Artificial Intelligence

Chapter 1

Outline

- ♢ What is AI?
- ♢ A brief history
- ♢ The state of the art

Artificial Intelligence

:: **Artificial Intelligence** explores and investigates various kinds of intelligent behavior by attempting to create it in the artificial.

– such as problem solving, learning, natural language understanding, deduction ...

– the research vehicle is (experimenting medium) is computer

:: according to **Marvin Minsky** Artificial Intelligence is a science of how to persuade computers to exhibit such a type of behaviour that conventionally requires Human Intelligence

:: Artificial Intelligence is sometimes called **machine intelligence** or **computational psychology**

What is AI?

Systems that think like humans	**Systems that think rationally**
Systems that act like humans	**Systems that act rationally**

Acting humanly: The Turing test

Turing (1950) "Computing machinery and intelligence":

- "Can machines think?" $\longrightarrow$ "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game

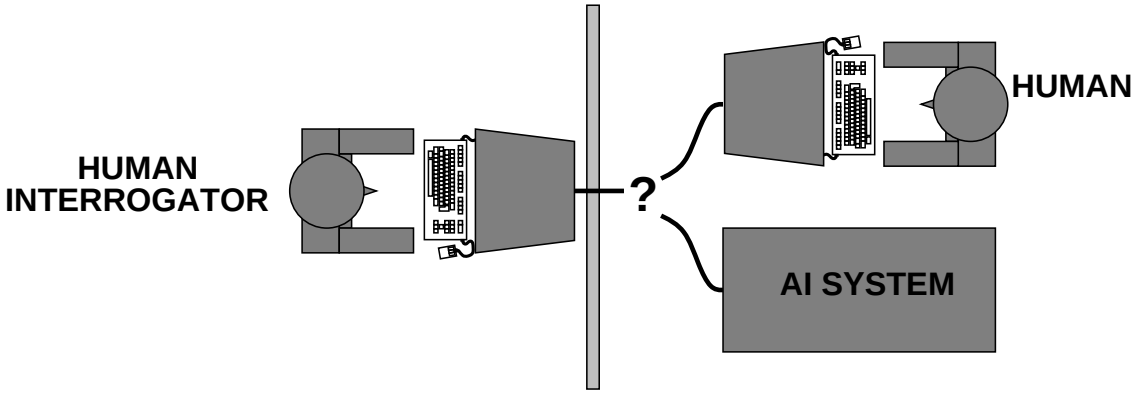

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning

Problem: Turing test is not **reproducible**, **constructive**, or amenable to **mathematical analysis**

Thinking humanly: Cognitive Science

1960s "cognitive revolution": information-processing psychology replaced prevailing orthodoxy of behaviorism

Requires scientific theories of internal activities of the brain

- What level of abstraction? "Knowledge" or "circuits"?
- How to validate? Requires
  1) Predicting and testing behavior of human subjects (top-down)
  or 2) Direct identification from neurological data (bottom-up)

Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI

Both share with AI the following characteristic:

the available theories do not explain (or engender) anything resembling human-level general intelligence

Hence, all three fields share one principal direction!

Thinking rationally: Laws of Thought

Normative (or prescriptive) rather than descriptive

Aristotle: what are correct arguments/thought processes?

Several Greek schools developed various forms of logic:
notation and **rules of derivation** for thoughts;
may or may not have proceeded to the idea of mechanization

Direct line through mathematics and philosophy to modern AI

Problems:
1) Not all intelligent behavior is mediated by logical deliberation
2) What is the purpose of thinking? What thoughts **should** I have out of all the thoughts (logical or otherwise) that I **could** have?

Acting rationally

Rational behavior: doing the right thing

The right thing: that which is expected to maximize goal achievement, given the available information

Doesn't necessarily involve thinking—e.g., blinking reflex—but thinking should be in the service of rational action

Aristotle (Nicomachean Ethics):

> **Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good**

Rational agents

An agent is an entity that perceives and acts

This course is about designing rational agents

Abstractly, an agent is a function from percept histories to actions:

$$f : \mathcal{P}^* \rightarrow \mathcal{A}$$

For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance

Caveat: **computational limitations make perfect rationality unachievable**

$\rightarrow$ design best program for given machine resources

AI prehistory

Philosophy	logic, methods of reasoning
	mind as physical system
	foundations of learning, language, rationality
Mathematics	formal representation and proof
	algorithms, computation, (un)decidability, (in)tractability
	probability
Psychology	adaptation
	phenomena of perception and motor control
	experimental techniques (psychophysics, etc.)
Economics	formal theory of rational decisions
Linguistics	knowledge representation
	grammar
Neuroscience	plastic physical substrate for mental activity
Control theory	homeostatic systems, stability
	simple optimal agent designs

Potted history of AI

1943	McCulloch & Pitts: Boolean circuit model of brain
1950	Turing's "Computing Machinery and Intelligence"
1952–69	Look, Ma, no hands!
1950s	Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
1956	Dartmouth meeting: "Artificial Intelligence" adopted
1965	Robinson's complete algorithm for logical reasoning
1966–74	AI discovers computational complexity Neural network research almost disappears
1969–79	Early development of knowledge-based systems
1980–88	Expert systems industry booms
1988–93	Expert systems industry busts: "AI Winter"
1985–95	Neural networks return to popularity
1988–	Resurgence of probability; general increase in technical depth "Nouvelle AI": ALife, GAs, soft computing
1995–	Agents, agents, everywhere . . .
2003–	Human-level AI back on the agenda

State of the art

Which of the following can be done at present?

♢ Play a decent game of table tennis

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web
- ♢ Buy a week's worth of groceries at Berkeley Bowl

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web
- ♢ Buy a week's worth of groceries at Berkeley Bowl
- ♢ Play a decent game of bridge

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web
- ♢ Buy a week's worth of groceries at Berkeley Bowl
- ♢ Play a decent game of bridge
- ♢ Discover and prove a new mathematical theorem

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web
- ♢ Buy a week's worth of groceries at Berkeley Bowl
- ♢ Play a decent game of bridge
- ♢ Discover and prove a new mathematical theorem
- ♢ Design and execute a research program in molecular biology

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web
- ♢ Buy a week's worth of groceries at Berkeley Bowl
- ♢ Play a decent game of bridge
- ♢ Discover and prove a new mathematical theorem
- ♢ Design and execute a research program in molecular biology
- ♢ Write an intentionally funny story

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web
- ♢ Buy a week's worth of groceries at Berkeley Bowl
- ♢ Play a decent game of bridge
- ♢ Discover and prove a new mathematical theorem
- ♢ Design and execute a research program in molecular biology
- ♢ Write an intentionally funny story
- ♢ Give competent legal advice in a specialized area of law

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web
- ♢ Buy a week's worth of groceries at Berkeley Bowl
- ♢ Play a decent game of bridge
- ♢ Discover and prove a new mathematical theorem
- ♢ Design and execute a research program in molecular biology
- ♢ Write an intentionally funny story
- ♢ Give competent legal advice in a specialized area of law
- ♢ Translate spoken English into spoken Swedish in real time

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web
- ♢ Buy a week's worth of groceries at Berkeley Bowl
- ♢ Play a decent game of bridge
- ♢ Discover and prove a new mathematical theorem
- ♢ Design and execute a research program in molecular biology
- ♢ Write an intentionally funny story
- ♢ Give competent legal advice in a specialized area of law
- ♢ Translate spoken English into spoken Swedish in real time
- ♢ Converse successfully with another person for an hour

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web
- ♢ Buy a week's worth of groceries at Berkeley Bowl
- ♢ Play a decent game of bridge
- ♢ Discover and prove a new mathematical theorem
- ♢ Design and execute a research program in molecular biology
- ♢ Write an intentionally funny story
- ♢ Give competent legal advice in a specialized area of law
- ♢ Translate spoken English into spoken Swedish in real time
- ♢ Converse successfully with another person for an hour
- ♢ Perform a complex surgical operation

What Belongs to Artificial Intelligence

:: **Artificial Intelligence** provides answers to the following questions:

- Automated Reasoning, Theorem Proving and Game Playing
- Expert and Knowledge-Based Systems
- Natural Language Processing and Semantic Modelling
- Robotics, Planning and Scheduling
- Recognition (machine vision, speech recognition)
- Machine Learning
- Evolutionary Techniques
- Neural Networks and Parallel Distributed Processing
- Distributed Artificial Intelligence and Multi-Agent Systems
- Philosophy Artificial Intelligence

Three Schools of Artificial Intelligence

:: Symbolic Functionalism

– intelligence represented in symbols and mutual manipulations

– formal models of reasoning, knowledge based

:: Connectionism

– inspired by natural processes

– emergency of intelligent behaviour

:: Robotics Functionalism

– computational implementation of behaviourism

Symbolic Functionalism

Good Old Fashioned Artificial Intelligence GOFAI

:: Functional Hypothesis

– Intelligent behaviour of a system is achieved through interaction among several systems components of different functionality

:: Physical Symbol Systems Hypothesis

– Physical Symbol Systems is an inevitable and satisfactory mean for implementation of Intelligent behaviour

Physical Symbol Systems Hypoythesis

:: **Physical Symbol Systems** (PSS) is a general machine providing evolution of population of symbol structures in time.

– Fundamental building primitives of the systems are symbols and operators of creation, destruction, modification and reproduction.

:: an example of PSS is a Turing Machine

– with the tape as the only symbol structure and single operator of modification

– or with a symbol as a single symbol structure and applying operators of destruction and creation

:: **Symbolic Functionalism** within AI concerns primarily methods for knowledge representation and reasoning models

Turing Machine

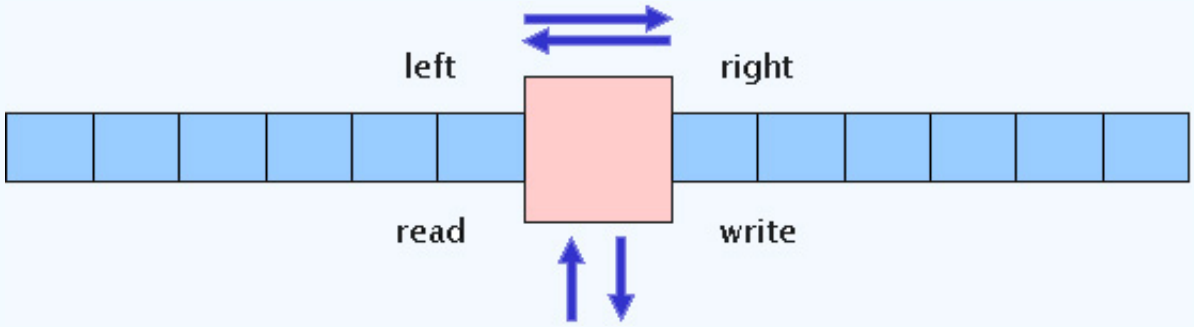

© Michal Pechoucek 2003 – Cybernetics and Artificial Intelligence
Czech Technical University and State University of New York

Connectionism

:: an essence of intelligence is rooted in static interconnection of huge number of simple processing units and in parallel behaviour of the system as a whole.

:: **connectionism** takes its inspiration from the **human brain** operations that is a unarguably a medium of intelligent behaviour

:: **neuron** as a this simple processing unit gets excited according to weighted sum of incoming precepts

:: **artificial neural networks** belong to the field of connectionism

Neuron and Artificial Neuron?

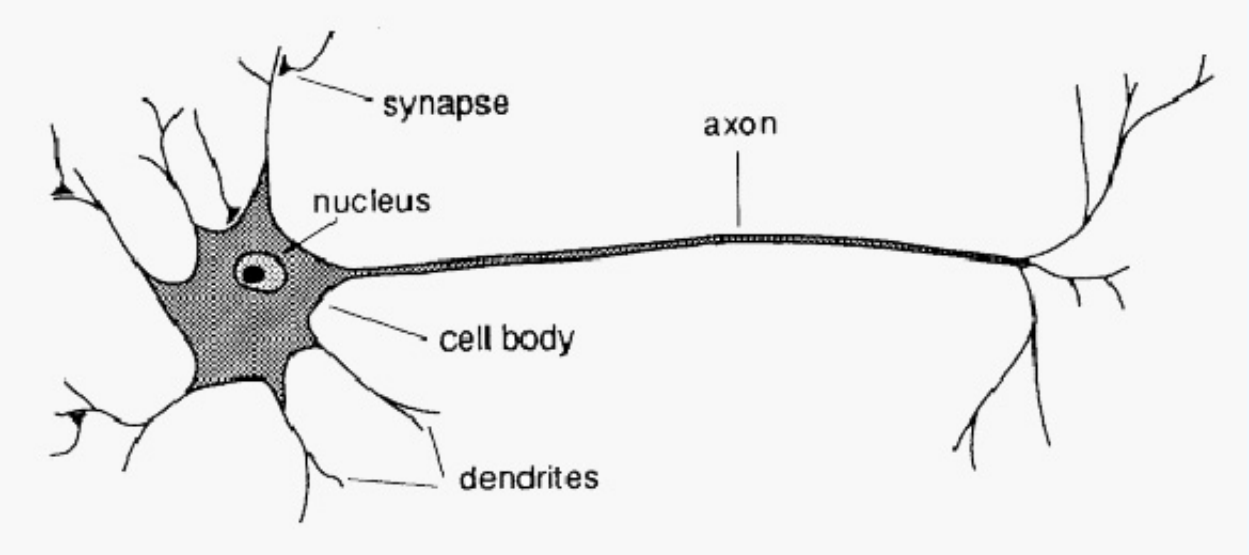

Neuron and Artificial Neuron?

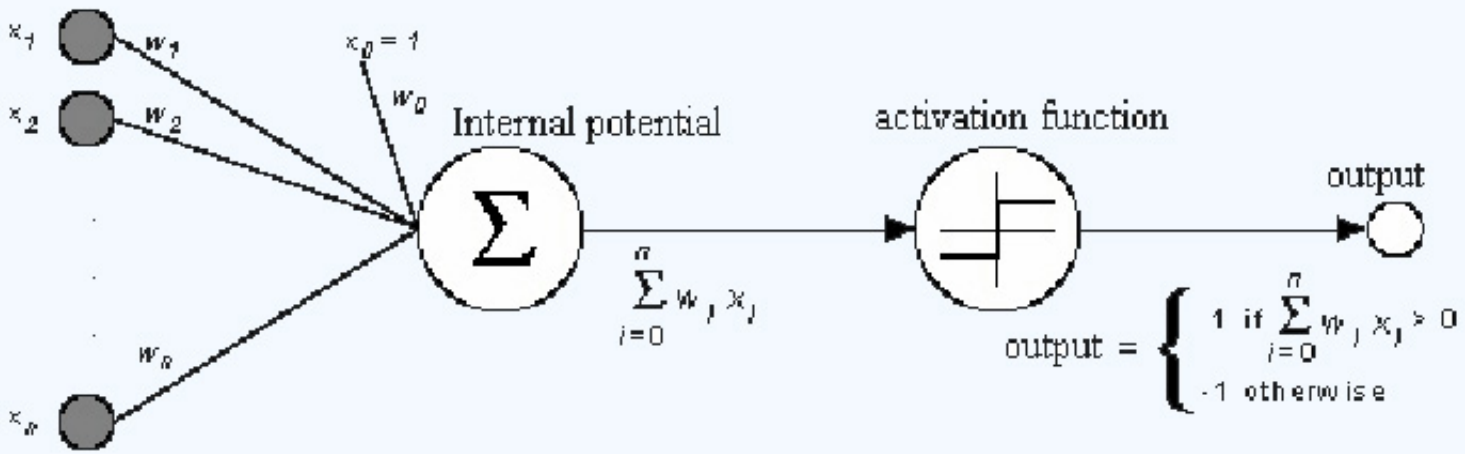

Connectionism cont.'

:: isnt it just an instance of physical symbol system?

– intelligence is based on interaction among components

– components are same and they are static in time

:: **black box paradigm** we do not need to understand how the brain functions when solving a problem

– good for classification type of reasoning rather than deduction

– knowledge is not modelled

A Robot?

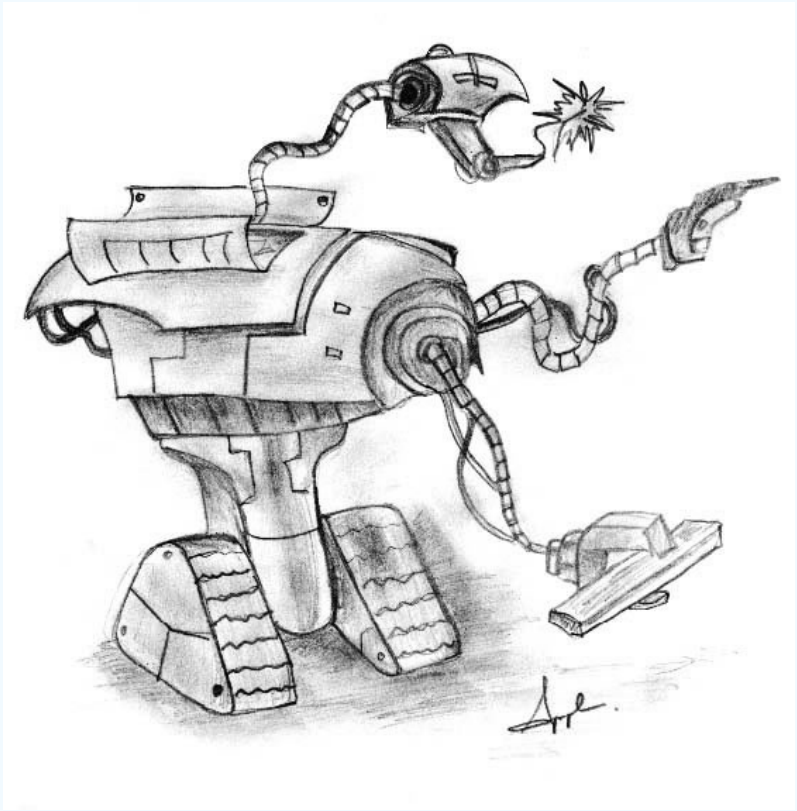

Reative Intelligence & Robotic Functionalism

- **Reactive Intelligence** is viewed as a sensible interaction between the intelligent system and the environment in which it is situated. Very often implemented as a huge collection of simple if-then rules
- Though the intelligence here is also rooted in interaction among systems components, unlike **connectionism** they are heterogeneous and hierarchically structured
- **Robotic Functionalism** is based on computational implementation of the behaviourism as a psychological school of thought
- While **symbolic functionalism** investigates reasoning background behind intelligent behaviour **robotic functionalism** rather explores and simulates outer of intelligent behaviour

Emerging Fields in Artificial Intelligence

:: Genetic Algorithms

– natural selection (crossover, mutation of population of solutions)

– due to its natural inspiration people classify genetic algorithms as a part of connectionism (it however does not seem to fit connectionism definition sounds more like PSS)

:: Distributed Artificial Intelligence

– concept of agency, collaboration between independent computational units

– resembles features of connectionism

:: Artificial Life

– simulation of emergent life-like properties in computer

Models of Understanding

let us view extent of artificial intelligence in terms of ability to understand
we distinguish **weak property** of a system (its specificity) and **strong property** (its generality)

:: weak understanding (Turing)

– such an understanding if inputs so that it makes the system to react intelligently (as human)

:: strong understanding (Brentano)

– such that the system feels the feeling of understanding as human do

this is why we distinguish strong and weak artificial intelligence
the concept of **turing test**

Middling Artificial Intelligence

:: According to Smith there is **middling artificial intelligence**, where the right type of behaviour reacting to given precepts is done through an appropriate knowledge structures and reasoning machinery.

- Robotics Functionalism Weak Artificial Intelligence
- Symbolic Functionalism Middling Artificial Intelligence
- Connectionism Strong Artificial Intelligence (???)

:: Example of natural language understanding

We claim that there is nothing like Strong Artificial Intelligence

Modeling Intelligent Reasoning

- **theoretical reasoning** - models of deduction, theorem proving strong approach high computational complexity (calculative rationality bottleneck)
- **knowledge-base reasoning** - (reactive) hardwired reasoning rules, sometimes with memory weak approach (case specific) reasonable complexity
- **state-space search** - searching the space of possible solution with a goal to find a solution high computational complexity state-space pruning by heuristics
- **practical reasoning** - reasoning in an intentional system based on models of intentions, beliefs and desires combination of knowledge-base reasoning (deliberation) and state-space search (planning)