context and history of artificial intelligence

lecture 1 of Research Seminar *Artificial Intelligence*

Maarten Lamers
Media Technology MSc programme, Leiden University

course webpage:
www.maartenlamers.com/RS

context and history of artificial intelligence

- follows the introduction chapter of

  Stuart Russell and Peter Norvig,
  "A*rtificial Intelligence, a Modern Approach*"
  (second edition), 2003

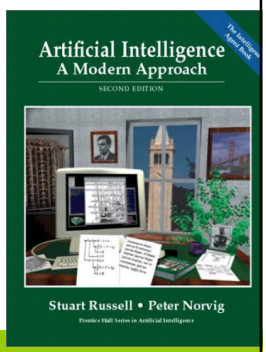

- topics
  1. what is AI?
  2. its foundations
  3. its history

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

context and history of artificial intelligence

artificial intelligence is one of three most mentioned topics, when scientists are asked what field they wish they worked in (Russel & Norvig)

what is AI?

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

what is AI?

- wanting to build *intelligent systems*
- there are four views of what “intelligent systems” do

	like humans	“ideally”
reasoning	*thinking humanly*	*thinking rationally*
behaving	*acting humanly*	*acting rationally*

- with “rationally” we mean coming up with the “best solution” given what the system knows

acting humanly: Turing Test

- Alan Turing (UK, 1912-1954)
- *“can machines behave intelligently?”*
- "*Computing Machinery and Intelligence*“, Mind Vol 49 Num 236, 1950.
- operational test for intelligent behavior

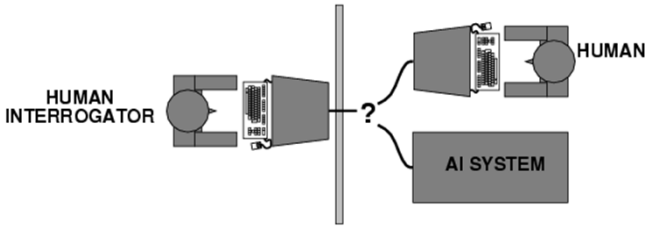

acting humanly: Turing Test

- he predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- the test is still relevant, after 60 years passed
- he suggested major components of AI
  - *natural language processing*
  - *representation of knowledge*
  - *automated reasoning*
  - *learning*

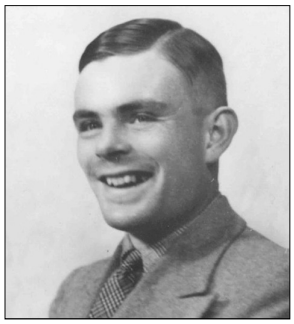

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

thinking humanly: cognitive modeling

- based on the 1960s "cognitive revolution“
- psychologists started viewing the brain as an information-processing device
- goal: *search for scientific theories of internal brain activities*
- how do you validate these theories?

1. predicting and testing behavior of humans (cognitive psychology)
2. direct identification from neurological data (cognitive neuroscience)

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

thinking rationally: the laws of thought

- Aristotle: *what are "correct" reasoning processes?*
- when given correct premises, how can we always reach a correct conclusion?

- by 1965, computers could *in principle* solve *any* logical problem described in logic notation

- problems
  1. in logic, how do you describe *informal knowledge* that may not be 100% certain?
  2. "in principle" is one thing, "in practice" is another! (*combinatorial explosion*)

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

acting rationally: rational agents

- an *agent* is something that *acts* (Latin *agere* = to do)
- rational behavior: come up with the "best solution" given what is known

- a rational agent acts to achieve the best outcome, or in the case of uncertainty, the best expected outcome

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

acting rationally: rational agents

- finding the best outcome does not necessarily require correct reasoning
  1. sometimes there is no provably correct action
  2. thinking is not always required, such as with reflexes

to understand the problem of “finding the best solution”,
you must know about

search space & *combinatorial explosion*

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

search space

- many problems can *in principle* be solved by *brute force*: just search the set of all possible solutions (*search space*) for the best solution
- for example, *traveling salesman problem* (TSP)

  what is the shortest route a salesman could take to visit a collection of cities only once and return to the starting point?

- every tour through all cities is a possible solution

search space

- can you search the complete search space?
  can you try all possible routes, and pick the shortest one?

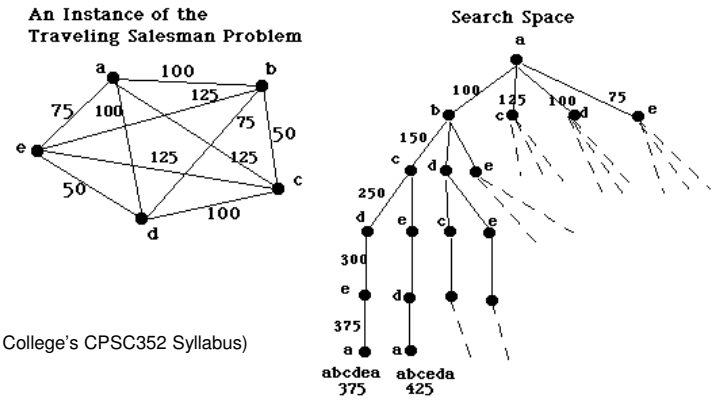

(images from Trinity College's CPSC352 Syllabus)

search space

- the strategy "*try all possible solutions, then pick the best one"* is called

  brute force approach

- it is when you search the *complete* search space, so that you *know* what is the best solution

combinatorial explosion

combinatorial explosion (a.k.a. *intractability)*

- when search space grows exponentially with the problem size
- for many problems, brute force searching becomes practically impossible!

combinatorial explosion

- problem size of TSP is the number of cities *n*
- number of TSP solutions for *n > 2* cities is *(n-1)!/2*

n = 10 → *181440 solutions*
n = 20 → *60822550204416000 solutions*
n = 100 → *4.66 • 10^{155} solutions*

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

heuristic approach

- rational agents often choose a *heuristic approach*
- search for a *reasonable* solution, using "rules of thumb"
  - educated guesses
  - trial and error
  - common sense

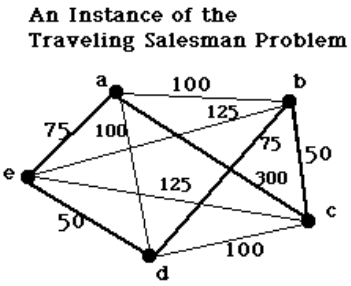

- example TSP heuristic:
  it seems un-logical to travel large distances between two cities → always go to nearest neighbor city.
- heuristics may lead to sub-optimal solutions!

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

a great part of AI, the part that tries to *act rationally*,
is concerned with

avoiding brute force search & combinatorial explosion

foundations of AI

which fields have contributed to what is now artificial intelligence research?

foundations of AI: philosophy

- can formal rules be used to draw valid conclusions?
- can mental mind arise from physical brain?
- where does knowledge come from?
- how does knowledge lead to action?

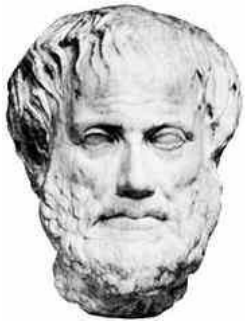

- logic + reasoning methods
- view of mind as a physical system
- foundations of learning

- concept of *induction*: acquiring general rules from repeated associations between their elements

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

intermezzo: induction versus deduction

- *inductive reasoning* (a.k.a. *generalization*)
  making generalizations based on individual instances

  "every time I touched ice, it was cold"
  → "all ice is cold"

- *deductive reasoning*
  deduce things that must be true when premises are true

  "all men are mortal" + "Socrates is a man"
  → "Socrates is mortal"

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

intermezzo: induction versus deduction

- humans are really good at generalization
- computers are really bad at it

	deductive reasoning	inductive reasoning (generalizing)
humans	difficult	easy
computers	easy	difficult

- *many interesting AI problems deal with making computers good at generalization*
- for example, recognizing dogs in photographs

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

intermezzo: generalization

foundations of AI: mathematics

- what are the formal rules to draw valid conclusions?
- what can be computed?
- how to reason with uncertainty?

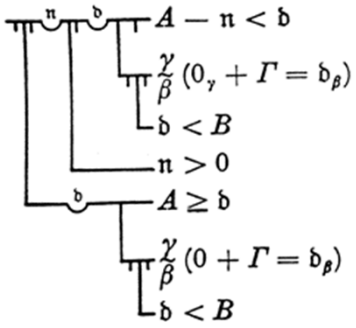

- logic and its representation
- probability theory

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

foundations of AI: mathematics

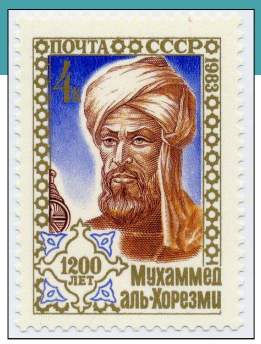

- *algorithms* (Al-Khwārizmī , 900AC)
- some functions on integers cannot be done in algorithm (Gödel, Penrose)
- Turing Machine can compute *any* computable function (*universal computation*), but some functions it cannot compute
- intractability, NP-completeness

foundations of AI: neuroscience

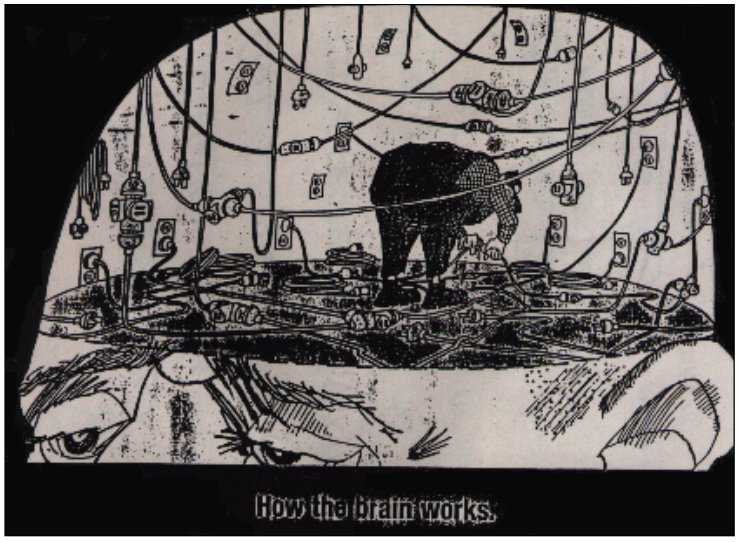

foundations of AI: neuroscience

- how do brains process information?
- neurons

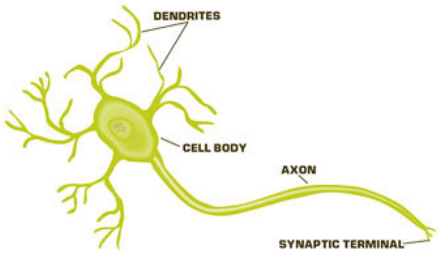

- *a collection of simple cells can lead to thought, actions and consciousness*
- we have no idea how single memories are stored
- computers switch faster, brains process faster

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

foundations of AI: psychology

- how do humans and animals think and act?

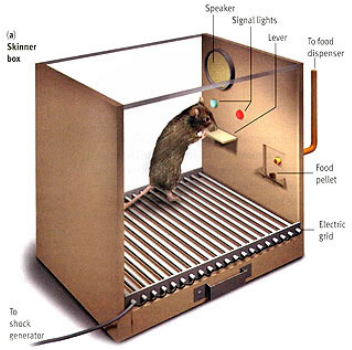

- experimental techniques
- notion of brain as information processing device
- computer models can address psychology of memory, language, learning, and reasoning (*cognitive science*)

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

foundations of AI: computer engineering

- how can we build an efficient computer?

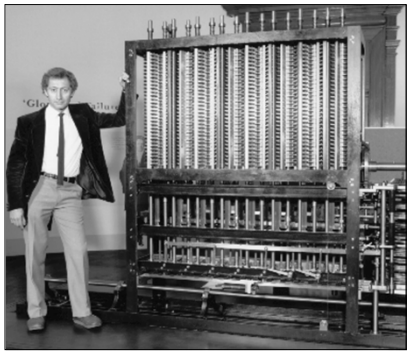

- programmable machines
- increasing speeds and memory
- operating systems and programming languages

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

foundations of AI: economics and linguistics

- decision theory:
  making decisions under uncertainty
- game theory:
  agent's actions may affect decisions of others
- operations research:
  making decisions now with benefit in the future

- language understanding
- knowledge representation
- grammar

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

a short history of AI

history of AI: the beginning

1943 McCulloch & Pitts developed on/off artificial neurons (*perceptron*)

1950 Alan Turing's "*Computing Machinery and Intelligence*"
- Turing Test
- Turing's vision of AI

1956 Dartmouth College meeting of the pioneers
- e.g. Marvin Minsky, John McCarthy, Claude Shannon
- name "*Artificial Intelligence*" adopted

history of AI: 1952-1969, great expectations

- early AI programs
- *General Problem Solver*, logical reasoning
- LISP
- microworlds, *Blocks World*

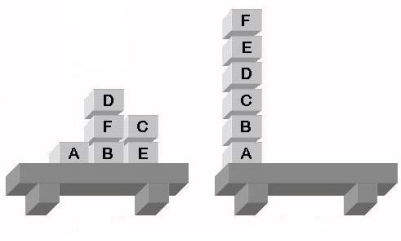

"[…] *there are now in the world machines that can think, that learn and that create. Moreover, […] in a visible future, the range of problems that they can handle will be coextensive with the range to which the human mind has been applied."*
(AI pioneer and Nobel laureate Herbert Simon, 1957)

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

history of AI: 1952-1969, great expectations

- faith that science would make a better future was *enormous*

- *"Just machines that make big decisions, programmed by fellas with compassion and vision."*

  from *I.G.Y.* by Donald Fagen about the International Geophysical Year 1957-1958

Research Seminar: Artificial Intelligence: *context and history of AI*
Maarten Lamers, Media Technology MSc programme, Leiden University

history of AI: 1952-1969, great expectations

- "A physical symbol system has the necessary and sufficient means for general intelligent action."
  – *physical symbol system hypothesis*
  – pioneers Alan Newell and Herbert Simon

necessary → human thinking is symbol manipulation
sufficient → machines can be intelligent

- John Searle's *Chinese Room* argument:
  "a program cannot be said to *understand* the symbols"
- Rodney Brooks' *Nouvelle AI* approach:
  "the world is its own best model"

history of AI: 1966-1973, reality sucks

- early AI systems failed miserably when applied to wider range of problems
  - programs are limited by lack of subject knowledge
  - outside microworlds, computational complexity became far too large (*combinatorial explosion*)
  - fundamental limitations of the used techniques
    e.g. Minsky & Papert prove perceptron limitations;
    neural networks research almost disappears

history of AI: knowledge based systems

- 1969-1979
- use subject-specific knowledge, so that larger reasoning steps can be made
- large number of special-purpose rules, instead of few general reasoning rules
- the rules come from experts, hence *expert systems*

- DENDRAL is first knowledge-intensive system (identifies unknown organic molecules)
- MYCIN expert system (to diagnose blood infections) can also reason under uncertainty

history of AI: growing up

1980-now	AI becomes an industry, mainly based on expert systems
1988	expert systems hype is over
1986-now	return of neural networks as result of *back-propagation* learning algorithm
1995	neural networks hype is over
1987-now	AI becomes a science: - building on existing theories - experimental evidence for hypotheses - real-world problems, instead of toy problems

history of AI: 1995-now, intelligent agents

- return to “whole agent” problem, after work on sub-problems of AI
- internet becomes environment for intelligent agents:
  - search engines
  - recommender systems
  - construction tools

amazon.com

history of AI: state of the art (*in 2003*)

- IBM’s *Deep Blue* defeated the world chess champion Garry Kasparov in 1997; IBM stock increased $18 billion

- in 1996, a program proved *Robbins conjecture*, unsolved for decades by mathematicians
- program solves crossword puzzles (dutch: *cryptogrammen*) better than most humans
- NASA's on-board autonomous planning program controlled the scheduling of operations 100 million miles from earth

history of AI: state of the art (*in 2003*)

- medical diagnosis programs perform at expert physician level
- computer-vision system steers minivan across USA, driving autonomously 98% of the time

- computer vision techniques and robotic assistants in microsurgery are common medical applications
- in 1991 Gulf War, logistics of 50.000 US vehicles, cargo and people was done by planning and scheduling program in few hours; which normally takes weeks

more recently

- In february 2011, IBM's "*Watson*" artificial intelligence system won the game-show "*Jeopardy!*" from the world's two best Jeopardy-players, twice.

  Watson processed the verbal question via natural language processing.

  (there's plenty of video's of this online)

history of AI: hypes

- history of AI has seen many hypes

  machine translation, “general problem solvers”, “artificial minds”, expert systems, neural networks, etcetera

→ this lead to bad reputation for AI →*AI winters*

- *“what is the current hype...?”*

follow-up reading

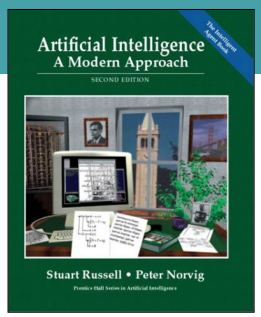

- most popular textbook:

  Stuart Russell and Peter Norvig,
  “A*rtificial Intelligence, a Modern Approach*”
  (second edition), 2003

- Chapter 1 of Rob Callan, “*Artificial Intelligence*”, 2003
  see course webpage for online resource

- interesting references on course webpage
- many, many more online resources available

context and history of artificial intelligence

questions?

context and history of artificial intelligence

see you next week,
don't forget the homework!