Deep Learning og Computer Vision

Chris Holmberg Bahnsen

AALBORG UNIVERSITET

Baggrund
Hvem er jeg?

- Cand. polyt. Elektronik & IT, 2013
- Ph.d.-afhandling i robust trafikovervågning, 2018
- Visual Analysis of People Laboratory, AAU
- Postdoc 2019 -

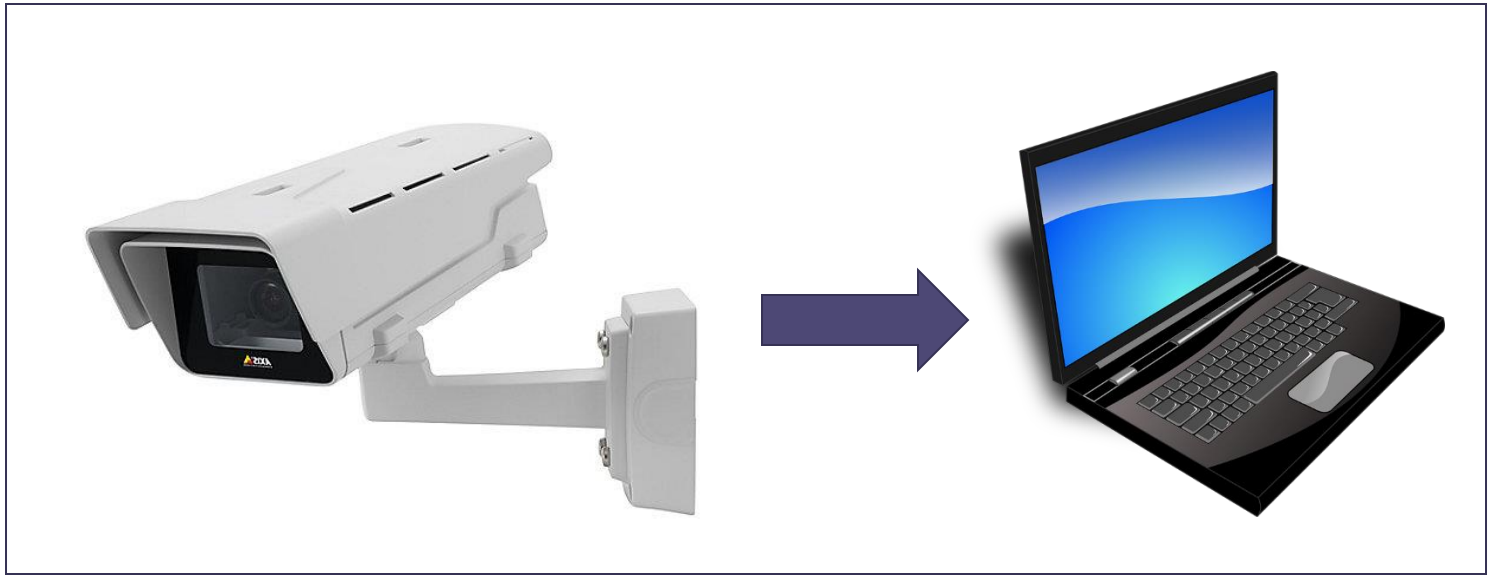

Hvad er AI?

- AI: ~1940s
- Machine Learning: ~1990s
- Deep Learning: ~2010s

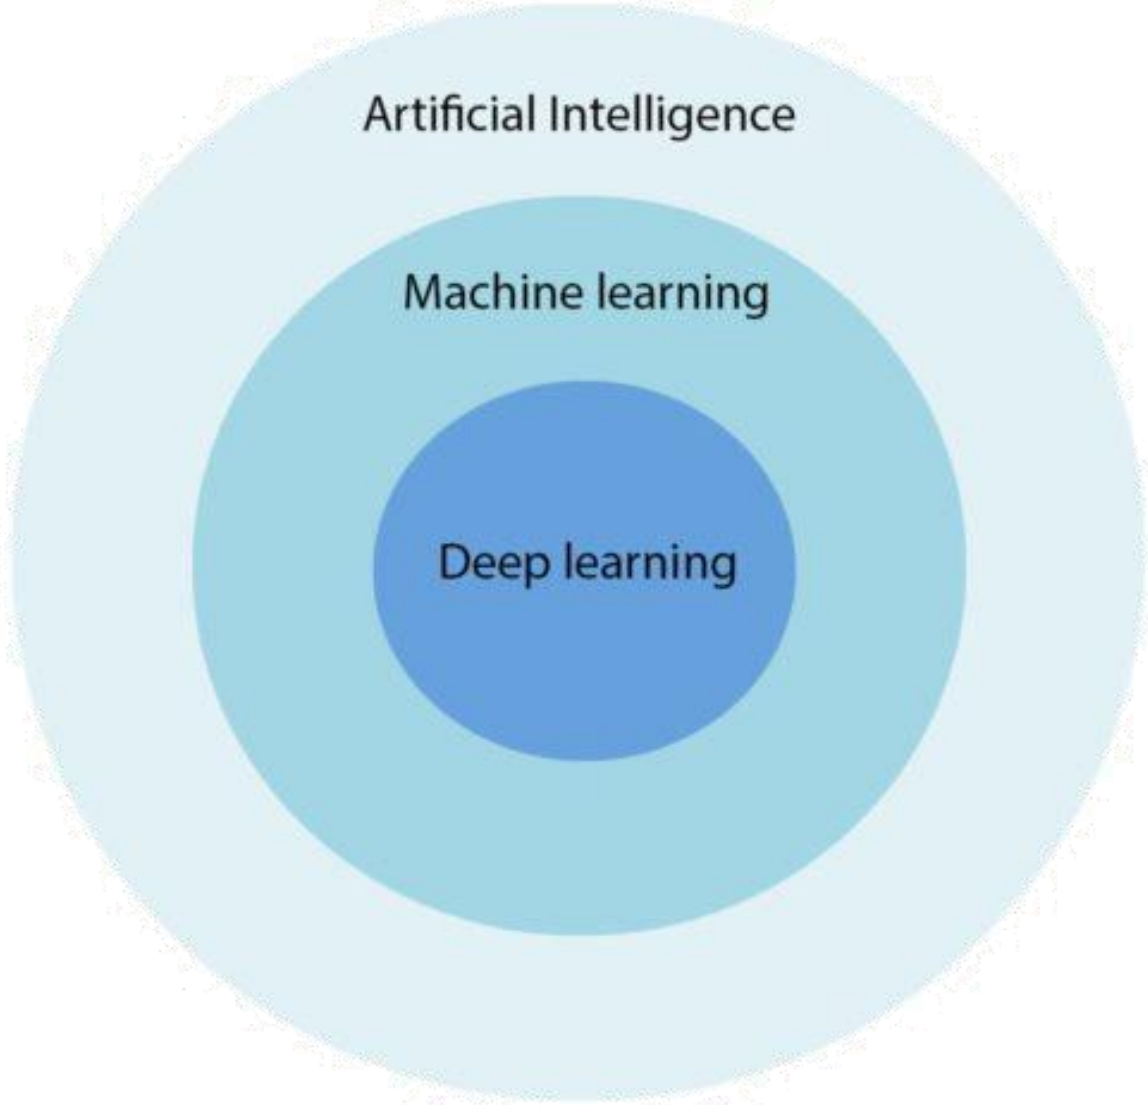

Hvad er AI?

- Opgave: Genkend en kat
- Løsning: Regelbaseret model af kattens dele

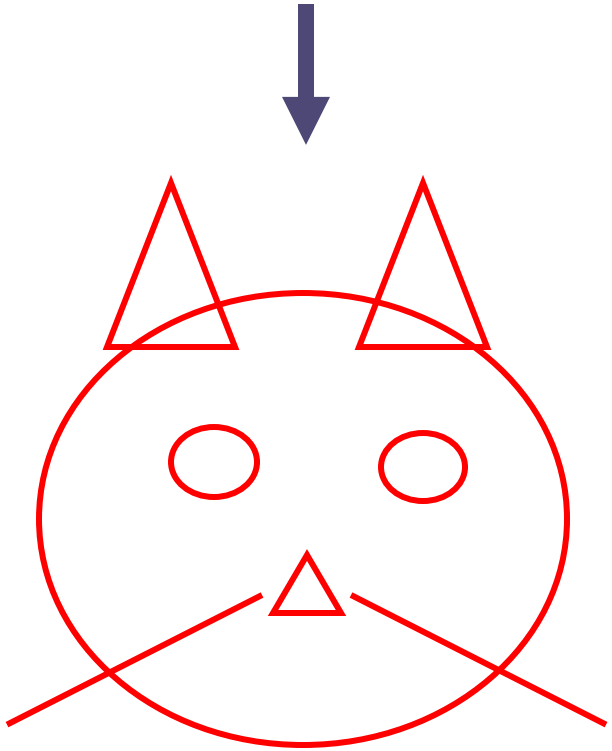

Hvad er AI?

- Opgave: Genkend en kat
- Løsning: Regelbaseret model af kattens dele

”Kat”

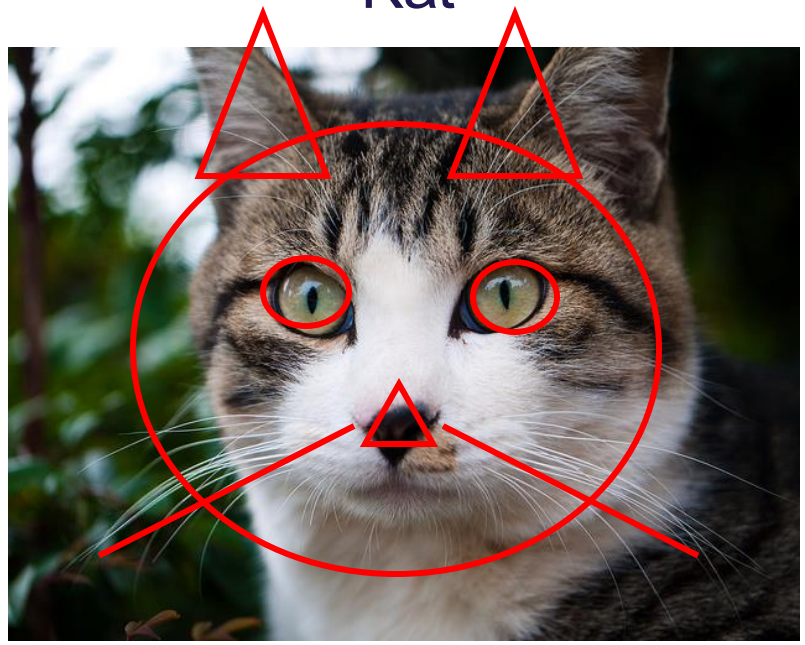

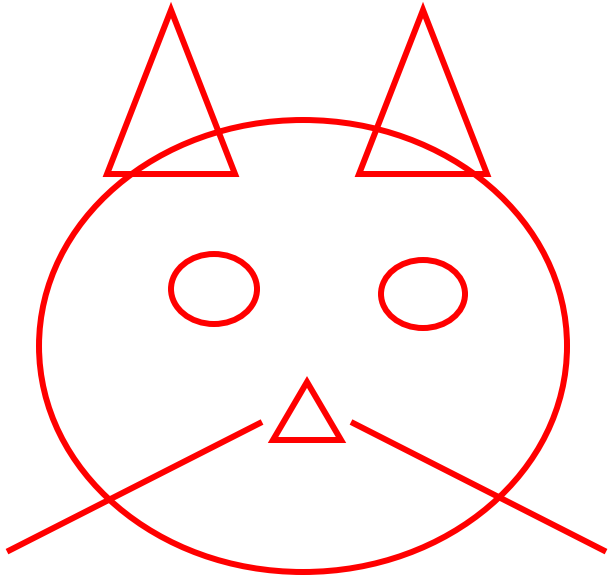

Hvad er AI?

- Duer ikke ved forskellige positurer
- Katte er utroligt deformerbare

?

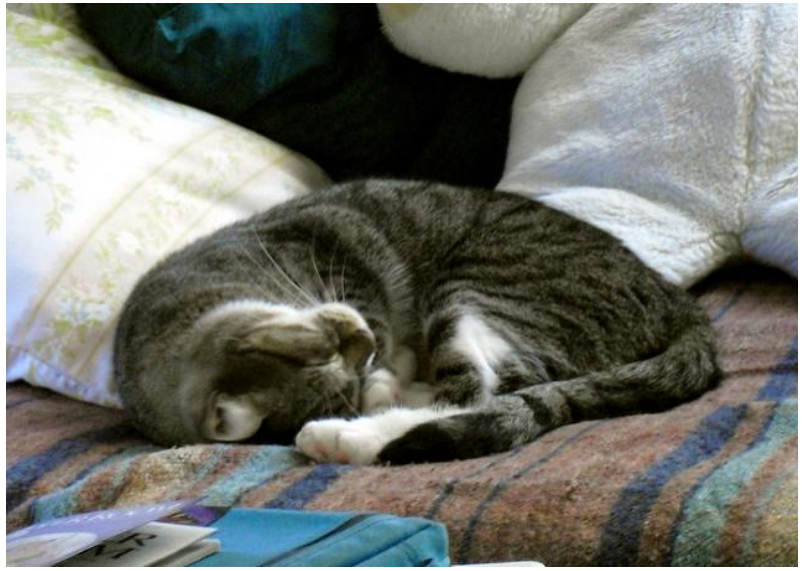

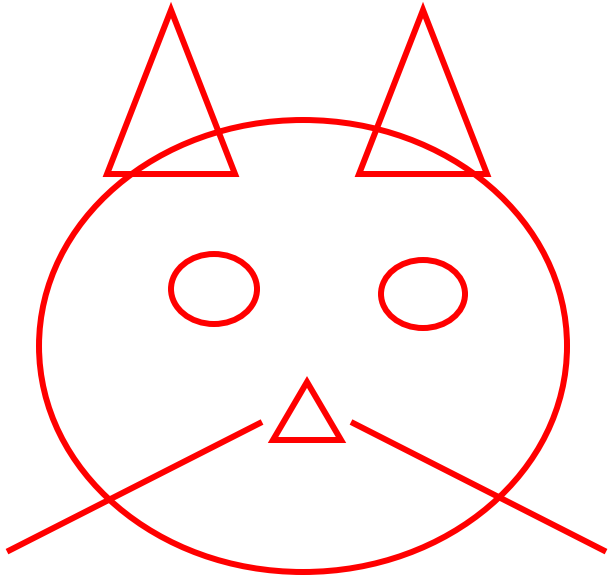

Hvad er AI?

- Umuligt at definere en god model for alle kattedeformationer
- Vi vil i stedet *lære* vores model i stedet for at hårdkode den

Hvad er AI?

- AI: ~1940s
- Machine Learning: ~1990s
- Deep Learning: ~2010s

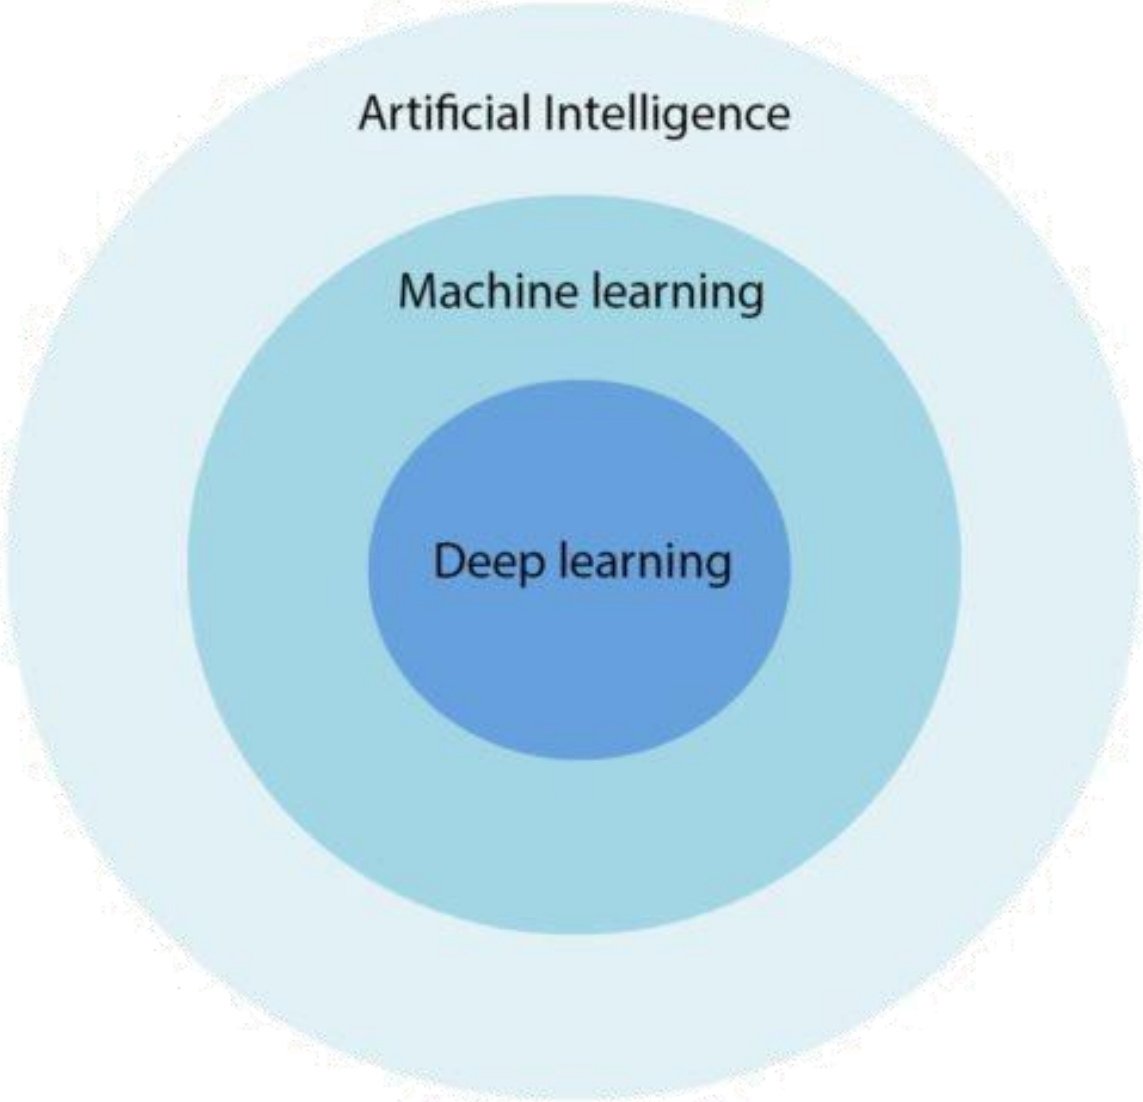

Hvad er machine learning?

- Eksempel: Kend forskel på Per and Lene
- Fremgangsmåde: Find karakteristiske features

”Per”

”Lene”

Hvad er machine learning?

- Eksempel: Kend forskel på Per and Lene
- Mængden af hår?

”Per”

”Lene”

Hår-ratio	0,09	0,24

Hvad er machine learning?

- Eksempel: Kend forskel på Per and Lene
- Mængden af hår?
- Ansigtets cirkularitet?

”Per”

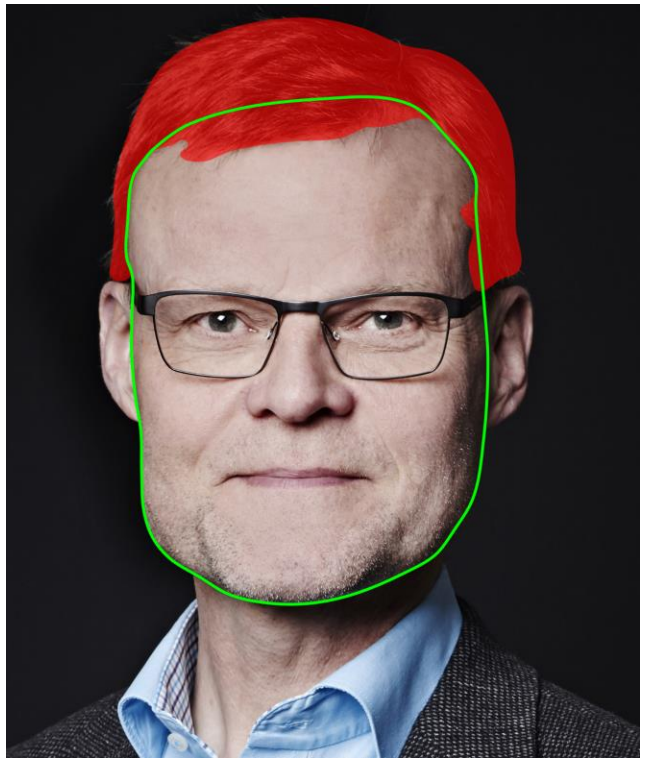

”Lene”

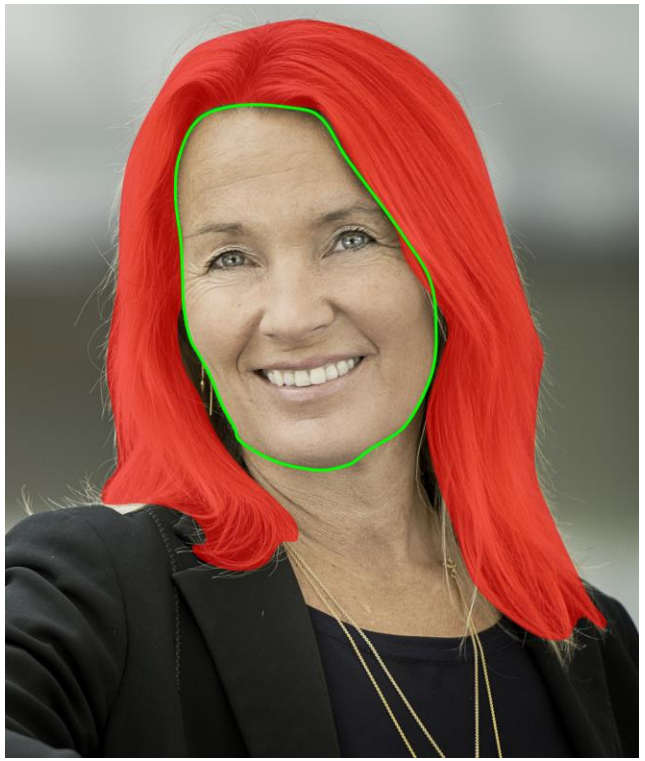

Hair-ratio	0,09	0,24
Circularity	0,79	0,87

Hvad er machine learning?

- Eksempel: Kend forskel på Per and Lene
- Mængden af hår?
- Ansigtets cirkularitet?
- Feature-rum

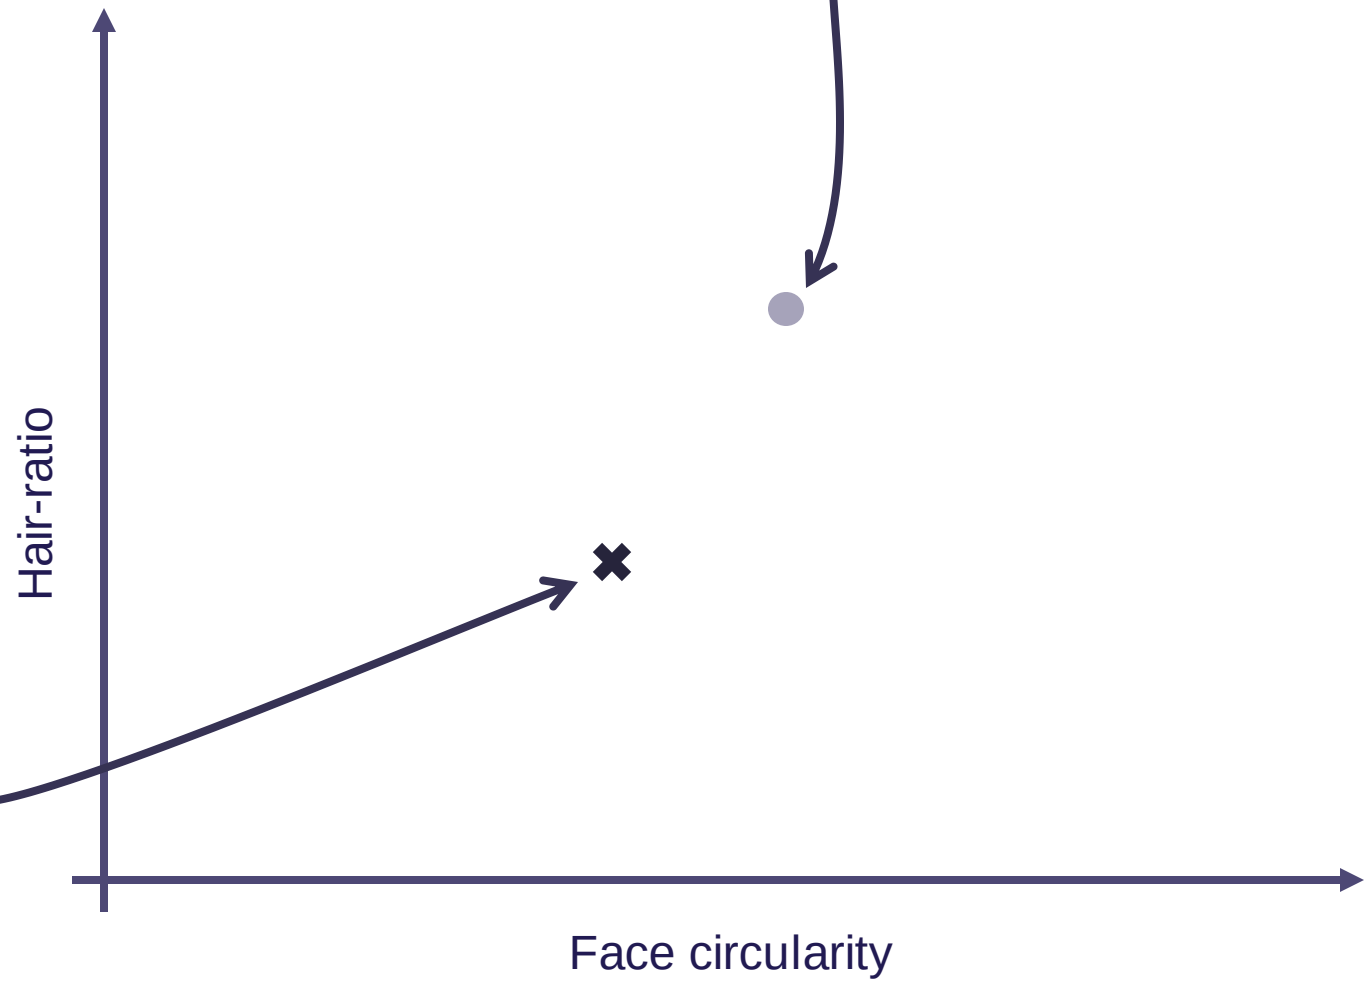

Hvad er machine learning?

- Eksempel: Kend forskel på Per and Lene
- Mængden af hår?
- Ansigtets cirkularitet?
- Flere billeder

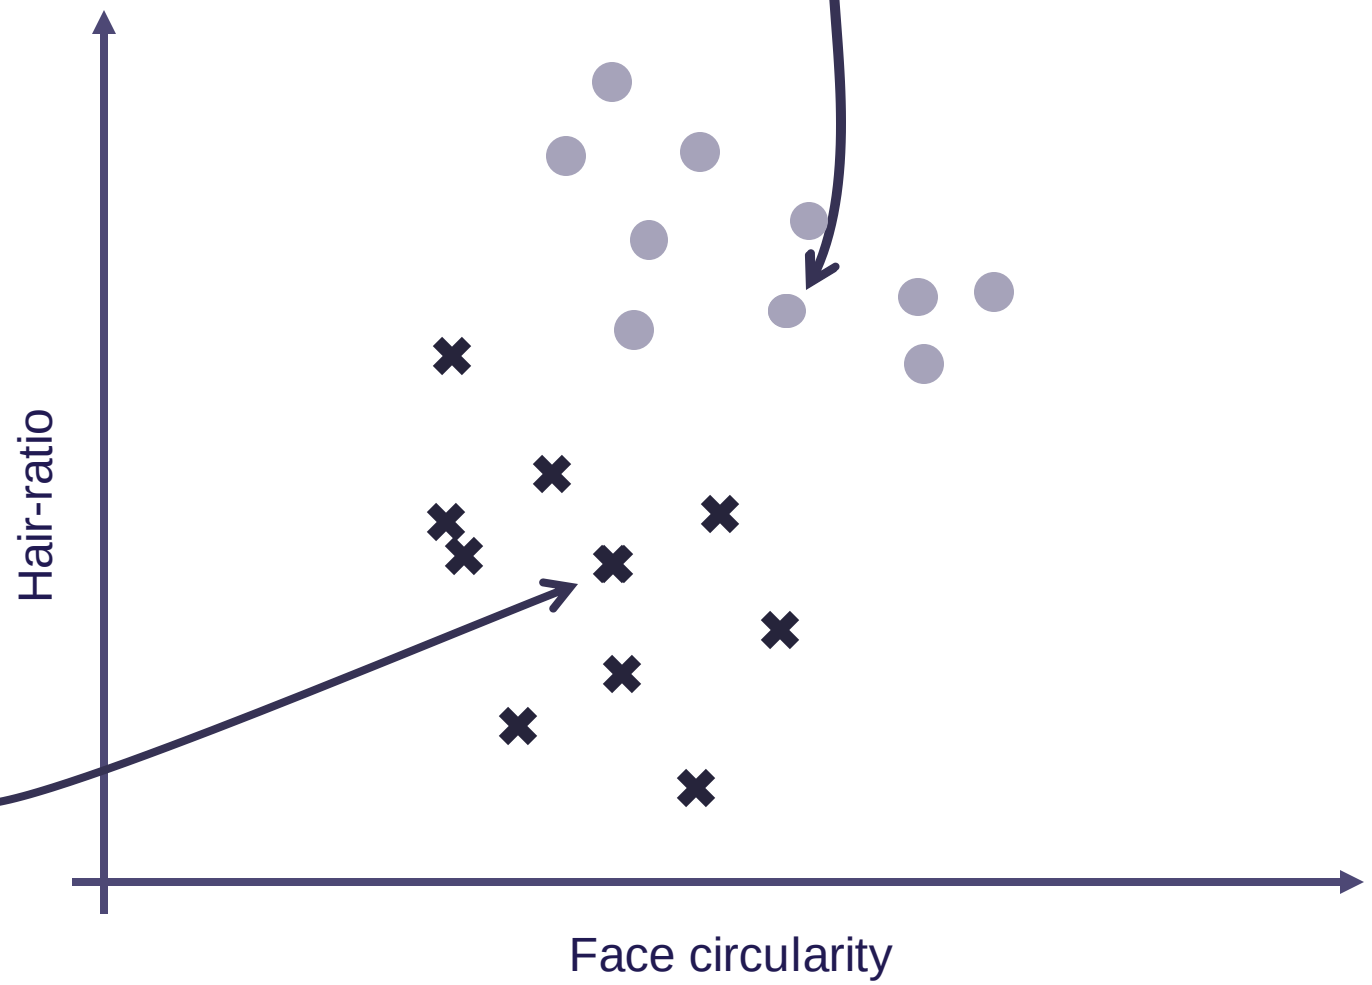

Hvad er machine learning?

- Eksempel: Kend forskel på Per and Lene
- Mængden af hår?
- Ansigtets cirkularitet?
- Find en linje der adskiller de to områder

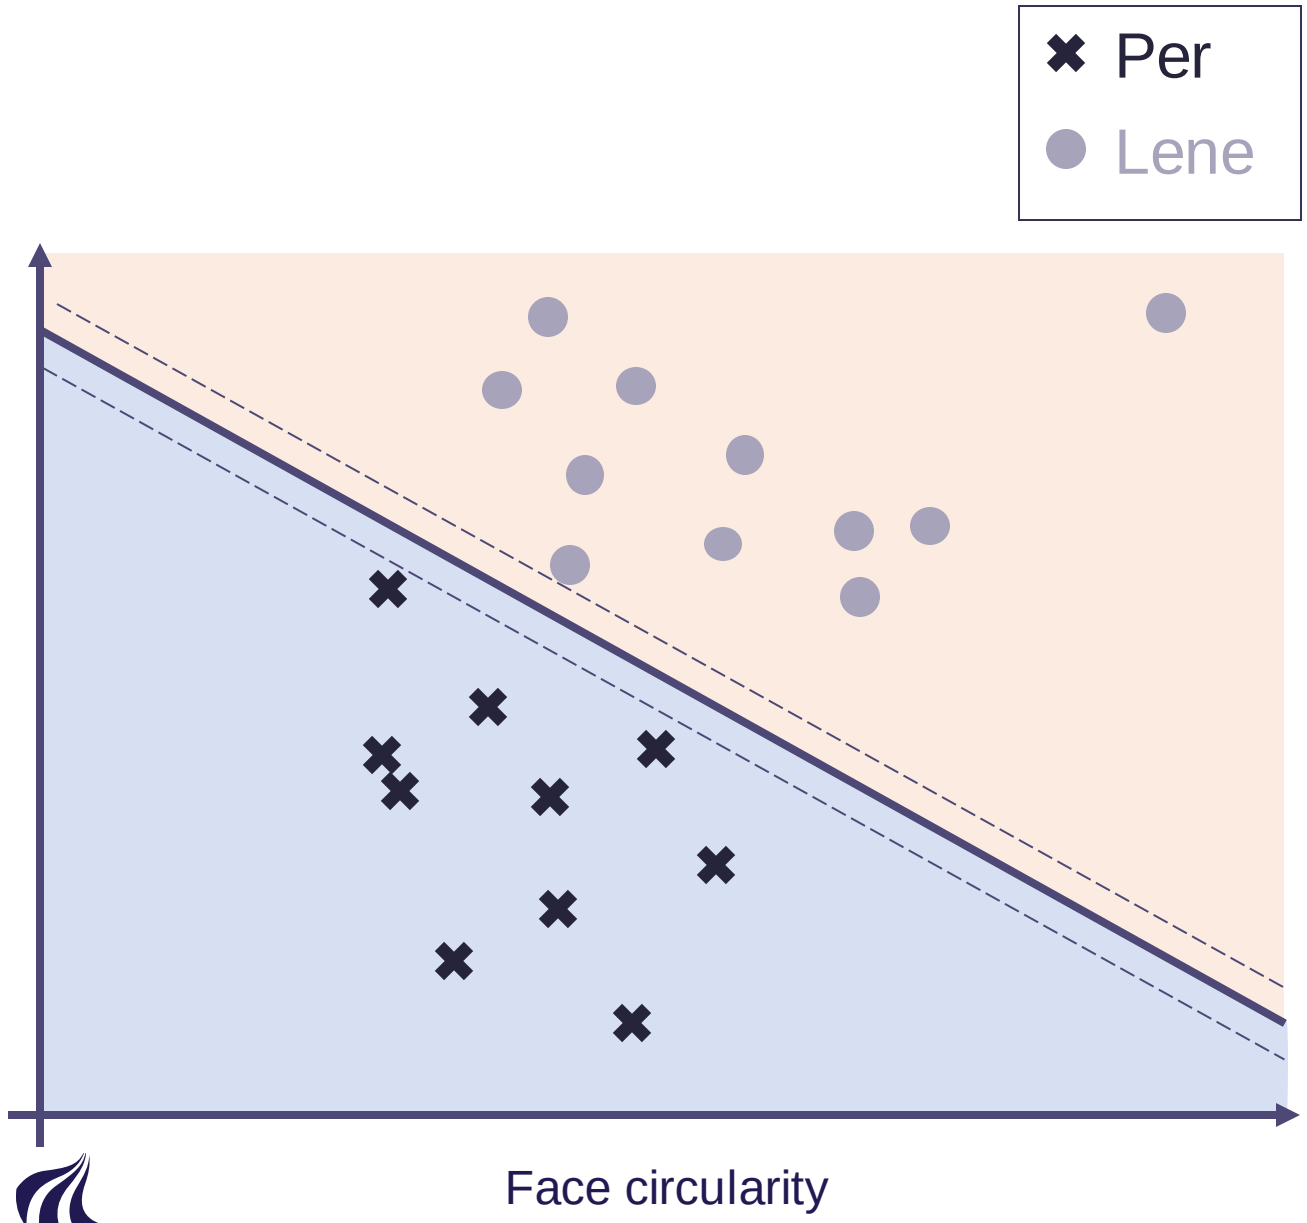

Hvad er machine learning?

- Nyt billede – hvem er det?
  - Udtræk features
  - Hvor er billedet i feature-rummet?
  - Sammenlign med linjen
  - Bestem hvem det er

- Det er machine learning

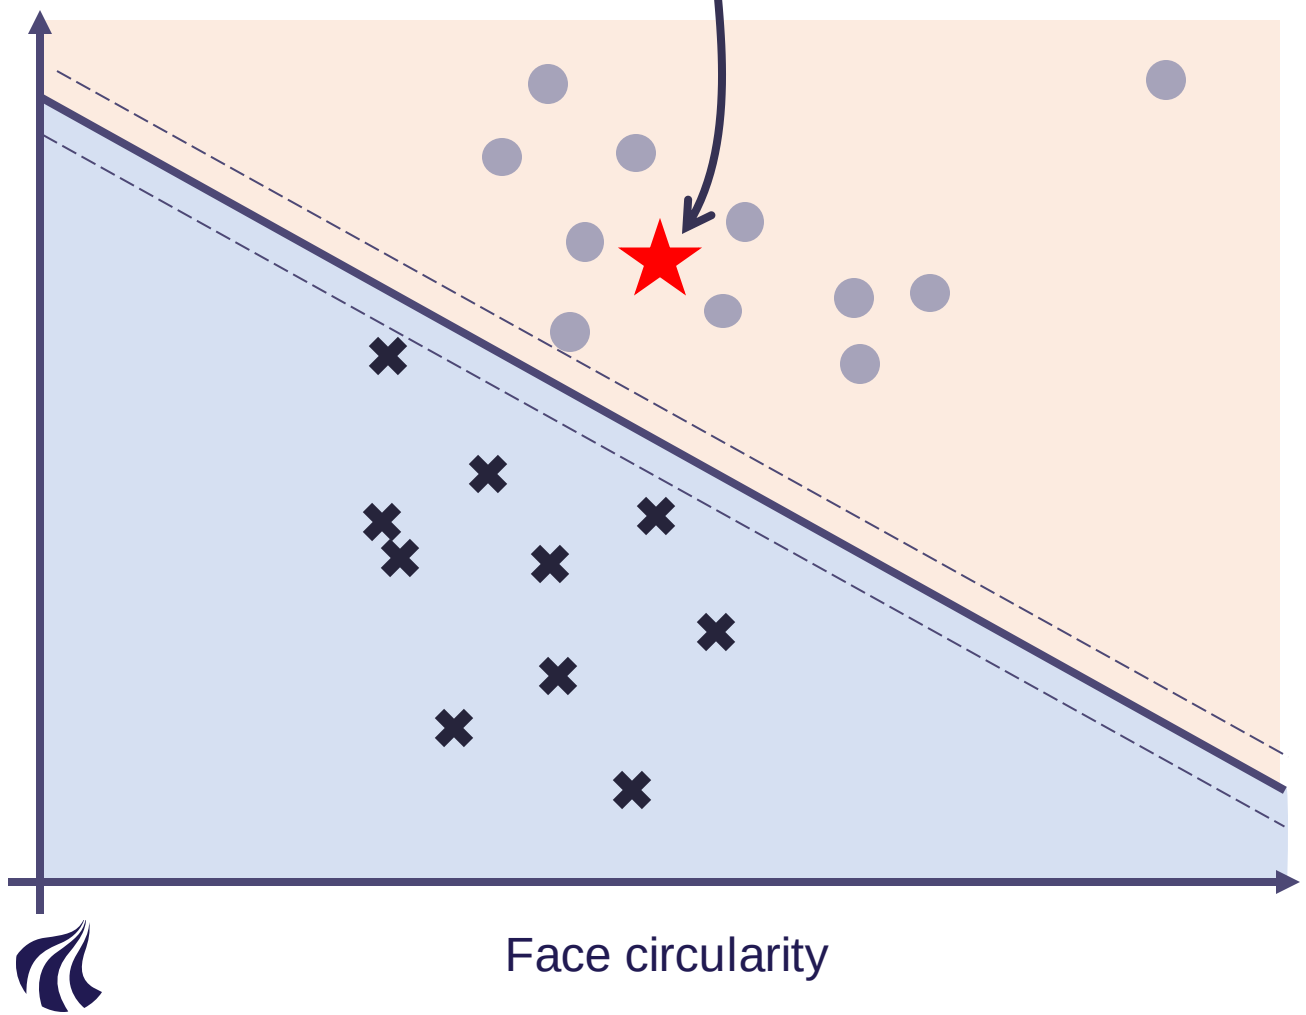

Hvad er AI?

- AI: ~1940s
- Machine Learning: ~1990s
- Deep Learning: ~2010s

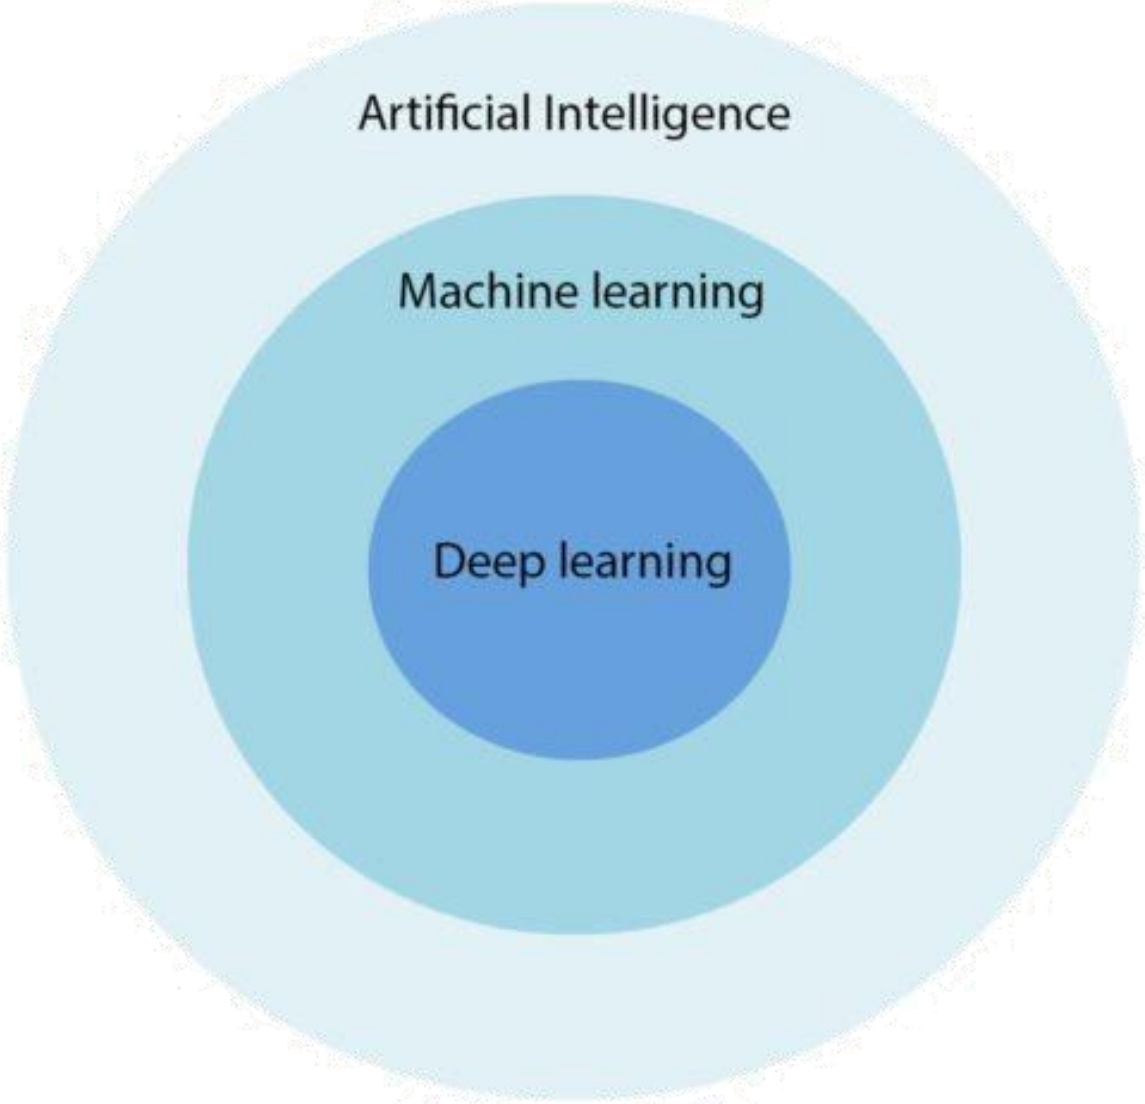

Deep Learning

Fra machine learning til deep learning

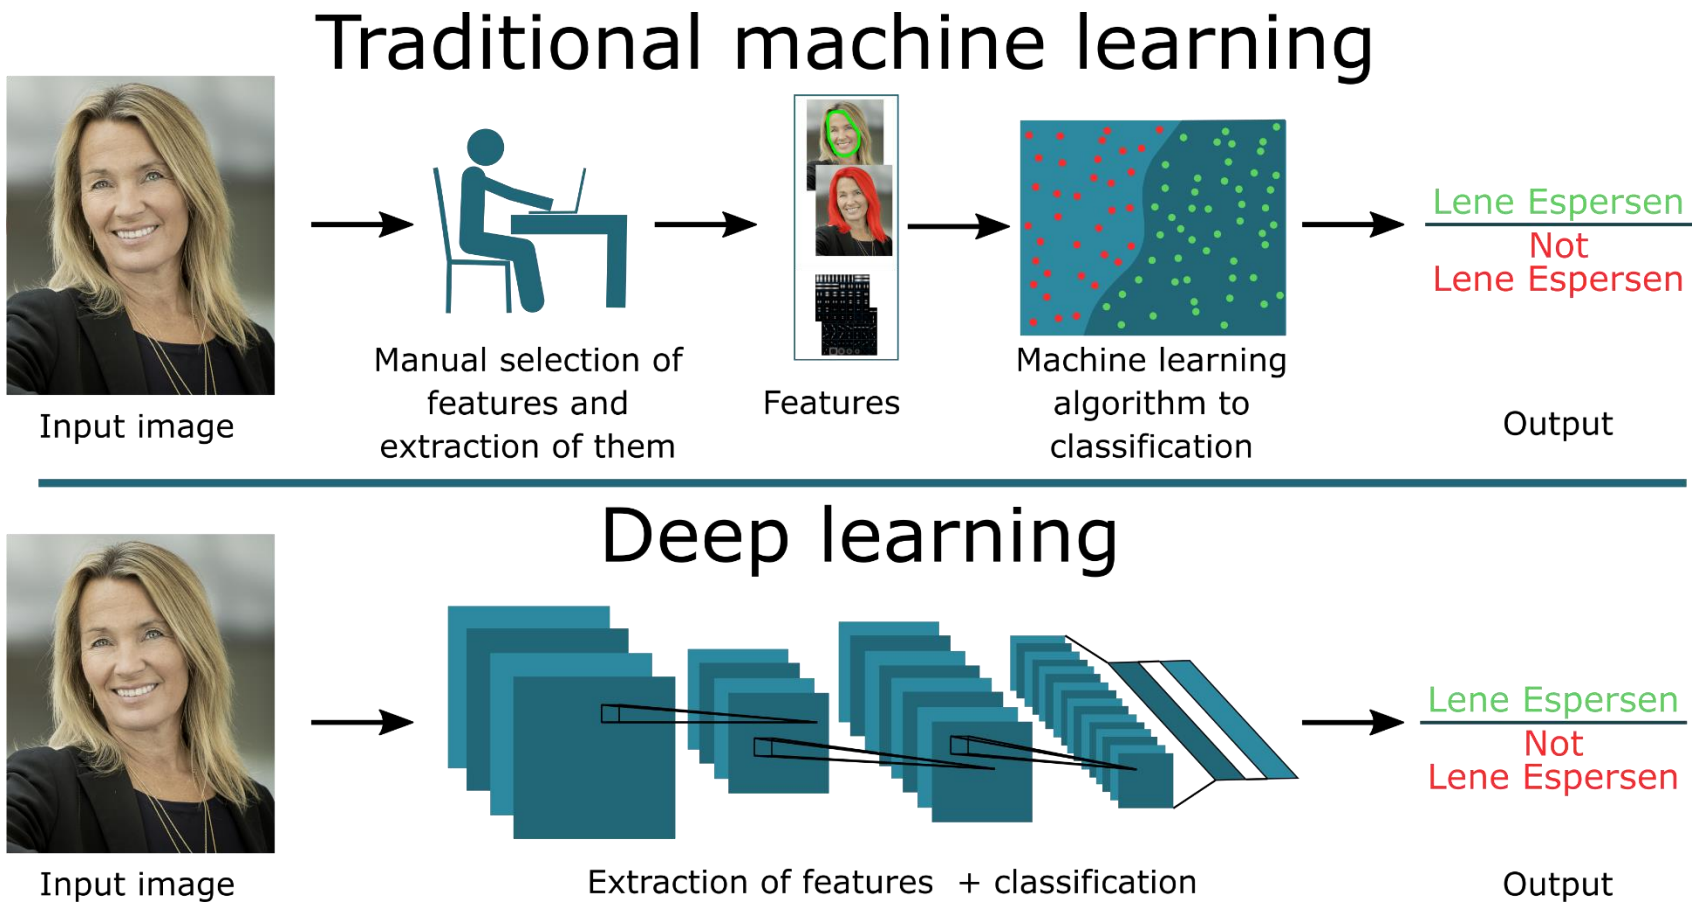

Deep Learning
Den menneskelige hjerne

- Vores hjerne indeholder ca. 100 milliarder neuroner, der er organiseret i et netværk
- Et neuralt netværk er en machine learning-metode, der er inspireret af den menneskelige hjerne

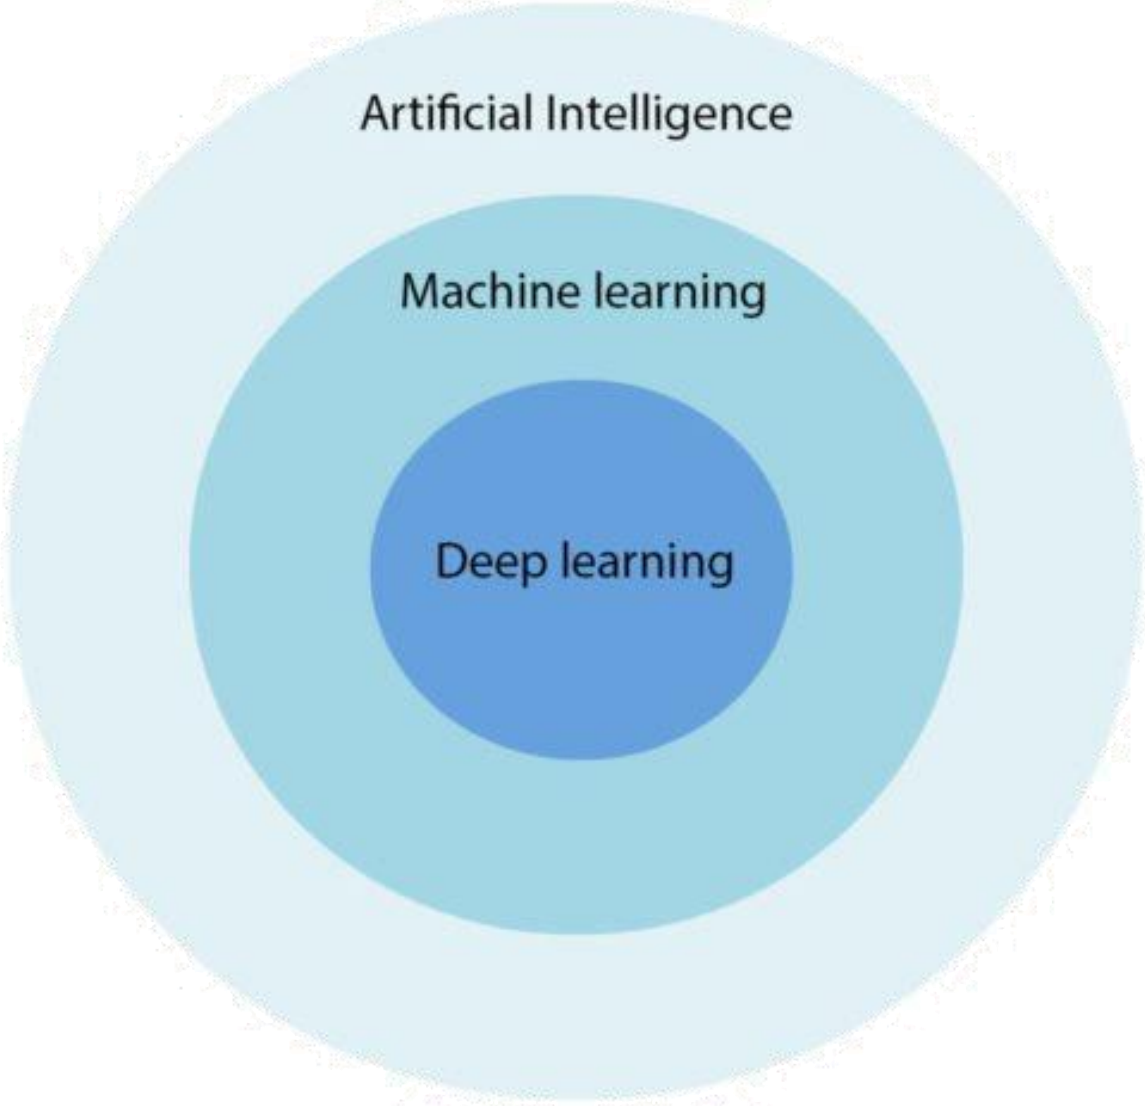

Deep Learning

Kunstigt neuralt netværk

Rigtig neuron

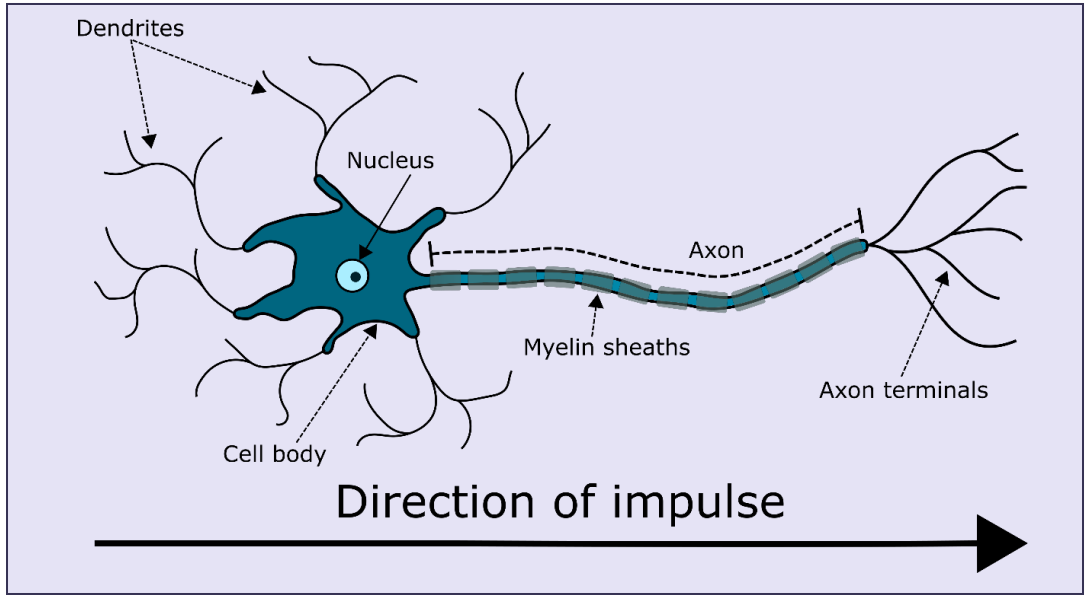

Kunstig neuron

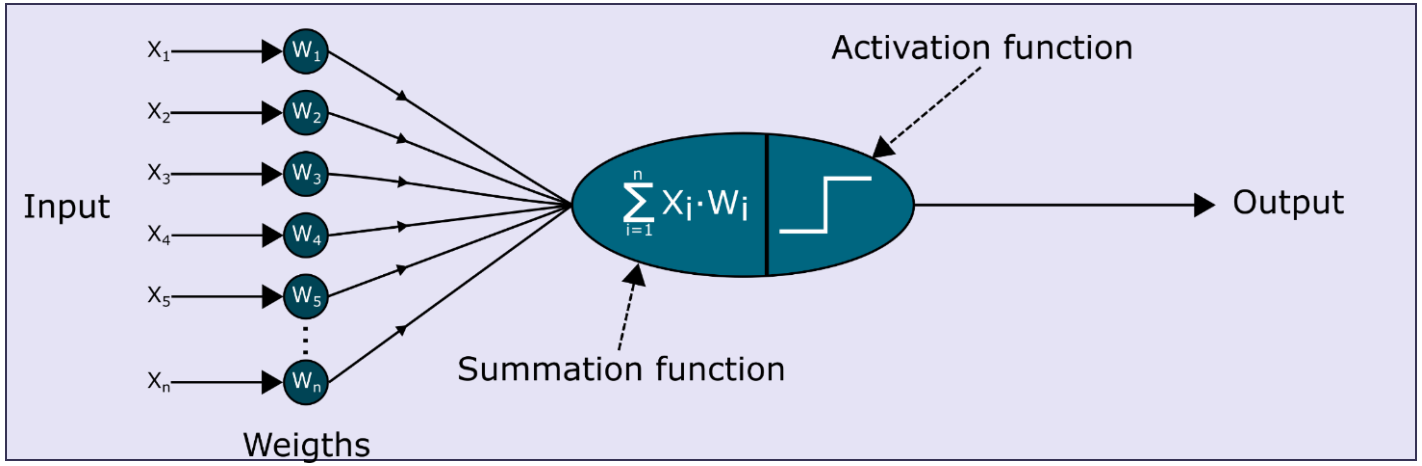

Deep Learning

Kunstigt neuralt netværk

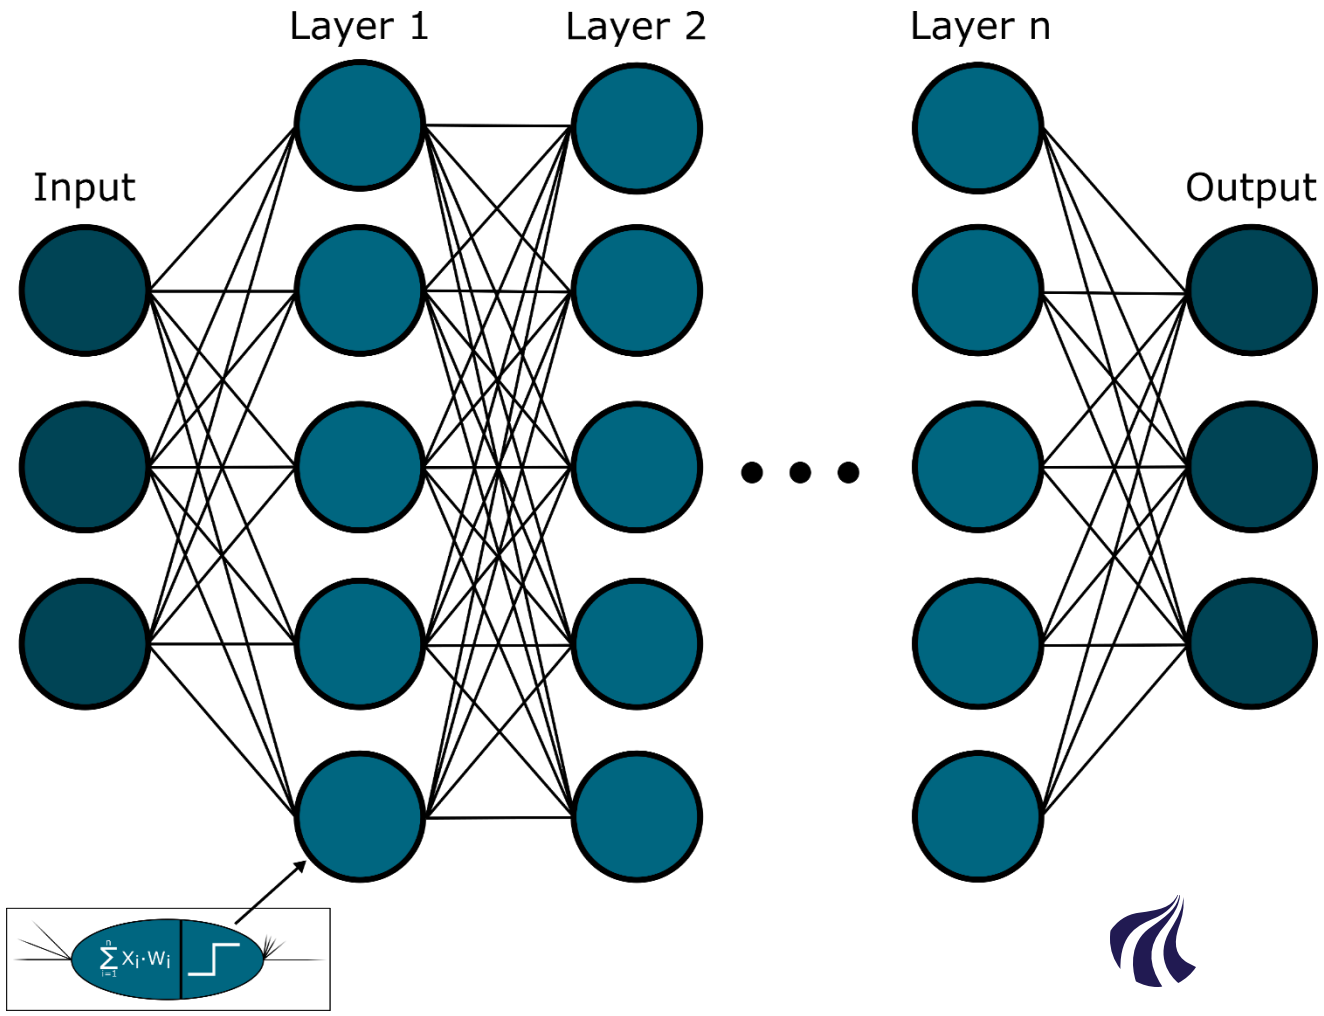

- Antallet af lag beskriver dybden af netværket
  - 5 – 150 lag
- 1 million – 100 millioner parametre

Deep Learning

Kunstigt neuralt netværk

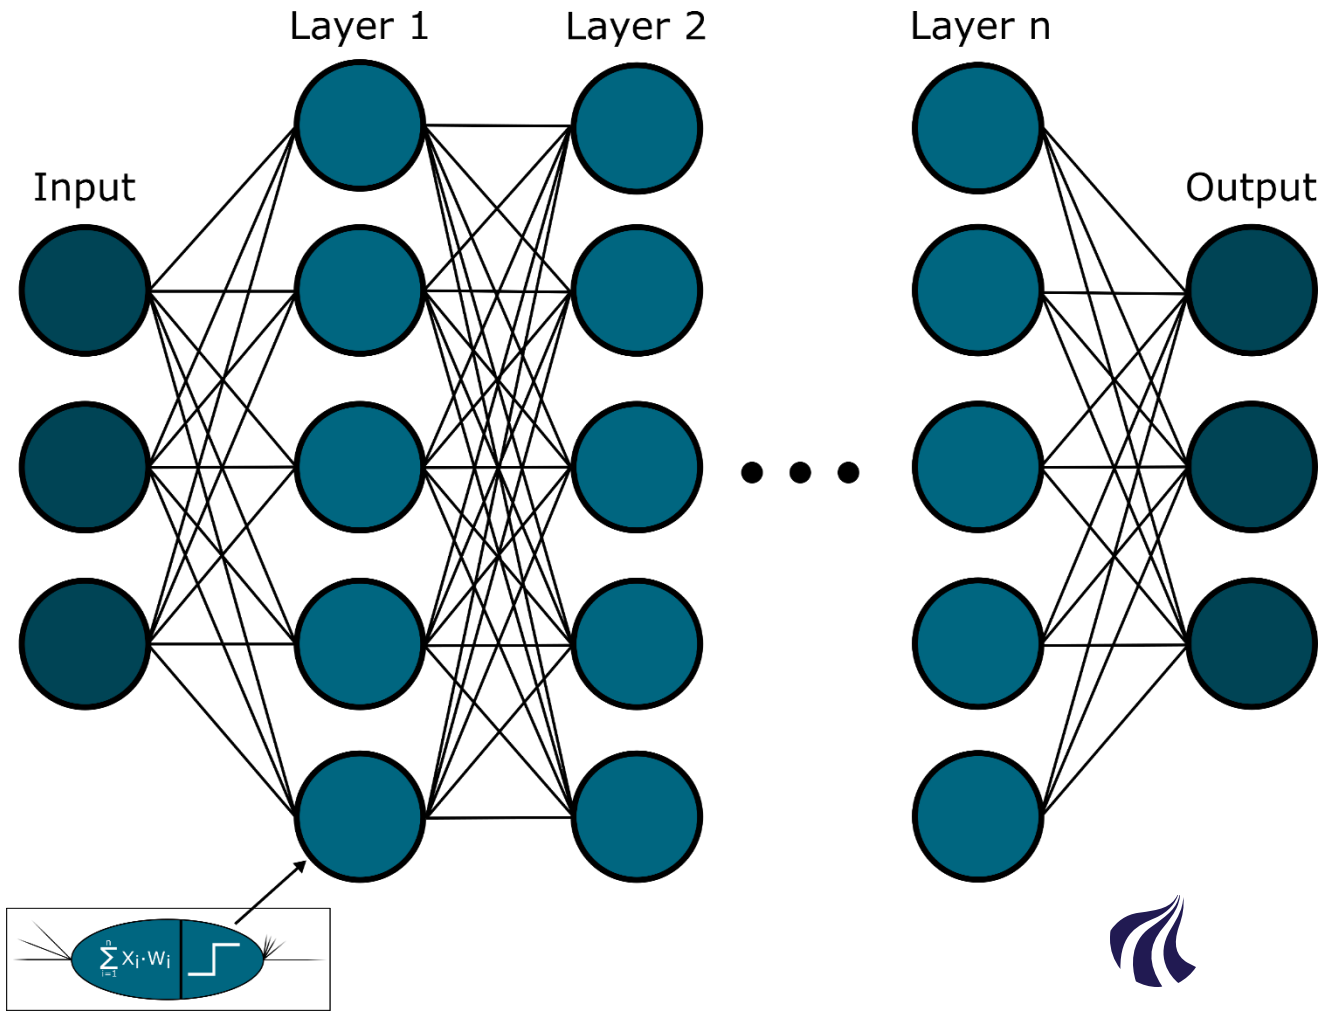

- Træning:
  - Presenter mange kendte input-outputrelationer

→ ”Per”

⋮

→ ”ikke Per”

  - Lad netværket iterere
    - Mange udregninger

Deep Learning

Kunstigt neuralt netværk

- Imponerende resultater

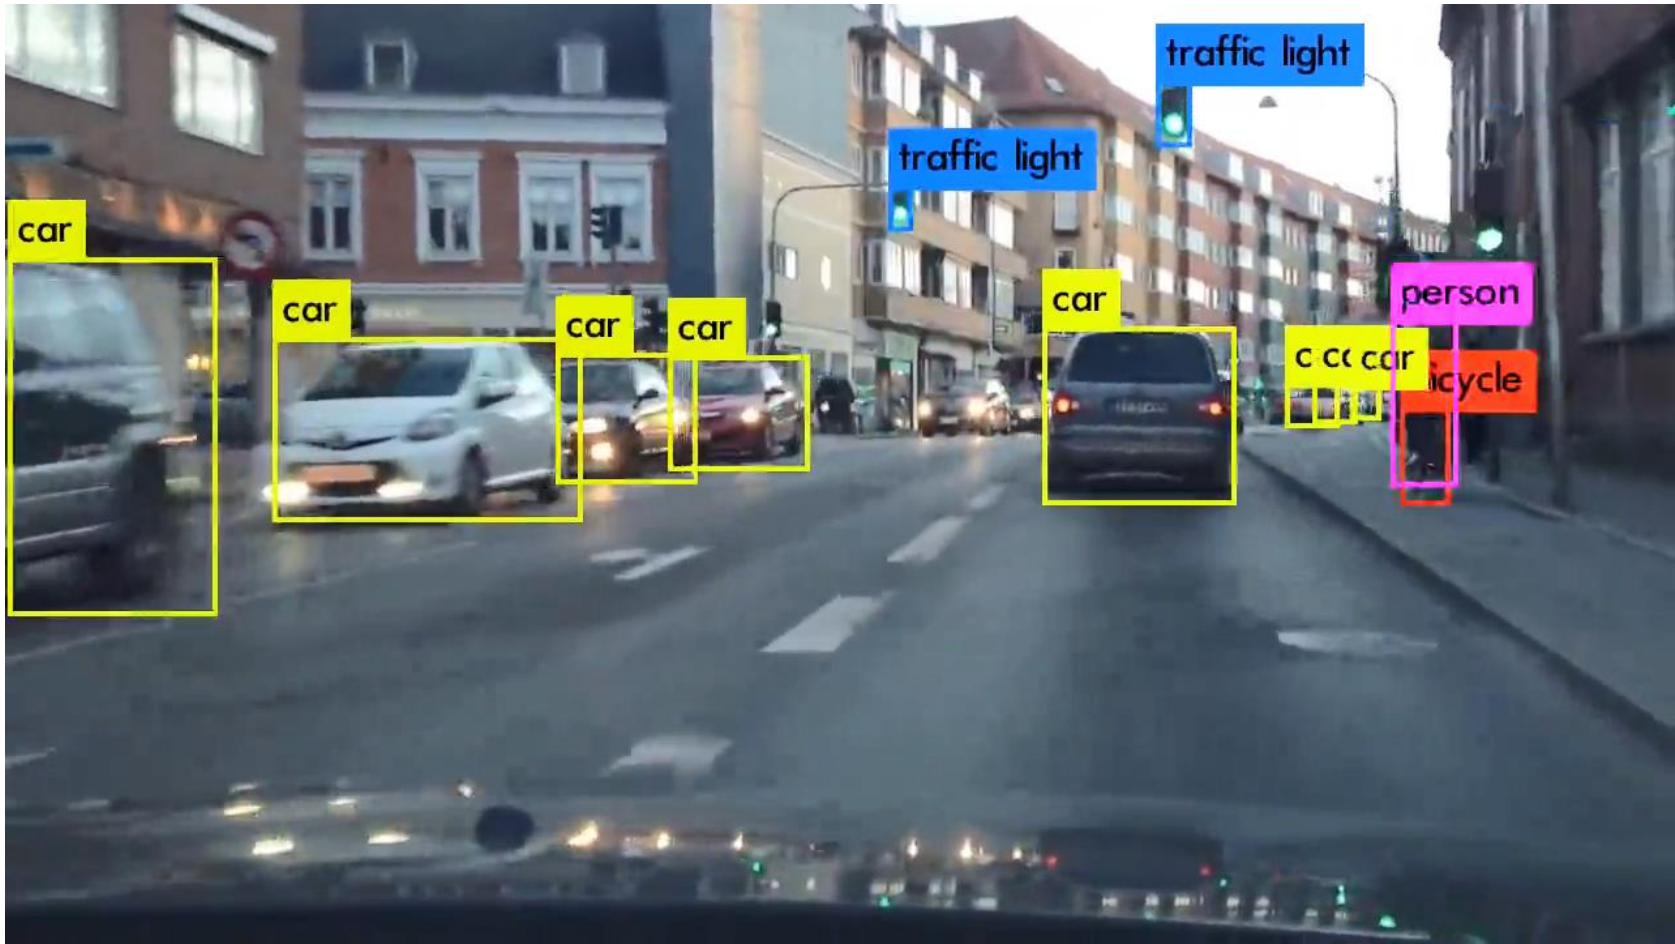

Deep Learning

Kunstigt neuralt netværk

- Imponerende resultater

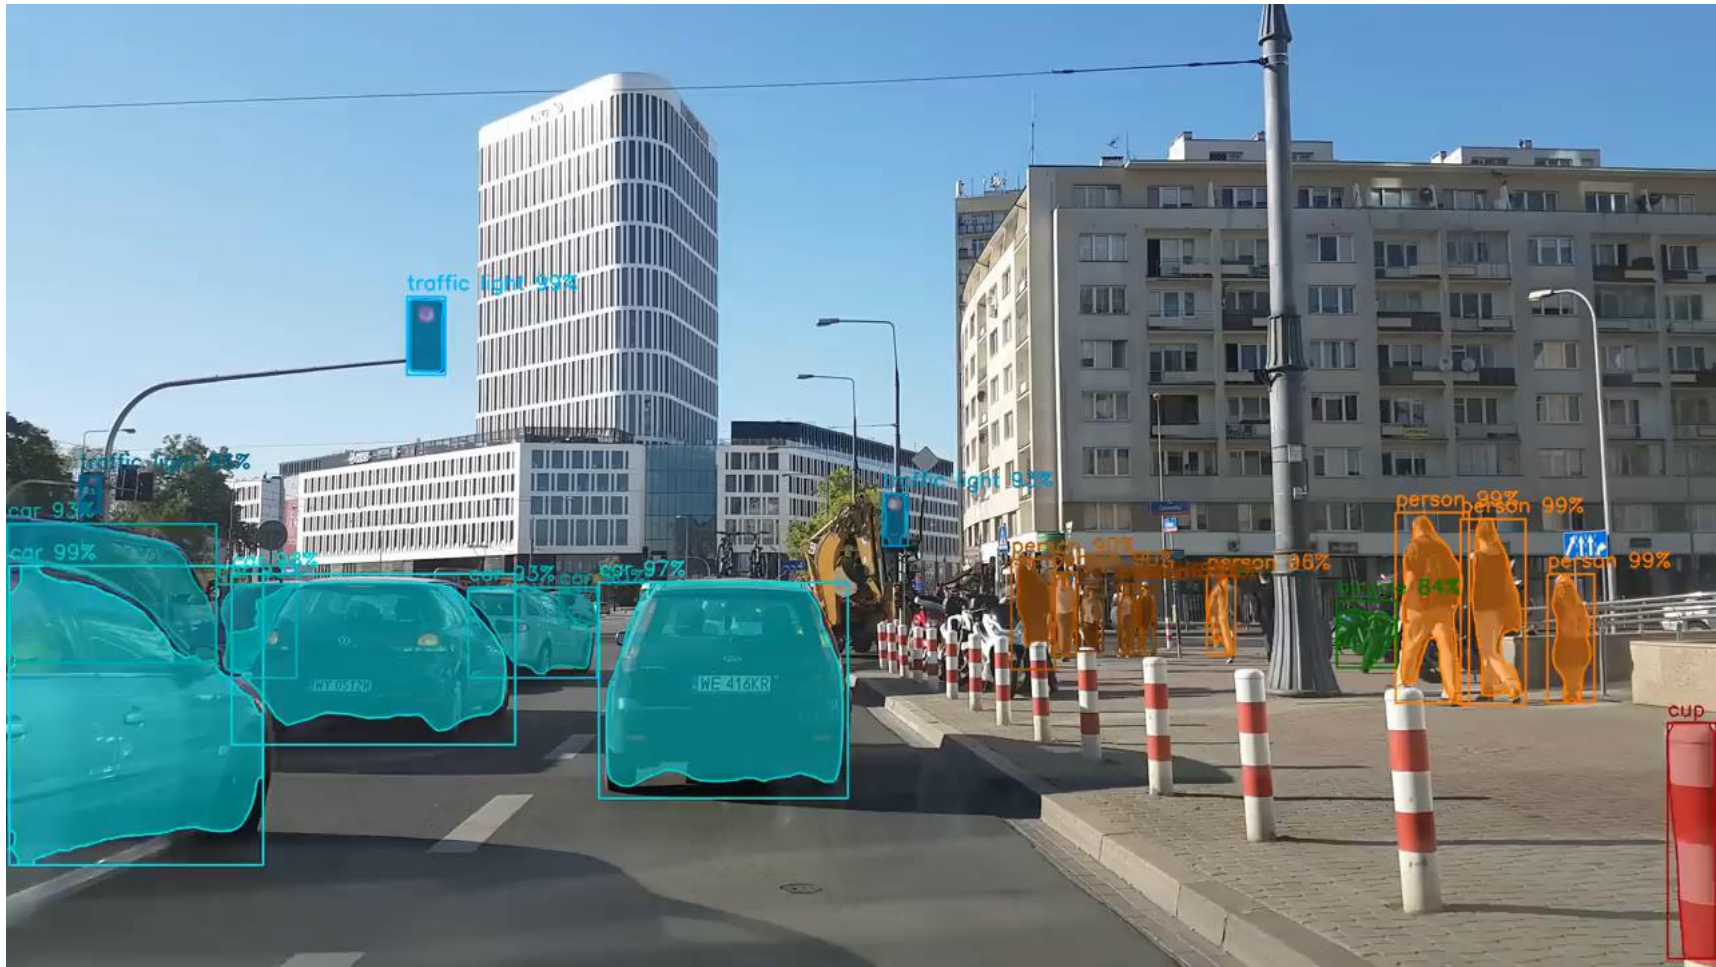

Deep Learning
Hvorfor først nu?

- Tilstrækkeligt store datasæt
- Massive mængder af regnekraft

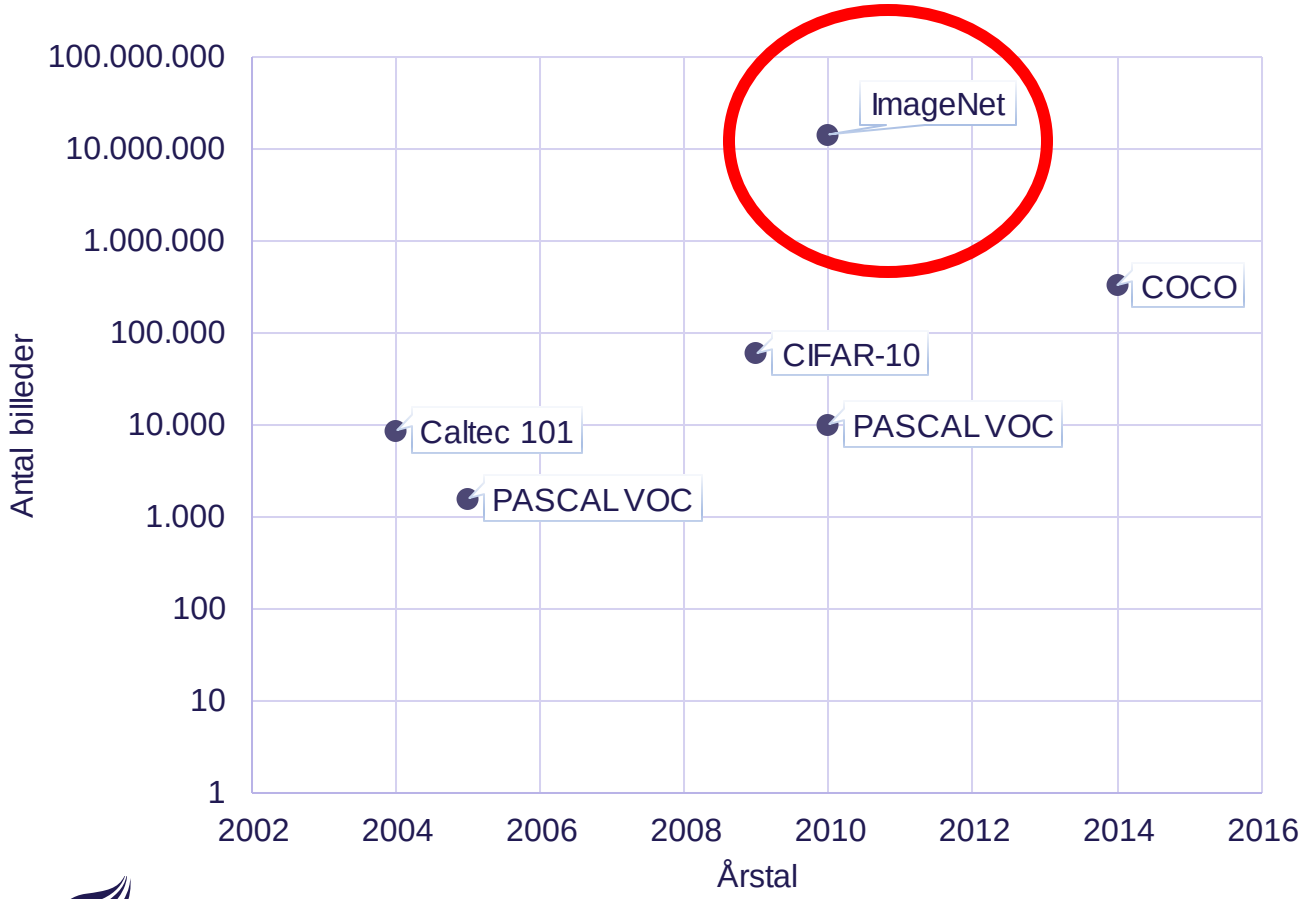

Hvordan kommer man i gang?

- Gør-det-selv

- Cloud-baserede tjenester

Google Cloud Vision API

Microsoft Cognitive Services

Hvordan kommer man i gang?

- Cloud-baserede tjenester

Google Cloud Vision API

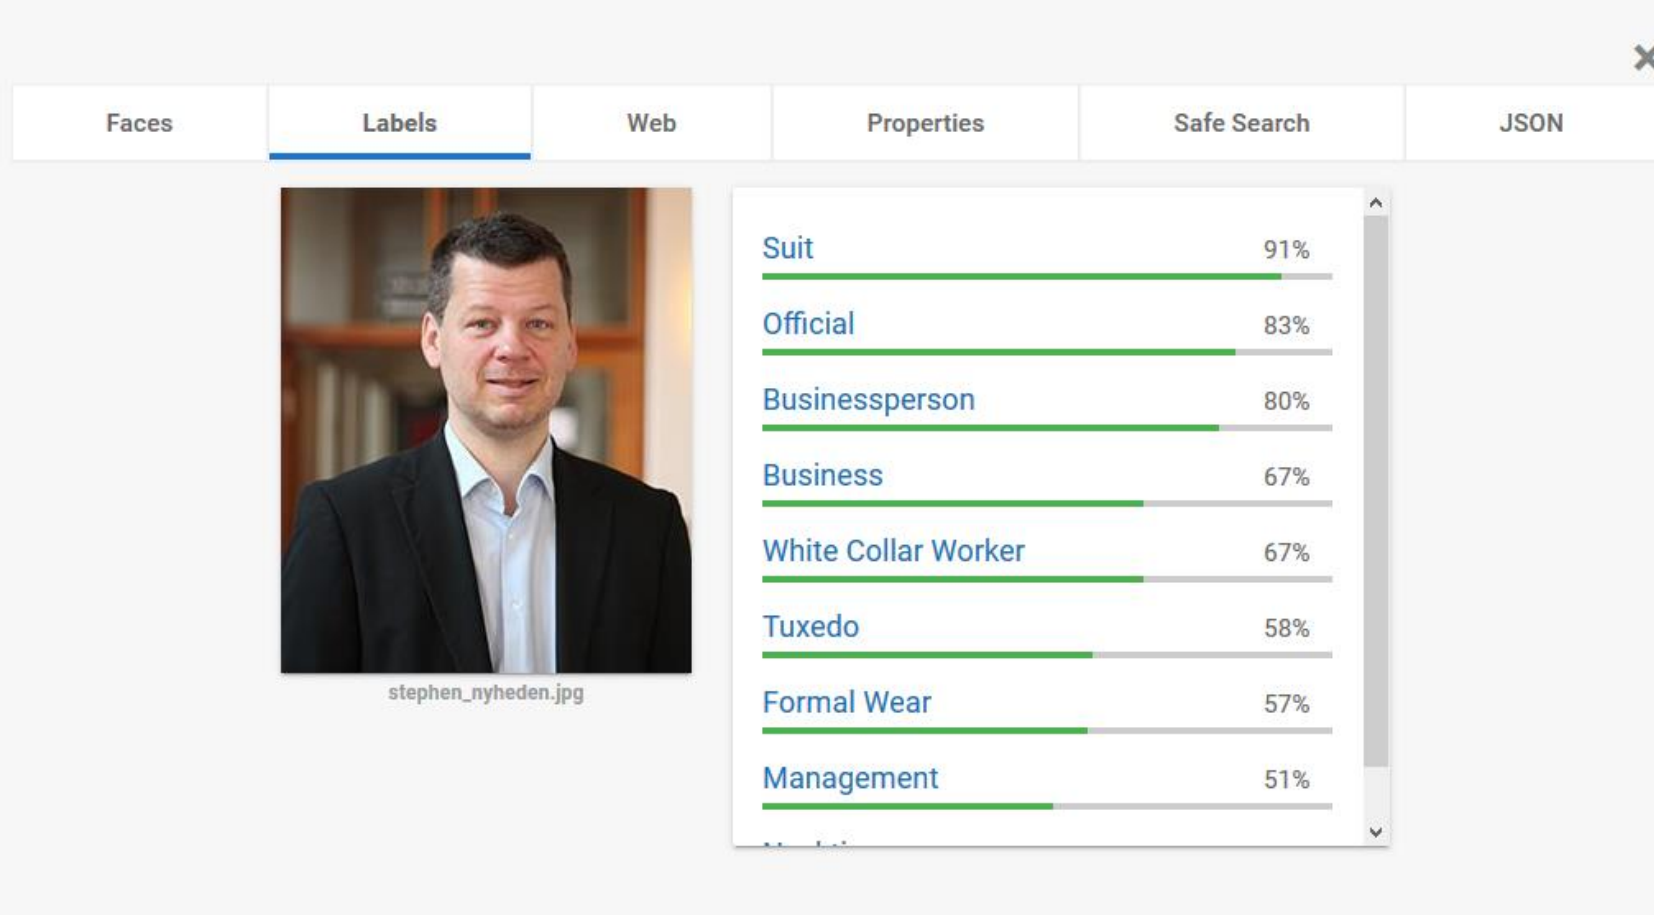

Hvordan kommer man i gang?

- Kræver træningsdata – jo mere, jo bedre

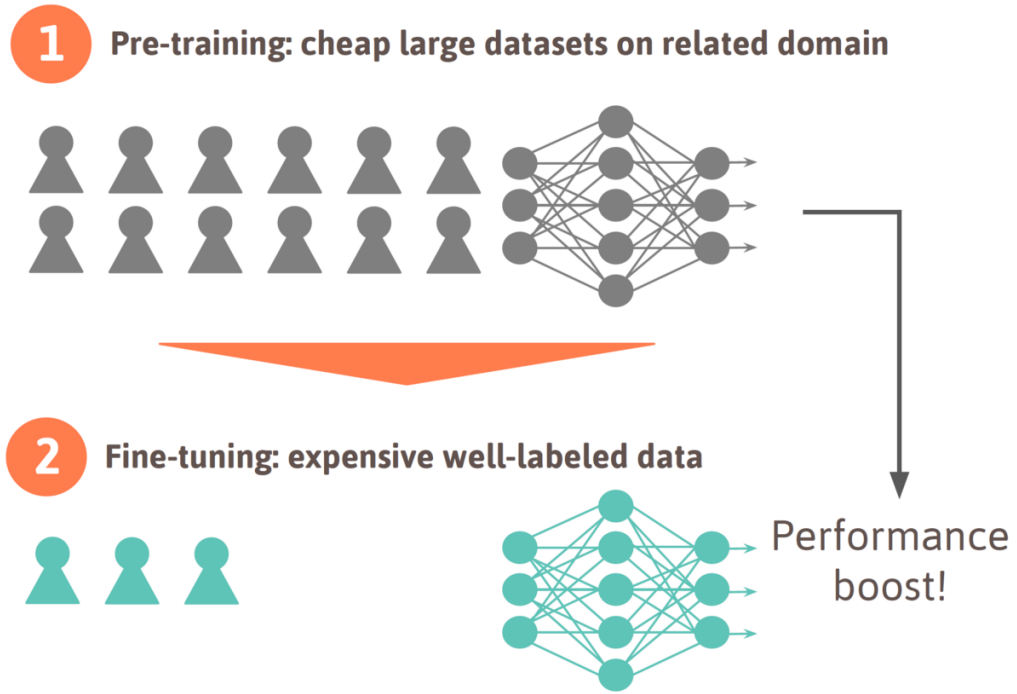

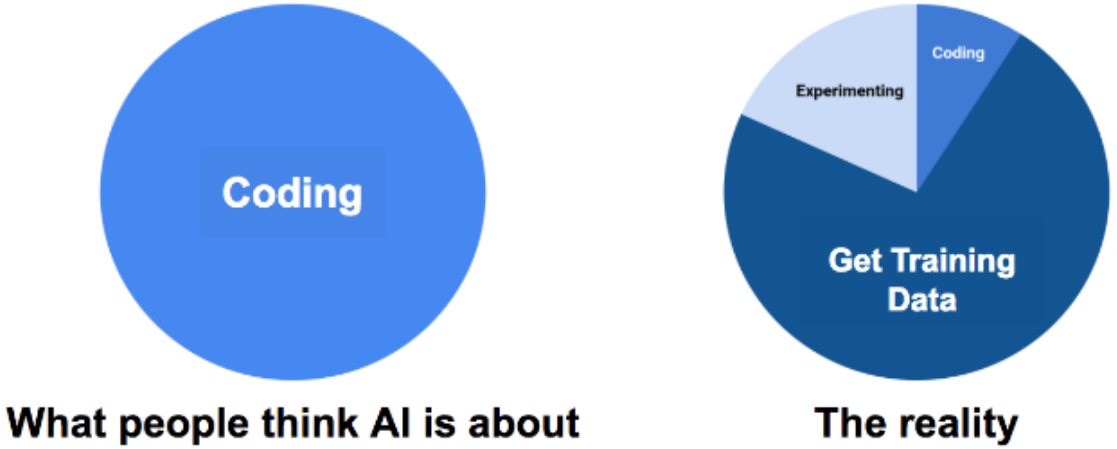

AALBORG UNIVERSITET

https://hackernoon.com/%EF%B8%8F-big-challenge-in-deep-learning-training-data-31a88b97b282

Hvordan kommer man i gang?

- Gør-det-selv
  - Find Github-repo med kode fra state-of-the-art framework
  - Download docker-billede
  - Finindstil netværket med eget træningsdata

- Generelt
  - Open-source code
  - Lukkede datasæt
  - Lukkede datahåndterings-værktøjer

Hvordan kommer man i gang?

Populære biblioteker

Keras

PyTorch

Programmeringssprog

python™

Onlinekurser

fast.ai

UDACITY

coursera

Hvad er prisen for en fejl?

- Et forkert resultat i en Google-søgning?
- Fejl i ansigtsgenkendelse?
- En industrikylling, der fejlagtigt bliver klassificeret som rask?
- Fejldetektion af en person i selvkørende biler?

ars TECHNICA

BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE

THE FUTURE IS NOT WHAT IT USED TO BE —

Amazon's Rekognition messes up, matches 28 lawmakers to mugshots

ACLU: "And running the entire test cost us $12.33—less than a large pizza."

CYRUS FARIVAR - 7/26/2018, 2:00 PM

AALBORG UNIVERSITET

Næste emne: Muligheder og begrænsninger i Deep Learning

AALBORG UNIVERSITET

AI & the road ahead

The road ahead

- Thanks for your attention
- Questions?

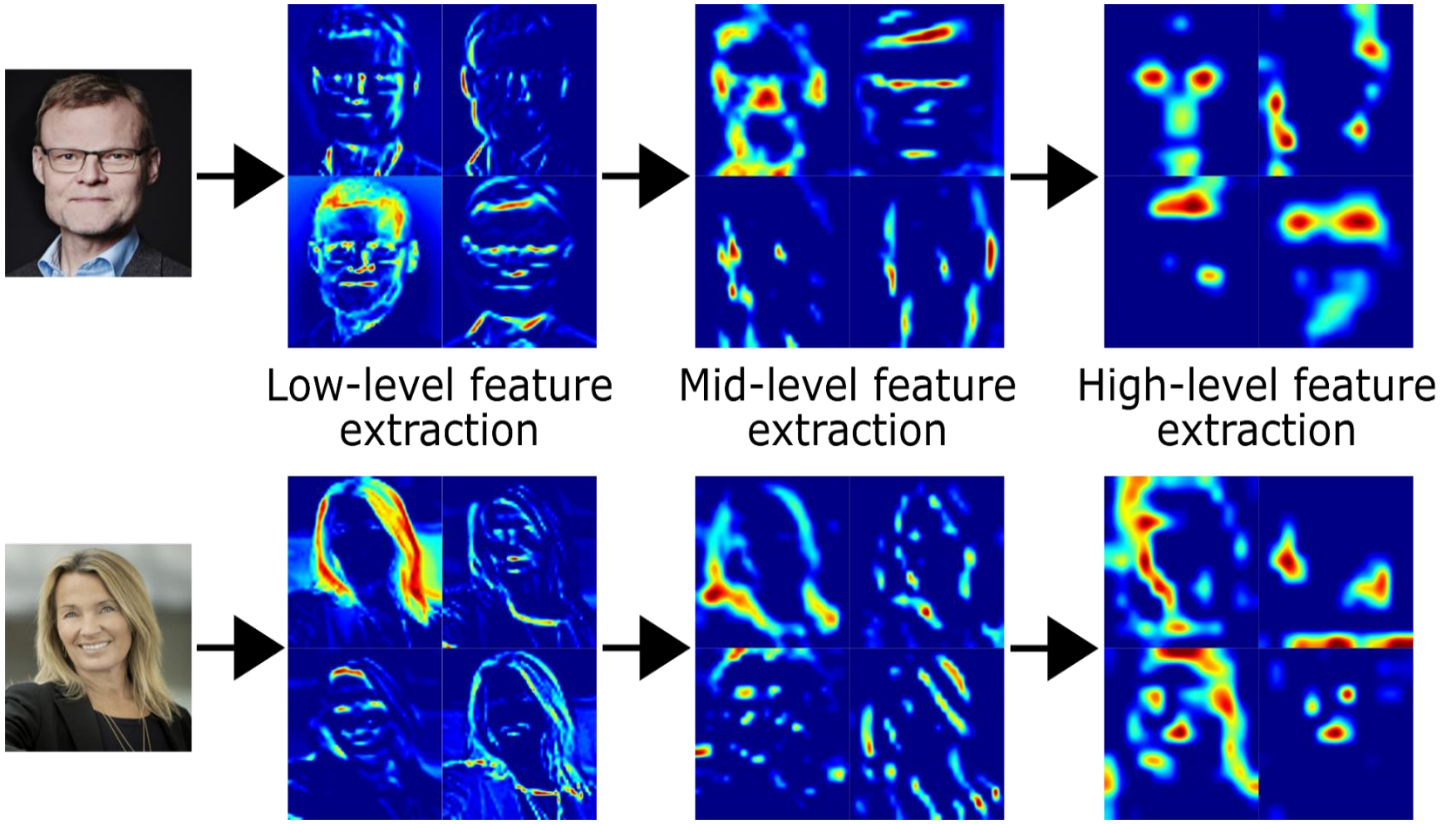

Background

Intelligence

- What does ”intelligence” mean?
- Is a machine intelligent?
- Alan Turing (1912 - 1954)
- The Turing-Test:

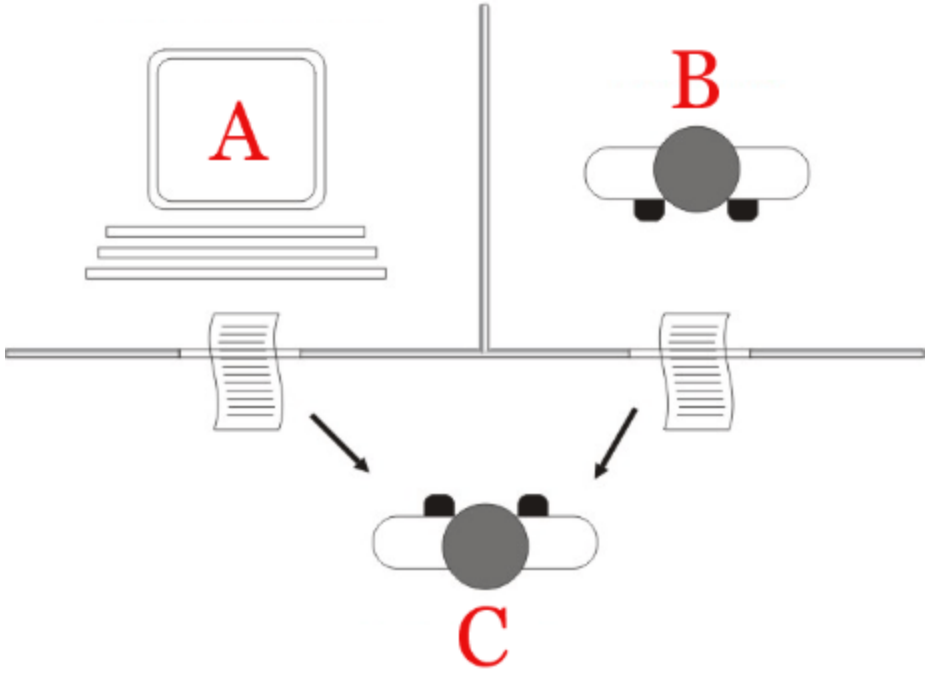

Background
Intelligence

- What does ”intelligence” mean?
- Is a machine intelligent?
- Alan Turing (1912 - 1954)
- The Turing-Test:

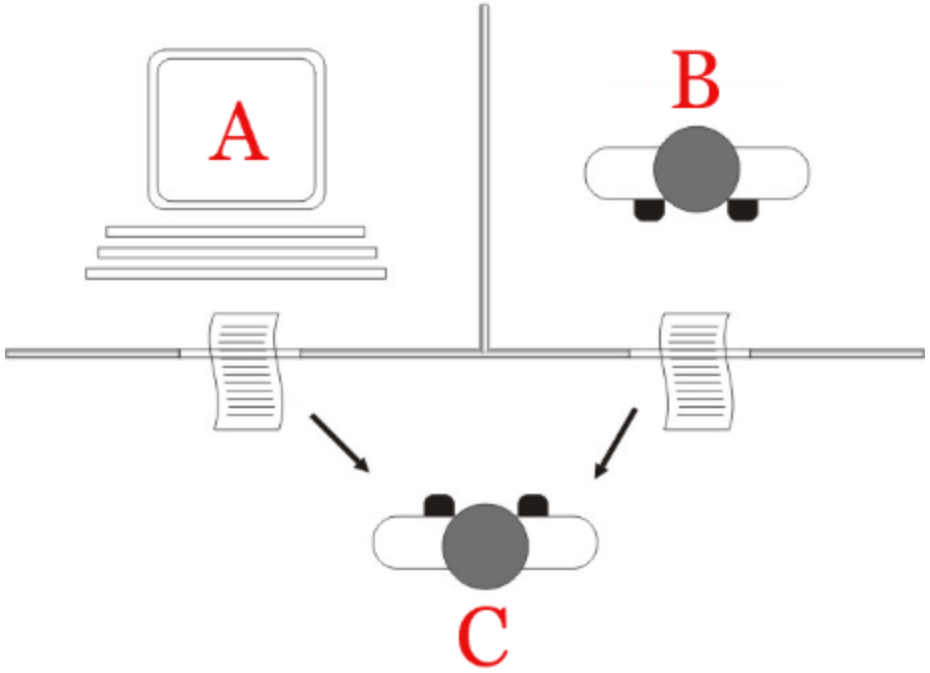

AALBORG UNIVERSITET

Background

Intelligence

- Earlier:
  - Calculator
  - Thermostats
  - Optical character recognition

Background
Intelligence

- Earlier:
  - Calculator
  - Thermostats
  - Optical character recognition

- Today:
  - People recognition

”Lene” ”Per”

-

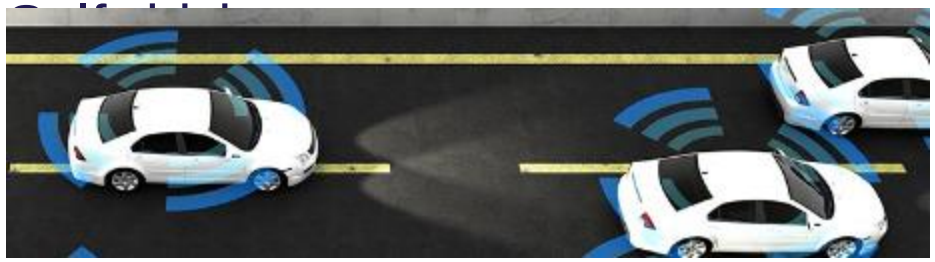

AALBORG UNIVERSITET

Background
Intelligence

- Human intelligence vs machine intelligence

	Machine	Human
Hard	 	
Easy		 

What is AI?

- AI based on rules

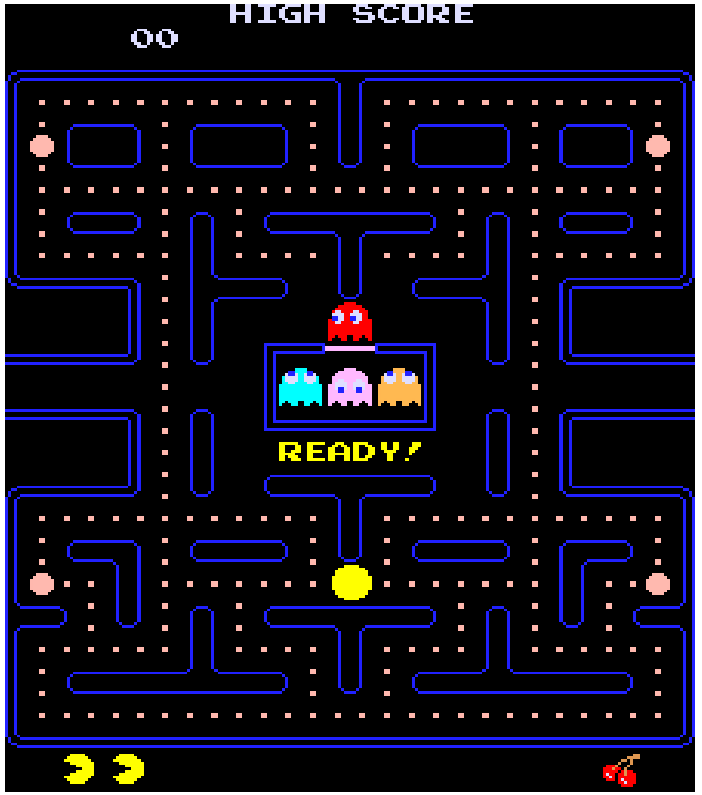

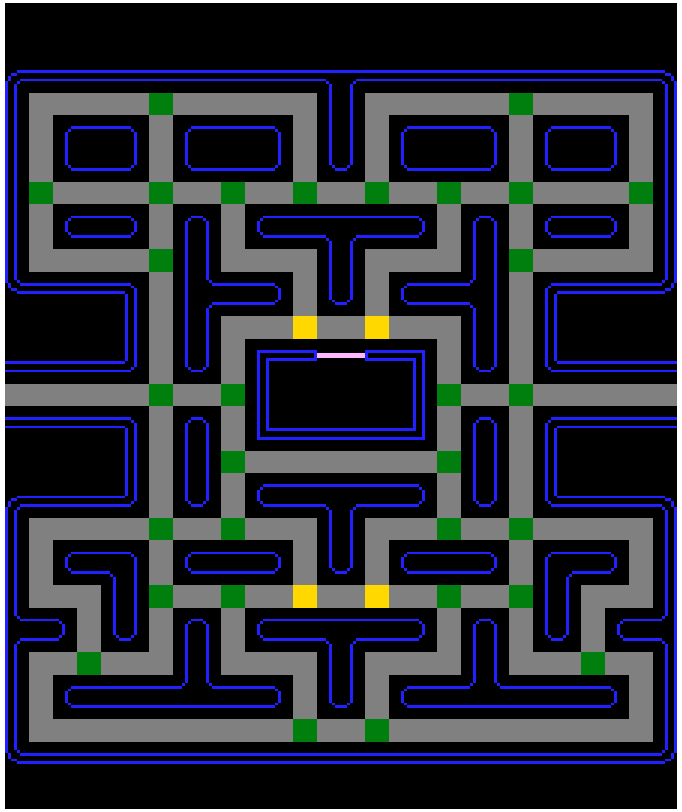

What is machine learning?

- Example: Differentiate between Per and Lene
  - Amount of hair?
  - Face circularity?
- Non-linear?

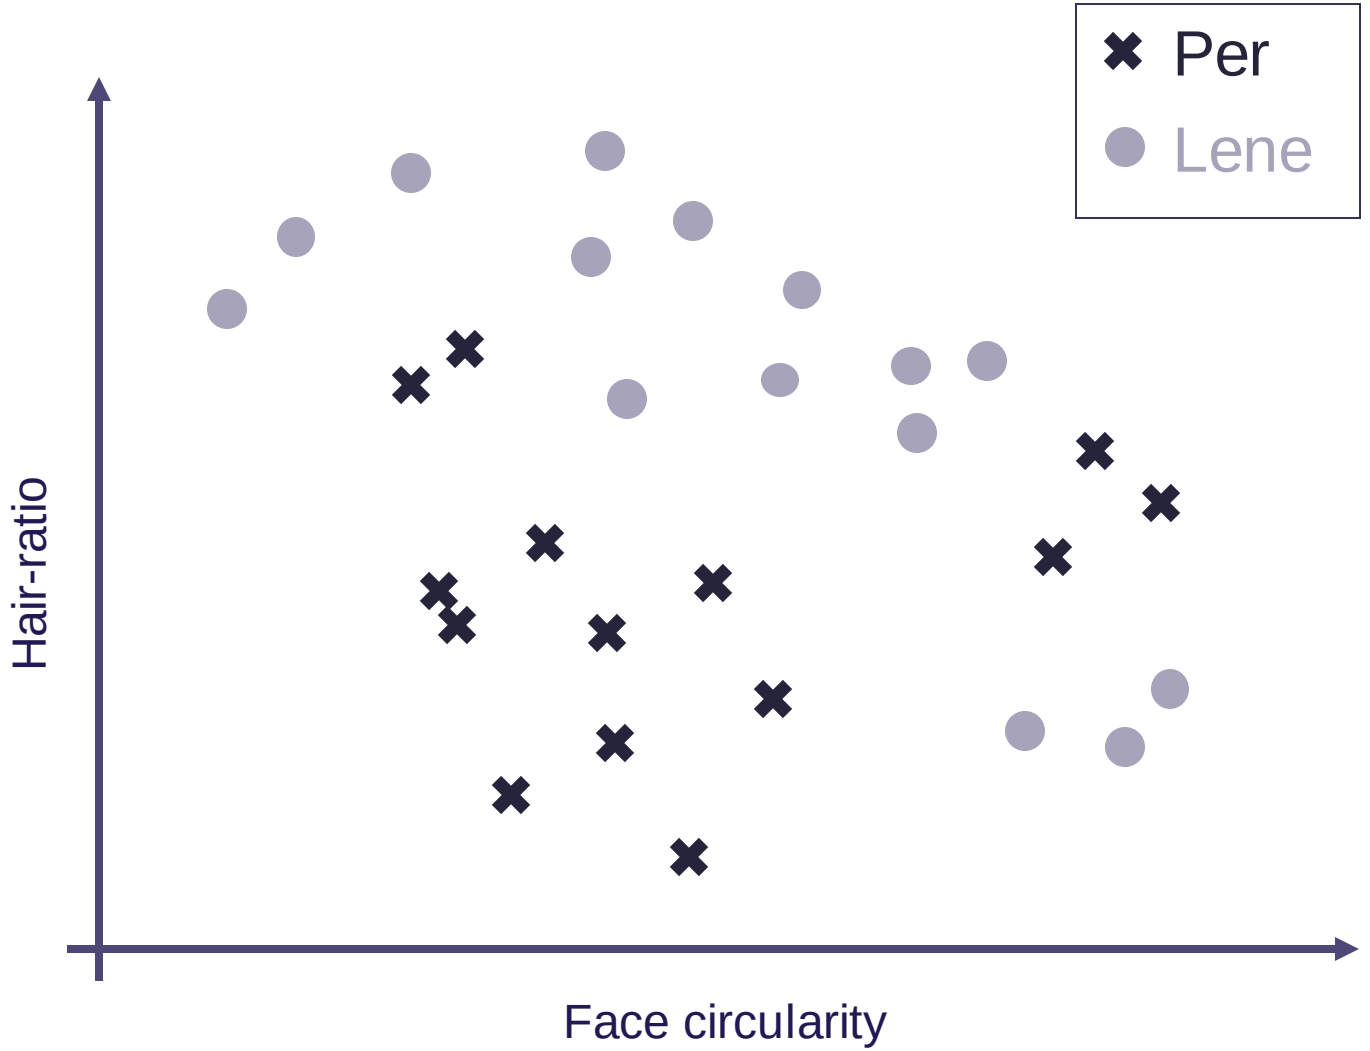

What is machine learning?

- Example: Differentiate between Per and Lene
  - Amount of hair?
  - Face circularity?

- Find curve separating the two categories
  - Non-linear?

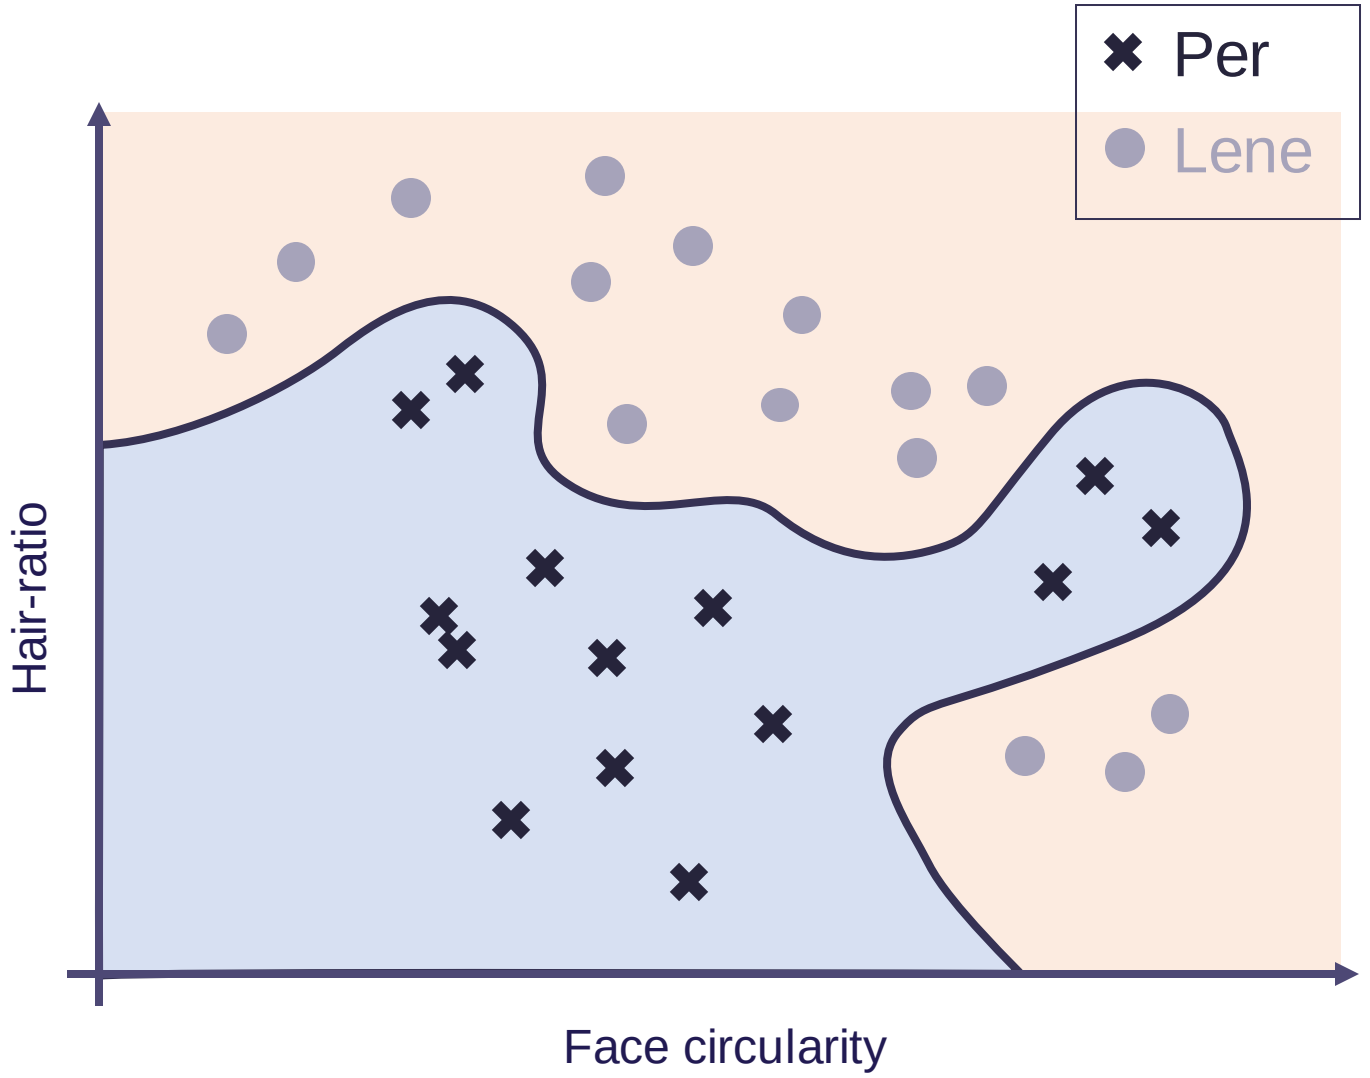

What is machine learning?

- New image – who?
  - Exctract features
  - Where in feature space?
  - Compare to decision line
  - Decide who it is

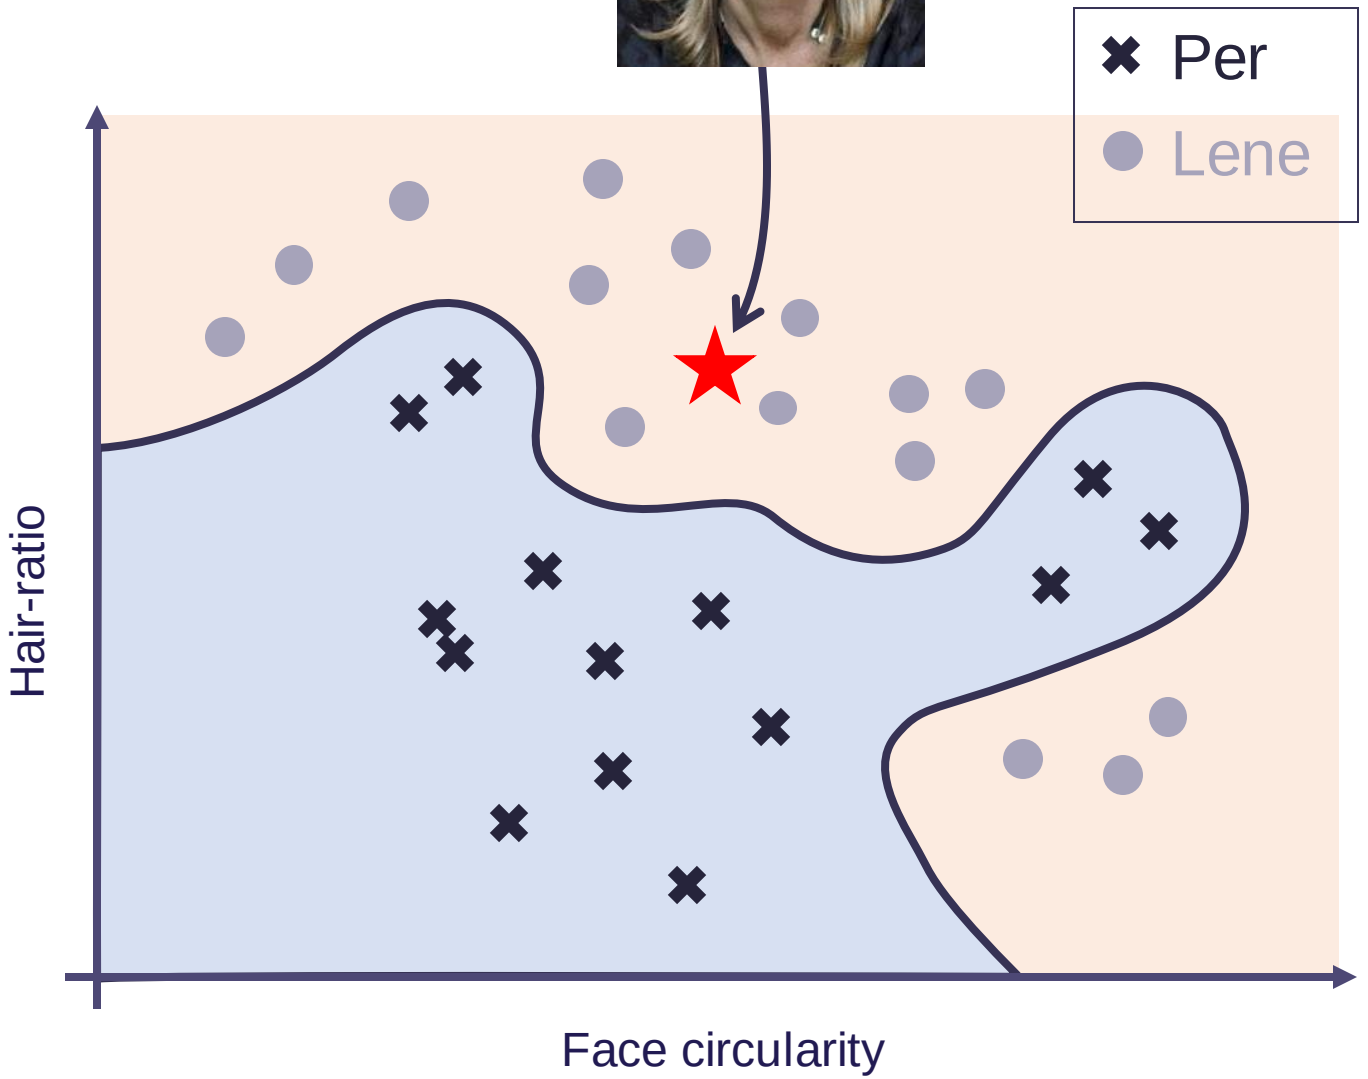

What is machine learning?

- New image – who?
  - Exctract features
  - Where in feature space?
  - Compare to decision line
  - Decide who it is

- This is Machine Learning

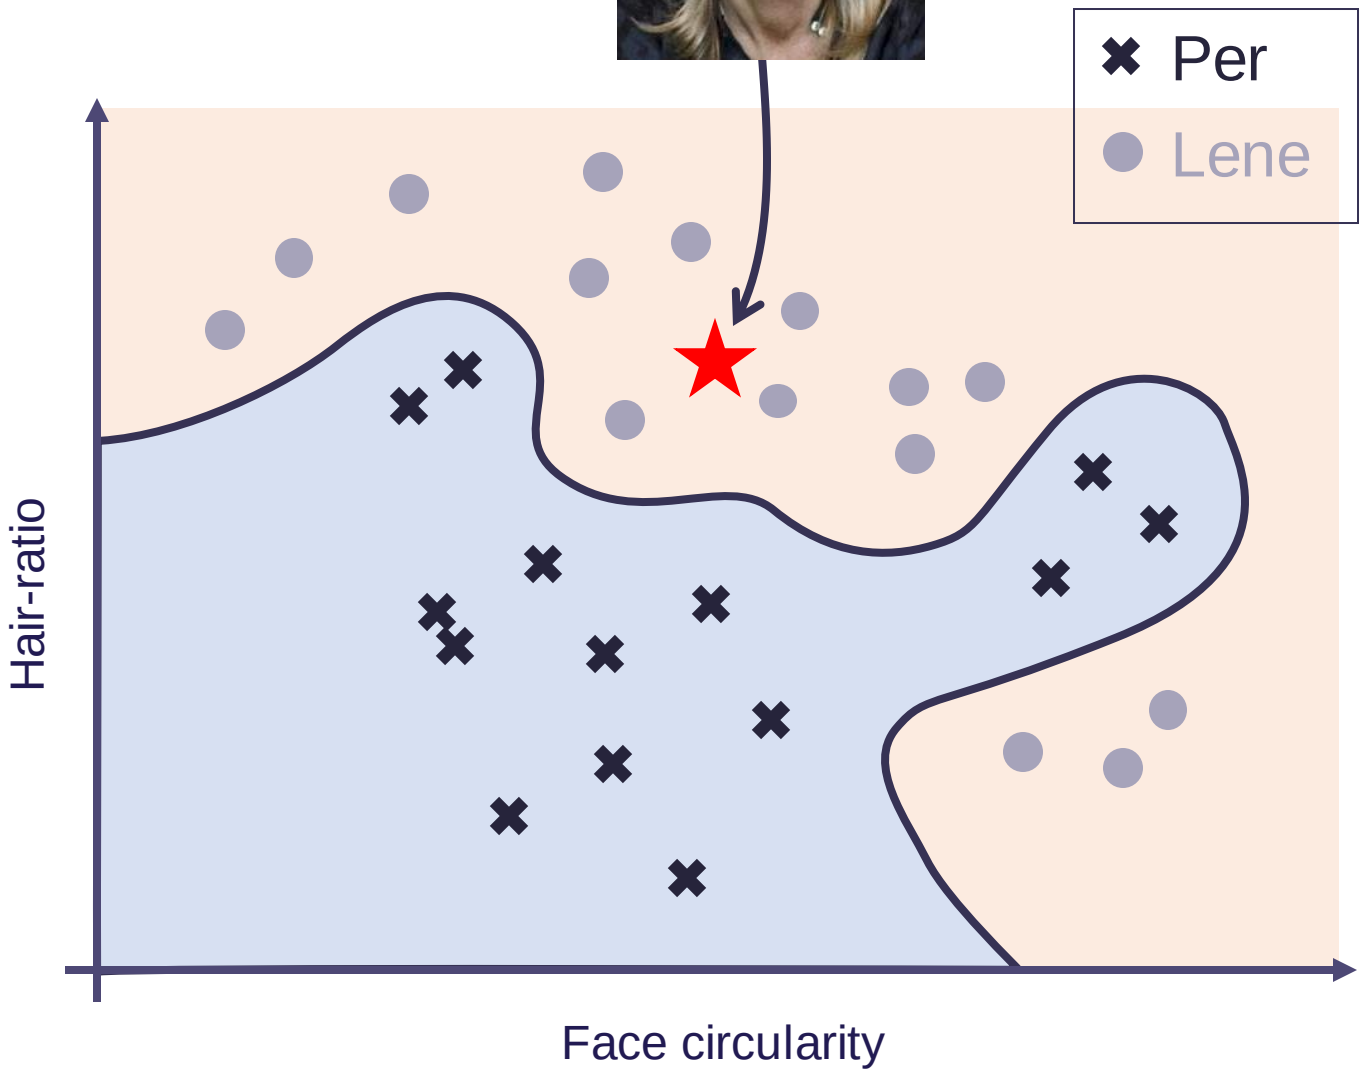

Deep learning

How does it work?

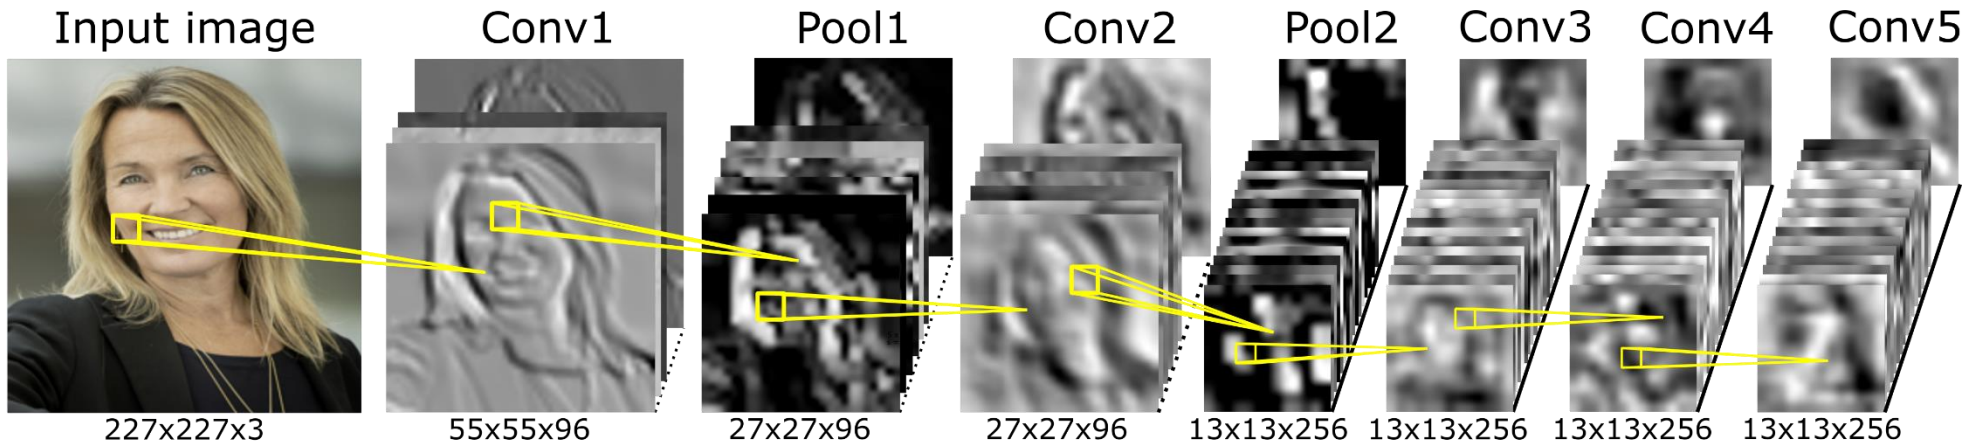

AI
The road ahead

- We may look to the past
  - As old jobs are lost to automation, new jobs are created

# of robots in Amazon Warehouses:
2014: 1.455
2017: 4.500