

Deep Learning

Lecture 0: Introduction


Prof. Gilles Louppe
g.louppe@uliege.be

Today

- Course outline
- Introduction to deep learning
- Fundamentals of machine learning

Outline

- Lecture 1: Fundamentals of machine learning
- Lecture 2: Multi-layer perceptron
- Lecture 3: Automatic differentiation
- Lecture 4: Training neural networks
- Lecture 5: Convolutional neural networks
- Lecture 6: Computer vision
- Lecture 7: Attention and transformer networks
- Lecture 8: GPT
- Lecture 9: Graph neural networks
- Lecture 10: Uncertainty
- Lecture 11: Auto-encoders and variational auto-encoders
- Lecture 12: Score-based diffusion models






My mission

By the end of this course, you will have a strong and comprehensive understanding of deep learning.

You will learn how to design deep neural networks for various advanced probabilistic inference tasks and how to train them.

The models covered in this course have broad applications in artificial intelligence, engineering, and science.

Why learning?



What do you see?



Sheepdog or mop?



Chihuahua or muffin?


The (human) brain is so good at interpreting visual information that the gap between raw data and its semantic interpretation is difficult to assess intuitively.



This is a mushroom.



This is a mushroom.



This is a mushroom.

```

array([[0.03921569, 0.03529412, 0.02352941, 1.          ],
       [0.2509804 , 0.1882353 , 0.20392157, 1.          ],
       [0.4117647 , 0.34117648, 0.37254903, 1.          ],
       ...,
       [0.20392157, 0.23529412, 0.17254902, 1.          ],
       [0.16470589, 0.18039216, 0.12156863, 1.          ],
       [0.18039216, 0.18039216, 0.14117648, 1.          ]]),

[[0.1254902 , 0.11372549, 0.09411765, 1.          ],
 [0.2901961 , 0.2509804 , 0.24705882, 1.          ],
 [0.21176471, 0.2          , 0.20392157, 1.          ],
 ...,
 [0.1764706 , 0.24705882, 0.12156863, 1.          ],
 [0.10980392, 0.15686275, 0.07843138, 1.          ],
 [0.16470589, 0.20784314, 0.11764706, 1.          ]]),

[[0.14117648, 0.12941177, 0.10980392, 1.          ],
 [0.21176471, 0.1882353 , 0.16862746, 1.          ],
 [0.14117648, 0.13725491, 0.12941177, 1.          ],
 ...,
 [0.10980392, 0.15686275, 0.08627451, 1.          ],
 [0.0627451 , 0.08235294, 0.05098039, 1.          ],
 [0.14117648, 0.2          , 0.09803922, 1.          ]]),


...,

```


This is a mushroom.

Writing a computer program that sees?









To extract semantic information, we need models with high complexity **that cannot be manually designed.**

However, we can write a program that learns the task of extracting semantic information.





The **machine learning approach** consists in:

- defining a parametric model
- optimizing its parameters, by "making it work" on the training data.

The deep learning revolution

Deep learning **scales up** the statistical and machine learning approaches by

- using larger models known as neural networks,
- training on larger datasets,
- using more compute resources.







Il y a eu
des hauts et des bas.

▶ 0:00 / 4:41



Specialized neural networks can be trained achieve super-human performance on many complex tasks that were previously thought to be out of reach for machines.



A person riding a motorcycle on a dirt road.



Two dogs play in the grass.



A group of young people playing a game of frisbee.



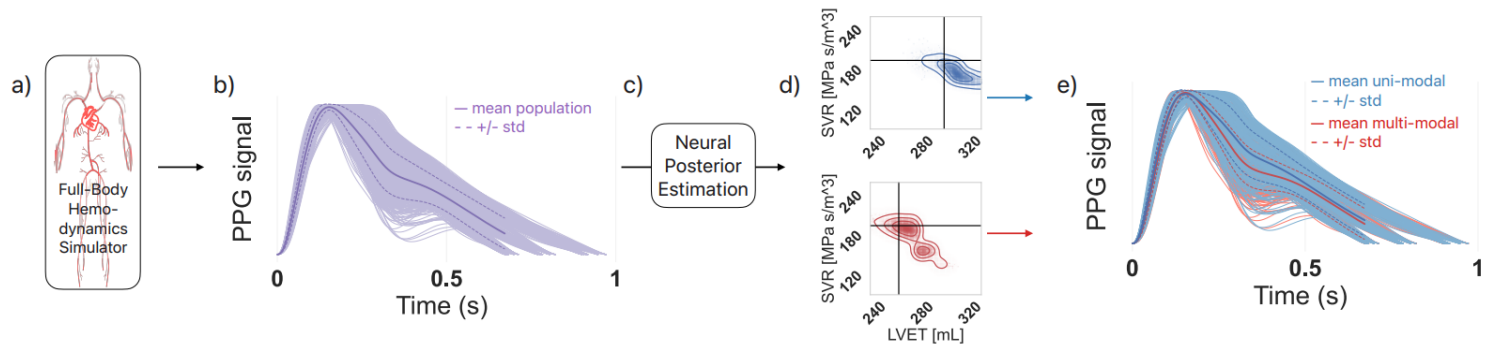



Two hockey players are fighting over the puck.



I: Jane went to the hallway.
I: Mary walked to the bathroom.
I: Sandra went to the garden.
I: Daniel went back to the garden.
I: Sandra took the milk there.
Q: Where is the milk?
A: garden

(Top) Scene understanding, pose estimation, geometric reasoning.
(Bottom) Planning, Image captioning, Question answering.

Neural networks form **primitives** that can be transferred to many domains.



(Top) Analysis of histological slides, denoising of MRI images, nevus detection.

(Bottom) Whole-body hemodynamics reconstruction from PPG signals.

Segment Anything Model - A Promptable Segmentation System #Shorts



AI at Meta

Abonnieren



Object detection, pose estimation, segmentation (Meta AI, 2023)



Sense, Solve, and Go: The Magic of the Wa...



Later bekij...



Delen



Building autonomous cars (Waymo, 2022)



Powering the Future of Clean Energy | I AM ...



Later bekij...



Delen



Powering the future of clean energy (NVIDIA, 2023)



Camels, Code & Lab Coats: How AI Is Adva...



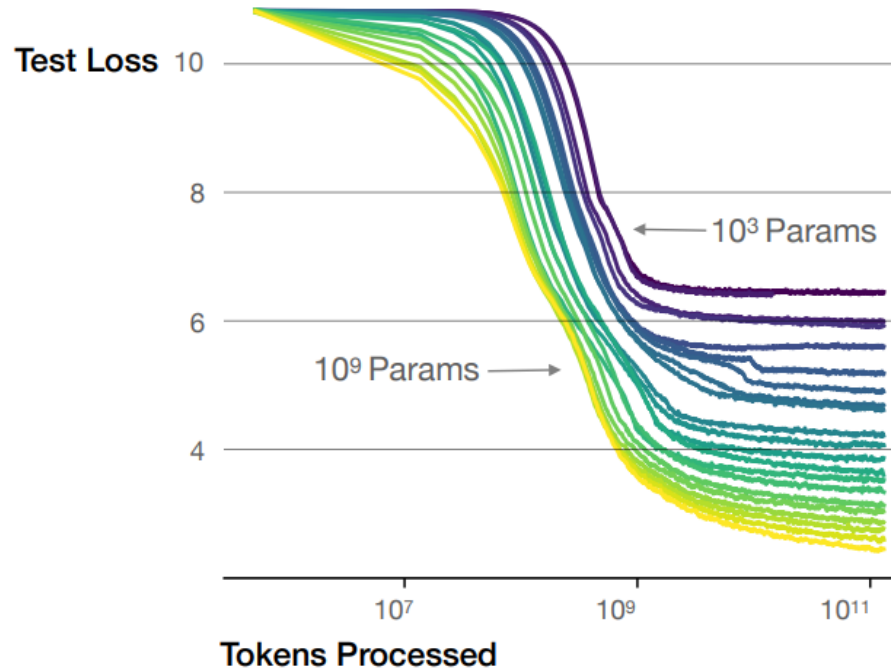
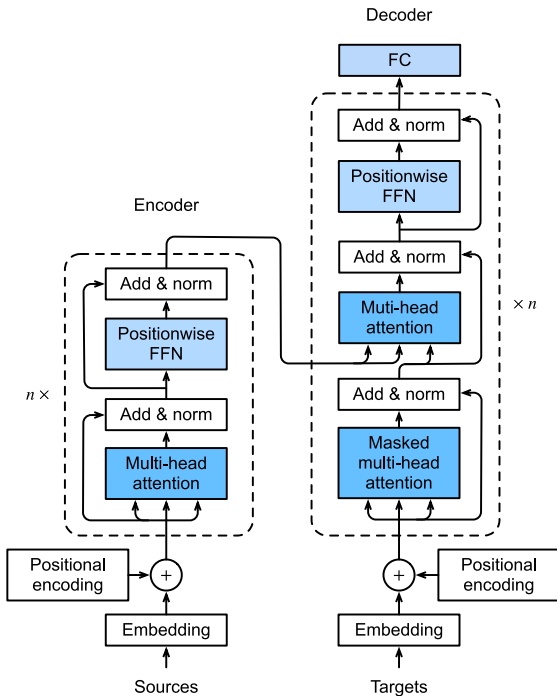
Later bekij...



Delen



How AI is advancing medicine (Google, 2023)



From large to enormous

Scaling up further to gigantic models, datasets, and compute resources keeps pushing the boundaries of what is possible, **with no sign of slowing down.**

ChatGPT PLUS



Help me pick

a gift for my dad who loves fishing

Brainstorm edge cases

for a function with birthdate as input, horoscope as ou...

Make up a story

about Sharky, a tooth-brushing shark superhero

Create a personal webpage for me

after asking me three questions




0:00 / 2:38

Send a message



Can you write me a report analyzing this chest X-ray?



Findings:

- Devices: None.
- Lungs: No pneumothorax. No substantial pleural effusion. Lungs appear clear.
- Cardiomeastinal: Normal heart size. Mediastinal contours within normal limits.
- Other: No acute skeletal abnormality.

Impression:

No active disease seen in chest.

Enter a question here

Med-PaLM 2 (Google) is a large language model tuned for the medical domain. It reaches 85%+ accuracy on US Medical Licensing Examination-style questions.

Deep learning can also **solve problems that no one could solve before.**

AlphaFold: From a sequence of amino acids to a 3D structure



nature

[Explore content](#) ▾ [About the journal](#) ▾ [Publish with us](#) ▾

[nature](#) > [articles](#) > [article](#)

Article | [Open access](#) | [Published: 15 July 2021](#)

Highly accurate protein structure prediction with AlphaFold

[John Jumper](#) , [Richard Evans](#), [Alexander Pritzel](#), [Tim Green](#), [Michael Figurnov](#), [Olaf Ronneberger](#), [Kathryn Tunyasuvunakool](#), [Russ Bates](#), [Augustin Židek](#), [Anna Potapenko](#), [Alex Bridgland](#), [Clemens Meyer](#), [Simon A. A. Kohl](#), [Andrew J. Ballard](#), [Andrew Cowie](#), [Bernardino Romera-Paredes](#), [Stanislav Nikolov](#), [Rishub Jain](#), [Jonas Adler](#), [Trevor Back](#), [Stig Petersen](#), [David Reiman](#), [Ellen Clancy](#), [Michal Zielinski](#), ... [Demis Hassabis](#) 


[+ Show authors](#)

[Nature](#) **596**, 583–589 (2021) | [Cite this article](#)

1.42m Accesses | **12k** Citations | **3493** Altmetric | [Metrics](#)

Abstract

Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function. Through an enormous experimental effort^{1,2,3,4}, the structures of around 100,000 unique proteins have been determined⁵, but this represents a small fraction of the billions of known protein sequences^{6,7}. Structural coverage is bottlenecked by the months to years of painstaking effort required to determine a single protein structure. Accurate computational approaches are needed to address this gap and to enable large-scale structural bioinformatics. Predicting the three-dimensional structure that a protein will adopt based solely on its amino acid sequence—the structure prediction component of the ‘protein folding problem’⁸—has been an important open research problem for more than 50 years⁹. Despite recent progress^{10,11,12,13,14}, existing methods fall far short of atomic accuracy, especially when no homologous structure is available. Here we provide the





AlphaFold: The making of a scientific breakt...



Later bekij...



Delen



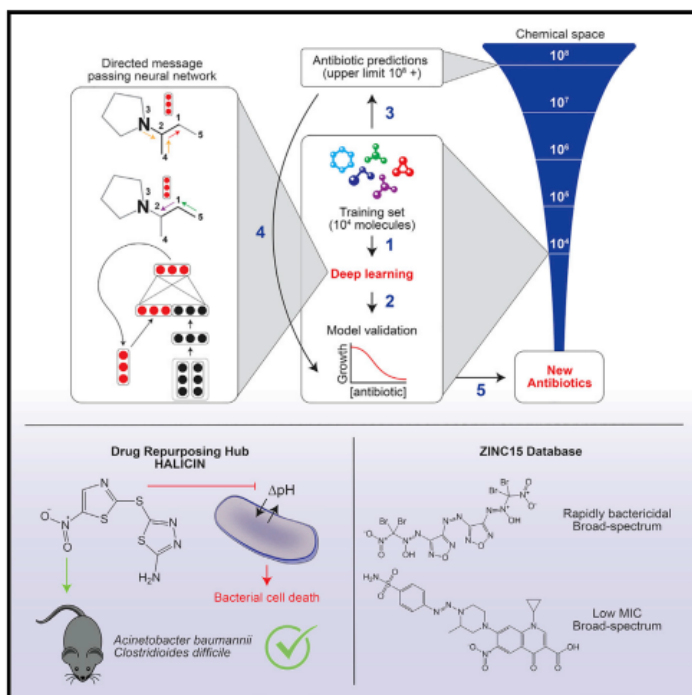
AI for Science (Deepmind, AlphaFold, 2020)

Drug discovery with graph neural networks

Cell

A Deep Learning Approach to Antibiotic Discovery

Graphical Abstract



Authors

Jonathan M. Stokes, Kevin Yang,
Kyle Swanson, ..., Tommi S. Jaakkola,
Regina Barzilay, James J. Collins

Correspondence

regina@csail.mit.edu (R.B.),
jimjc@mit.edu (J.J.C.)

In Brief


A trained deep neural network predicts antibiotic activity in molecules that are structurally different from known antibiotics, among which Halicin exhibits efficacy against broad-spectrum bacterial infections in mice.

Predicting cardiovascular risks from retinal images

Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning

Ryan Poplin^{1,4}, Avinash V. Varadarajan^{1,4}, Katy Blumer¹, Yun Liu¹, Michael V. McConnell^{2,3}, Greg S. Corrado¹, Lily Peng^{1,4*} and Dale R. Webster^{1,4}

Traditionally, medical discoveries are made by observing associations, making hypotheses from them and then designing and running experiments to test the hypotheses. However, with medical images, observing and quantifying associations can often be difficult because of the wide variety of features, patterns, colours, values and shapes that are present in real data. Here, we show that deep learning can extract new knowledge from retinal fundus images. Using deep-learning models trained on data from 284,335 patients and validated on two independent datasets of 12,026 and 999 patients, we predicted cardiovascular risk factors not previously thought to be present or quantifiable in retinal images, such as age (mean absolute error within 3.26 years), gender (area under the receiver operating characteristic curve (AUC) = 0.97), smoking status (AUC = 0.71), systolic blood pressure (mean absolute error within 11.23 mmHg) and major adverse cardiac events (AUC = 0.70). We also show that the trained deep-learning models used anatomical features, such as the optic disc or blood vessels, to generate each prediction.



Source Li, Abner, and Abner Li. 2018. "Alphabet's Verily Analyzing Retinas W/ Machine Learning To Predict Heart Disease". 9to5google.com/2018/02/19/alphabet-verily-eyes-heart-disease/.



For the last forty years we have programmed computers; for the next forty years we will train them.

Chris Bishop, 2020.

