AI & Law – a technical perspective

WIRTSCHAFTS
UNIVERSITÄT
WIEN VIENNA
UNIVERSITY OF
ECONOMICS
AND BUSINESS

Prof. Dr. Axel Polleres

OPENDATA
data.wu.ac.at

Institute for Information Business. data.wu.ac.at

EFMD EQUIS ACCREDITED AACSB ACCREDITED ASSOCIATION OF MBAS AMBA ACCREDITED

What is (the state of the art in) AI?

- Types of (Artificial) Intelligence
- (Deep) Neural Networks

- Some examples:
  - Question Answering
  - Surveilance, face recognition (mood, sexual orientation)
  - Diagnosis based on images
  - Image Creation
  - Text Creation
  - Deep Fakes

What is Intelligence?

What is Intelligence?

- A very general mental capability that, among other things, involves the ability to **reason**, **plan**, **solve problems**, think **abstractly**, **comprehend** complex ideas, learn quickly and **learn from experience**. It is not merely book learning, a narrow academic skill, or test-taking smarts. Rather, it reflects a broader and deeper capability for comprehending our surroundings — "catching on," "**making sense**" of things, or **"figuring out" what to do** - Wall Street Journal, 1994

- Intelligence ...
  - Perception (sound, images…)
  - Abstraction (Language understanding, Vision: situation, object detection)
  - Decision Making
  - Problem solving
  - Planning
  - Creativity
  - Emotional knowledge
  - Self-awareness/Consciousness

Involves Increasingly complex tasks

What is (the state of the art in) AI?

“Types” of AI:

- **Weak AI** refers to narrow embodiments of an AI – AI as a tool
- **Strong AI** refers to a a machine with consciousness, sentience and mind
- **Artificial general intelligence (AGI)** (a machine with the ability to apply intelligence to any problem, rather than just one specific problem).

We are not yet there… and it will probably still take a while!

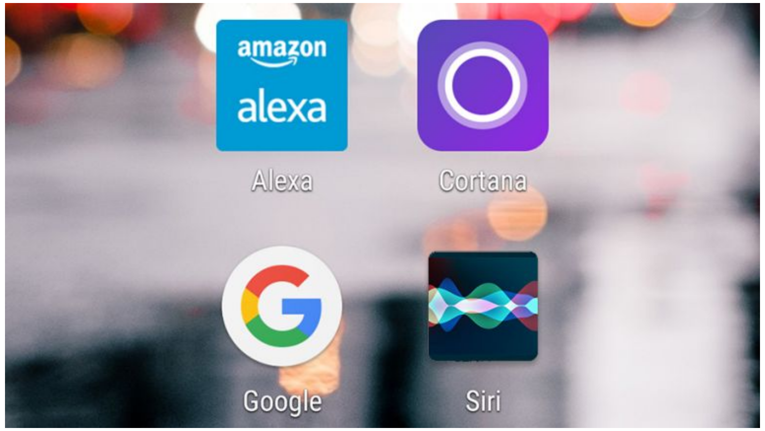

Have you ever tried having a reasonable conversation with Alexa?

What is (the state of the art in) AI?

- Distinction between “AI main strands”
  - **Model-based** vs. **Function-Based AI**
- Boost of Machine Learning/Deep Learning:
  - **Data, Hardware, Algorithms**
    - Lots of data and the ability to store/process it
    - GPUs
    - New Neural eEtwork architectures Deep Neural Networks, LSTMs …etc.
- *"What just happened is the successful employment of AI technology in some widespread applications, aided greatly by developments in related fields, and by new modes of operation that can tolerate lack of robustness or intelligence."*

A. Darwiche. Human-Level Intelligence or Animal-Like Abilities?

https://cacm.acm.org/magazines/2018/10/231373-human-level-intelligence-or-animal-like-abilities/fulltext

Model-Based AI (or Rule-based AI)

- Up to large-scale real-world applications:

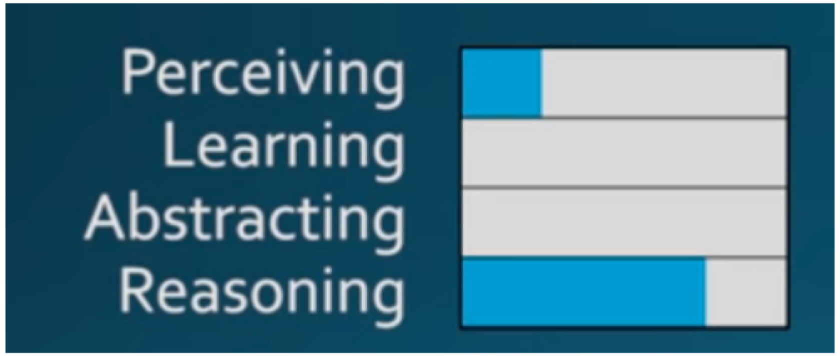

Works very well in structured, regulated domains:

- Automated Verification of chips
- Scheduling and logistics,
- Production planning in automated plants
- Playing games with a finite, discrete search space and finite rules (Chess, but not Go – why?)
- Tax declarations
- → **big advantage:** Declarative Rules and constraints are “explainable by design”.

Function-Based AI (or (Deep) Learning)

The original ideas are not so new, starting around the 1950s, around the idea of the **perceptron** (a simplified mathematical model of a neuron):

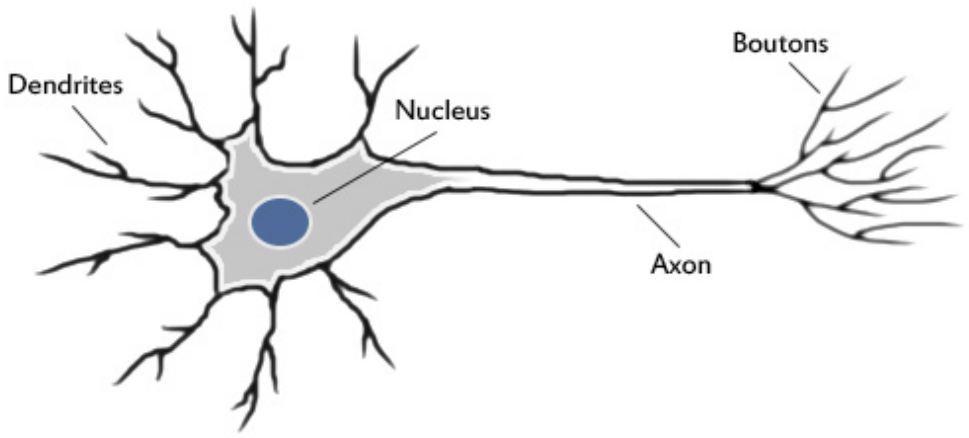

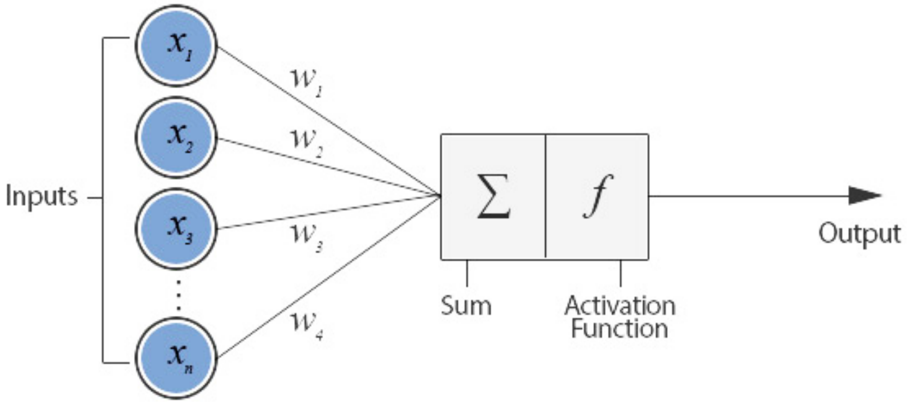

Rosenblatt, Frank (1957), The **Perceptron -** a perceiving and recognizing automaton. Report 85-460-1, Cornell Aeronautical Laboratory.

Function-Based AI (or (Deep) Learning)

- (Deep) Neural Networks Learning:

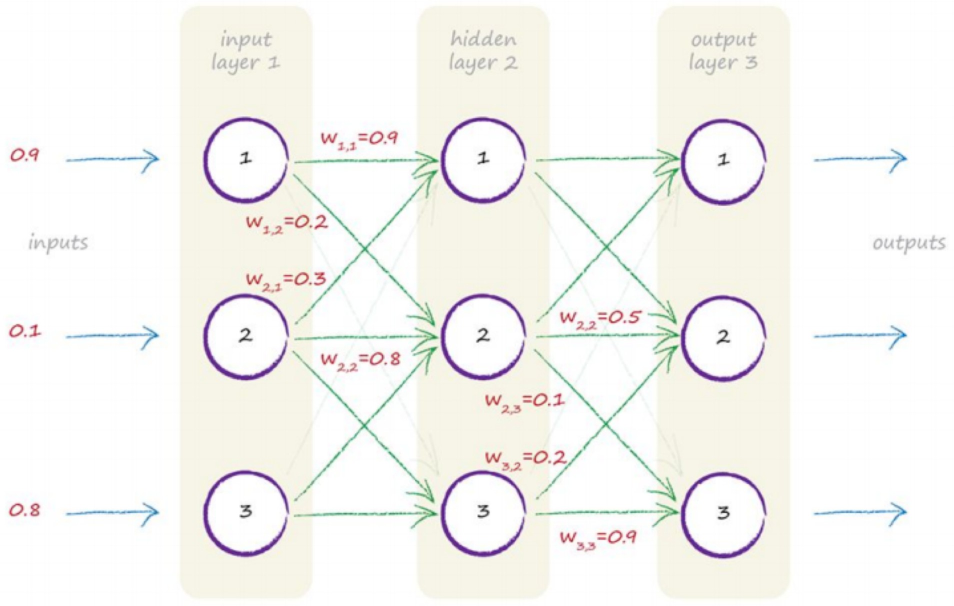

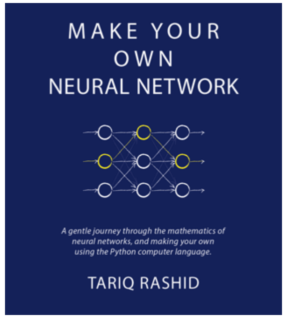

MAKE YOUR OWN NEURAL NETWORK

A gentle journey through the mathematics of neural networks, and making your own using the Python computer language.

TARIQ RASHID

Source:
http://makeyourownneuralnetwork.blogspot.com/

- Pro:
  - Given enough example data, a neural network can "learn" any function [1]
  - Works great for: predictions from large amounts of training data, can even re-create new data itself!
- Con: Can't **explain** what it does, not easy to "train" constraints, works *only* with enough training data

[1] Kurt Hornik, Maxwell B. Stinchcombe, Halbert White: Multilayer feedforward networks are universal approximators. Neural Networks 2(5): 359-366 (**1989**)

What is (the state of the art in) AI?

- Use of (function-based) AI got a lot easier (many new software tools, easy access to cloud infrastructure)

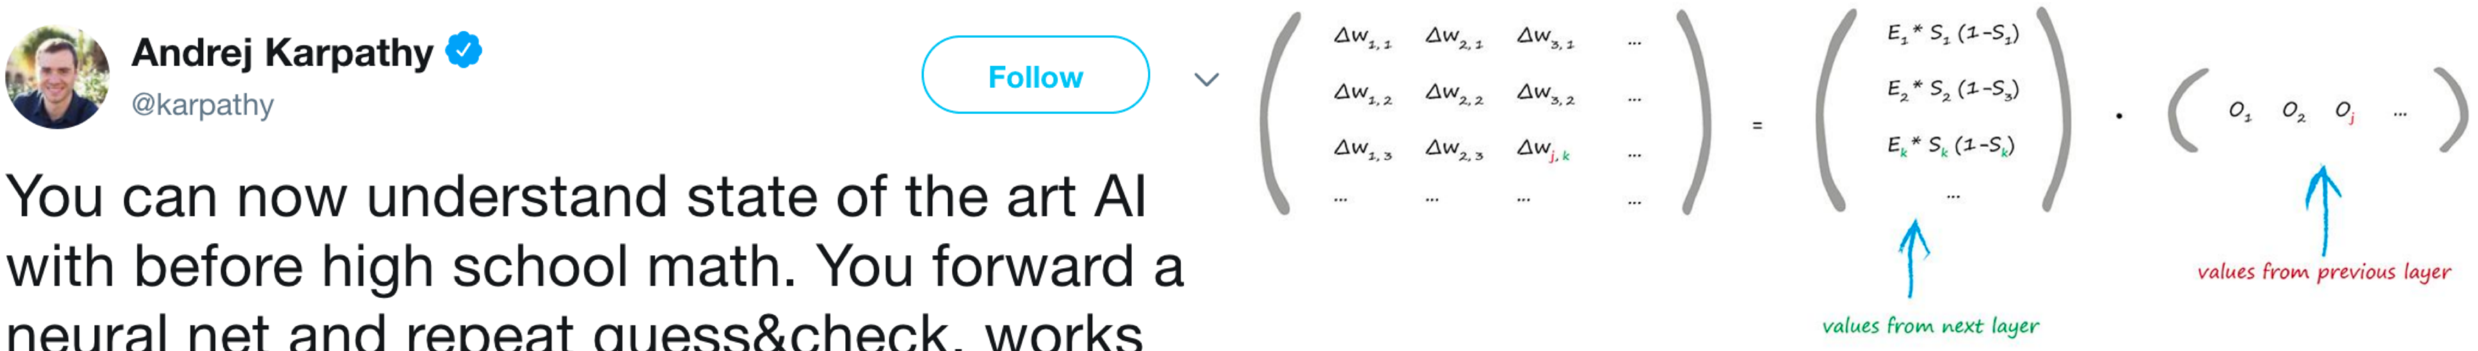

Andrej Karpathy
@karpathy

You can now understand state of the art AI with before high school math. You forward a neural net and repeat guess&check. works well enough.

$$\begin{pmatrix} \Delta w_{1,1} & \Delta w_{2,1} & \Delta w_{3,1} & \dots \\ \Delta w_{1,2} & \Delta w_{2,2} & \Delta w_{3,2} & \dots \\ \Delta w_{1,3} & \Delta w_{2,3} & \Delta w_{j,k} & \dots \\ \dots & \dots & \dots & \dots \end{pmatrix} = \begin{pmatrix} E_1 * S_1 (1-S_1) \\ E_2 * S_2 (1-S_3) \\ E_k * S_k (1-S_k) \\ \dots \end{pmatrix} \cdot \begin{pmatrix} O_1 & O_2 & O_j & \dots \end{pmatrix}$$

- https://twitter.com/karpathy/status/841739127796125696?lang=en

Why (now) Neural Networks?
Data, Hardware, Algorithms!

- Some Milestones (non-chronological):

- Multi-Layer (feedforward) networks with one hidden layer have been proven to be able to approximate essentially **any** functions**:**

  - Kurt Hornik, Maxwell B. Stinchcombe, Halbert White: Multilayer feedforward networks are universal approximators. Neural Networks 2(5): 359-366 (**1989**)

Specialized architectures with multiple layers - inspired by nature:

- **Convolutional Neural Networks (CNN) (nowadays often used for image recognition)**
  - Hubel, David H., and Torsten N. Wiesel. "Receptive fields of single neurones in the cat's striate cortex." *The Journal of physiology* 148.3 (**1959**): 574-591.
  - Fukushima, Kunihiko. "Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position." *Biological cybernetics* 36.4 (1980): 193-202.

- **Long-Short-Term Memory Networks (LSTM) (nowadays often used for sequential data (speech, text)**
  - Hochreiter, Sepp, and Jürgen Schmidhuber. "Long short-term memory." *Neural computation* 9.8 (**1997**): 1735-1780.

- **Generative Adversarial Networks (GAN):**
  - Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. "Generative adversarial nets." In *Advances in neural information processing systems*, pp. 2672-2680. 2014.

Function-Based AI: Example

Thanks to Prof. Georg Dorffner (MedUni Wien)

Image recognition in diagnosis. Typical example use of Convolutional NNs

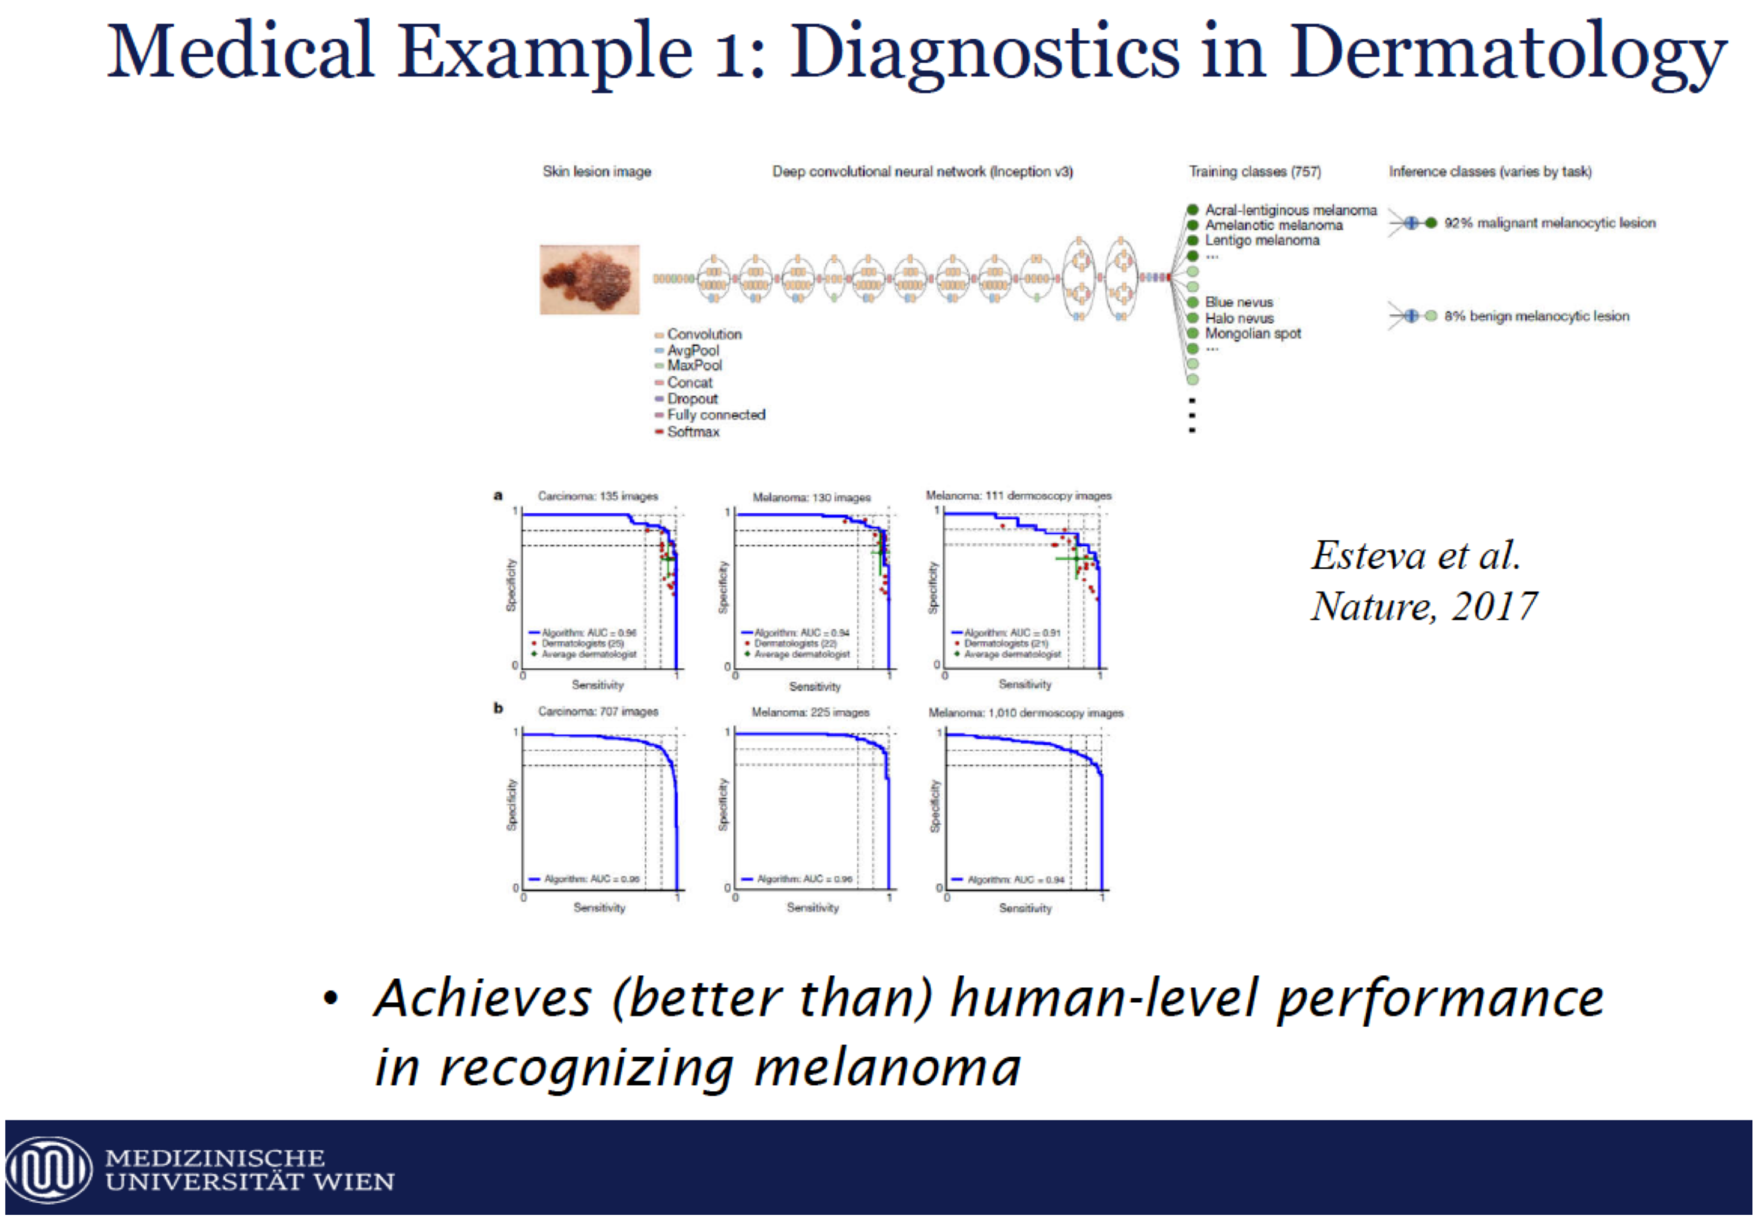

Medical Example 1: Diagnostics in Dermatology

Esteva et al. Nature, 2017

- *Achieves (better than) human-level performance in recognizing melanoma*

MEDIZINISCHE
UNIVERSITÄT WIEN

Function-Based AI other examples:

- Many useful applications:

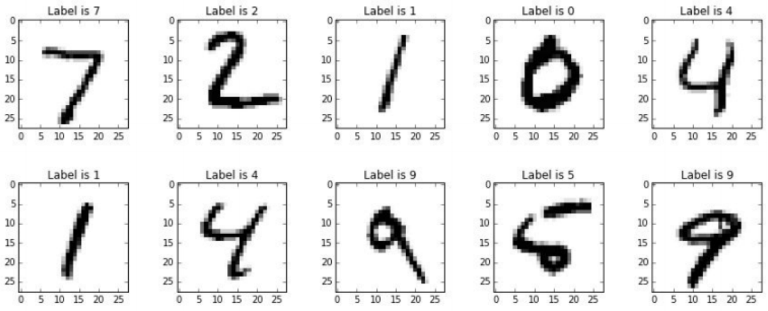

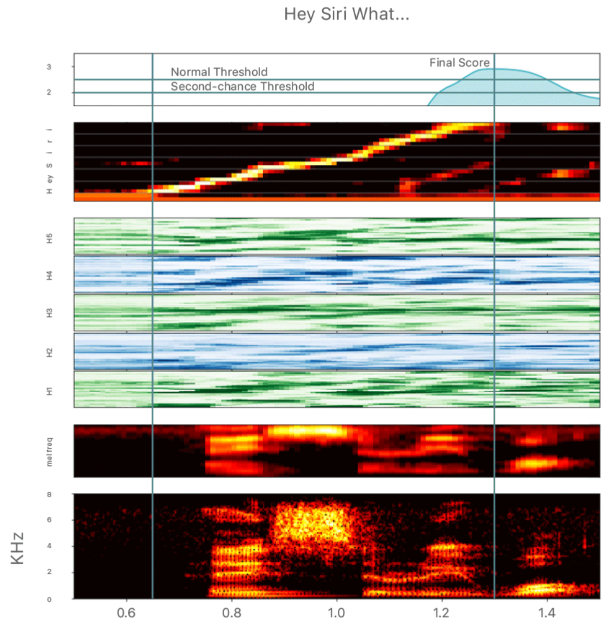

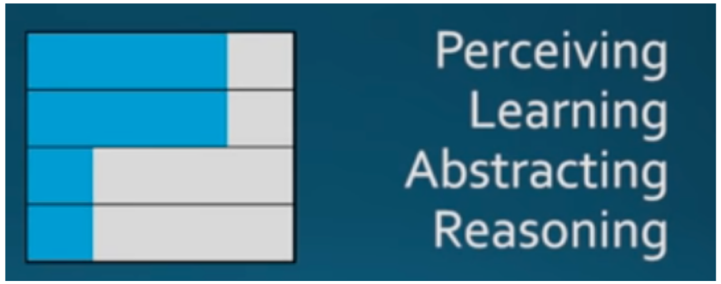

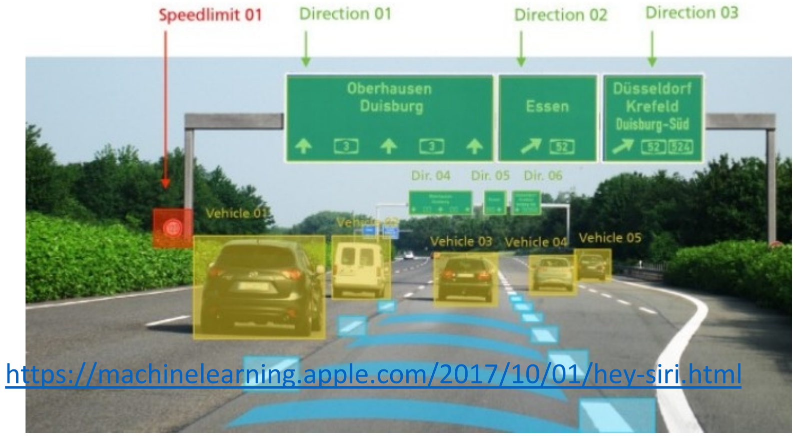

https://machinelearning.apple.com/2017/10/01/hey-siri.html

Function-Based AI : “Creating” Data… Example:

- Examples from https://twitter.com/drbeef for usage of a Generative Adversarial Networks (GAN): ***Image Generation***

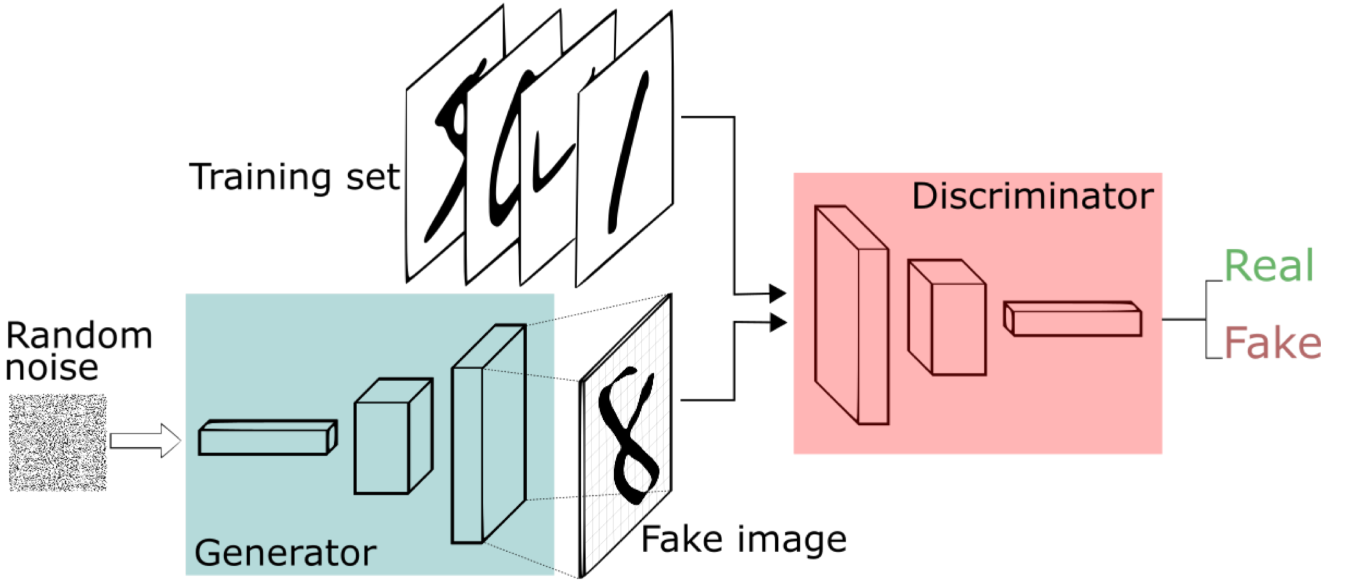

Example - Landscape paintings

Example – Nude paintings

Pro:

Not only learns patterns, but also re-creates (similar) new data itself!
Works quite well already for texts!

Con:

not easy to “train” constraints
Again, needs a lot of training data

Function-Based AI other examples:

- Also problematic applications (fake manipulation of video content or generation of voice characteristics: https://en.wikipedia.org/wiki/Deepfake

So, what is current AI?

- Learning **patterns** from massive data.
- Making **predictions** from learnt **patterns**.
- **Generating similar data** according to these patterns.

- What about making **decisions** from learnt patterns?
    - Machines make less error than humans vs.
    - Machines make different errors than humans vs.
    - Machines make errors because of humans (programmers or data)

Where’s the IP?

Traditional view:

- Programming = Data + Algorithms

vs.

- AI = **Data** + (Deep) Neural Networks

  + Data Preparation + Parameter Tuning

Related Questions:
Where we need better regulations…

- Data ownership:
    - Data Ownership and IP … what about synthesized data?
    - Data Licenses

- Personal Data Protection and Bias:
    - GDPR is a good start, but (how) can we automate it?
        - Challenges: Consent, Trusted Computing, Sticky Policies

- Explainability and Responsibility
    - Knowledge Graphs & Hybrid AI

Model-based AI for Data ownership & Licensing: **Machine-readable policies**

Digital Rights Management

A Data Licensing Perspective…

https://pixabay.com/en/data-computer-businessmen-monitor-2899902/

dalicc

DATA LICENSES CLEARANCE CENTER

https://www.dalicc.net/

DALICC is funded by the Austrian Federal Ministry of Transport, Innovation and Technology (BMVIT) under the program "ICT of the Future" between November 2016 - October 2018. More information https://iktderzukunft.at/en/

fh/// st.pölten

WU WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

SEMANTIC WEB COMPANY

STI · INNSBRU

Höhne In der Maur & Partner Rechtsanwälte

Model-based AI for Personal Data Protection: **Machine-readable policies**

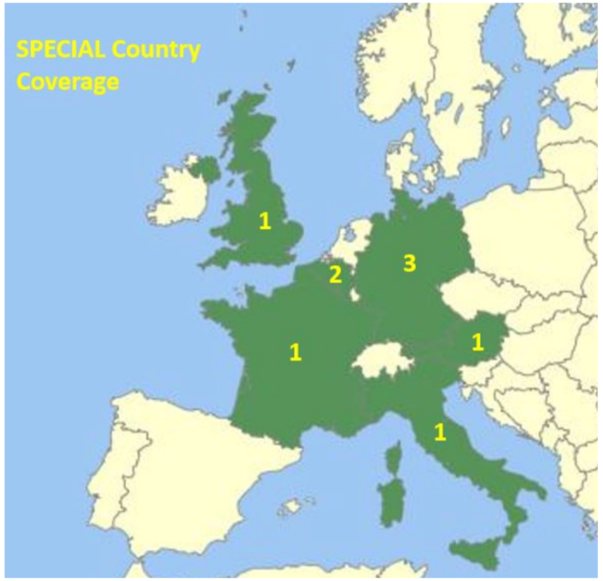

Research and innovation Action
9 partners from 6 countries
January 2017 to December 2019
3,991,389 €

SPECIAL

Scalable Policy-awarE Linked Data arChitecture for prIvacy, trAnsparency and compLiance

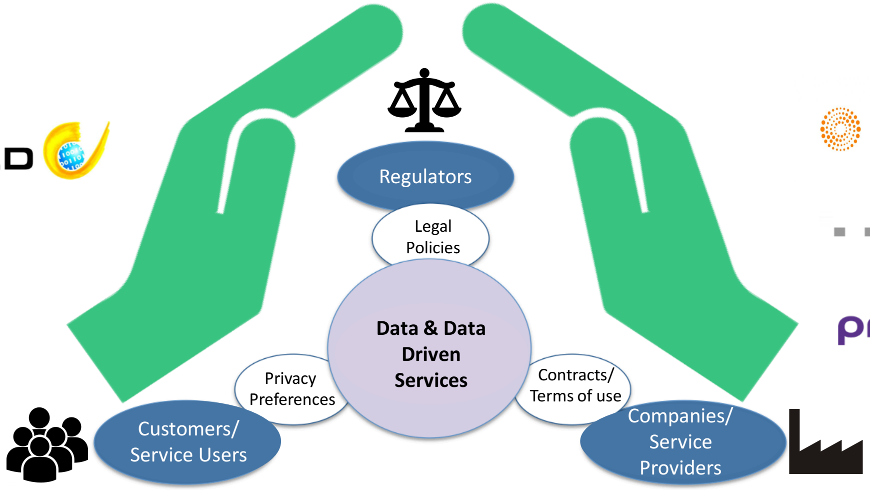

Horizon 2020
European Union funding
for Research & Innovation

This project receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 731601

https://www.specialprivacy.eu/

Function-Based AI for Personal Data Protection:

Can GANs be used as a way to anonymize data?

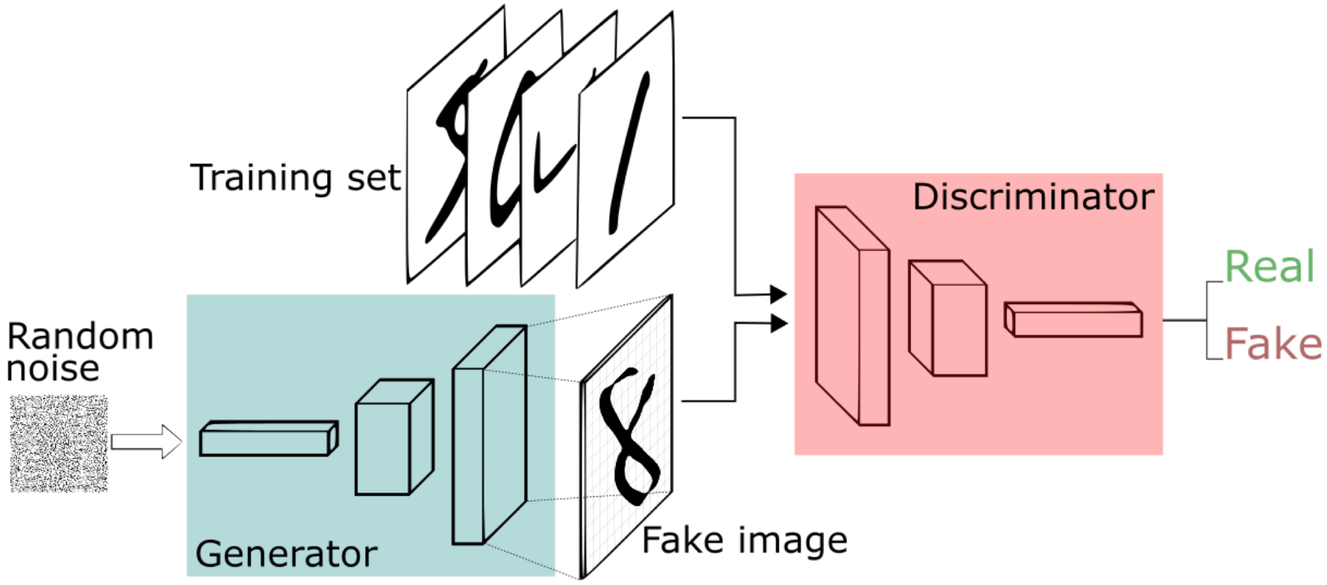

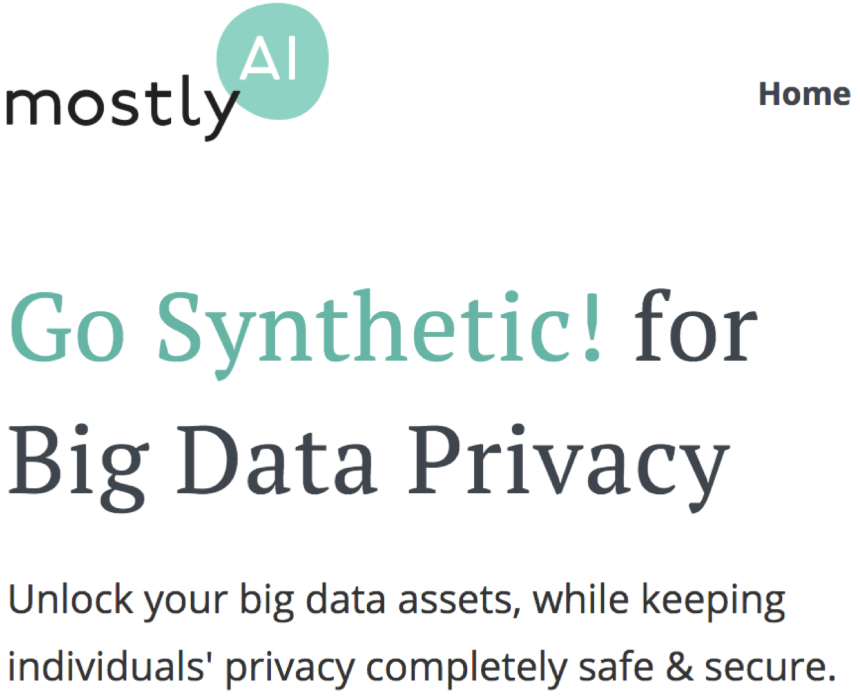

Go Synthetic! for Big Data Privacy

Unlock your big data assets, while keeping individuals' privacy completely safe & secure.

Last, but not least: Bias and Explainability

https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

This article relates to a computer program that generates a score predicting the likelihood of criminals committing a future crime.

In 2014, then U.S. Attorney General Eric Holder warned that the risk scores might be injecting bias into the courts.

ProPublica did, as part of a larger examination of the powerful, largely hidden effect of algorithms in American life.

Still largely unresolved – Subject of research:

Solution? ***Hybrid AI*** – i.e., combining Model-based and funciton-based AI

Image references:

- http://makeyourownneuralnetwork.blogspot.com/
- https://twitter.com/drbeef
- https://www.youtube.com/watch?v=cQ54GDm1eL0

- John Launchbury: A DARPA Perspective on Artificial Intelligence - **https://www.youtube.com/watch?v=-O01G3tSYpU**

- A. Darwiche. Human-Level Intelligence or Animal-Like Abilities? Communications of the ACM, October 2018, Vol. 61 No. 10, Pages 56-67 https://cacm.acm.org/magazines/2018/10/231373-human-level-intelligence-or-animal-like-abilities/fulltext

Acknowledgements:

- Thanks to Sabrina Kirrane, Georg Dorffner for some of the slides.