EPFL EXTENSION SCHOOL

AI Tutorial

Rémi SABONNADIERE

Analytics & Marketing Manager

Presented to: ITU

Date: 11/10/2019

Plan

Part 1 - Overview of AI

- Why AI has rised up these last decade ?
- What is AI? Narrow vs. General AI
- Some of today's major AI applications

Part 2 - Applications

- The importance of data
- AI in practice
- AI's impact on society and economy

Part 3 - How to become an AI driven organization ?

- Why jobs evolve ?
- The Organizational challenge
- Education is key

What is AI?

Mathematical techniques exist since decades

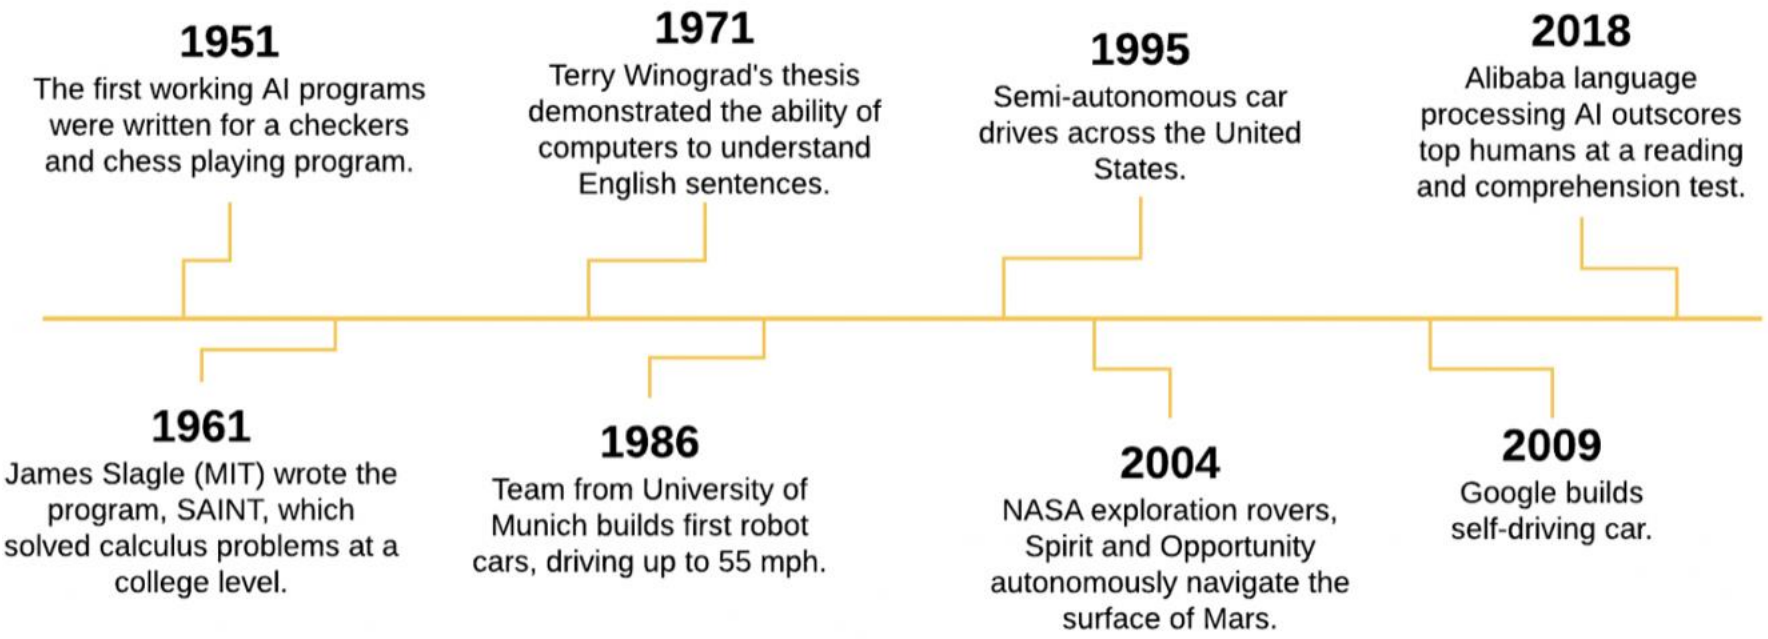

Breakthrough comes from data volume and quality

- 2.5 quintillion bytes of data created each day
- The last two years, 90 percent of the data in the world was generated
- 3.7 billion humans use the internet (7.5% CAGR)
- Emergence and democratization of sensors technologies

The 4Vs rule

Data Volume

Data Variety

Data Velocity

Data Veracity

Why data volume matters ?

Does the maximum speed of a vehicles increased with horse power of an engine ?

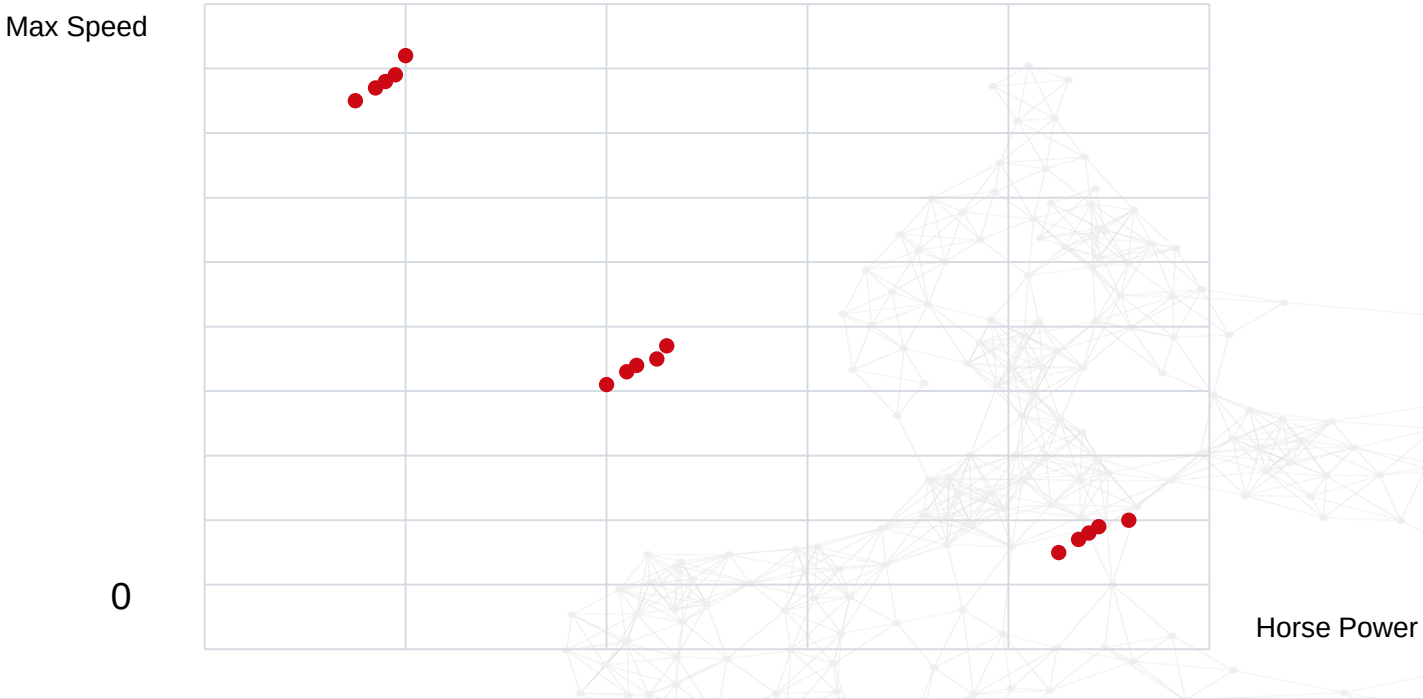

Why data volume matters ?

Predicting the maximum speed of vehicles with respect to horse power of an engine

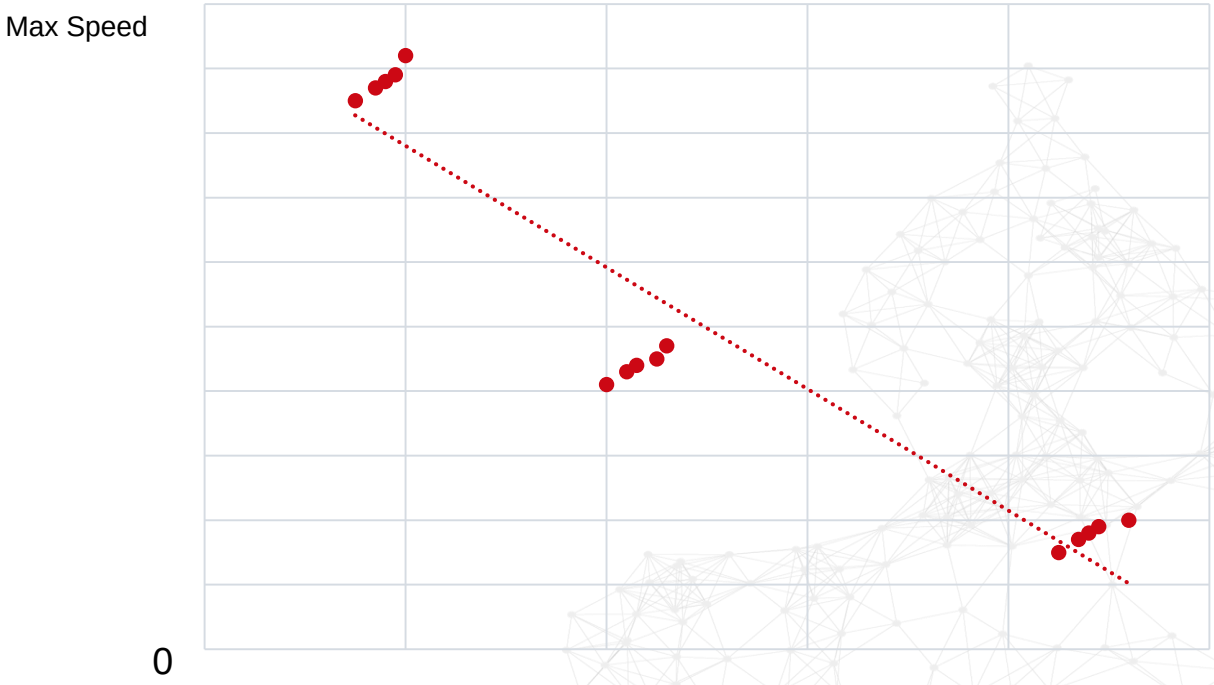

NO ?!

Why data volume matters ?

Does the maximum speed of a vehicles increased with horse power of an engine ?

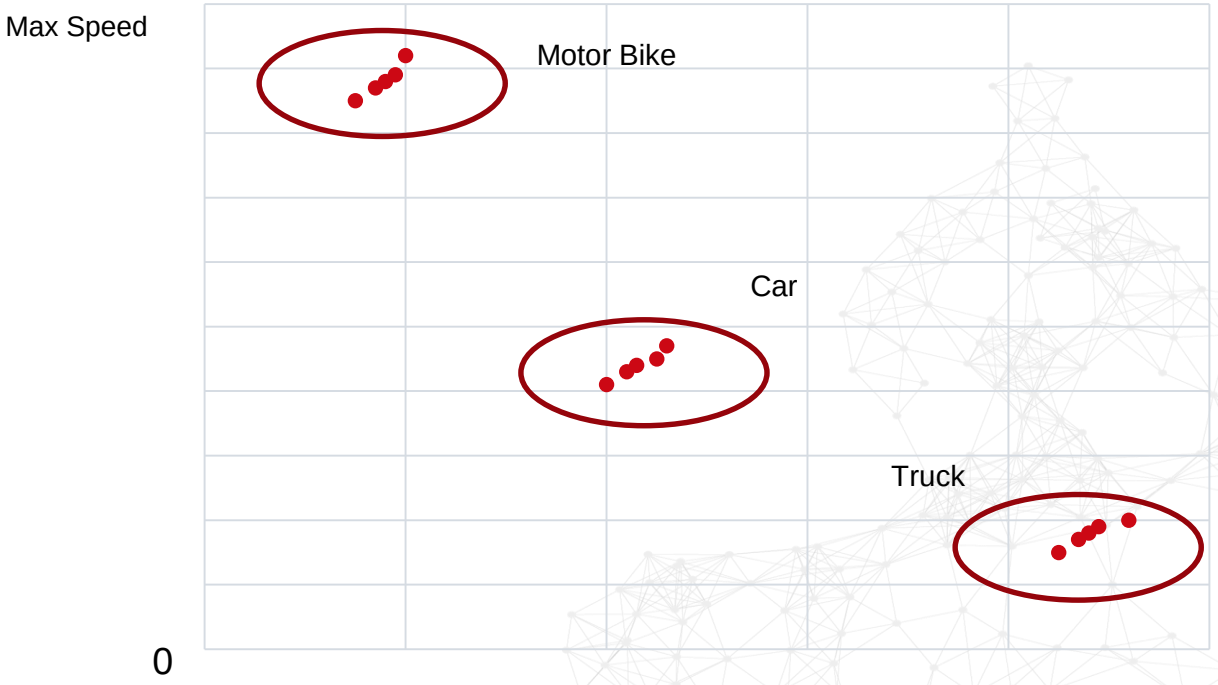

Why data volume matters ?

Does the maximum speed of a vehicles increased with horse power of an engine ?

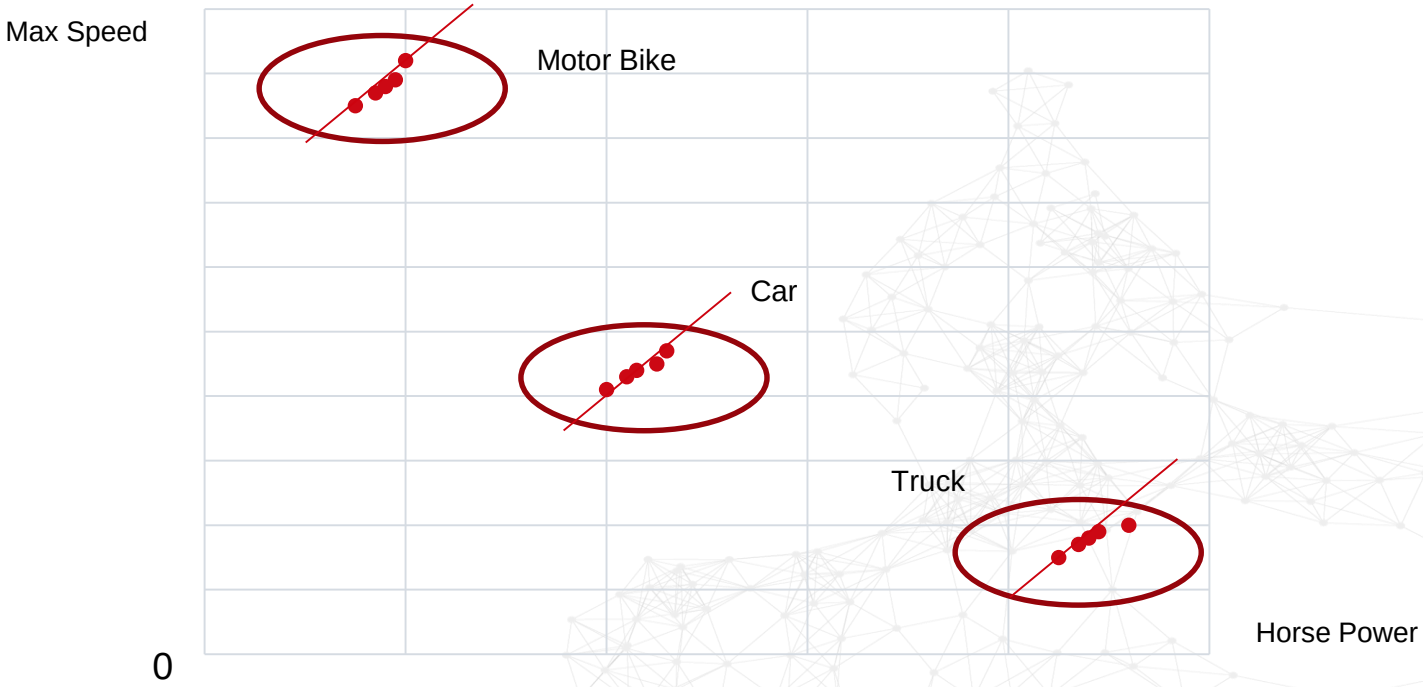

Why data volume matters ?

We can repeat the exercise within the car category for instance (SUV, 4*4) and prediction will get better and better until sample is too small

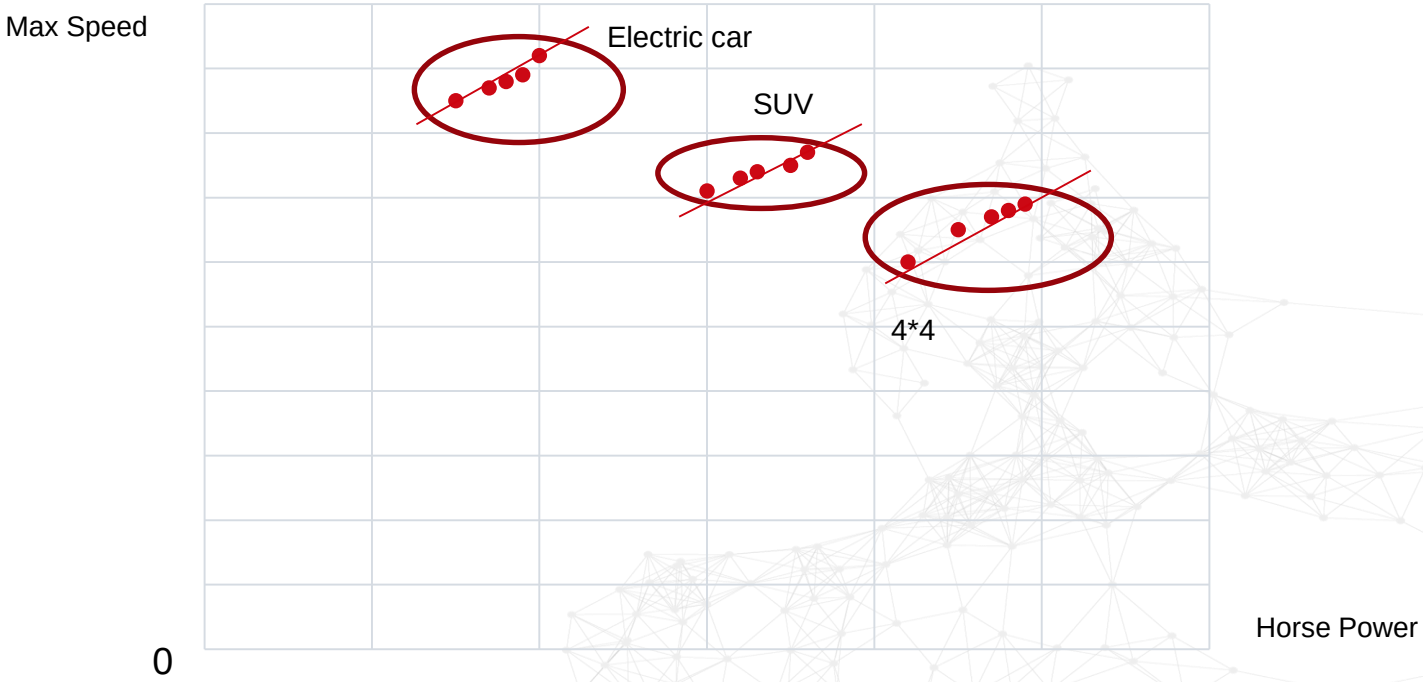

Terminology of AI - Narrow vs. General AI

Narrow artificial intelligence

- what we currently call “AI”
- examples: self driving cars, speech recognition, smart advertising
- effective on only one specific problem
- no generalization beyond the task

General artificial intelligence

- what Hollywood calls “AI”
- can do anything what a human can do
- translates experience to unseen tasks (holistic view)

When will we have General AI?

- *might still be 100 years away* - Andrew Ng
- *the intelligence singularity will be reached in 2045* - Ray Kurzweil

Terminology of AI - Machine Learning

Programming
- instruct the machine how to carry out a specific task
- steps are defined by human expert

Machine Learning (ML)
- let a machine learn how to best complete a specific task
- steps are learned from data / experiences

Current state of machine learning
- human provides goal, framework and data
- main focus of AI, but only a subset thereof
- fast evolving (exponentially)

Terminology of AI - Machine Learning

Concretely, how do you distinguish them?

- Programming - handcrafted rules, recipe designed by a human expert
- Machine learning - automatically learns how to solve the task from data

2006 - Honda ASIMO

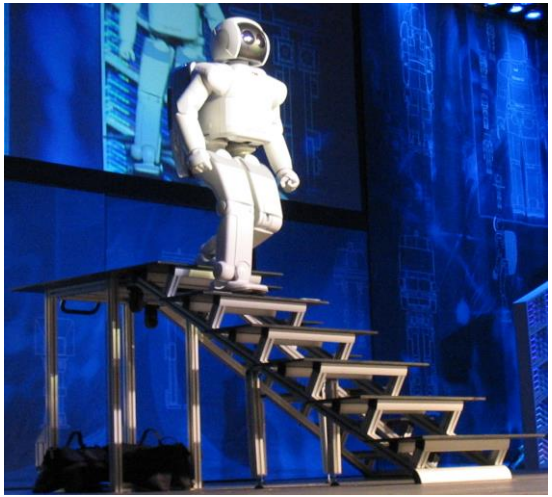

2018 - Boston Dynamics

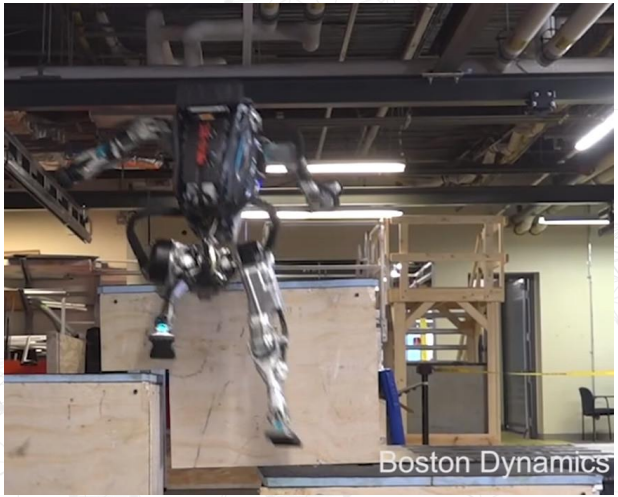

Terminology of AI - Neural Networks and Deep Learning

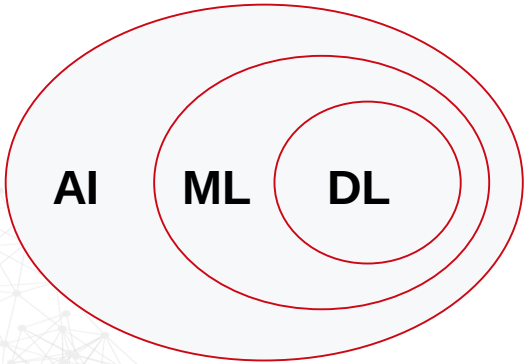

Neural Networks

- framework for deep learning
- historically inspired by brain
- many different architectures

Deep Learning (DL)

- Multi-layer Neural Network that progressively extracts higher level features from raw input
- ML learns to solve a task, while DL learns the steps to solve a task
- Examples: image classification & segmentation, speech recognition, text analysis

Context

- Concept of Neural Networks exists since early 1940s
- 2010 “deep learning” era due to data availability and computational power
- driving the current AI hype

Terminology of AI - Types of Machine Learning

Supervised Learning

- Learn with the help of the correct answer
- **Example** - Detect patients that are likely to develop a disease

Unsupervised Learning

- Learn the structure, there is no correct answer
- **Example** - Discover groups with shared risks

Reinforcement Learning

- Learn actions that lead to some reward
- **Example** - Improve treatment by successive iterations

Examples of what AI can do

Simple cognitive tasks (things we can perform within a few seconds)

Recognize a multitude of objects in an image (ex. Mask RCNN)

Language translation (ex. Google Translate)

Speech to Text (ex. Siri, Alexa, Google Assistant)

Smart and personalized assistant (Google Duplex)

… many more

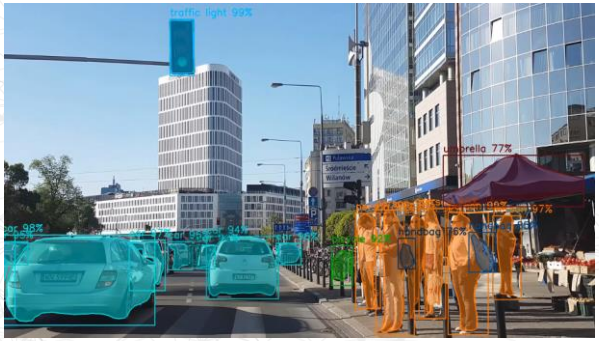

What is likely to be achieved soon?

Karras, Laine, & Aila (2018) by NVIDIA

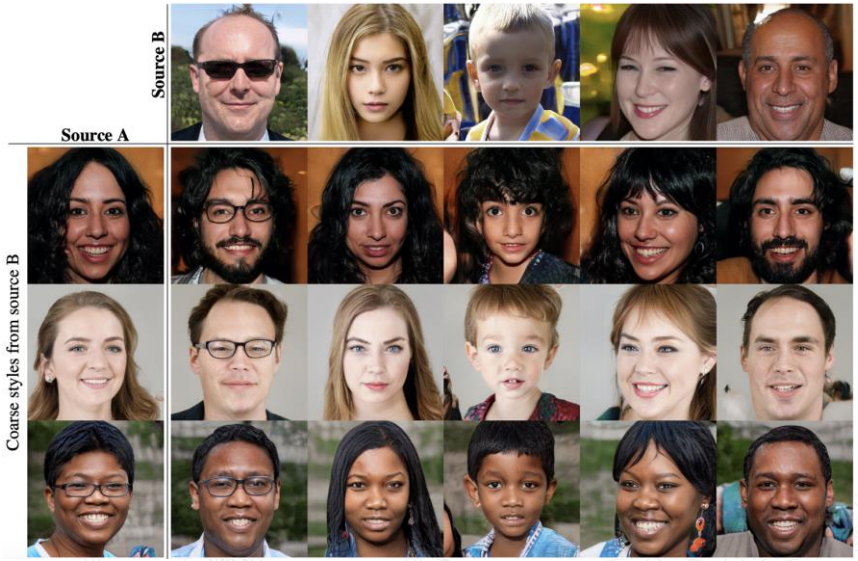

Simultaneous speech translation

- End-to-End Speech-to-Speech Translation
- Preserving Vocal Characteristics
- Modulating emotional prosody

Content generation

- text, speech, images, video, chatbots
- ex. thisPersonDoesNotExist.com
- ex. Text editing of video

Personalized education

- adapt teaching technique to learner
- tailor exercise difficulty to learner

Fully autonomous driving

- Society of Automotive Engineers (SAE) defines 5 levels of automation
- Tesla currently is at level 2 (partial automation)
- Level 4: car can handle all aspects of driving in certain conditions

What AI cannot do for now?

High-level cognitive tasks

- Reasoning from complex, unstructured, abstract knowledge
- Abstracting - have a global view
- Argument, summarize
- Translate experience to unseen tasks (ex. read upside down)

Examples

- Having an open conversation
- Make complex hypotheses / invent new things
- Recognize its own limitations
- Have free will or make moral choices

Some of today's major AI applications

Software industry

Recommendation systems ex. ebay, amazon, netflix, youtube

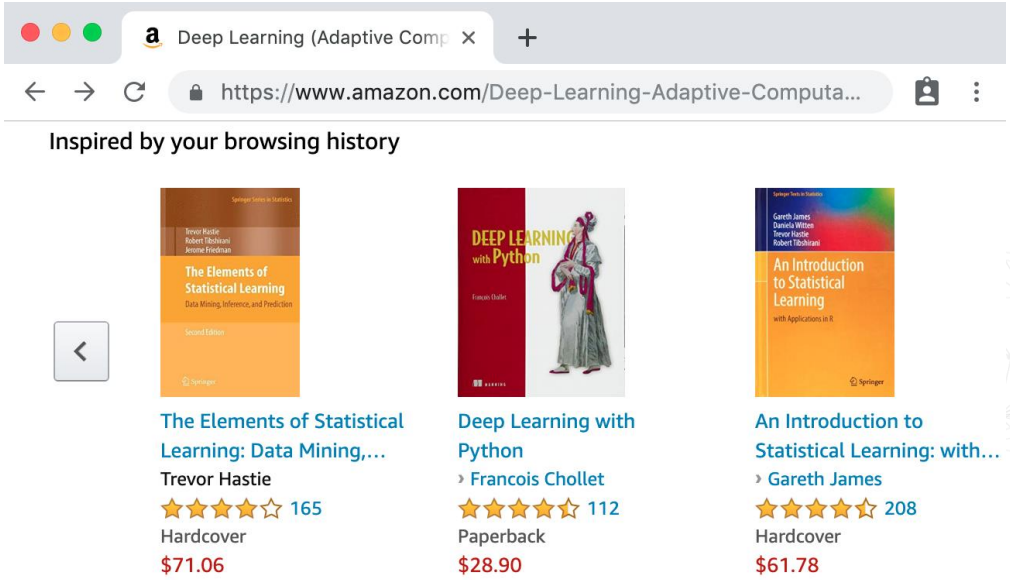

Software industry

Additional examples

- Spam filtering
- Smart assistants ex. smart reply, compose in Gmail
- System security ex. intruders, virus detection

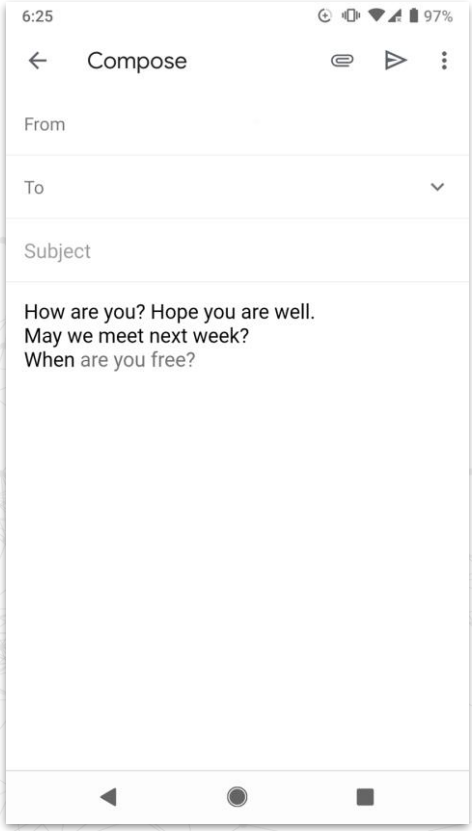

Humanitarian Assistance and Sustainability

- Predicting earthquake based on seismic data
- Predicting / modelling of famine
- Rare animals detection
- Energy distribution efficiency

Communication

Domains involving Natural Language Processing (NLP)

- Machine translation
- Document understanding and classification
- Personal Assistants and Chatbots

Text and speech synthesis

- Dialog Assistants
- But also, fake and targeted news generation

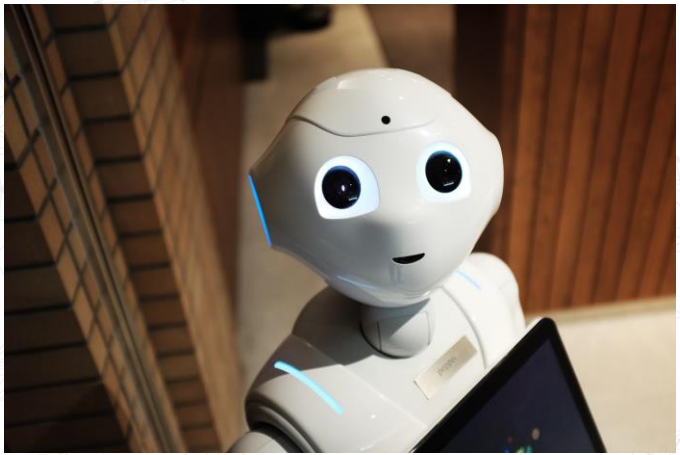

Assisted Content Creation

Entertainment industry

- Film - Actors controlling virtual avatar
- AI for video games

Generative Adversarial Networks (GANs)

- GauGAN by NVIDIA Research
- Pix2Pix network

Travel

- Predicting delays, traffic jams
- Planning optimal routes

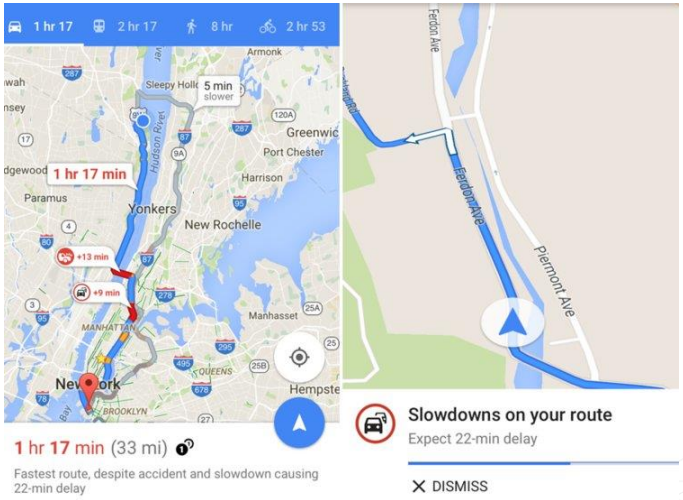

- Supply and demand forecasting
- Dynamic real-time pricing

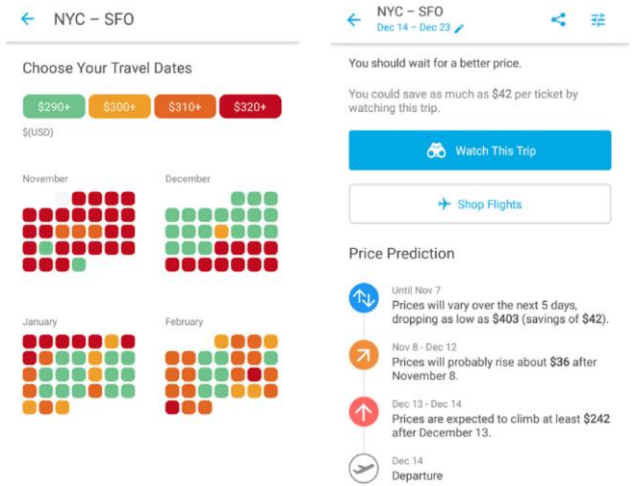

Transportation

- Autonomous transportation
  ex self driving cars, drive assist

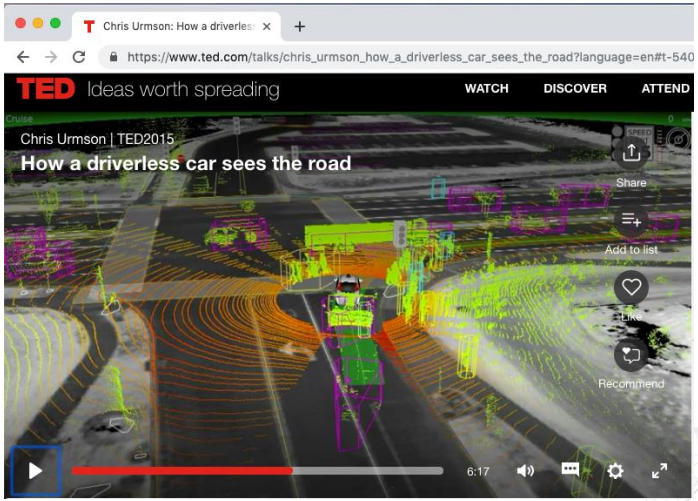

- Connected or collaborative driving
  including machines

Transportation

- Optimize flight routes
- Adaptive scheduling
  ex reacting to heavy storms

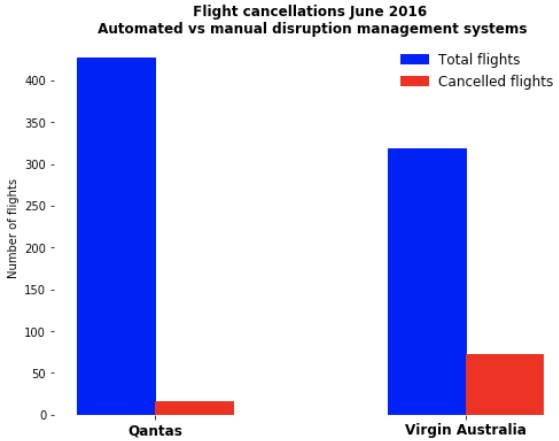

Insurance

- Pricing premiums
- Personalise products
- Claim processing
- Fraud detection

Insurance

- Remote damage estimation and assessment

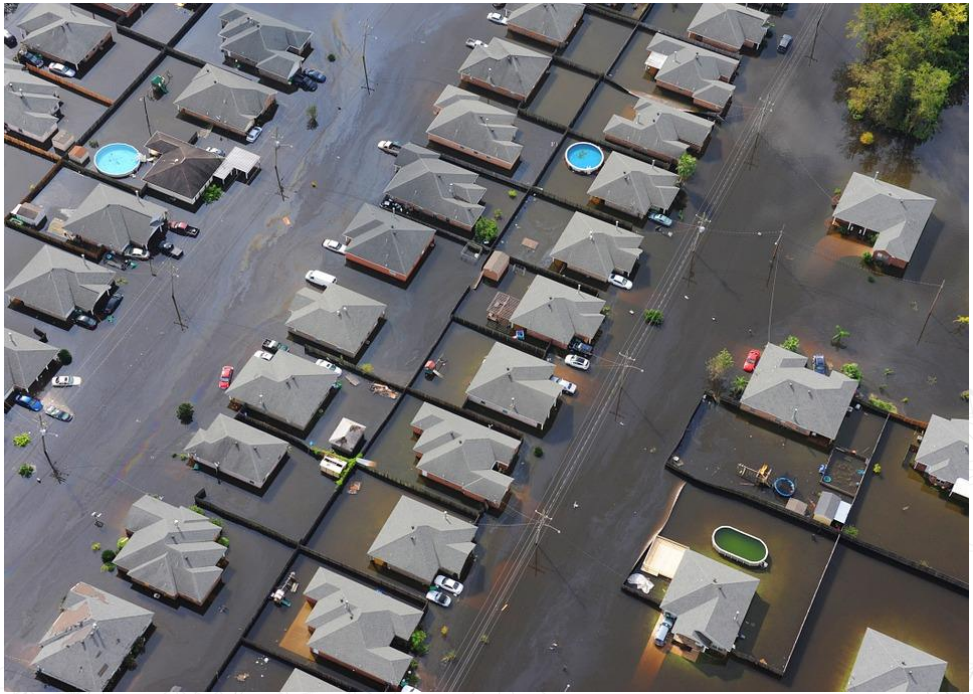

Energy sector

- Forecast demand and supply side
  - Predict production of solar and wind from weather data
  - Autonomous grid management (feed into grid vs storage)

- Smart energy consumption
  - Adapt to users behavioural patterns
  - Adapt to prices
  - Improved efficiency, ex Google servers

- Detecting and assessing gas and oil deposits from seismic images

Farming

- Animal health and welfare monitoring
- ‘AI assistants’ for veterinarians

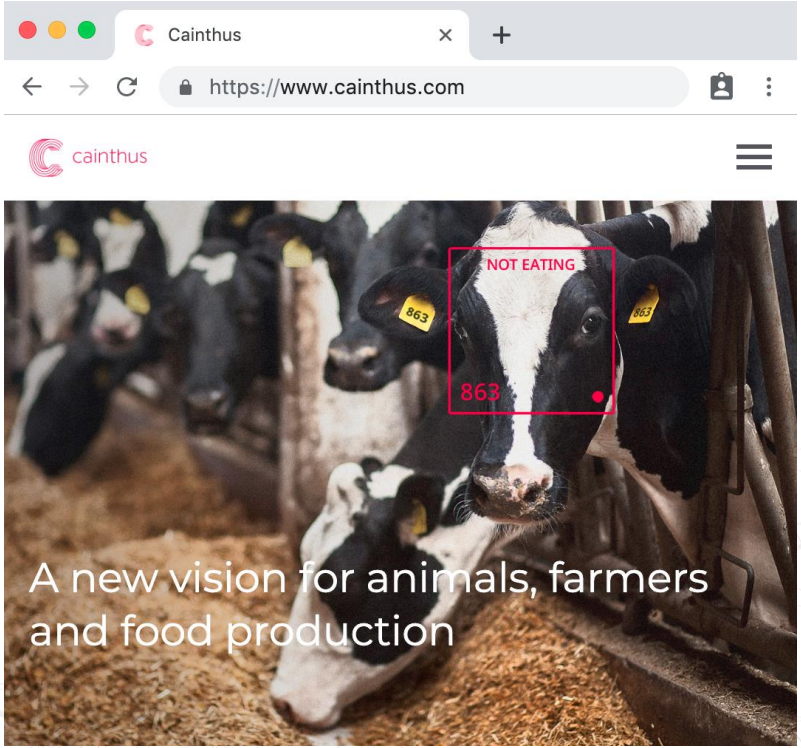

Farming

- Crops disease monitoring and targeted interventions i.e. pesticides, fertilizers
- Filtering low quality grains, foreign bodies from images
- Determine ideal harvesting time
- Yield prediction

Maintenance

- Predictive maintenance
- Failure anticipation
- Service scheduling
- Product inspection, e.g. detect defects from images or sensors

Healthcare

Automated diagnostics

- Lack of access to doctor
- Anonymised access

Illness detection

- Based on a large collection of past cases
- Could be used as additional indicators
- Better understand the factors

Support

- Emotional support through avatars

Google DeepMind AlphaFold

Guessing how a protein folds from its genetic sequence

- Team of AI experts, external to the field
- Few prior knowledge about the task

Real impact in treating diseases

- Alzheimer
- Parkinson
- Huntington

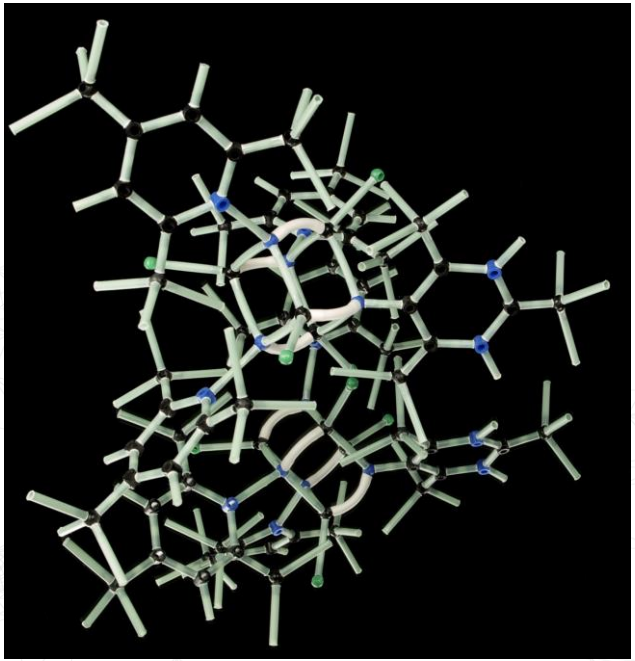

The importance of data

Ingredients of Machine Learning

ML = data + computational power + expertise

It takes all three to do Machine Learning well

“Data is the new oil”

- Need for quality data
- Need for quantities of data
- Data pipelines

- Data has become a commodity
- Not a limited resource

Data Quality, Consistency and Accessibility

- 80% of the work of a data scientist is data cleaning
- Different data sets are inconsistent with each other
- Limited amount of open data available (incl. scientific data)
- Don’t have a “Web of Data”

Data as a limiting factor

- Lack of quality and quantity result in poor results
- Lack of access prevent new AI services from being developed
- Data may not even exist yet
- General AI needs high quality data across all tasks
- Transfer Learning make use of already learned tasks

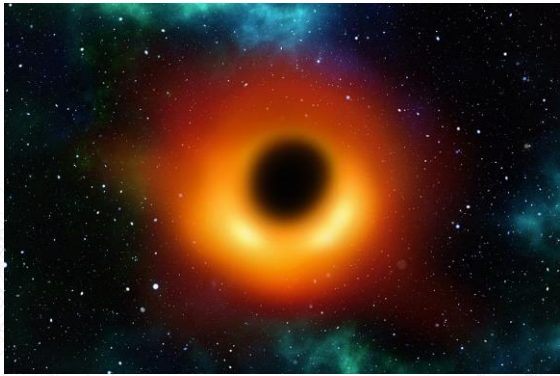

Examples of where AI goes wrong

- AI does not answer your question
  - Recognising camouflaged tanks? No, checking the weather!
- A technical blackbox sensitive to changes
  - Self-driving car misreads STOP sign with ‘graffiti’
  - Face recognition software fooled by masks or stickers
- Stereotypes and biases hidden within data
  - Biases discovered in face recognition software
  - From friendly millennial chatbot to racist within 24hrs
  - Gender-dependent association of words learned

How is AI done in practice?

Data Science, Machine Learning, AI?

The Data Science Pipeline

- Data preparation
- Data exploration
- Analysis and modeling
- Communication and reporting

Analysis and modeling (the AI part)

- Statistical analysis
- Machine Learning
- Deep Learning

Preparation

Exploration

Analysis

Communication

Overview of AI tools

Data science

- Python, R
- Jupyter notebooks

Machine Learning

- Scikit-Learn
- Custom platforms

Deep learning

- Google TensorFlow, Keras
- Facebook Pytorch

What is AI’s impact on society and economy?

What is AI bias?

- Biases
  - Statistical bias - systematically missing the intended target
  - Data science bias - data that is not representative of the population
  - Human biases - over 180 defined and classified

- AI bias = All of the above
  - AI is statistical learning
  - Imbalanced data leads to poor performance on rarer examples
  - Human biases are reflected in the data we create and use

- Blame game
  - Headlines: “AI found to be racist/elitist/misogynistic/…”
  - Humans responsible for the input, AI just executes

The “Black Box” problem: Interpretability and trust

- AI are vastly complex systems with millions of parameters dependent on data used for training
- So how do we know what an AI system really learned?
  Did the AI system learn as intended?
- How can we understand/explain decisions made by AI?
  Should we trust them?

Some of the pitfalls of using AI

- Overreliance on AI tools
- Assigning human traits to AI
- On the user side: Lack of understanding/awareness of limitations
- Unintended consequences for society

Ethical and moral questions around AI

- Who does AI serve: individuals, society, companies ?
- Is today’s data collection an automated version of stalking?
- AI designed to manipulate or harm
- Who is responsible for the failure of AI and its consequences?
- Different standards for AI and humans

AI and personal data

- Personal data can be analysed and used to predict and potentially manipulate people's behavioural patterns
- Combining unrelated data sets containing personal data makes it possible to identify the individual and hence is subject to data protection
- Countries are concerned about uncontrolled collection, storage and processing of personal data which could violate people's right to data protection
- EU Law: "Everyone has the right to the protection of personal data concerning him or her."
- Countries have put in place regulations on data protection and the processing of data where it is possible to directly or indirectly identify people.

Why everyone needs to learn how to use data – AI and jobs

Overall

- By 2030, DL based AI is estimated to add 13 trillion US dollars of value annually
- AI to influencing every sector of the economy
- Automation of standard tasks frees human to apply higher skills and expertise
- Loss of jobs compensated by new jobs created (ex internet)

Individuals

- Impacting everyone’s job and their private life
- Need basic understanding of the “Importance of data”
- Need to understand benefits and dangers of AI

JOBS ARE CHANGING, WHY ?

Paper → Desktop tools

Desktop tools → Code

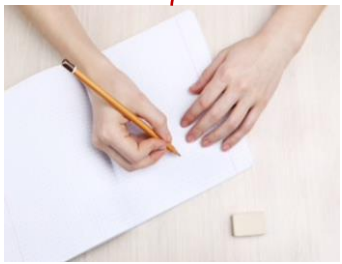

Productivity

IT JOBS AS WELL

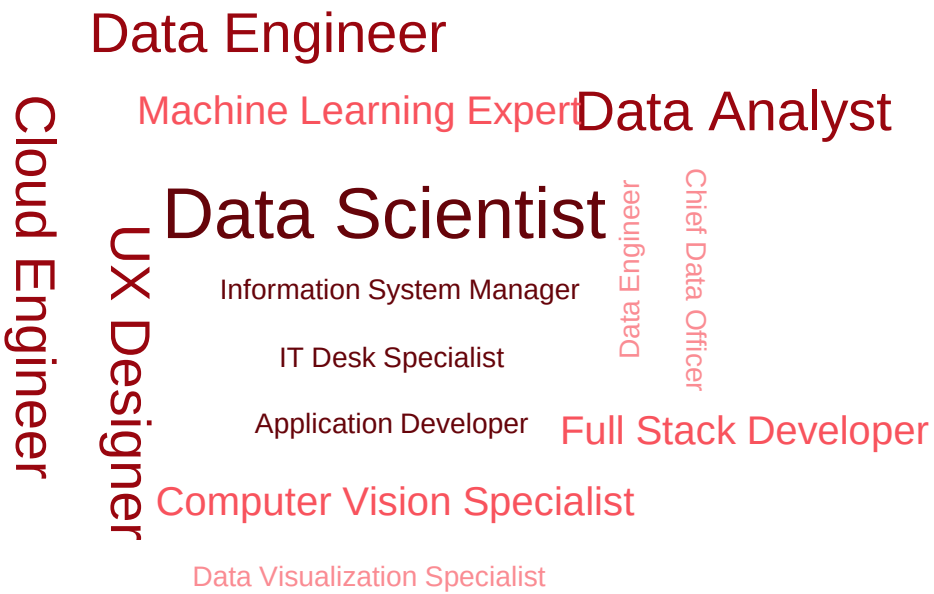

MAIN CHALLENGE IS AROUND ORGANIZATION AND CULTURE

IT teams

"Did you try to switch it off and on ?"

Associates

« Who modified the Excel ? »

Management

« I've read on LinkedIN that we should do blockchains ! »

Meanwhile. …

External providers

TRANSFORMATION HAPPENS AT ALL LEVEL

IT teams

« Yes Boss, product wil be ready in 90 days »

Associates

« I am going to share the visualisation of my analysis »

Innovation teams

« Let's prepare a technical test to asses this start-up's technology »

Management

HOW TO DO IT ?

- R&D
- Partnerships
- Acquisitions

- Open space
- Collaborative tools
- Flexibility

- Training
- Attraction
- Culture

QUALITY TRAINING ENABLES COST EFFICIENT DIGITAL SKILLS ACQUISITION

	Trained Employee	New Employee	Contracted Consultant
Company's culture	✓	✗	✗
Industry Knowledge	✓	✗	✓
In-house capabilities	✓	✓	✗
Bring fresh vision	✓	✓	✓

CHOOSE WISELY

MUST-HAVE OF DIGITAL EDUCATION

Online

Support 1-1

Projects

BENEFITS OF CORPORATE DIGITAL LEARNING

Talent development / attraction

Organic Innovation

Internal Communication

Four options for successful implementation

Call for volunteers: Commit to a limited number of subscriptions and communicate to staff so they can submit an application for the training of their choice

New hires: Train new employees to ensure incoming workforce is properly prepared to tackle digital challenges

Ambassador program: Identify one staff per team who can train and then advise and motivate future potential learners

Department-centric: Train a large group of people within the same department to build synergies; replicate it to others when successful

Q&A