Adaptive Systems

Inman Harvey
Informatics

Lecture 18:
Course Overview and Summary

Focussing on first part of the course, analytic framework, that came before the robotics examples in latter stages

Objectives

- To gain some familiarity with a number of different approaches to modelling and understanding adaptive processes in natural and artificial systems. In particular, to gain some understanding of approaches (old and recent) to generating adaptive behaviours in autonomous robots.

Rationale

- This course will cover theoretical aspects of biological adaptation and recent work in AI which is geared towards understanding intelligence in terms of the generation of adaptive behaviour in autonomous agents, acting in dynamic uncertain environments. Adaptation will be studied at both the evolutionary and the lifetime scale.

Rationale

- Lectures will give a general coverage. Seminars and exercises will guide you deeper into certain topics. You are expected to engage in background reading and follow up references mentioned in the lectures.
- No single textbook. But lots of books, book chapters, articles, online material, etc.

http://www.cogs.susx.ac.uk/users/inmanh/adsys10/

Contents: Cybernetics

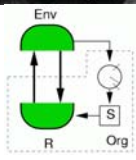

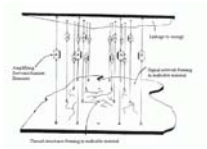

Cybernetics

- Cybernetics (from Greek for `steerman', the title of Wiener's 1948 book) is the study of the mathematics of machines, their control and communication.
- It aims at offering unifying principles for general classes of phenomena. For this reason it works at abstract levels introducing concepts such as control, feedback, signal, information, and so on, in models that could be applied to brains, ecologies, economies, societies, etc. Black box models.
- AI: a subset of the central aims of cybernetics.

What came out of cybernetics?

- AI/Robotics
- Control Theory
- Systems thinking
- Complex systems
- Information theory
- CAs, NNs, GAs.
- Self-organisation,
- Ultrastability,
- Requisite variety,
- Order from noise,
- Autopoiesis,
- Neurorobotics.

It influenced the work of many people (Piaget, Arbib, Pattee, Rosen, Maturana, Varela, J.Z. Young, Braitenberg, Rosenblatt, Cariani, J.J. Gibson, Edelman, etc.)

Ashbyan approach …

- W. Ross Ashby
- Organisms as state-determined systems
- Adaptation as stability
- Ultrastability and the 'homeostat'
- An example: Adaptation to visual inversion

Adaptation as stability

- *An organismic criterion*
- Definition: Behaviour is adaptive if it contributes to the maintenance of the essential variables within viable limits.
- Homeostasis is a low-level example of self-correcting mechanism.

Ultrastability

- Sensorimotor interaction
- **R** represents a subsystem of the organism responsible for overt behaviour/ perception.
- **S** represents those parameters affecting **R**. We assume that relevant features of behaviour do not change unless there is a change in **S**.

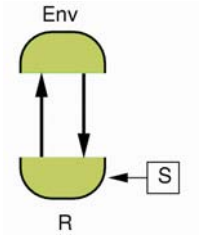

Ultrastability

- When essentials variables go out of bounds (system ceases to be stable) they induce changes in **S**.
- IF the whole system finds a new equilibrium, it will have adapted.

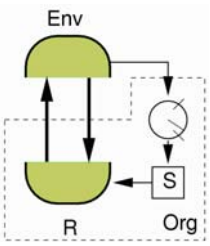

Ultrastability

- Double feedback:
  - Sensorimotor coupling.
  - Through essential variables acting on parameters.
- How do essential variables affect parameters? Depends on the system. Ashby proposed step-functions as a possibility.

Ultrastability

In the unstable case, state trajectories will reach a critical condition (right). If parameters were different (left) the system could still be stable under the new environmental pressure.

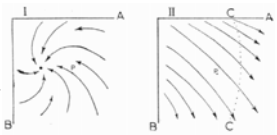

Steps functions acting through secondary feedback could take the dynamics from one field to the other.

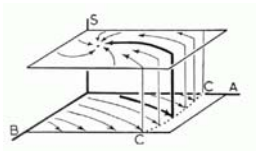

The Homeostat

- Electromagnetic device consisting of 4 ultrastable units that could be coupled in different ways
- Many experiments including habituation, reinforcement learning.

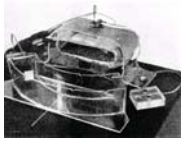

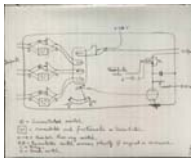

Autopoiesis, minimal living systems

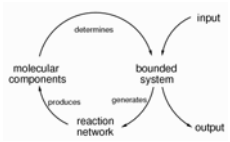

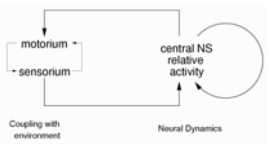

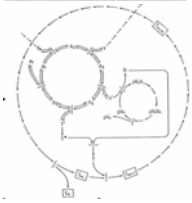

Αυτόσ = self, Ποιειν = to produce

- Autopoietic system: a system organised as *"a network of processes of production, transformation, and destruction of components that produces components which:*
  - *(i) through their interactions and transformations regenerate and realize the networks of processes (relations) that produced them; and*
  - *(ii) constitute it* [the system] *as a concrete unity in the space in which they* [the processes] *exist by specifying the topological domain of its realization as a network".*

- In essence
- A kind of organisational homeostasis. The *organisation* of living systems is such that this organisation is actively maintained. Its *structure* can (must) change in order maintain itself. (A virus is excluded.)

metabolism + boundary
↓
produces
↓
metabolism + boundary
↓
produces
↓

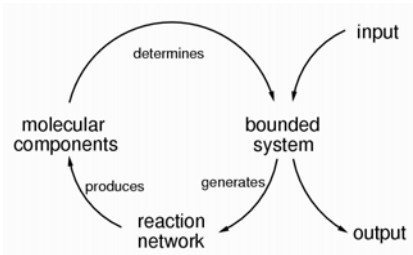

Evolution

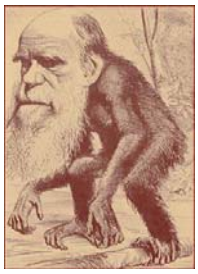

Major Evolutionary Transitions

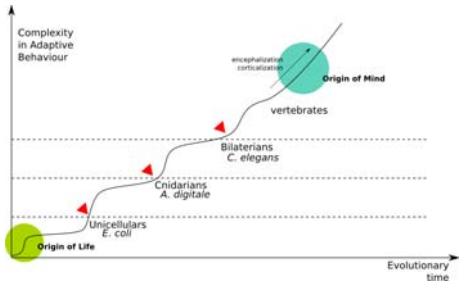

Stage 1: Protocellular Autonomy

- What is fascinating about living organization is that what the systems *does* (function) feeds back to what the system *is* (structure).
- Living systems are *precarious*, organized far-from-thermodynamic equilibrium, their parts operate so as to maintain other parts, none of which can stand in isolation.
- Function is not imposed from outside but *defined from within*. Self-maintenance is the mother of all values. What is good or bad is defined by the system.
- This is why living organization is considered autonomous. *Auto* = Self, *Nomos* = Norms.

Minimal adaptivity: membrane transport

Osmosis: passive transport

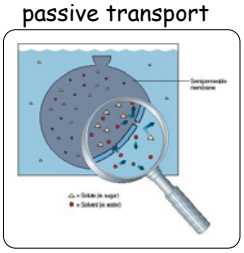

Ion-pumping: active transport

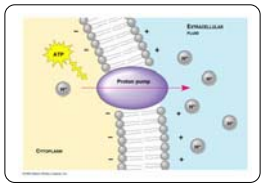

Chemotaxis in *E. coli*

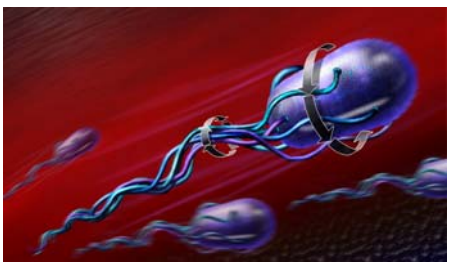

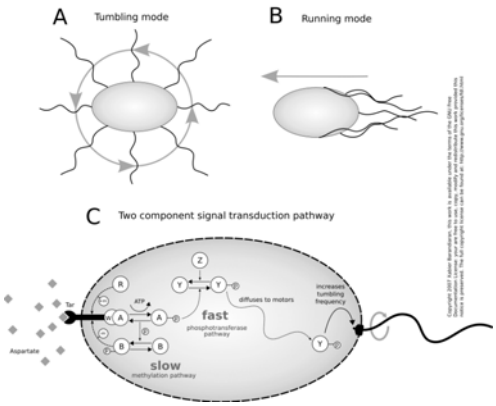

Origins of the Nervous System

- The origin of the NS:
  - Epithelial conduction is internalized
  - *Hydra*: neural nets coordinate unitary mouvement by generating internal rhythms

- The *decoupling* of the Nervous System
  - The dynamics of electric conduction are isolated from other metabolic processes
- *A. Digitale*
  - Differentiated circuits (rings)
  - Behavioural sequences

Adaptive Behaviour in *A. digitale*

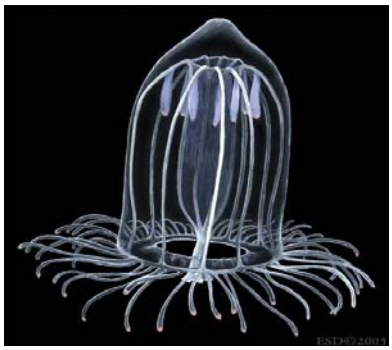

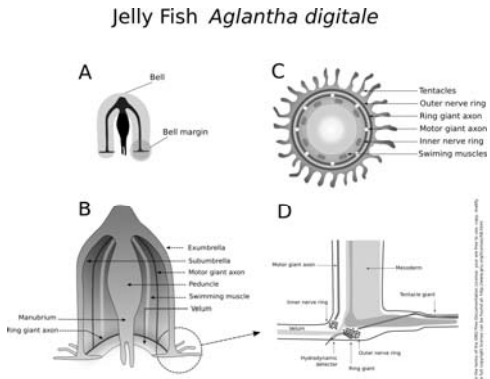

Properties of Neurons relevant for adaptive behaviour

- FACT: there is an evolutionary trend towards:
- 1. *Diversity* of multiple neurodynamic structures,
- 2. *Integration* of this diversity into a unified whole, and
- 3. *Plasticity* to modify and reorganize neurodynamic structures.

Levels of analysis

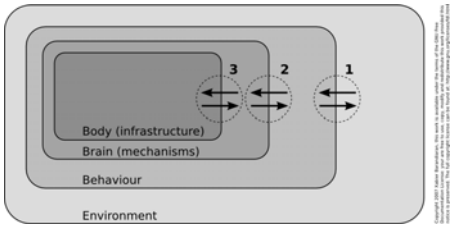

1. Organism-environment

- More diverse, integrated and plastic behaviour was enabled-required by
- Ecological environments in which high-energy resources were available to agents through increasingly sophisticated sensorimotor correlations
- Social co-evolutionary traps through hostility and cooperation

Godfrey-Smith

Types of Learning

- Supervised learning (we'll ignore this today)
  - "Teacher" provides required response to inputs. Assumes that desired behaviour is known.
  - E.g., backpropagation
- Unsupervised learning
  - Learner looks for patterns in the input. No "right" answer.
  - E.g., Hebbian learning, as well as ...
- Reinforcement learning
  - Learner not told which actions to take, but receives reward/feedback after each action.
  - E.g., Temporal difference learning

Learning 'trade-offs'

- Stability vs. plasticity
  - Need to be able to change
  - Don't want to destroy ***all*** previous learning
- Exploitation vs. exploration
  - Previously well rewarded options may be well rewarded in the future
  - It is always possible that a previously unexplored option may provide even higher reward

Hebbian learning

- Adjust strength of each connection in proportion to the product of their simultaneous activation.

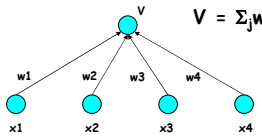

$V = \Sigma_j w(j)x(j)$

$\Delta w(i) = \alpha V x(i)$

Hebbian learning

- 'plain' Hebbian learning can allow V to grow to infinity.
- One solution is to normalize so that $|w|=1$.
- Another, more elegant solution is Oja's rule:

$$\Delta w(i) = \alpha V(x(i) - V w(i))$$

The additional term causes the weight vector to approach 1 asymptotically

Reinforcement learning (RL)

- In between supervised and unsupervised learning.
- Learning a mapping from **situations** to **actions** in order to maximize a (scalar) **reward/reinforcement** signal.
- Divide system into "situations", "actions", and "goals".
- Trial-and-error approach is required.
- Online book: Sutton & Barto (1998).

Delayed reward

- RL faces the problem of temporal delays between actions and rewards – the ***credit assignment*** problem
- One solution is to wait until the 'final reward' before updating the value function estimate.
  - But when is the 'final' reward?
  - ... and can't we be more efficient?
- Another solution is to have a complete and accurate internal model of agent-environment interaction
  - This is close to explicit planning
  - What if we don't want to assume the prior existence of such a model?

The second trick: temporal difference learning

- Temporal difference (TD) learning is a way to estimate the value function without having to wait until the final reward.
- At each time-step, we update the value function based on the error between successive predictions of the final reward value.
- Because it bases its update in part on an existing estimate, it is a form of 'bootstrapping'.

RL in context

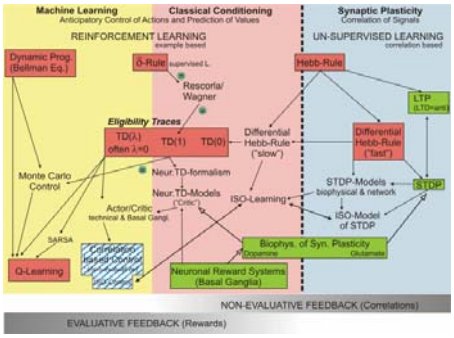

Worgotter & Porr (2005). *Neural Computation*

Sensory substitution

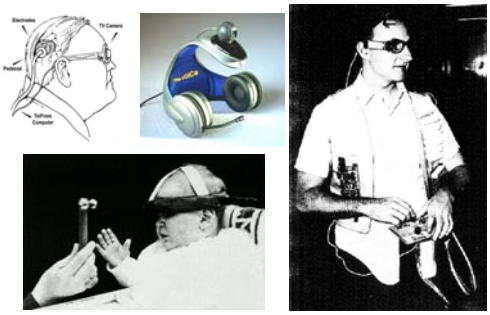

Distorted perception

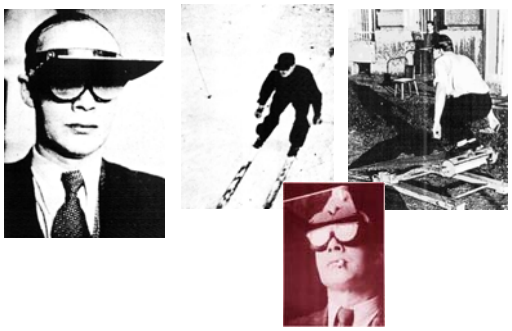

Adaptation is linked to *patterns of activity*, and may be confined only to those patterns (i.e., no adaptation by sitting and waiting).

Adaptation to grasping objects in an inverted visual field can happen by repeated trials at grasping. Adaptation to walking, by trying to walk. These work more or less independently.

Re-adaptation is needed after taking off the distorting device.

Distortions: up/down, left/right, prisms, bi-colour, etc.

Similar experiments with auditory perception.

Local vs global adaptation

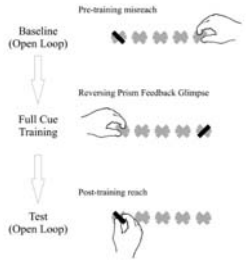

- Marotta et al., *J. Neurophysiol.*, (2005).

Circular processes and invariants

- An aside.
- As a rule of thumb, or general hypothesis, whenever you have complex, durable processes with circular or reciprocal causality, you will observe the formation of some *spontaneous invariant organisation*. (Otherwise the processes will not be durable and/or distinguishable).
- Examples: Nest building, pheromone trails (products constrain production). Stigmergy. (also called historical processes)

Sensorimotor coordination

- Can be viewed as a process of reciprocal causation. But it's difficult to be *aware* of it.
- Fluid behaviour becomes less and less conscious as it starts to rely on coherences of SM coordination.
- Surprise originates only when SM coordination is broken. (An unexpected obstacle, a contradiction between movement and proprioception, etc.)
- SM coordination can be organised into layers (individual movements, whole acts, chains of activity, etc.)

Adaptation: habit re-organisation

- Ashbyan-like theory: an invariant guiding change.
- Instead of essential variables for survival, essential variables that maintain the residual organisation (e.g., the culminating event of grasping). Keep changing until these variables return to their "normal" values (until you start grasping again).

Robotics

Applications of Adaptive Systems Part 2 of Course (not reviewed again here!)

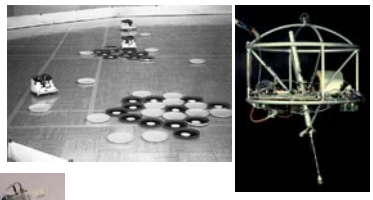

Evolutionary robotics

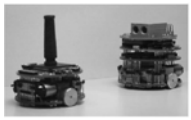

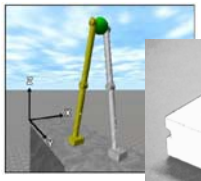

Minimal cognition, neural systems

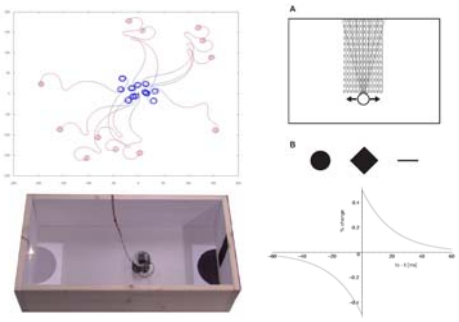

Neurorobotics

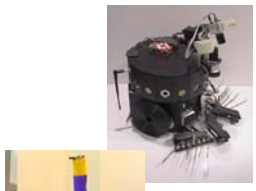

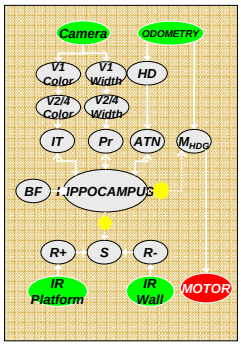

Assessment (Undergraduates)

- Exercise 1 (50%)
  - Programming exercise: Based on a robot or GA project. A 1500-word report to be handed in.
- Exercise 2 (50%)
  - A 2000-word essay on a relevant topic of your choice (list of topics will be made available).

NB: word-counts have been revised to these figures

Assessment (Postgraduates)

- Programming project (100%)
  - A 3500-word term paper (topic to be agreed) based on programming or robotic project and containing essay elements.
- Advice
  - You are encouraged to seek feedback on your choice of topic. Suitable topics and format advice will be made available.

NB: word-count has been revised to this figure

Your Feedback on this Course

- To be done via your Study Direct a/c
- Due by 28th March, online
- IF you are participating in a Student boycott of the NSS survey, because of current University actions, fair enough
- But in so far as this AdSys course feedback is different, please do it online, OR …
- If you have a principled reason for not doing it, please email the authorities (eg vc@sussex.ac.uk) to say so …
- But don't merely forget to do anything !!!!!

Questions