Artificial Intelligence is it finally arriving?

Leslie Smith
Computing Science and Mathematics
University of Stirling
May 2 2013.

UNIVERSITY OF STIRLING

Artificial Intelligence is it finally arriving?

Are we nearly there yet?

Leslie Smith
Computing Science and Mathematics
University of Stirling

UNIVERSITY OF STIRLING

Overview

- What *is* Artificial Intelligence?
  - A slippery customer….
- What do we actually *want* from Artificial Intelligence?
- …and when are we likely to get it?

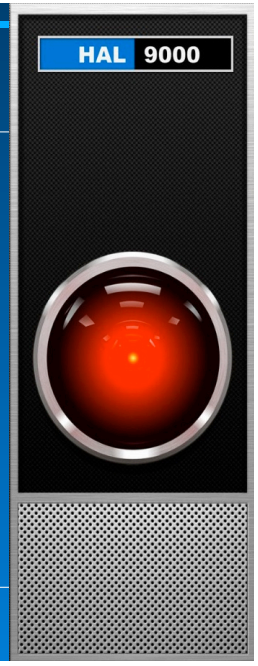

Why we need to know what AI is

"How can I go forward when I don't know which way I'm facing?" (John Lennon song, How, 1971)

Does climbing a tree get you closer to the stars?

What *is* intelligence?
Is animal intelligence the same as synthetic intelligence or is it simply different?

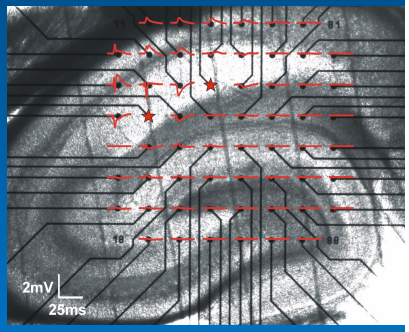

An aside on Natural Intelligence

- Difficult to define, and source of many conflicts
- Eysenck, Intelligence, and Intelligence tests
  - What do intelligence tests measure… or
  - Is intelligence defined by what intelligence tests measure?
- Racial and cultural issues
  - Cultural stereotyping in tests: test including questions on cars or planes, given to groups who hardly ever drive or fly
    - Has been a major issue …

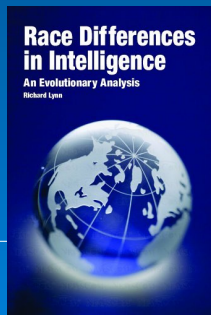

Easier question:

- What problems require intelligence in order to solve them?

Easier question:

- What problems require intelligence in order to solve them?

…But

Artificial Intelligence is a moving target

- As we solved problems believed to require human-like intelligence We have redefined quite what we mean by intelligence
  - Chess-playing?
    - Brute-force and clever searching techniques
  - Character and face recognition?
    - Data reduction, followed by a recogniser
  - Other game playing?
    - E.g. Backgammon and Tesauro's TDGammon
- Once we solve these problems
  - We tend to stop considering them to need intelligence.

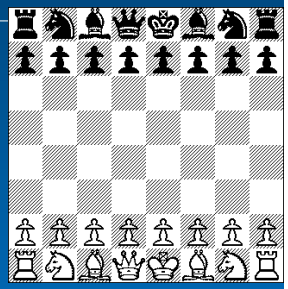

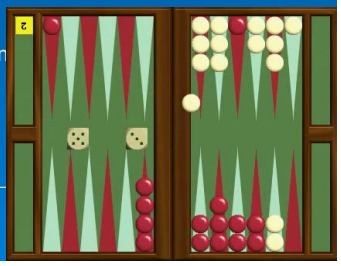

The Turing Test

Turing test

A test to empirically determine whether a computer has achieved intelligence

Alan Turing

An English mathematician wrote a landmark paper in 1950 that asked the question: *Can machines think?*

He proposed a test to answer the question "How will we know when we've succeeded?"

The Turing Test: setup

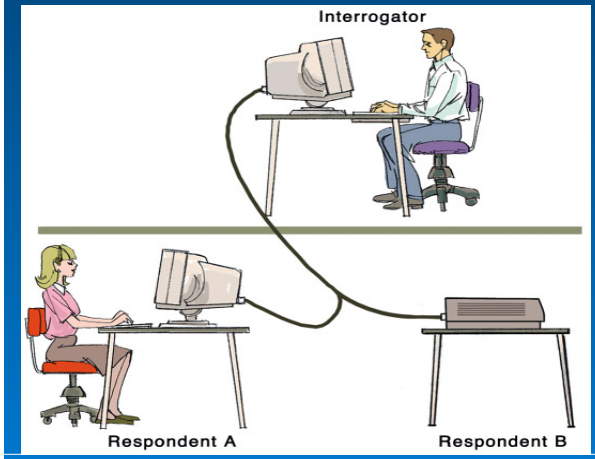

In a Turing test, the interrogator must determine which respondent is the computer and which is the human

The Turing Test: equivalences

Weak equivalence

Two systems (human and computer) are equivalent in results (output), but they do not arrive at those results in the same way

- Perhaps the computer is purely using logic

Strong equivalence

Two systems (human and computer) use the same internal processes to produce results

- Perhaps using software that works in the same way as the brain

Does the underlying mechanism matter for intelligence?

The Turing Test remains popular

Loebner prize

The first formal instantiation of the Turing test, held annually

Chatbots

A program designed to carry on a conversation with a human user

first one: Weizenbaum's Eliza.

This program identified specific key words in the users input and generated "appropriate" responses. Many people thought they were communicating with a person.

Try it http://www.manifestation.com/neurotoys/eliza.php3

What does the Turing test *really* test?

- What enables a machine to carry on a conversation?
  - Understanding what's been said
    - (actually typed, but we'll ignore that problem!)
  - Having the ability to synthesize a grammatical …
    - Or at least human-like
      - … response
  - Applying appropriate context to enable the generation of an appropriate response
- (and that's all assuming that there's no simpler mechanism for response generation in operation)
  - Could one achieve this algorithmically?
    - Does this require intelligence?

Searle's Chinese Room

Two different views of AI

- Top-down view:

$$p(\theta|\mathbf{X},\alpha) = \frac{p(\mathbf{X}|\theta)p(\theta|\alpha)}{p(\mathbf{X}|\alpha)} \propto p(\mathbf{X}|\theta)p(\theta|\alpha)$$

  - We can produce a nice clean mathematical abstraction of intelligence and implement it in electronics.
    - Logic, mathematics, rationality
- Bottom-up view:
  - We need to start by considering low-level issues: the construction of the brain, neuroscience, etc. Intelligence arises from the tissues
    - And it's not just a mathematical construct

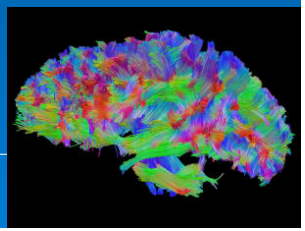

Of Brains

–Human brain
  – about 10,000,000,000 neurons

–C.elegans brain
  –323 Neurons

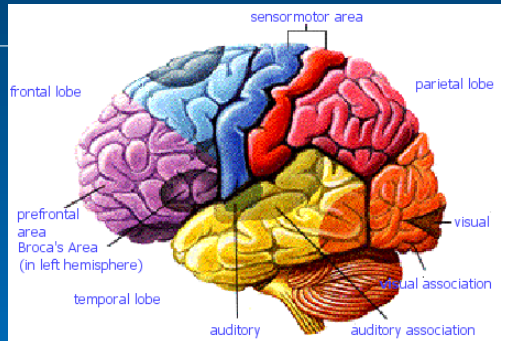

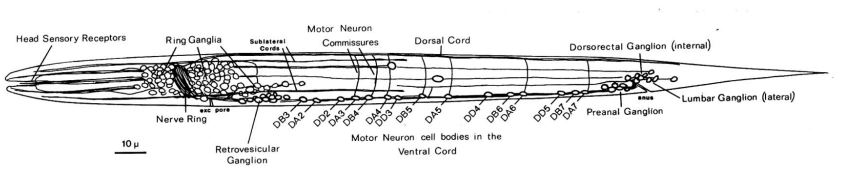

Of neurons

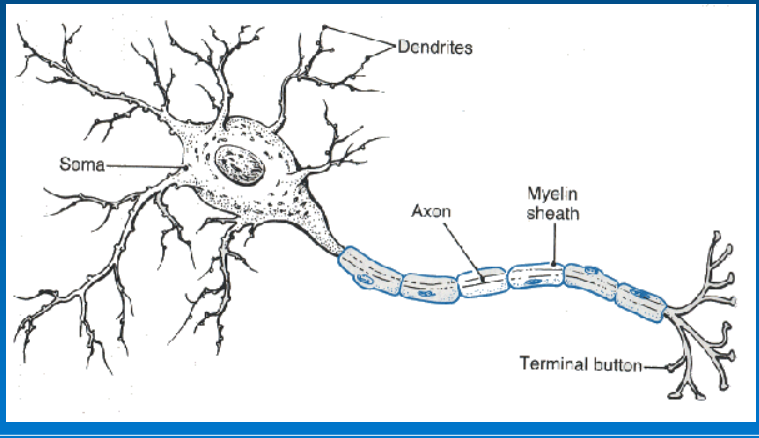

Real neural systems

- are highly parallel
  - All neurons and glial cells are all operating at once
  - (whether they are all actually in use at once is harder to answer!)
- are plastic
  - In the sense that they change with usage
    - Like most biological systems
- …are learning systems
  - So we may not need to program in all the eventualities that we expect them to meet
- A different view of intelligence
  - Ability to cope with novel situations

What gives brains intelligence?

- A hard question!
- One answer: a mixture of
  - Rapid capabilities, such as reflexes, and fast reactions
    - Probably using the brainstem and midbrain
  - Deliberative capabilities, recognising, interpreting, remembering, recalling, …
    - Probably using the cortex

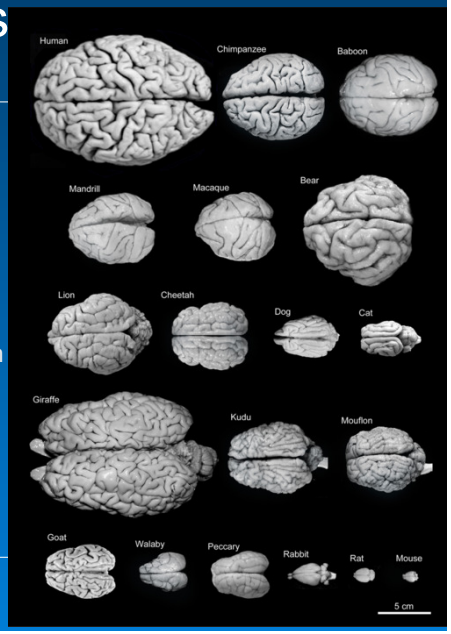

Artificial neural networks

- "Intelligence" modeled on the brain

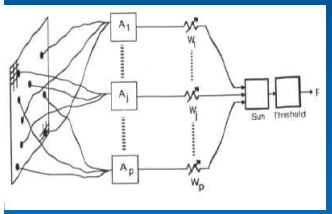

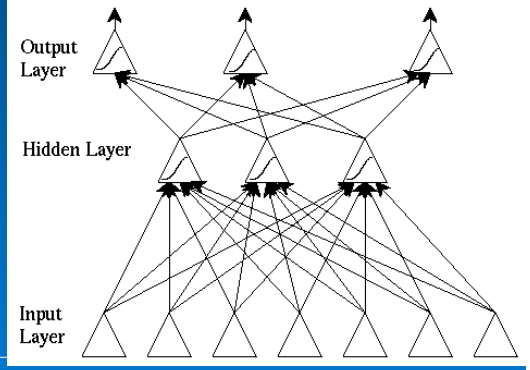

Neural approaches

- There are many neurally based algorithms
  - Some for classification
  - Some for function approximation
  - Some for identifying statistical regularities in input
- Sometimes seen as a branch of statistics
- Sometimes criticised for lacking explanation for their behaviour
  - But nonetheless used (e.g. in data mining, in real-time face detection in cameras, …)

But is this intelligence?

- Or just a collection of interesting techniques for doing
  - Pattern recognition
  - Game playing …
- Which we might combine with top-down techniques to provide a powerful system
  - But would it be intelligent?
  - Or more like something from Searle's Chinese Room?
- Could we do better by following the large-scale architecture of the brain?

What do we actually *want* from AI?

- A robotic helper for older people?

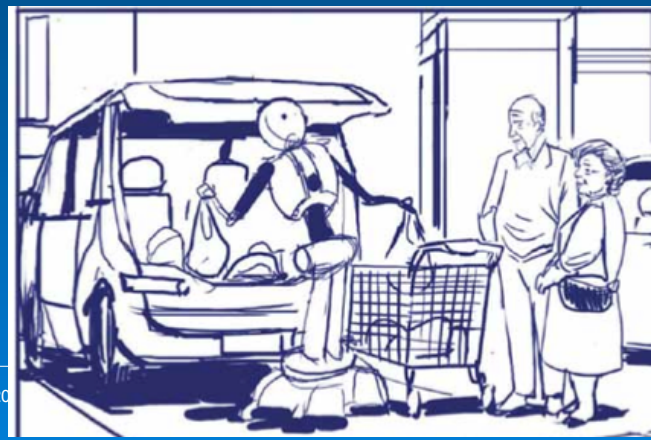

A humanoid robot that can do things?

- iCub

Something to solve real problems?

- In the 1980's Japan had the 5th generation initiative
  - Aim: to solve real (political, social) problems using top-down robotics
- Or something to perform difficult or tedious or dangerous jobs without supervision
  - Undersea inspection
  - Decommissioning nuclear power stations
  - Cleaning floors?

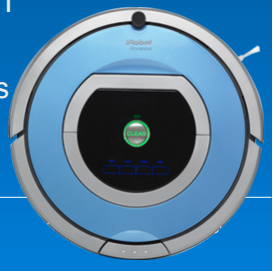

What's missing?

- These machines need
  - To be able to *understand* their environment
    - Though they might be able to use what's been learned by another robot
  - Autonomy
    - Do we really want to have to tell them everything that they have to do?
- These appear to imply *learning.*
  - Otherwise we need to supply an impossibly large amount of information
- Can we achieve this? And, if not, why not?

Understanding an environment

- Embodiment
  - Intelligence requires a universe to operate in!
- Sensory processing
  - From a camera image to a set of objects, perhaps moving objects, in an environment.
  - From what a microphone picks up to a set of sound sources in an environment.
- Scene Analysis
  - Visual, auditory, tactile, olfactory, …
  - Multimodal

Physical interaction.

- There's many levels to this
  - *Physical interaction.*
  - Picking up an egg, or picking up a cup, or picking up a block of wood.
    - These are different from each other!

Higher level interactions

- Which egg to pick up?

The everyday stuff…

- Smoothly walking over a rough surface
  - Like Boston Dynamics "Big Dog"
- Walking through a doorway without colliding with it
- Finding one's way to this lecture
  - And home again
- All of these really require understanding
  - The environment
  - And what one is trying to achieve…one's *goal*

Autonomy and AI

- Autonomy is the ability to make decisions
  - without human intervention
- The problem: unpredictability
  - One could attempt to program in decisions
    - If xxx then yyy …..
  - But the world is full of unpredictable and unexpected situations
    - Yet we (and AI machines) still need to take decisions
      - Often quite rapidly

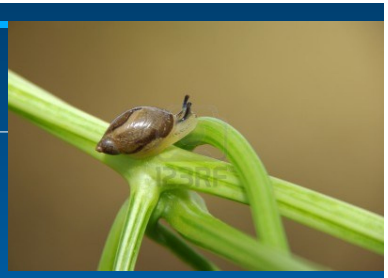

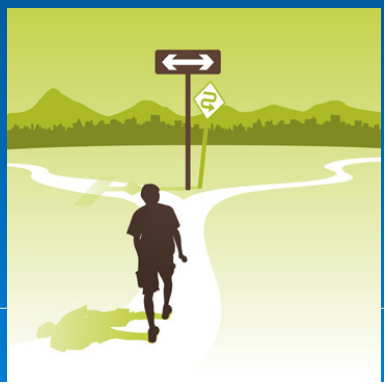

Goals and autonomy

- Making decisions implies knowing what the system is trying to do
  - Knowing its *goal*.
- But the goal may be too high level
  - *Keep the corridor free of obstructions*
- So we may need sub-goals
  - *Move objects out of the corridor*
    - *But which objects?*
- It can be hard to identify exactly what to do purely from the system's goals
  - Decisions are often very low-level, and sometimes the subgoals are less than obvious.

… and talking of goals…

- Robocup soccer: autonomous robots
- http://www.youtube.com/watch?feature=player_embedded&v=qDUOqtJvueM

Military robots

- Often not autonomous, but actually remote controlled
- Sometimes with a degree of autonomy
  - E.g. Boston Dynamics big dog.
- Ethics and military robotics

from whom is it ethical to accept research and development money? What attributes, such as weaponization, autonomy or intelligence, should I design into my technology? Which organizations and individuals should be allowed to buy and use my technology? Who should own or be able to access information gathered by my technology? If someone is harmed in association with the technology, who is responsible, and how is this determined? (Singer P.W., Nature 477 399-401, 22/9/2011)

Civilian AI systems

- Helpers for people?
- (Unsuccessful) EU Flagship proposal Robot Companions for Citizens
  - "Sustainable welfare through *sentient machines*: technology that integrates perception, cognition, emotion, and action with a contextual awareness of self, others, and the environment."
  - See http://www.robotcompanions.eu

Toyota's Robina

Building more human-like intelligence

- Historically, the human brain has often been compared to the most complex pieces of equipment that existed:
  - Mechanical, electrical systems, like a telephone exchange
  - Analogue and later digital computers
- But any analysis of brains quickly shows that these *may* be good analogies, but that's all they are.

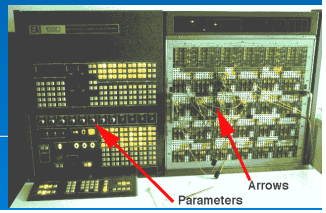

What makes the brain intelligent?

- A large number of highly integrated capabilities.
  - Fast reactions from the brainstem
    - Reflexes, plus short paths (few synaptic junctions) from input to output

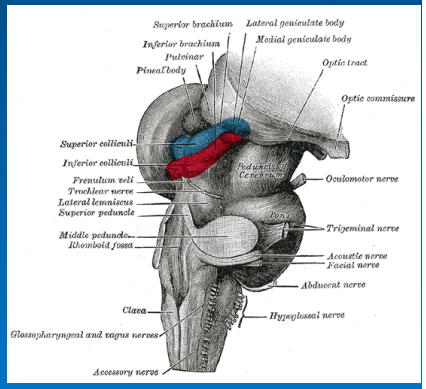

Midbrain and cortex

- Slightly slower reactions, introducing context
  - And cross modality integration
  - From the mid-brain.
- Slower, more deliberative capabilities from the cortex
  - Picking out statistical regularities in sensory input
  - Analysing these, organising these, remembering, recalling, …

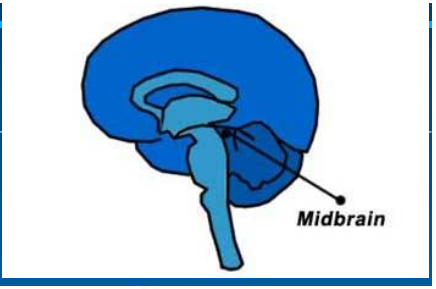

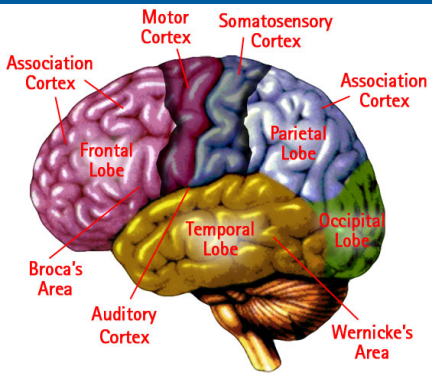

Intelligence is ..

- Making use of the large expanse of neural circuitry in the cortex.
  - Suggesting that the much increased size of the human cortex is critical for human intelligence
    - Perhaps linking this to relatively recent evolutionary changes
      - » Though these have many effects
- Note that these concepts assume that intelligence is embodied.
  - And not an abstract entity

Can we build this?

- Build a multi-level system with
  - Fast reactions (immediate reactions, analogue of brainstem: probably single sensory)
  - Slower reactions (quick reactions, analogue of midbrain: include correlation across modalities)
  - A large additional general-purpose data analysis and storage system (analogue of cortex)
    - Capable of extracting structure from single and multi-sensory datasets (reducing dimensionality, reducing free energy)

And then integrate these effectively (perhaps that's the hardest bit).

What's *still* missing?

- All this still doesn't help us to understand creativity, or insight.
  - Though insight might come from synchronisation across a number of different cortical modules

The Singularity

- Thesis: build an intelligent machine, and it will build an even more intelligent machine
- Issues: what *is* an intelligent machine?
  - If it's a question of integrating General Purpose slow and fast reaction systems, then there's no reason to believe that the thesis applies
  - But if there really is an identifiable abstract meaning to "intelligence" then perhaps there is a real issue

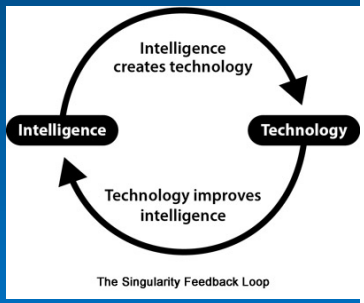

The Singularity Feedback Loop

Is AI finally arriving?

- There are some autonomous systems now being produced
  - And these can work in limited situations
- "Real" AI … (sometimes known as Artificial General Intelligence, AGI)
  - With all the connotations of human intelligence
- … is still a long way off
  - But the military are *very* interested, as are health organisations.
- I suggest we will see a gradual increase in "intelligent" autonomous systems in gradually increasingly complex environments.
  - And I don't believe in the Singularity!

End of lecture

- Questions, comments invited.