

Baron Wolfgang von Kempelen's *Chess-Playing Turk* (1769)

Lecture 1a


An Introduction to Artificial Intelligence

Alan Dorin

FIT3094 Artificial Life, Artificial Intelligence and Virtual Environments

artificial, adjective

made or produced by human beings rather than occurring naturally



The Terminator, James Cameron 1984

life, noun

the condition that distinguishes animals and plants from inorganic matter, including the capacity for growth, reproduction, functional activity, and continual change preceding death

intelligence, noun

the ability to acquire and apply knowledge and skills

Reading

Russell, S. J. and P. Norvig (2003). *Artificial Intelligence, A Modern Approach*, Prentice Hall. (chapter 1)

Langton, C. G. (1988). Artificial Life. *Artificial Life*. C. G. Langton, Addison-Wesley: 1-47

Levy, S. (1992). *Artificial Life, the quest for a new creation*. London, Penguin Books.

What is the field of *Artificial Intelligence* and what are its aims?

An AI system might...

think like humans

The exciting new effort to make computers think... machines with minds, in the full and literal sense – Haugeland 1985

[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning – Bellman 1978

act like humans

The art of creating machines that perform functions that require intelligence when performed by people – Kurzweil 1990

The study of how to make computers do things at which, at the moment, people are better – Rich and Knight 1991

The Thinker - Rodin 1902



think rationally


The study of mental faculties through the use of computational models – Charniak and McDermott 1985

The study of the computations that make it possible to perceive, reason, and act – Winston 1992

act rationally

Computational Intelligence is the study of the design of intelligent agents – Poole et al. 1998

AI... is concerned with intelligent behaviour in artifacts – Nilsson 1998





Paramecium

Acting like a human, The *Turing Test*



Getting this right requires:

- Natural language processing
- Knowledge representation
- Automated reasoning
- Machine learning




Acting rationally, the *autonomous agent*

A computer *agent* is able to act autonomously, perceive its environment, persist for an extended period of time, adapt to change, and adopt goals.

A rational agent is an agent that is able to act to achieve the expected best outcome (e.g. to persist or reach a goal).

How can we make the characters and opponents in a game *autonomous, rational agents*?

Is my
opponent a
person?



The philosophy of machine intelligence

1. Can formal rules be used to draw valid conclusions?

2. How does the mental mind arise from a physical brain?

3. Where does knowledge come from?

4. How does knowledge lead to action?



1. Can formal rules be used to draw valid conclusions?

Aristotle formulated rules for logical reasoning.

Ramon Lull built a machine for generating combinations of ideas based on an earlier Arabic device called the *zairja*.

Leonardo da Vinci designed a mechanical calculator.

Schickard and Pascal built calculating machines.

Hobbes proposed that reasoning was like numerical computation.

Leibniz constructed a simple mechanical concept manipulation device.

(384-322 BCE)



(1232-1315)



(1452-1519)



(1592-1635)



(1623-1662)



(1588-1679)



(1646-1716)



This is *history*. Why is history important?

2. How does the mental mind arise from a physical brain?

Descartes believed that mind and matter were distinct (*dualism*) – if mind is governed by physical law, what of free will? Only humans have a soul and free will he thought. (Animals are “just” machines)

If mind and matter are distinct, how do thoughts control a physical body?



Phoenix by Kare Halvorsen

Materialism: mind is matter operating within physical/chemical laws

(1596-1650)



*I think
therefore
I am*



Vaucanson's duck (1735)
– we'll meet it again!



A cockroach – can we replicate this level of intelligence in our machines?

Materialism: mind is matter running only on physical/chemical processes but...

3. Where does knowledge come from?

Francis Bacon → **John Locke**, empiricism - "Nothing is in the understanding, which was not first in the senses"

David Hume, proposed that repeated associations between elements leads to the formation of general rules - induction

Charles Babbage, devised a series of mechanical engines, predecessors to the modern computer, automating the calculation of complex mathematical tables.

Moritz Schlick, **Otto Neurath**, **Rudolf Carnap** developed logical positivism - knowledge is the connection of logical theories that are tied to sensory inputs.

(1561-1626)

(1632-1704)

(1711-1776)

(1791-1871)

(1882-1936)

(1882-1945)

(1891-1970)



Carnap's book, *The Logical Structure of the World*, (1928) defined an explicit computational procedure for extracting knowledge from experiences... mind as computational process.

4. How does knowledge lead to action?

I need covering; a cloak is a covering; I need a cloak.

What I need, I have to make; I need a cloak.
I have to make a cloak.



cloak

Aristotle
(384-322 BCE)

Aristotle's algorithm, now called *regression planning*:

We deliberate not about ends but about means. For a doctor does not deliberate whether he shall heal, nor an orator whether he shall persuade...

They assume the end and consider how and by what means it is to be attained; and if it seems to be produced by several means they consider by which it is most easily and best produced, while if it is achieved by one only they consider how it will be achieved by this and by what means this will be achieved, till they come to the first cause...

... and what is last in the order of analysis seems to be first in the order of becoming. And if we come on an impossibility, we give up the search, e.g. if we need money and this cannot be got; but if a thing appears possible we try to do it.

Nicomachean Ethics (Book III. 3, 1112b), Aristotle

The Church-Turing thesis

A Universal Turing machine (symbolised U) is an “idea” for a basic abstract symbol-manipulating device with certain properties...


U is capable of computing all and only outcomes of definite procedures.



Alonzo Church
(1903-1995)

I agree with you my student,
although I'd say it a bit
differently.

There are however, some functions
that are not computable.



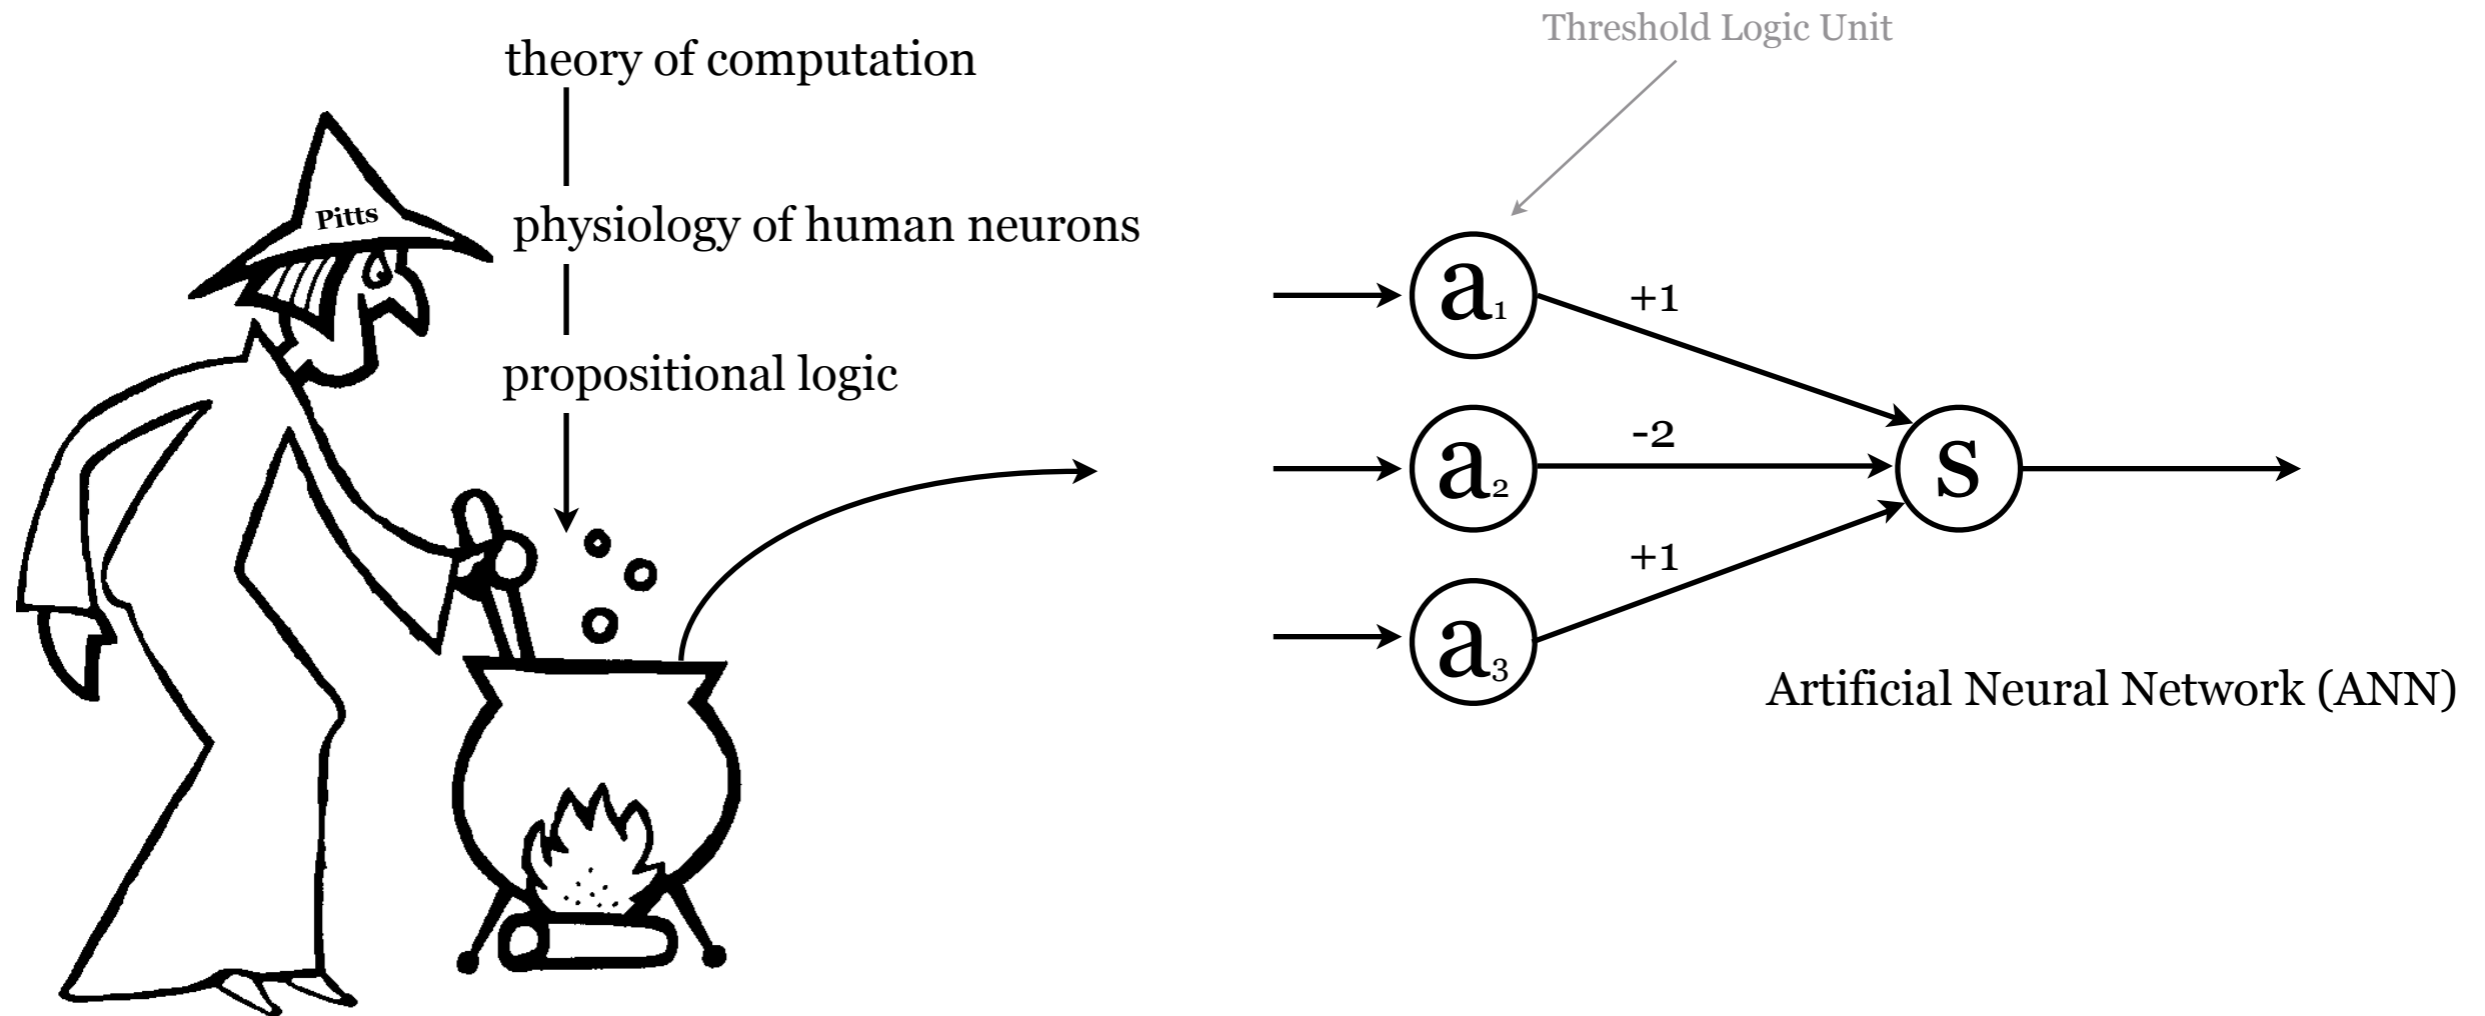
Alan Turing
(1912-1954)

Is *intelligence* a computable function?
Can rational thought be captured in an algorithm?

If so, is the algorithm **tractable**?



The beginnings of the field of Artificial Intelligence.



McCulloch and **Pitts** threw some ideas together and came up with the idea of an *artificial neural network* (1943)

Hebb later demonstrated an updating rule for the weights between neurons that allowed the neural network to *learn* (1949)

Minsky and **Edmonds** built the first neural network computer (1951)

ANNs reappear later in the course.

Dartmouth Summer Research Workshop (1956)

McCarthy, Minsky, Shannon, Rochester, More, Samuel, Solomonoff, Selfridge, Newell & Simon attended.

They shaped AI in the USA over the next 20 years.

“We should call this field
Artificial Intelligence”



Trenchard More John McCarthy Marvin Minsky Oliver Selfridge Ray Solomonoff

Reunion at the July AI@50 conference, Dartmouth.
Photo by Joseph Mehling.

<http://www.dartmouth.edu/~vox/0607/0724/ai50.html>



Herbert
Simon

Allen
Newell



Allen Newell
(1927 - 1992)



Herbert Simon
(1916 - 2001)

At the time of the Dartmouth workshop, Newell & Simon already had a reasoning program, *Logic Theorist*, capable of “thinking non-numerically”.

Later they developed *General Problem Solver*.

Their success lead them to the **Physical Symbol System Hypothesis***...

A physical symbol system has the necessary and sufficient means for general intelligent action



Phew!



blip, blip-bleep.
(I'm not so sure!)

* *Computer Science as Empirical Inquiry: Symbols and Search*, Turing Award Lecture, 1975

The 50s and 60s

Programs were made that could prove theorems
tricky for many students – Gelernter

A learning program was written to play draughts (thereby disproving the hypothesis that computers could only do what they were told to do) – Samuel

The AI programming language LISP was defined - McCarthy

```
(defun factorial (n)
  (if (<= n 1)
      1
      (* n (factorial (- n 1)))))
```

Minsky supervised some students tackling limited problems operating in domains called *microworlds*. These seemed to require intelligence to solve...



SHRDLU

Terry Winograd's PhD thesis (1971) described a program SHRDLU for understanding natural language in a microworld of blocks on a table top.

Person: PICK UP A BIG RED BLOCK.

Computer: OK. (does it)

Person: GRASP THE PYRAMID.

Computer: I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN.

Person: FIND A BLOCK WHICH IS TALLER THAN THE ONE YOU ARE HOLDING AND PUT IT INTO THE BOX.

Computer: BY "IT", I ASSUME YOU MEAN THE BLOCK WHICH IS TALLER THAN THE ONE I AM HOLDING.

Computer: OK. (does it)

Person: WHAT DOES THE BOX CONTAIN?

Computer: THE BLUE PYRAMID AND THE BLUE BLOCK.

Person: WHAT IS THE PYRAMID SUPPORTED BY?

Computer: THE BOX.

Person: HOW MANY BLOCKS ARE NOT IN THE BOX?

Computer: FOUR OF THEM.

Person: IS AT LEAST ONE OF THEM NARROWER THAN THE ONE WHICH I TOLD YOU TO PICK UP?

Computer: YES, THE RED CUBE.

Person: IS IT SUPPORTED?

Computer: YES, BY THE TABLE.

Person: CAN THE TABLE PICK UP BLOCKS?

Computer: NO.

Person: CAN A PYRAMID BE SUPPORTED BY A BLOCK?

Computer: YES.

Person: CAN A PYRAMID SUPPORT A PYRAMID?

Computer: I DON'T KNOW.

Person: STACK UP TWO PYRAMIDS.

Computer: I CAN'T.

Person: WILL YOU PLEASE STACK UP BOTH OF THE RED BLOCKS AND EITHER A GREEN CUBE OR A PYRAMID?

Computer: OK. (does it)

Person: CALL THE BIGGEST BLOCK "SUPERBLOCK".

Computer: OK.

Person: HAVE YOU PICKED SUPERBLOCK UP SINCE WE BEGAN?

Computer: YES.

Person: WHY DID YOU DROP IT?

Computer: BY "IT", I ASSUME YOU MEAN THE SUPERBLOCK.


I CAN'T EXPLAIN A NON-EXISTENT EVENT.

Person: IS THERE ANYTHING TO THE RIGHT OF THE RED PYRAMID?

Computer: YES, FOUR THINGS: THE BOX, THE RED CUBE, THE BLUE PYRAMID AND THE BLOCK NAMED SUPERBLOCK

Person: THANK YOU.

Computer: YOU'RE WELCOME!



When will (some) intelligence be artificial?

Alan Turing (1950): by 2000

Herbert Simon (1959): within 10 years...

A computer will be chess champion and a significant mathematical theorem will be proven by machine.

Ray Kurzweil (2008): 2029...

I've made the case that we will have both the hardware and the software to achieve human level artificial intelligence with the broad suppleness of human intelligence including our emotional intelligence by 2029.

When will (some) intelligence be artificial?



Lost In Space, Irwin Allen (83 episodes, 1965-1968) 1997

In 1997, the Earth suffers from overpopulation and natural resource depletion. Professor John Robinson, his wife and children, their friend and pilot, Major Don West, are chosen to travel on the Jupiter 2 to Alpha Centauri to search for a habitable planet for mankind to colonize.

The Robot is a Model B-9, Class M-3 General Utility Non-Theorizing Environmental Control Robot, which had no given name. Although a machine endowed with superhuman strength and futuristic weaponry, he often displayed human characteristics such as laughter.



2001: A Space Odyssey, Stanley Kubrick, 1969 2001

Dave Bowman: Hello, HAL do you read me, HAL?

HAL: Affirmative, Dave, I read you.

Dave Bowman: Open the pod bay doors, HAL.

HAL: I'm sorry Dave, I'm afraid I can't do that.

Dave Bowman: What's the problem?

HAL: I think you know what the problem is just as well as I do.

Dave Bowman: What are you talking about, HAL?

HAL: This mission is too important for me to allow you to jeopardize it.

Dave Bowman: I don't know what you're talking about, HAL?

HAL: I know you and Frank were planning to disconnect me, and I'm afraid that's something I cannot allow to happen.

Dave Bowman: Where the hell'd you get that idea, HAL?

HAL: Dave, although you took thorough precautions in the pod against my hearing you, I could see your lips move.



STANLEY KUBRICK'S
2001:
a space odyssey



The Terminator, James Cameron, 1984 2029

An advanced model of the T-600 is created, with cloned skin made to look like ordinary humans. The new T-800 is now made of living tissue over a metal endoskeleton. They can still be detected by dogs. Cyberdyne Machine Network sends a T-800, Model 101 back to 1984, and Tech-Com sends Kyle Reese to follow it.

AI has had some measured success to date, especially where it is possible to amass human knowledge and use it methodically to solve problems... e.g. expert systems for medical diagnosis.

AI has had *many* millions of dollars invested in it. This has not resulted in the general success that was anticipated by researchers and science fiction.

AI turned out to be hard!

Still, AI is a helpful ingredient of many modern software systems.



For now though, let's look at

Artificial *Life*...

**GHOST IN THE SHELL 2:
INNOCENCE**

TM and © 2004 DreamWorks, LLC. All rights reserved.