

Non-Symbolic AI lecture 5

EASU

We shall look at 2 alternative non-symbolic AI approaches to robotics

- Subsumption Architecture
- Evolutionary Robotics

Classical AI

EASU

When building robots, the Classical AI approach has the robot as a scientist-spectator, seeking information from outside.

"SMPA" -- so-called by Brooks (1999)

- S sense
- M model
- P plan
- A action

Brooks' alternative

EASU

Brooks' alternative is in terms of many individual and largely separate **behaviours** – where any one behaviour is generated by a pathway in the 'brain' or control system all the way from Sensors to Motors.

No Central Model, or Central Planning system.

Subsumption architecture (1)

EASU



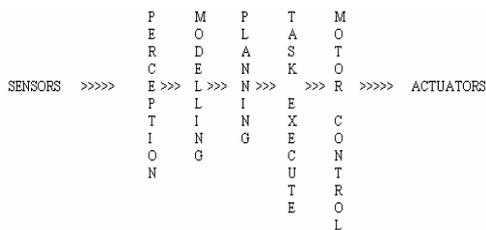
Fig. 1. Traditional decomposition of a mobile robot control system into functional modules.



Fig. 2. Decomposition of a mobile robot control system based on task-achieving behaviors.

(1a)

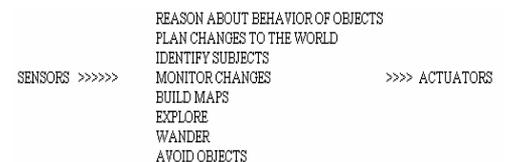
EASU



Traditional decomposition of a mobile robot control system into functional modules

(1b)

EASU



Decomposition of a mobile robot control system based on task-achieving behaviors

Subsumption architecture (2)

EASU

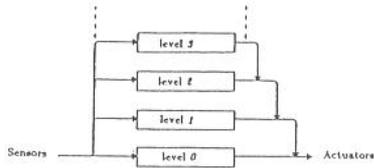
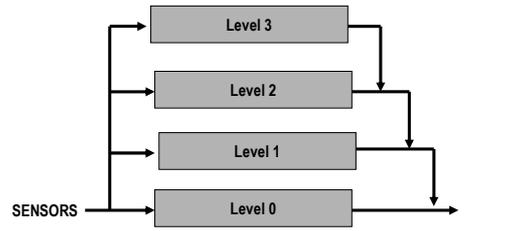


Fig. 3. Control is layered with higher level layers subsuming the roles of lower level layers when they wish to take control. The system can be partitioned at any level, and the layers below form a complete operational control system.

(2a)

EASU



Control is layered with higher levels subsuming control of lower layers when they wish to take control.

Subsuming

EASU

'Subsume' means to take over or replace the output from a 'lower layer'.

The 2 kinds of interactions between layers are

1. Subsuming
2. Inhibiting

Generally only 'higher' layers interfere with lower, and to a relatively small extent – this assists with an incremental design approach.

Subsumption architecture (3)

EASU

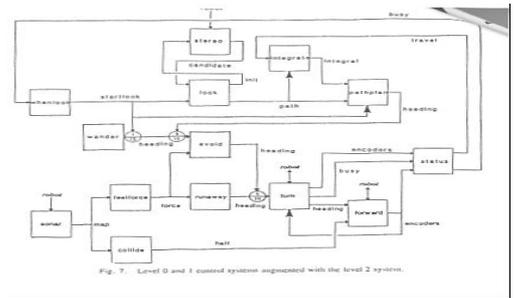


Fig. 7. Level 0 and 2 control system augmented with the level 2 system.

Subsumption architecture (4)

EASU

That looked a bit like a Network – except rather than (artificial) Neurons the components are versions of

- AFSMs
- Augmented
- Finite
- State
- Machines

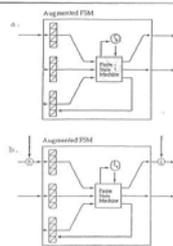


Figure 2.5. An augmented finite state machine (a) consists of registers, alarm clocks, a combinatorial network and a regular finite state machine. Input messages are delivered to registers, and messages can be generated on output wires. AFSMs are used together in networks using message passing wires. As new wires are added to a network (b), they can connect to existing registers, inhibit outputs, or suppress inputs.

AFSMs

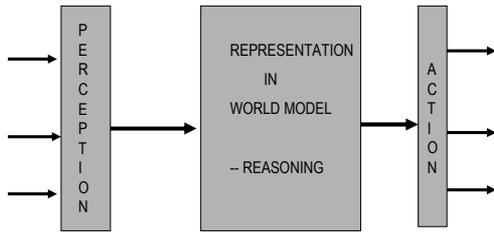
EASU

An AFSM consists of registers, alarm clocks (time!), a combinatorial network and a regular finite state machine. Input messages are delivered to registers, and messages can be generated on output wires.

As new wires are added to a network (lower figure before), they can connect to existing registers, inhibit outputs, or suppress inputs.

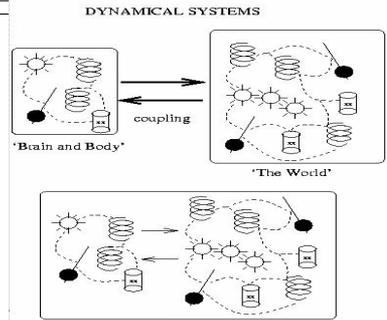
... could be a classical approach ?

EASU



... or a Dynamical Systems Approach

EASU



DS approach to Cognition

EASU

cf R Beer 'A Dynamical Systems Perspective on Autonomous Agents' Tech Report CES-92-11. Case Western Reserve Univ. Also papers by Tim van Gelder.

In contrast to Classical AI, computational approach, the DS approach is one of 'getting the dynamics of the robot nervous system right', so that (coupled to the robot body and environment) the behaviour is adaptive.

Brook's subsumption architecture, with AFSMs (Augmented Finite State Machines) is one way of doing this.

Dynamic Recurrent Neural Networks

EASU

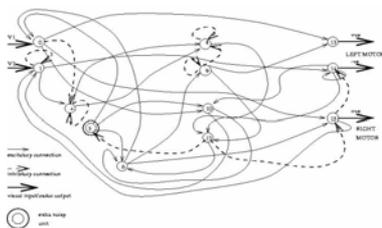
DRNNs (or CTRNs = Continuous Time Recurrent Networks) are another (really quite similar way).

You will learn about other flavours of Artificial Neural Networks (ANNs) in Adaptive Systems course.
-- eg ANNs that 'learn' and can be 'trained'.

These DRNNs are basically different -- indeed basically just a convenient way of specifying a class of dynamical systems -- so that different genotypes will specify different DSs, giving robots different behaviours.

One possible DRNN, wired up

EASU



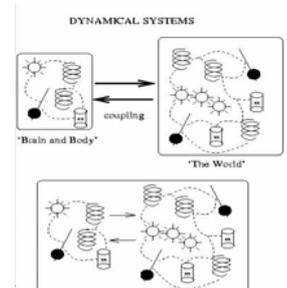
This is just ONE possible DRNN, which ONE specific genotype specified.

Think of it as ...

EASU

Think of this as a nervous system with its own Dynamics.

Even if it was not connected up to the environment (i.e. it was a 'brain-in-a-vat'), it would have its own dynamics, through internal noise and recurrent connections)



DRNN Basics

EASU

The basic components of a DRNN are these
(1 to 4 definite, 5 optional)

1.  Nodes
2.  Directed links
3.  Signals with time-delays
4.  Non-linear threshold function
5.  Noise

ER basics

EASU

The genotype of a robot specifies
(through the encoding genotype->phenotype that WE decide on as appropriate)
how to 'wire these components up' into a network connected to sensors and motors.

(Just as there are many flavours of feedforward ANNs, there are many possible versions of DRNNs – in a moment you will see just one.)

Then you hook all this up to a robot and evaluate it on a task.

Evaluating a robot

EASU

- When you evaluate each robot genotype, you
- ✓ Decode it into the network architecture and parameters
 - ✓ Possibly decode part into body/sensor/motor parameters
 - ✓ Create the specified robot
 - ✓ Put it into the test environment
 - ✓ Run it for n seconds, scoring it on the task.

Any evolutionary approach needs a selection process, whereby the different members of the population have different chances of producing offspring according to their **fitness**

Robot evaluation

EASU

(Beware - set conditions carefully!)

Eg: for a robot to move, avoiding obstacles – have a number of obstacles in the environment, and evaluate it on how far it moves forwards.

Have a number of trials from random starting positions

- ✓ take the average score, or
- ✓ take the worst of 4 trials, or
- ✓ (alternatives with different implications)

Deciding on appropriate fitness functions can be difficult.

DSs -> Behaviour

EASU

The genotype specifies a DS for the nervous system

Given the robot body, the environment, this constrains the behaviour

The robot is evaluated on the behaviour.

The phenotype is (perhaps):

- ✓ the architecture of the nervous system(/body)
- ✓ or ... the behaviour
- ✓ or even ... the fitness

Robustness and Noise

EASU

For robust behaviours, despite uncertain circumstances, noisy trials are needed.

Internal noise (deliberately put into the network) affects the dynamics (eg self-initiating feedback loops) and (it can be argued) makes 'evolution easier'

-- 'smooths the fitness landscape'.

Summarising DSs for Robot Brains

EASU

They have to have **temporal** dynamics.
Three (and there are more...) possibilities are:

- (1) Brook's subsumption architecture
- (2) DRNNs as covered in previous slides
- (3) Another option to mention here: Beer's networks

see Beer ref. cited earlier, or "Computational and Dynamical Languages for Autonomous Agents", in Mind as Motion, T van Gelder & R. Port (eds) MIT Press

Beer's Equations

EASU

Beer uses **CTRNNs** (continuous-time recurrent NNs), where for each node ($i = 1$ to n) in the network the following equation holds:

$$\tau_i \frac{dy_i}{dt} = -y_i + \sum_{j=1}^n w_{ji} \sigma(y_j - \theta_j) + I_i(t)$$

y_i = activation of node i

τ_i = time constant, w_{ji} = weight on connection from node j to node i

$\rho(x)$ = sigmoidal = $(1/1+e^{-x})$

η_i = bias,

I_i = possible sensory input.

Applying this for real

EASU

One issue to be faced is:

- Evaluate on a real robot, or
- Use a Simulation ?

On a real robot it is expensive, time-consuming -- and for evolution you need many many evaluations.

Problems of simulations

EASU

On a simulation it should be much faster
(though note -- may not be true for vision)
cheaper, can be left unattended.

BUT AI (and indeed Alife) has a history of toy, unvalidated simulations, that 'assume away' all the genuine problems that must be faced.

Eg: grid worlds "move one step North"

Magic sensors "perceive food"

Principled Simulations ?

EASU

How do you know whether you have included all that is necessary in a simulation?

-- only ultimate test, **validation**, is whether what works in simulation ALSO works on a real robot.

How can one best insure this, for Evolutionary Robotics ?

'Envelope of Noise' ?

EASU

Hypothesis: -- "if the simulation attempts to model the real world fairly accurately, but where in doubt extra noise (through variations driven by random numbers) is put in, then evolution-in-a-noisy-simulation will be more arduous than evolution-in-the-real-world"

ie put an envelope-of-noise, with sufficient margins, around crucial parameters whose real values you are unsure of.

"Evolve for **more robustness** than strictly necessary"

Problem: some systems evolved to rely on the existence of noise that wasn't actually present in real world!

Jakobi's Minimal Simulations

EASU

See, by Nick Jakobi:

- (1) Evolutionary Robotics and the Radical Envelope of Noise Hypothesis and
- (2) The Minimal Simulation Approach To Evolutionary Robotics

available on <http://www.cogs.susx.ac.uk/users/nickja/>

Minimal simulation approach developed explicitly for ER -- the problem is often more in simulating the environment than the robot.

Minimal Simulation principles

EASU

Work out the minimal set of environmental features needed for the job -- the **base set**.

Model these, with some principled envelope-of-noise, so that what uses these features in simulation will work in real world
-- **'base-set-robust'**

Model everything ELSE in the simulation with wild, unreliable noise -
- so that robots cannot evolve in simulation to use anything other than the base set
-- **'base-set-exclusive'**