

L'IA sans mystère

Les réseaux de neurones expliqués

Prof. Gilles Louppe
g.louppe@uliege.be




L'intelligence artificielle moderne ...

- se programme elle-même;
- résoud des problèmes complexes;
- lit, écrit, entend, parle et voit.

L'IA se programme elle-même




Comment estimer la pression artérielle d'un patient en l'absence de tensiomètre?



Les programmes classiques


On écrit un programme qui prend l'âge d'un patient et retourne une pression artérielle en suivant une suite d'instructions pré-définies.





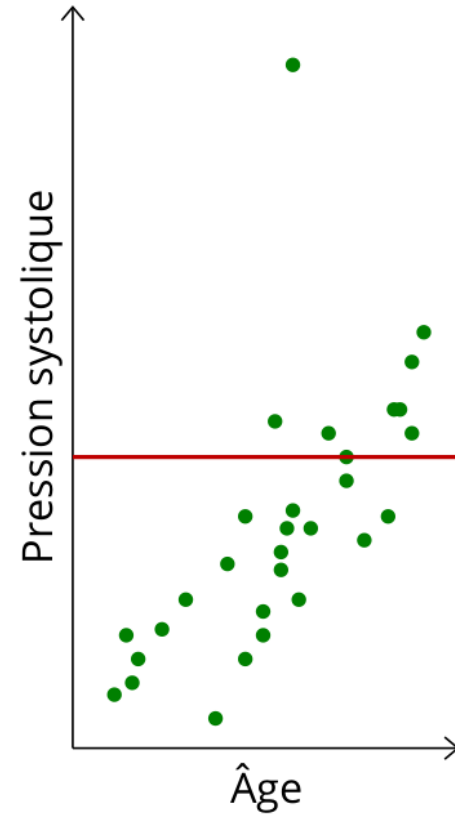

```
def blood_pressure(age):  
    if age < 20:  
        return 120  
    elif age < 40:  
        return 130  
    else:  
        return 140
```



Le machine learning



On écrit un programme qui apprend, de lui-même et sans instructions pré-établies, à prédire la pression artérielle à partir d'exemples.





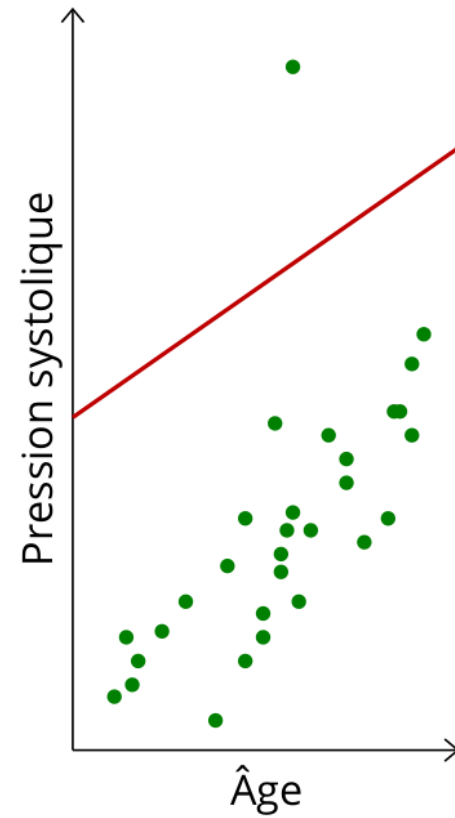



















Les réseaux de neurones


Cette stratégie rappelle la plasticité des réseaux de neurones naturels permettant au cerveau de créer, renforcer ou défaire les connexions entre les neurones pour apprendre.



Les réseaux de neurones
avant l'entraînement




Les réseaux de neurones
après 2 semaines de
stimulation cognitive




Les réseaux de neurones
après 2 mois de stimulation
cognitive

En poussant cette analogie, des neurones artificiels peuvent être définis comme des fonctions mathématiques qui prennent des entrées x , les transforment, et retournent des sorties y .

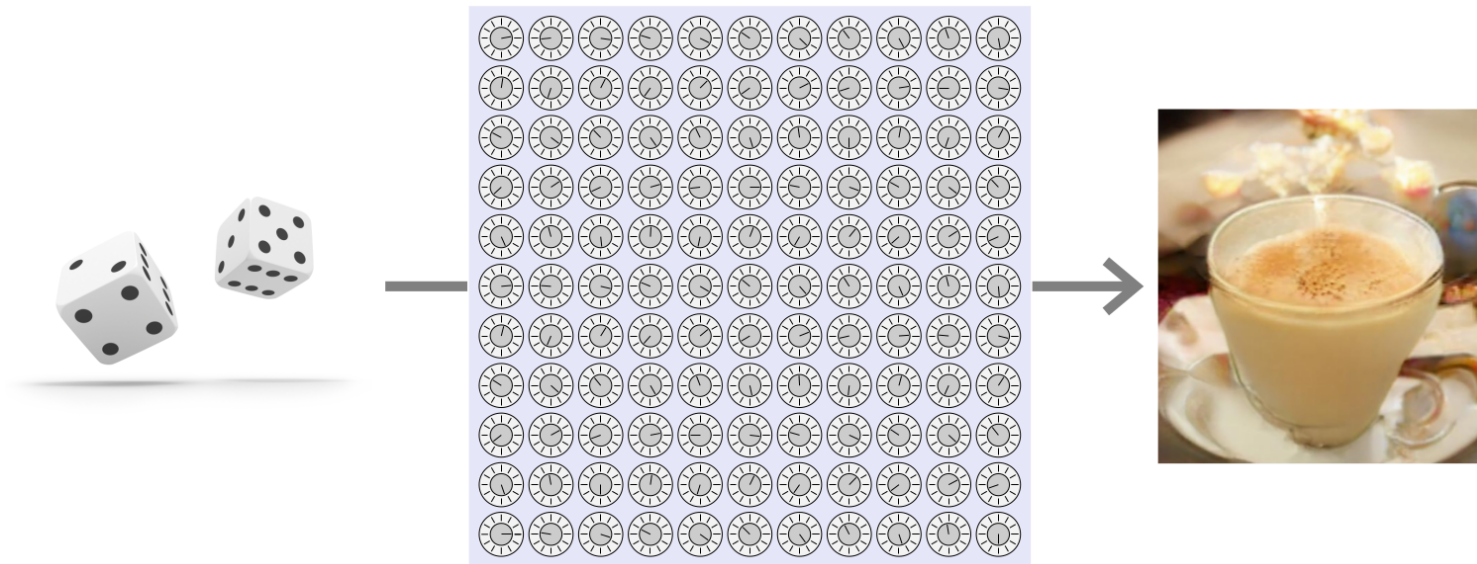
Ces neurones peuvent ensuite être organisés en couches interconnectées, pour former des **réseaux de neurones artificiels**.




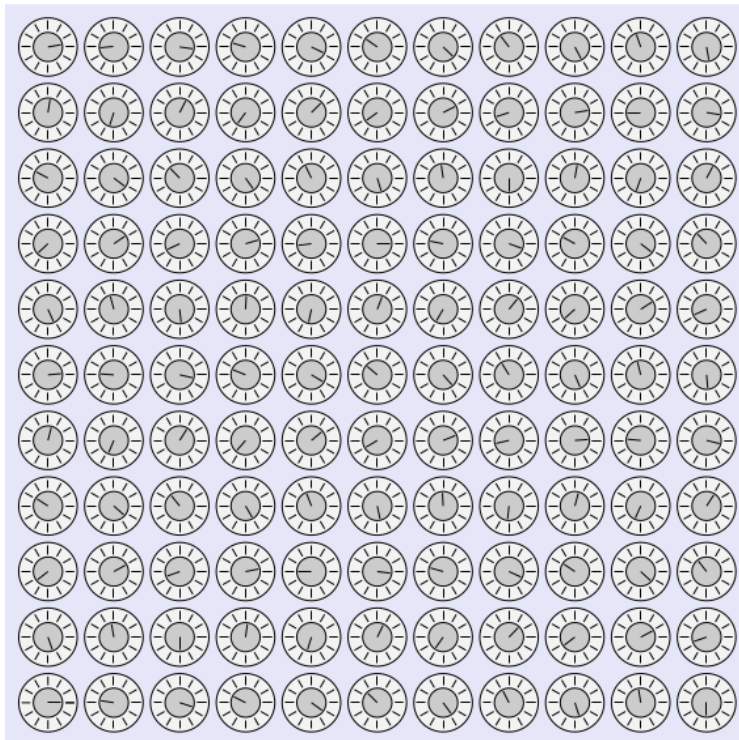
À grande échelle, ces modèles sont capables d'extraire de l'information à partir d'image, de son, ou de texte.



... ou de produire des images, du son, ou du texte.




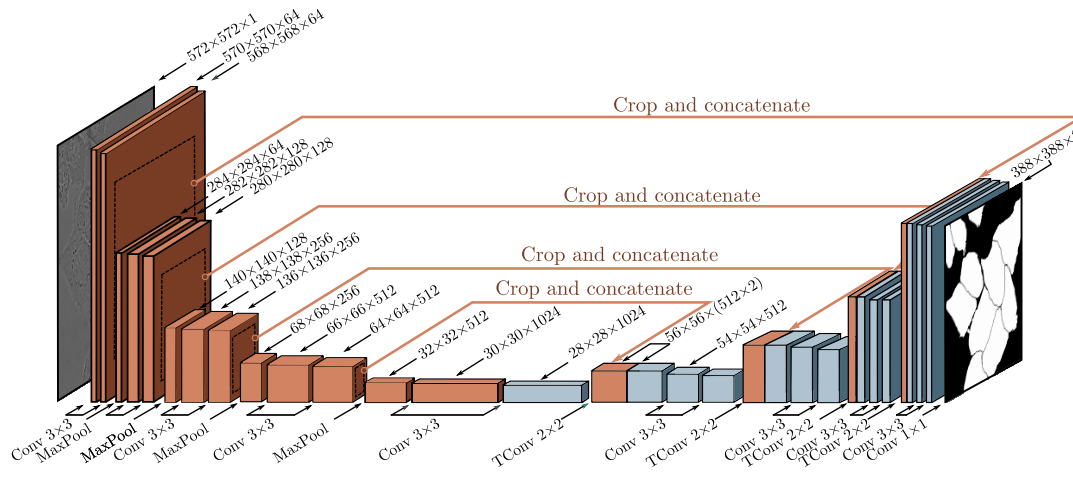
Comme précédemment, les réseaux de neurones sont **entraînés** en ajustant leurs paramètres pour minimiser l'erreur entre leurs prédictions et les vraies valeurs.



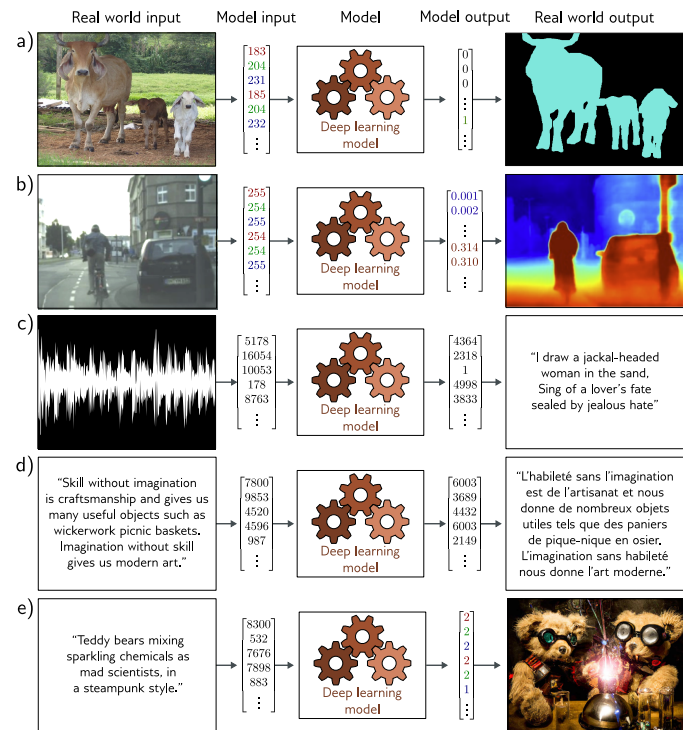
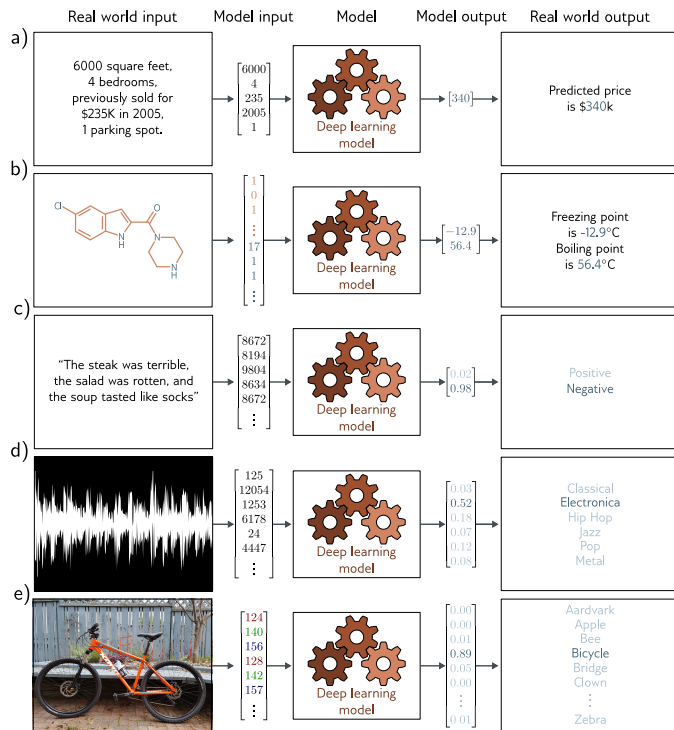
$$\theta_{t+1} = \theta_t - \gamma \nabla \mathcal{L}(\theta_t)$$

Le **deep learning** (ou **apprentissage profond**) applique la même stratégie, mais à une échelle beaucoup plus grande encore:

- en utilisant des réseaux de neurones gigantesques (jusqu'à 10^{11} paramètres),
- en s'entraînant sur des données les plus larges possibles (TBs ou PBs),
- en utilisant des ressources de calcul dédiées (des data centers entiers).




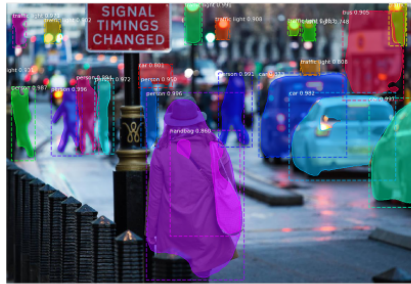


L'IA résoud des problèmes



La plupart des problèmes d'IA peuvent être formulés comme des tâches de **prédiction** où l'objectif est de prédire une sortie y à partir d'une entrée x .

Des **réseaux des neurones spécialisés** peuvent ainsi être entraînés et atteindre des performances supérieures à celles de l'homme sur des tâches complexes.

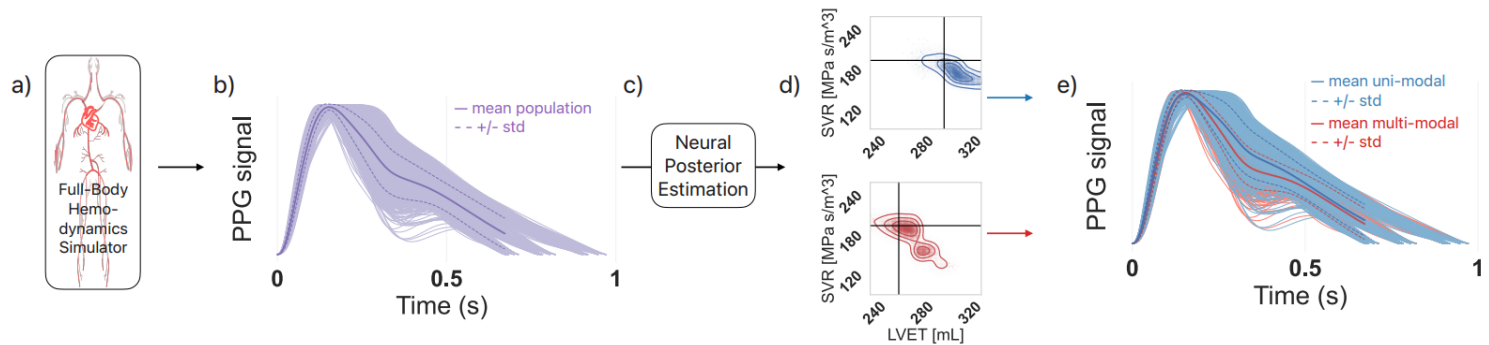


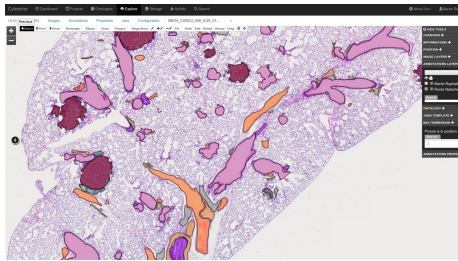


I: Jane went to the hallway.
I: Mary walked to the bathroom.
I: Sandra went to the garden.
I: Daniel went back to the garden.
I: Sandra took the milk there.
Q: Where is the milk?
A: garden

(Haut) Compréhension de scène, estimation de pose, raisonnement géométrique.

(Bas) Planification, annotation d'images, compréhension à la lecture.

Les réseaux de neurones forment des **primitives** qui peuvent être transférées à de nombreux domaines.



(Haut) Analyse des lames histologiques, débruitage des images IRM, détection des nævus. (Bas) Reconstruction de l'hémodynamique corporelle à partir de PPGs.



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



Later bekij...



Delen



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, 2018)



Le deep learning peut également **résoudre des problèmes que personne ne pouvait résoudre auparavant.**

AlphaFold: D'une séquence d'acides aminés à une structure 3D



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






AlphaFold: The making of a scientific breakt...



Later bekij...



Delen



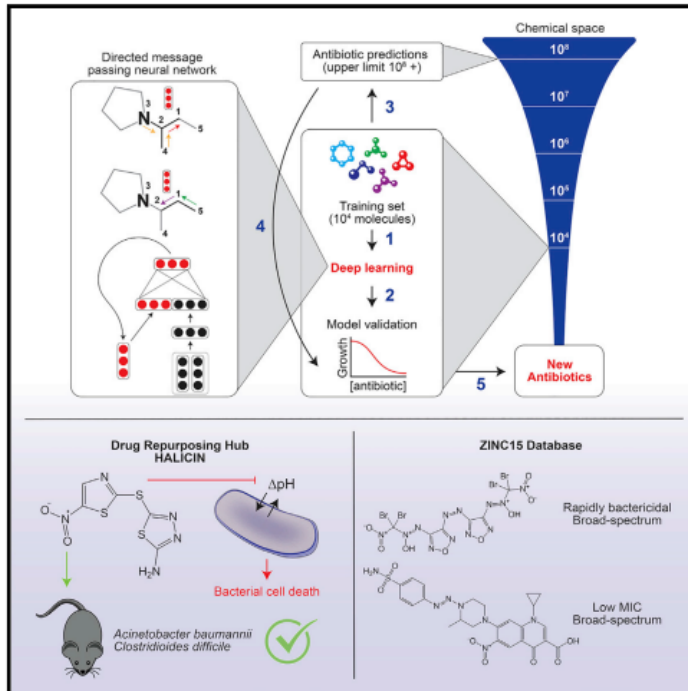
AI for Science (Deepmind, AlphaFold, 2020)

Découverte de médicaments avec des 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

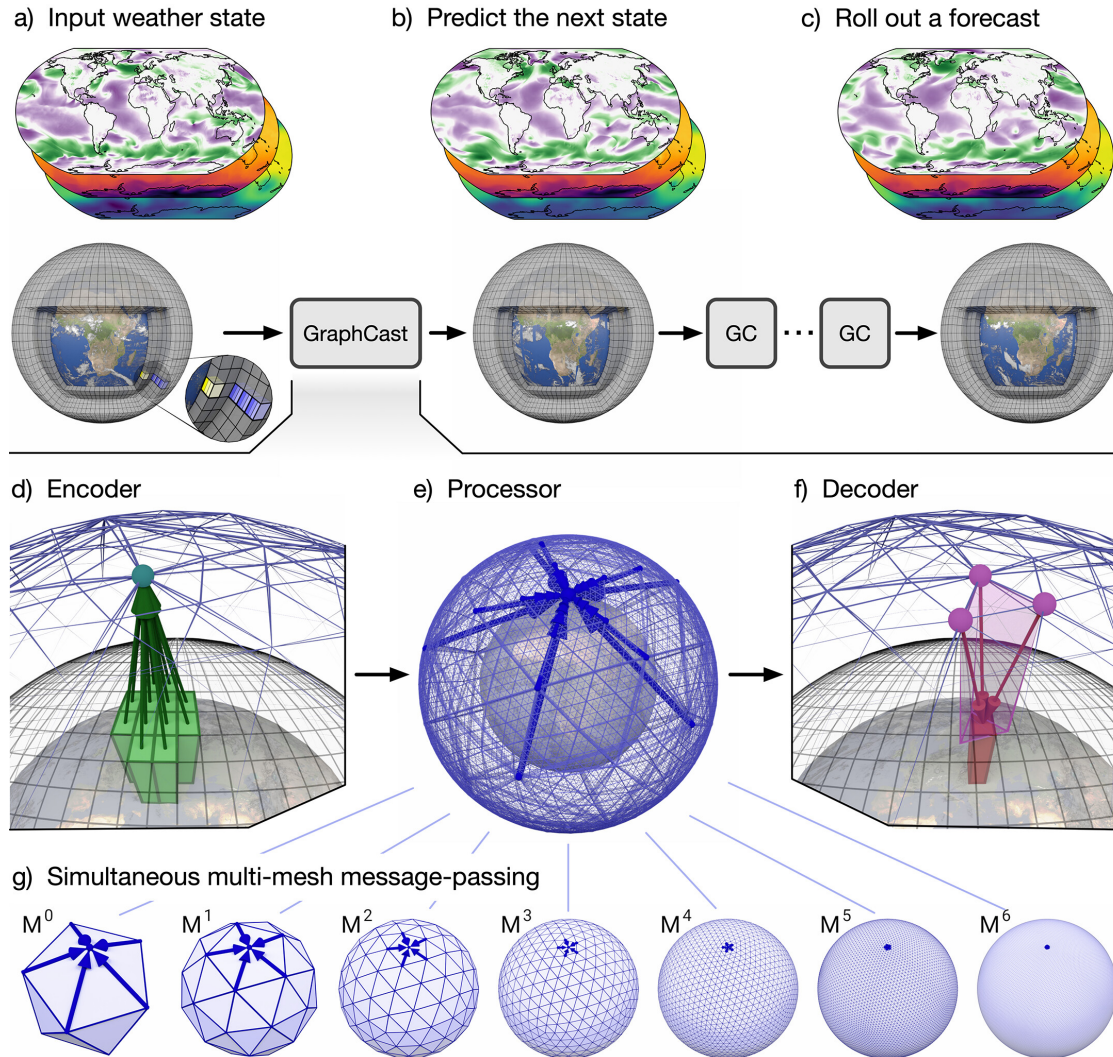
Correspondence

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.

GraphCast: des prédictions météorologiques rapides et précises



L'IA lit, écrit, entend, parle et voit

The breakthrough

Attention Is All You Need

Ashish Vaswani*
Google Brain
avaswani@google.com

Noam Shazeer*
Google Brain
noam@google.com

Niki Parmar*
Google Research
nikip@google.com

Jakob Uszkoreit*
Google Research
usz@google.com

Llion Jones*
Google Research
llion@google.com

Aidan N. Gomez* †
University of Toronto
aidan@cs.toronto.edu


Lukasz Kaiser*
Google Brain
lukasz.kaiser@google.com

Illia Polosukhin* ‡
illia.polosukhin@gmail.com

Abstract

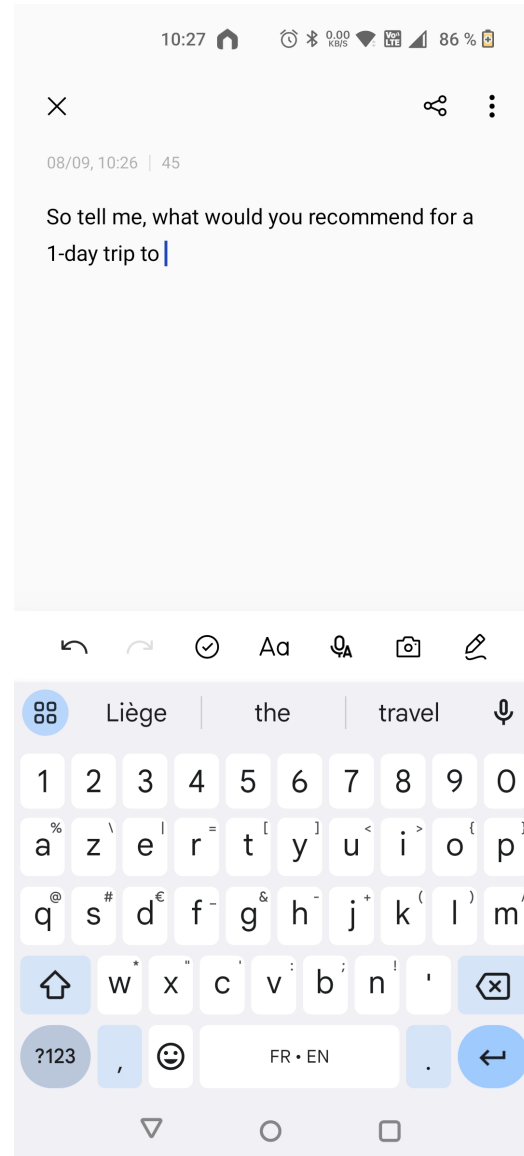
The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.

Vaswani et al., 2017.



Une idée simple:

Devinez le mot suivant



In the 1960s, Armstrong _____

In the 1960s, Armstrong performed ____

In the 1960s, Armstrong performed a moonwalk ____

In the 1960s, Armstrong performed a moonwalk on the ____


In the 1960s, Armstrong performed a moonwalk on the lunar ____

In the 1960s, Armstrong performed a moonwalk on the lunar surface and said ____



What does it mean to predict the next token well enough? It's actually [...] a deeper question than it seems. Predicting the next token well means that you understand the underlying reality that led to the creation of that token.

Ilya Sutskever, 2023.



Une simplicité brutale: De plus grands modèles, entraînés sur plus de données, et plus de ressources de calcul continuent de repousser les limites de ce qui est possible, **sans signe de ralentissement**.

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

+ Send a message





Claude 3.5 Sonnet as a writing partner



Later bekij...



Delen

The screenshot shows the Claude 3.5 Sonnet interface. The main chat window contains the following text:

potentially influence their current relationship and work together. It adds depth to the story by suggesting multiple avenues for backstory and character development.

Would you like me to explain any part of the genogram in more detail or explore how these family connections might influence the story?

5 Fantastic. I'm curious how best to showcase the arthropod <-> human interaction. Could you ideate on a sample of the translation?

Certainly! Let's explore a sample of how the arthropod-human interaction and translation might work between Claw'd and Dr. Wave. I'll create a brief dialogue that demonstrates this unique communication system.

Sample Dialogue: Claw'd and Dr. Wave's Translated Interaction
Click to open document

Reply to Claude...

Claude 3.5 Sonnet

An inset window titled "Sample Dialogue: Claw'd and Dr. Wave's Translated Interaction" displays the following content:

Claw'd and Dr. Wave: A Translated Interaction

Setting: An underwater research station. Dr. Wave is analyzing data on a holographic display. Claw'd approaches, his movements captured by an array of advanced sensors.

Claw'd: [Performs a series of rapid claw clicks, antennae vibrations, and releases a pattern of small bubbles]

AI (Translation Interface): Analyzing crab communication patterns... Processing...

Translated output: "Hey doc,

Last edited just now

Copy

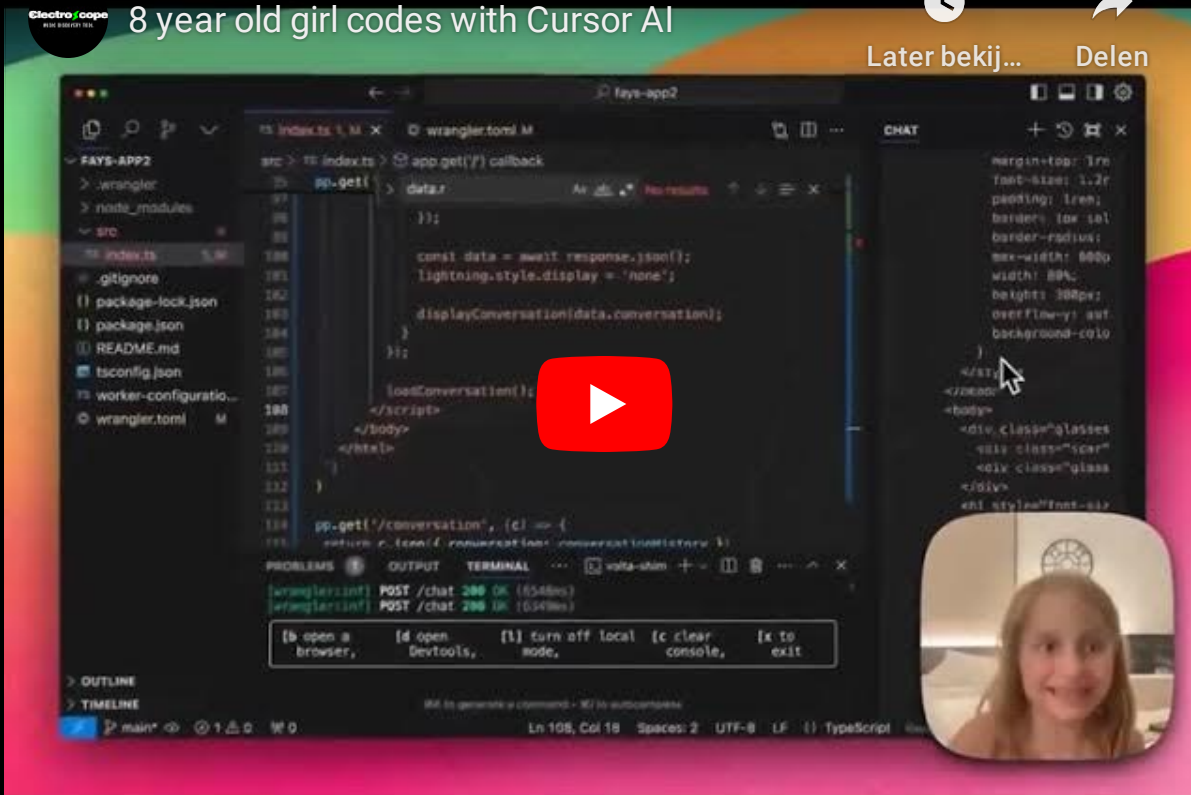
Conversational AI assistants (Anthropic, 2024)



Later bekij...



Delen



Code assistants (Cursor, 2024)



Rock, Paper, Scissors with GPT-4o



Later bekij...



Delen



Not just text, but also images and sounds (OpenAI, 2024)

Et maintenant?



The good

- L'IA est devenue un outil capable de résoudre des problèmes complexes et de faire des découvertes.
- L'interface homme-machine est en train de changer radicalement, avec des assistants conversationnels facilitant l'utilisation du numérique.
- Les progrès continuent à un rythme effréné.



The bad

- Les modèles de deep learning sont souvent difficiles à interpréter et à expliquer.
- L'IA se trompe, et peut se tromper de manière catastrophique malgré les apparences de précision.
- Les modèles de deep learning sont devenus très gros et nécessitent des ressources de calcul importantes, avec des conséquences environnementales significatives.



The ugly

- Les modèles de deep learning sont biaisés et peuvent perpétuer des discriminations.
- Les utilisations malveillantes de l'IA sont de plus en plus fréquentes (deep fakes, bots, manipulation, etc).
- Dans tous les domaines de la société, une dépendance à l'IA s'installe, avec des risques de déshumanisation et de perte de contrôle.

