Introduction to artificial intelligence

Arthur Ruder ¦ Enclustra GmbH ¦ AI seminar EPFL Lausanne & ZHAW Winterthur ¦ 19 & 21/11/2019

ENCLUSTRA
FPGA SOLUTIONS

Introduction to artificial intelligence

...more specifically: machine learning and artificial neural networks

Arthur Ruder ¦ Enclustra GmbH ¦ AI seminar EPFL Lausanne & ZHAW Winterthur ¦ 19 & 21/11/2019

Motivation

Motivation

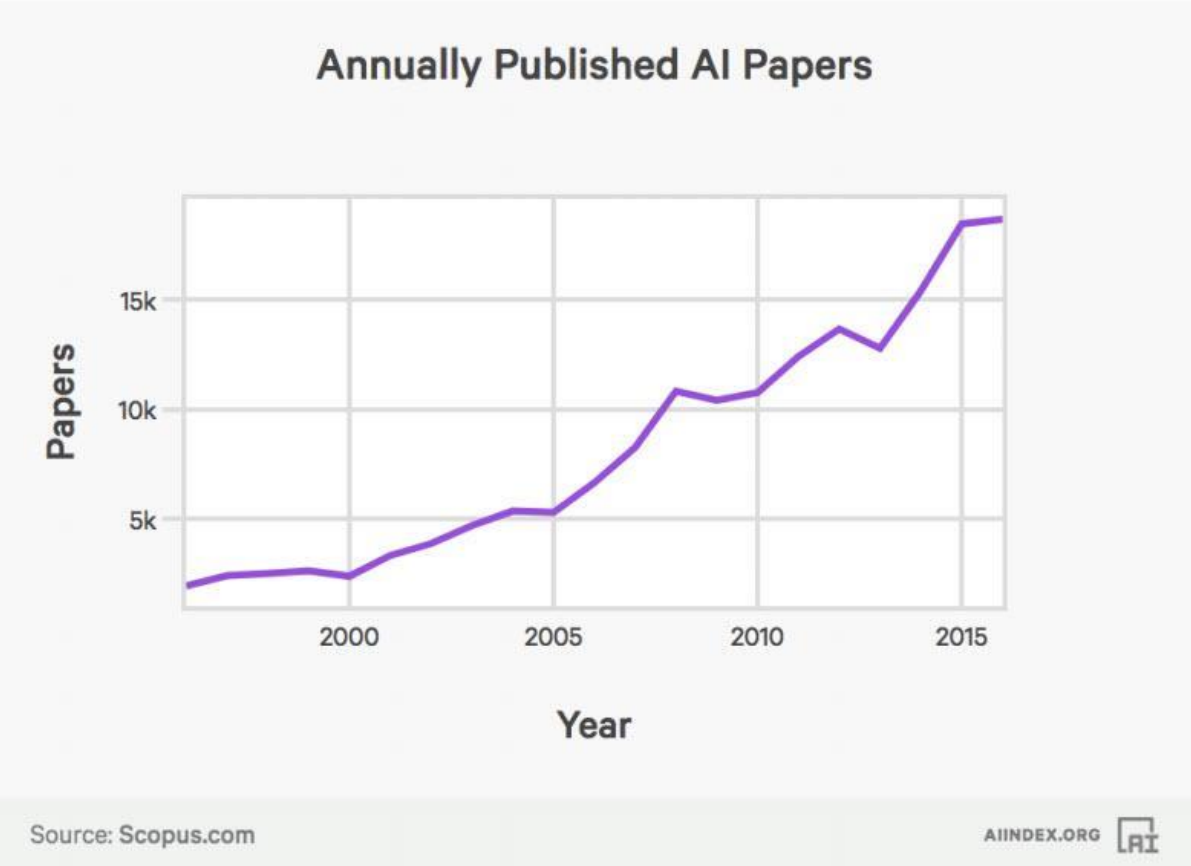

https://www.forbes.com/sites/louiscolumbus/2018/01/12/10-charts-that-will-change-your-perspective-on-artificial-intelligences-growth/#589705f94758

Motivation

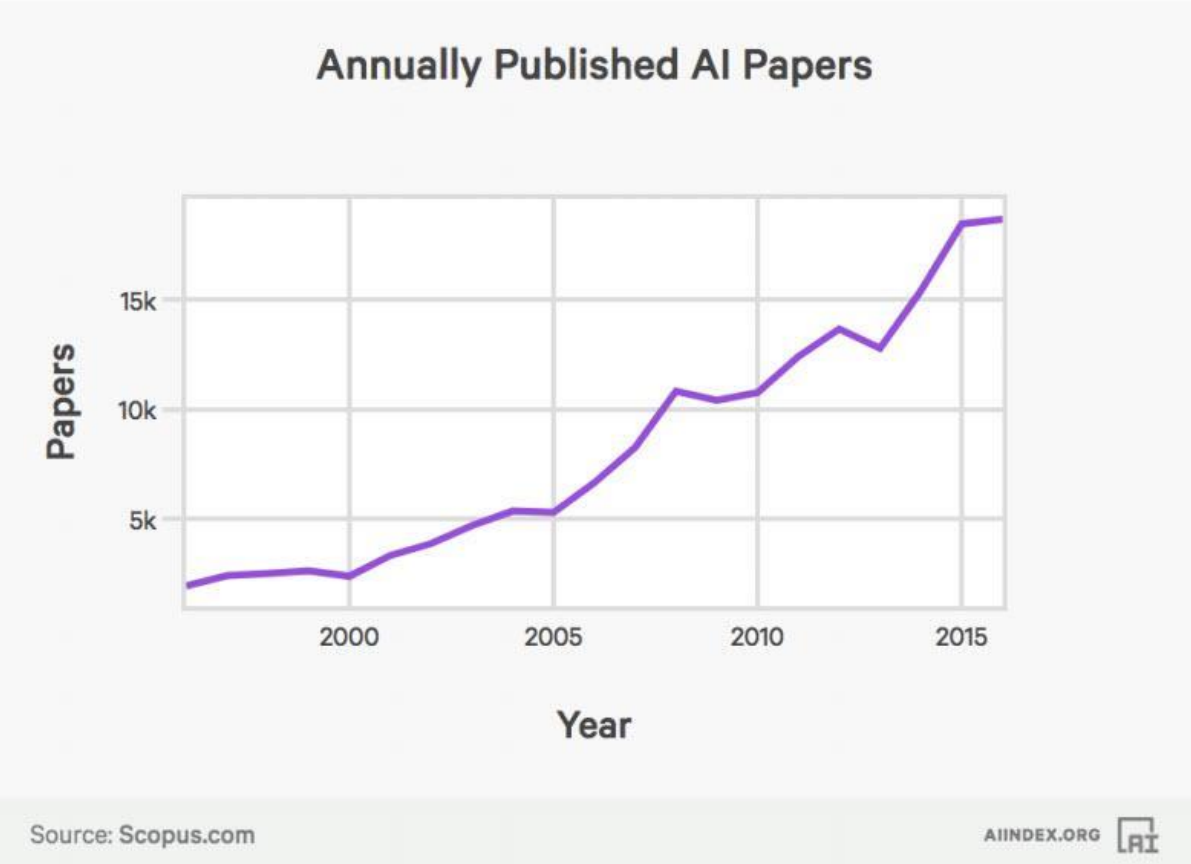

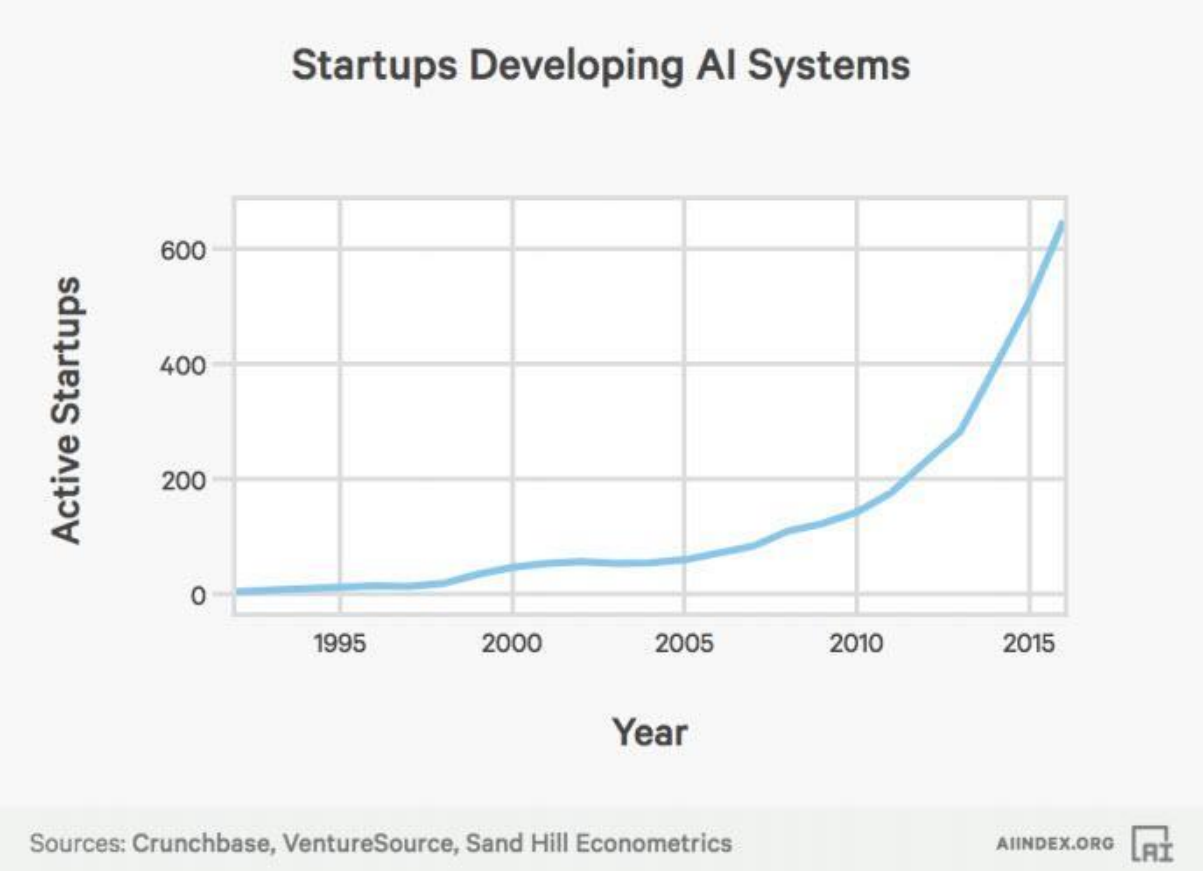

https://www.forbes.com/sites/louiscolumbus/2018/01/12/10-charts-that-will-change-your-perspective-on-artificial-intelligences-growth/#589705f94758

Applications

Applications

- Computer vision
  - Image detection
  - Image classification
  - …

Applications

- Computer vision
  - Image detection
  - Image classification
  - …
- Language processing
  - Speech recognition
  - Translation
  - …

Applications

- Computer vision
  - Image detection
  - Image classification
  - …
- Language processing
  - Speech recognition
  - Translation
  - …
- Recommendation systems
- …

Overview: What is machine learning?

ENCLUSTRA FPGA SOLUTIONS

Artificial Intelligence

Overview: What is machine learning?

ENCLUSTRA FPGA SOLUTIONS

Artificial Intelligence

Machine Learning

Overview: What is machine learning?

ENCLUSTRA FPGA SOLUTIONS

Artificial Intelligence

More advanced concepts

Machine Learning

Overview: What is machine learning?

Artificial Intelligence

More advanced concepts

Machine Learning

Supervised Learning

Regression
- Linear
- Logistic

Artificial neural networks
- Convolutional
- Recurrent
- LSTM
- …

Overview: What is machine learning?

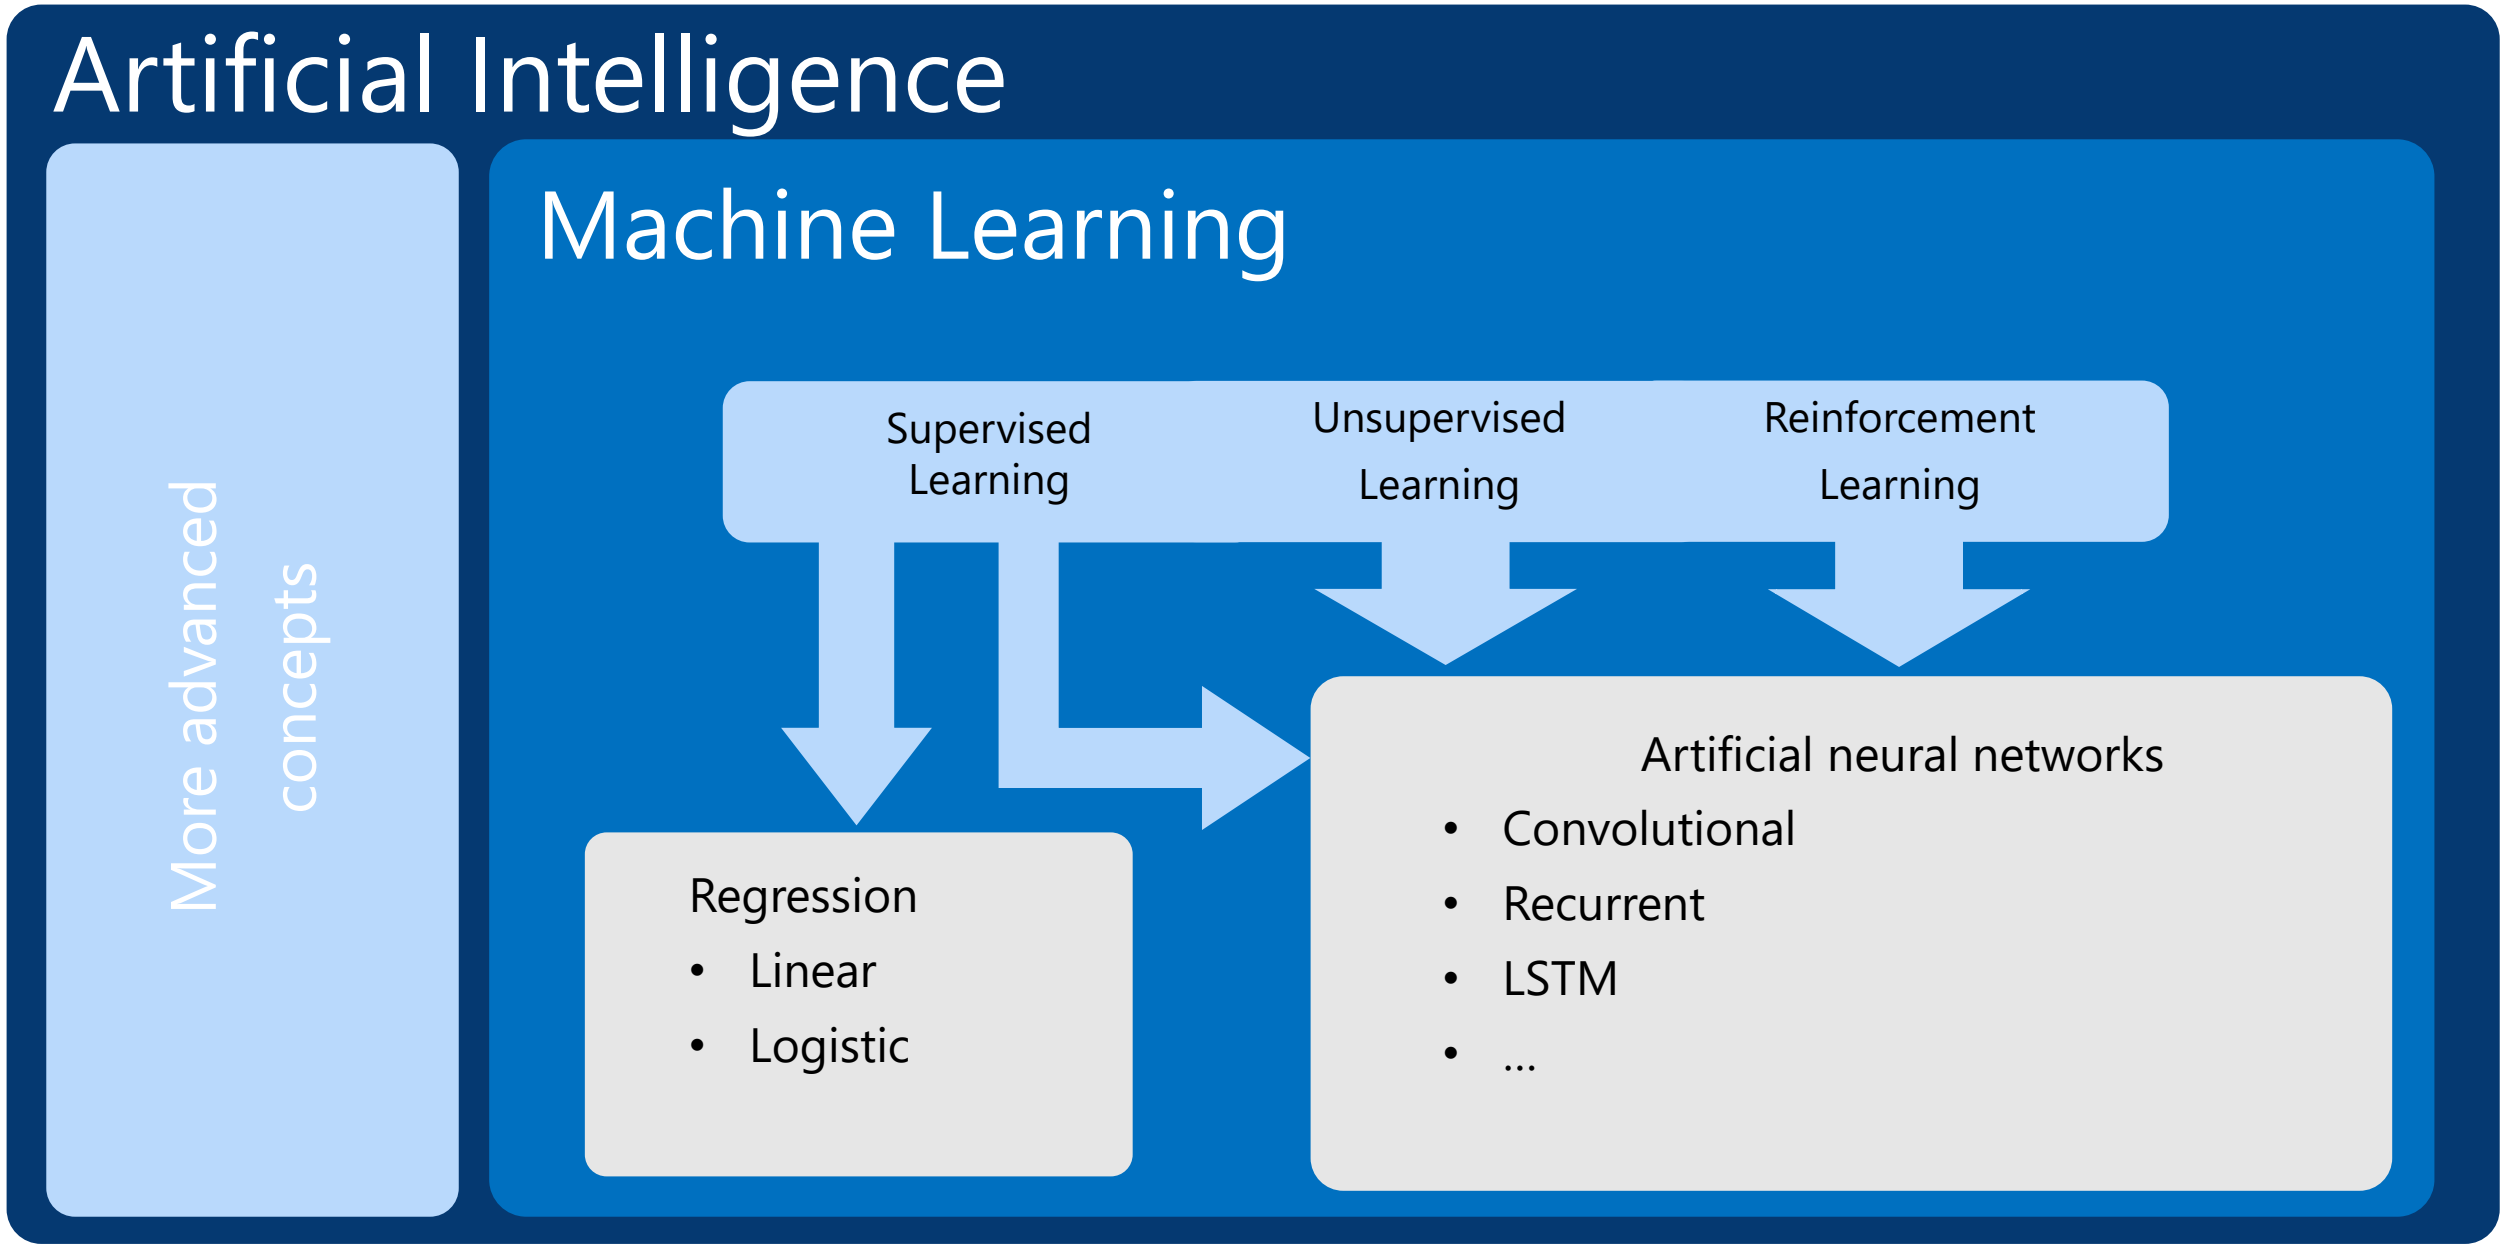

Supervised learning

Supervised learning

Problem: predict real estate prices

Living area [m^2]	Price [CHF]
80	700.000
30	250.000
⋮	⋮

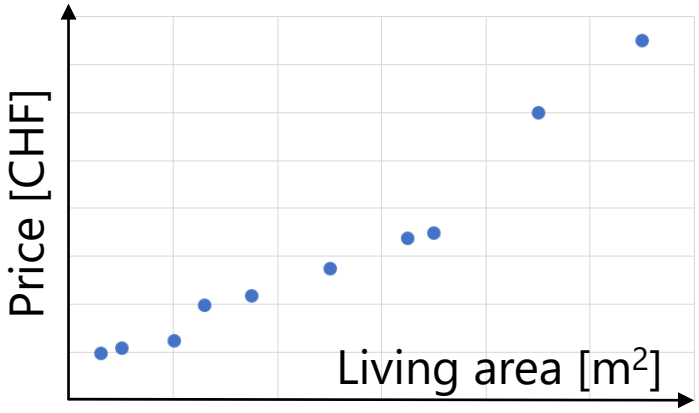

Supervised learning

Problem: predict real estate prices

Living area [m^2]	Price [CHF]
80	700.000
30	250.000
⋮	⋮

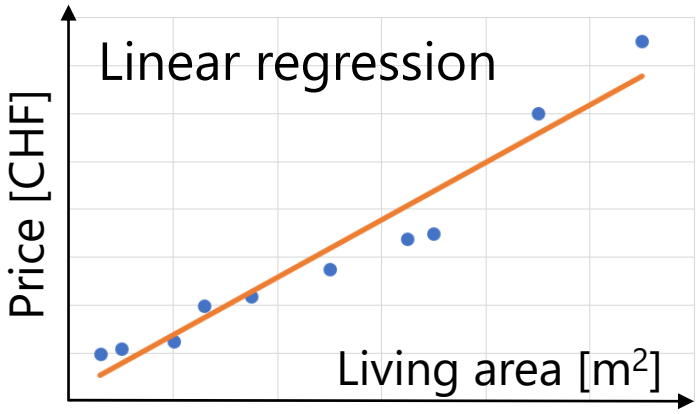

Supervised learning

Problem: predict real estate prices

Living area [m²]	Price [CHF]
80	700.000
30	250.000
⋮	⋮

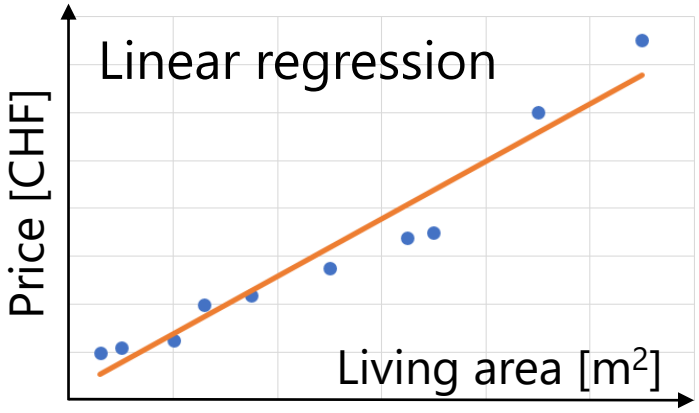

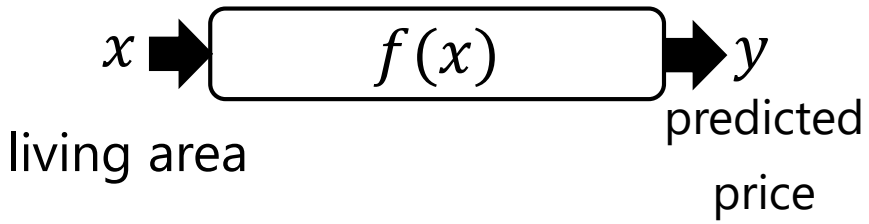

Supervised learning

Problem: predict real estate prices

Living area [m²]	Price [CHF]
80	700.000
30	250.000
⋮	⋮

Concept

Data set

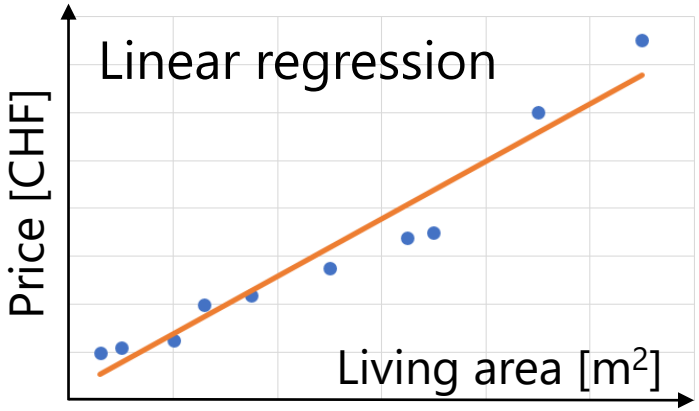

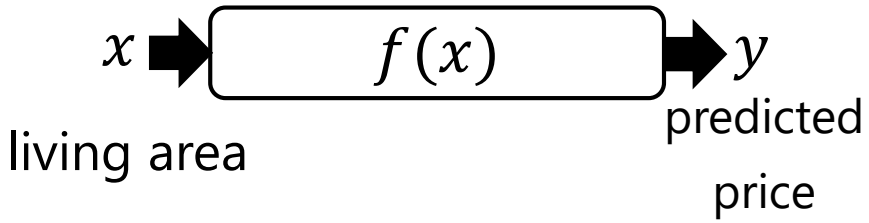

Supervised learning

Problem: predict real estate prices

Living area [m²]	Price [CHF]
80	700.000
30	250.000
⋮	⋮

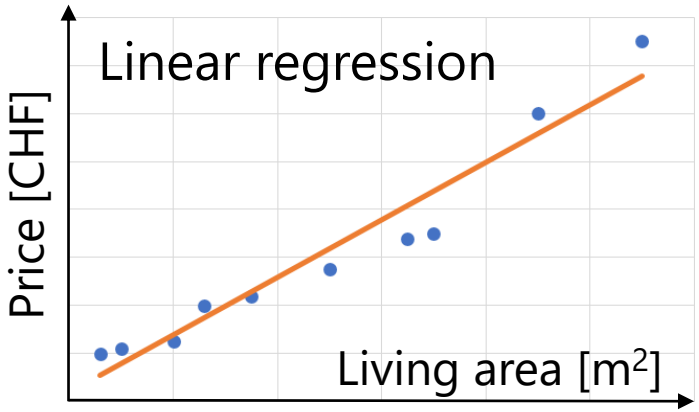

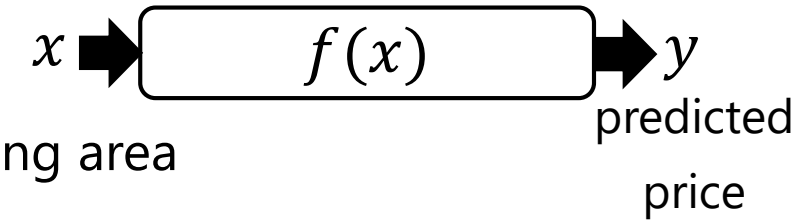

Concept

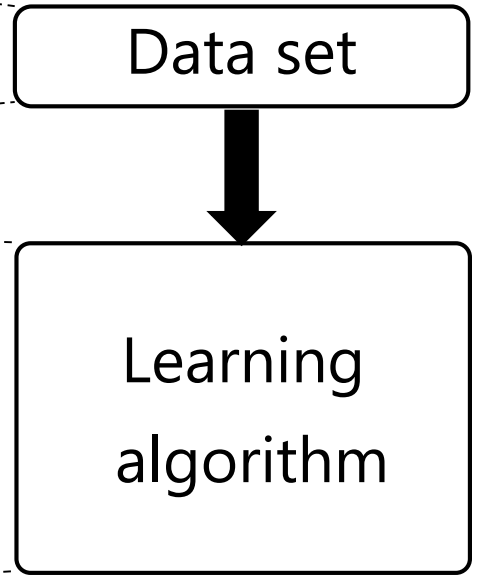

Supervised learning

Problem: predict real estate prices

Living area [m²]	Price [CHF]
80	700.000
30	250.000
⋮	⋮

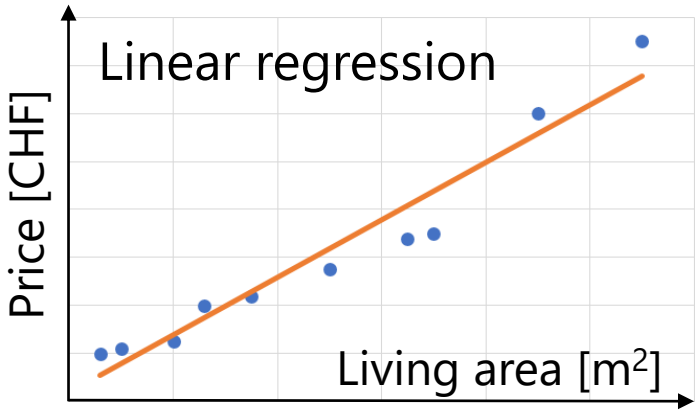

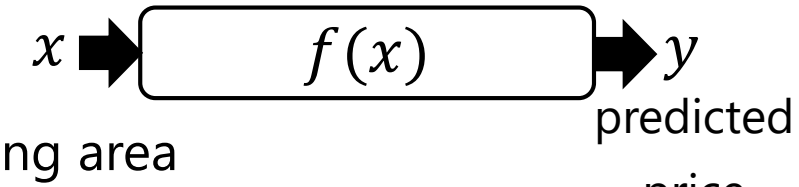

Concept

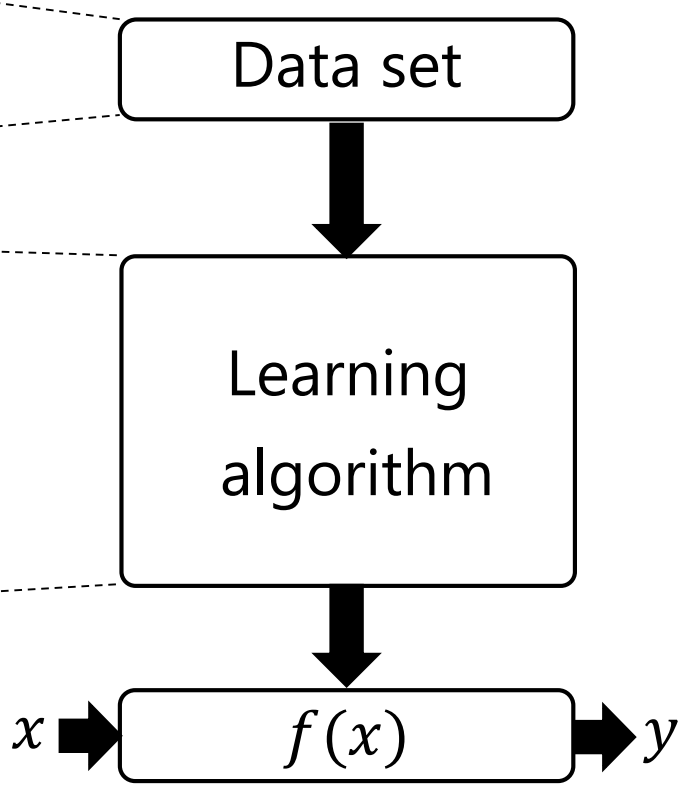

Artificial neural networks

Artificial neural networks

Input: pictures (→ pixels)

⋱

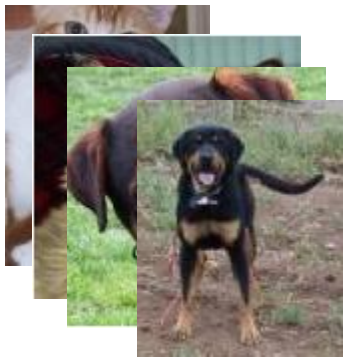

Artificial neural networks

Input: pictures (→ pixels)

∴

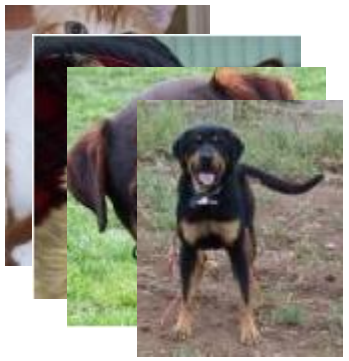

Artificial neural network

Artificial neural networks

Input: pictures (→ pixels)

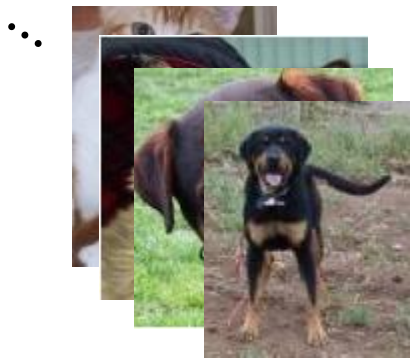

Artificial neural network

Output: probability → dog or cat?

Artificial neural networks

Input: pictures (→ pixels)

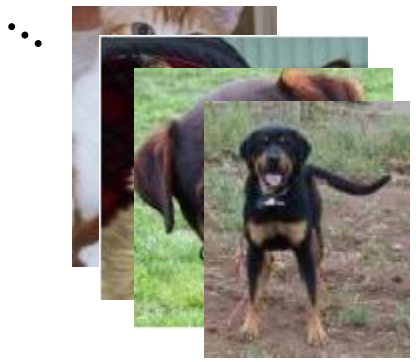

Artificial neural network

Output: probability → dog or cat?

What is inside of this black box?

Neuron

Neuron

Biological

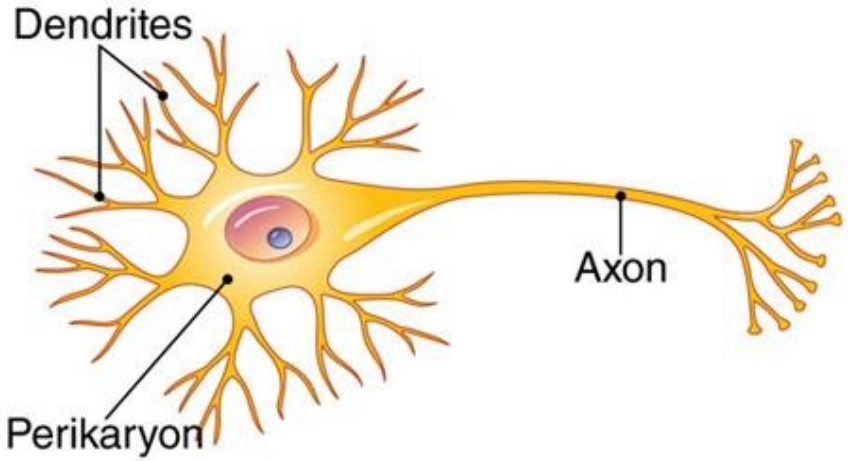

http://ice.uthscsa.edu/dentalhistologyprepcourse/09%20nerve/Nerve_print.html

- Dendrites can receive signals
- Once threshold reached → send signal down the axon

Neuron

Biological

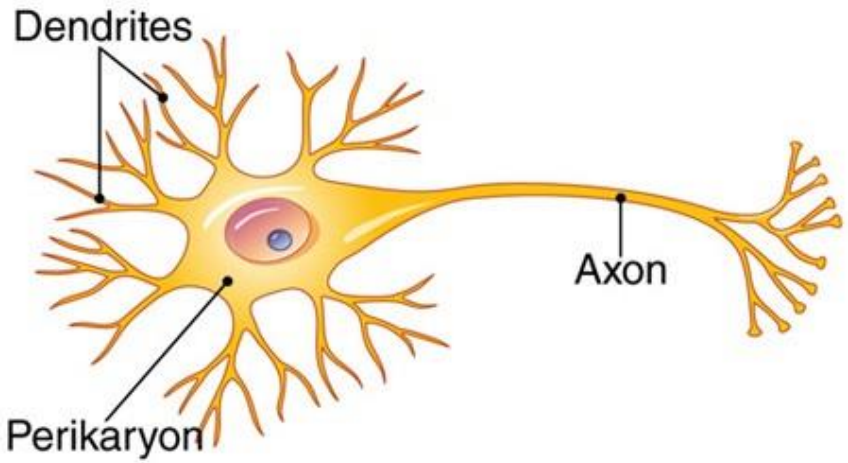

http://ice.uthscsa.edu/dentalhistologyprepcourse/09%20nerve/Nerve_print.html

- Dendrites can receive signals
- Once threshold reached → send signal down the axon

Artificial

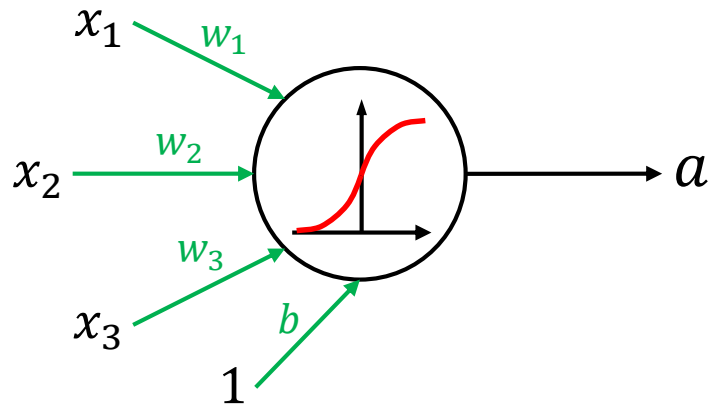

- $\vec{x}$: inputs vector $(x_1, \dots, x_n)$
- $\vec{w}$: weights vector $(w_1, \dots, w_n)$
  → adjustable
- b : bias term
  → adjustable
- a : activation

Artificial neural network architecture

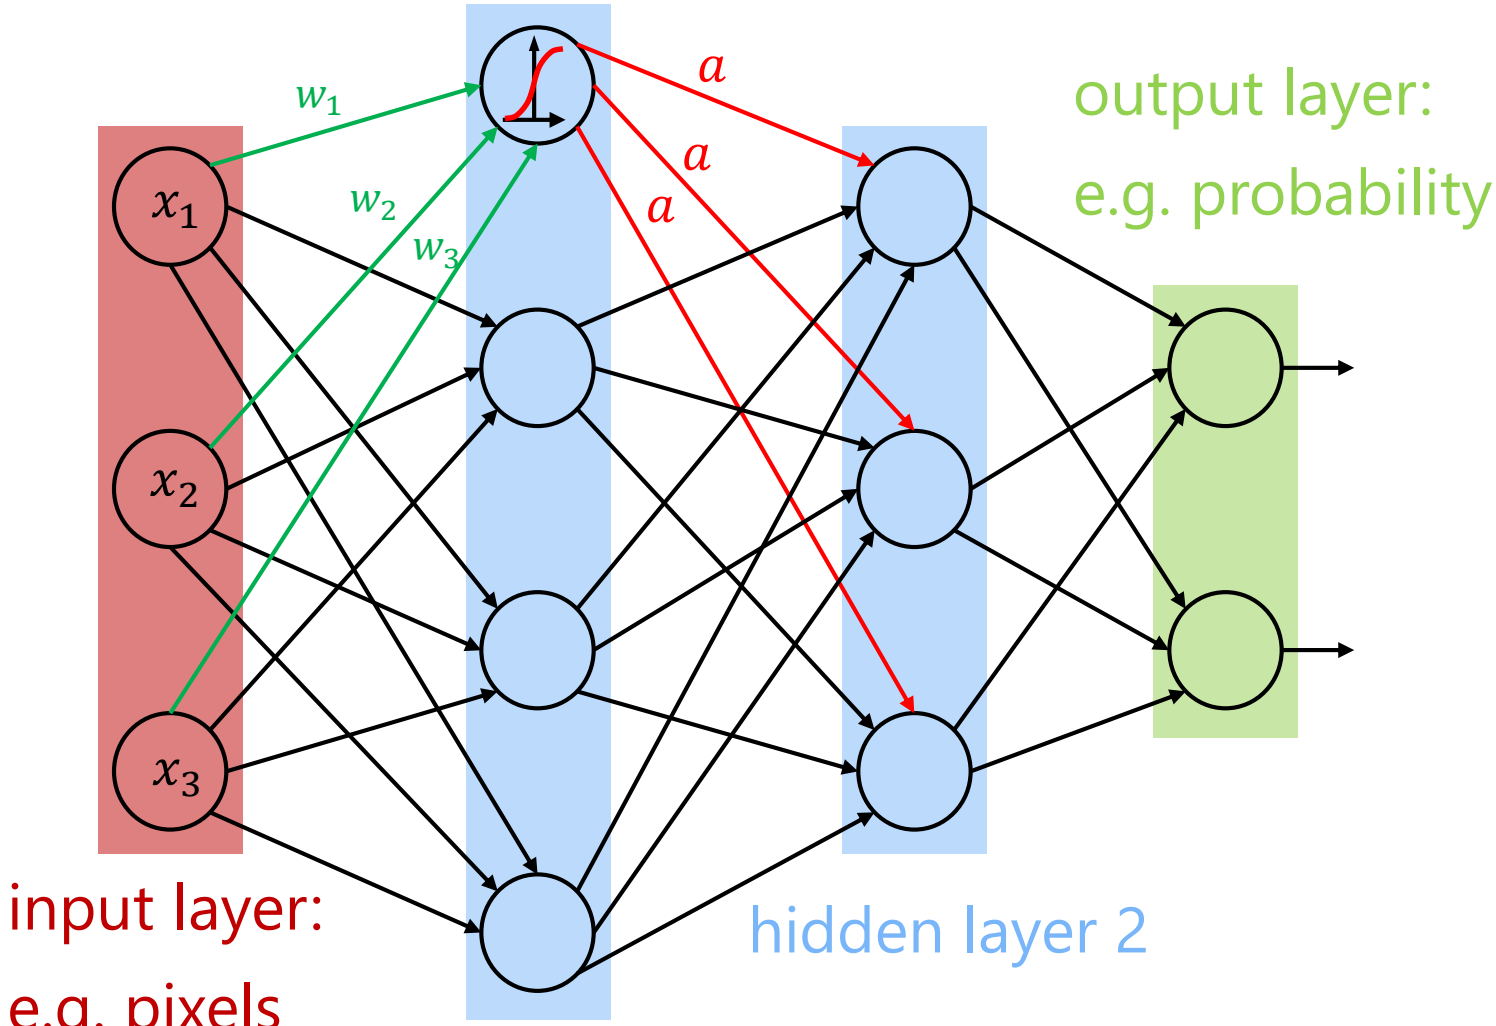

Machine learning concepts: training phase

- Goal: obtain trained weights

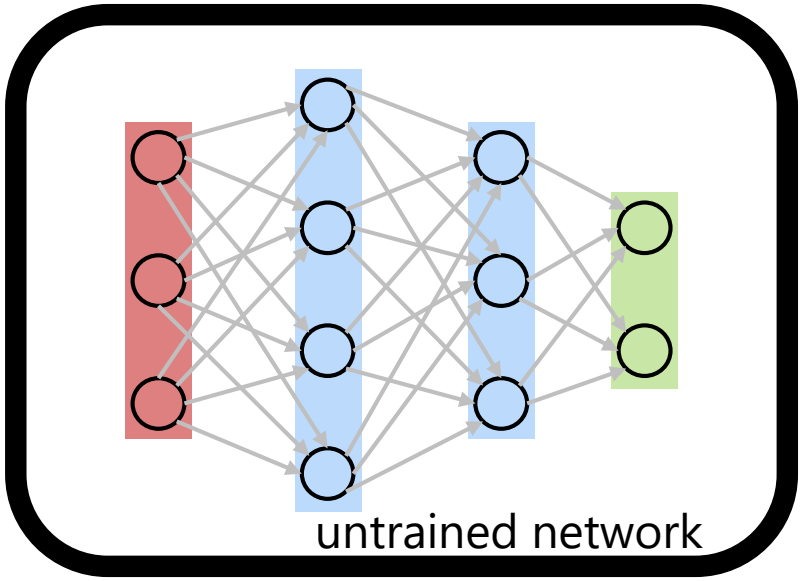

Machine learning concepts: training phase

- Goal: obtain trained weights

Inputs: training set

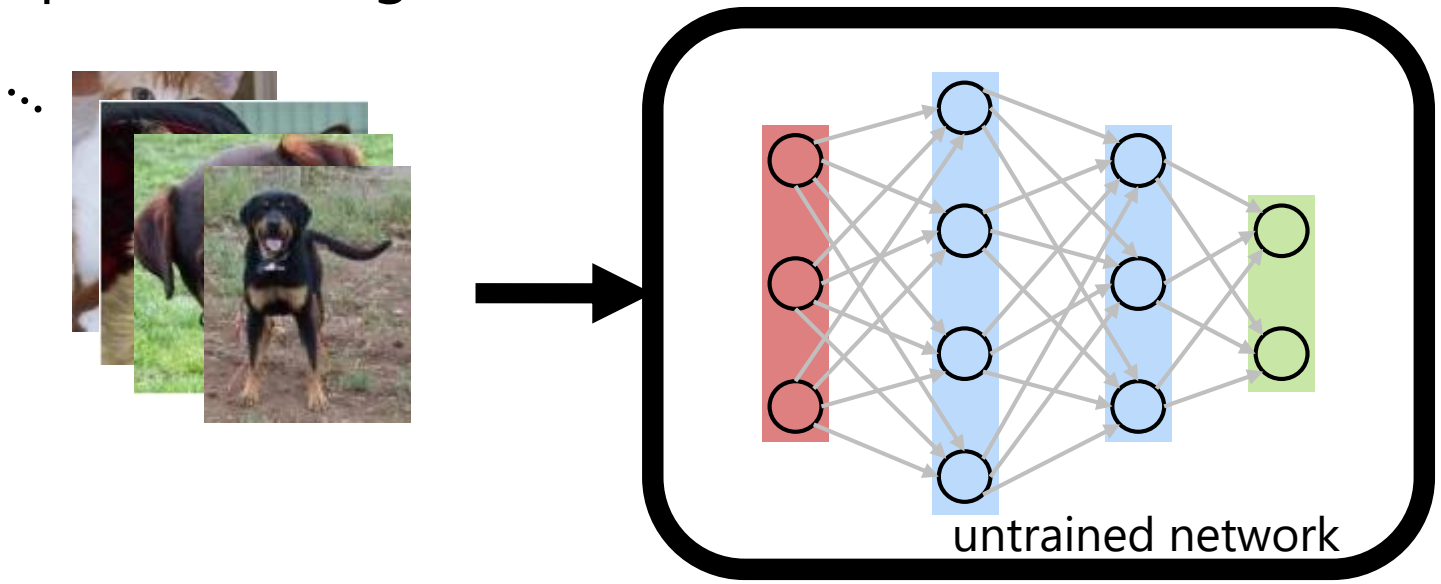

Machine learning concepts: training phase

- Goal: obtain trained weights

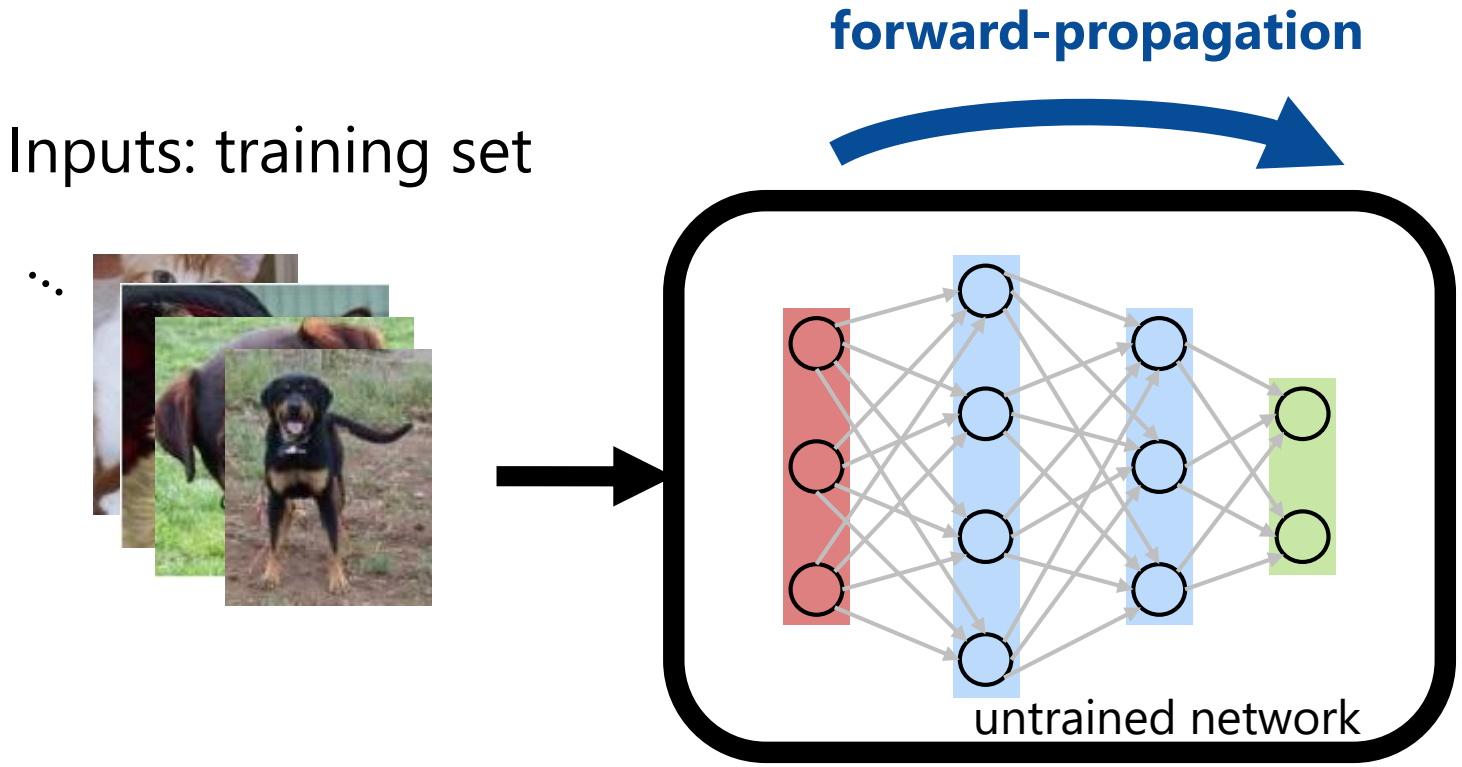

Machine learning concepts: training phase

- Goal: obtain trained weights

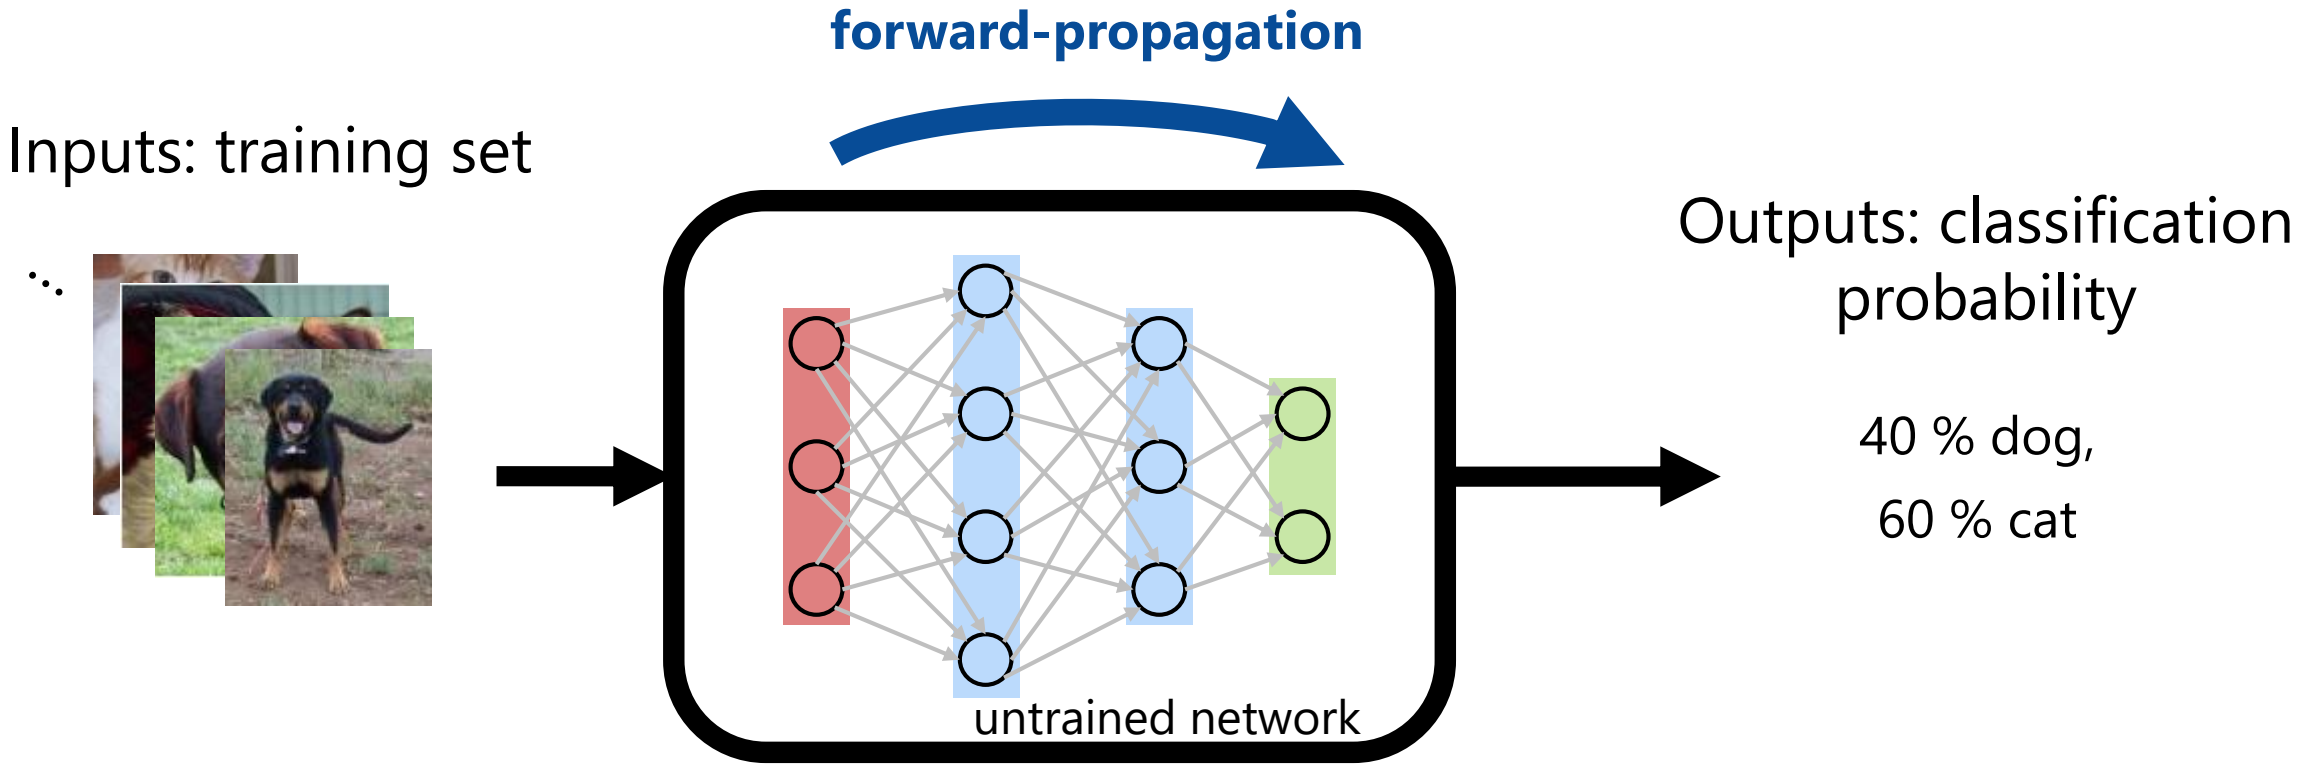

Machine learning concepts: training phase

- Goal: obtain trained weights

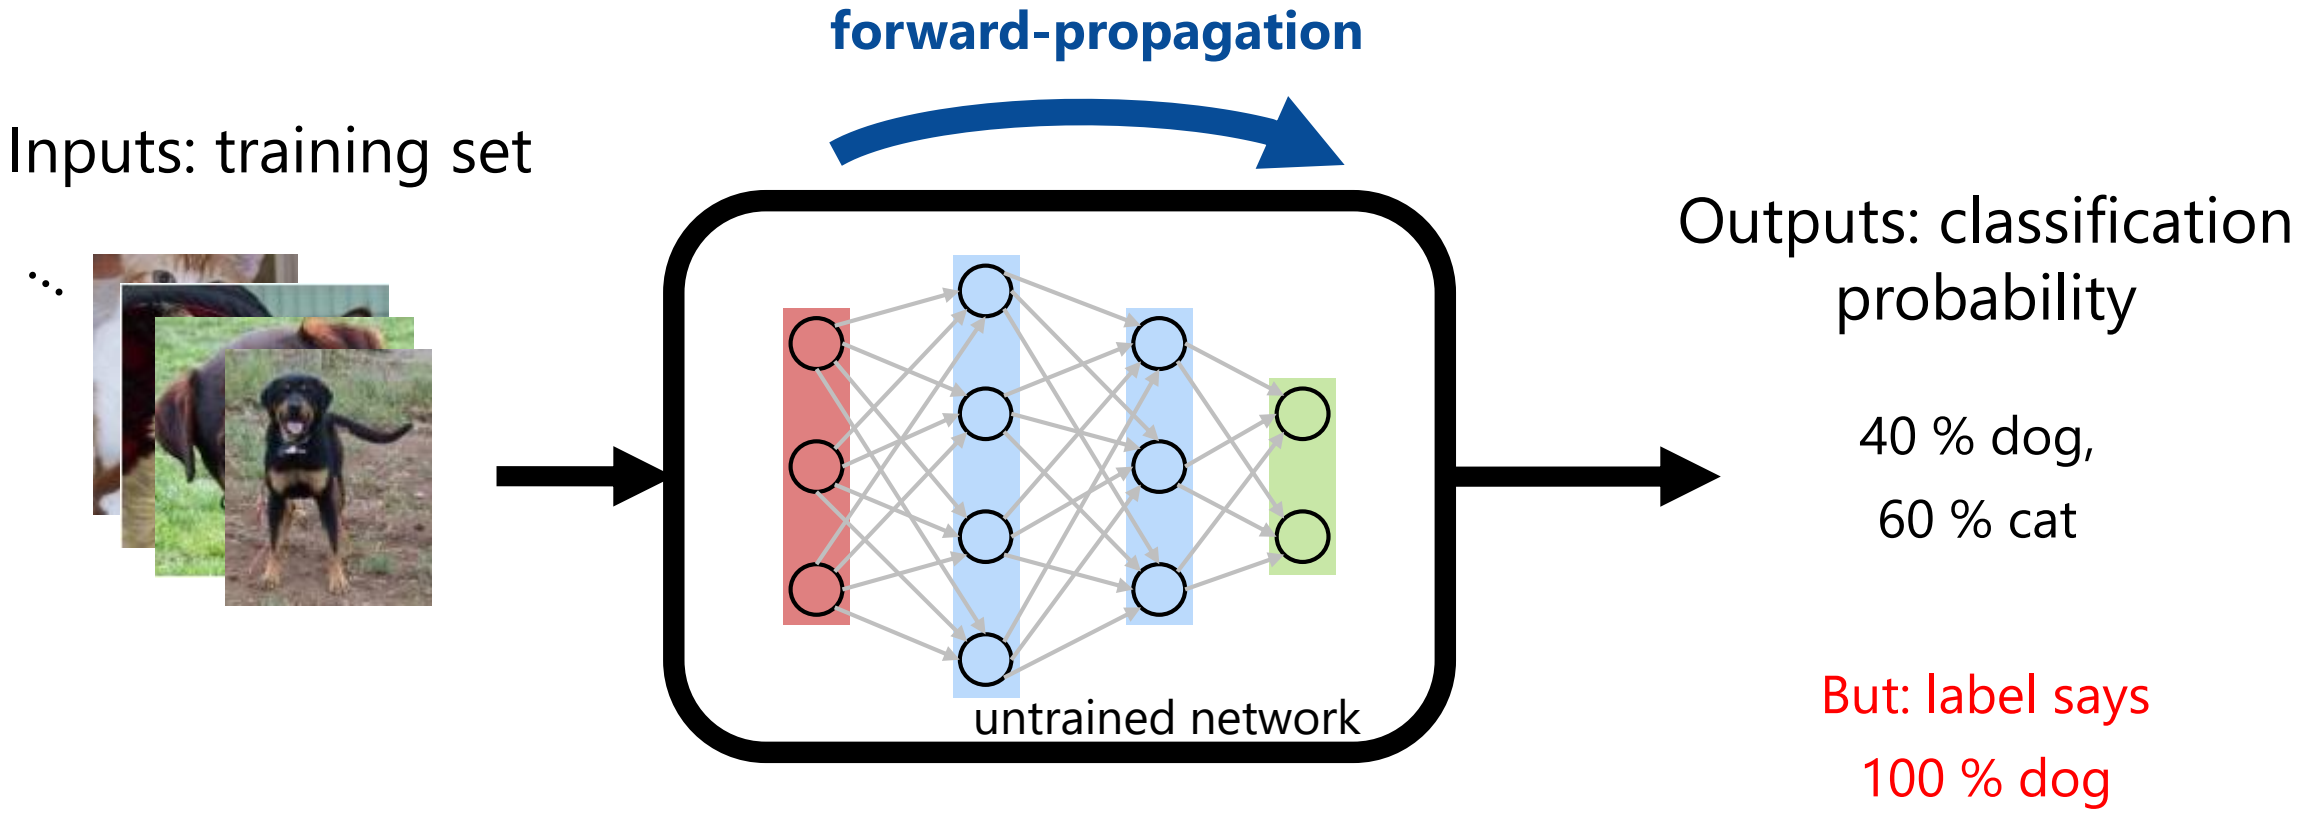

Machine learning concepts: training phase

- Goal: obtain trained weights

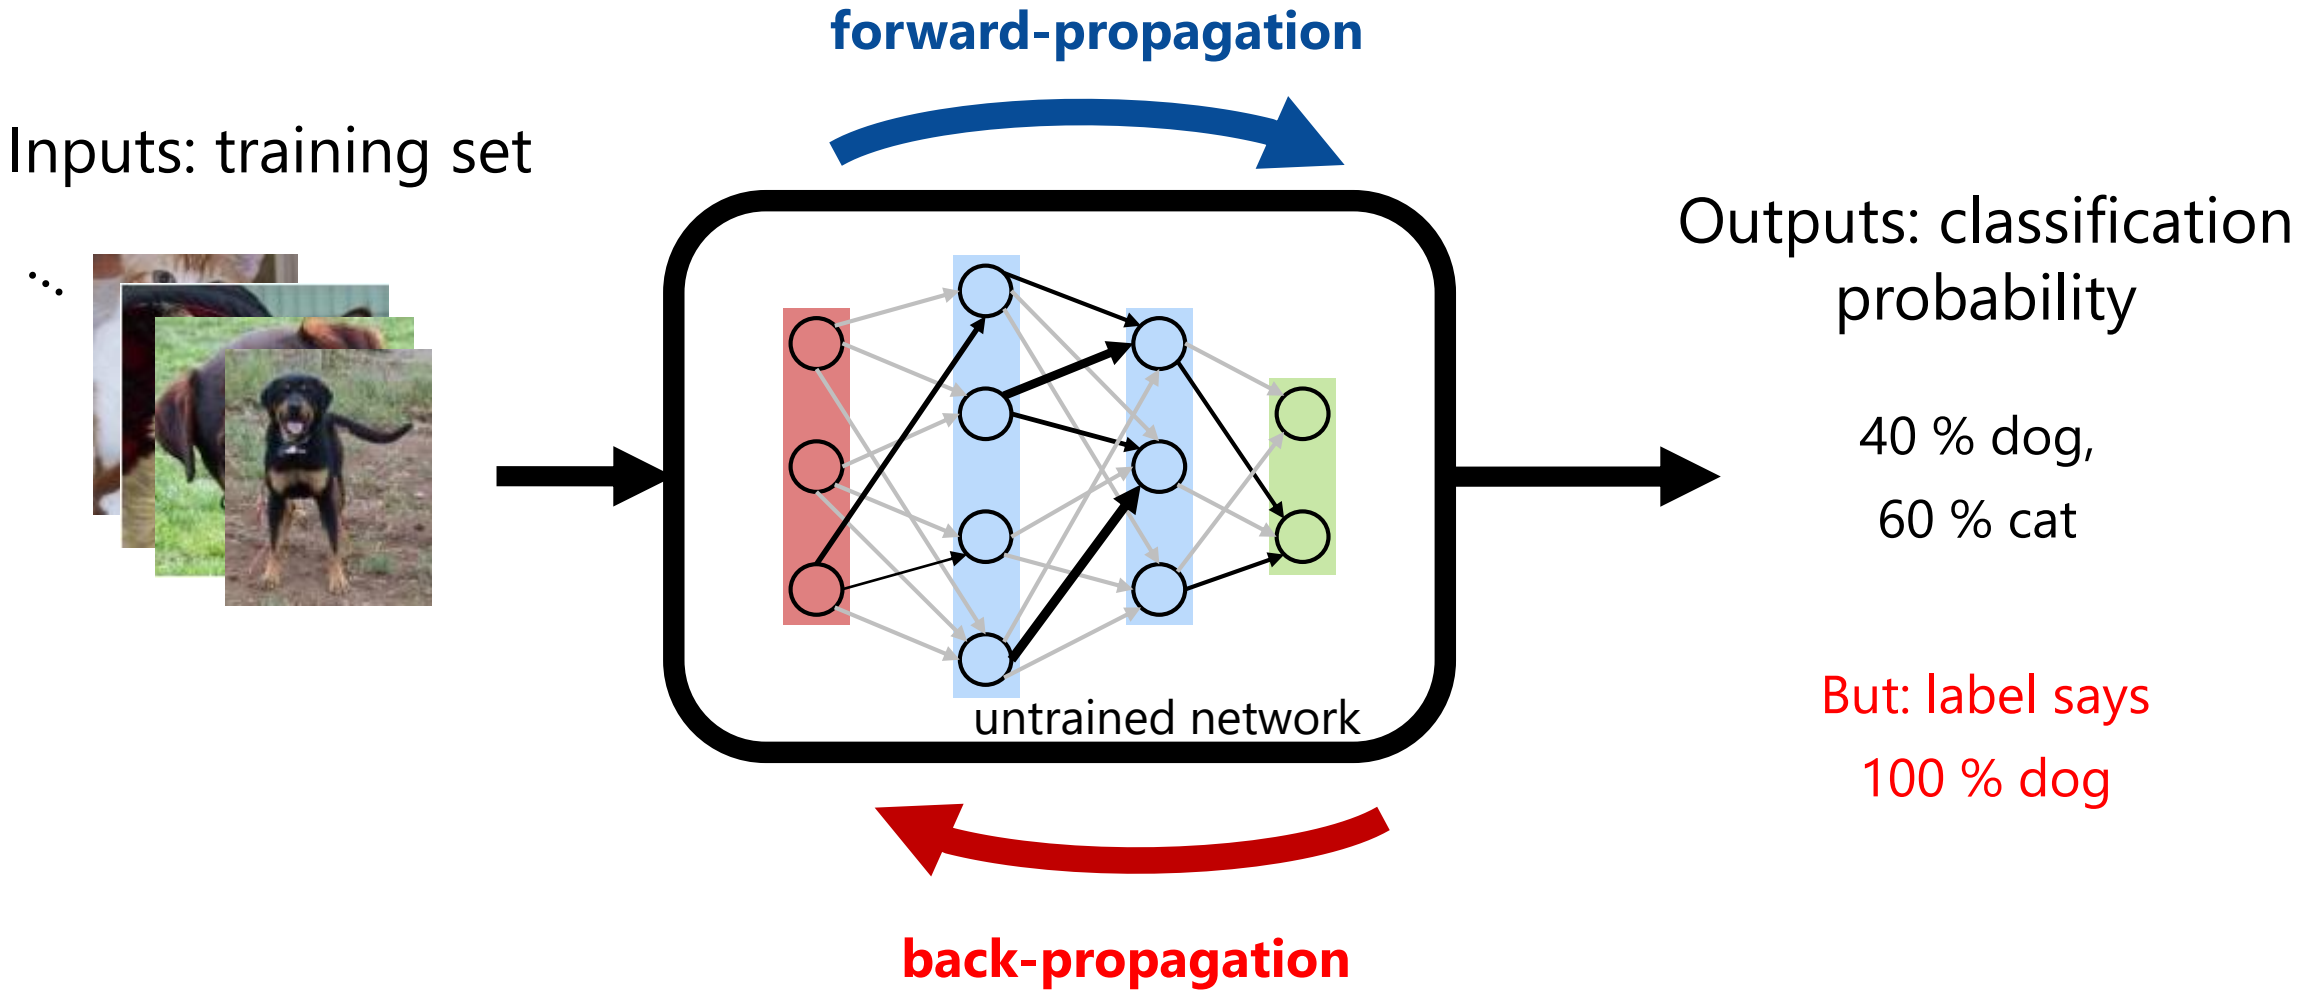

Machine learning concepts: training phase

- Goal: obtain trained weights

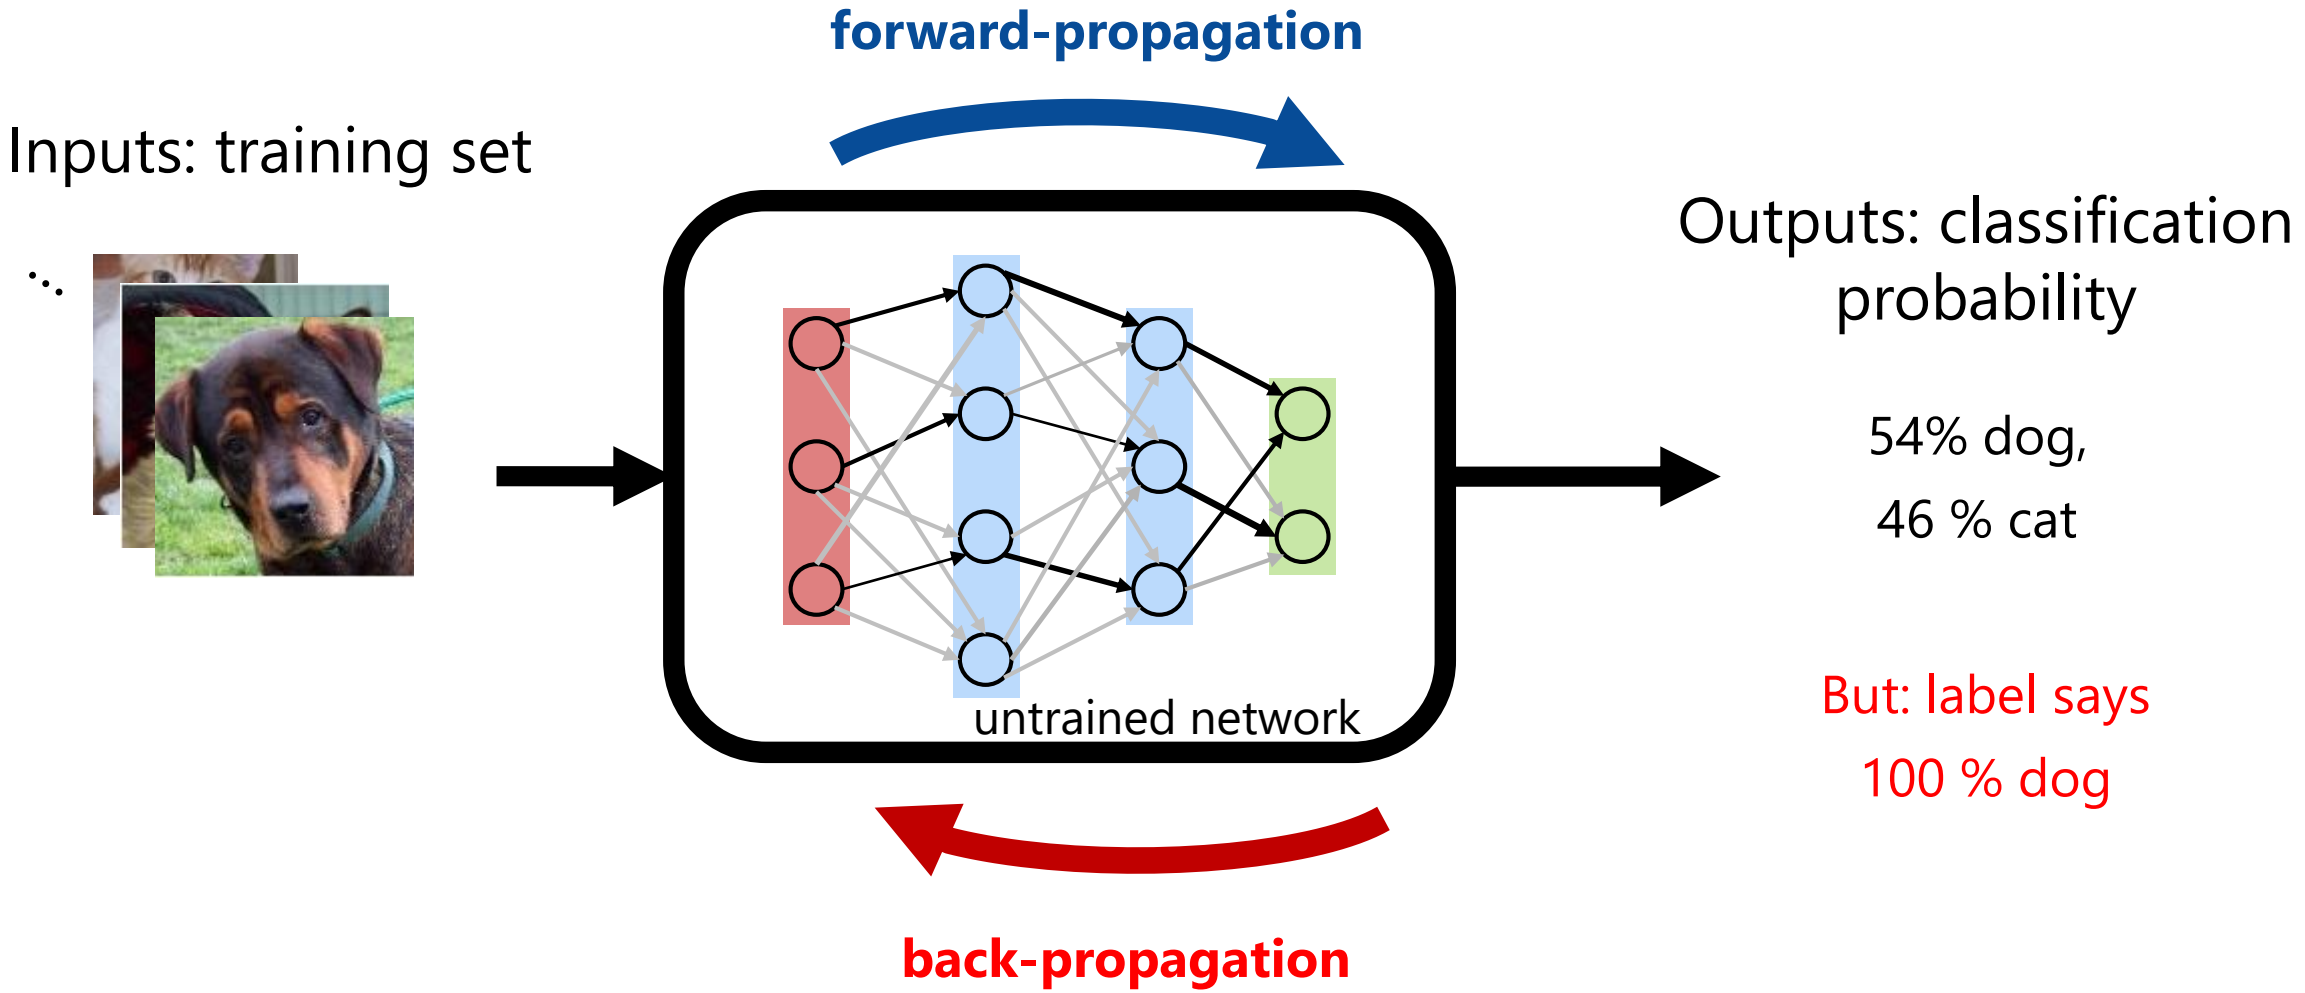

Machine learning concepts: inference

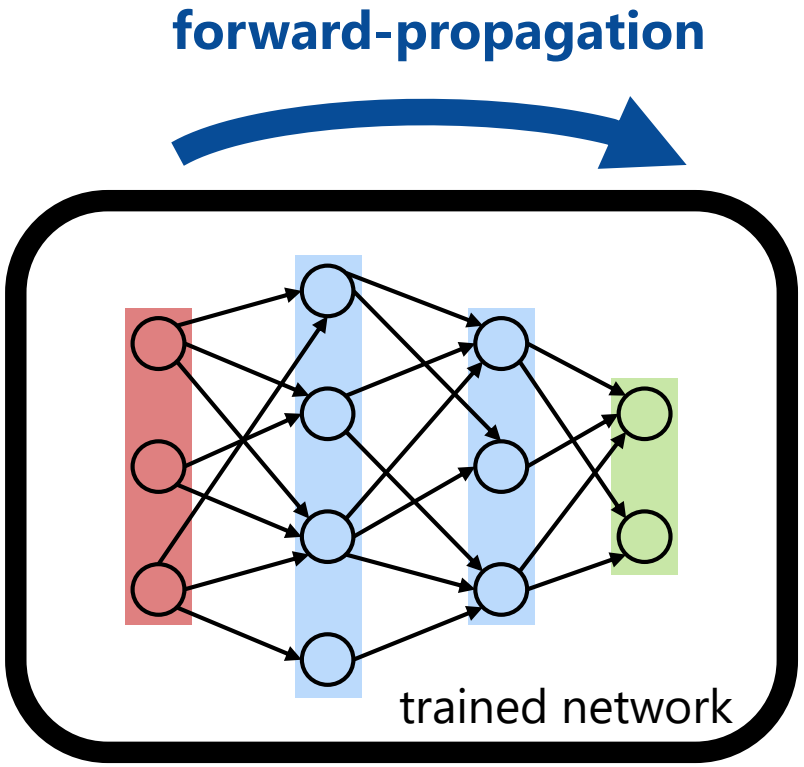

Machine learning concepts: inference

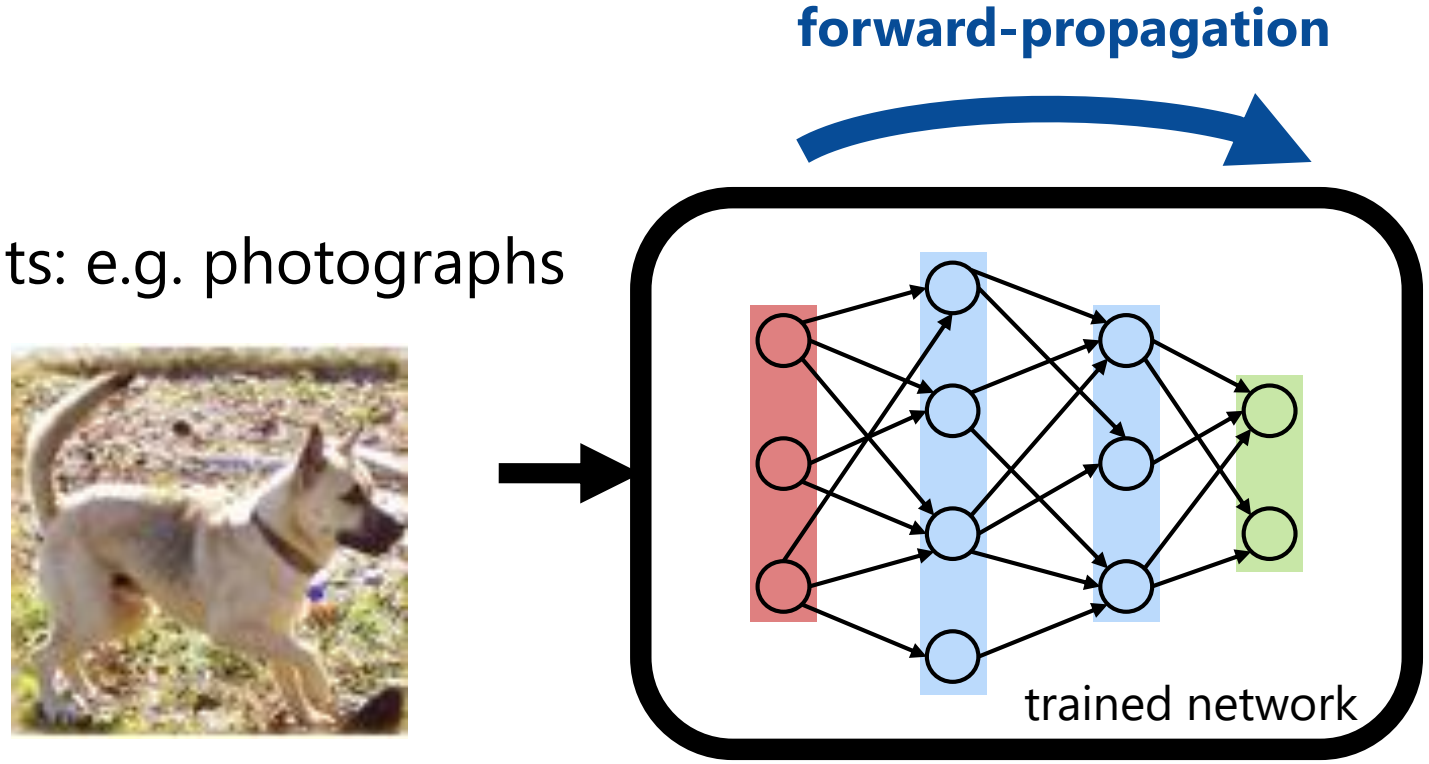

Machine learning concepts: inference

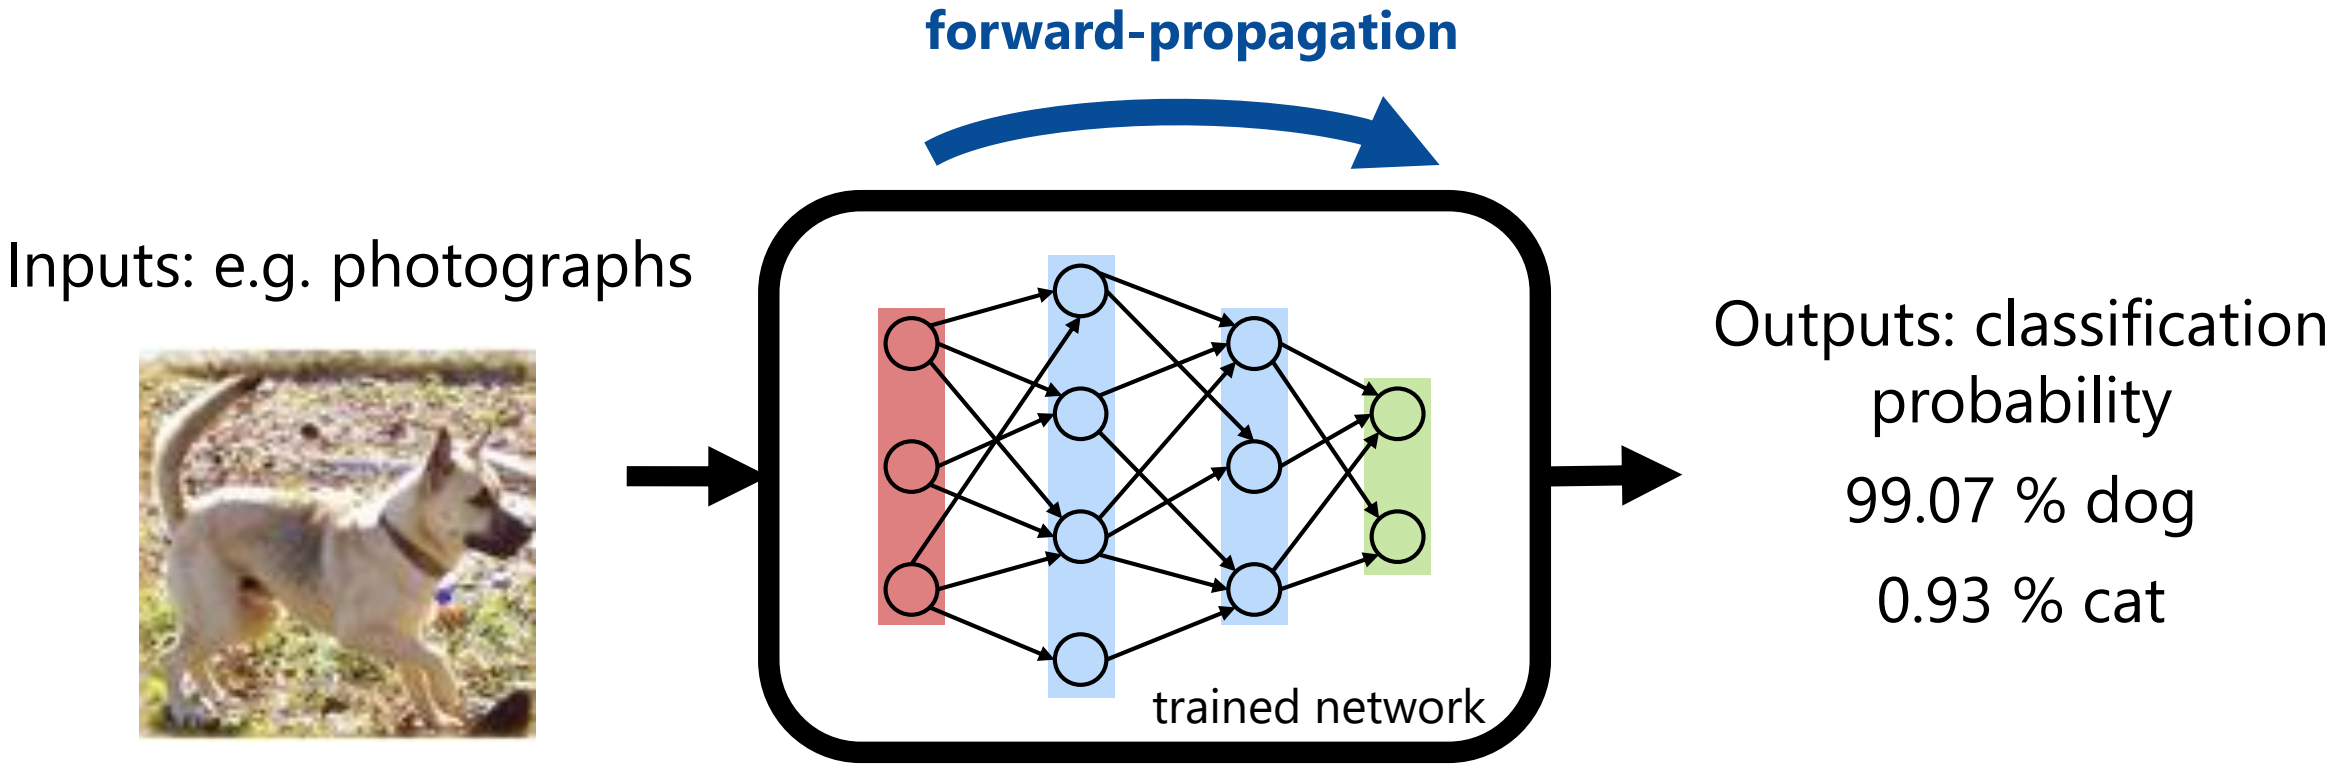

Machine learning concepts: feature extraction

Original image

Machine learning concepts: feature extraction

Original image

5x5 input image

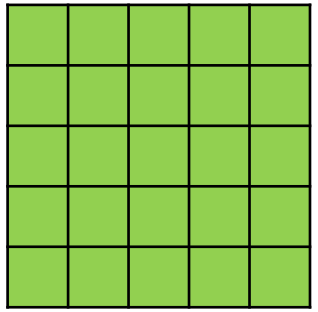

3x3 filter

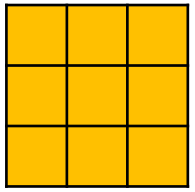

Machine learning concepts: feature extraction

Original image

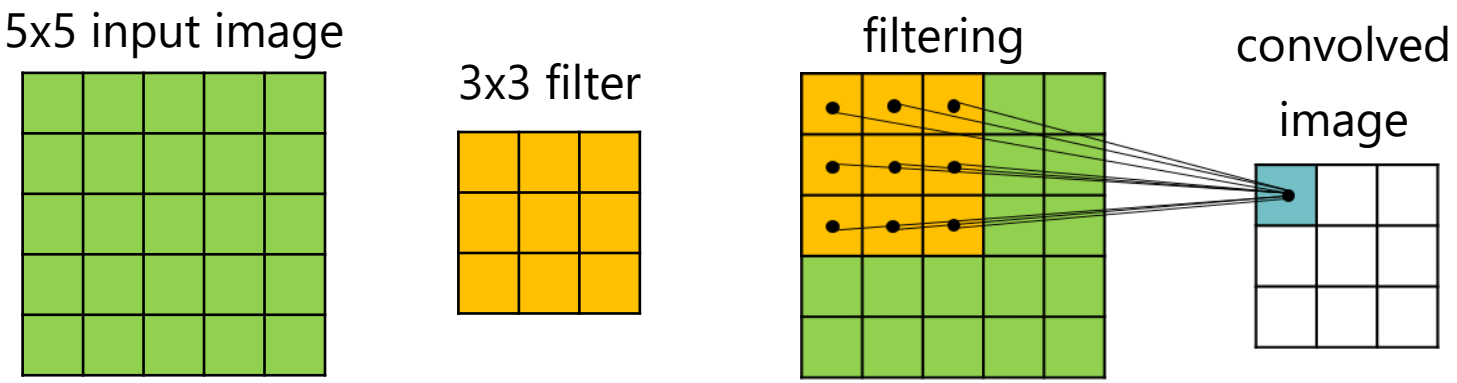

Machine learning concepts: feature extraction

Original image

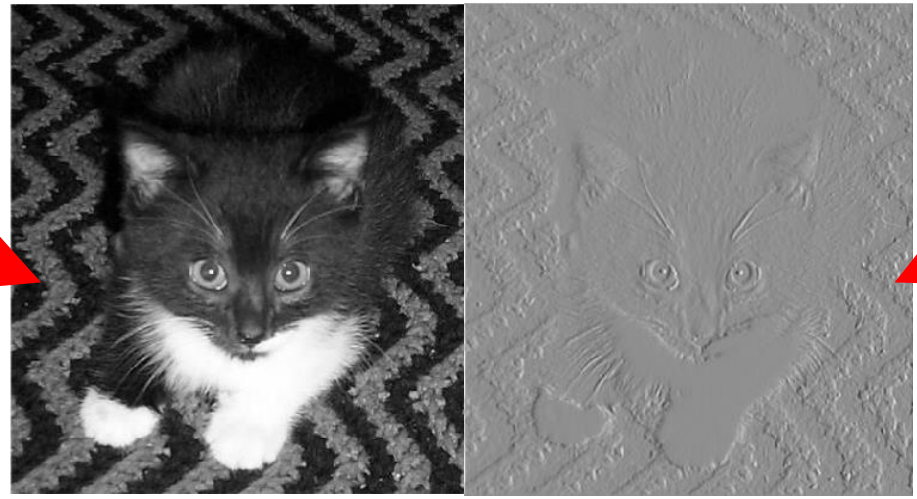

Convolved image

→ e.g. vertical edge detection

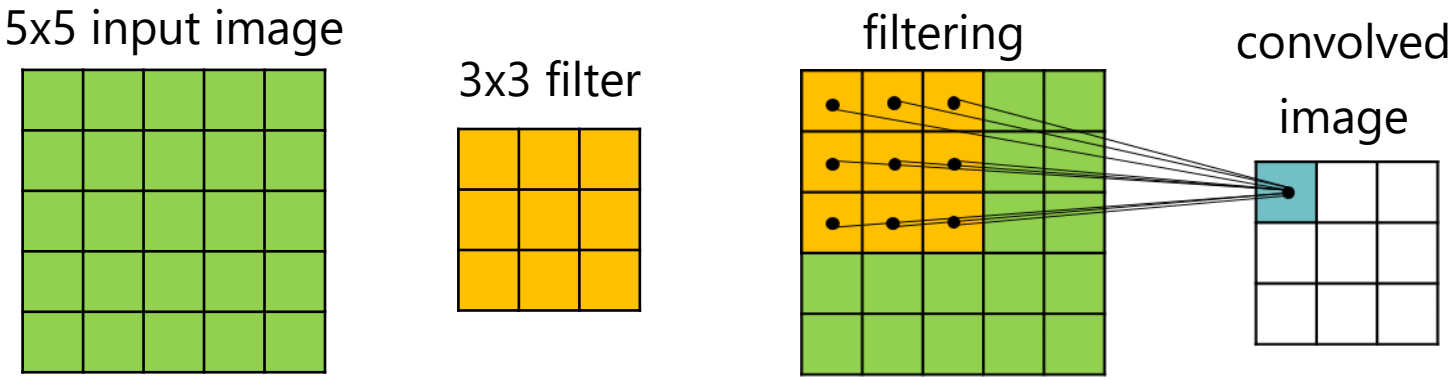

Machine learning concepts: feature extraction

Original image

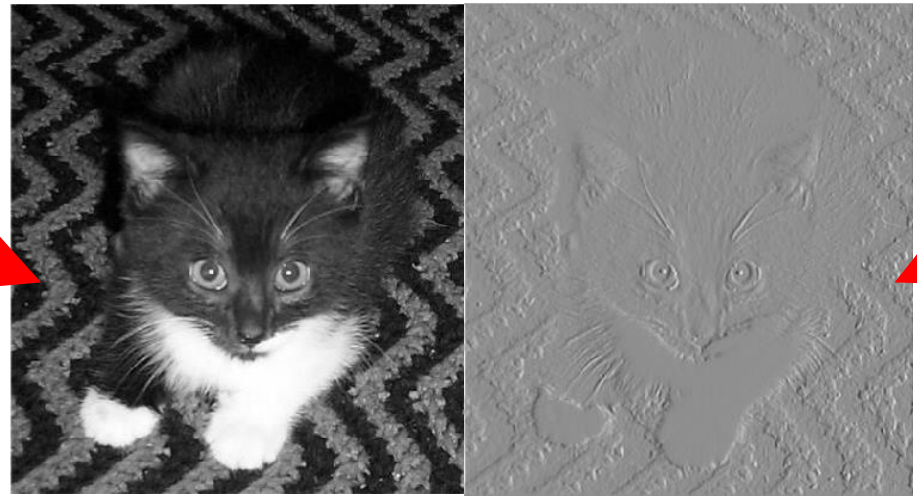

Convolved image

→ e.g. vertical edge detection

5x5 input image

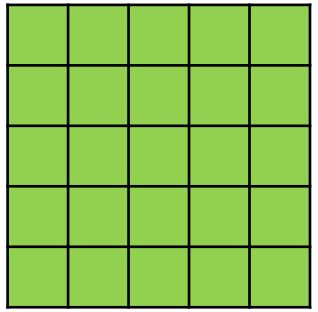

3x3 filter

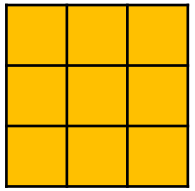

filtering

convolved image

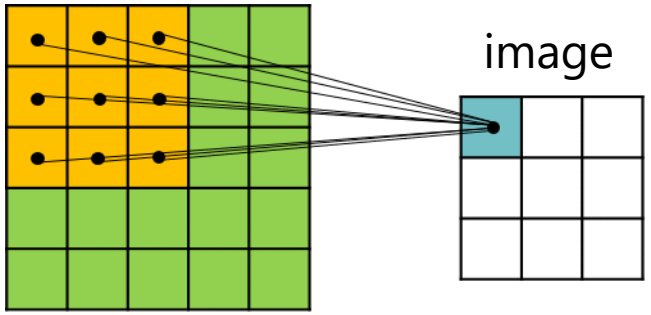

Neural network input layer

5x5 input image

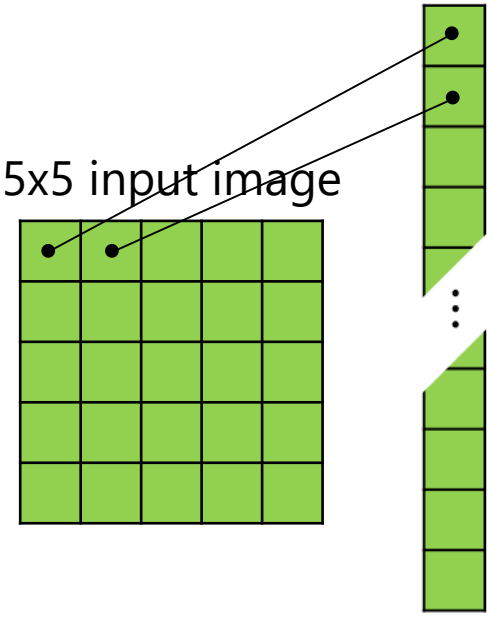

Why “deep” learning?

Why “deep” learning?

- First artificial neural networks were able to solve relatively simple tasks

Why “deep” learning?

- First artificial neural networks were able to solve relatively simple tasks
  - E. g. handwritten character recognition

Why “deep” learning?

- First artificial neural networks were able to solve relatively simple tasks
  - E. g. handwritten character recognition
- Next step: solve more demanding tasks like …

Why “deep” learning?

- First artificial neural networks were able to solve relatively simple tasks
  - E. g. handwritten character recognition
- Next step: solve more demanding tasks like …
  - Coloured image classification with lots of output classes
  - Real-time object detection
  - Pose estimation
  - Segmentation
  - Etc.

Why “deep” learning?

- First artificial neural networks were able to solve relatively simple tasks
  - E. g. handwritten character recognition
- Next step: solve more demanding tasks like …
  - Coloured image classification with lots of output classes
  - Real-time object detection
  - Pose estimation
  - Segmentation
  - Etc.

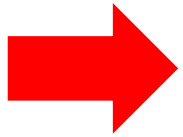

How do we do this?

Why “deep” learning?

- First artificial neural networks were able to solve relatively simple tasks
  - E. g. handwritten character recognition
- Next step: solve more demanding tasks like …
  - Coloured image classification with lots of output classes
  - Real-time object detection
  - Pose estimation
  - Segmentation
  - Etc.

→ How do we do this?

How “deep” exactly?

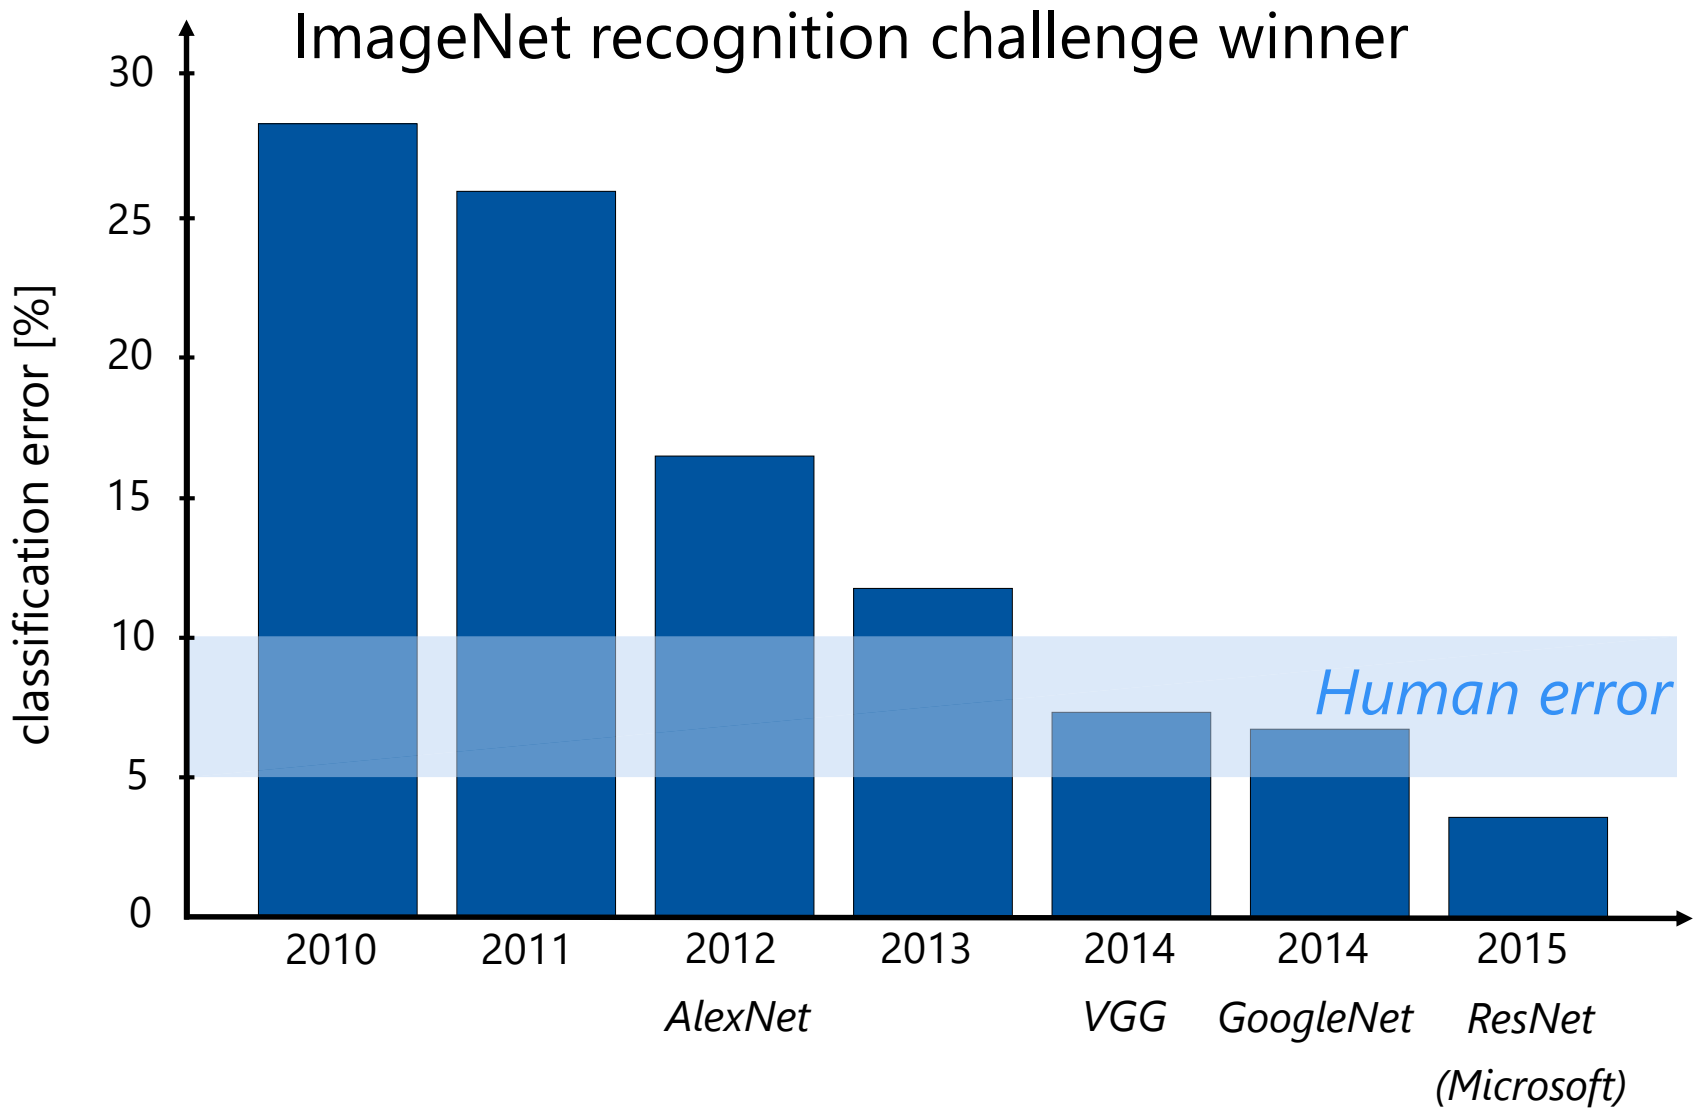

How “deep” exactly?

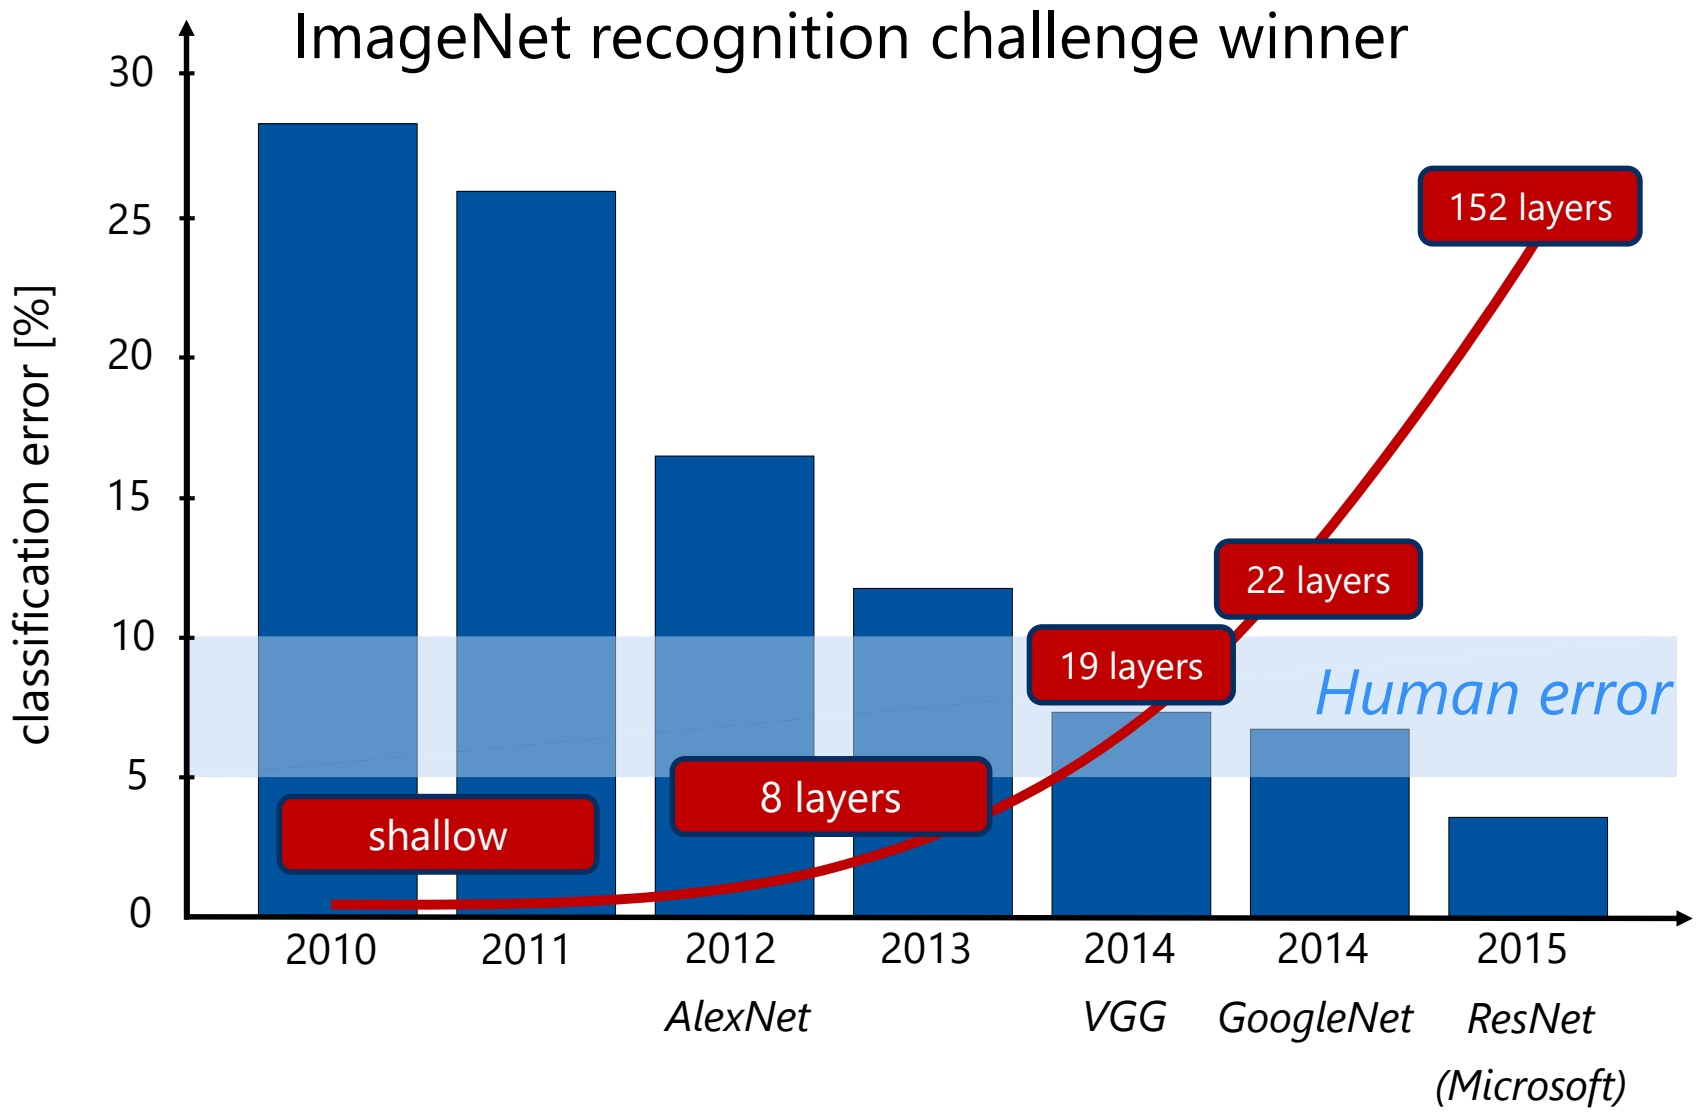

State-of-the-art: GoogleNet

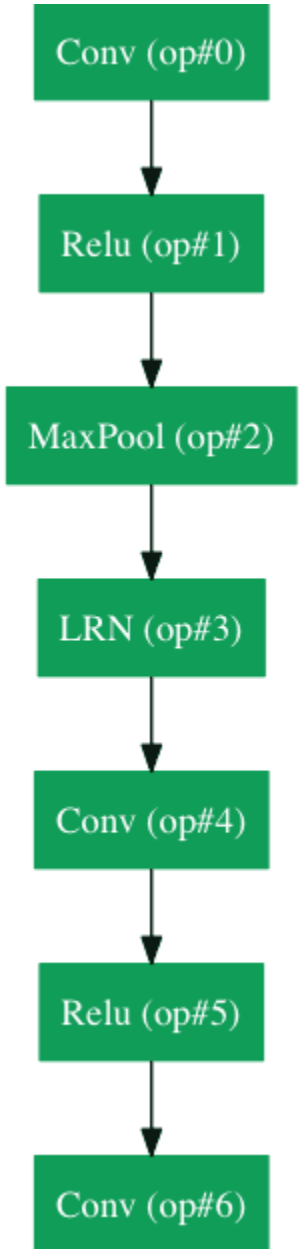

Summary

- Overview of machine learning and neural networks

Summary

- Overview of machine learning and neural networks
- Key concepts for current and future applications
  - Supervised learning
  - Feature extraction
  - Deep neural networks

Summary

- Overview of machine learning and neural networks
- Key concepts for current and future applications
  - Supervised learning
  - Feature extraction
  - Deep neural networks
- One question remains though…

Summary

- Overview of machine learning and neural networks
- Key concepts for current and future applications
  - Supervised learning
  - Feature extraction
  - Deep neural networks
- One question remains though…

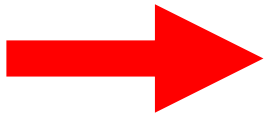

What hardware do we need for this?

CPUs, GPUs, FPGAs, ASICs??