Intelligent Autonomous Robotics

1. Introduction

1.1 What do we mean by ‘autonomy’ and ‘intelligence’?

Barbara Webb

Aim

- Machines that autonomously perform intelligent tasks in the real world


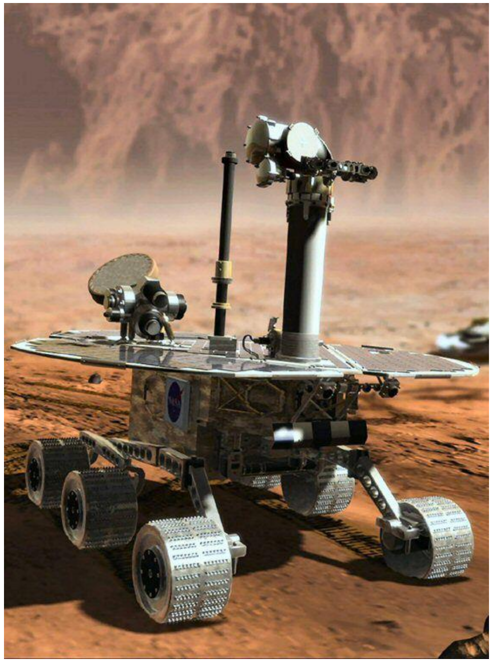


- But see the DARPA rescue challenge 2015

What does ‘autonomous’ mean for a robot?

- Take a few minutes to think for yourself what an *autonomous* robot should be capable of…

What does ‘autonomous’ mean for a robot?

No human in the control loop
(automatic – “self-moving”)

Not attached to anything for power or processing
(self-contained in operation)

Capable of maintaining behaviour against disturbance
(autopilot – “self-regulating” – cybernetic)

Generates own capabilities (self-organising)

Not dependent on human intervention to survive
(self-sufficient)

Generates own goals (self-governing - autonomous)

Generates own existence (autopoietic – “self-producing”)

Autonomy

Crucial aspects of autonomy for this course are:

- The system can achieve a task on its own
- The system is affected by and affects the real world around it *directly*, with no intervention (at least for the duration of its task)

As a consequence we have a closed loop:

- Output affects subsequent input (and task achievement) in ways governed by real world physics (e.g. time, forces, materials…)

What does ‘intelligent’ mean for a robot?

- Again, take a few minutes to think about what you would expect to see in an *intelligent* robot…

What does ‘intelligent’ mean for a robot?

- Can carry out a task that requires more than a pre-programmed sequence, e.g., with decision points depending on the real state of the world
- Adapts to dynamic environments
- Can plan (and re-plan) appropriate actions given high-level goals
- Learns to improve performance from experience

Intelligence

Crucial aspect of intelligence for this course is:

- System is adaptive to the situation

As a consequence:

- In contrast to traditional AI, much of the ‘intelligent’ competence we seek is common to humans and animals
- Intelligence is not just ‘in the head’

Intelligence

- Robotics addresses the crucial roles of *embodiment* and *situatedness* in intelligence
  - We frequently use interaction with the world to help solve ‘cognitive’ problems, e.g., sorting, writing, external memory.
  - Even our ‘off-line’ thinking is strongly body-based

A quick personality test

What is your immediate response to the following simple question?

Wednesday’s meeting has been moved forward two days. Which day will it happen?

Movement through space as a metaphor for time

Duffy, S. E., & Feist, M. I. (2014) Cognitive Linguistics, 25 (1), 29–54.

"Friday" = Ego-moving metaphor

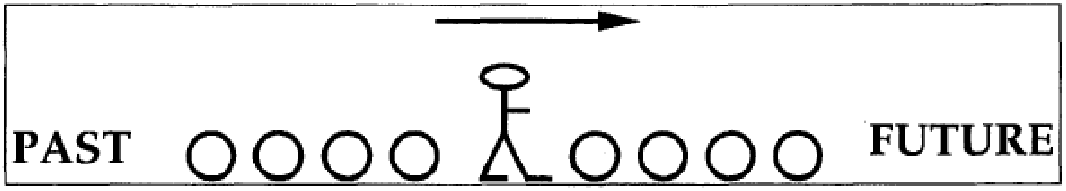

- Correlates with procrastination

"Monday" = Time-moving metaphor

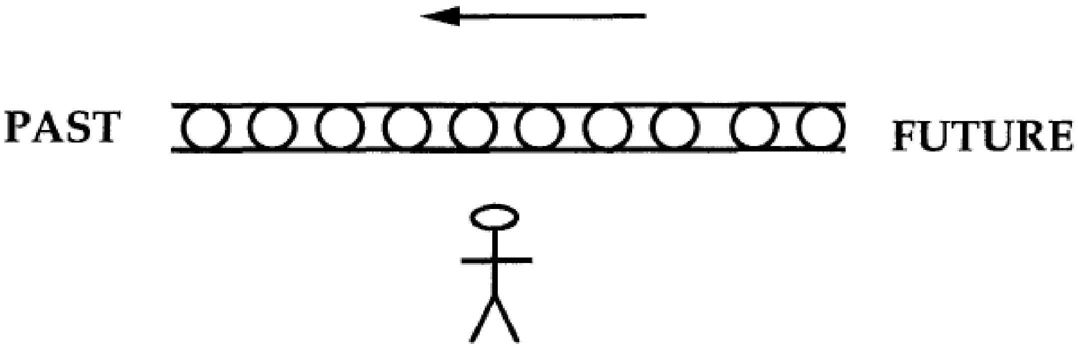

- Correlates with conscientiousness

Intelligence

- Robotics addresses the crucial roles of *embodiment* and *situatedness* in intelligence
  - We frequently use interaction with the world to help solve ‘cognitive’ problems, e.g., sorting, writing, external memory.
  - Even our ‘off-line’ thinking is strongly body-based, e.g., metaphors of time as space.
  - Many believe we will not be able to build a real AI system unless it in some way shares our physical experience

Practicals

- ‘Embodied cognition’ means you learn more easily by physical interaction with a real system: so need hands on experience with ideas that will be covered in lectures!
- Task is to programme Khepera robots to collect ‘food’ and take it home, details here: www.inf.ed.ac.uk/teaching/courses/iar/practicals.html
- Worth 50% of final mark, formative feedback provided along the way.
- Practical time: 4-6 Tues FH
- Practical partners

1.2 A biological example of autonomous intelligence

The desert ant

The ant environment

1st excursion

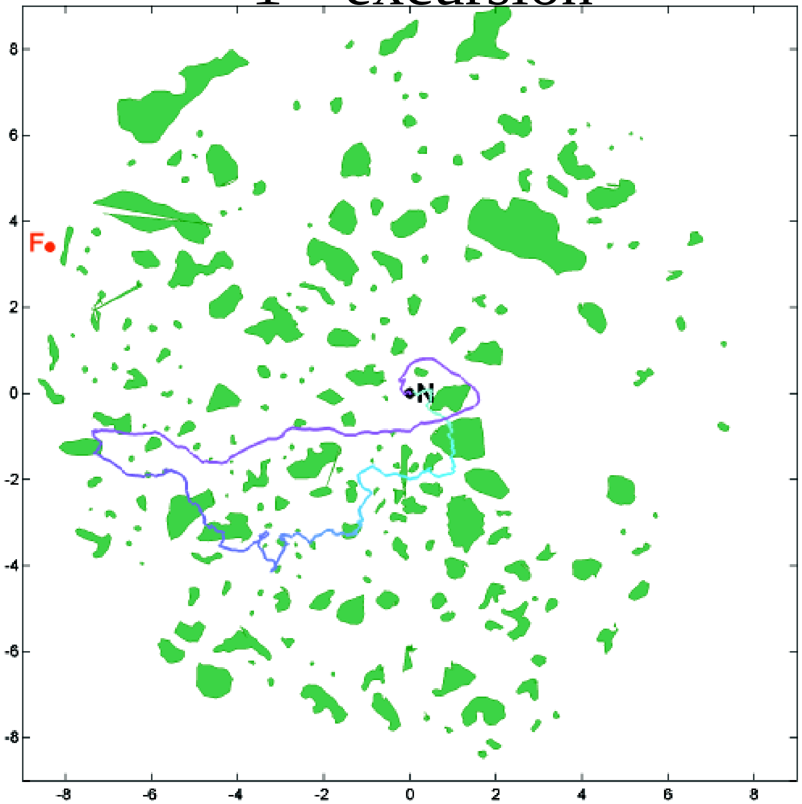

2nd excursion

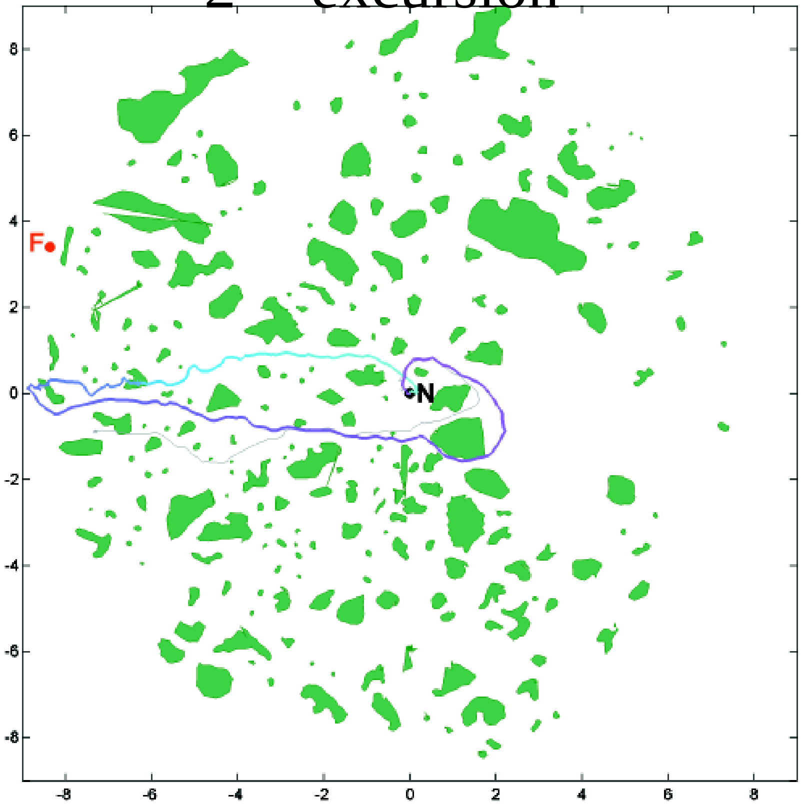

3rd excursion

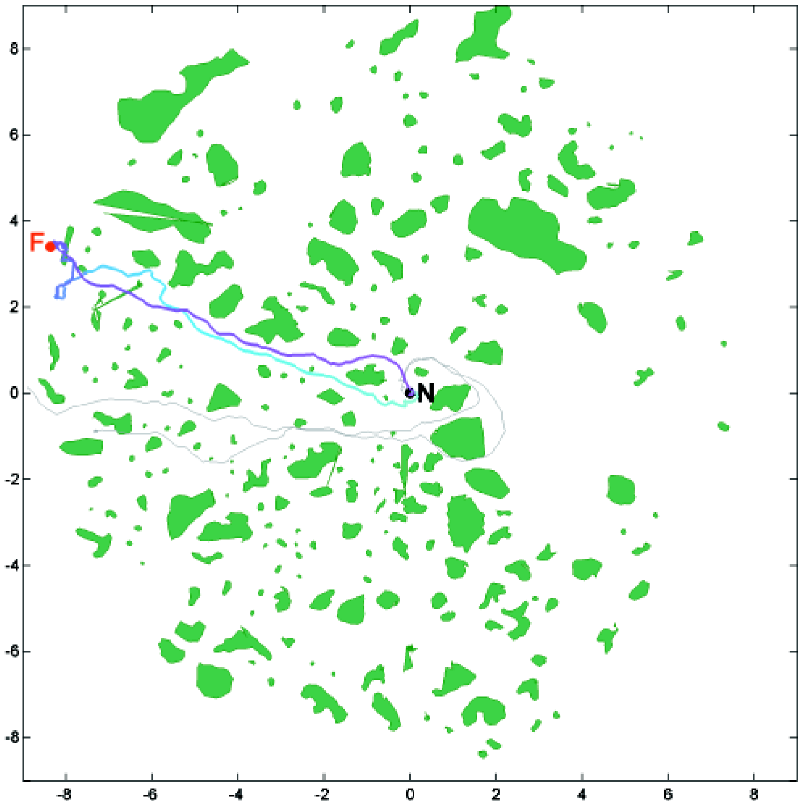

4th excursion

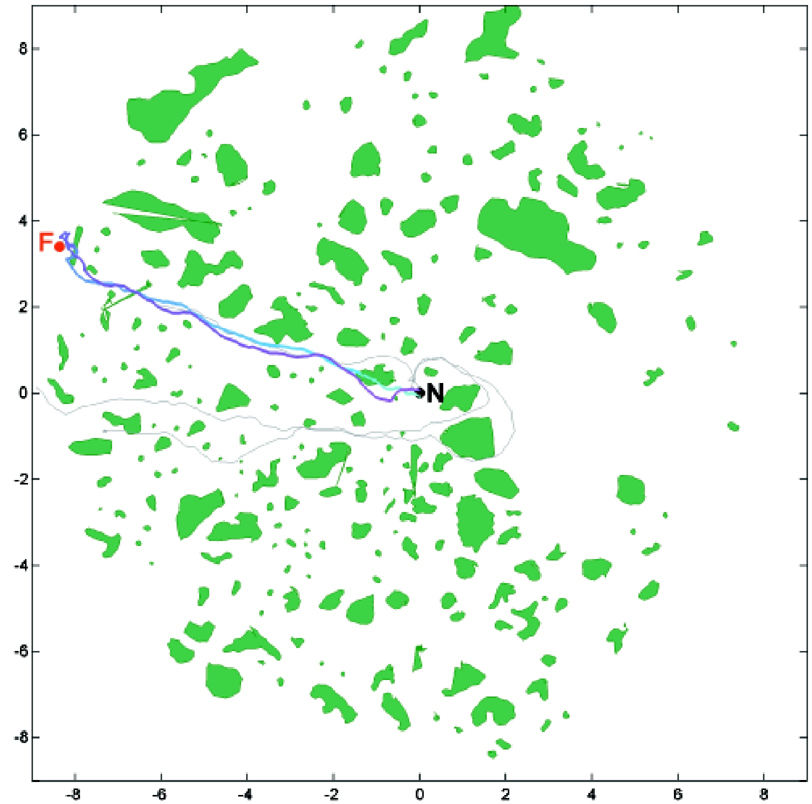

Zero vector test

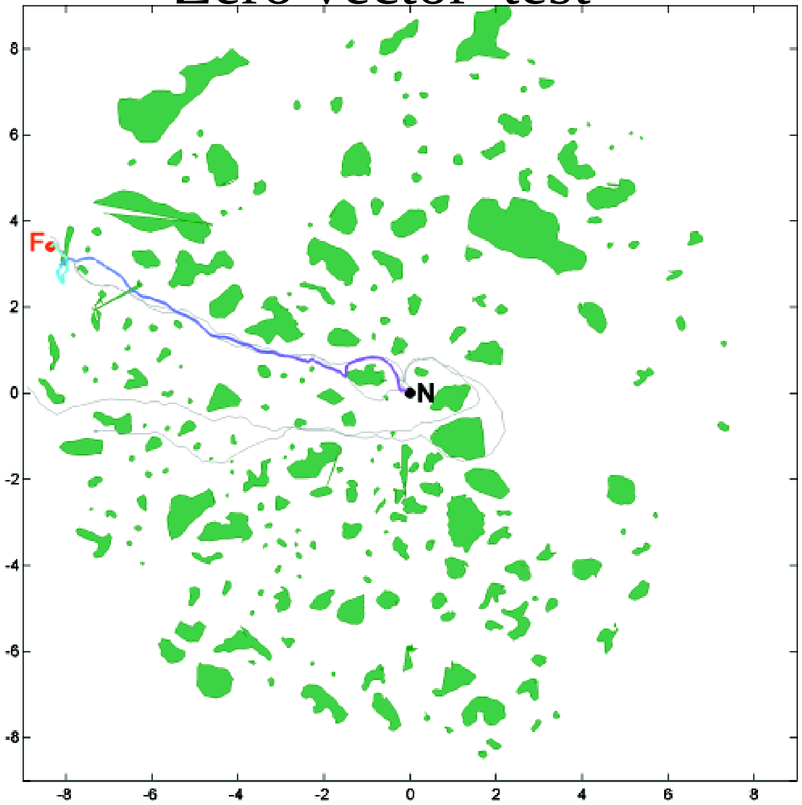

Zero vector to new location

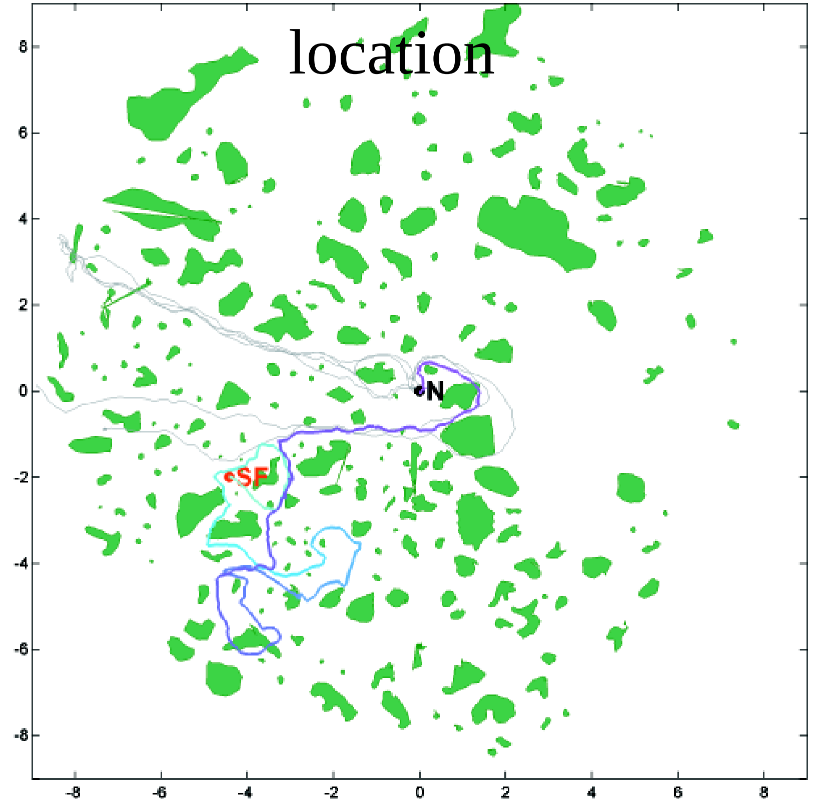

Through the eyes of an ant

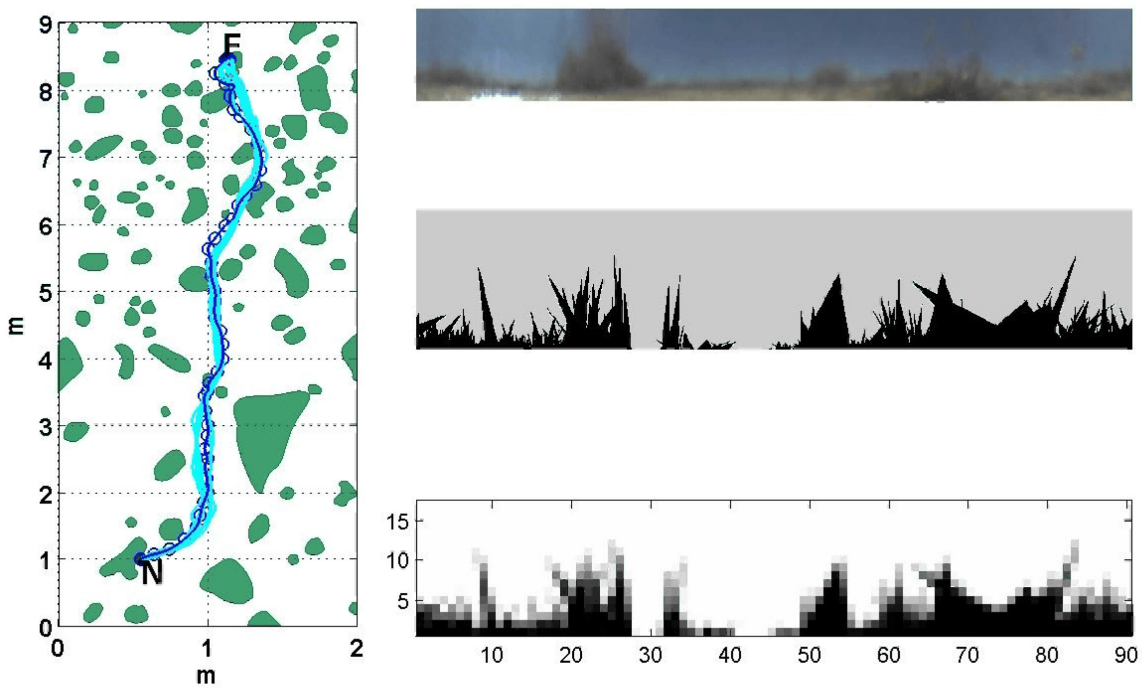

Before the next section:

Sketch out what capabilities you think a robot would need to be able to navigate like an ant. For example:

- What sensory systems would it need?
- What information would it need to extract from the environment?
- What information would it need to store?
- How would it use the stored and current information to decide what to do?

(N.B. we will come back in later lectures to how the ant actually does it…)

1.3 The planning/control problem

- What should our robot do next?
  - N.B. could refer to short or long time horizon
- How can we bring about a desired state of the robot and/or world?
  - Complete a task, probably against disturbances.
- What control policy will satisfy the robot’s goals within the robot and world constraints?

The planning/control problem

Some typical examples:

- Get robot from A to B, within certain time
- Complete a mission within power constraints
- Map an area to a given level of accuracy
- Decide between alternative routes, e.g., uncertain shortcut vs. well-known path
- Stay on the road and don’t collide with anything

Consider problem of steering a car on a racetrack.
Might have:

- Input: distance from edge, y
- (Internal) state: heading, x
- Output: steering angle, u
- Disturbances: undulating track

Want to determine a *policy*: $u = \pi(x, y)$

Multiple possible approaches e.g.:

- Open loop: pre-programmed sequence of actions
- Feedback: Turn wheel based on distance from edge
- Feedforward: Make corrections based on upcoming turn

The planning/control problem

Planning and control essentially refer to the same thing, i.e., deciding what the robot will do, but at different levels:

Deliberative reasoning	Typically offline	Far time horizon	Sensing only if plan fails
Trajectory planning	Look ahead but modify	Near future	Sensing used to monitor
Low level control	Online	Immediate	Sensing used directly

May use offline planning to construct an executable controller.

Why planning/control is difficult

- Intrinsic uncertainty is inherent to robotics
- A robot's knowledge of the problem is limited to what it has been told and what its sensors can tell it
  - Typically high level prior info
  - Typically limited sensor range
- The actual effect of a robot's actions is usually uncertain
  - And the world might change

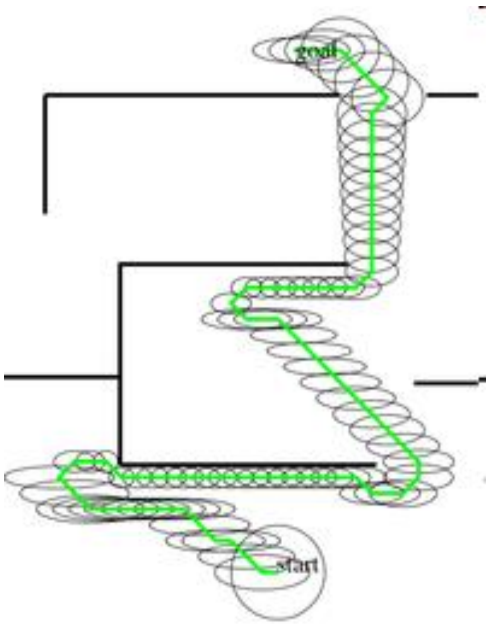

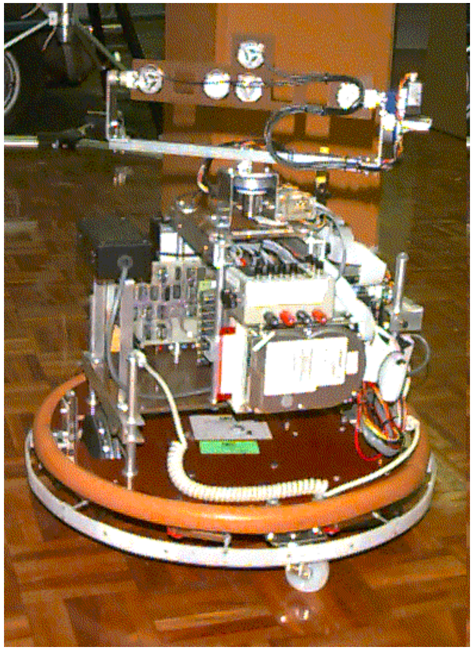

Historically, there have been different approaches to dealing with this inherent uncertainty

Model-based Principled but brittle	Assume everything is known, or engineer robot or situation so this is approximately true	sense→plan→act
Reactive Robust and cheap but unprincipled	Assume nothing is known, use immediate input for control in multiple tight feedback loops	sense→act sense→act
Hybrid Best and worst of both ?	Plan for ideal world, react to deal with run-time error	plan ⇩ sense→act
Probabilistic Principled, robust but computationally expensive	Explicitly model what is not known	sense→ plan → act with uncertainty

Before the next lecture:

- Think of a specific robot you have heard about recently (or have browse on youtube!). Look up a bit more information about what it is capable of and how it works. Is it intelligent? Is it autonomous?
- Email me a link to your chosen example, and we will discuss them in the next lecture.

“Vehicles”

- Thought-provoking book by Braitenberg
- Essential reading for the course (recommended purchase, also in library).
- Some copies to borrow, but must return for exchange by next lecture.

References

Valentino Braitenberg, “Vehicles: experiments in synthetic psychology”, MIT Press, Cambridge MA, 1984

Robin R. Murphy, “Introduction to AI Robotics”, MIT Press, Cambridge MA, 2000

Sebastian Thrun, Wolfram Burgard and Dieter Fox, “Probabilistic Robotics”, MIT Press, Cambridge MA, 2005

Rudiger Wehner “The architecture of the desert ant’s navigational toolkit” Myrmecological News 12:85-96, 2009