up to 67 times per second, ensuring that Roomba cleans effectively, intelligently and safely.

Roomba

Rodney Brooks

CSAIL, MIT
[people.csail.mit.edu/brooks/]
iRobot
[www.irobot.com]



What is the essential difference between physical interaction and informative ordering when an autonomous and distributed system forms the morphology? In this research, we have been trying to realize complex informative ordering, exploiting physical interaction. Recent advances in robotics reveal the importance of autonomous self-construction and embodiment for building intelligent systems. While currently most robot construction and repair is performed manually, this will be quite difficult when (a) the complexity of the systems exceeds a certain threshold, and (b) if these systems have to be truly adaptive. With conventional engineering hitting a complexity barrier it seems very useful to draw inspiration from natural systems, such as cells. Through natural evolution they have come up with many interesting solutions for some of the problems that future robotics will have to deal with, like self-organization and adaptivity to changing environments, fault tolerance and self-repair, self-programming and self-replication, to name but a few.

Scalable Self-Assembling Robots

Shuhei Miyashita

Artificial Intelligence
Laboratory, University of
Zurich
[www.ifi.unizh.ch/ailab]



Intelligent and complex agent behaviour can sometimes arise from very simple rules of interaction with the environment or other agents. Here, the physical interaction plays an important part. A famous example for such emergent behaviour are the vehicles from Braitenberg's thought experiments. The behaviours of these vehicles were even interpreted as expressions of love or fear. The other direction is much more complicated: How can an observed or experienced agent behaviour be understood in terms of rules, goals, or even intentions?

Sony AIBO

Verena Hafner

DAILabor, TU Berlin
[www.verena-hafner.de]



The Stumpy Project explores the fundamental design principles of locomotion on the basis of our biological knowledge. However, we do not simply copy the design of the biological systems, but we try to extract the underlying principles. One of the most fundamental challenges in this project is how to enhance the behavioral diversity of a robot by conserving the simplicity of the morphological and physiological design. Given this perspective, in this project, we are investigating the interplay between the oscillation based actuation, the material properties, and the interaction with the environment. Stumpy uses inversed pendulum dynamics to induce bipedal hopping gaits. Its mechanical structure consists of a rigid inverted T-shape mounted on four compliant feet. An upright "T" structure is connected to this by a rotary joint. The horizontal beam of the upright "T" is connected to the vertical beam by a second rotary joint. Using this two degrees of freedom mechanical structure, with a simple oscillatory control, the robot is able to perform many different behavior controls for the purpose of locomotion including the gait controls of hopping, walking and running.

Stumpy

Fumiya Iida
Gabriel J. Gomez
Rolf Pfeifer

Artificial Intelligence
Laboratory, University of
Zurich
[www.ifi.unizh.ch/ailab]



To understand the physics of water striders to model their characteristics of floating on the surface of water. We are using micro-actuators to simulate water striders' movements. We are also investigating different materials to improve the robot's ability to float on water.

Water Strider Robot

Metin Sitti

NanoRobotics Lab,
Carnegie Mellon
University
[www.cs.cmu.edu/~msitti/]



Humanoid robots are fascinating from two points of view, firstly their construction and secondly because they lend life to inanimate objects. The combination of biology and robots leads to smoother and compliant movement which is more pleasant for us as people. Biologically inspired robots embody non-rigid movement which are made possible by special joints or actuators which give way and can both actively and passively adapt stiffness in different situations.

ZAR5

Ivo Boblan
Rudolf Bannasch

Bionics Group, TU Berlin
[www.bionik.tu-berlin.de]
Evologics
[www.evologics.de]