UPPSALA
UNIVERSITET

Flexible learning systems driven by data

Thomas Schön
Uppsala University

NordicPOP Annual Meeting 2021
Online
March 16, 2021.

"Machine learning gives computers the ability to **learn without being explicitly programmed** *for the task at hand."*

What we do in the team — who we are

Solid basic research and the **most interesting applied research** (often with companies).

1. Create **probabilistic models** of dynamical systems and their surroundings.
2. Develop methods to **learn** models from data.

The models can then be used by machines (or humans) to **understand** or **make decisions** about what will happen next.




AI – ML – DL

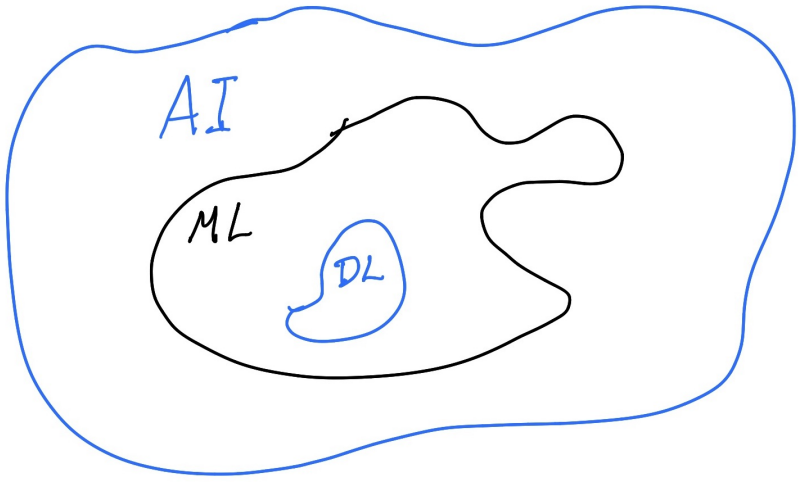

Artificial Intelligence (AI) refers to systems that **display intelligent behaviour by analysing their environment and taking actions**—with some degree of autonomy—**to achieve specific goals**.

Machine learning (ML) gives computers the ability to **learn without being explicitly programmed** for the task at hand.

Machine learning is one of today's most rapidly growing technical fields, representing a creative combination of **mathematics** and **programming**.

Deep learning (DL)

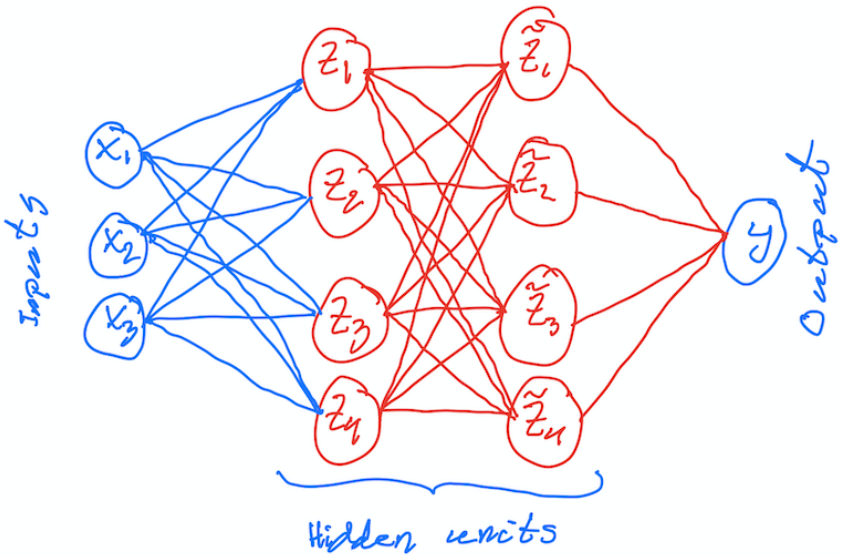

Underlying idea: when representing a function, a deep, hierarchical model can be **exponentially more efficient** than a shallow model.

With enough training data the machine can be trained to make very good predictions from previously unseen data.

Outline

Aim: Make use of concrete examples to give a feeling for the use of AI/ML as a **tool in research**.

Outline:

1. Introduction
2. **Flexible models**
3. Ex. Medicine
4. Ex. Generative ML – Democracy, freedom and justice?
5. Ex. Autonomous systems
6. AI4Research – Uppsala University AI project

Key lesson from contemporary machine learning

It is sometimes (often...) easier to solve a problem by starting from examples of input-output data than trying to manually program it.

Key lesson from modern machine learning:

Flexible models often gives the best performance.

How can we build and work with these flexible models?

1. Models that use a large number of parameters compared with the data set. (**parametric**, ex. deep learning)

   LeCun, Y., Bengio, Y., and Hinton, G. **Deep learning**, *Nature*, Vol 521, 436–444, 2015.

2. Models that use more parameters as we get access to more data. (**non-parametric**, ex. Gaussian process)

   Ghahramani, Z. **Probabilistic machine learning and artificial intelligence**. *Nature* 521:452-459, 2015.

Ex. (medicine) ECG diagnosis

Aim: Predict abnormalities based on a short-duration 12-lead electrocardiogram (ECG) recording.

Current situation: The automated diagnosis that is currently available is not good enough.

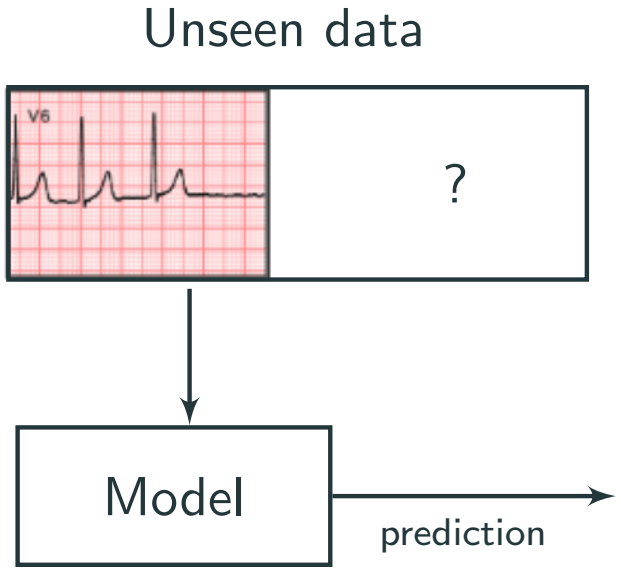

Background: Joint work with cardiologists and ML engineers from Brazil with an urgent need for automated analysis due to the **vast distances** between the patient and a cardiologist with full expertise in ECG diagnosis.

The existing telehealth network provides the data (more than 2 300 000 ECGs), implying some clinical relevance.

Ex. (medicine) ECG diagnosis

Supervised learning of a model from labelled data.

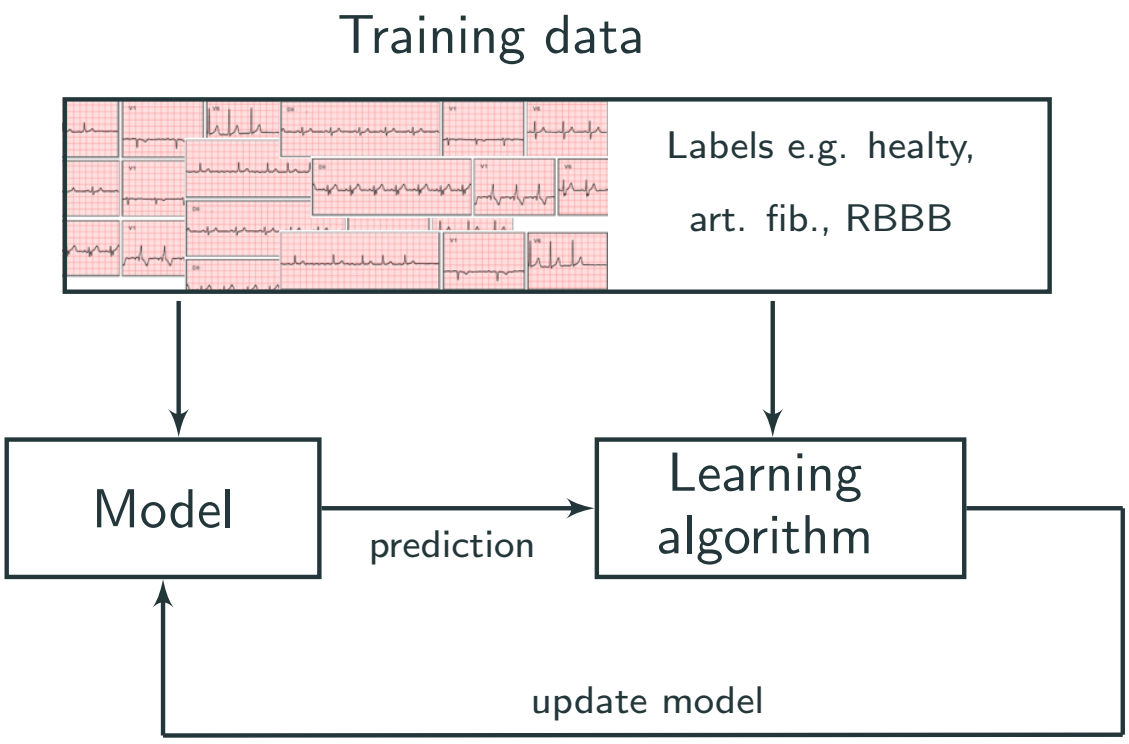

Ex. (medicine) ECG diagnosis

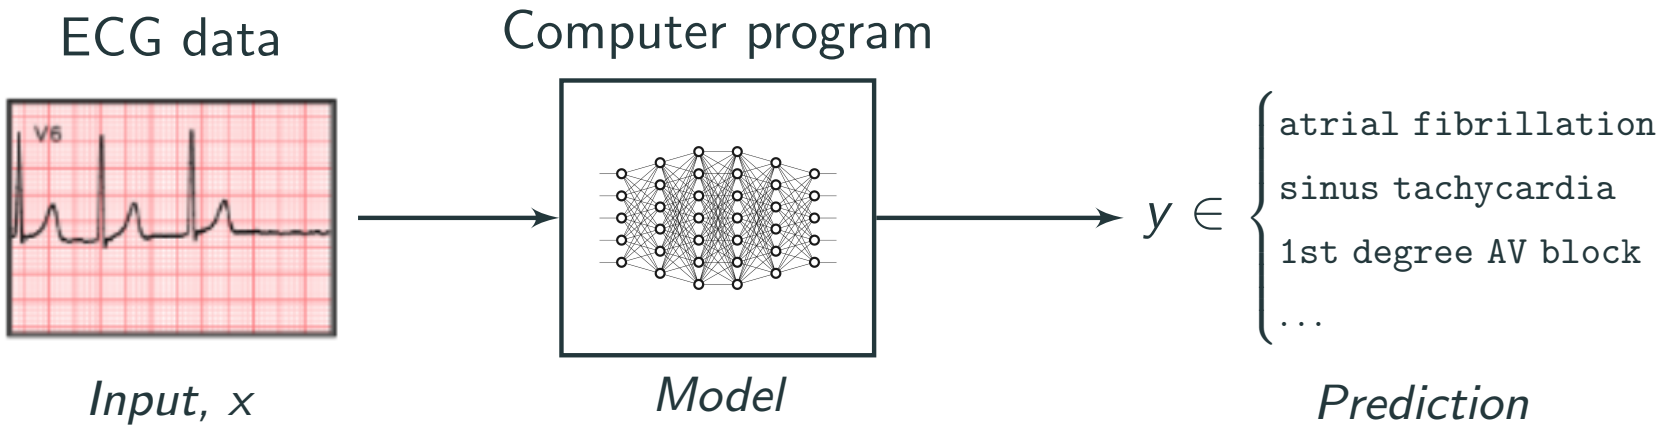

We are now reaching human level (medical doctor) performance on certain specific tasks.

Key difference to classical approach: The model is **not** derived based on our ability to mathematically explain what we see in an ECG.

Instead, a generic model is **automatically learned** based on data.

Antonio H. Ribeiro, Manoel H. Ribeiro, Gabriela M.M. Paixao, Derick M. Oliveira, Paulo R. Gomes, Jessica A. Canazart, Milton P. S. Ferreira, Carl R. Andersson, Peter W. Macfarlane, Wagner Meira Jr., TS, Antonio Luiz P. Ribeiro. **Automatic diagnosis of the 12-lead ECG using a deep neural network**. *Nature Communications*, 11(1760), 2020.

Adding creativity to the machines

Inspiration from Katja de Vries (Faculty of Law, Uppsala University).

Using **generative models** (VAE and GAN) we have added a very basic form of imagination to the machines.

Vahdat, A, and Kautz, J. **NVAE: A deep hierarchical variational autoencoder**. *34th Conference on Neural Information Processing Systems (NeurIPS)*, Vancouver, Canada, December, 2020.

Democracy, freedom and justice?

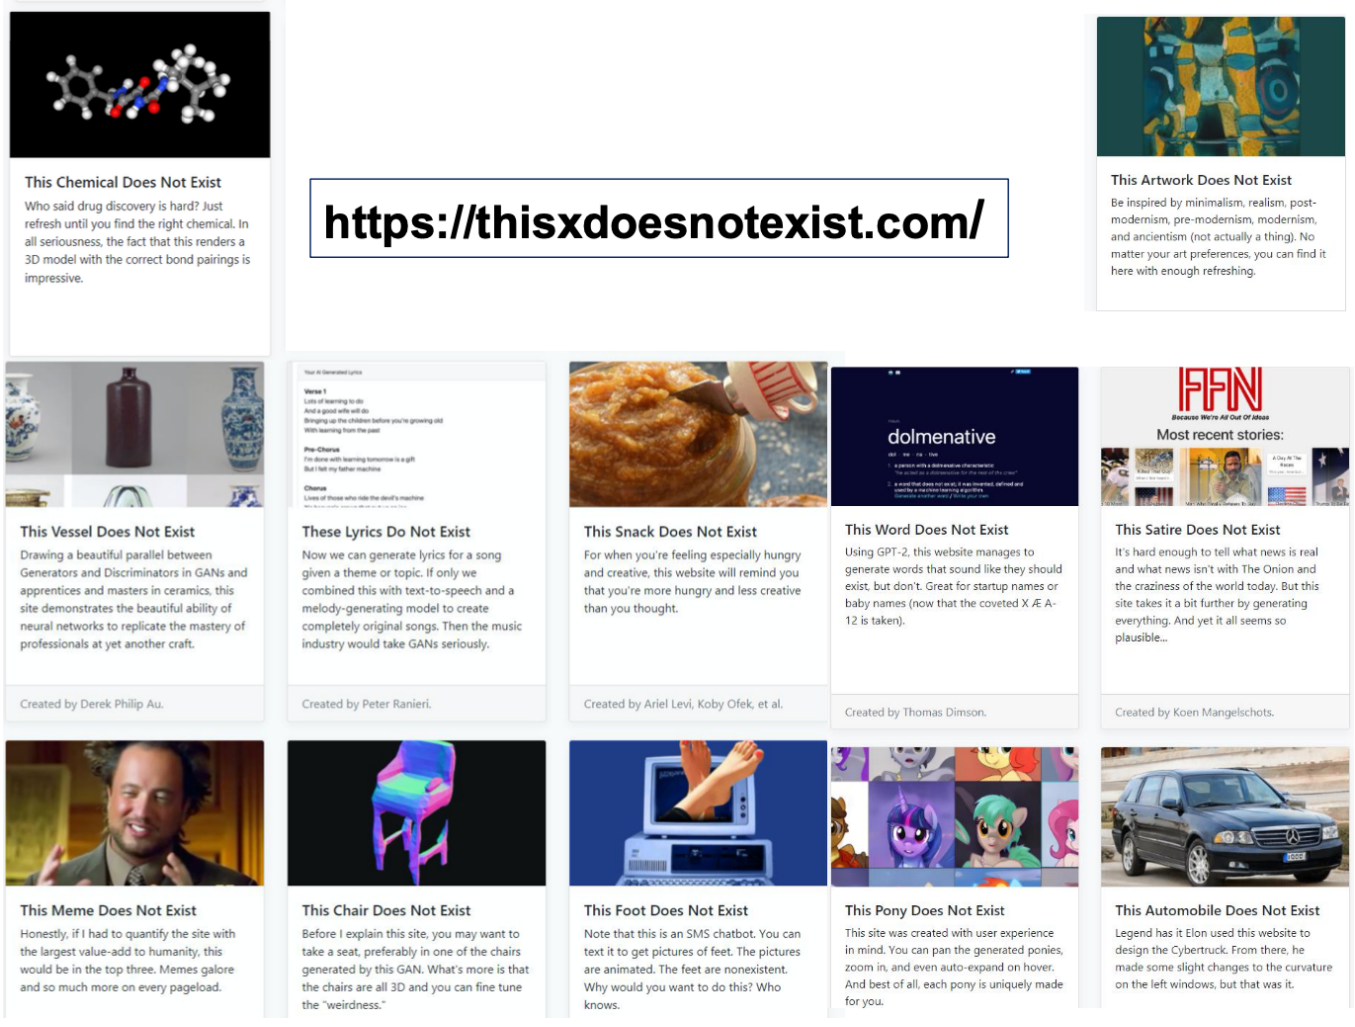

How can creative AI affect democracy, freedom and justice?

The problem we are interested in – regression using DNNs

Supervised regression: learn to predict a continuous output (target) value $y^\star \in \mathcal{Y} = \mathbb{R}^K$ from a corresponding input $x^\star \in \mathcal{X}$, given a training set $\mathcal{D}$ of i.i.d. input-output data

$$\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N, \qquad (x_n, y_n) \sim p(x, y).$$

Deep neural network (DNN): a function $f_\theta : \mathcal{X} \to \mathcal{Y}$, parameterized by $\theta \in \mathbb{R}^P$, that maps an input $x \in \mathcal{X}$ to an output $f_\theta(x) \in \mathcal{Y}$.

Generally applicable, but we have (so far) mainly worked with examples from computer vision and robotics.

Input space $\mathcal{X}$: Space of images or point clouds.

Output space $\mathcal{Y} = \mathbb{R}^K$: $\mathcal{Y} = \mathbb{R}^2$ for image-coordinate regression, $\mathcal{Y} = \mathbb{R}_+$ for age estimation, $\mathcal{Y} = \mathbb{R}^4$ for 2D bounding-box regression.

Ex. Visual tracking

Task: Estimate a bounding box of a target object in every frame of a video. The target object is defined by a given box in the first video frame.

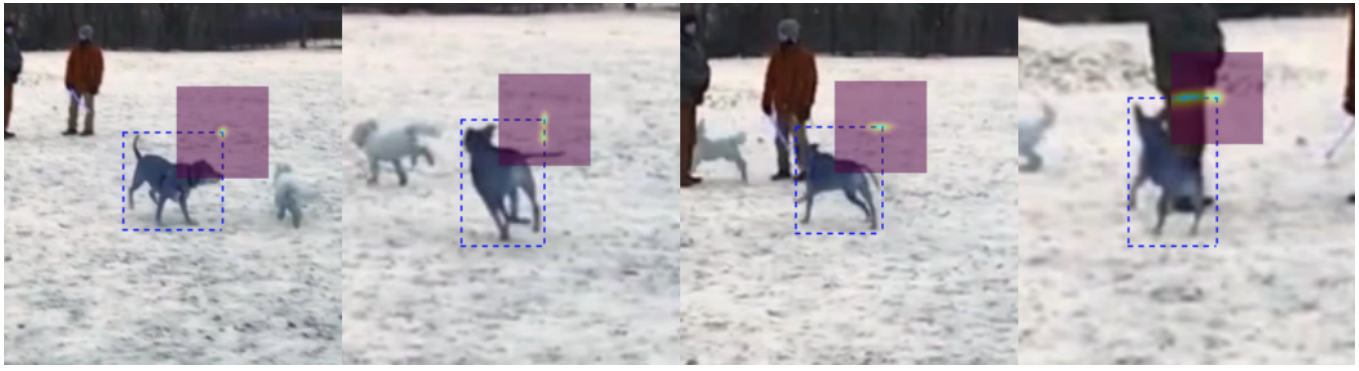

Fredrik K. Gustafsson, Martin Danelljan, Goutam Bhat, TS. **Energy-based models for deep probabilistic regression**. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Online, August, 2020.

Code available from `github.com/fregu856/ebms_regression`

Ex. 3D object detection from laser data

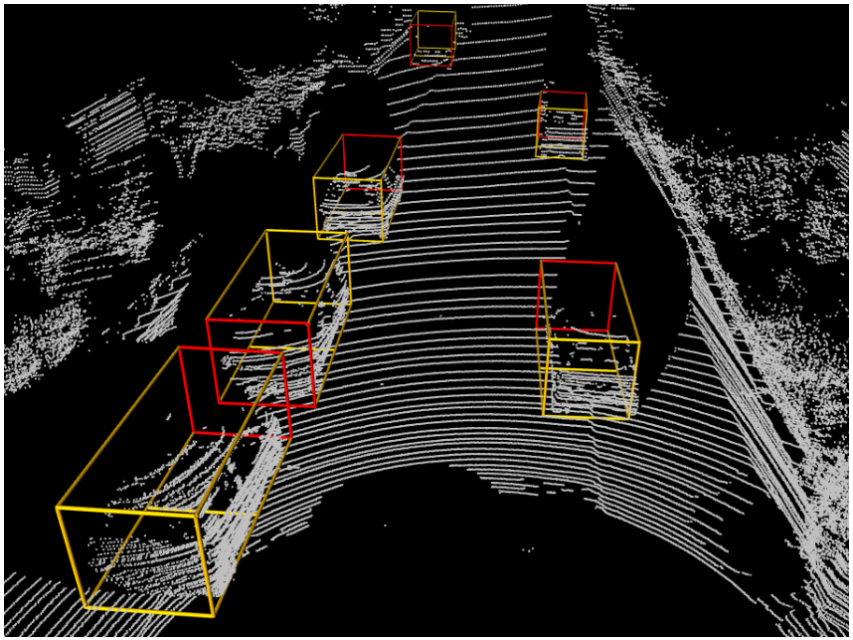

Task: Detect objects from sensor data (here laser), estimate their size and position in the 3D world.

Key perception task for self-driving vehicles and autonomous robots.

New probabilistic regression model based on deep neural networks.

The **combination** of **probabilistic models** and **deep neural networks** is very exciting and promising.

Fredrik K. Gustafsson, Martin Danelljan, and TS. **Accurate 3D object detection using energy-based models**. *arXiv:2012.04634*, October, 2020.

Two comments on the key lesson

Flexible models often gives the best performance.

---

1. **Blending** prior knowledge and data

Create flexible model building blocks containing the basic knowledge we have about the phenomenon we are studying.

2. Be careful as flexible models can be **deceptive**!

AI4Research – university-wide AI project

At Uppsala University we **develop and make use of AI for research**.

A **time-limited five year effort** consisting of an **antidisciplinary entity** from the entire university.

Located in newly refurbished premises at our main library Carolina Rediviva.

Key mechanism: **Internal AI sabbatical periods**

- Funded 50% by the entity and the rest by the department where the fellow remains employed/external grants.
- Duration: around 12 months.
- The fellows brings along 1-2 of their PhD students/post-docs.

AI4Research – team this year

Read about the research from the project website

www.uu.se/forskning/ai4research

Co-located with the Anders Wiklöf Institute for Heart Research.

Conclusions

While AI techniques are used more and more in industry, scientists are—for good reasons—becoming aware of the potential in using AI in fundamental research.

The best predictive performance is currently obtained from **highly flexible learning systems**.

Remember to talk to people who work on **different problems** with **different tools**!! (Visit other fields!)