Intelligence as The Dynamics of Interaction with the World

(following on from Lecture 1)

Passive Dynamic Walking

See Human Power & Robotics Lab, Theoretical & Applied Mechanics, Cornell University

- No motors
- No sensors
- No control system
- Just a slope and a body

Dynamic skills all the way up?

- We can readily imagine evolution crafting simple behaviours like this, but
- When we think about it, it seems like we are in rational, deliberative (linguistic?) control.
- Remember the trap of introspection (eg. Maier's "two strings" expt)?

Could the same principles of dynamical interaction with the world apply not just for simple reflex-like responses, but be universal?

Rodney Brooks
MIT AI Lab & iRobot Corp.

and his students, in the mid 1980's to early 1990's, developed "behaviour-based" robotics and suggested the "subsumption architecture" as one practical design method.

Provocative papers such as "intelligence without representation" and "intelligence without reason" deliberately sought to disturb the traditional AI establishment.

Central Ideas

- **<u>Situatedness</u>**: the robots do not deal with abstract descriptions, but with the "here" and "now“ of the environment that directly influences the behavior of the system.
  - *slogan*: “The world is it’s own best model”

- **<u>Embodiment</u>**: The robots have bodies and experience the world directly - their actions are part of a dynamic with the world and the actions have immediate feedback on the robots' own sensations.

Some implications: (1) The robot can only sense the world from its own perspective – no cameras stuck on the ceiling to give a fixed frame of reference. (2) The “symbol grounding problem” of traditional AI disappears.

Traditional decomposition into sequential functional (information processing) modules

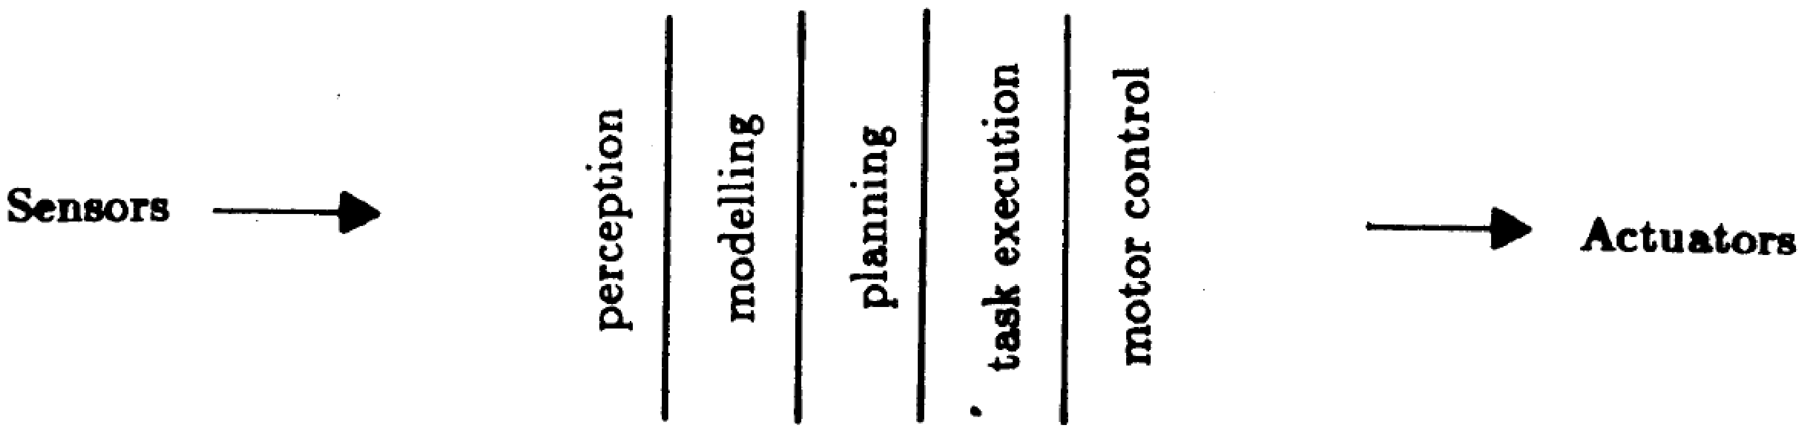

Figure 1. A traditional decomposition of a mobile robot control system into functional modules.

Brooks called this the SMPA approach: Sense, Model, Plan, then Act

Decomposition into parallel task-achieving competences ("behaviours")

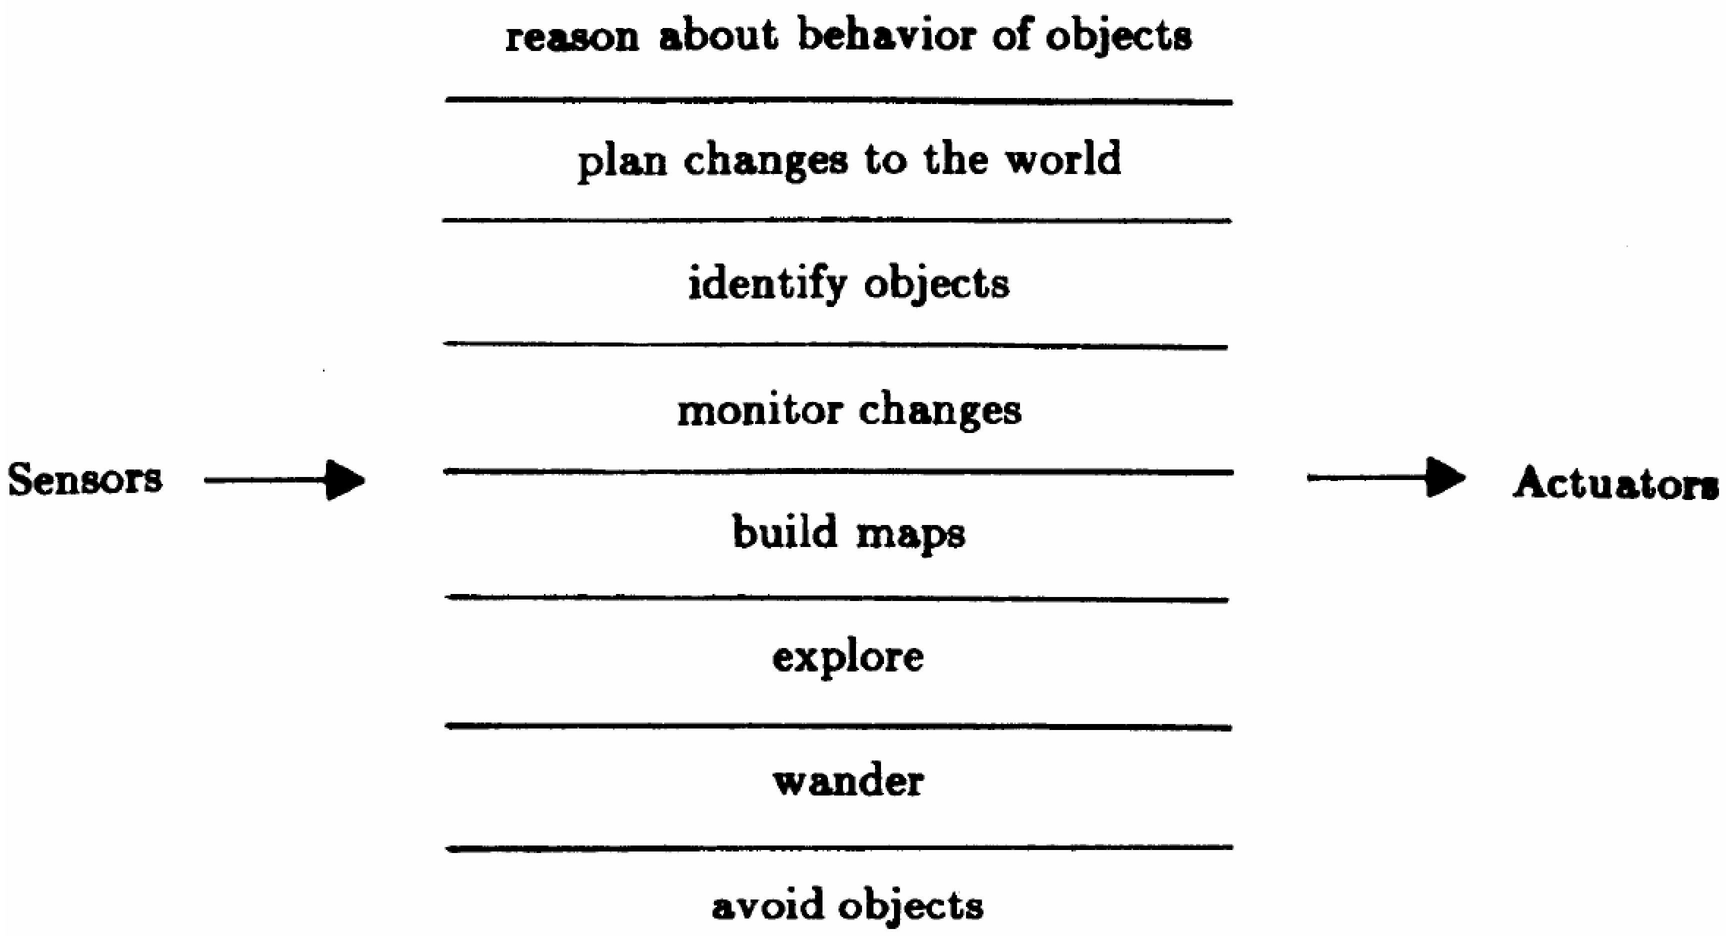

Subsumption Architecture

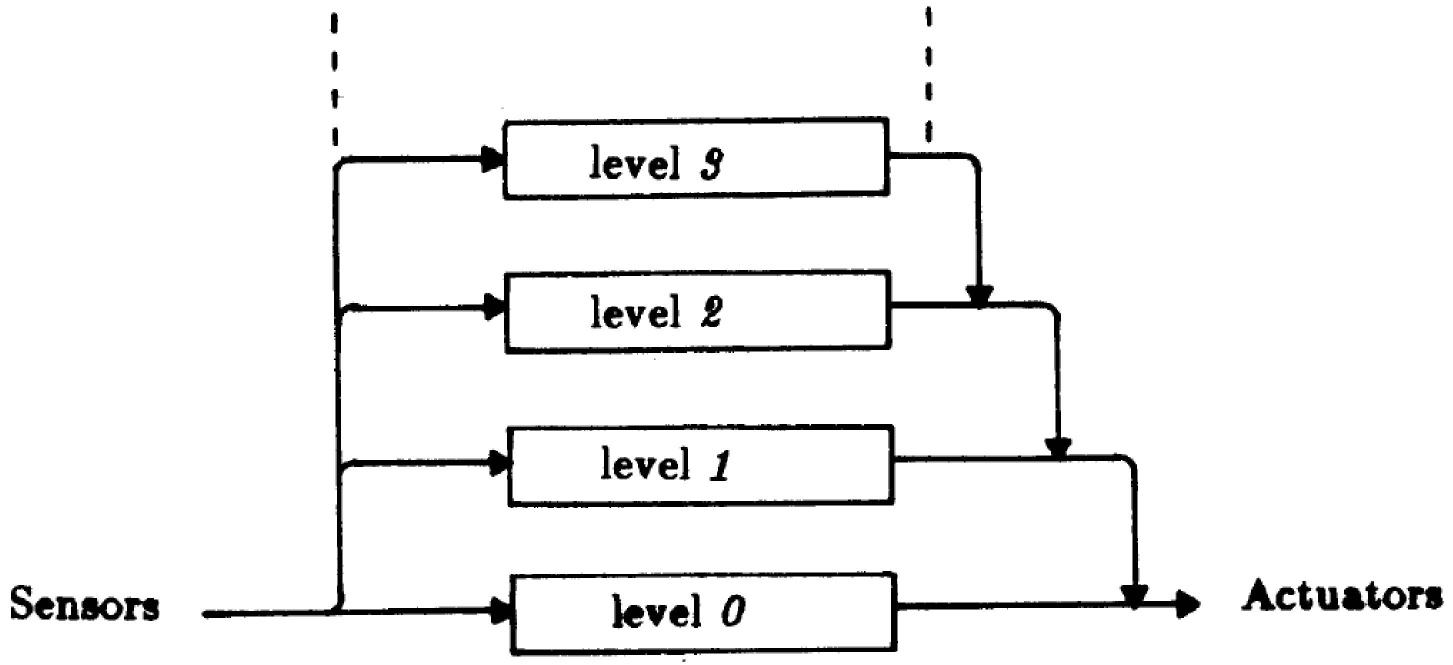

Figure 3. Control is layered with higher level layers subsuming the roles of lower level layers when they wish to take control. The system can be partioned at any level, and the layers below form a complete operational control system.

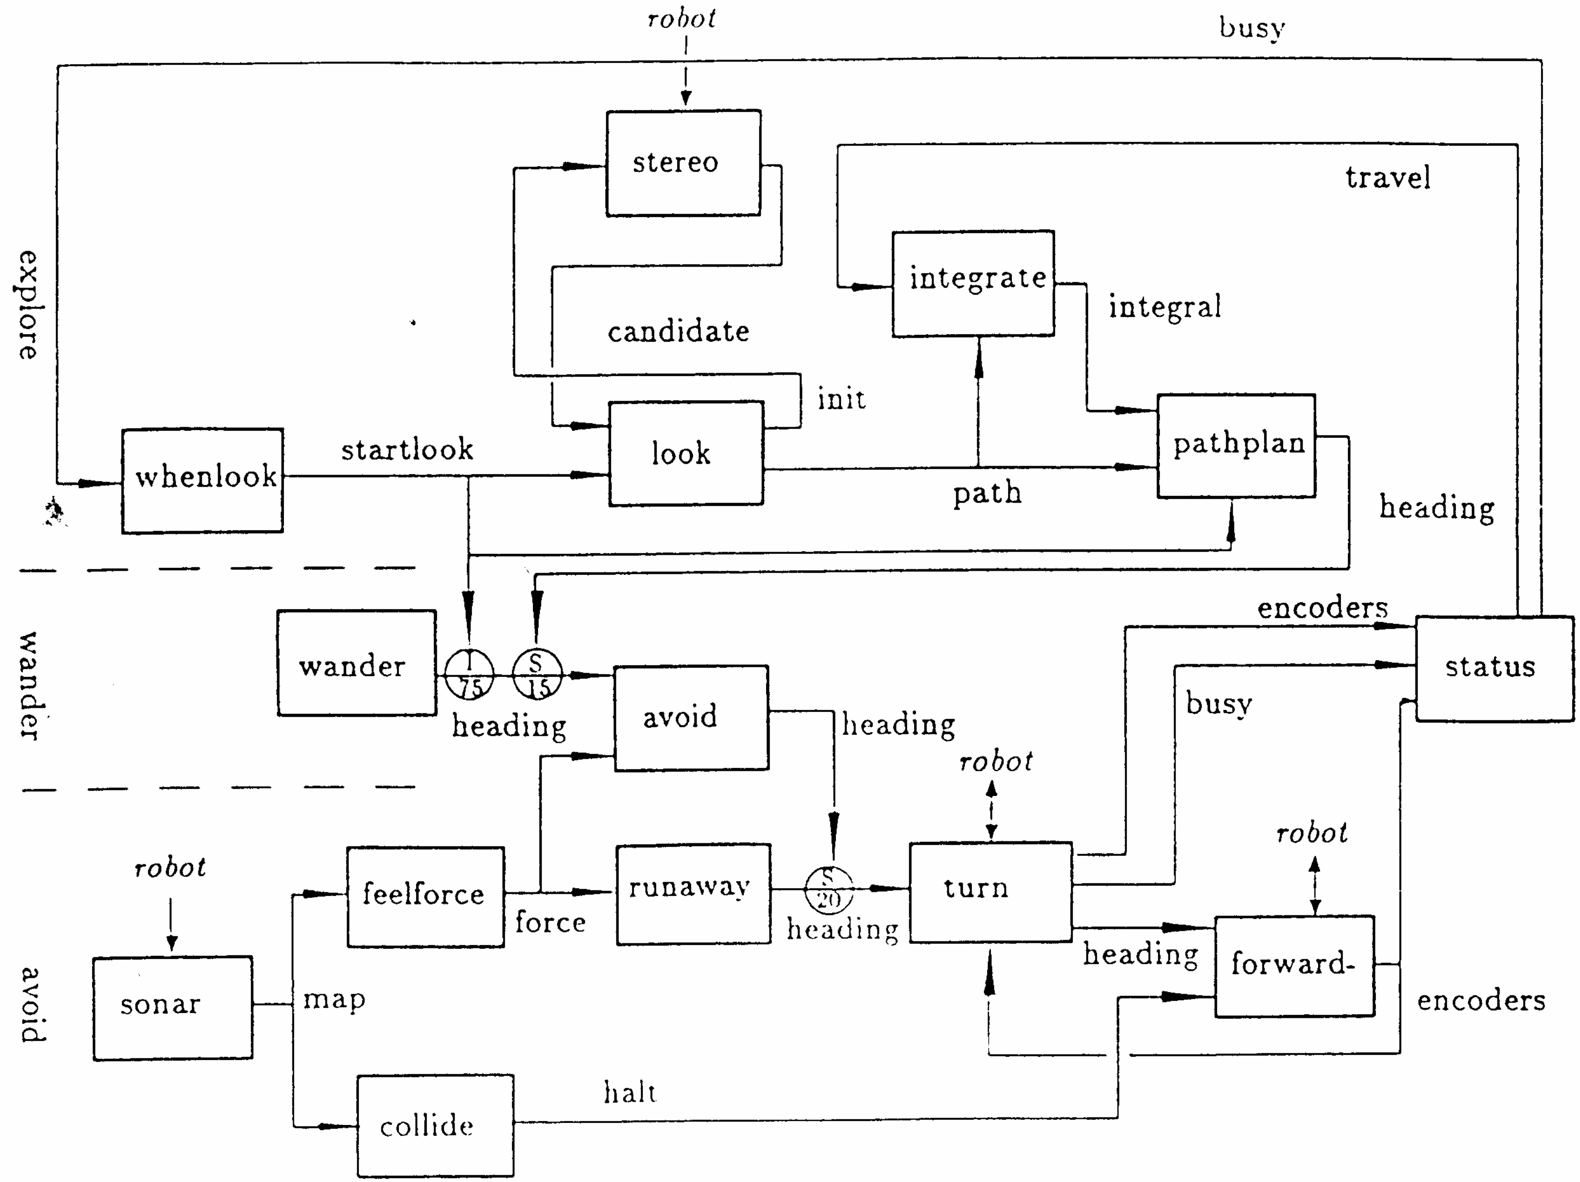

Figure 11.4
We wire finite state machines together into layers of control. Each layer is built on top of existing layers. Lower-level layers never rely on the existence of higher-level layers.

Principles

- There is no central locus of control: behaviours autonomously become active when conditions are appropriate. Each layer just does its own thing as best it can.
- There is no central storage. Individual layers may have “representations” and information processing steps, but these are specific to that behaviour.
- Keep it simple. Complex behaviour can *emerge* from interaction with a complex environment, and from the interactions between the modules and layers.

More Principles

- Lower layers are built first and fully debugged on the real robot in the real world, and are then frozen before a higher layer is added.
- Higher layers can rely on lower layers, but not vice-versa: at any time you could slice off some higher layers and still have a functioning robot. Lower layers keep functioning if a higher one fails to operate (eg. if the perceptual conditions do not match its preconditions for operating).
- Keep each added layer as a short connection between perception and actuation.
- Minimise interaction between layers.

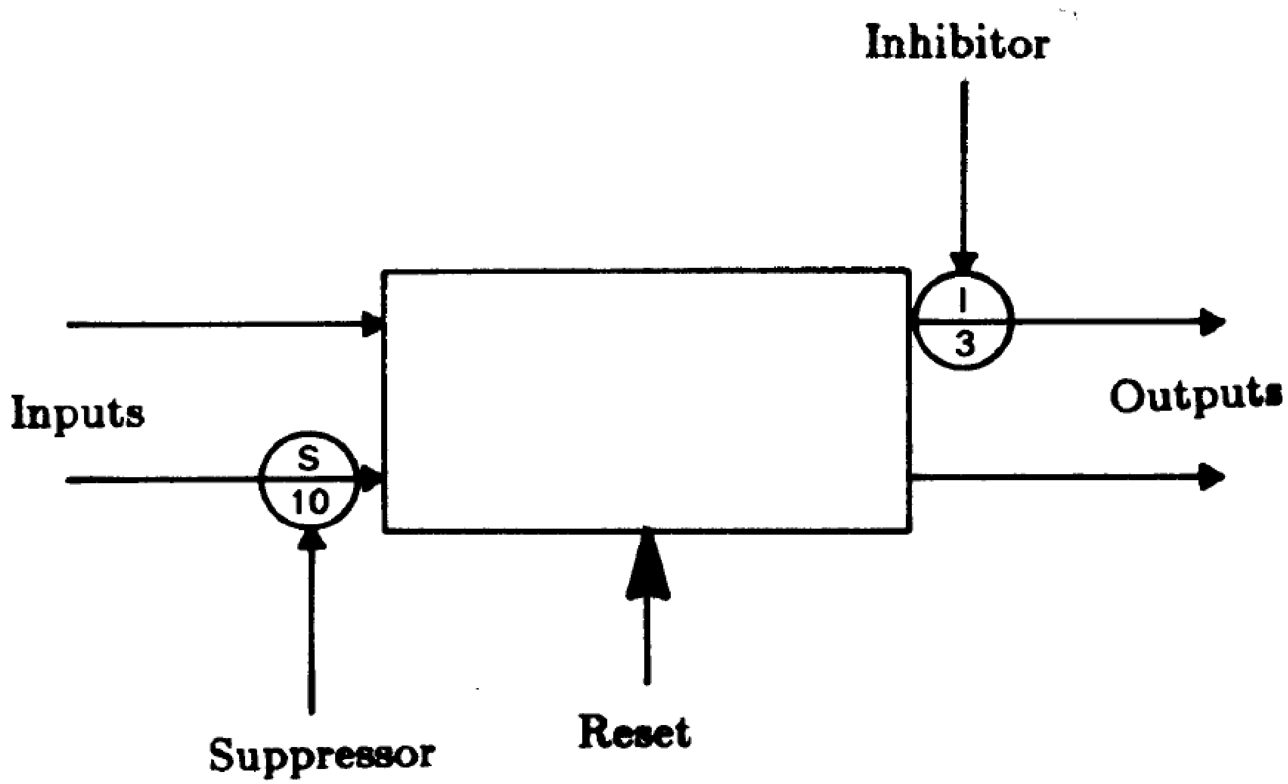

Figure 4. A module has input and output lines. Input signals can be suppressed and replaced with the suppressing signal. Output signals can be inhibited. A module can also be reset to state NIL.

The numbers in the circles are time constants. These modules run asynchronously, and send fixed length messages – eg 1 bit, or 24 bit, depending on the robot – down their “wires” which may be physical or virtual (if the modules are running in software on the same processor)

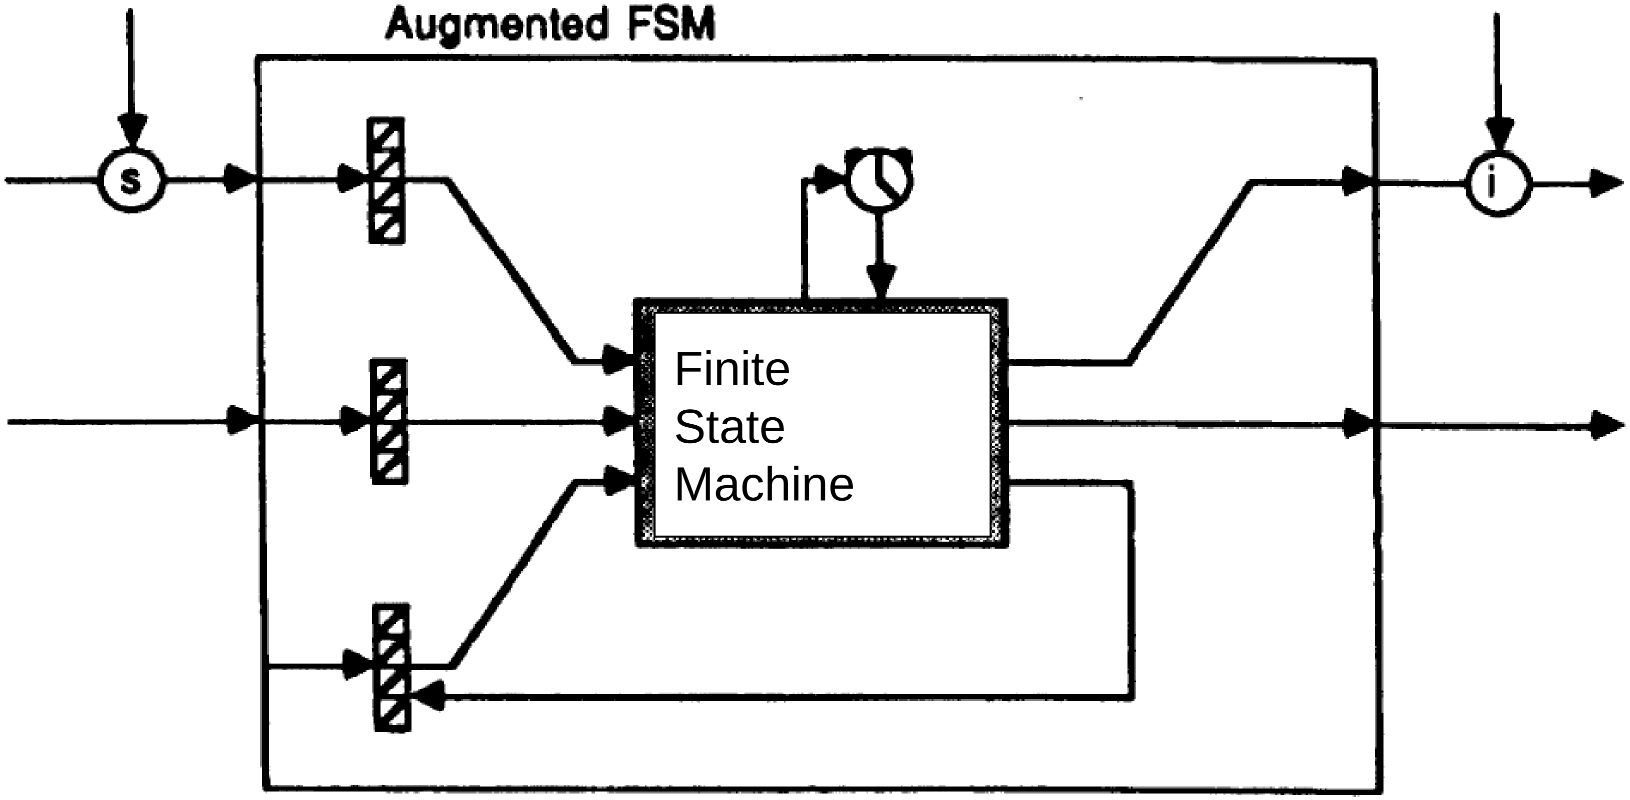

Figure 1. An augmented finite state machine consists of registers, alarm clocks, a combinatorial network and a regular finite state machine. Input messages are delivered to registers, and messages can be generated on output wires. AFSMs are wired together in networks of message passing wires. As new wires are added to a network, they can be connected to existing registers, they can inhibit outputs and they can suppress inputs.

(from "Elephants don't play chess" Robotics and Autonomous Systems 6 (1990) 3-15)

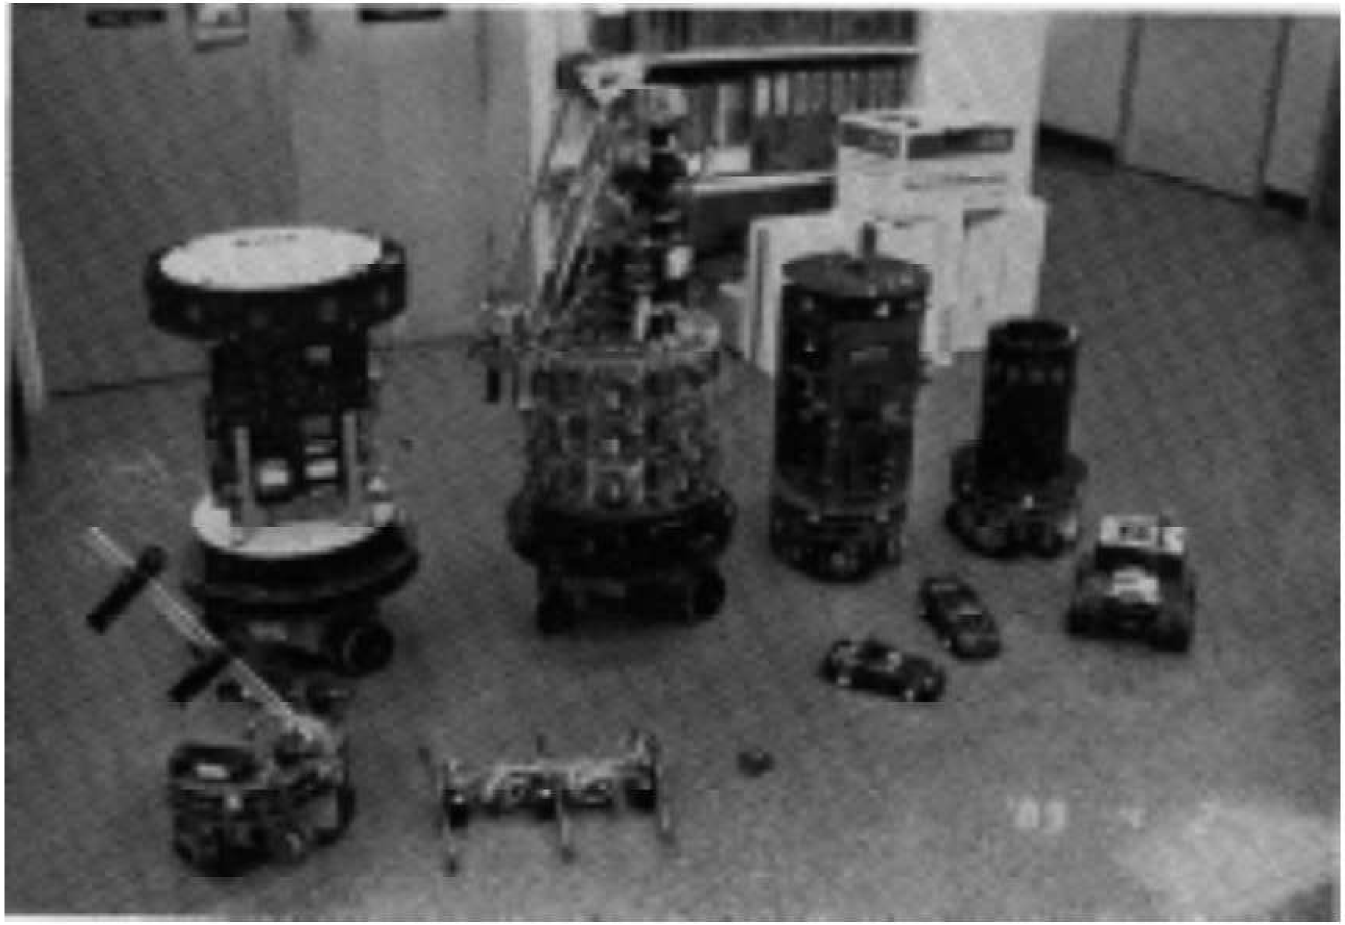

Fig. 1. The MIT Mobile Robots include, in the back row, left to right; Allen, Herbert, Seymour and Toto. In front row are Tito, Genghis, Squirt (very small) Tom and Jerry, and Labnav.

Genghis

- 1kg robot, highly distributed control system
- 12 motors, 12 force sensors, 6 pyroelectric sensors, one inclinometer and 2 whiskers.

- Successfully walks over rough terrain/obstacles
- Follows co-operative humans.

Ghengis Subsumption Layers

1. stand up
2. walk without any sensing
3. use force measurements to comply with rough terrain
4. use force measurements to lift its legs over obstacles
5. use inclinometer measurements to selectively inhibit rough terrain compliance when appropriate
6. use whiskers to lift feet over obstacles
7. use passive infrared sensors to detect people and to walk only when they are present
8. use the directionality of infrared radiation to modulate the backswing of particular leg sets so that the robot follows a moving source of radiation.

What about hybrids? Brooks says it may help today but eventually it's still doomed.

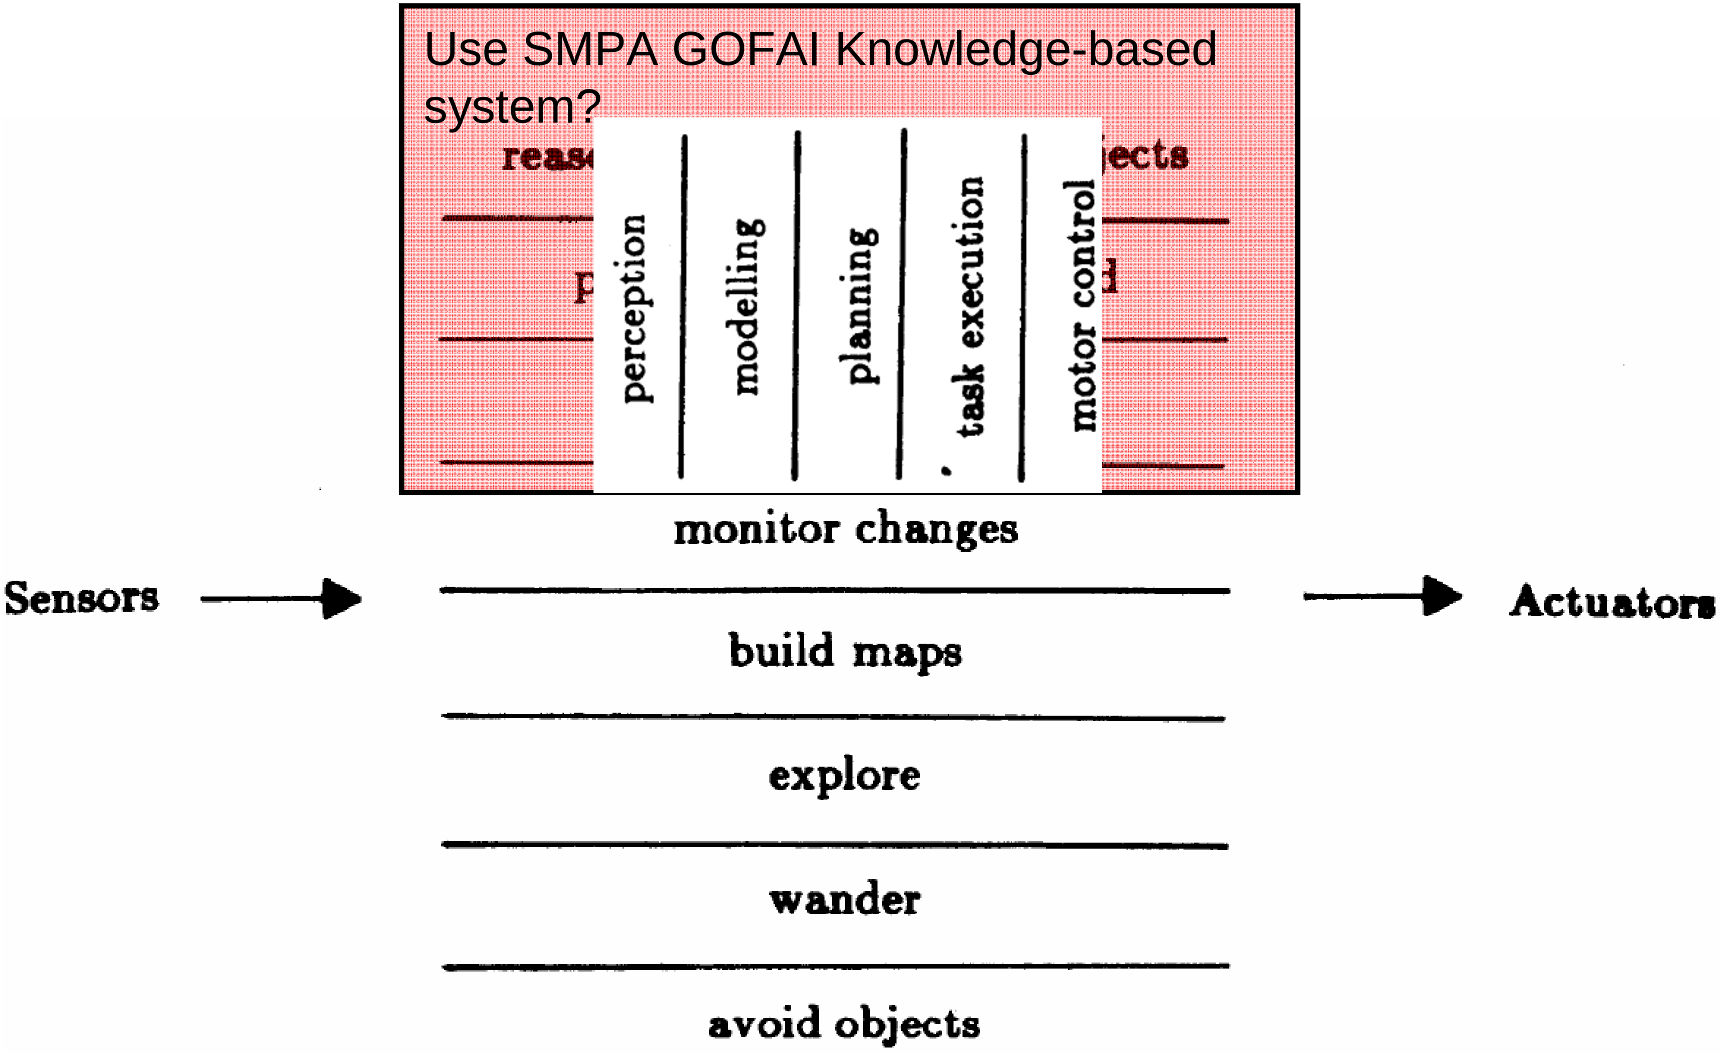

Criticisms/challenges 1

- It's still hard to scale up: have we just moved the goal-posts?
  - it does simple things in the real world rather than "clever" things in a simplified world

- How can we design all those behaviours to integrate properly?
  - Extended methods can include things like a network of spreading activation/inhibition between competing behaviours to achieve "action selection" (Pattie Maes)
  - Hormone models
  - Tempting to sneak-in a central locus of control/storage but stick a biological-sounding label on it!

Criticisms/challenges 2

- If each behaviour has its own private task-specific representations, how could we have “concepts”, learn from watching others, transferring knowledge from one situation to another, etc? (D. Kirsh “Today the earwig, tomorrow man”)
- “Behaviours” are categories in a description given by an external observer. They may not make sense from the agent’s perspective. This makes having a separate bit of brain for each “behaviour” questionable. (Hence Brooks tries to talk of “competences.”)

Comparison to Natural Evolution

- For reflex arcs, it’s sensor-signal > spine > motor signal: the brain only finds out what happened later!
- Vital basic capabilities are associated with ancient areas of the brain near the top of the spinal cord (eg. heartrate, blood pressure, breathing)
- Advanced capabilities are associated with areas added later in evolution
- It doesn’t carve-up as cleanly as a SA robot! Evolution can break some of Brooks’ principles. Let’s try it – next lecture, Evolutionary Robotics.

Patellar reflex

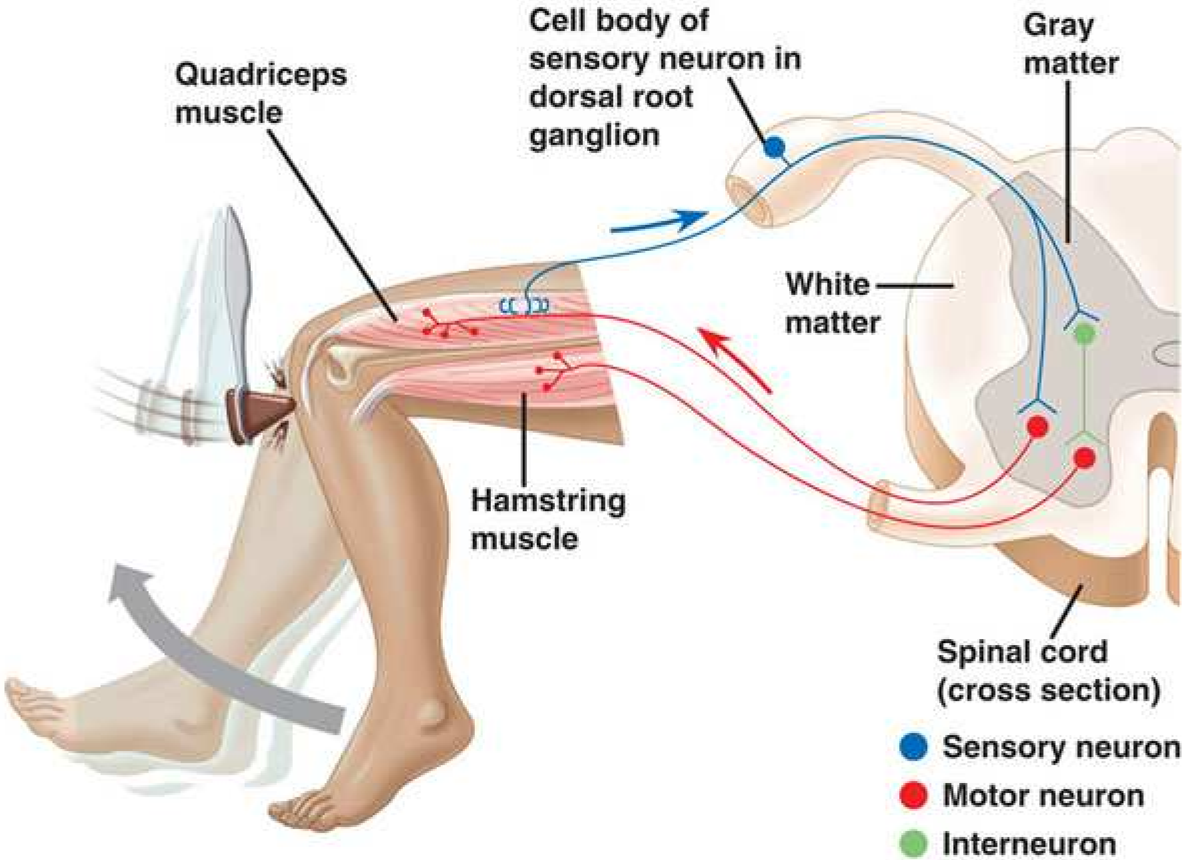

If we start to fall from upright this response is triggered, and it kicks us back, rapidly restabilising us.