



Tekoälyn algoritmit ovat vain tyhmää perusmatematiikkaa

Erkki Oja

Previously: Aalto University

Department of Computer Science

AI and Machine Learning at Aalto

- Research started in the 1970's
- Several national Centers of Excellence in Research, continuously from 1995 till 2017
- First Master's Program in the field
- Strong ties with Helsinki University, the new Finnish Artificial Intelligence Center FCAI
- Other universities have strong groups, too.

Contents of my talk:

- What are Artificial Intelligence, Neural Network, and Deep Learning
- How and why does deep learning work
- Discriminative and Generative AI
- The future?



**What is AI, Machine Learning (ML),
and Neural Networks (NN)?**

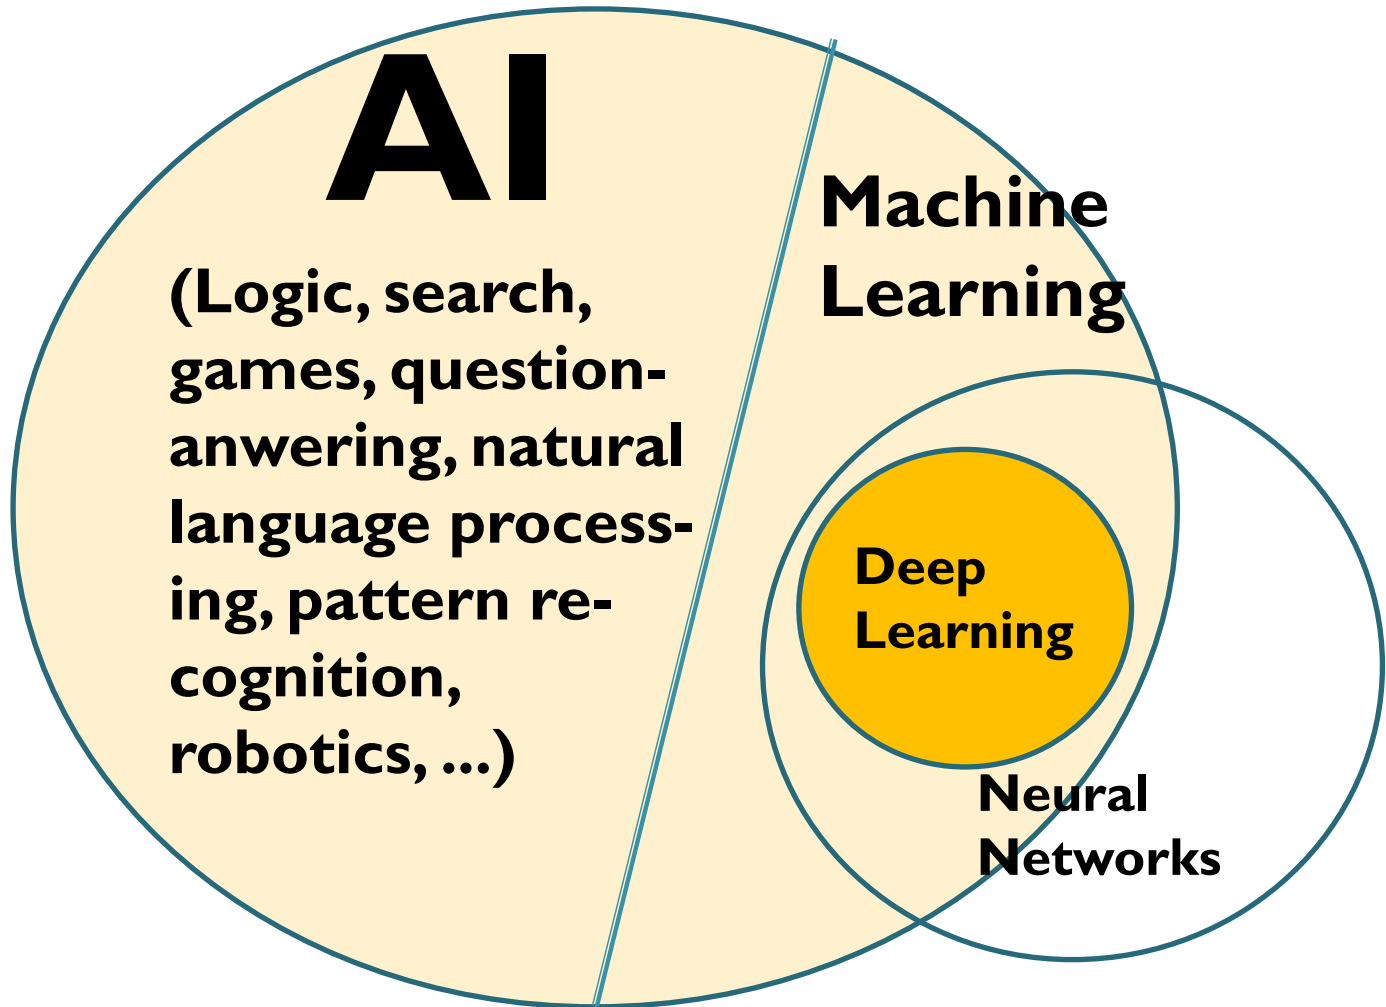
AI

(Logic, search, games, question-answering, natural language processing, pattern recognition, robotics, ...)

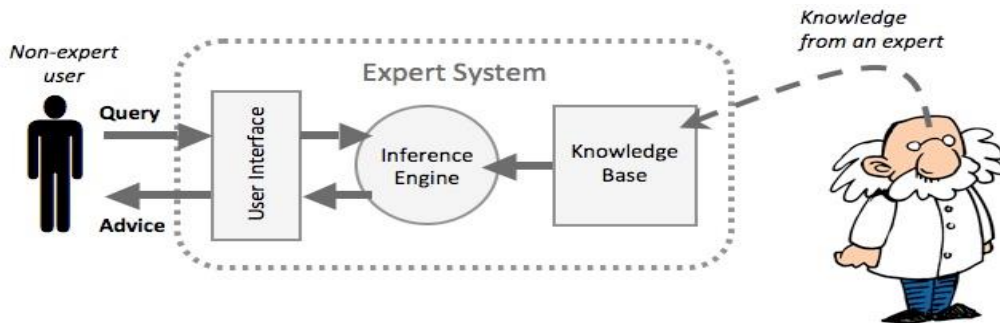
Machine Learning

Deep Learning

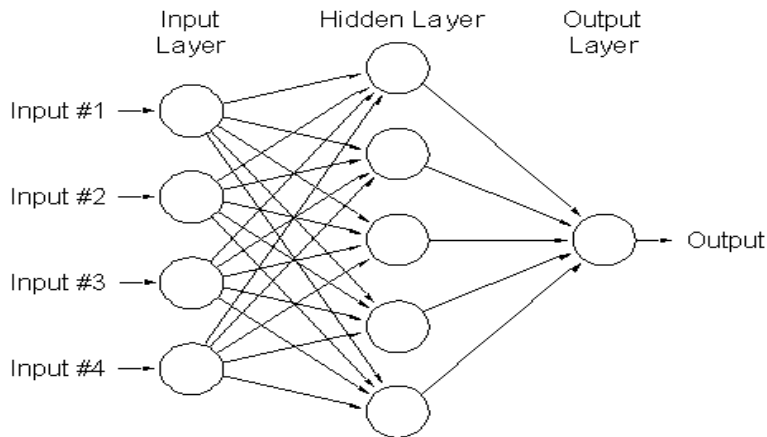
Neural Networks



Two alternative / competing ways to build AI

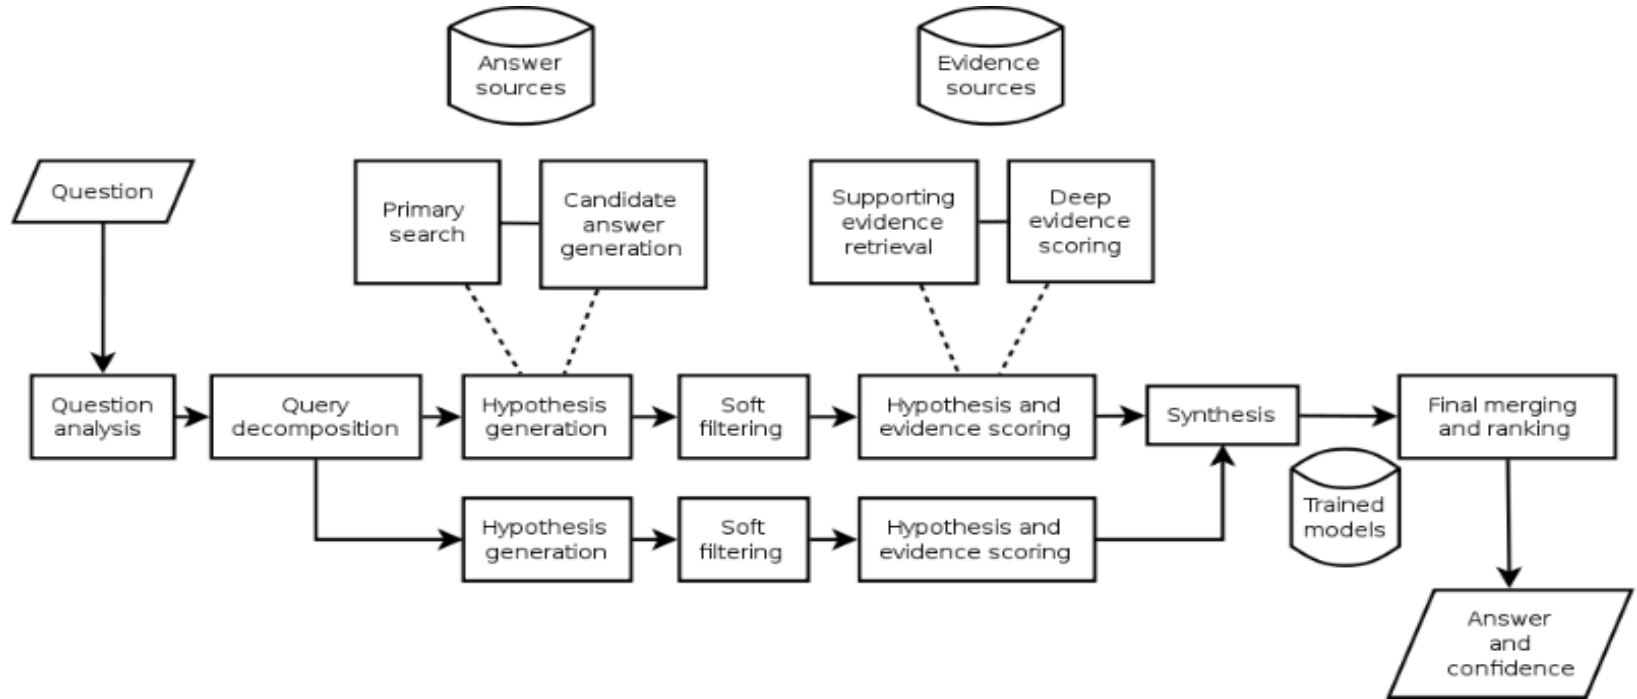


Logic /
symbolic
AI



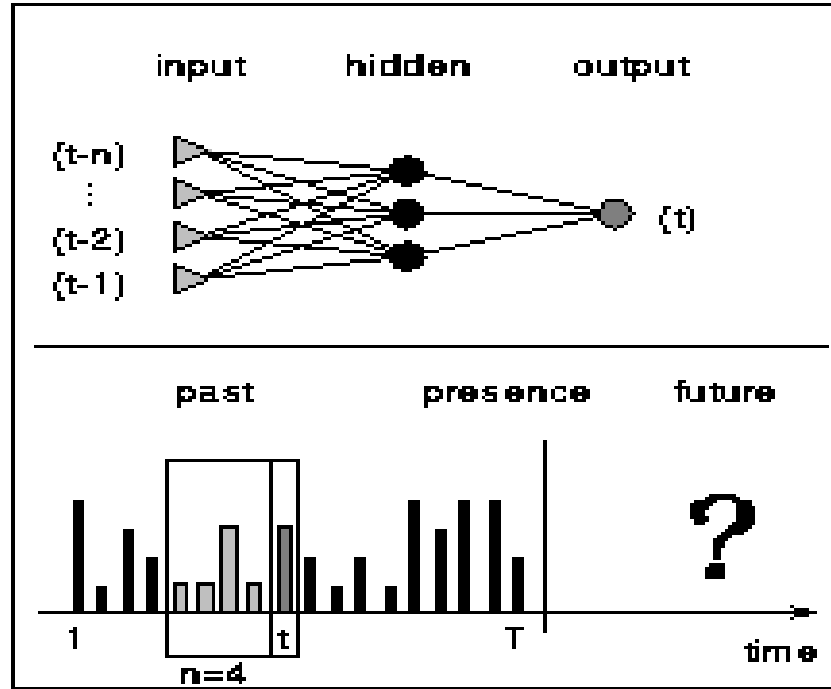
Neural
network with
statistical
machine
learning

Recent "traditional" AI system: IBM Watson. Note the complexity!



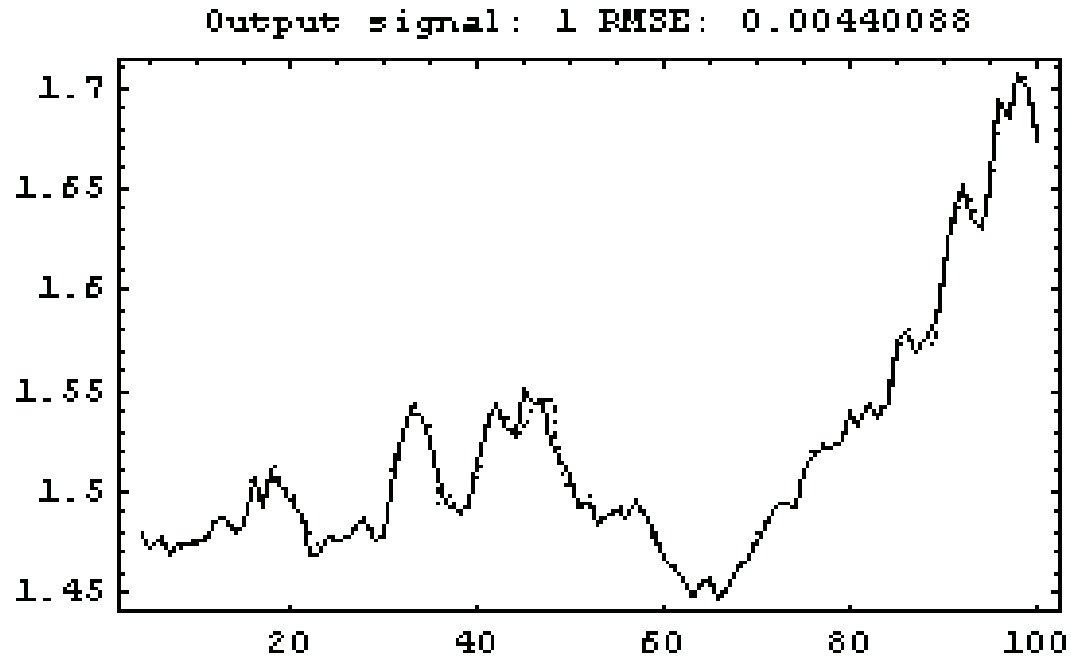
An old example of neural network: **time series prediction**
(Today, this would be called "generative AI")

Training



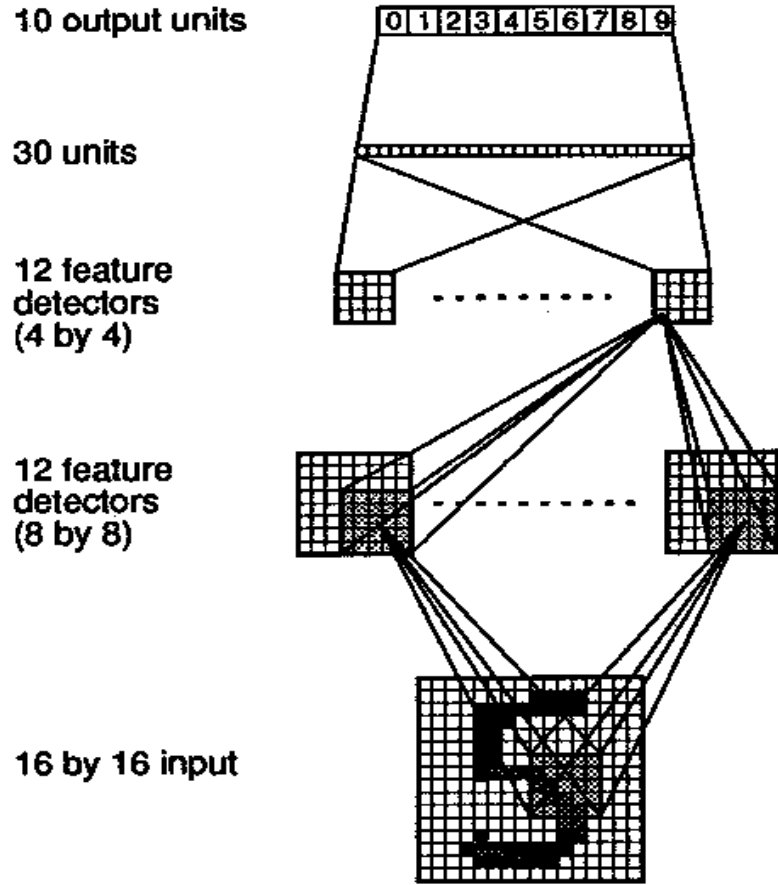
$3 \times (4 + 1)$
parameters
in hidden
layer,
4
parameters
in output
neuron:
19
parameters
altogether

Neuroverkon käyttö aikasarjan ennustamisessa



Tyypillinen ennustustulos (valuuttakurssin kehittyminen)

A truly ground-breaking early NN classifier: **LeNET** by Yann leCun, 1989



Neuroverkkoluokitin käsinkirjoitettujen postinumeroitten luokitteluun

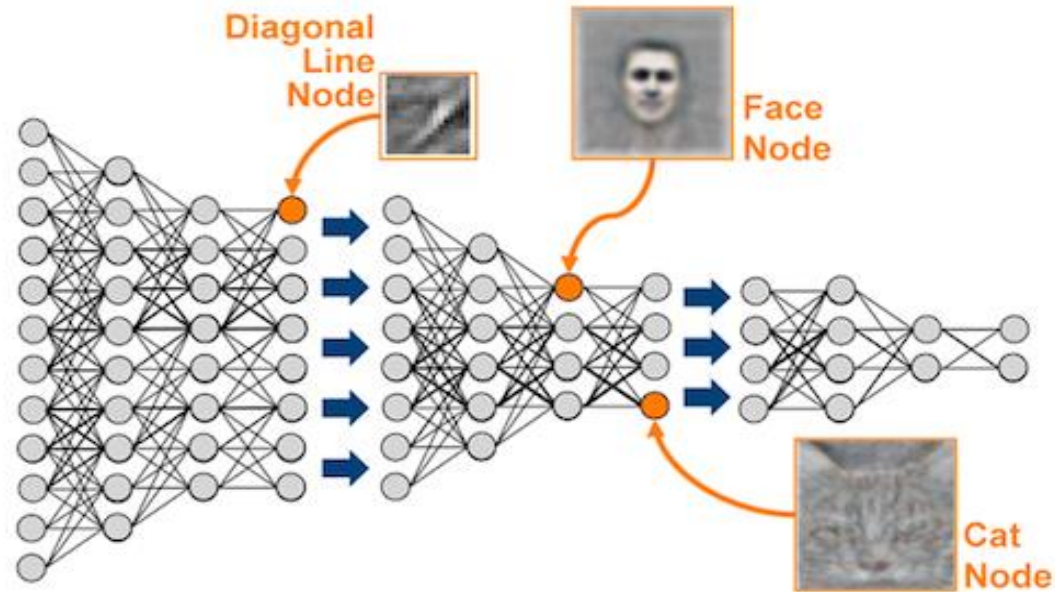
Hottest topic in AI today: Deep Learning


You want to build a computer program that will recognize either human or feline (cat) faces



You use a deep neural network:

Large
training
material
(lots of
faces)





A demo using a much larger image collection (the ImageNet benchmark): five most likely classes shown



kit fox
red fox
dingo
coyote
grey fox



minivan
recreational vehicle
limousine
police van
beach wagon



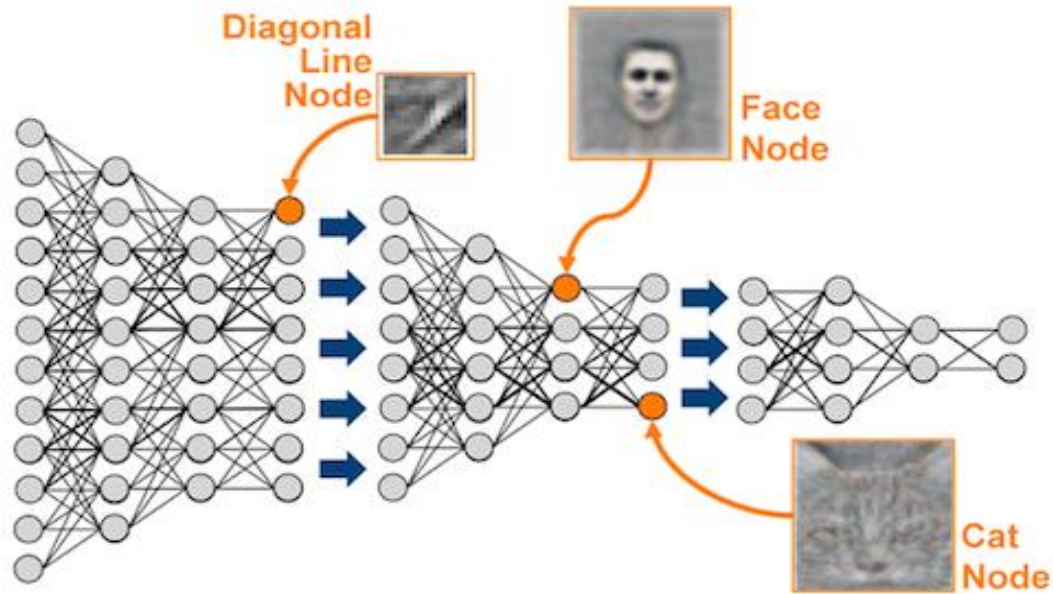
drum
cornet
banjo
collar
toaster



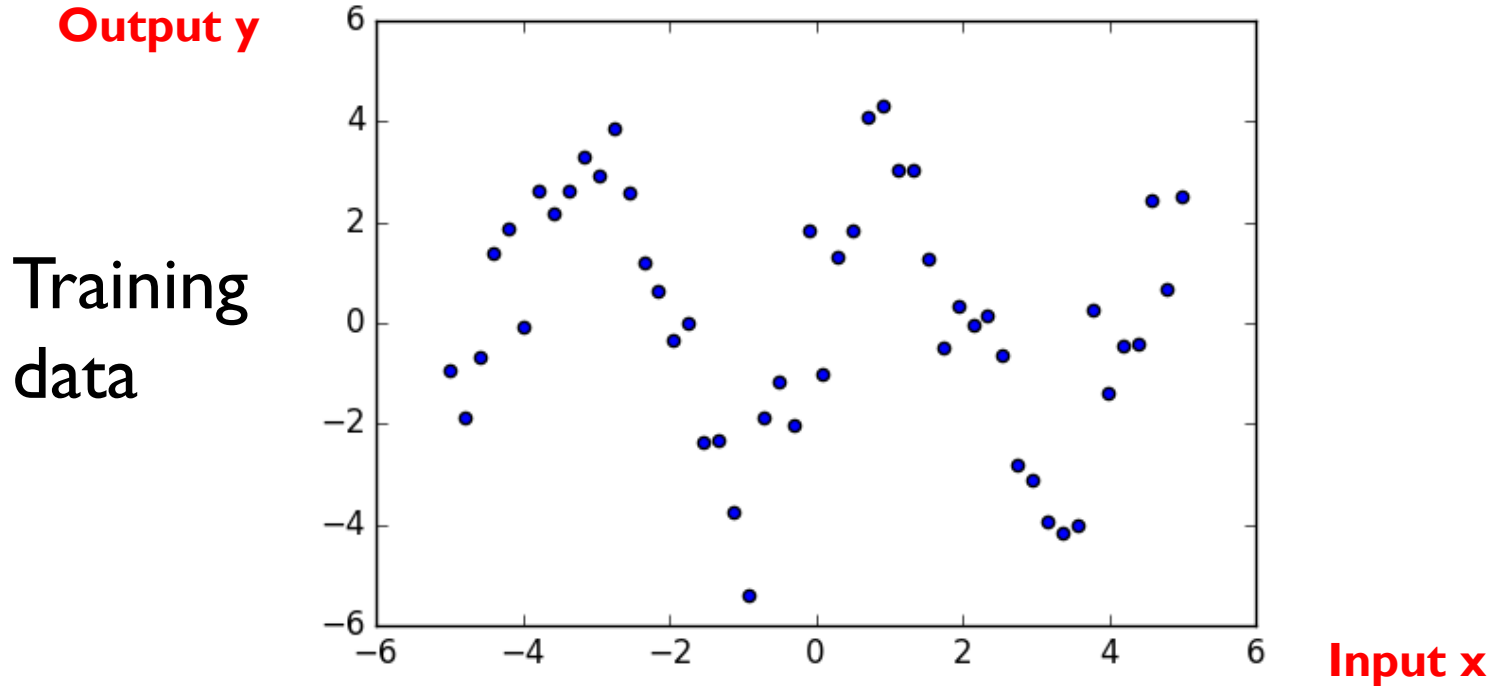
liner
fireboat
drilling platform
aircraft carrier
dock

How and why does deep learning work?

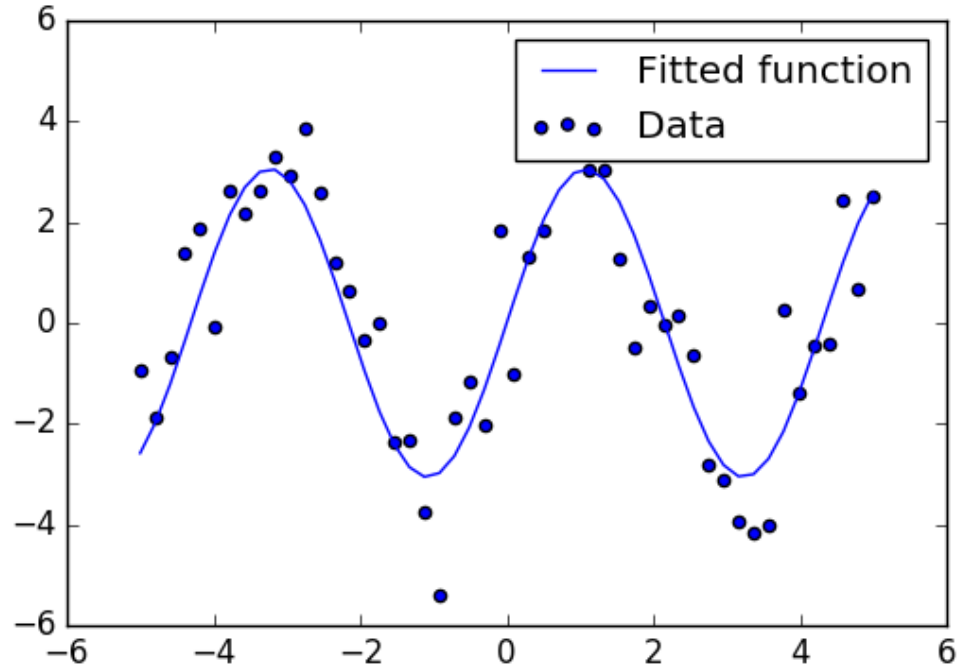
Deep Learning: why it works




Remember "curve fitting" from school

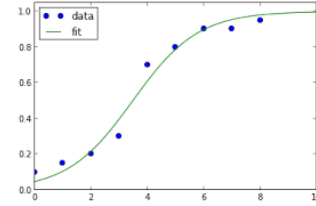


C. F. Gauss, 1809: Least Mean Squares solution

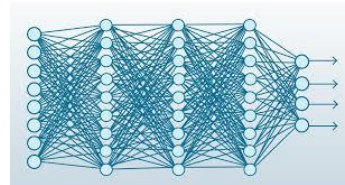


- 
- Easily done by regression analysis software
 - Note: the curve (function) gives you the output value for **any** input value
 - Easy to generalize for many inputs, many outputs
 - Generic ***input-output machine!***

Curve fitting :



Deep neural network :

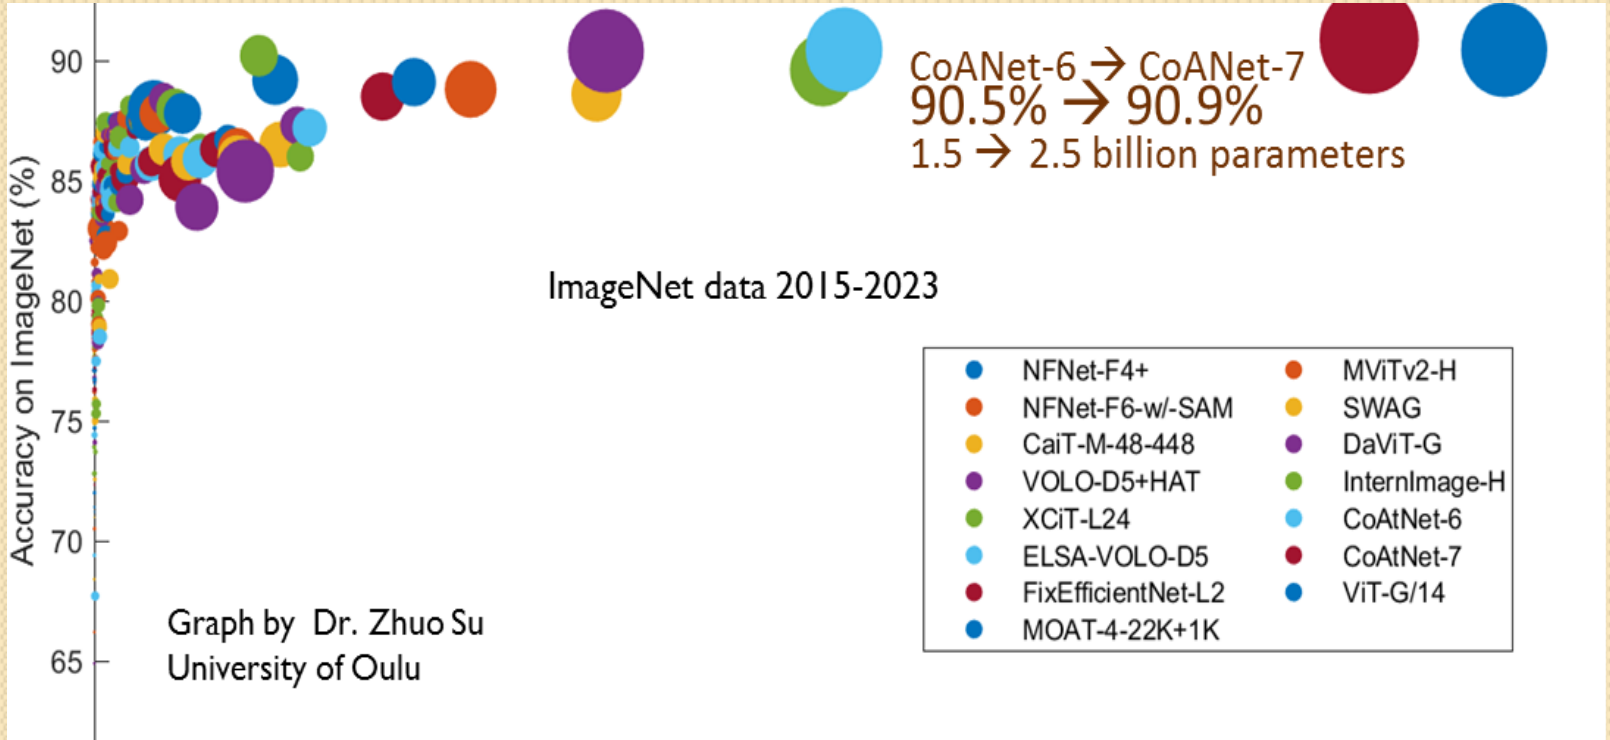


Starting-point for the success: ImageNet (2010)

- In the ImageNet competition, a database of 14 million images was collected and manually classified into 1000 classes
- A major breakthrough: AlexNet by G. Hinton's group (2012). Very deep network (8 layers) and GPU units for computation. Training time 6 days.
- This essentially started the Deep Learning revolution.



Accuracy saturates with network size



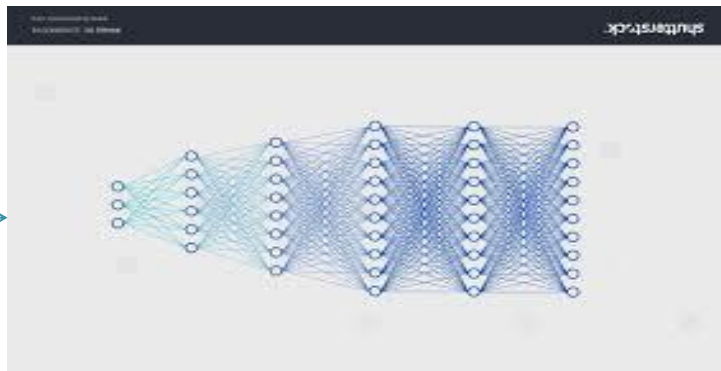


Discriminative and Generative AI

- Classification of "natural" data (images, videos, speech, signals, even text) is today called **Discriminative AI**
- Note that in the input-output function that AlexNet deep network is computing, the input is much larger (higher dimensional) than the output: for images, input image is about $3 \times 1000 \times 1000 = 3$ million pixels, and output is one of 1000 class labels (which can be just numbers from 1 to 1000).
- Mathematically, this is "easy" in theory.

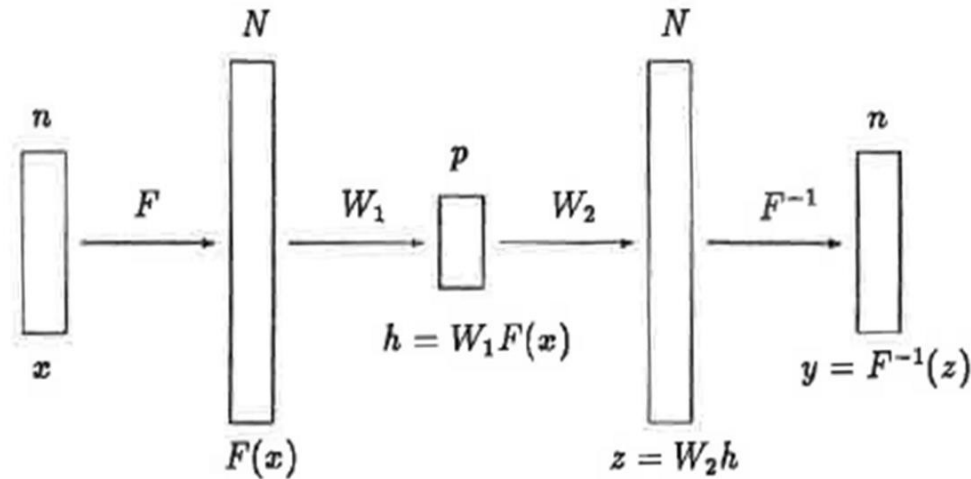
- The new really surprising discovery is that you can turn this around: go from labels to images

CAT →




This is called **Generative AI**.

What is needed is a multi-layer neural network, **deep auto-encoder**. A very early version was suggested by Oja (1991).



: A 5-layer network with linear and nonlinear layers. The boxes denote layers; the number in the layer is given above the box, and the output vector of the layer is given under the e arrows give the transformation between two layers.

- 
- Today, auto-encoders have many more hidden layers.
 - The "latent variables" in the bottle-neck have the same dimensionality as the text label.
 - More advanced: Generative Adversarial Network (GAN).

Basic principle behind ChatGPT

- This is the most famous AI system today
- Given a query, it generates text
- How does it work?
- First, just an example:



RA

You

Could you explain the term 'Information Architecture' in a way that anyone can understand without using jargon? Use an example to explain in a simple manner.

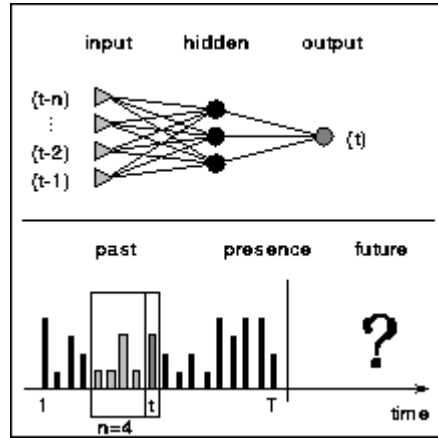


ChatGPT

Information Architecture is like a library's layout. Imagine walking into a library to find a book about dogs. A well-organized library has clear signs and sections. You'd find a 'Pets' section, then look for 'Dogs'. Everything is sorted to help you find books easily. Information Architecture does this for websites. It organizes content so you can find information quickly, just like books in a library.

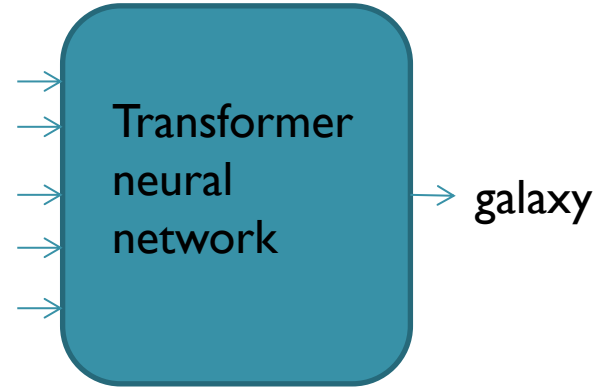


A very simplified explanation




Remember this?

Together
we
will
rule
the



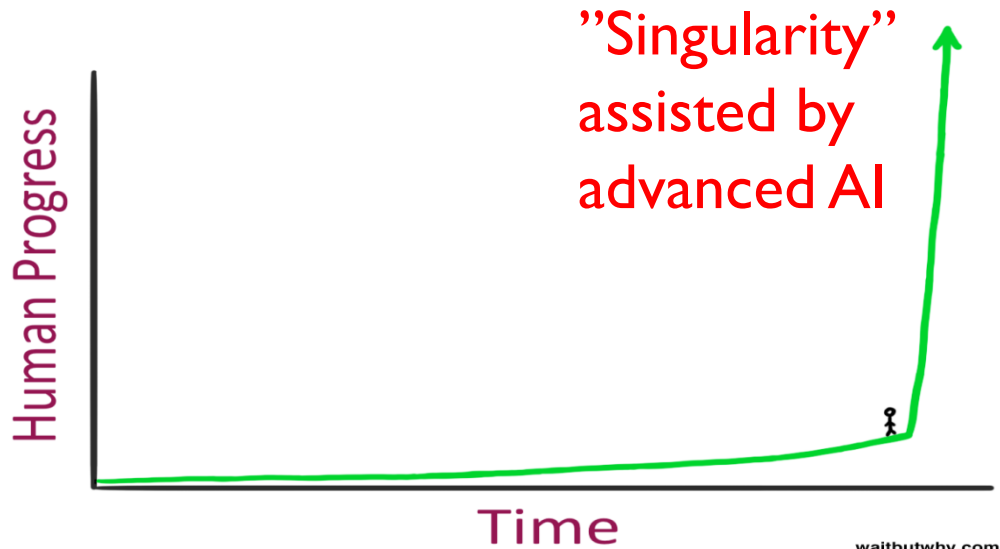
Generative Pre-trained
Transformer (GPT), also
called Large Language
Model (LLM)

- 
- The Transformer deep neural network in the GPT is huge: about 175 billion parameters, about 500 billion word training material (all text in WWW)
 - Note that the training is **self-supervised** (automatic): take a sequence of words, e.g. a sentence, use first part as training input, second part as training output
 - Size matters: increasing the size improves the results. This is a **brute force** approach, only possible for the richest IT companies.



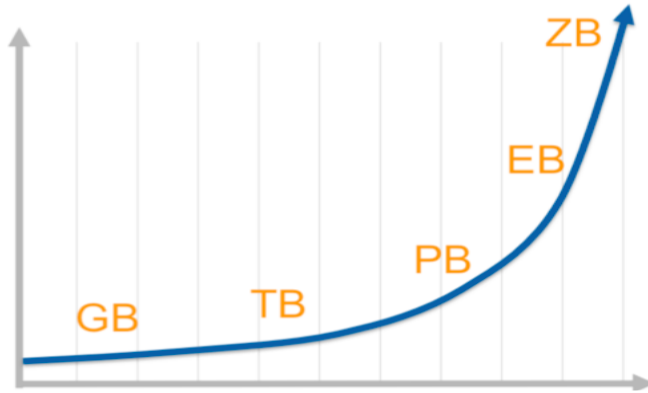
The future?

- With applications like ChatGPT, the future of AI looks bright, although with many problems
- Even this?



Two basic building blocks for AI (even if algorithms are "stupid")

Training data



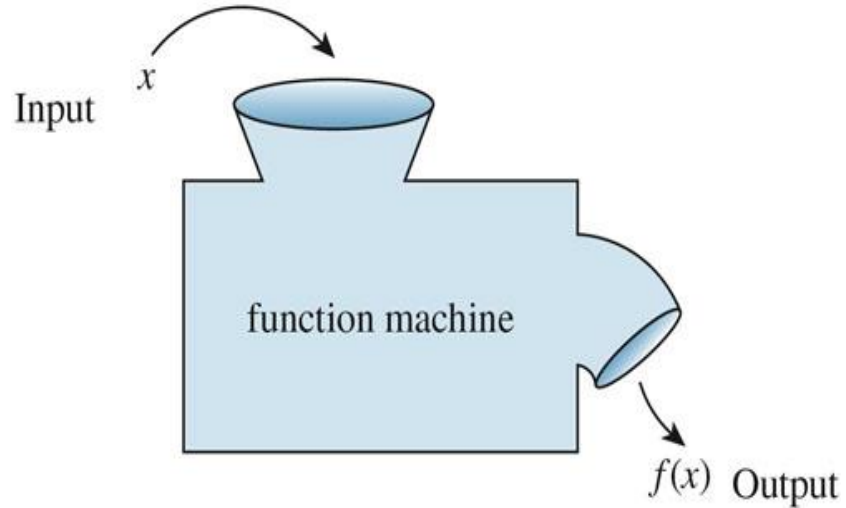
Computing power




Brute force!

Short-time goal

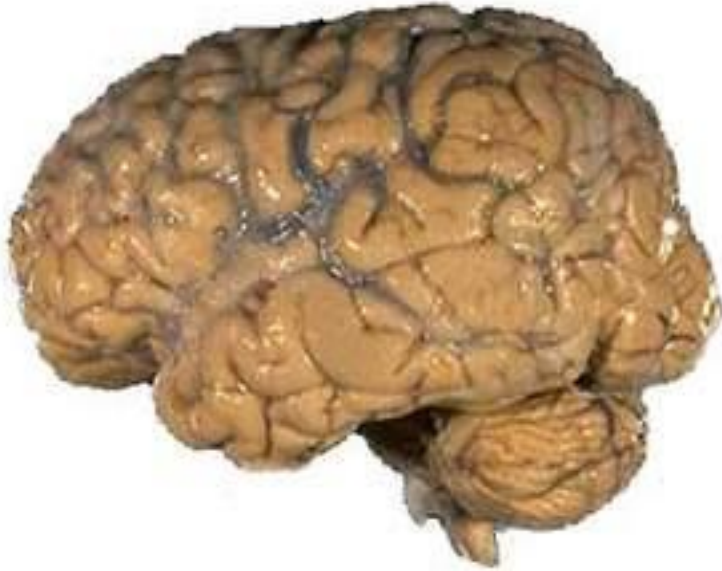
- Extending Machine Learning to more practical applications
- Theory: e.g. learnability of complex cognitive functions (such as inference, recognition, planning)



- 
- Moving cognitive routines from human brains to machines will be the new wave of automation
 - In the long run, Deep Learning will probably not be enough, but a synthesis of symbolic/logic AI and statistical machine learning will be needed.

The big question:

- What does the unexpected success of Deep Learning tell us about the world?
- What does it tell us about human/animal intelligence?



Thank you
for your
attention!