AI in Engineering

Albert van Breemen
a.j.n.v.breemen@tue.nl
Program manager Ai, EAISI & HTSC

HTSC
HIGH TECH SYSTEMS CENTER
TU/e

2 July 2020

Artificial Intelligence is like a Swiss army knife …

... as it can be applied to many different applications/areas, such as machine vision, logistics, planning, robotics, optimization, predictive maintenance, sales, …

HTSC HIGH TECH SYSTEMS CENTER TU/e

High Tech prepares for AI

Consumer 'database'-driven AI

Google
Uses ML to predict traffic density based on anonymous mobile phone position data.
https://ai.google/

Netflix
Uses ML as part of their movie recommender system.
https://www.wired.co.uk/article/how-do-netflixs-algorithms-work-machine-learning-helps-to-predict-what-viewers-will-like

Facebook
Uses ML-based algorithms to detect and recognize faces in photo's.
https://research.fb.com/category/machine-learning/

Spotify
Uses ML for recommending new music.
https://www.oreilly.com/ideas/machine-learning-at-spotify-you-are-what-you-stream

Apple
Developed speech recognition service Siri with ML.
https://machinelearning.apple.com/

Paypal
Uses ML to detect and combat fraud
https://www.paypal-engineering.com/tag/machine-learning/

2nd wave of AI applications

Industrial 'sensor data'-driven AI

General Electric
Uses ML develop Digital Twins to understand, predict and optimize performance of assets.
https://www.ge.com/digital/predix/digital-twin

Tesla
Uses ML to develop autonomous driving cars.
https://www.youtube.com/watch?v=2aVc84mGJfE

Siemens
Uses AI a.o. for Industry 4.0, medical and traffic optimization applications
https://www.siemens.com/global/en/home/company/innovation/pictures-of-the-future/artificial-intelligence.html

PTC
Applies AR and ML for maintenance.
https://www.ptc.com/en

SemioticsLab
Dutch startup using ML for predictive maintenance.
https://www.semioticlabs.com/nl/

Freenome
Startup using AI for early cancer detection.
https://www.freenome.com/

High Tech prepares for AI

Data Science	**AI (in) Engineering**
DATA	DATA
↓ *is a*	↓ *generated by*
NUMBER	SENSOR
↓ *has a*	↓ *has a*
Numerical/Logical meaning	Physical meaning
↓ *analysed by*	↓ *analysed by*
Mathematics, Computer Science	Engineering, Science

Designing high tech systems

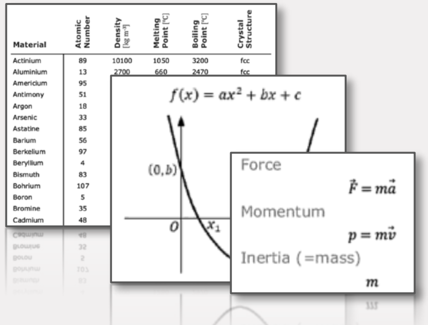

Engineer

First principles

MODEL

Data

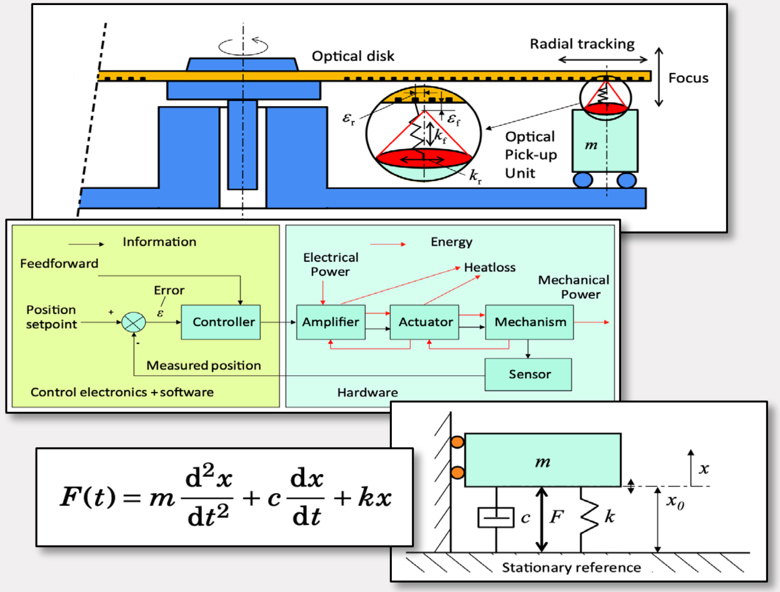

$$F(t) = m\frac{d^2x}{dt^2} + c\frac{dx}{dt} + kx$$

Optical disk system

SYSTEM

Source: 'The Design of High Performance Mechatronics', R.M. Schmidt et al, 2014

HTSC HIGH TECH SYSTEMS CENTER TU/e

Growing design challenges

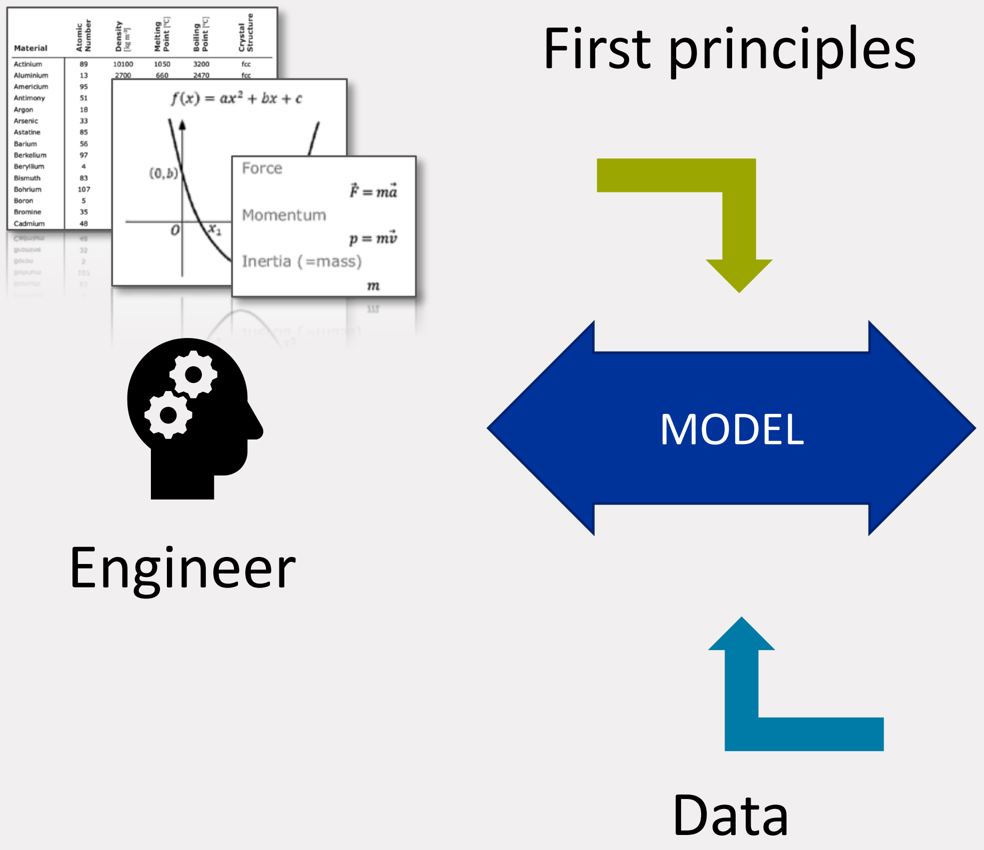

MORE COMPLEXITY

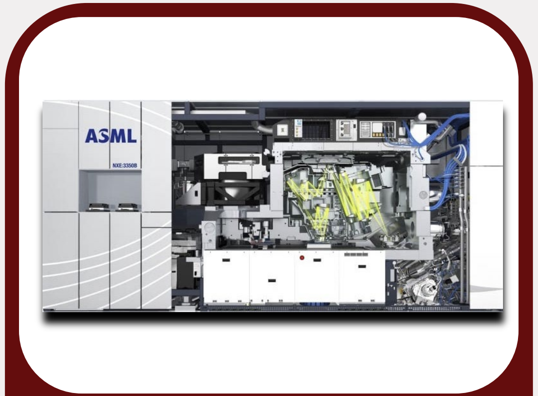

MORE PERFORMANCE

MORE INTELLIGENCE

MORE INTEROPERABILITY

Growing design challenges

Complexer models and bigger data needed to deal with challenges

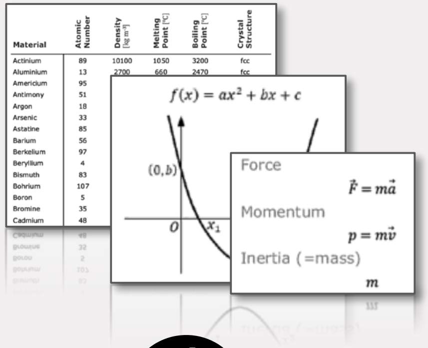

First principles

Engineer

COMPLEXER
MODEL

BIGGER Data

MORE COMPLEXITY

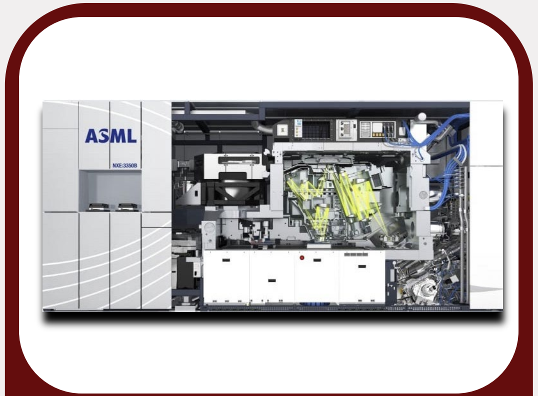

MORE PERFORMANCE

MORE INTELLIGENCE

MORE INTEROPERABILITY

HTSC HIGH TECH SYSTEMS CENTER TU/e

Growing design challenges

Artificial Intelligence is a technology platform to deal with bigger datasets and complexer models

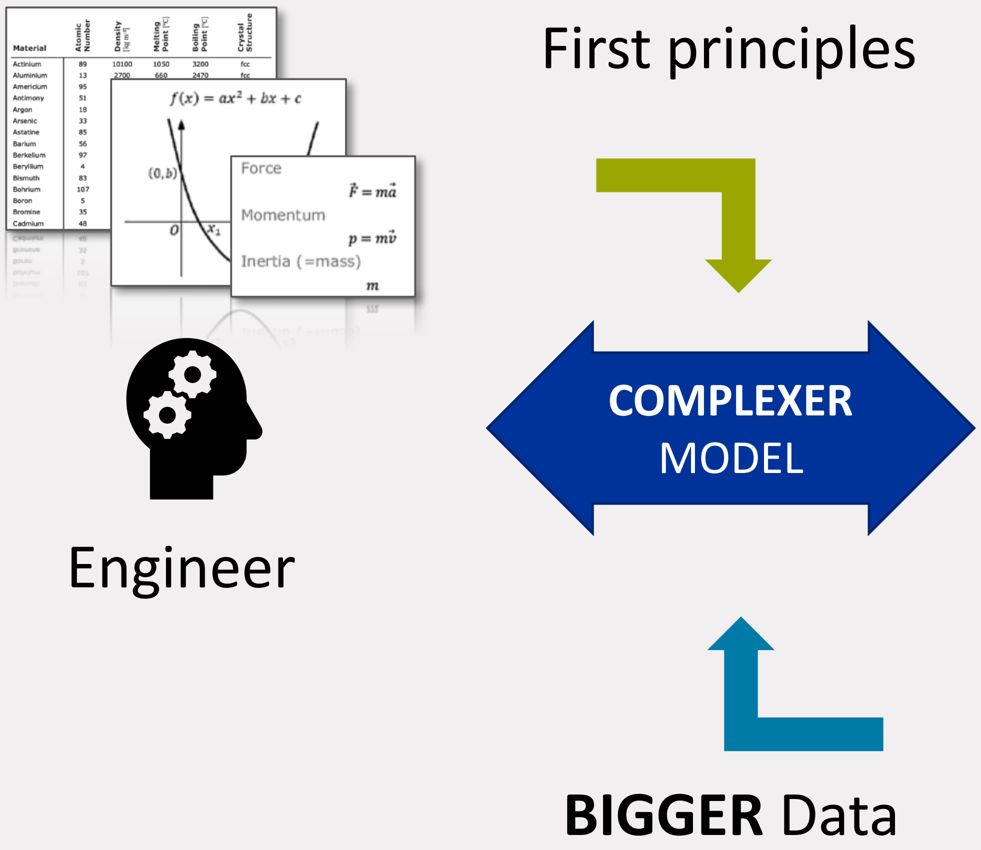

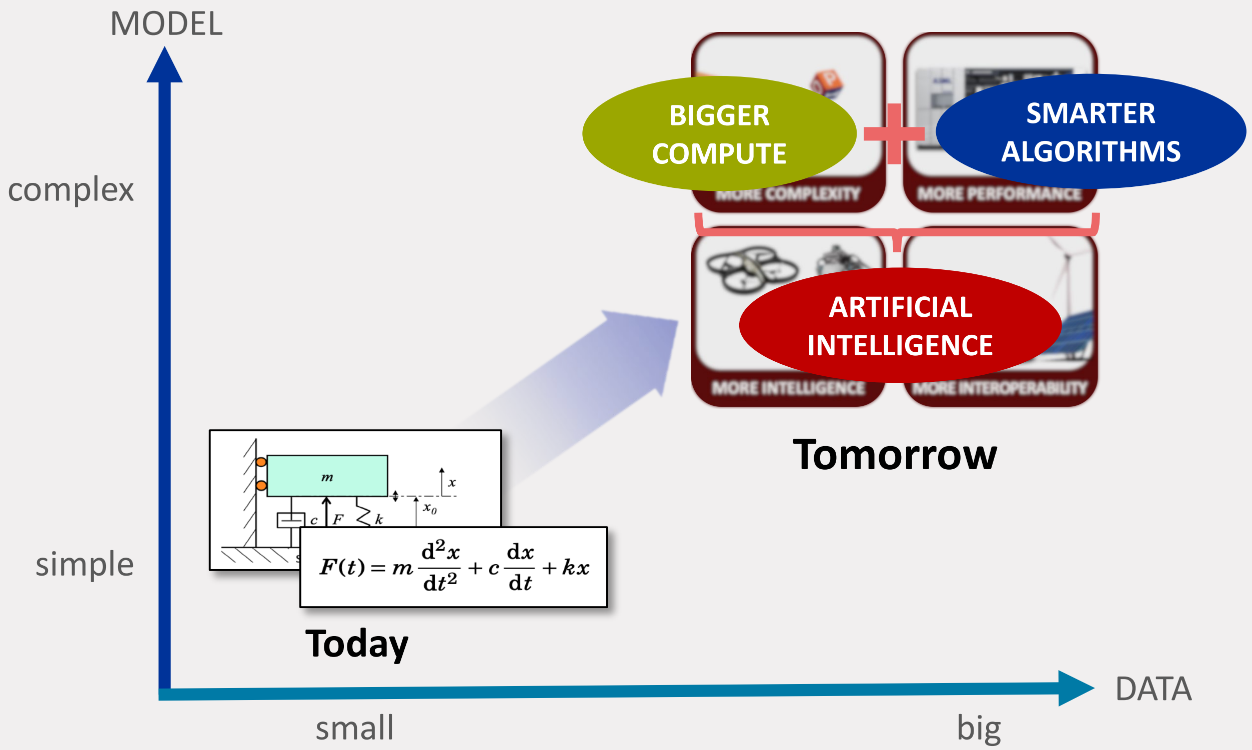

Bigger compute

GPU computing drives AI

GPU COMPUTING ROADMAP

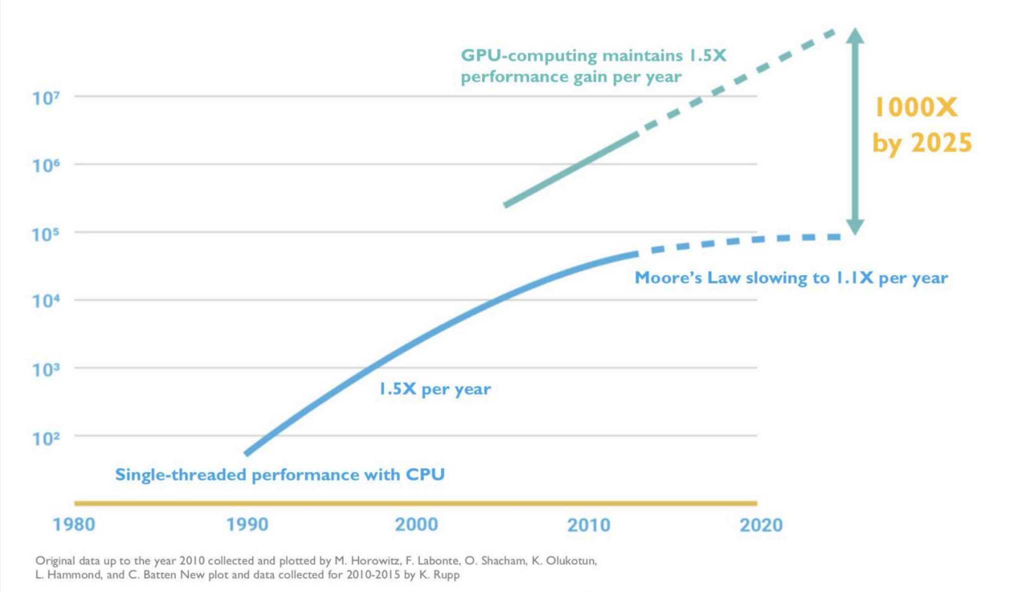

Source: Nvidia, 2018

AI COMPUTE USAGE FOR VISION, PLANNING, SPEECH

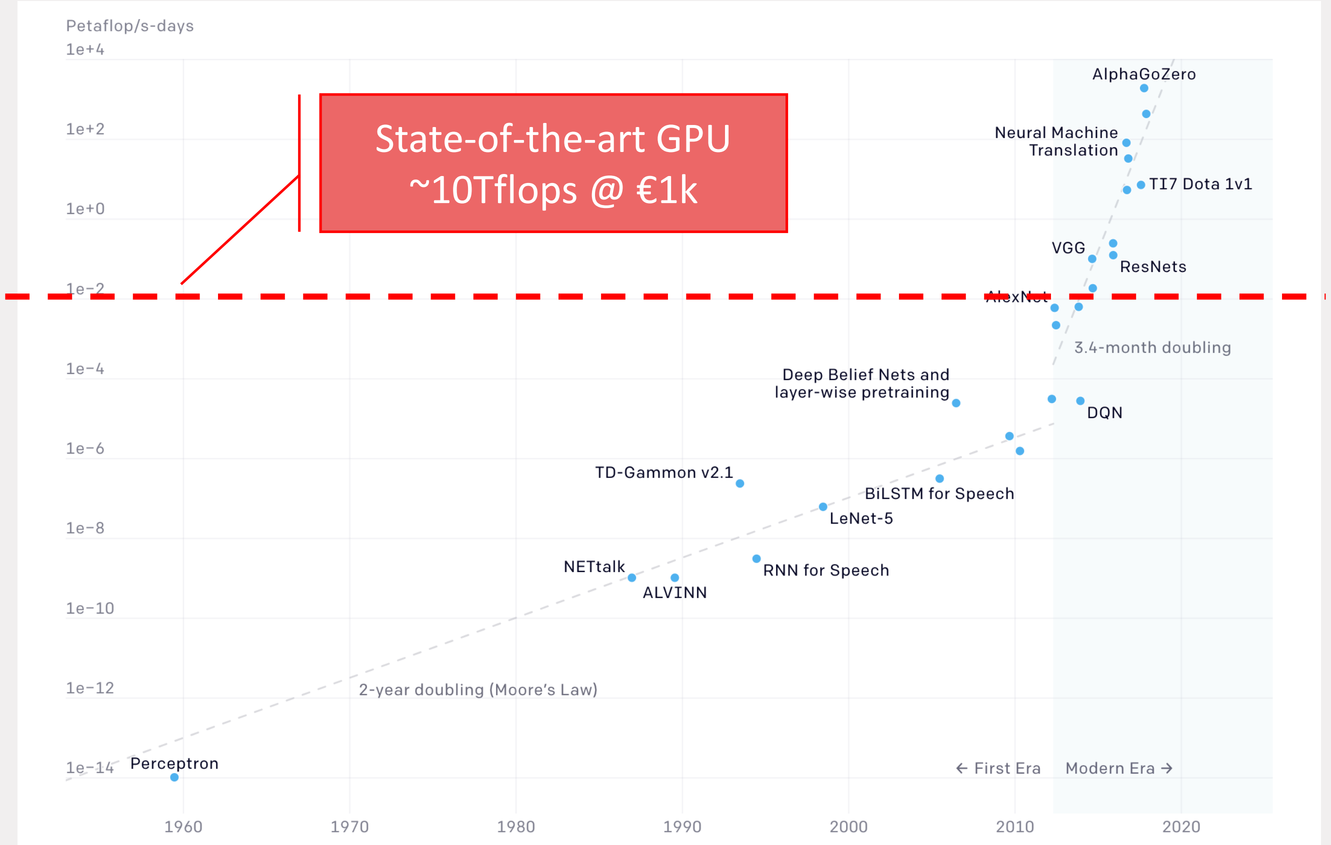

Source: https://openai.com/blog/ai-and-compute/

Smarter algorithms

Scale drives performance

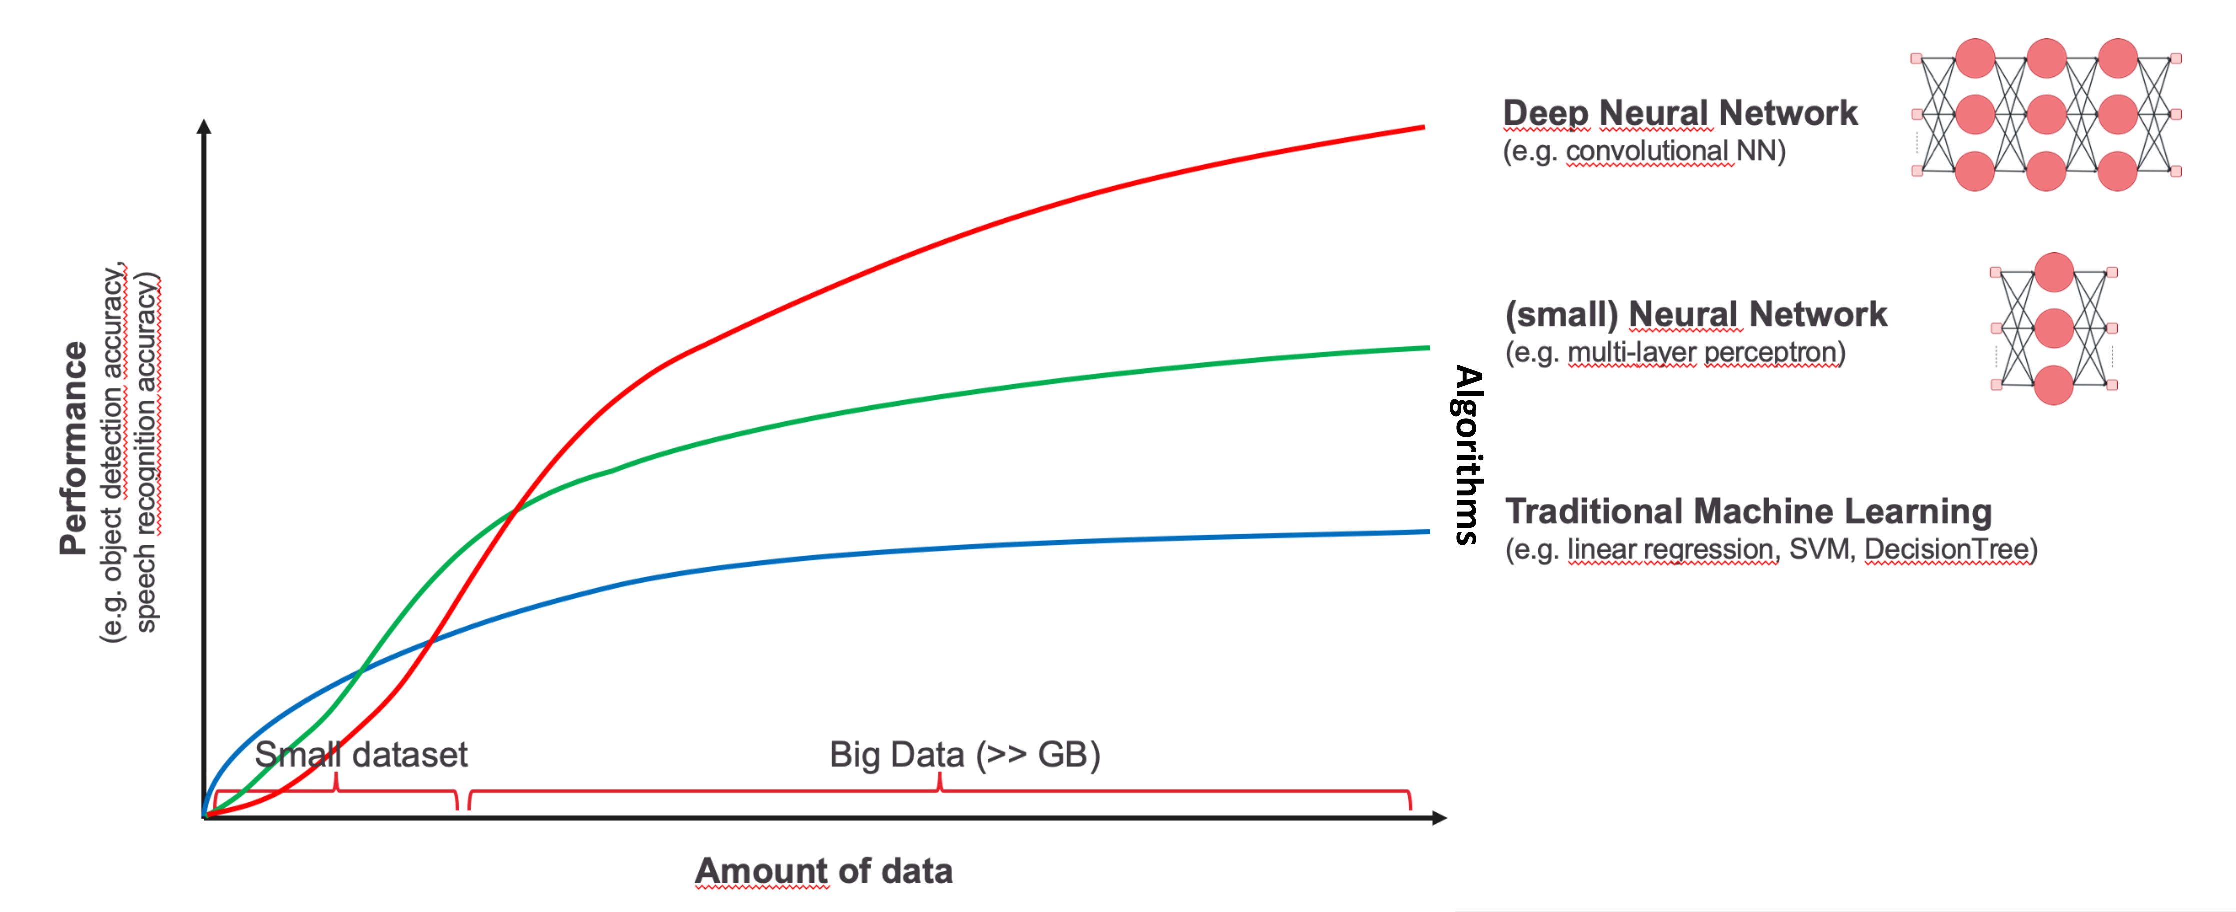

Source: 'Masterclass Deep Learning', VBTI 2018

Artificial Intelligence

Artificial Intelligence (AI) is a multidisciplinary field of science whose goal is to create intelligent machines.

	Knowledge/Logic	Machine Learning	Deep Learning

AI Technology

Knowledge/Logic

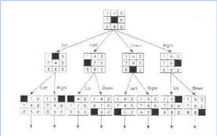

tree search

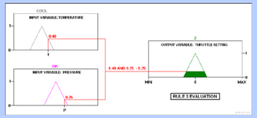

fuzzy logic

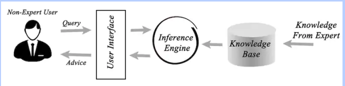

expert system

Machine Learning

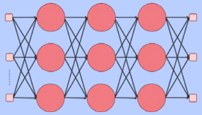

neural networks

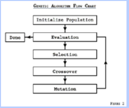

genetic algorithms

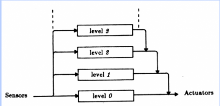

subsumption

Deep Learning

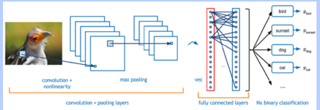

convolutional nn

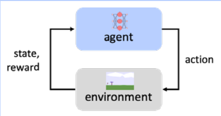

reinforcement learning

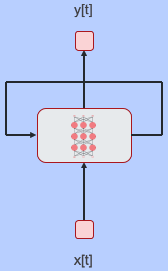

recurrent nn

AI Application

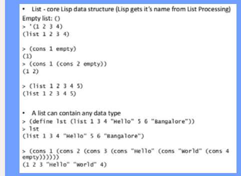

- List - core Lisp data structure (Lisp gets it's name from List Processing)

Empty list: ()

```
> '(1 2 3 4)
(list 1 2 3 4)

> (cons 1 empty)
(1)
> (cons 1 (cons 2 empty))
(1 2)

> (list 1 2 3 4 5)
(list 1 2 3 4 5)
```

- A list can contain any data type

```
> (define lst (list 1 3 4 "Hello" 5 6 "Bangalore"))
> lst
(list 1 3 4 "Hello" 5 6 "Bangalore")

> (cons 1 (cons 2 (cons 3 (cons "Hello" (cons "World" (cons 4 empty))))))
(1 2 3 "Hello" "World" 4)
```

Lisp, Prolog

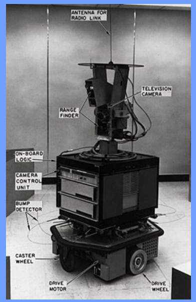

Shakey robot
STRIPS planning system

Fuzzy Logic
Rice Cooker

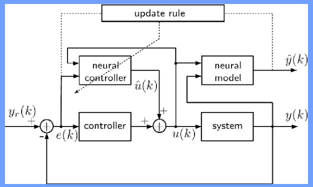

neural control

(consumer) robotics

self-driving car

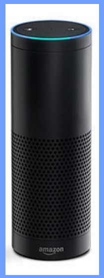

HomePods

1950's | 1960's | 1970's | 1980's | 1990's | 2000's | 2010's | 2020's

HTSC HIGH TECH SYSTEMS CENTER TU/e

Deep learning: Convolutional Neural Networks (CNNs)

- A CNN is a deep neural network optimized for processing image data
- CNNs became popular in 2012 when 'AlexNet' won the ImageNet competition
  - **13 layers** deep, ± **62mln parameters** to train
  - Trained the network on **15 mln annotated** images (**22,000 categories**)
  - **1.4 ExaFLOP** (= 1.4e6 TFLOP) needed for training
  - 7.8 GFLOP per forward pass per image
  - Trained **on two GTX 580 GPUs** for **5 to 6 days**
  - 1500 GFLOPS / GPU, 3000 GFLOPS total

  https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf
- CNNs are outperforming traditional CV systems in many visual domains

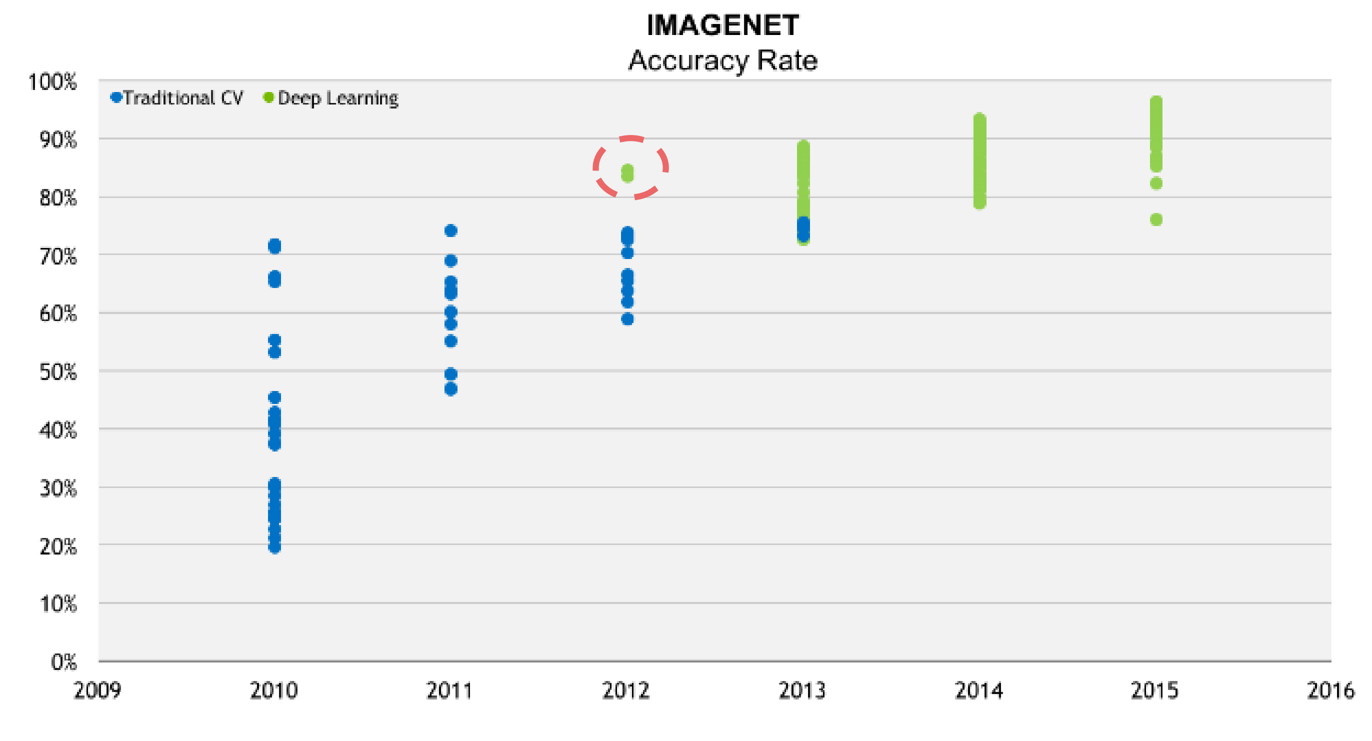

feature map

Source: 'Masterclass Deep Learning', https://vbti.nl/masterclasses/ , VBTI 2018

Deep learning: Deep Reinforcement Learning (DRL)

- RL is a machine learning technique that learns to optimize a reward by interacting with (virtual/real) environments.
- Performance of recent RL algorithms on games surpass human-performance on many games, including Atari games, Dota 2, poker, chess, Shogi, checkers, and Go.

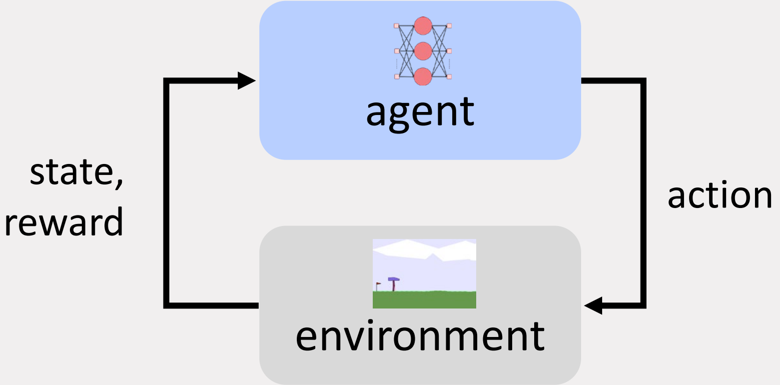

OpenAI gym benchmark environment

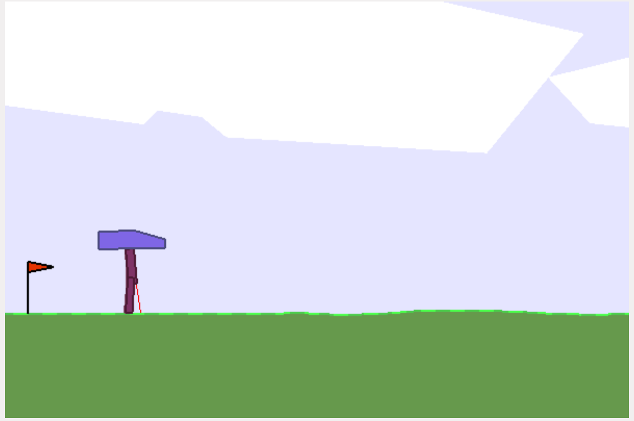

Agent not trained

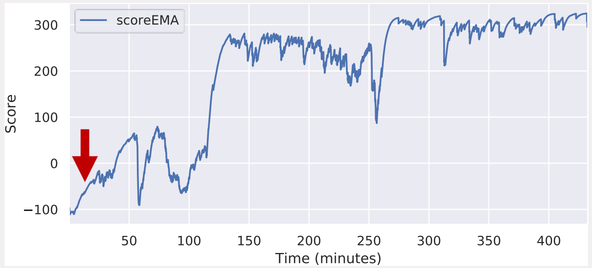

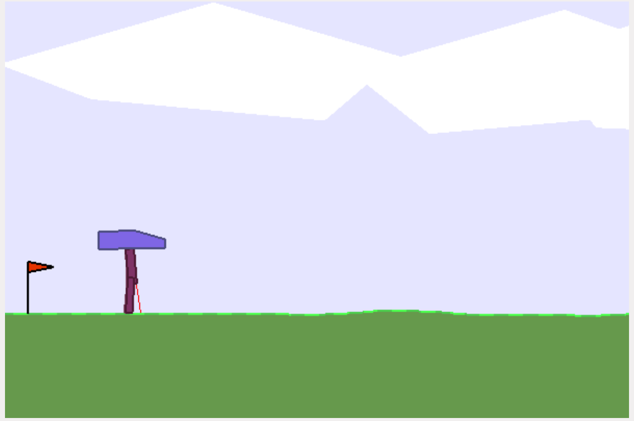

Agent trained after 8 hours

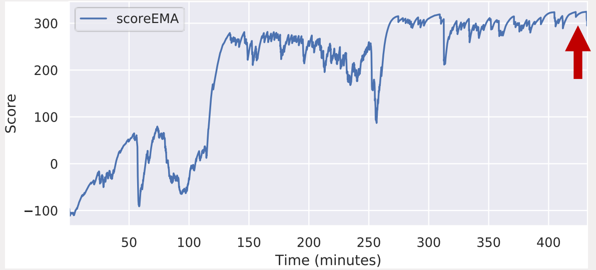

Source: ‘Masterclass Deep Learning’, https://vbti.nl/masterclasses/ , VBTI 2018

AI research & trends

- **Computer vision**: Deep Learning outperform human-expert based algorithms and are the ‘low hanging fruit’ for industry, for a.o. robotics, inspection and monitoring applications.
- **Natural language processing**: Deep Learning enabled mature level of speech recognition and text processing. Could be applied for developing machine interface (e.g. in combination with Augmented Reality) or for processing engineering documents.
- **Planning & Problem Solving**: By combining different Deep Learning techniques such as convolutional neural networks and reinforcement learning, new systems are build that ourperform humans in solving specific tasks (and games).
- **Robotics & Control**: Training AI algorithms on (virtual) Digital Twins of systems results in robust control solutions to problems virtually impossible to humanly solve.

AI solving challenges

CHALLENGE 1
Fingers of gripper are glued together and robot needs to pickup part.

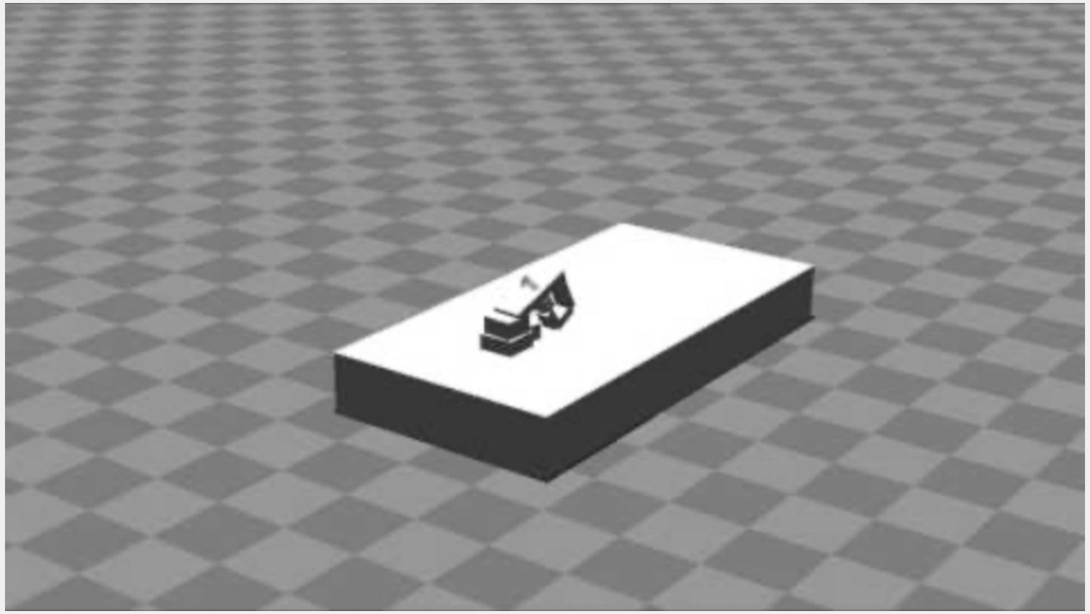

CHALLENGE 2
Robot need to walk around while minimizing ground-feet contact.

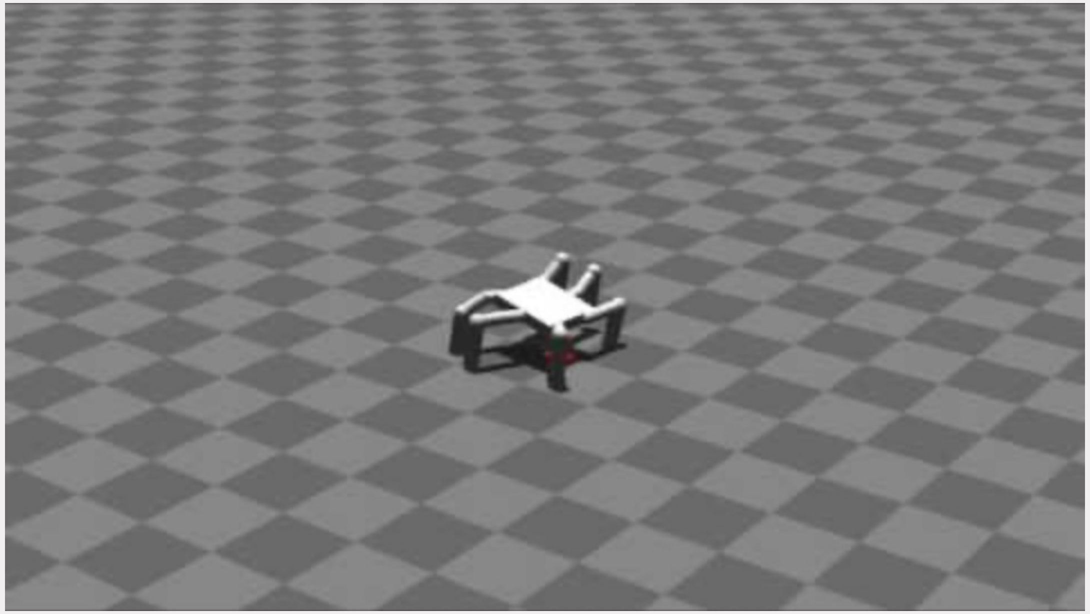

Robotics example: learning dexterity

Problem

- Robot hand should reposition an object
- High-dimensional control (24 DoF)
- Noisy and partial observations (friction, slippage)
- Manipulating more than one object

Source: https://arxiv.org/abs/1808.00177

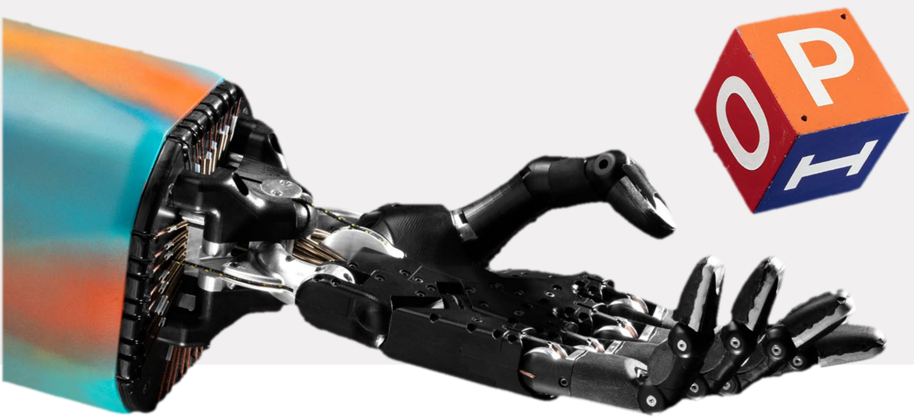

Solution

- Train in simulation
  - CNN to estimate object pose from 3 camera's
  - RL + RNN (LSTM) to determine fingertip actions
- Transfer to real world
- OpenAI's Rapid platform is used with 6144 CPU cores and 8 GPUs, collecting about one hundred years of experience in 50 hours.

Robotics example: learning dexterity

Problem

- Robot hand should reposition an object
- High-dimensional control (24 DoF)
- Noisy and partial observations (friction, slippage)
- Manipulating more than one object

Source: https://arxiv.org/abs/1808.00177

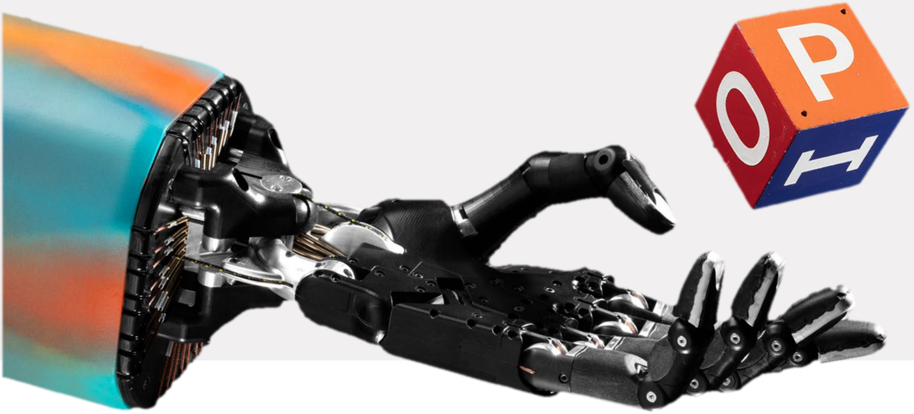

Result

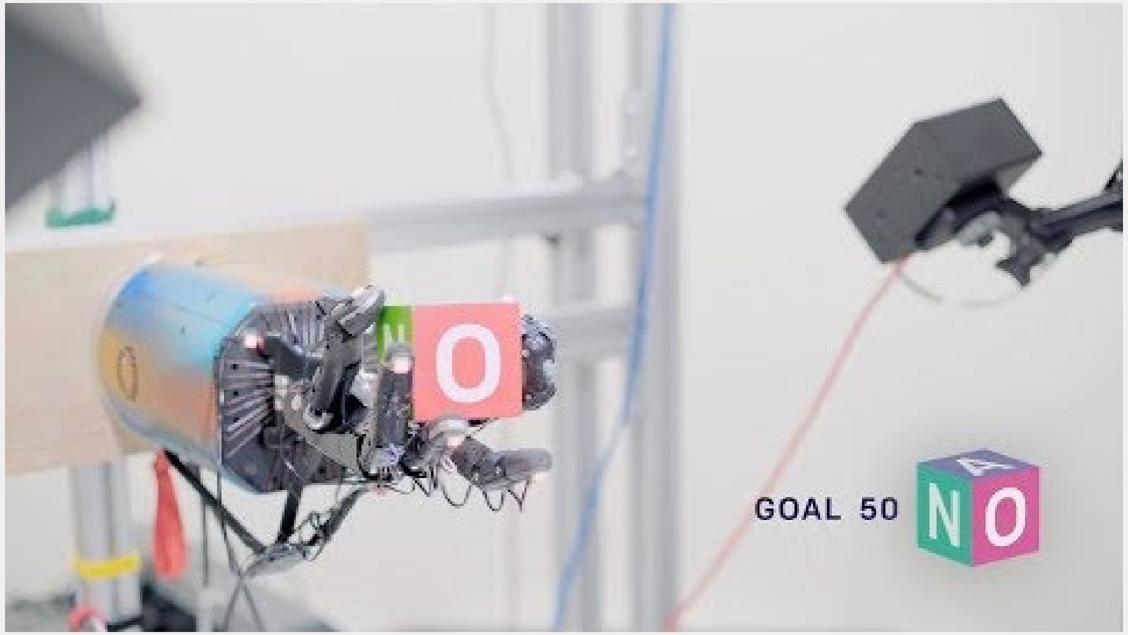

https://www.youtube.com/watch?feature=oembed&v=DKe8FumoD4E

Summary

- Artificial Intelligence technology is **maturing rapidly** the last few years and **outperforms human (solutions)** in various areas.
- There is a **gap** between the world of **engineering** and **data science** that can be bridged by:
  - **Combining** first principle techniques with machine learning
  - Making Deep Learning algorithms **explainable**
- AI performance is driven by **compute power**, **big data** and **algorithms**
- Still **many barriers** need to be taking by industry to adopt AI, a.o.
  - Complex AI technology stack
  - Many number of different AI algorithms (which to use when?)
  - Transfer results from virtual to real world

AI Engineering Lab

Support industry with adopting AI technology and create research consortia

- Part of the Eindhoven AI Systems Institute (EAISI)
- Multidisciplinairy
  - Brings together researchers across the areas of mechatronics, data science, mathematics, computer science, computer vision, embedded systems, robotics, control engineering and more.
    - Partners include FME

aie-lab@tue.nl

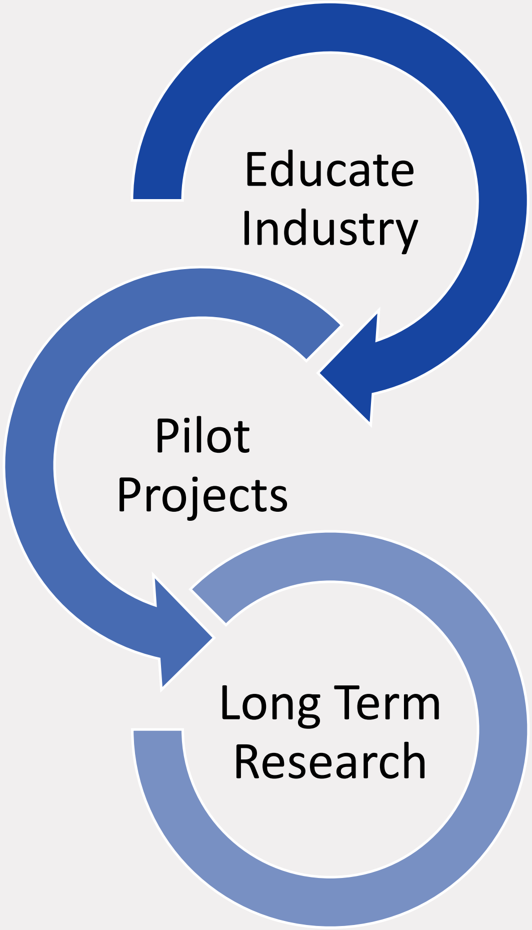

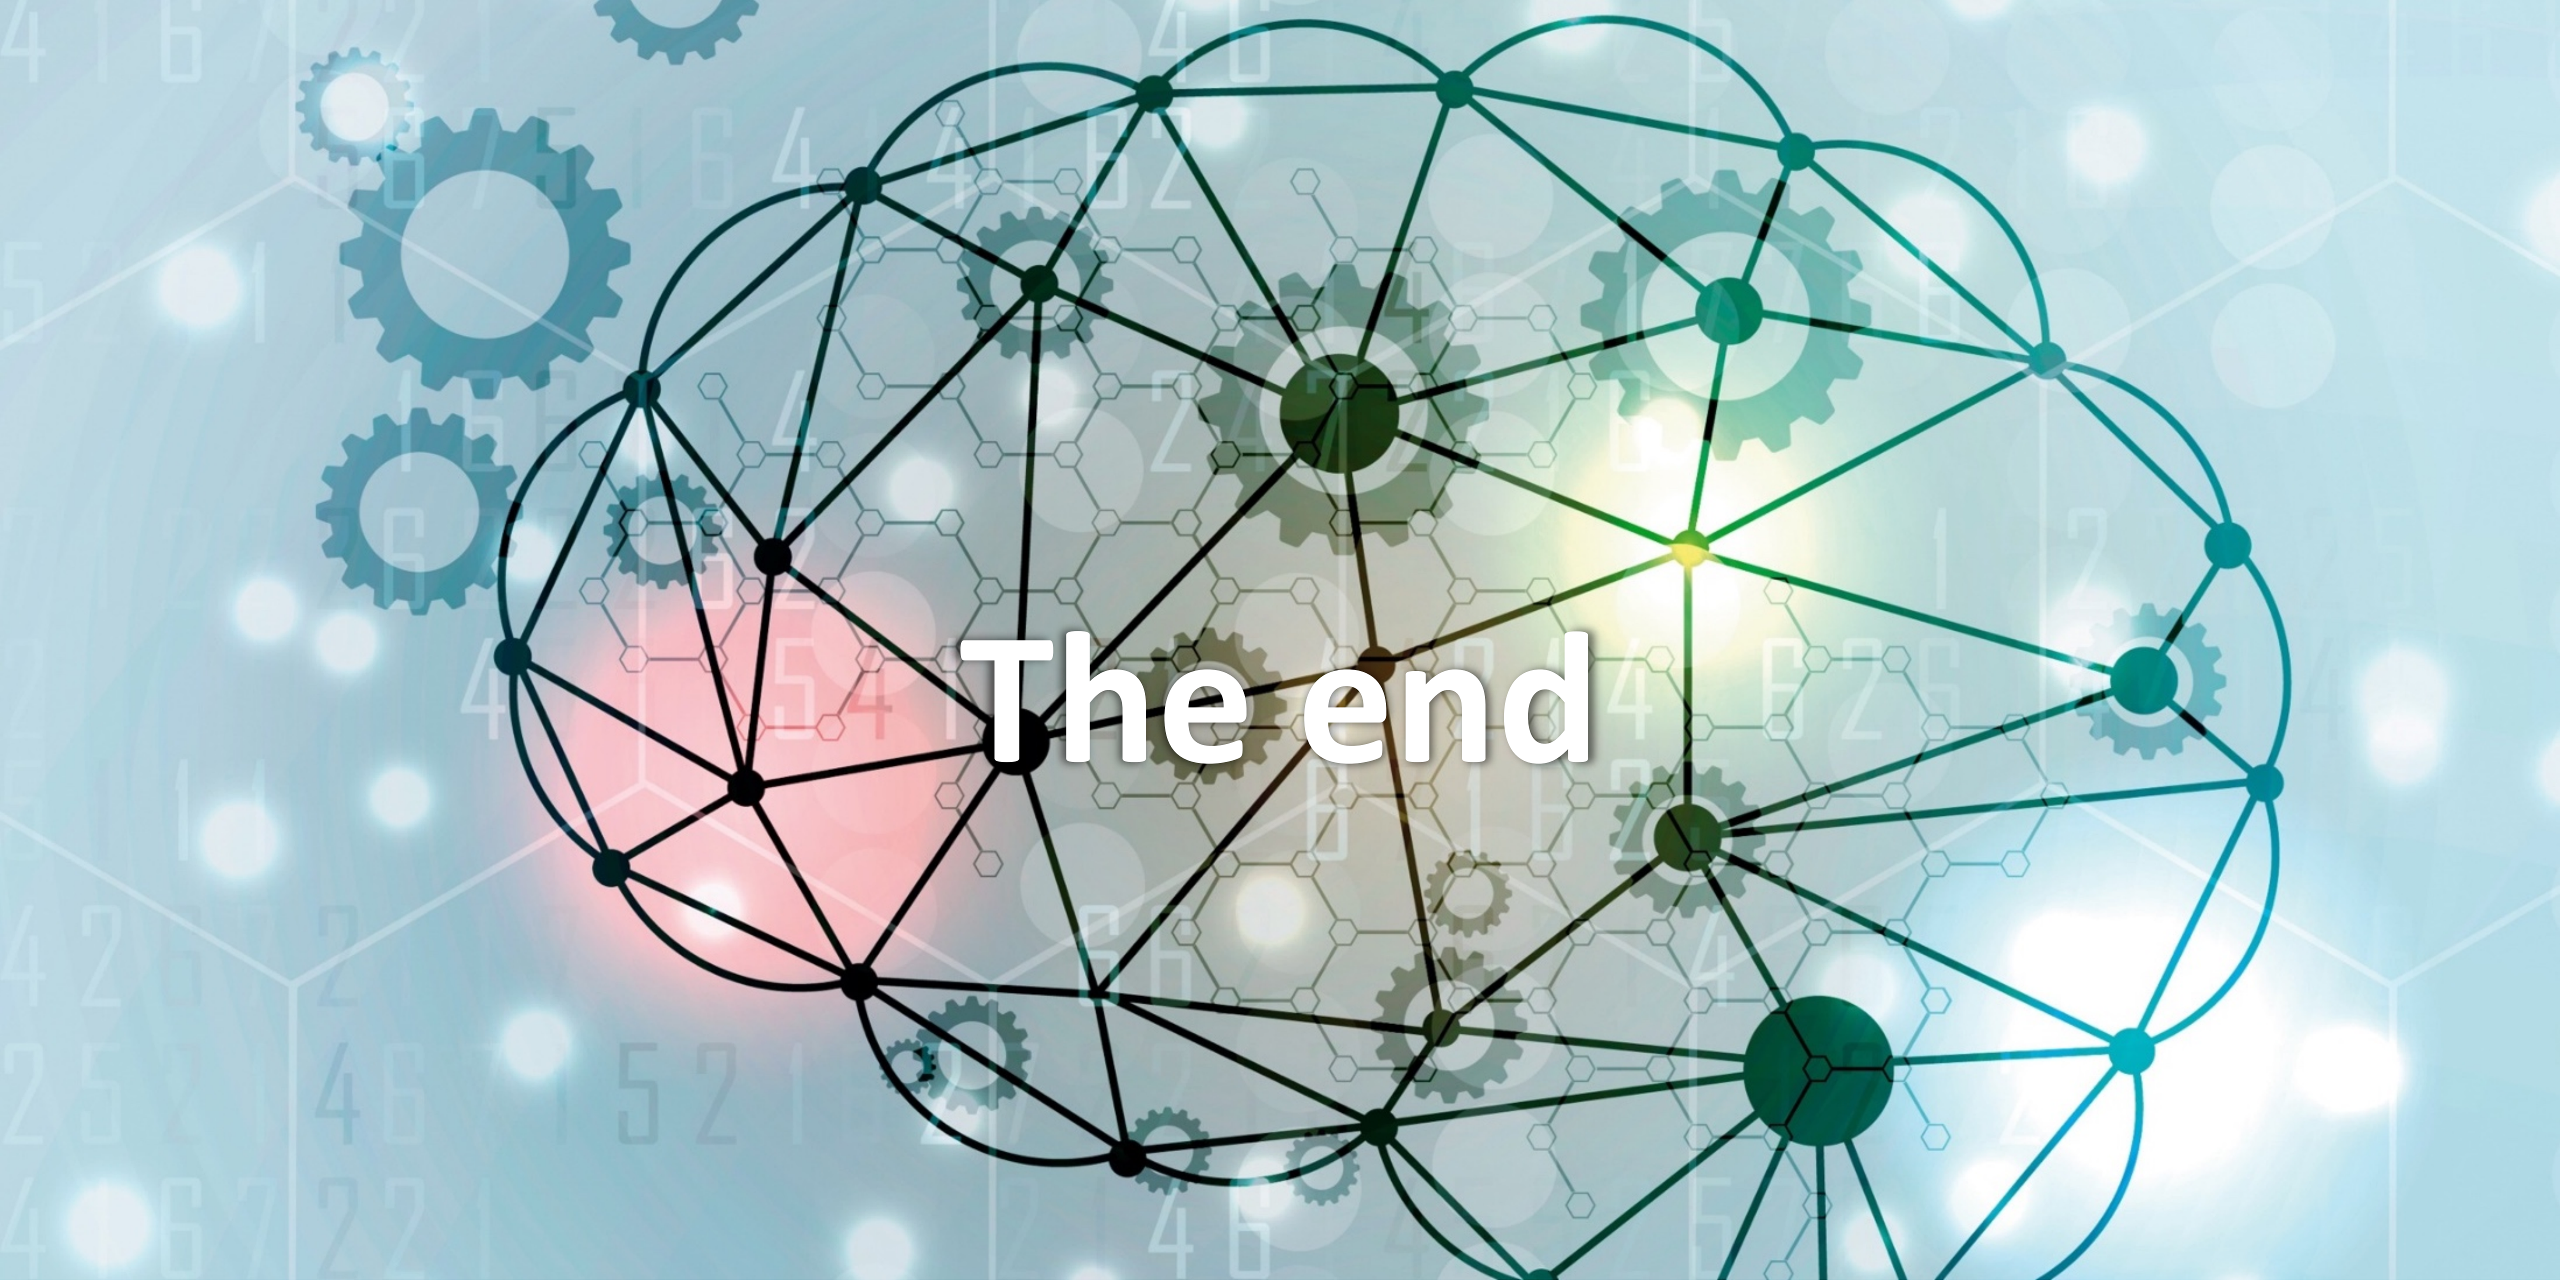

The end

Deep learning: Recurrent Neural Networks (RNNs)

- RNN uses feedback loops from the output of a layer to the input of a previous layer
- RNN have a memory for past events
- RNN can learn correlations between signals at different time moments
- A special type of cell structure, called Long Short Term Memory, enabled learning
- RNN's are success behind many speech recognition and natural language processing systems.

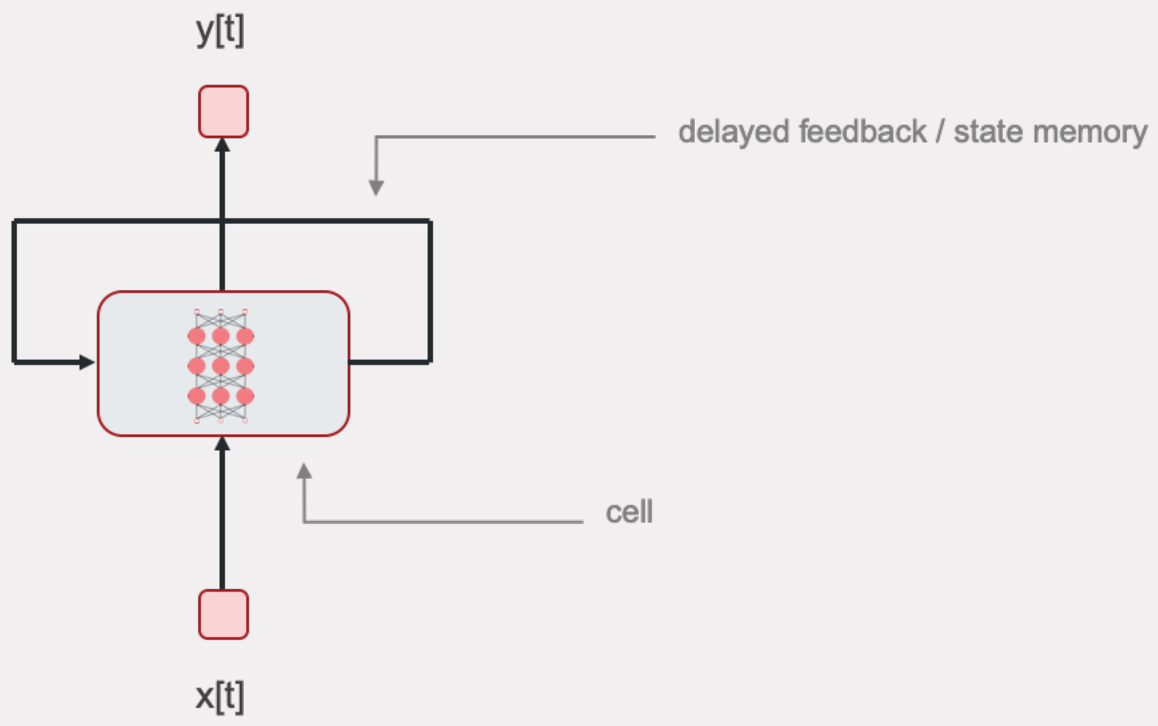

Source: 'Masterclass Deep Learning', https://vbti.nl/masterclasses/ , VBTI 2018

Drone example: learning to fly

Problem

- Learn a drone to navigate indoor and avoid objects
- Reduce gap between simulated and real-world training

Source: https://arxiv.org/pdf/1704.05588.pdf

Solution

- Make robust drone to collect real-world data to learn from
- Use pre-trained CNN (AlexNet, transfer learning) to map camera input to actions
- AR Drone 2.0 flew in 20 different indoor environments, racking up 11,500 collisions over the course of 40 hours of flying time.

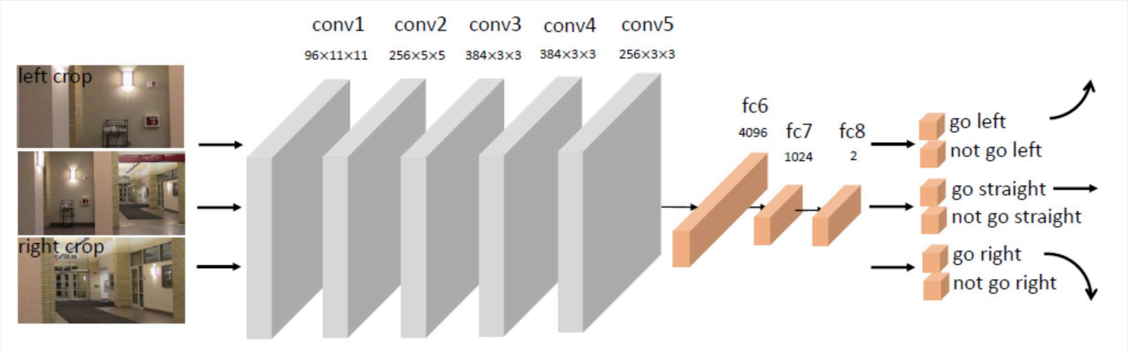

Drone example: learning to fly

Problem

- Learn a drone to navigate indoor and avoid objects
- Reduce gap between simulated and real-world training

Source: https://arxiv.org/pdf/1704.05588.pdf

Result

https://www.youtube.com/watch?feature=oembed&v=HbHqC8Himol

Metrology example: improve semicon production yield

Problem

- Reducing the influence known contributors have toward the on product overlay budget when moving semicon manufacturing toward 7nm node for logic and 15nm node for memory (2015!)
- Inline metrology is costly (duration)

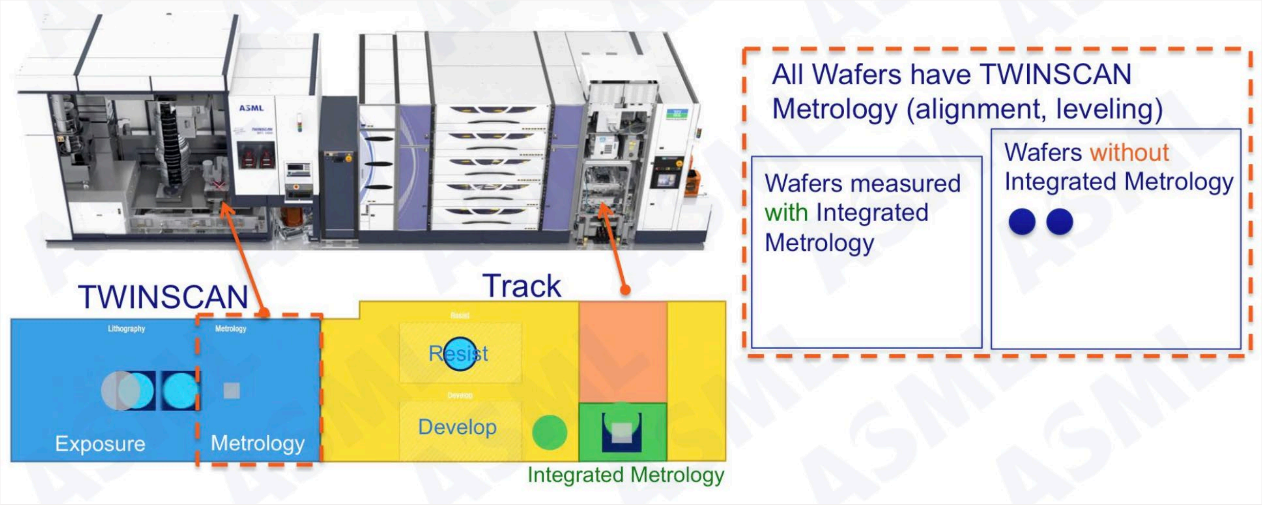

Solution

- Use **TWINSCAN metrology** to measure all wafers and use **inline metrology** to measure overlay for some wafers
- Train a model to approximate overlay vector maps for the entire lot of wafers.
- With the approximated overlay vector maps for all wafers coming off the track, a process engineer can redirect wafers or lots with overlay signatures outside the standard population to offline metrology for excursion validation.

“Virtual overlay metrology for fault detection supported with integrated metrology and machine learning”, Emil Schmitt-Weaver, MATLAB Expo 2016 Benelux – June 28th

“Virtual overlay metrology for fault detection supported with integrated metrology and machine learning”, Hong-Goo Lee et al., SPIE, 19 March 2015

Metrology example: improve semicon production yield

Problem

- Reducing the influence known contributors have toward the on product overlay budget when moving semicon manufacturing toward 7nm node for logic and 15nm node for memory (2015!)
- Inline metrology is costly (duration)

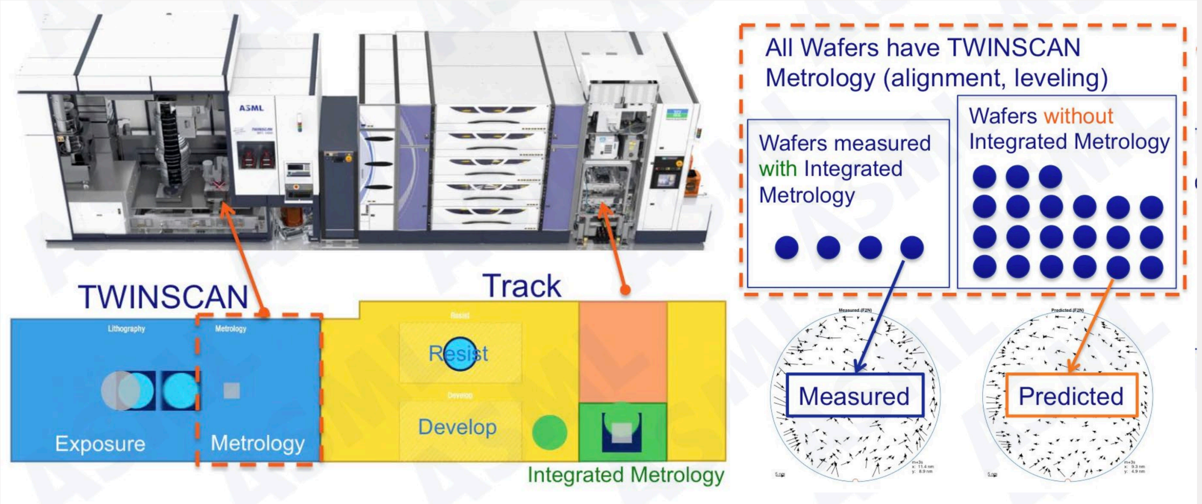

Result

- The model identifies systematic and random overlay errors, improving overlay performance.

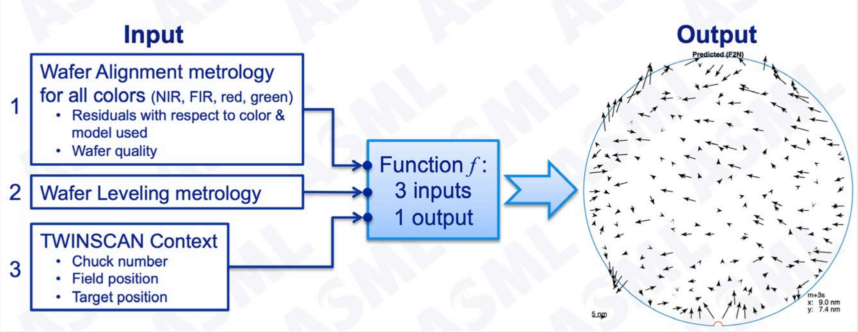

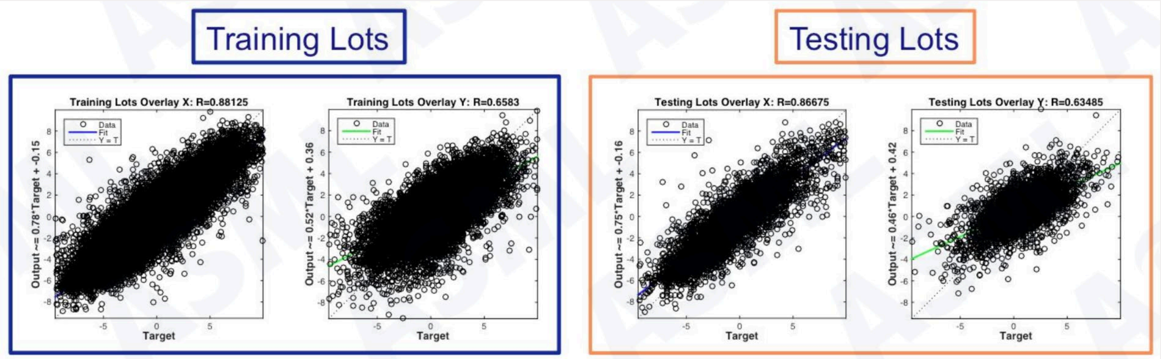

Problem solving example: generative design

Problem

- Optimizing products by taking design and manufacturing constrains simultaneously into account.

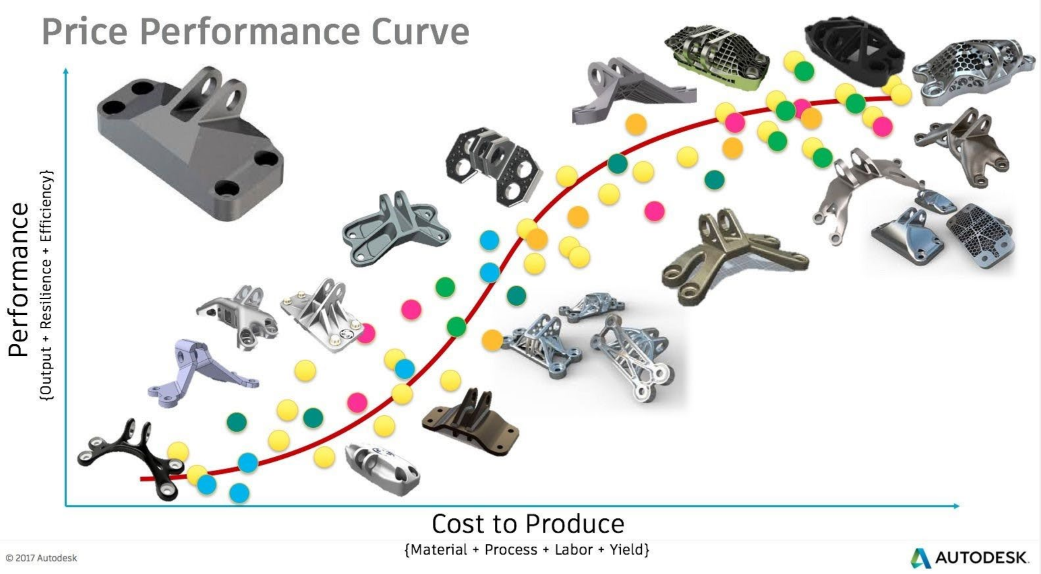

Solution

- Use AI as a tool during the engineering design process
- Autonomously creates optimal designs from a set of system design requirements.
- Engineers can interactively specify their requirements and goals, including preferred materials and manufacturing processes—and a **generative engine** will automatically produce a manufacture-ready design as a starting point or as a final solution.

Source: https://www.autodesk.com/solutions/generative-design/manufacturing

Problem solving example: generative design

Problem

- Optimizing products by taking design and manufacturing constrains simultaneously into account.

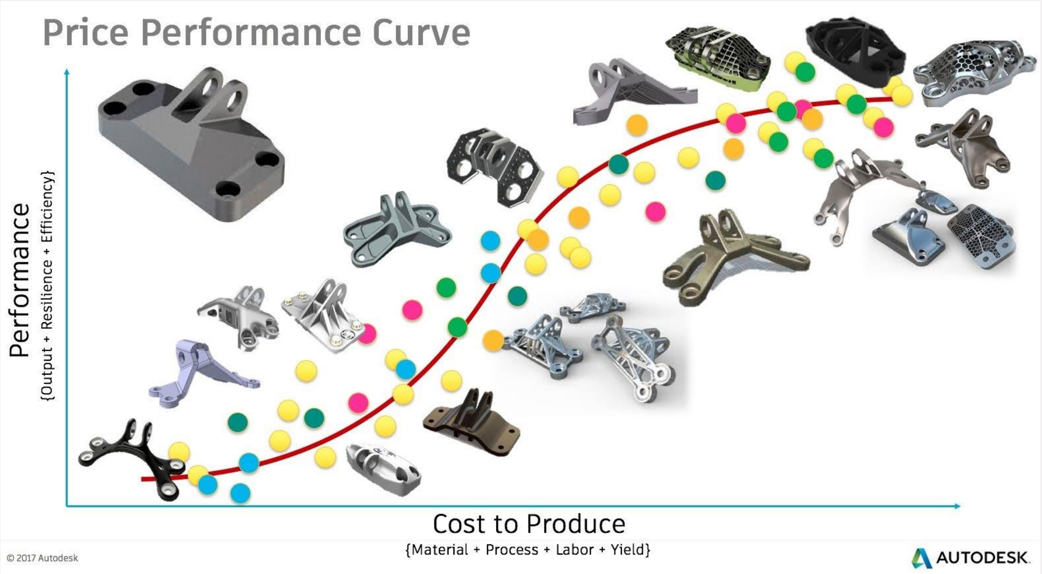

Result

https://www.youtube.com/watch?feature=oembed&v=vtfNlWEJxw4

HTSC HIGH TECH SYSTEMS CENTER TU/e