








Artificial Intelligence: When algorithms meet Science



March 7, 2024



Prof. Gilles Louppe
g.louppe@uliege.be

































Deep Learning

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.



Specialized neural networks can be trained to achieve super-human performance on complex tasks.



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.



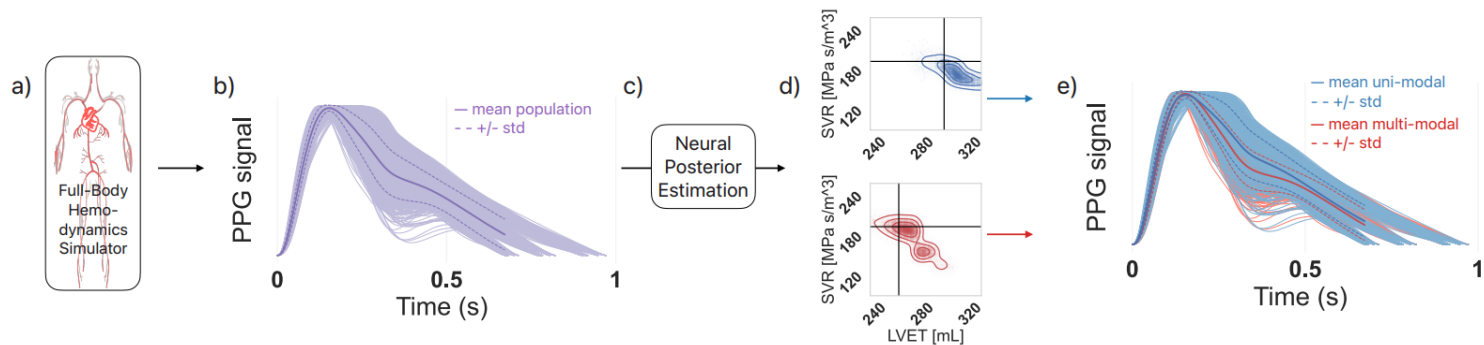

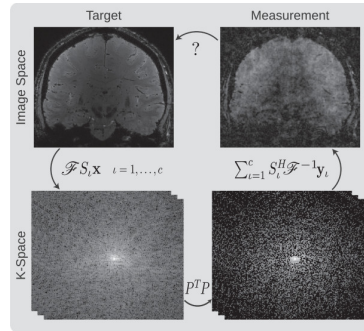
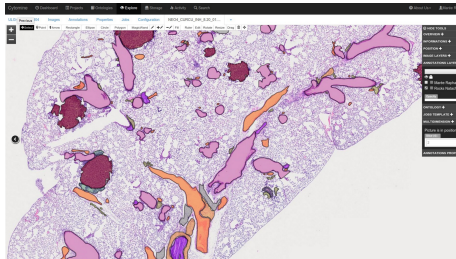
Two hockey players are fighting over the puck.



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


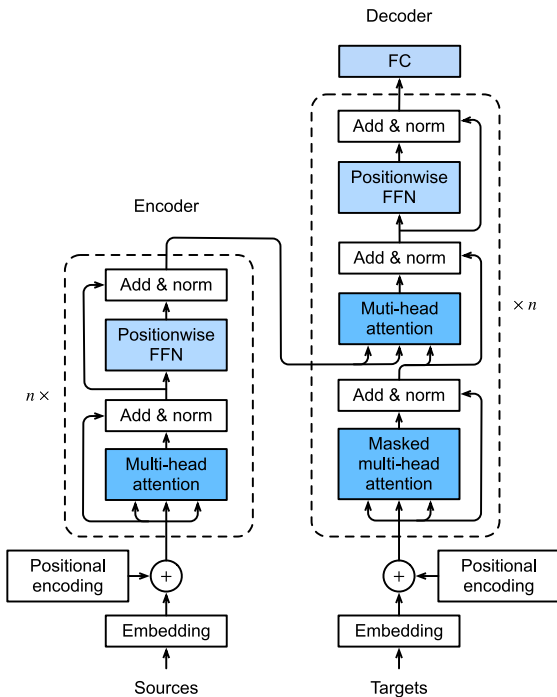
(Top) Scene understanding, pose estimation, geometric reasoning.
 (Bottom) Planning, Image captioning, Question answering.

Immediate applications of specialized neural networks in health and medicine have been around for a while.



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

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



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.**

Guess the next _____

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



(demo)



Video niet beschikbaar

Deze video is niet beschikbaar



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




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

nature
biomedical engineering


ARTICLES

<https://doi.org/10.1038/s41551-018-0195-0>

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

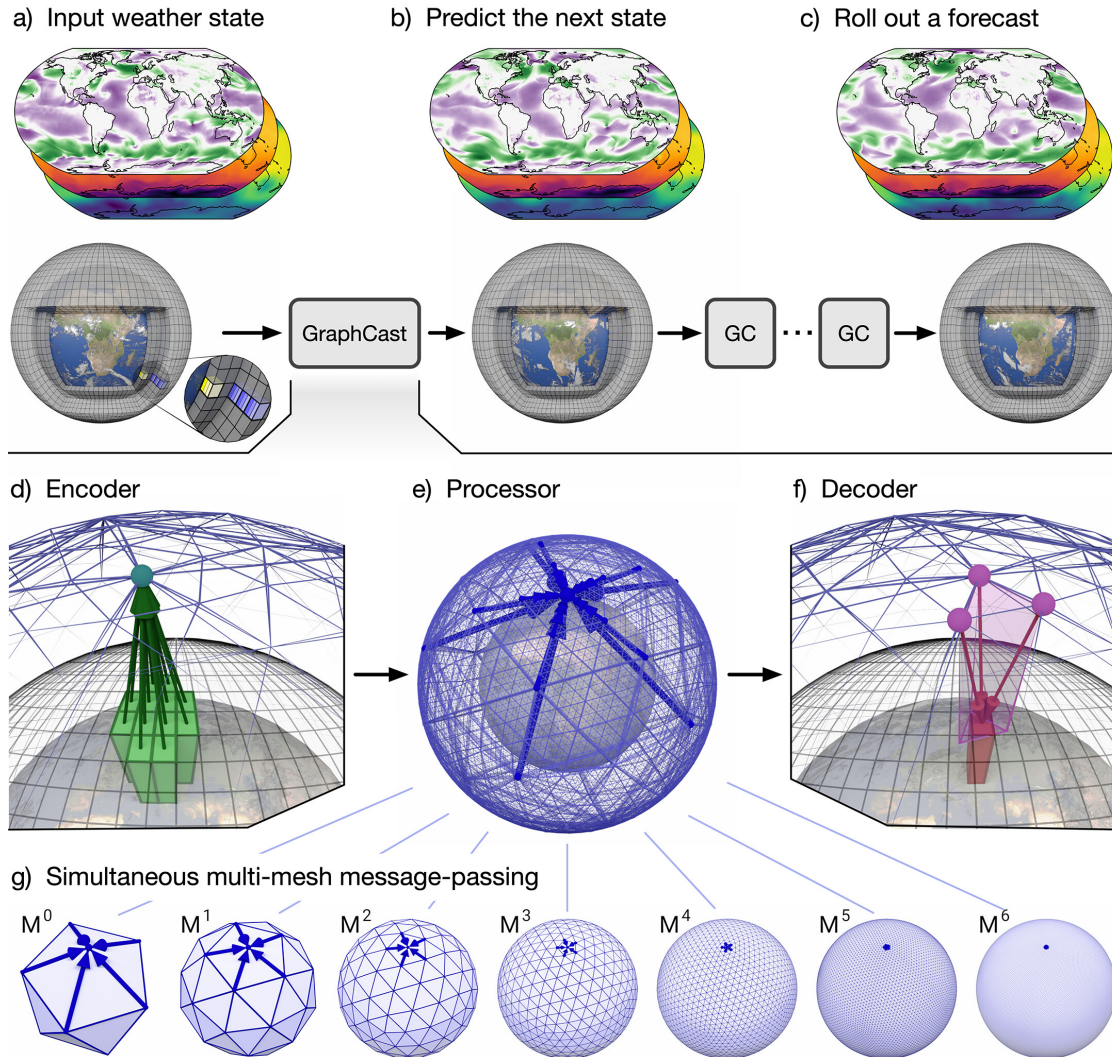
Ryan Poplin^{1,4}, Avinash V. Varadarajan^{1,4}, Katy Blumer¹, Yun Liu¹, Michael V. McConnell^{2,3}, Greg S. Corrado¹, Lily Peng^{1,4*} and Dale R. Webster^{1,4}


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". *9To3google*. <https://9to3google.com/2018/02/19/alphabet-verily-eyes-heart-disease/>.

GraphCast: fast and accurate weather forecasts







Characterization of exoplanet atmospheres



B - A



Limitations

- Neural networks are **black boxes**. Their internal workings are difficult to understand, and their predictions difficult to explain.
- Certifying the robustness of neural networks is also difficult. How to use them in safety-critical applications, like health and medicine?
- Modern neural networks require **large amounts of data and compute**. The resources required to train and run them can be prohibitive.
- Data is often **biased**. How to ensure that a neural network will not reproduce these biases?



Summary

- Deep neural networks learn to solve complex tasks by themselves, given enough data and compute.
- They are already used in health and medicine, and will be used more and more in the future with the advent of conversational and generative AI systems.
- They are not perfect, and their limitations must be taken into account.