11. Visualization of data & artificial intelligence

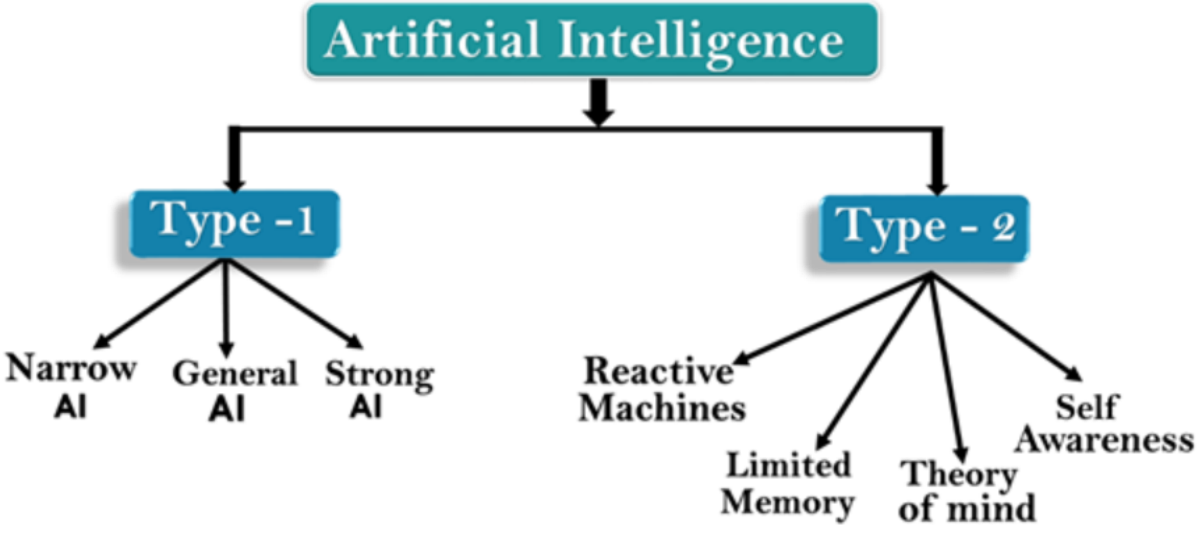

Revision: Data vs. information

- "*Data*" is just numbers or text without context
- When we add context to data, it becomes *information*
- Correctly interpreted information enables advances in *knowledge*

- Are we visualizing data or information?
  - The nature of this question is more philosophical, because the computer doesn't understand the difference
- We might start with visualization of data – which then becomes visualization of information, when context is added
  - Removal of outliers
  - Axis labels, headers, scaling

Types & dimensions of data

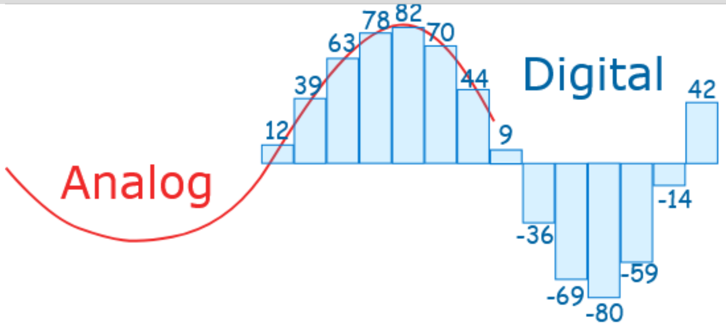

- Our data can be numerical or symbolic
- Numeric data signal can be analog or digital
  - Analog is continuous by nature (number data, but recording must be done by discretization)
  - Digital is discrete by nature, consists of 1's and 0's
  - Signals differ in susceptibility to disturbances (analog is prone to noise)
- Data has different dimensions – based on what we're recording:
  - 1-dimensional (single variable – for example, voltage measurement)
  - 2-dimensional (x- & y-coordinates – for example, black and white photo)
  - 3-dimensional (for example, a color image – RGB)
  - 4-dimensional (3D + time; for example, a color video)
  - n-dimensional

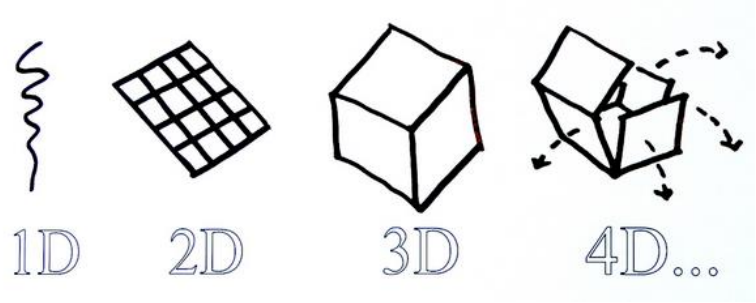

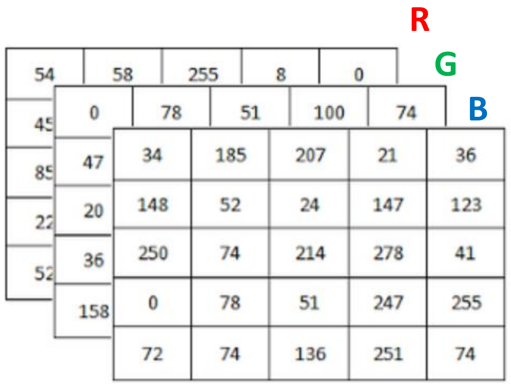

Information processing or calculation?

- Terminology-wise, what are we actually doing?
- The borderlines of scientific disciplines are not clear
- One take at the subject based on data type & methods of problem-solving:
  - Note: this is a simplified view and has been contested

	Algorithmic	Heuristic
Numeric data	Technical mathematics calculations	Simulation & signal processing
Symbolic data	Economical/governmental information processing	Artificial intelligence & knowledge engineering

Importance of visualization

- Measurements provide us lots of data, usually in numeric form
- It is very hard to see trends from hundreds of rows of numbers/text
- When data is plotted to a graph, relationships between variables become easier to see
- Essential for formulation of hypotheses

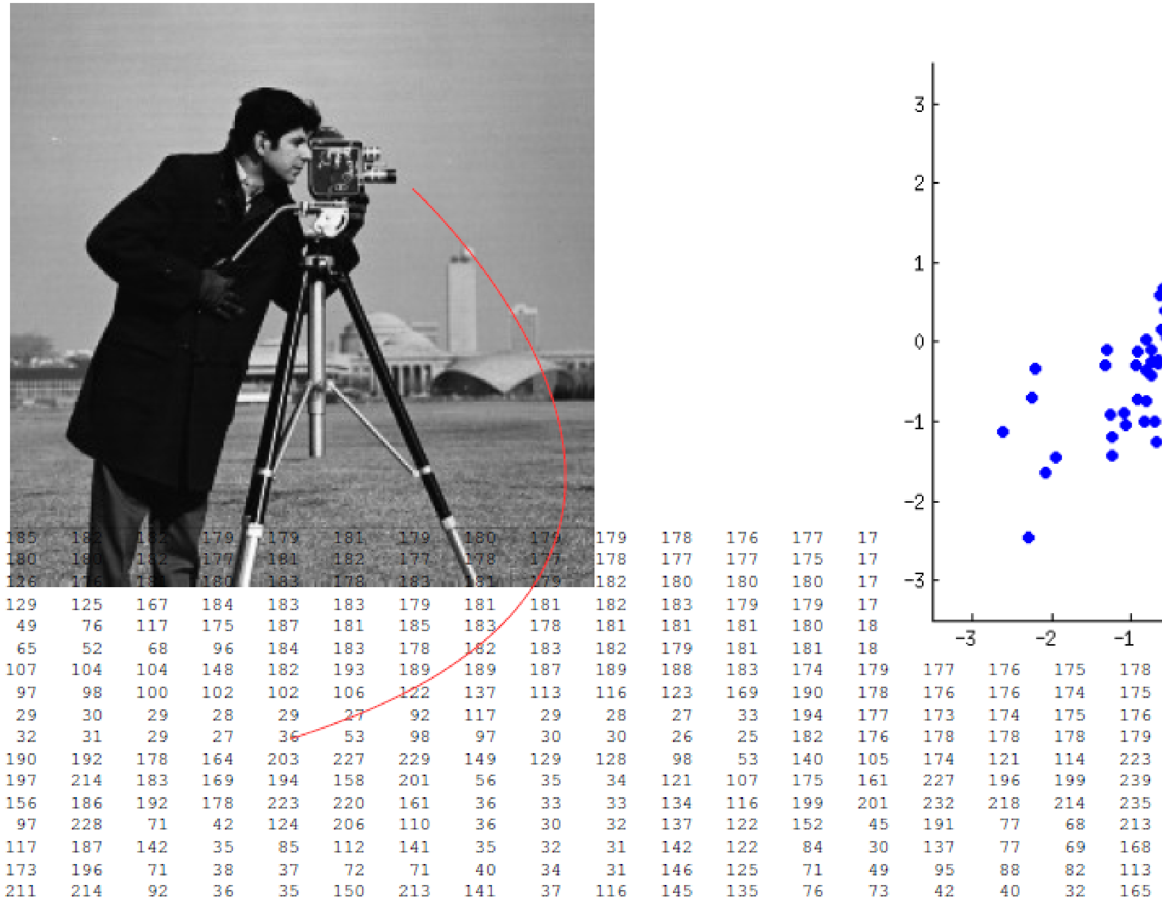

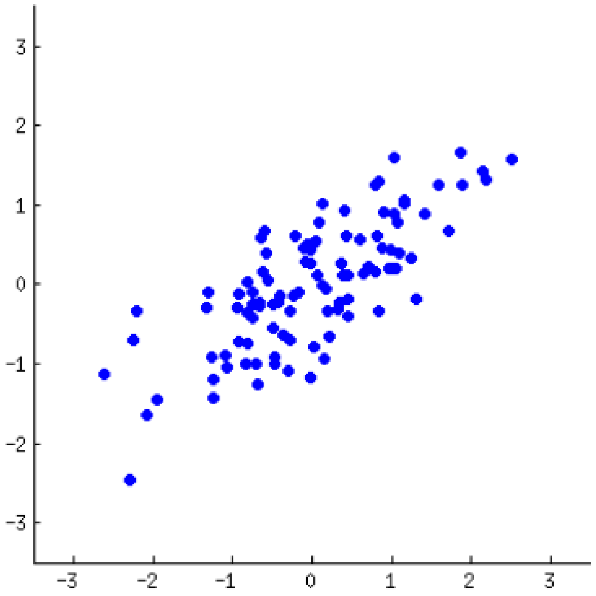

n-dimensional data

- If we conduct large-scale measurements for some machine/phenomenon, we’ll usually have multiple sensors in various places and record dozens of things
- As a result, do we have dozens of one-dimensional data streams or one n-dimensional data stream?
  - If all measurements are done at the same time, theoretically the latter – because the data allows us to study the relations between variables!
  - Difficulties in visualization of n-dimensional data often forces us to decrease the number of dimensions – at least at once
- This requires tools from the toolbox of information processing
  - Selection of essential variables for examination of the desired task using heuristics or specific calculation methods (for example principal component analysis, PCA)
  - Signal processing & linear algebra might be needed, too (data cleanup)

From data to geometry: production

- Computer graphics is all about how geometric shapes are formed from data
  - Group of data points are told to form an object – but this is nothing yet
  - Object can be given attributes, which specify the connection method of points & possibly some other features, too (like color, for example)
  - Visibility of the object is dependent on player's orientation & interaction with other objects

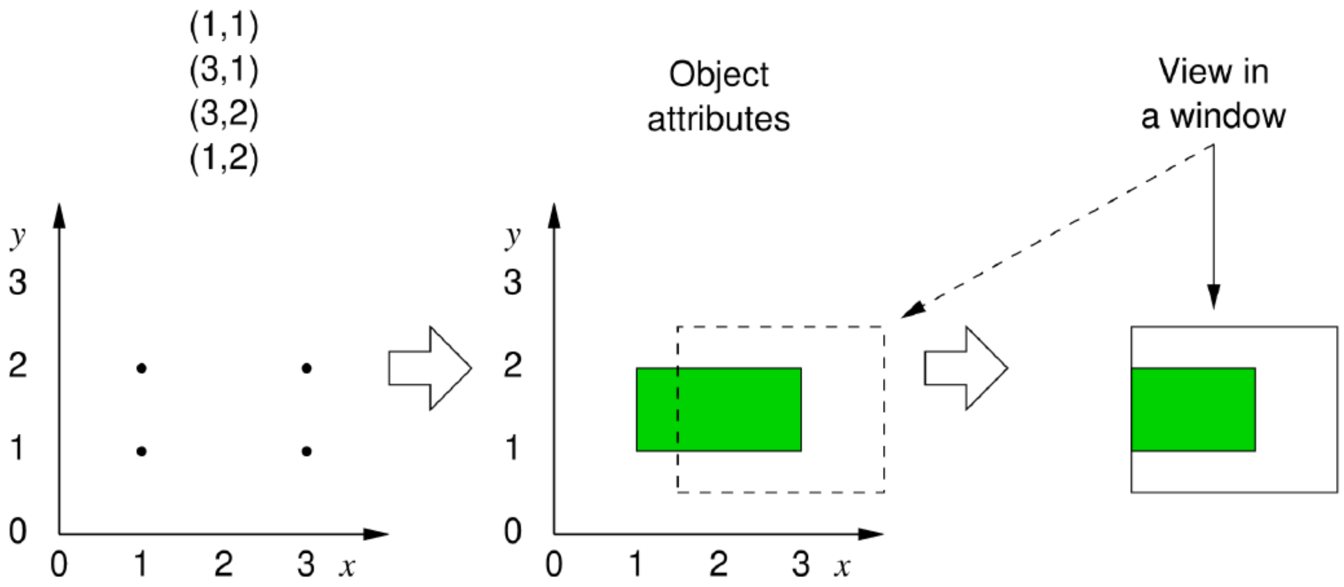

From data to geometry: reproduction

- If we want to reproduce a real-life geometry on a computer, we have to be able to measure it properly
  - 3D laser-scanning is a good option, but expensive & work-heavy
- A cheaper alternative is to use a “regular” camera and photometric stereo
  - One camera, two pictures of the sample – one lighted from left, one from right
  - By combining the pictures, we can separate intensity and gradient
  - Surface shape can then be extracted by integrating the gradient

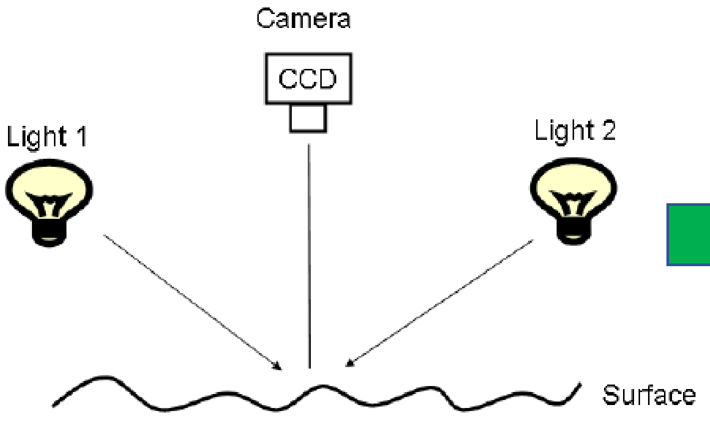

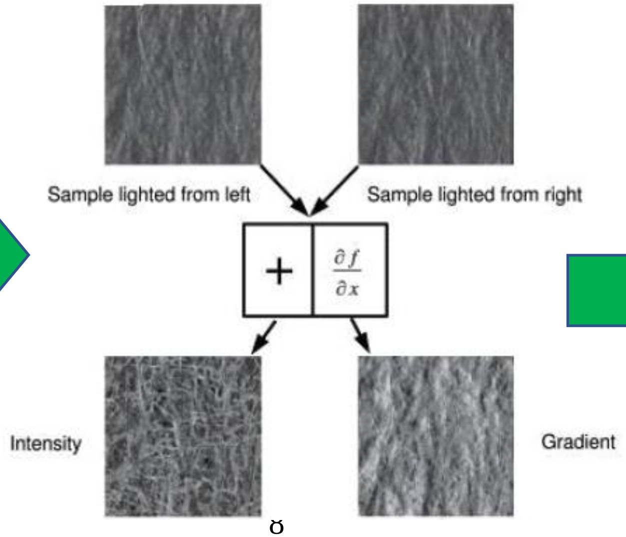

$$\int \frac{\partial f}{\partial x} dx$$

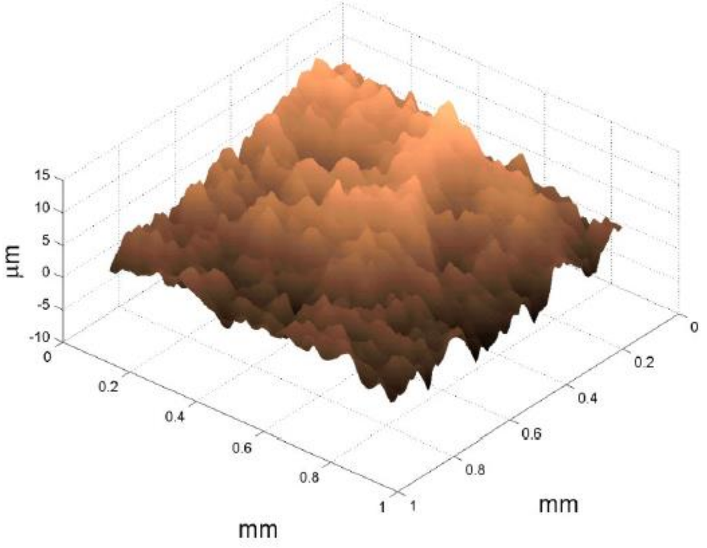

From data to geometry: reproduction

- More refined version of the previous approach is to use structured light
- Widely used in facial reconstruction
    - “Shape from shading” technique is used to convert one image to a 3D surface
    - Structured light is then projected at the 3D reconstruction
    - Produces a smoother result

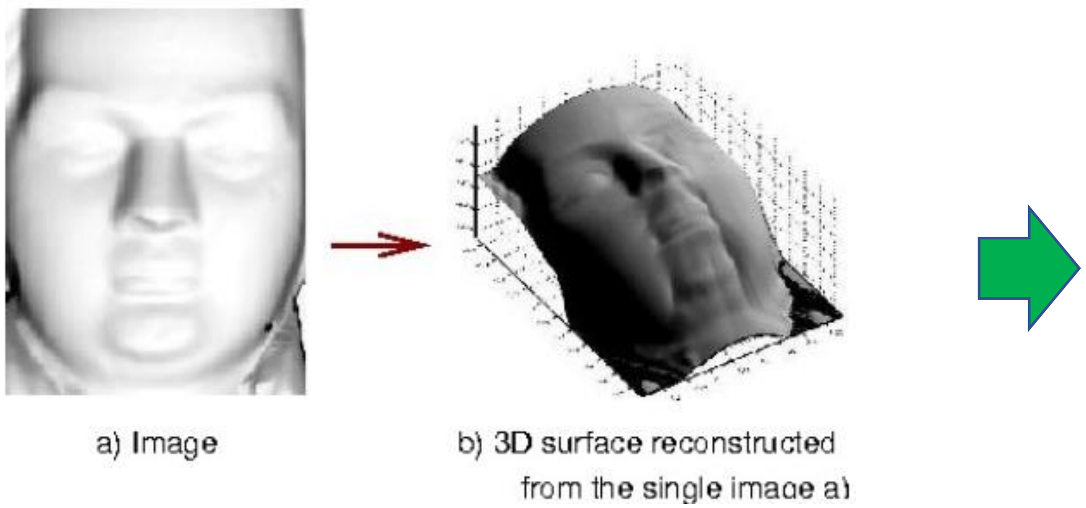

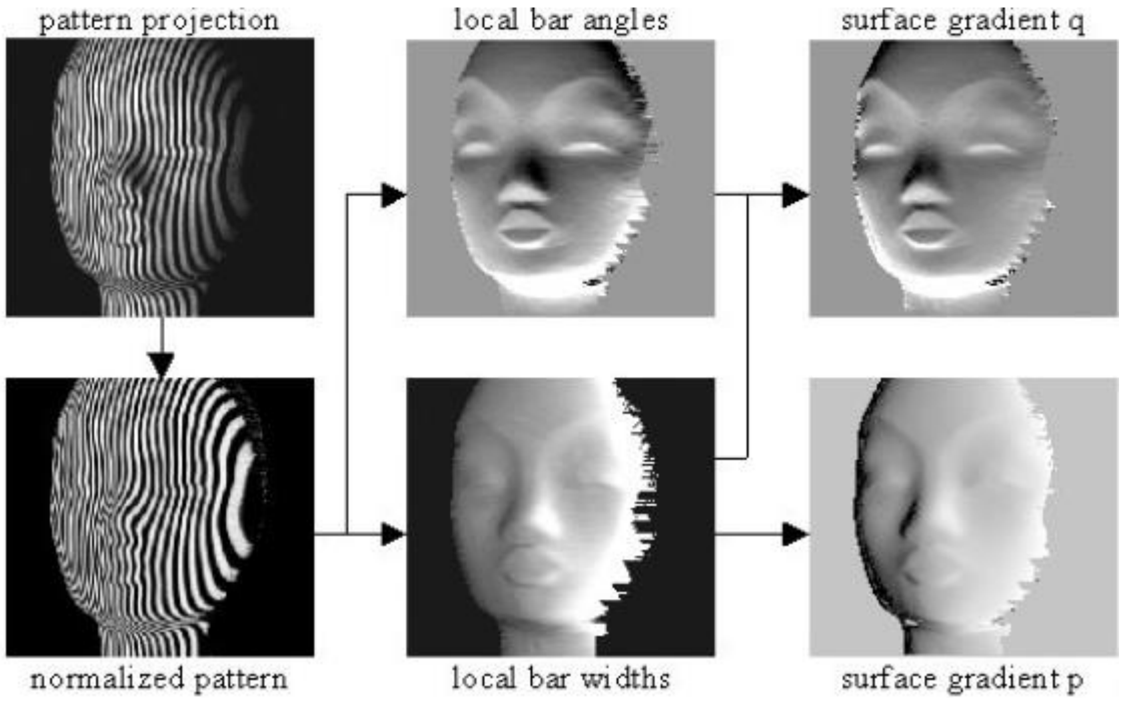

Power & limitations of the human mind

- Human mind is a very powerful problem-solving tool
- Its power lies mostly in its ability to create connections between different subjects
  - Think on multiple levels of abstraction
  - Formulate analogies and use them to advantage when learning new things
  - Create innovations by combining knowledge from multiple areas
  - Understand the context of things and draw conclusions from subtle hints
- On the other hand, human mind has many limitations
  - Low amount of “random-access memory” (especially)
  - Constant multitasking is problematic (interruptions are a bad thing here)
  - Variation in performance due to fatigue & boredom
- Conclusion: some problems are more suited for humans, some for computers
- In order to enlarge the latter group, we need artificial intelligence

What is intelligence?

- What can be classified as intelligence?
  - Reflexes vs. learned behavior models vs. innovative problem-solving
- How do we know if animals are intelligent?
  - “Blanket test”
  - Communication
  - Ability to speak
  - Learning from mistakes (trial and error)
  - Abstract thinking

- How do we know that humans are intelligent?
  - We’ve set the standards by ourselves, is that fair?
- How can we then define when a machine can be considered intelligent?
  - A computer can solve certain problems very quickly – when specified by user

Measurement of intelligence

- There are no “correct” answers to previous questions due to their philosophical nature; let’s take a more engineering-like approach and try to quantify
- Most well-known measure of intelligence is the intelligence quotient (IQ)
  - IQ score of 100 represents a median result, larger score means more intelligent
  - Distributions around the 100 differ by author (several scales)
  - …but what do we use as a population based on which we specify the median?
- Measurement of IQ is not 100% objective and has received criticism
  - IQ tests are based on logical thinking and finding patterns
  - A person can improve their IQ test score by rehearsing such problems; does this improve the person’s intelligence or just the test result?
  - Is logical thinking the only sign of intelligence, or can there be something else?
- Current understanding: there are 8 types of intelligence (logico-mathematical is just one of them)

LUT University

Concepts of artificial intelligence

- Artificial intelligence is a field of computer science that aspires to build autonomous machines that can carry out tasks without human intervention
- Alongside CS, also other fields of science are needed
  - Psychology (desired behavior)
  - Linguistics (communication)
  - Philosophy (ethics)

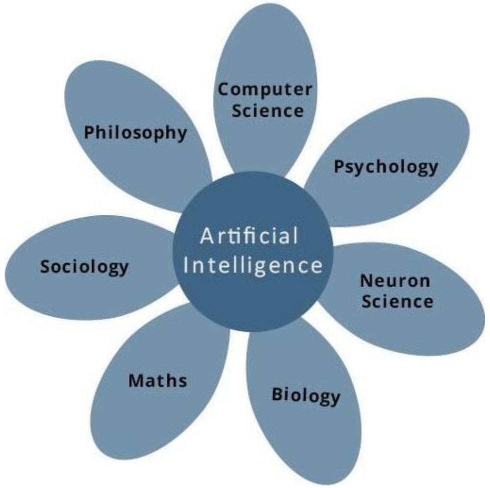

- Some general terms and classifications:
  - Weak AI: only developed for one single action
  - Strong AI: superior to humans (not created yet)
  - Symbolic AI ("GOFAI"): created using "traditional" methods, creator sets the rules
  - Connectionist AI: created using neural networks, AI learns the rules by itself using sample data

Goals of artificial intelligence

- AI research is trying to create machines that are capable to
  - Detect and measure (basis of cognition)
  - Process measurements and information (basis for learning, deduction and planning)
  - Learn (learning algorithms, computational intelligence)
  - Deduce (classification & problem-solving)
  - Plan (optimization, assessing quality of forthcoming actions)
  - Communicate (information retreaval & exchange)
  - Move and handle objects (robotics)
- The advances are pursued along two research methodologies:
  - Engineering approach, which is performance-oriented (”just make it work”)
  - Theoretical approach, which is simulation-oriented (computational understanding)

Turing test

- Turing test was proposed by Alan Turing in 1950
- Idea: a human interrogator communicates with two “test subjects” via a typewriter program for a period of time, and after this is asked to tell which test subject was a real human and which one a machine portraying a human
- What counts as a “passed” test?
  - Turing predicted a 30% pass rate (5min test) by year 2000
  - This was reached in 2014 by ”Eugene Goostman”
  - ”Eugene” was claimed to be a 13-year-old Ukrainian boy
  - Was this ”identity” a key factor in passing the test?
- Also criticism has been voiced towards the test
  - John Searle’s ”Chinese room”
  - Does the test tell anything about intelligence of the machine?

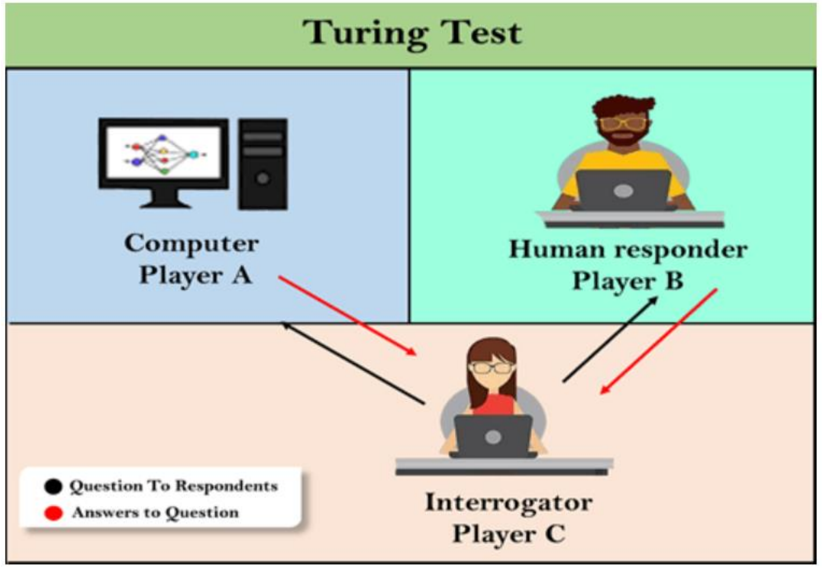

Methods for implementation of AI

- AI can be implemented using one or several of the following methods:
  - Machine vision
  - Search trees
  - Heuristics
  - Expert systems
  - Neural networks
  - Fuzzy logic
  - Fractals & chaos theory
  - Evolutionary computation
  - Swarm intelligence

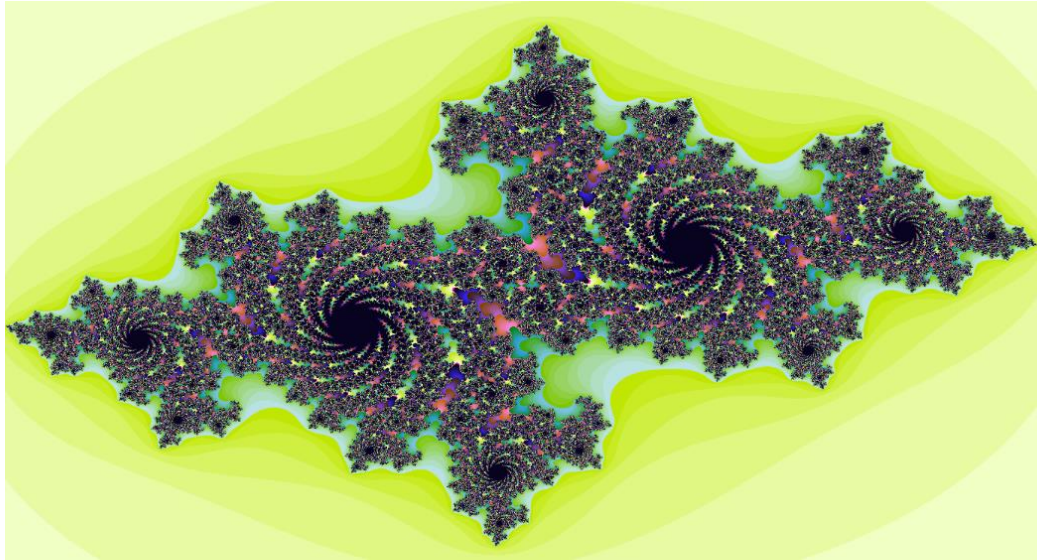

- In order to create a strong AI, it is mandatory to combine these methods

Intelligent agents

- Intelligent agent = a “device” that responds to stimuli from its environment
  - A robot / game character / self-driving car / chatbot / etc.
- Actions of the agent must be rational responses
- These actions are classified to different levels:
  - Reflex action: pre-programmed responses to specific inputs
  - Contexted action: responses depend on inputs AND the current environment
  - Goal-based action: agents’ responses are results of following a plan to seek a goal
  - Utility-based action: agent is able to measure what option most likely leads to the goal
- Agent is able to learn if its responses improve over time; this can be done by
  - Developing procedural knowledge (via trial and error)
  - Storing declarative knowledge (new principes added to the bank)

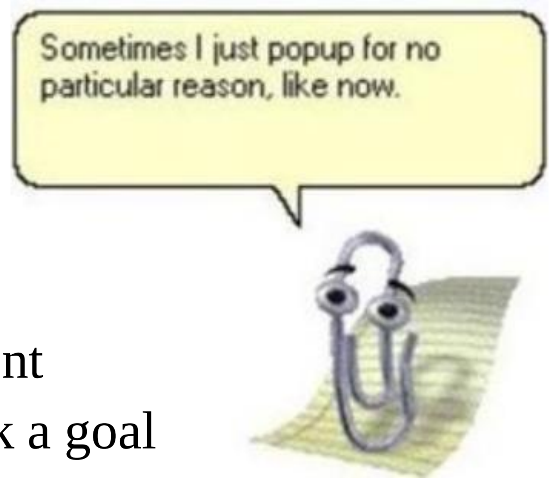

Symbolic AI example: eight-puzzle

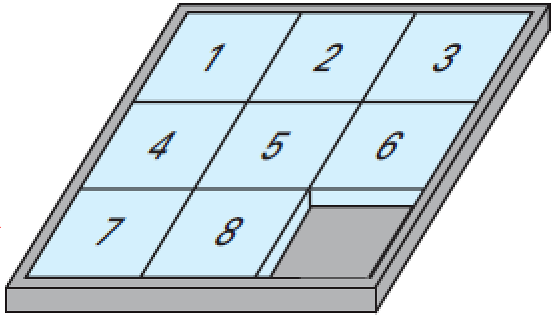

- Suppose we try to get AI to solve us a puzzle that has 3x3 tile matrix with 8 numbered tiles and one free slot
  - "Solved" state presented in the picture
- First, we need some kind of a machine that can both sense the current state of the puzzle as well as make moves
  - One possible option for such a machine pictured here

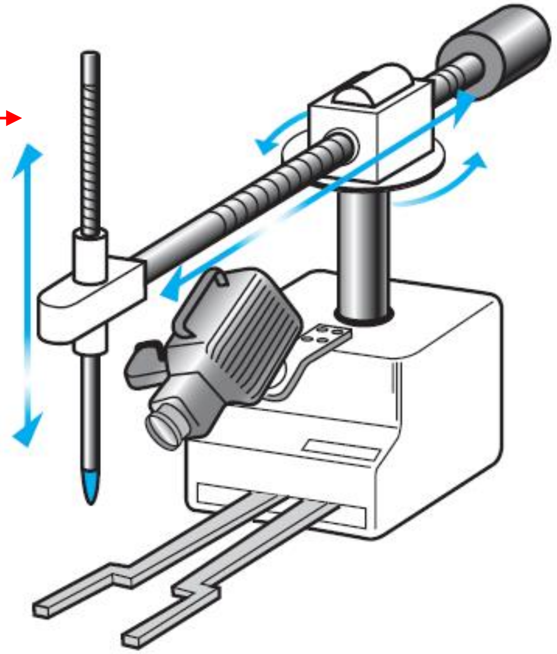

- Extracting the positions of tiles is easy, because the geometry is so simple:
  - Image processing = identify geometric features (tiles, numbers)
  - Image analysis = identify what these features mean
- …but what if our puzzle numbers use different font?

Symbolic AI example: eight-puzzle

- When the current state has been found out, it’s time to make moves
- For this, we need a production system that has 3 main components:
  - Collection of states (start state and goal state as the most important ones)
  - Collection of productions (possible movements of tiles)
  - Control system (decision-making on which tile to move)
- We can search for a solution by using a breadth-first search tree:
  - From start state, branch to all possible follow-up states
  - Continue this branching until one branch reaches the goal state
- Downside of breadth-first method: if we’re far away from the solution, the search tree grows to immense size quite quickly
- Alternative: depth-first search tree
  - Explore each possibility until the end

Symbolic AI example: eight-puzzle

- Depth-first search can be improved if we can formulate a fitness function and then use this as a heuristic to decide which would be the best move
- In this example, two options come to mind:
  a) Sum of tiles that are not in their correct place
  b) Sum of distances of tiles from their desired location
- Both of these are easy to calculate
- Let's use b) as a heuristic now
- In each node, we select the branch that has the smallest fitness value
  - If this ends up in a worse situation, we back off

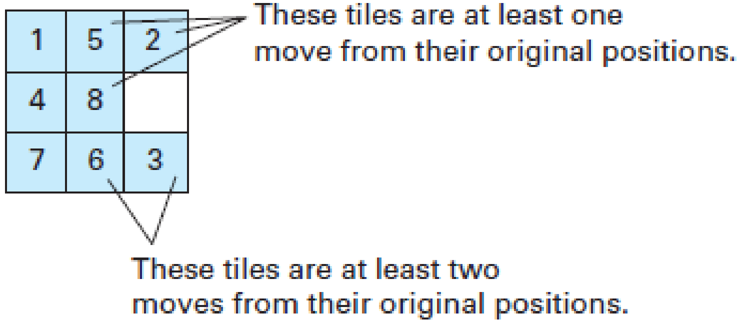

Symbolic AI example: eight-puzzle

- Comparison of search trees

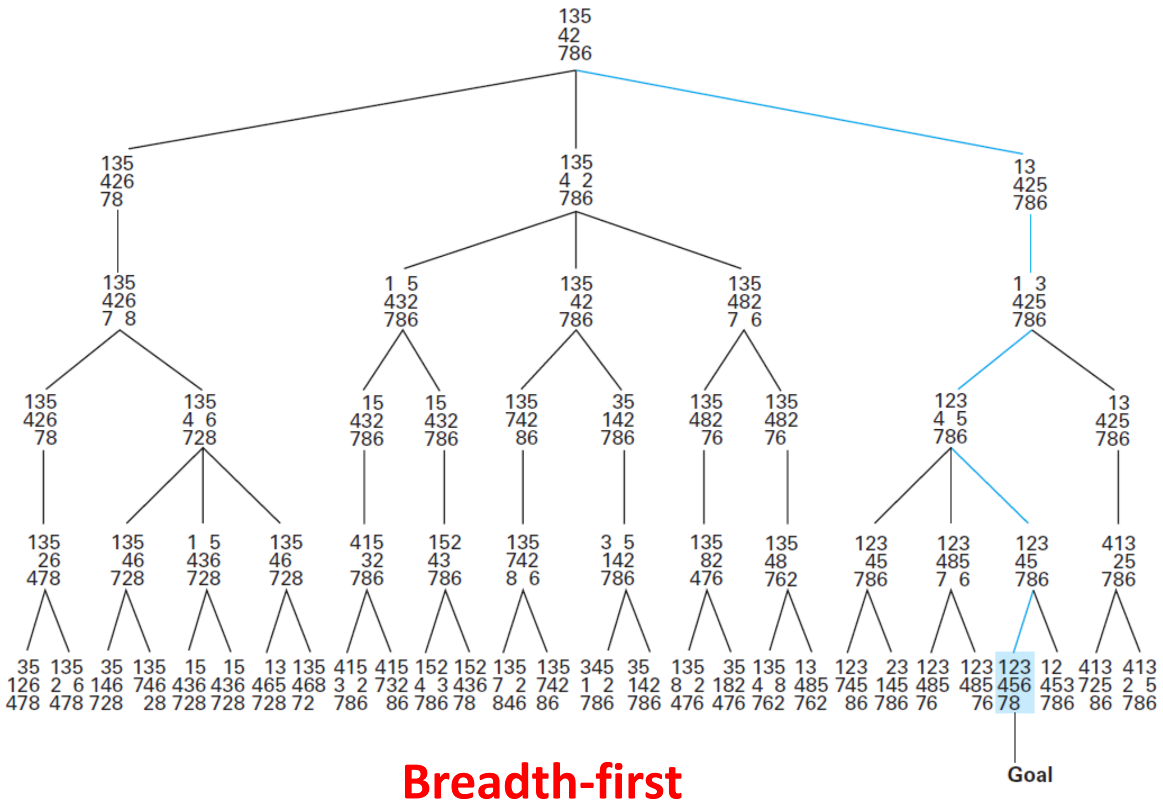

Breadth-first

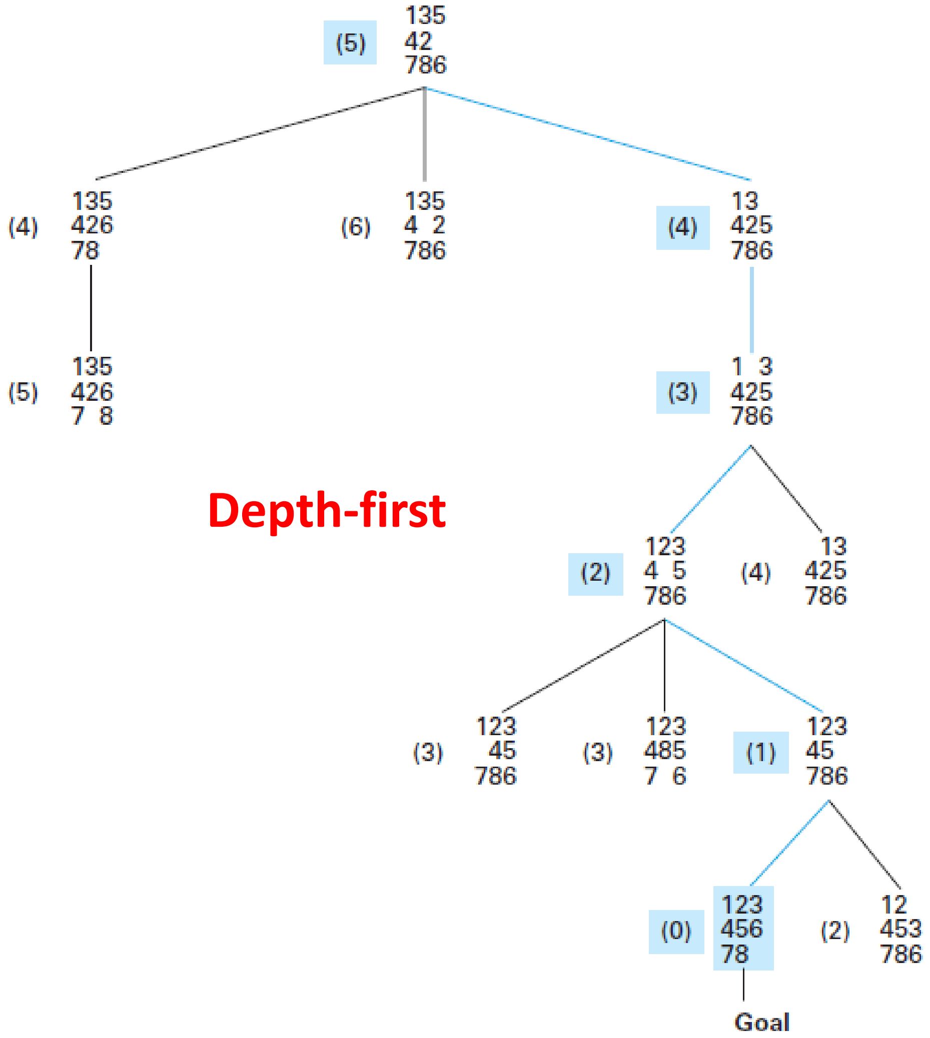

Depth-first

Expert systems

- Computer programs which aim to provide help for decision-making
- Experts put together a system that can be used by non-experts
- Expert system consists of three main parts:
  - Knowledge base = big storage of knowledge put together by experts (note: knowledge can be factual or heuristic)
  - Inference engine = rules that are used to make deductions (note: can be deterministic or probabilistic; latter one takes into account the uncertainty)
  - User interface = program that helps the user to communicate with the program
- Rules are independent of each other
- System is quick and reliable
- Often programmed using Prolog

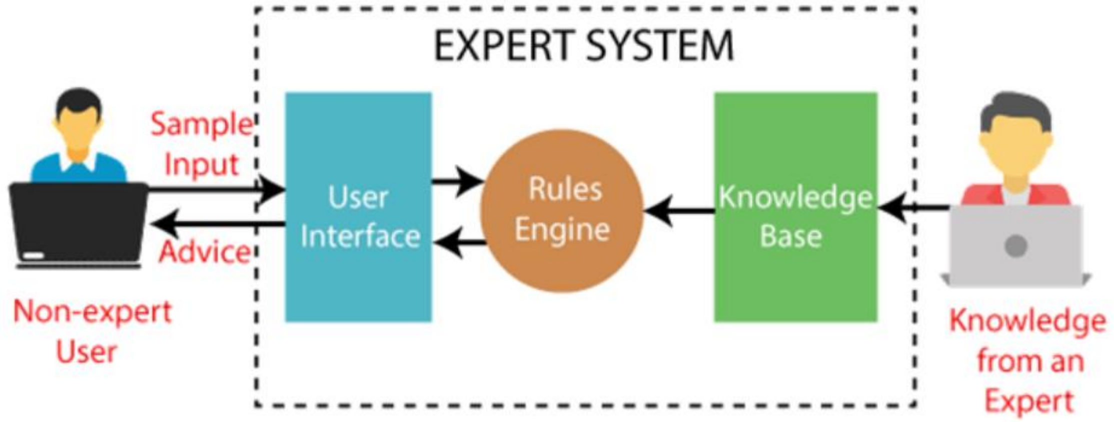

Learning

- One requirement for good AI is capability to learn new things
- Approaches to computer learning can be classified based on the level of human intervention needed:
  - Learning by imitation (human demonstrates, agent mimics)
  - Learning by supervised training (machine performs the action, human grades the success)
  - Learning by reinforcement (machine performs the actions and is able to assess the “goodness” of those actions by itself)
- Imitation is often used in computer programs (“did you mean: ___”)
- Supervised training is common, used for example in e-mail spam filters
- Reinforcement is a good method when the “goodness” can be calculated or otherwise assessed (game won vs. lost)

Neural network

- The most common (albeit not the only one) method of creating a connectionist AI are artificial neural networks
- This network is a model that mimics networks of neurons in biology
- Neurons are connected to each other via synapses
  - Dendrites = “input” tentacles
  - Axon = “output” tentacle
  - Synapses = small gaps between the two
- Neuron cell has two possible states
  - Excited state
  - Inhibited state
- State is transmitted to others via the axon

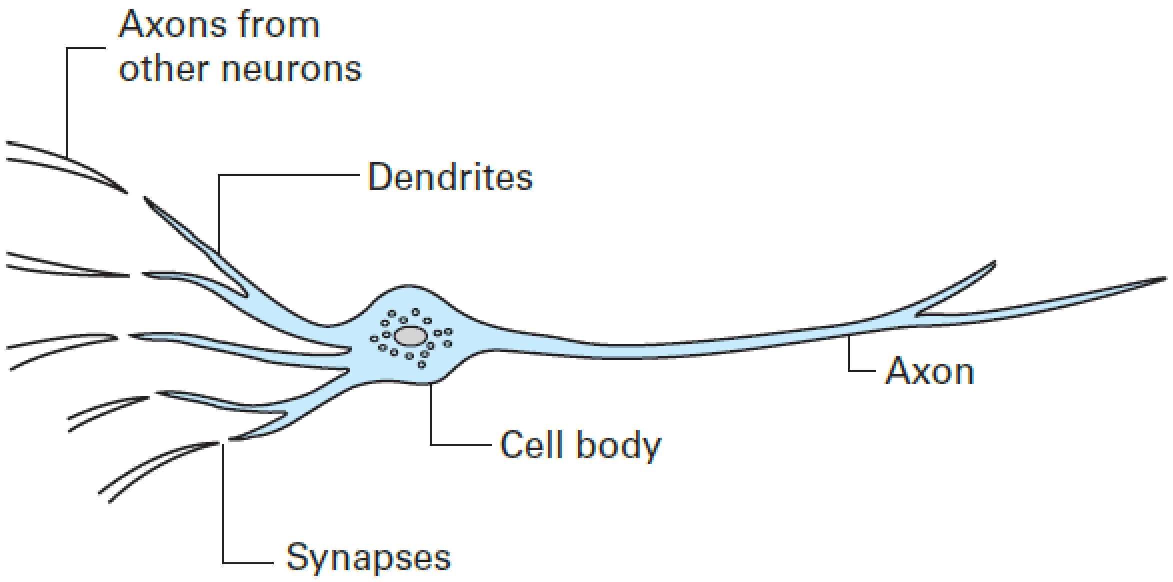

Perceptron

- A single-layer neural network is called a perceptron
  - Input values + bias are given weights
  - Perceptron calculates a weighted sum of these ("effective input")
  - Effective input is compared to the threshold value
  - If threshold value is exceeded, step function gives 1 as the output value (if not, the output value is 0)
  - Training is done by adjusting the weights (network adjusts!)

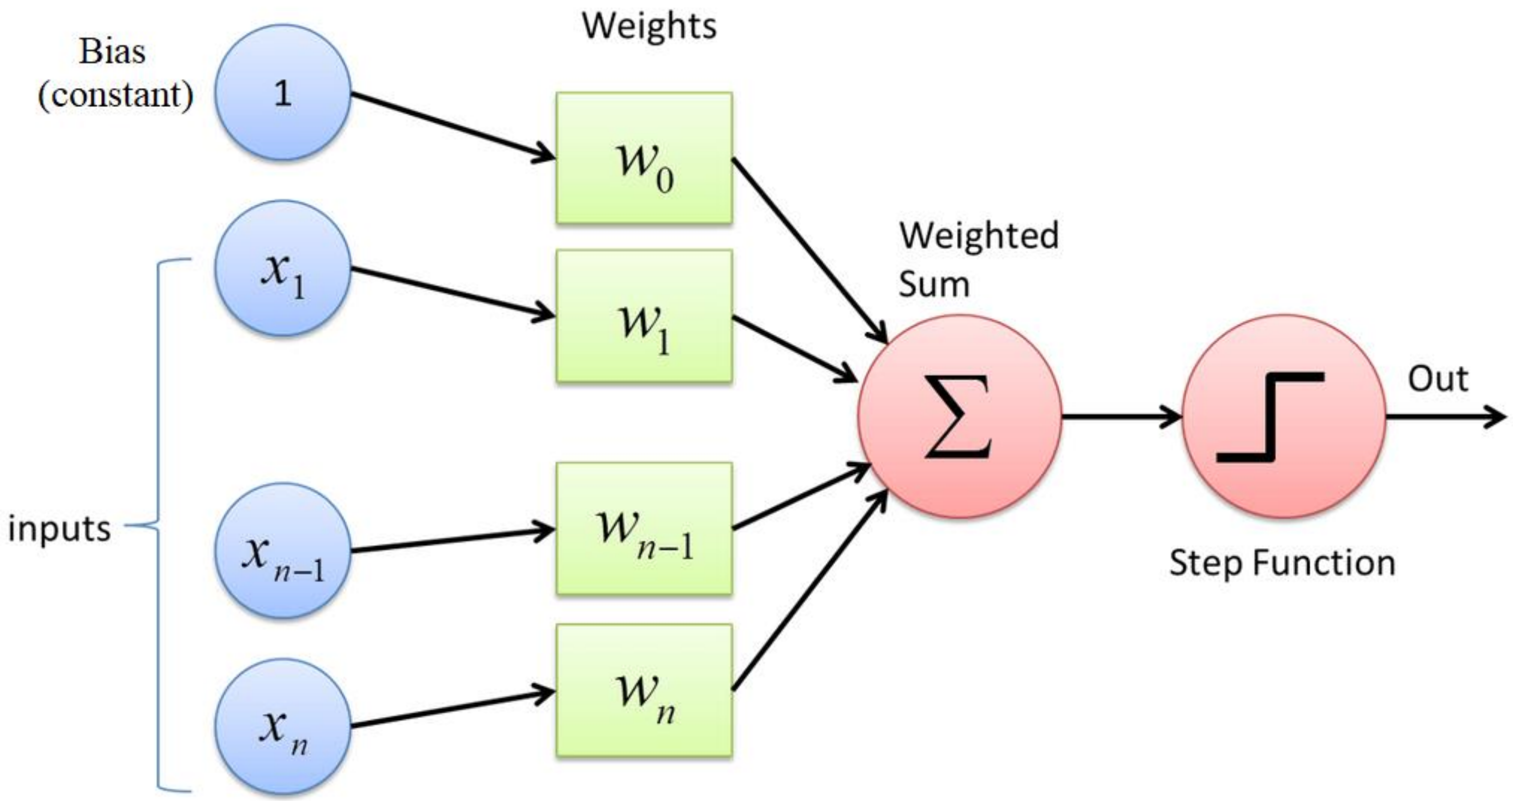

Multi-layer perceptron

- Human body is estimated to have approximately 1000 neurons, each with around 100 synapses – so there can be multiple layers
- A multi-layer perceptron is called neural networks

- Consists of
  - Input layer
  - Hidden layer (one or several)
  - Output layer
- Increasing the number of variables in the problem increases the need for more neurons and more layers
- When properly trained, is able to learn and model the behavior of any function

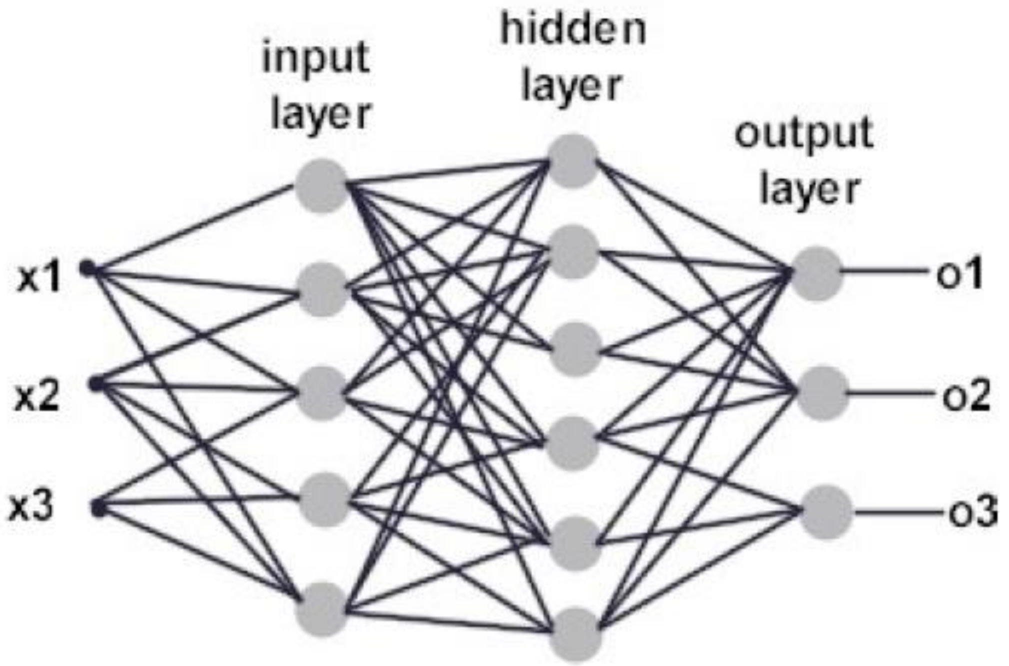

Fuzzy logic

- In fuzzy logic, truth table values for inputs may be any number in [0,1]
- Good for taking into account the impreciseness of information – and also the impreciseness of concepts!
- Example: what is classified as a “young” person?
  - Depends on who you’re asking from, but the transition age from young to older is 22…32
- Information: John is 28 years old
  - John is ”rather young”
- Useful tool when
  - Concepts are vague
  - Data is subjective

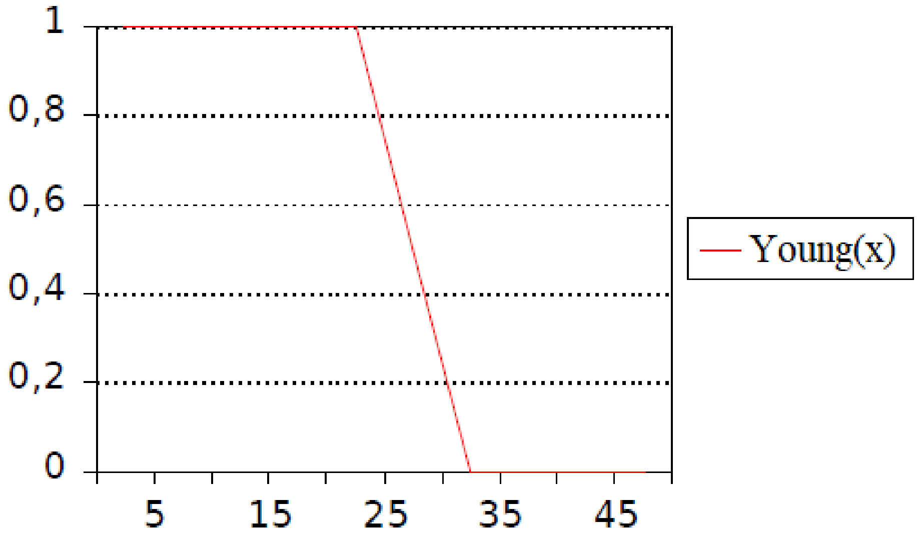

Asimov’s three laws of robotics

- Intelligent agents in AI are often considered as robots
- Here, ethics comes to play
    - What should robots be allowed to do and what they shouldn’t?
- Isaac Asimov (mostly known for his sci-fi books) presented the three laws of robotics in 1942
- The order of laws is important!
- Later, Asimov also added a 0th law “Don’t harm humanity or allow humanity to come to harm”

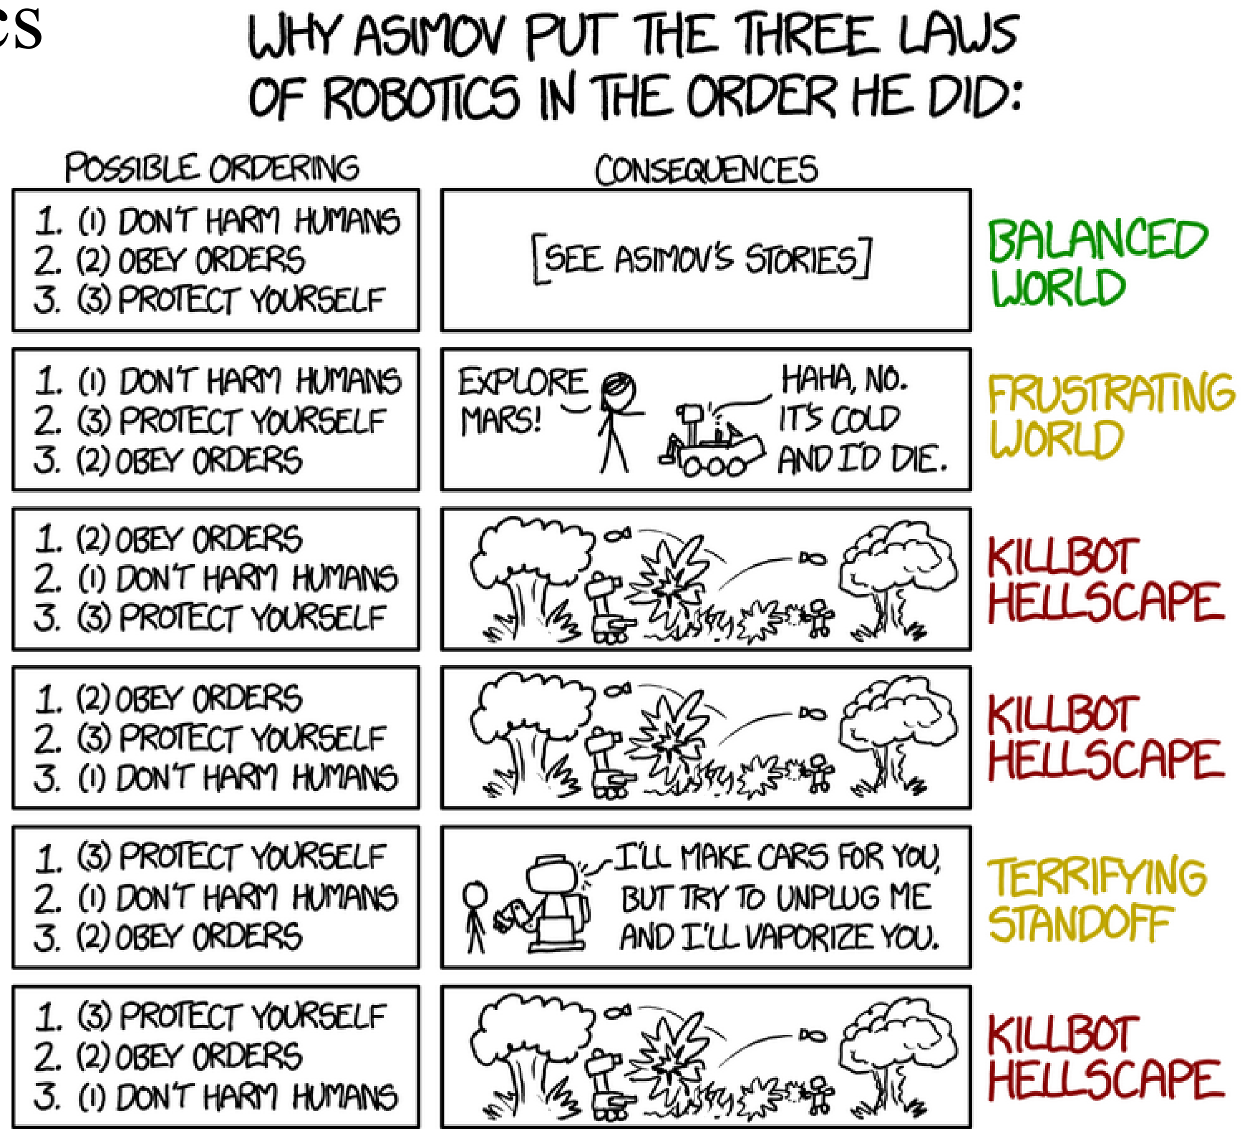

Challenges of AI

- Despite massive leaps during the last decades, AI is still miles away from being even close to the level of humans
- Many challenges exist – some technological, some not so much:
  - AI still needs a LOT of computing power
  - Private funding is not easy to get, because aside from large technology companies, implementation possibilities of AI are not well understood in the business sector
  - It's hard to get people excited about computers that can perform tasks which almost any human being can do better (especially a problem with computer vision)
  - Trust issues - hackers may find surprisingly simple ways to fool AI
  - Ethical and legislational questions about rights and responsibilities of agents need serious consideration
  - If strong AI can be developed, will science fiction become science fact?

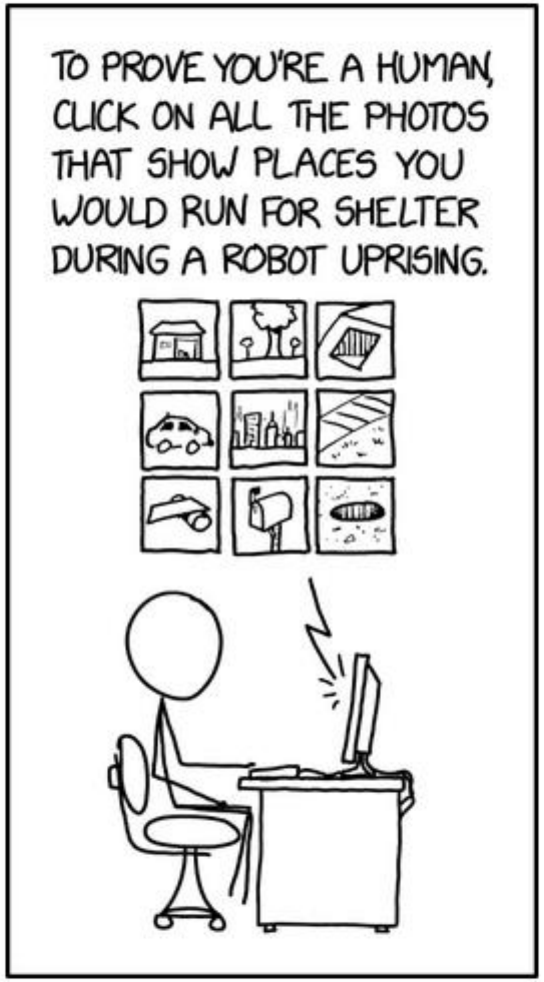

Thank you for listening!