Artificial Intelligence and Machine Learning

Artificial Intelligence

1950: Alan Turing's "Computing Machinery and Intelligence"

1955: McCarthy, Minsky, Rochester, and Shannon:

> *"We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."*

Artificial Intelligence

Definition:
Capability of a machine to imitate intelligent human behavior

Artificial Intelligence in the classical sense has by far not been realized (if ever), has not had any practical impact, and is therefore not studied in this course.

What is studied in this course?

Artificial Intelligence

Definition:
Capability of a machine to imitate intelligent human behavior

Artificial Intelligence in the classical sense has by far not been realized (if ever), has not had any practical impact, and is therefore not studied in this course.

What is studied in this course?
Augmented Intelligence (AI) or Intelligence Augmentation (IA)

Artificial Intelligence

Augmentated Intelligence is about empowering humans with tools that make them more capable, while traditional AI has been about removing humans fully from the loop.

Throughout this course, Artificial Intelligence = Augmentated Intelligence

AI has had enormous impact in the last few years

Many tools that make AI useful in practice are fairly recent

What made recent breakthrough in AI possible?
Combination of Big Data with advances in machine learning, fast algorithms, and computer power

Machine Learning

Machine Learning is a subfield within Artificial Intelligence that builds algorithms, which allow computers to learn to perform tasks from data instead of being explicitly programmed.

Supervised Learning: uses a known dataset (the training dataset) to make predictions. The training dataset includes input data and labeled responses. From it, the supervised learning algorithm seeks to build a model that can make predictions of the response values for a new dataset.
A test dataset is often used to validate the model.

Unsupervised Learning: is used to draw inferences from datasets consisting of input data without labeled responses.

Machine Learning

Main tasks of supervised learning:

- Regression
- Classification

Main tasks of unsupervised learning:

- Clustering
- Density estimation
- Dimension reduction

Regression

Learning the relationship between independent variables (predictors) and dependent variables.
Used for prediction and forecasting.

Estimation of housing price (say, based on location, number of bedrooms, ...), prediction of stock price, weather forecast, ...

Methods:

- Linear regression (least squares, ...)
- Kernel regression
- Gaussian process regression

Clustering

Finding natural groupings of data and a label associated with each of these groupings.

Marketing (consumer groups), Netflix, Amazon, text mining, image segmentation, ...

Methods:

- k-means
- Spectral clustering
- Hierarchical clustering

Dimension reduction

Reducing the number of variables under consideration

Data visualization, faster processing of data, reducing storage,

Methods:

- Principal component analysis
- Manifold learning
- Random projections
- Compressive sensing

Density estimation

Construction of an estimate of an unobservable underlying probability density function based on observed data. Finding likelihood or frequency of objects.

Finance (risk estimation), medical diagnostics, outlier detetion

Methods:

- Histograms
- Kernel density estimation
- Mixture of Gaussians

Classification

Organizing data into gategories, predicting a category of a data.

Does a person have a certain illness or not?
Classifying an image according to the objects in the image
Anomaly detection: detecting if a transaction is a fraud or not
Spam filtering, News vs Fake News

Methods:

- Support vector machines
- Random forests
- Deep Learning

Deep Learning

Deep Learning is one driving force behind recent success of AI

What is Deep Learning?

Deep Learning

Deep Learning is one driving force behind recent success of AI

What is Deep Learning?

Convolutional neural network with many layers

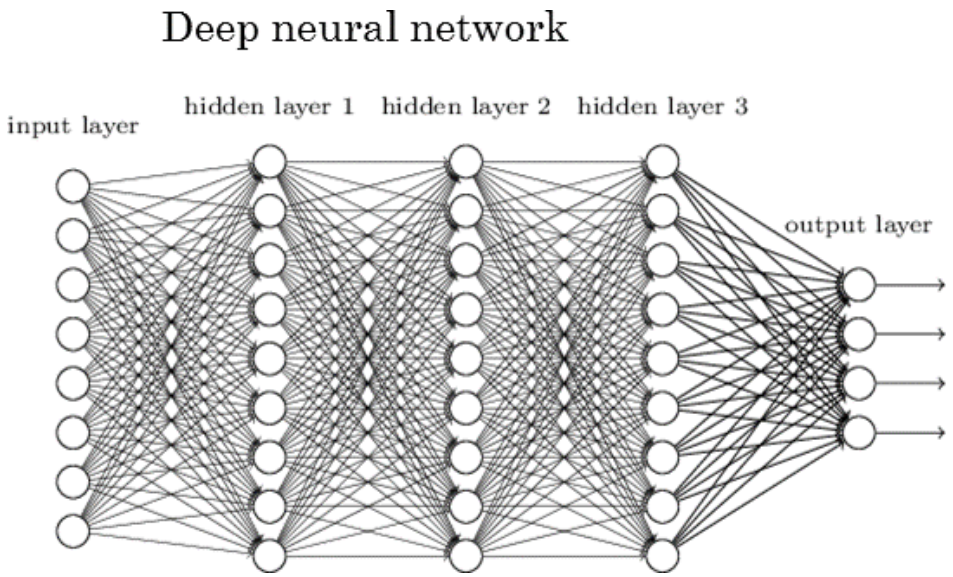

Applications of AI

DeepMind's AlphaGo beats a world class Go player

Image classification: Google, Facebook (DeepFace), ...

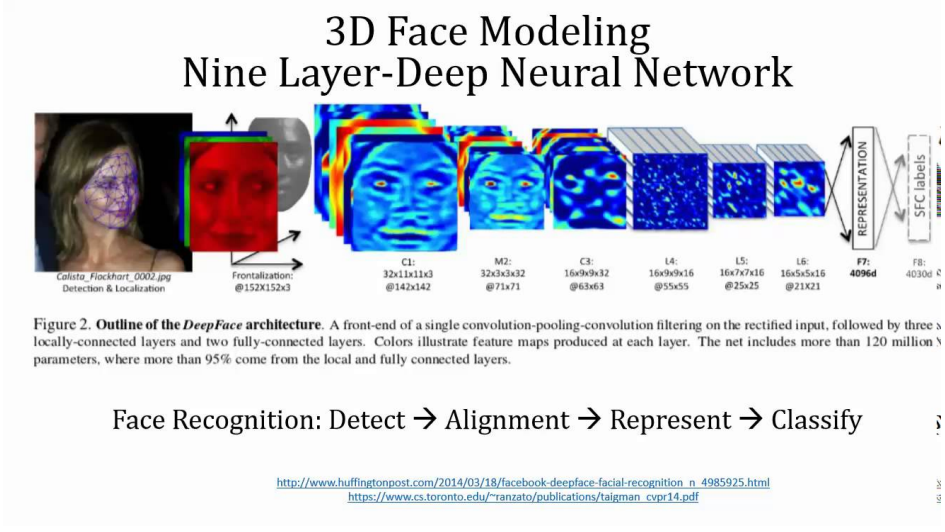

Figure 2. **Outline of the *DeepFace* architecture**. A front-end of a single convolution-pooling-convolution filtering on the rectified input, followed by three locally-connected layers and two fully-connected layers. Colors illustrate feature maps produced at each layer. The net includes more than 120 million parameters, where more than 95% come from the local and fully connected layers.

Face Recognition: Detect → Alignment → Represent → Classify

http://www.huffingtonpost.com/2014/03/18/facebook-deepface-facial-recognition_n_4985925.html
https://www.cs.toronto.edu/~ranzato/publications/taigman_cvpr14.pdf

Applications of AI

How can a computer learn concepts?

Currently AI heavily relies on vast amounts of training data.

Applications of AI

Self-driving cars

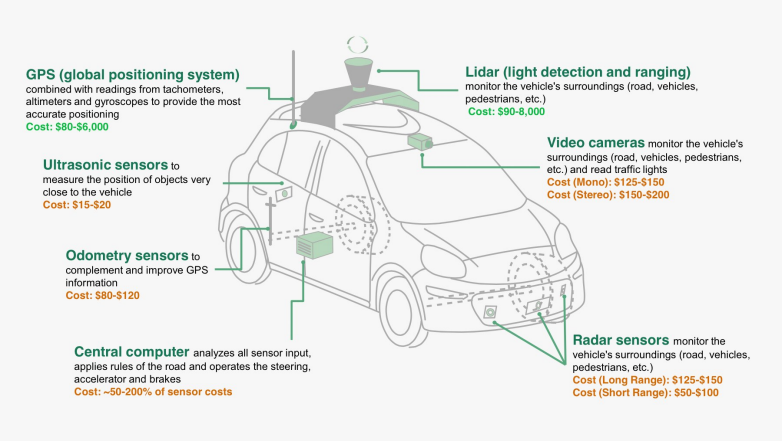

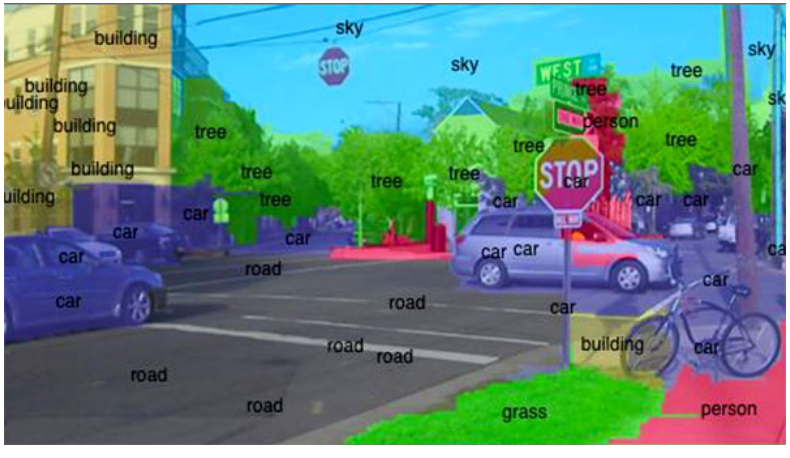

Uber Ends Self-Driving Car Test in San Francisco

Uber Suspends Tests of Self-Driving Vehicles After Arizona Crash

Applications of AI

Self-flying airplanes: easier than self-driving cars.

Airplanes use autopilots for many years

Airbus plans a self-flying car to be in operation by 2018

Applications of AI

1997 chess victory of IBM's Deep Blue over Garry Kasparov

2011 Jeopardy! victory of IBM's Watson system over two human contestants. Watson needed to be able to extract semantic meaning from the questions.

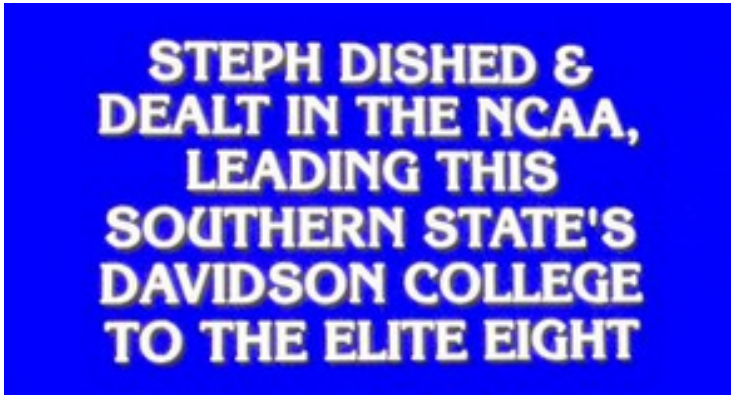

Speech recognition, automatic translation

Medical Diagnostics

- Machine Learning can help find patterns in large amounts of data to detect markers for diseases
- Assume we have a 3-dim. MRI of the brain:
  Doctors can look at 3-dim. MRIs only one slice at a time.
  They may miss patterns that can much better be detected by looking at the entire 3-dim. data simultaneously.
  Mathematical algorithms can easily analyze 3-dim. data
- The point is not to replace the doctor (as classical AI might attempt to do), but to assist the doctor with information that may be difficult to access - this is augmented intelligence.
- AI must be trustworthy. Will we develop trust as we interact with AI systems over time, as we have done with ATMs?

Challenges of AI

Surprise:

What is difficult for humans is “easy” for AI
(playing chess, detecting patterns in complex data, ...)

What is easy for humans is very difficult for AI
(moving around, language, common sense reasoning, ...)

What AI can or cannot do

- Play a decent game of table tennis?
- Play a decent game of Jeopardy?
- Drive safely along a curving mountain road?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at Whole Foods?
- Discover and prove a new mathematical theorem?
- Converse successfully with another person for an hour?
- Perform a surgical operation?
- Put away the dishes and fold the laundry?
- Translate spoken Chinese into spoken English in real time?
- Write an intentionally funny story?

What AI can or cannot do

- Play a decent game of table tennis?
- Play a decent game of Jeopardy?
- Drive safely along a curving mountain road?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at Whole Foods?
- Discover and prove a new mathematical theorem?
- Converse successfully with another person for an hour?
- Perform a surgical operation?
- Put away the dishes and fold the laundry?
- Translate spoken Chinese into spoken English in real time?
- Write an intentionally funny story?

Green = Yes, Red = No, Blue=?

Consequences of AI

AI will lead to economic disruptions

Consequence: many people will lose their jobs, social turmoil

Inequalities in society will increase much further

Changes may be at least as forceful as during the Industrial Revolution

Steps to reduce massive negative impact of AI need to be taken already now before AI fully kicks in