# What is AI – and where is it heading?
# Part II: Symbolic and subsymbolic AI

Thomas Bolander, Professor, DTU Compute

*DigHumLab, 28 November 2019*

# What is subsymbolic AI?

- A great part of AI concerns making **predictions** (predicting the weather, customer behaviour, flu epidemics, heart-attacks, etc.).
- Using a computer to make predictions requires it to have a **mathematical model** of the phenomenon it's modelling.
- This model can either be **explicitly represented** (e.g. via formulas or rules) or **implicitly represented** (learned from experience with no symbolic representation of rules or properties).
- The former is **symbolic AI**, the latter is **subsymbolic AI**.

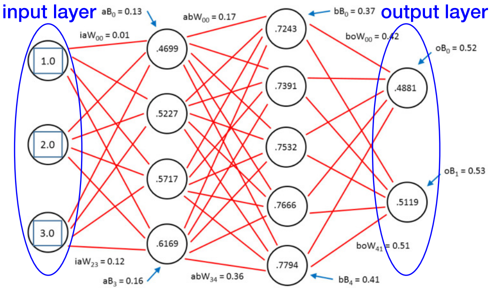


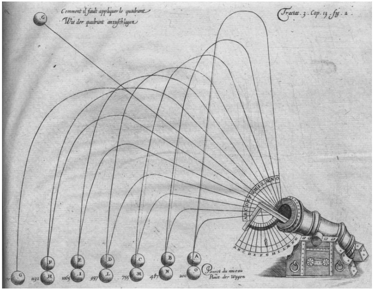

# Currently most prominent in subsymbolic AI: (Deep) neural networks

A neural networks produces an **output** (one or several numbers) from an **input** (one or several numbers).

Most often used for **classification**: Which class does the input belong to. E.g. is it a cat or dog (input is a picture), a good or bad customer (input is all data about customer), skin cancer or not (input is a picture). **Perception** rather than **higher-order cognition**.

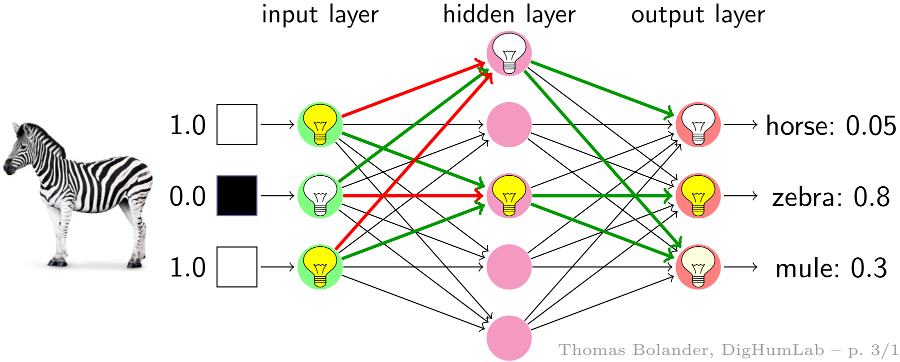

# Demo of deep neural networks

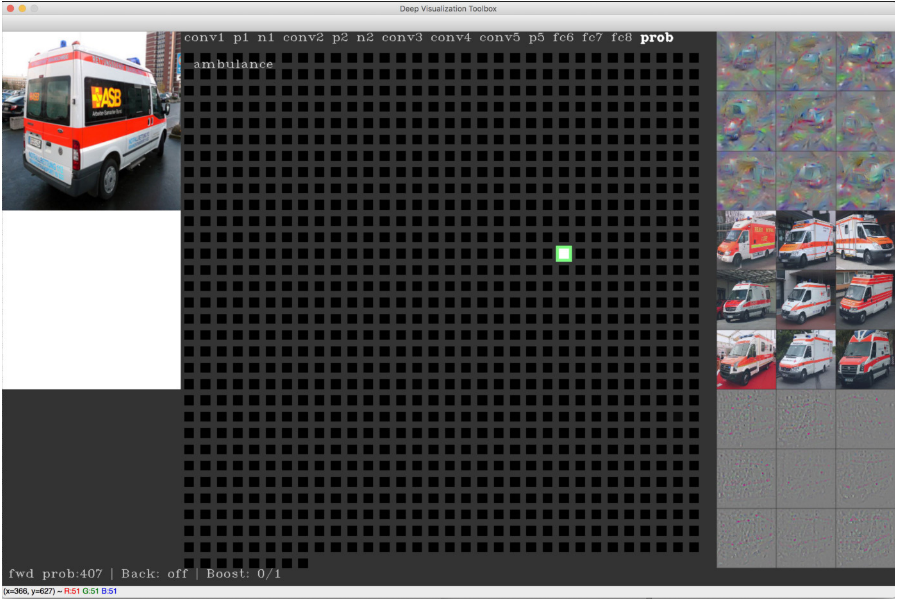

# Examples of successful Danish AI (deep neural networks)

corti


LEO Innovation Lab

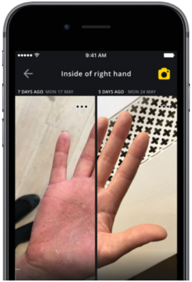

# The other side of the coin...

*For example, we cannot argue that a pedestrian detector is safe simply because it performs well on a large data set, because that data set may well omit important, but rare, phenomena (for example, people mounting bicycles). We wouldn't want our automated driver to run over a pedestrian who happened to do something unusual.*

(Russell & Norvig: Artificial Intelligence—A Modern Approach, 3ed, 2010.)


Tesla crash
June 2016

Uber Volvo crash
March 2018

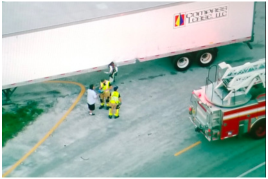

Tesla crash
March 2019

# Uber Volvo crash 2018

*Despite the fact that the car detected Herzberg with more than enough time to stop, it was traveling at 43.5 mph when it struck her and threw her 75 feet. When the car first detected her presence, 5.6 seconds before impact, it classified her as a vehicle. Then it changed its mind to "other," then to vehicle again, back to "other," then to bicycle, then to "other" again, and finally back to bicycle.*

(Marshall & Davies: Uber's Self-Driving Car Didn't Know Pedestrians Could Jaywalk, Wired, 5 November 2019)

# Challenges in (sub-symbolic) AI

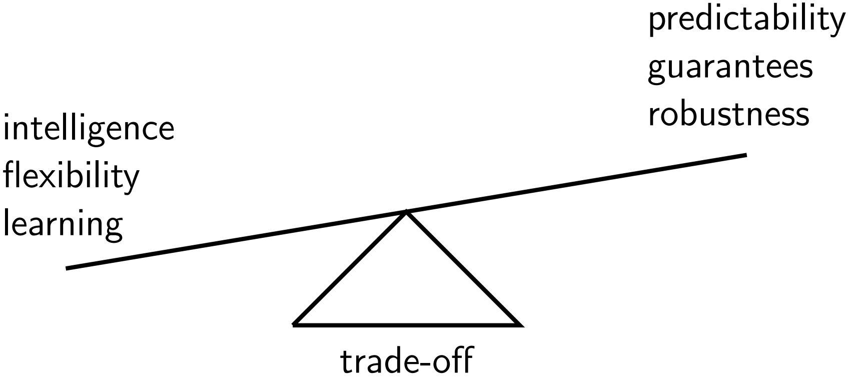

# Image filtering at Facebook and Instagram

Deep neural networks in practice:
**300 million pictures a day** are filtered through neural network image recognition algorithms at Facebook and Instagram. **Classification task**: appropriate/not appropriate.

Do you know, or can you guess, what the image recognition algorithms did with these pictures?

# The evolution of neural networks: scalability

- **1950s**: The first artificial neural networks.
- **1970s**: They develop essentially into their current form.
- **1980s**: Only small, simple tasks like recognising hand-written letters.
- **2018**: Industrial scale applications, better than humans at some pattern recognition tasks (e.g. recognising skin cancer).

So what happened since the 1980s:

1. Dramatic increase in **computer power**. Powerful GPUs.
2. Dramatic increase in **available data** (needed for training).
3. Small, but powerful, **modifications**: convolutional layers, dropout etc.

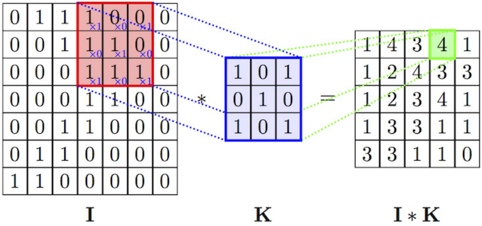

# What is symbolic AI?

- The computer's model of its environment is **symbolically represented**, e.g. via rules or formulas (can e.g. be logical formulas representing the world in terms of objects, relations and logical relationships).
- These rules and formulas are **often hardcoded**, but can also sometimes be **learned**.
- **Examples** of symbolic AI: **Siri on iPhone** (explicit models of possible user interactions) and **Google Maps** (explicit maps + rules for navigation + search).

$$a_x = \frac{-kv_x}{m} = \frac{dv_x}{dt} \text{ (1),}$$

and

$$a_y = \frac{1}{m}(-kv_y - mg) = \frac{-kv_y}{m} - g = \frac{dv_y}{dt} \text{ (2)}$$

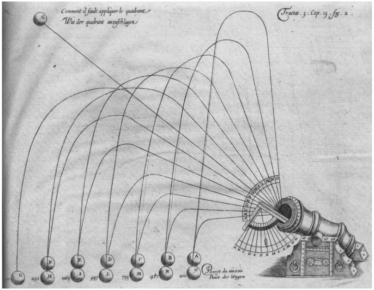

# Reinforcement Learning

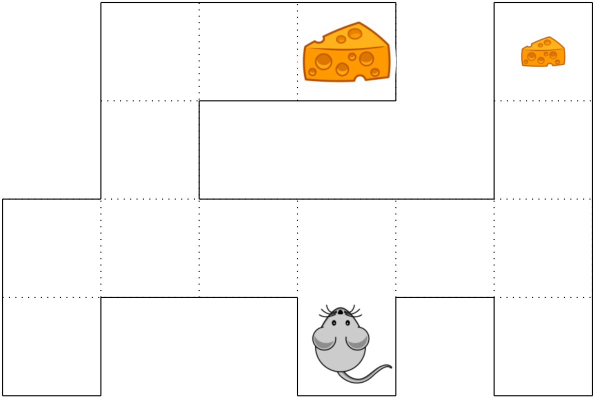

*Essentially, the animal makes more-or-less random movements and selects, in the sense that it subsequently repeats, those which produced the "desired result".*

Pringle, J. W. S. (1951)

# Learning of movements with reinforcement learning (Google DeepMind 2017)

http://www2.compute.dtu.dk/~tobo/deepmind_walking_nosound.mp4
(Nicolas Heess et al.: Emergence of Locomotion Behaviours in Rich Environments, Arxiv, 7 July 2017)

# Automated planning/problem-solving by searching

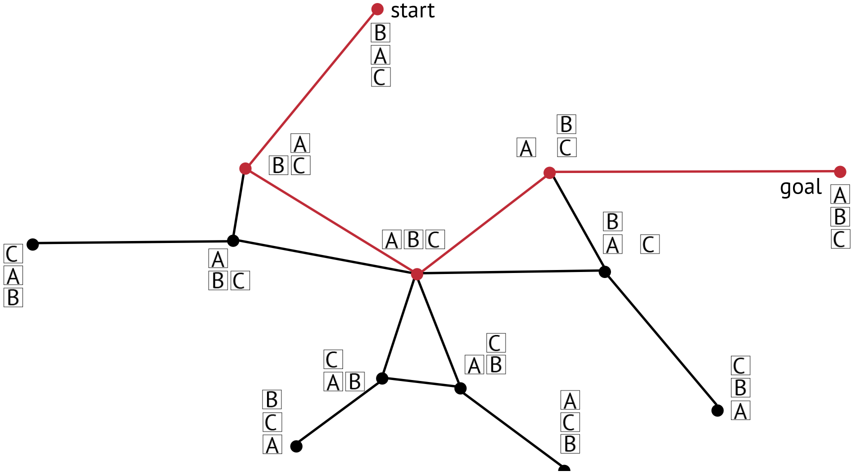

# Robots anno 1969 vs 2016

http://www2.compute.dtu.dk/~tobo/shakey_short_gremlin.mp4 Shakey the Robot, 1969

http://www2.compute.dtu.dk/~tobo/amazon_kiva.mp4 Amazon KIVA robots, 2016