Introduction to Artificial Intelligence

(c) Marcin Sydow

What is AI? (AI - Artificial Intelligence)

What is the goal of AI?

What is AI? (AI - Artificial Intelligence)

What is the goal of AI?

To create machines/algorithms that can “think”.

What is AI? (AI - Artificial Intelligence)

What is the goal of AI?

To create machines/algorithms that can "think".

A Vision present in culture for ages and in modern literature:

- myths and legends (e.g. automatons in Ancient Egipt, China, Greece (e.g. Talos), Europe (e.g. Golem), etc.)
- science-fiction literature (e.g. "Cyberiada" and many other books by Stanisław Lem and innumerable other SF writers, "Alien", "Star wars", "Blade runner", and innumerable other SF movies, etc.)

AI today

Today it is a rapidly growing interdisciplinary field of research and diverse practical applications with numerous successes.

The contemporary chapter of this history since 60's (first computers)

"The science and engineering of making intelligent machines" (John McCarthy 1955)

Topisc of this course

- Introduction
- Elements of Machine Learning
  - Basics of ML, taxonomy of ML
  - Perceptron and Neural Networks
  - Knowledge representation
  - Overfitting and evaluation
  - Naive Bayes Classifier and other classifiers and regressors
  - Clustering
- Elements of discrete optimisation
  - discrete optimisation problems
  - brute-force method
  - greedy algorithms
  - local search heuristics
  - other approaches (e.g. genetic algorithm, etc.)

What are the features of intelligence?

What are the features of intelligence?

- learning (on examples)

What are the features of intelligence?

- learning (on examples)
- solving complex problems

What are the features of intelligence?

- learning (on examples)
- solving complex problems
- adaptation (to dynamic situation)

What are the features of intelligence?

- learning (on examples)
- solving complex problems
- adaptation (to dynamic situation)
- reasoning (based on knowledge and rules)

What are the features of intelligence?

- learning (on examples)
- solving complex problems
- adaptation (to dynamic situation)
- reasoning (based on knowledge and rules)
- perception (vision, hearing)

What are the features of intelligence?

- learning (on examples)
- solving complex problems
- adaptation (to dynamic situation)
- reasoning (based on knowledge and rules)
- perception (vision, hearing)
- knowledge (representation)

What are the features of intelligence?

- learning (on examples)
- solving complex problems
- adaptation (to dynamic situation)
- reasoning (based on knowledge and rules)
- perception (vision, hearing)
- knowledge (representation)
- generalisation (of observed cases)

What are the features of intelligence?

- learning (on examples)
- solving complex problems
- adaptation (to dynamic situation)
- reasoning (based on knowledge and rules)
- perception (vision, hearing)
- knowledge (representation)
- generalisation (of observed cases)
- communication (language)

What are the features of intelligence?

- learning (on examples)
- solving complex problems
- adaptation (to dynamic situation)
- reasoning (based on knowledge and rules)
- perception (vision, hearing)
- knowledge (representation)
- generalisation (of observed cases)
- communication (language)
- planning
- and other ...

Wide spectrum: sensing-reasoning-acting

- perception (artificial "senses": computer vision, speech recognition)
- **knowledge** (knowledge representation: rules, decision tables, decision trees, ontologies)
- reasoning (logics, automated proving)
- **learning** (machine learning (ML): supervised (classification, regression), unsupervised (clustering))
- communication (natural language processing (NLP): e.g. information retrieval, text mining, query answering, machine translation, automatic knowledge acquisition)
- **task solving & planning** (searching, heuristics, multi-agent systems, cooperation, competition, evolution, swarm intelligence)
- motion and object manipulation (robotics)

AI is a very large field now

AI has numerous inter-related fields, e.g.:

- computational intelligence
- evolutionary and other bio-inspired optimisation heuristics
- knowledge acquisition and management
- pattern recognition
- machine learning
- NLP (natural language processing)
- deep learning
- reinforcement learning

Machine Learning

It is an approach to solve problems that are hard to precisely specify by making that computers/algorithms:

- learn on examples (e.g. data)
- apply automatically learnt (hierarchy of) concepts to "understand" the domain of solved problem

Such approach makes it possible to avoid the problem of precisely specifying the task to be solved (e.g. how to recognize a face?, etc.)

AI is interdisciplinary

- computer science
- mathematics
- linguistics
- philosophy
- (neuro)psychology
- robotics
- biology

Tools

- logics
- probability calculus
- optimisation
- economics and game theory, etc.

Other aspects:

- philosophical (Can machine *really* think?, etc.)
- theoretical (limitations based on computational theory, Goedel's theorem, etc.)
- ethical (Is AI development only advantageous for humans? Can it harm? Can it be dangerous for our civilisation or humankind? May be it already partially is?)

Divisions by:

- approaches and tools
- problems solved, applications

Strategies to approach AI

1. symbolic (logic, representation)
2. computational (intensively searching the solution space)
3. data-centric (data, data science, statistics and probability)

Examples of applications of AI

- proving mathematical theorems
- playing chess or other game
- solving a puzzle
- finding a way out of a maze
- medical diagnostics, early cancer detection, etc.
- weather prediction
- grouping similar objects
- recognising voice commands
- face identification (photo or video)
- understanding natural languate (translation, summarisation, querying, etc.)
- autonomous Mars mission robot

OK, but when can we say a machine is intelligent?

For example when it is capable of:

- playing and winning a game (checkers, chess, etc.)
- quickly solving a complex puzzle
- predicting weather based on observing atmospheric conditions
- autonomously moving in a hard terrain (desert, city, etc.)
- recognising a human face or emotions
- proving a mathematical theorem, etc.

Turing Test

(How to verify whether a machine/algorithm is intelligent?)

Turing Test:

- A i B communicates whith each other in natural language (text)
- one of them is human
- the other is a machine pretending to be a human
- C observes the communication
- can C figure out who is a human and which is a machine?

(Turing Test concerns only some aspects of AI)

No system passed Turing Test (so far)

Strong and Weak AI

- weak AI (narrow aspects, particular problems, etc.)
  - ability to solve particular complex problems
  - knowledge representation
  - adaptivity
  - learning
  - reasoning
- strong AI (general intelligence)
  all the above plus:
  - consciousness
  - creativity
  - awareness of self-boundedness
  - evolution

Strong AI is still not achieved (and one can ask what could be a reason to create strong AI)

We focus on weak AI here.

How did contemporary AI started

In the beginning phase, AI was capable of solving some problems that were difficult for human, but easy for computers, e.g. chess playing, and other tasks based on a set of simple and precise rules.

The real challenge was how to make computers solve problems that are easy and intuitive for humans (even children) but hard to precisely specify for computers algorithms (e.g. face recognition, etc.)

Beginnings of contemporary AI

among others, Alan Turing:

- theory of computations
- Turing machine (a programmable bit-manipulating machine capable of universal computations)

Beginnings of contemporary AI

1956: Darthmouht College Conference: (among others) John McCarthy, Marvin Minsky, Allen Newell, Arthur Samuel, and Herbert Simon...

They started creating programs that:

- won checkers with people
- proved mathematical theorems
- communicated in simple English

History: XX century

- 1958 perceptron (Rosenblatt)
- 1960's: DARPA funding
- 1969 Minsky published "Perceptrons" (several limitations of perceptrons were discovered)
- 1970's: "pessimism" ("AI winter")
- 1980's: "renaissance" (expert systems, decision support systems, backpropatation algorithm for neural networks, Hopfield networks, etc.
- 1990's: data mining, "intelligent" medical diagnostics, etc.

Brief History (subjective and simplistic view)

- pre-history (before ca. 1960 - first "modern" computers)
- romantism (60-65) - optimistic view that AI will reach human in 10 years...
- darkness (65-70) - pessimism
- renaissance (70-75) - first built practical expert systems that worked
- collaboration (75-80) - interdisciplinary research: natural sciences, theory, industry, linguistics
- commercialisation (80-)
- the second renaissance (2000-) - successes in deep learning, etc.

Examples of Success of AI

Examples of Success of AI

- 97 deep blue won chess with human master

Examples of Success of AI

- 97 deep blue won chess with human master
- 2005 DARPA grand challenge (131 miles on desert!)

Examples of Success of AI

- 97 deep blue won chess with human master
- 2005 DARPA grand challenge (131 miles on desert!)
- 2007 DARPA urban challenge (55 miles in city, recognising traffic lights, road signs, pedestrians, etc.!)

Examples of Success of AI

- 97 deep blue won chess with human master
- 2005 DARPA grand challenge (131 miles on desert!)
- 2007 DARPA urban challenge (55 miles in city, recognising traffic lights, road signs, pedestrians, etc.!)
- 2011 "IBM Watson System" wins on-line TV-quiz "Jeopardy!"
- ...

It is everywhere...

It is everywhere...

- smartphones

It is everywhere...

- smartphones
- omnipresent surveillance cameras

It is everywhere...

- smartphones
- omnipresent surveillance cameras
- search engines

It is everywhere...

- smartphones
- omnipresent surveillance cameras
- search engines
- games

It is everywhere...

- smartphones
- omnipresent surveillance cameras
- search engines
- games
- intelligent cars

It is everywhere...

- smartphones
- omnipresent surveillance cameras
- search engines
- games
- intelligent cars
- intelligent buildings

It is everywhere...

- smartphones
- omnipresent surveillance cameras
- search engines
- games
- intelligent cars
- intelligent buildings
- intelligent cities

It is everywhere...

- smartphones
- omnipresent surveillance cameras
- search engines
- games
- intelligent cars
- intelligent buildings
- intelligent cities
- intelligent things (internet of things)

It is everywhere...

- smartphones
- omnipresent surveillance cameras
- search engines
- games
- intelligent cars
- intelligent buildings
- intelligent cities
- intelligent things (internet of things)

(where is the limit of this process ?)

It is everywhere...

- smartphones
- omnipresent surveillance cameras
- search engines
- games
- intelligent cars
- intelligent buildings
- intelligent cities
- intelligent things (internet of things)

(where is the limit of this process ?)

More importantly:

It is everywhere...

- smartphones
- omnipresent surveillance cameras
- search engines
- games
- intelligent cars
- intelligent buildings
- intelligent cities
- intelligent things (internet of things)

(where is the limit of this process ?)

More importantly:

Is the uncontrolled development of AI obviously good for us?

Literature

Neural Networks (e.g.):
M. Negnevitsky "Artificial Intelligence"

Machine Learning (e.g.):
P. Cichosz "Systemy Uczące się"
Witten et al. "Data Mining"

General AI (e.g.):
G. Luger "Artificial Intelligence"

NP-completess (e.g.):
Cormen et al. "Introduction to algorithms"

Optimisation (e.g.):
C.Papadimitriou "Combinatorial Optimisation"

Approximation algorithms:
V.Vasirani "Approximation algorithms"

Complexity:
C.Papadimitriou "Complexity Theory"

Questions/Problems:

- list the key aspects of intelligence
- Turing's test
- strong AI vs weak AI
- short history of AI
- list 3 different modern applications of AI
- positive and negative aspects of development of AI

Thank you for attention