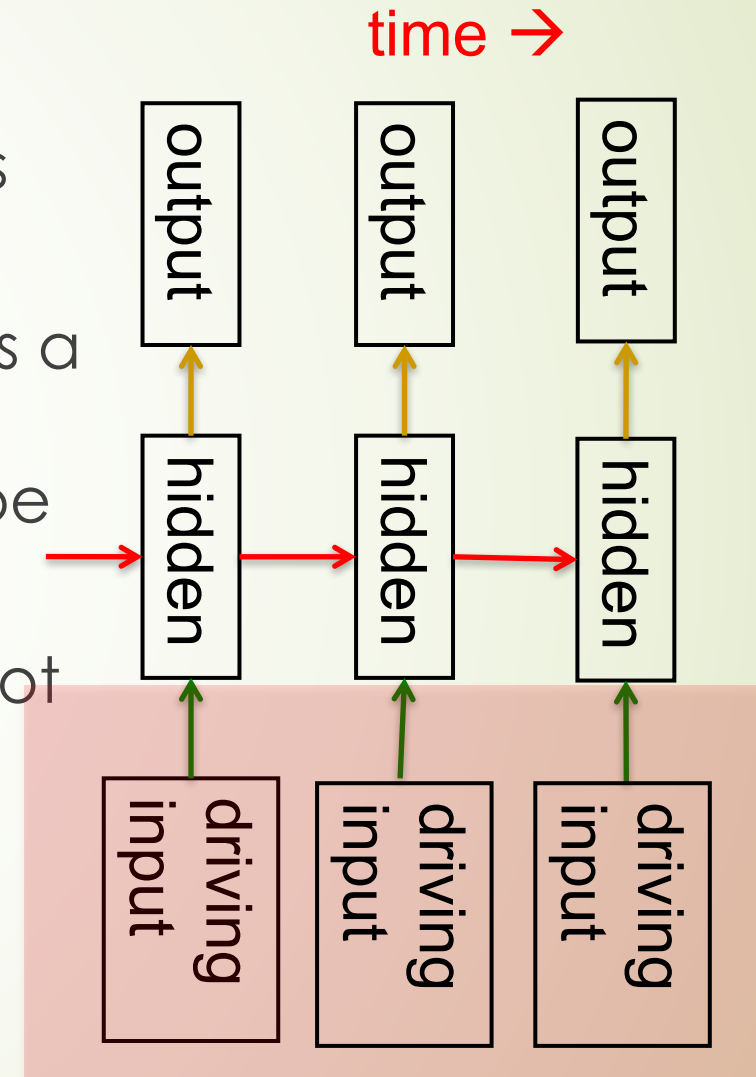


Time series: Linear Dynamical Systems (1940s-)

- ▶ The hidden state has linear dynamics with Gaussian noise and produces the observations using a linear model with Gaussian noise.
- ▶ Kalman Filter: A linearly transformed Gaussian is a Gaussian. So the distribution over the hidden state given the data so far is Gaussian. It can be computed using “Kalman filtering”.
- ▶ To predict the next output (so that we can shoot down the missile) we need to infer the hidden state.


$$h_t = W_{hh}h_{t-1} + W_{hx}x_t + \epsilon_t^h$$

$$y_t = W_{yh}h_t + W_{yx}x_t + \epsilon_t^y$$




Hidden Markov Models (1970s-)

- Hidden Markov Models have a discrete one-of-N hidden state. Transitions between states are stochastic and controlled by a transition matrix. The outputs produced by a state are stochastic.
- We cannot be sure which state produced a given output. So the state is “hidden”.
- It is easy to represent a probability distribution across N states with N numbers.
- To predict the next output we need to infer the probability distribution over hidden states.
 - HMMs have efficient algorithms (**Baum-Welch or EM Algorithm**) for inference and learning.
 - **Jim Simons** hires Lenny Baum as the founding member of Renaissance Technologies in 1979



time →



Lenny Baum became a devoted Go player despite his deteriorating eyesight.

Recurrent Neural Networks (1986-)

- **The issue of a hidden Markov model (HMM):**


- At each time step it must select one of its hidden states. So with N hidden states it can only remember $\log(N)$ bits about what it generated so far.

- RNNs are very powerful, because they combine two properties:

- Distributed hidden state that allows them to store a lot of information about the past efficiently.
 - Non-linear dynamics that allows them to update their hidden state in complicated ways.

- Rumelhart et al. enables training by **BP** algorithm

- With enough neurons and time, RNNs can compute anything that can be computed by your computer.




$$h_t = \sigma_h(W_{hh}h_{t-1} + W_{hx}x_t)$$

$$y_t = \sigma_y(W_{yh}h_t)$$

Long-Short-Term-Memory (LSTM)


- [Sepp Hochreiter; Jürgen Schmidhuber \(1997\). "Long short-term memory". *Neural Computation*. 9 \(8\): 1735–1780. \(https://www.bioinf.jku.at/publications/older/2604.pdf\)](https://www.bioinf.jku.at/publications/older/2604.pdf)
- Introduction of short path to learn deep networks without vanishing gradient problem.



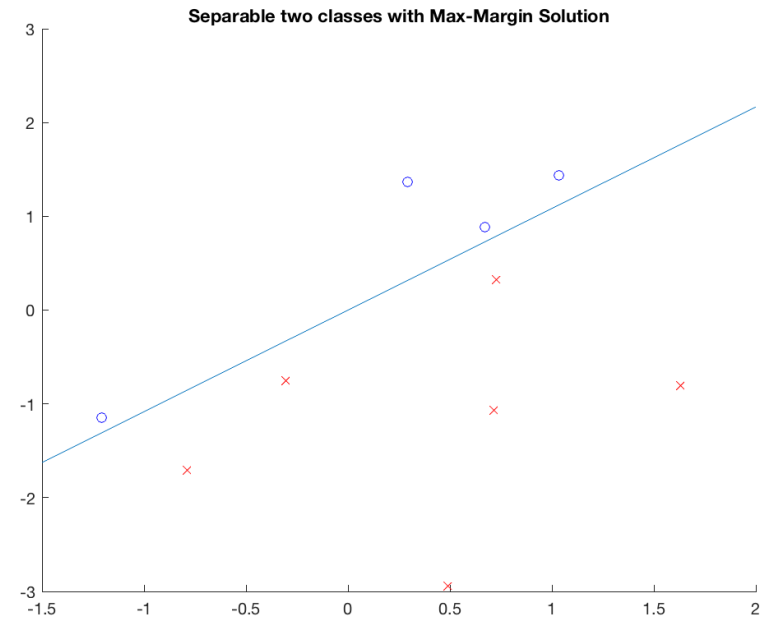
Max-Margin Classifier (SVM)

$$\text{minimize}_{\beta_0, \beta_1, \dots, \beta_p} \|\beta\|^2 := \sum_j \beta_j^2$$

subject to $y_i(\beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip}) \geq 1$ for all i



Vladimir Vapnik, 1994



MNIST Dataset Test Error

LeCun et al. 1998



Simple SVM performs
as well as Multilayer
Convolutional Neural
Networks which need
careful tuning (LeNets)


Dark era for NN: 1998-2012



2000-2010: The Era of SVM, Boosting, ... as nights of Neural Networks



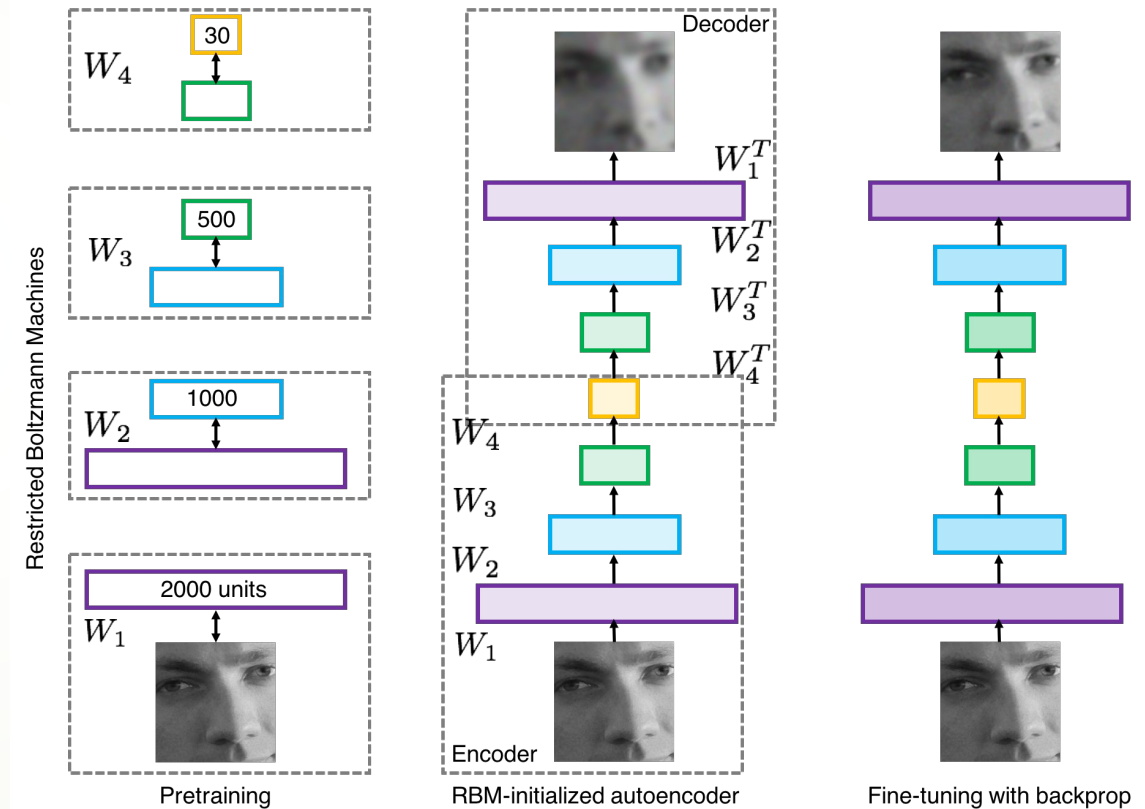
Decision Trees and Boosting



- Breiman, Friedman, Olshen, Stone, (1983): CART
- “The Boosting problem” (M. Kearns & L. Valiant): **Can a set of weak learners create a single strong learner?** (三个臭皮匠顶个诸葛亮?)
- Breiman (1996): Bagging
- Freund, Schapire (1997): **AdaBoost** (“the best off-the-shelf algorithm” by Breiman)
- Breiman (2001): **Random Forests**


Restricted Boltzman Machine, 2006 (Deep Learning)

- **Hinton and Salakhutdinov**, Reducing the Dimensionality of Data with Neural Networks, **Science, 2006**
- Reinvigorating research in Deep Learning
- Shows importance of **pretraining** (**greedy layer-wise, a.k.a. block coordinate descent**)




Around the year of 2012: return of NN as 'deep learning'

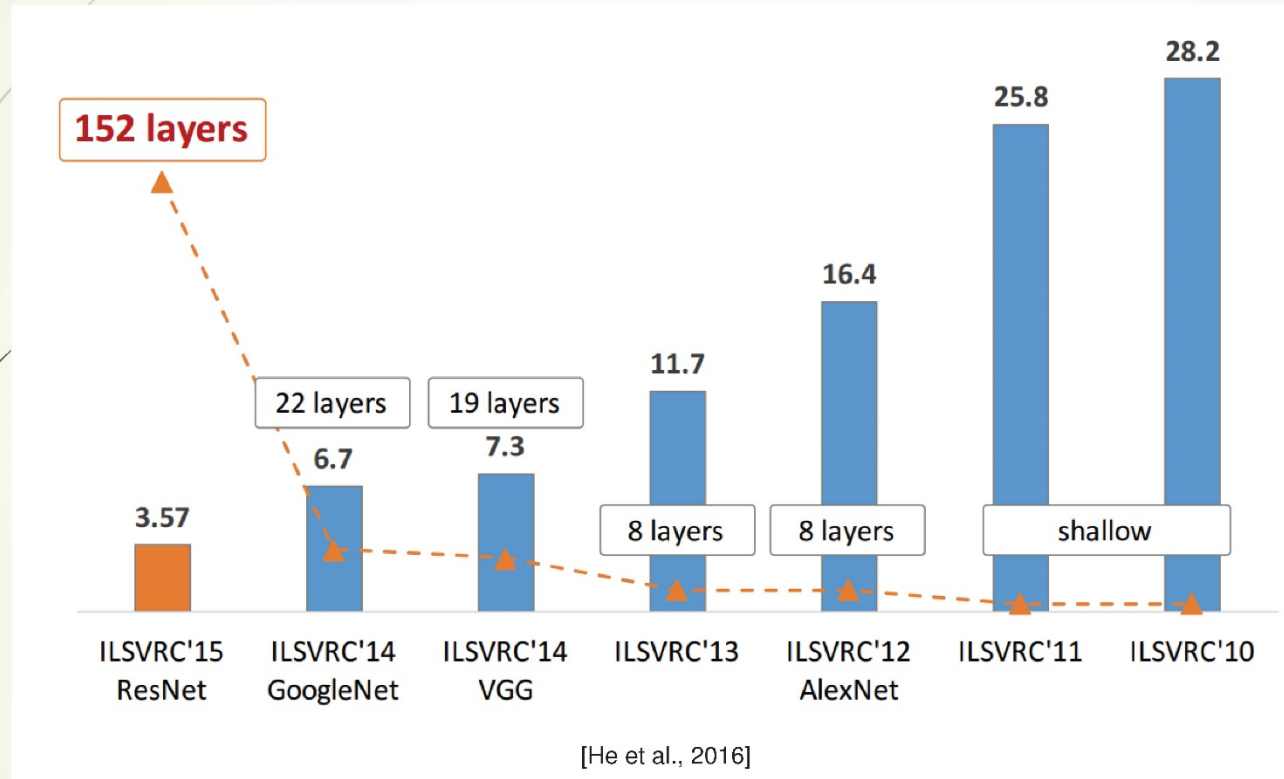
Speech Recognition: TIMIT



Computer Vision: ImageNet




Depth as function of year



ILSVRC ImageNet Top 5 errors


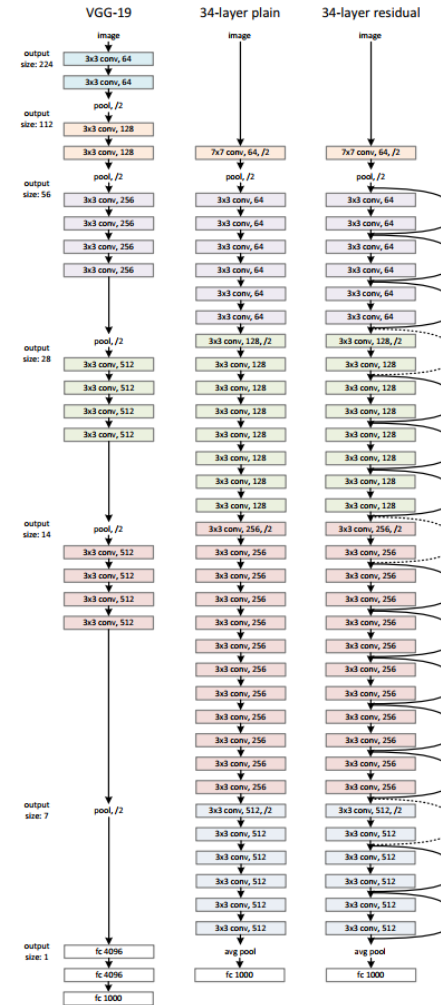

- ImageNet (subset):
 - 1.2 million training images
 - 100,000 test images
 - 1000 classes
- ImageNet large-scale visual recognition Challenge



ResNet (2015)

[He-Zhang-Ren-Sun, 2015]

- Solves problem by adding skip connections
- Very deep: 152 layers
- No dropout
- Stride
- Batch normalization



Source: Deep Residual Learning for Image Recognition

GPU + Big labeled data

"We're at the beginning of a new day...
This is the beginning of the AI revolution."
— Jensen Huang, GTC Taiwan 2017



兩股力量驅動電腦的未來

深度學習點亮人工智慧紀元。



受到人腦的啟發，深度神經網路具備上億的類神經連結，藉由巨量資料來學習，這仰賴極大量的運算。

同時，摩爾定律已到了尾聲 - CPU已不可能再擴張成長。

程式設計人員無法創造出可以更有效率發現更多指令級並行性的CPU架構。

電晶體持續每年增長50%，但是CPU效能僅能成長10%。

TWO FORCES DRIVING THE FUTURE OF COMPUTING




NVIDIA


Reaching Human Performance Level in Games



Deep Blue in 1997




AlphaGo "LEE" 2016




Natural Language Processing (NLP) and Machine Translation

- ▶ In **2013-2015**, **LSTMs** started achieving state-of-the-art results
 - ▶ Successful tasks include: handwriting recognition, speech recognition, machine translation, parsing, image captioning
 - ▶ LSTM became the dominant approach
- ▶ In **2019**, other approaches (e.g. **Transformers**) have become more dominant for certain tasks.
 - ▶ For example in **WMT** (a MT conference + competition):
 - ▶ In WMT 2016, the summary report contains "RNN" 44 times
 - ▶ In WMT 2018, the report contains "RNN" 9 times and "Transformer" 63 times
 - ▶ **Source:** "Findings of the 2016 Conference on Machine Translation (WMT16)", Bojar et al. 2016, <http://www.statmt.org/wmt16/pdf/W16-2301.pdf>
 - ▶ **Source:** "Findings of the 2018 Conference on Machine Translation (WMT18)", Bojar et al. 2018, <http://www.statmt.org/wmt18/pdf/WMT028.pdf>

Rapid Progress for NLP Pretraining (GLUE Benchmark)



More compute, more better?



Protein Folding Structure Prediction

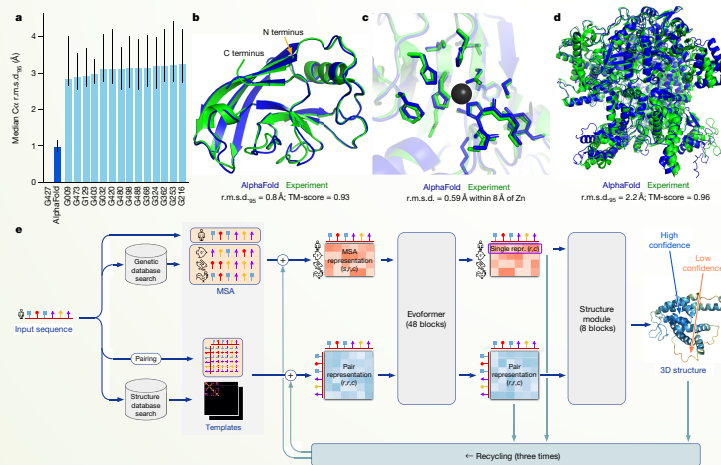



Fig. 1 | AlphaFold produces highly accurate structures. **a**, The performance of AlphaFold on the CASP14 dataset ($n = 87$ protein domains) relative to the top 15 entries (out of 146 entries), group numbers correspond to the numbers assigned to entrants by CASP. Data are median and the 95% confidence interval of the median, estimated from 10,000 bootstrap samples. **b**, Our prediction of CASP14 target T1049 (PDB 6Y4F, blue) compared with the true (experimental) structure (green). Four residues in the C-terminus of the crystal structure are B-factor outliers and are not depicted. **c**, CASP14 target T1056 (PDB 6Y1J). An example of a well-predicted zinc-binding site (AlphaFold has accurate side chains even though it does not explicitly predict the zinc ion). **d**, CASP target T1044 (PDB 6VR4) – a 2,180-residue single chain – was predicted with correct domain packing (the prediction was made after CASP using AlphaFold without intervention). **e**, Model architecture. Arrows show the information flow among the various components described in this paper. Array shapes are shown in parentheses with s , number of sequences (N_{seq} in the main text); r , number of residues (N_{res} in the main text); c , number of channels.

Article

Highly accurate protein structure prediction with AlphaFold

<https://doi.org/10.1038/s41586-021-03819-2>

Received: 11 May 2021

Accepted: 12 July 2021

Published online: 15 July 2021

Open access

Check for updates

John Jumper^{1,4,5*}, Richard Evans^{1,4}, Alexander Pritzel^{1,4}, Tim Green^{1,4}, Michael Figurnov^{1,4}, Olaf Ronneberger^{1,4}, Kathryn Tunyasuvunakool^{1,4}, Russ Bates^{1,4}, Augustin Zidek^{1,4}, Anna Potapenko^{1,4}, Alex Bridgland^{1,4}, Clemens Meyer^{1,4}, Simon A. A. Kohli^{1,4}, Andrew J. Ballard^{1,4}, Andrew Cowie^{1,4}, Bernardino Romera-Paredes^{1,4}, Stanislav Nikolov^{1,4}, Rishub Jain^{1,4}, Jonas Adler¹, Trevor Back¹, Stig Petersen¹, David Reiman¹, Ellen Clancy¹, Michal Zielinski¹, Martin Steinegger^{2,3}, Michalina Pacholska¹, Tamas Berghammer¹, Sebastian Bodenstein¹, David Silver¹, Oriol Vinyals¹, Andrew W. Senior¹, Koray Kavukcuoglu¹, Pushmeet Kohli¹ & Demis Hassabis^{1,4,5*}

Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function. Through an enormous experimental effort^{1–4}, 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^{10–14}, existing methods fall far short of atomic accuracy, especially when no homologous structure is available. Here we provide the first computational method that can regularly predict protein structures with atomic accuracy even in cases in which no similar structure is known. We validated an entirely redesigned version of our neural network-based model, AlphaFold, in the challenging 14th Critical Assessment of protein Structure Prediction (CASP14)¹⁵, demonstrating accuracy competitive with experimental structures in a majority of cases and greatly outperforming other methods. Underpinning the latest version of AlphaFold is a novel machine learning approach that incorporates physical and biological knowledge about protein structure, leveraging multi-sequence alignments, into the design of the deep learning algorithm.

The development of computational methods to predict three-dimensional (3D) protein structures from the protein sequence has proceeded along two complementary paths that focus on either the physical interactions or the evolutionary history. The physical interaction programme heavily integrates our understanding of molecular driving forces into either thermodynamic or kinetic simulation of protein physics¹⁶ or statistical approximations thereof¹⁷. Although theoretically very appealing, this approach has proved highly challenging for even moderate-sized proteins due to the computational intractability of molecular simulation, the context dependence of protein stability and the difficulty of producing sufficiently accurate models of protein physics. The evolutionary programme has provided an alternative in recent years, in which the constraints on protein structure are derived from bioinformatics analysis of the evolutionary history of proteins, homology to solved structures^{18,19} and pairwise evolutionary correlations^{20–24}. This bioinformatics approach has benefited greatly from

the steady growth of experimental protein structures deposited in the Protein Data Bank (PDB)⁵, the explosion of genomic sequencing and the rapid development of deep learning techniques to interpret these correlations. Despite these advances, contemporary physical and evolutionary-history-based approaches produce predictions that are far short of experimental accuracy in the majority of cases in which a close homologue has not been solved experimentally and this has limited their utility for many biological applications.

In this study, we develop the first, to our knowledge, computational approach capable of predicting protein structures to near experimental accuracy in a majority of cases. The neural network AlphaFold that we developed was entered into the CASP14 assessment (May–July 2020; entered under the team name ‘AlphaFold2’ and a completely different model from our CASP13 AlphaFold system¹⁰). The CASP assessment is carried out biennially using recently solved structures that have not been deposited in the PDB or publicly disclosed so that it is a blind test

¹DeepMind, London, UK. ²School of Biological Sciences, Seoul National University, Seoul, South Korea. ³Artificial Intelligence Institute, Seoul National University, Seoul, South Korea. *These authors contributed equally: John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Zidek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohli, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Demis Hassabis. [†]e-mail: jumper@deepmind.com; dhcontact@deepmind.com

AI for Science

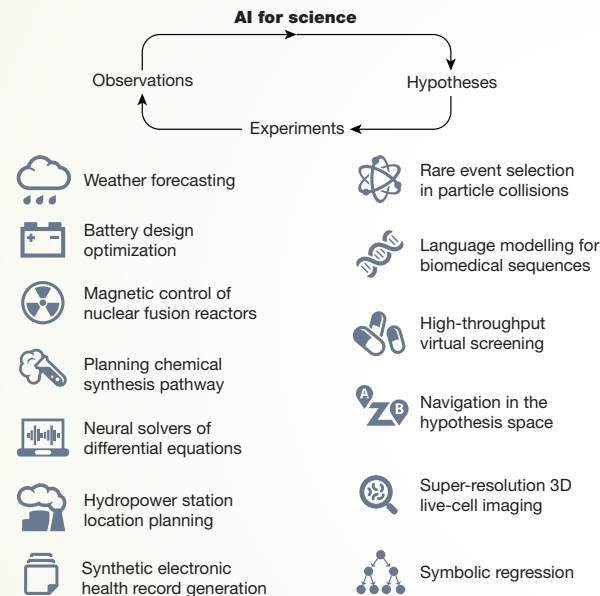


Fig. 1 | Science in the age of artificial intelligence. Scientific discovery is a multifaceted process that involves several interconnected stages, including hypothesis formation, experimental design, data collection and analysis. AI is poised to reshape scientific discovery by augmenting and accelerating research at each stage of this process. The principles and illustrative studies shown here highlight the contributions to enhance scientific understanding and discovery.

Review

Scientific discovery in the age of artificial intelligence

<https://doi.org/10.1038/s41586-023-06221-2>

Received: 30 March 2022

Accepted: 16 May 2023

Published online: 2 August 2023

Check for updates

Hanchen Wang^{1,2,3,7,38,39}, Tianfan Fu^{3,39}, Yuanqi Du^{4,39}, Wenhao Gao⁵, Kexin Huang⁶, Ziming Liu⁷, Payal Chandak⁸, Shengchao Liu^{9,10}, Peter Van Katwyk^{11,12}, Andreea Deac^{9,10}, Anima Anandkumar^{2,13}, Karianne Bergen^{11,12}, Carla P. Gomes⁴, Shirley Ho^{14,15,16,17}, Pushmeet Kohli¹⁸, Joan Lasenby¹, Jure Leskovec⁶, Tie-Yan Liu¹⁹, Arjun Manrai²⁰, Debora Marks^{21,22}, Bharath Ramsundar²³, Le Song^{24,25}, Jimeng Sun²⁶, Jian Tang^{9,27,28}, Petar Veličković^{17,29}, Max Welling^{30,31}, Linfeng Zhang^{32,33}, Connor W. Coley^{5,34}, Yoshua Bengio^{9,10} & Marinka Zitnik^{20,22,35,36,37}

Artificial intelligence (AI) is being increasingly integrated into scientific discovery to augment and accelerate research, helping scientists to generate hypotheses, design experiments, collect and interpret large datasets, and gain insights that might not have been possible using traditional scientific methods alone. Here we examine breakthroughs over the past decade that include self-supervised learning, which allows models to be trained on vast amounts of unlabelled data, and geometric deep learning, which leverages knowledge about the structure of scientific data to enhance model accuracy and efficiency. Generative AI methods can create designs, such as small-molecule drugs and proteins, by analysing diverse data modalities, including images and sequences. We discuss how these methods can help scientists throughout the scientific process and the central issues that remain despite such advances. Both developers and users of AI tools need a better understanding of when such approaches need improvement, and challenges posed by poor data quality and stewardship remain. These issues cut across scientific disciplines and require developing foundational algorithmic approaches that can contribute to scientific understanding or acquire it autonomously, making them critical areas of focus for AI innovation.

The foundation for forming scientific insights and theories is laid by how data are collected, transformed and understood. The rise of deep learning in the early 2010s has significantly expanded the scope and ambition of these scientific discovery processes¹. Artificial intelligence (AI) is increasingly used across scientific disciplines to integrate massive datasets, refine measurements, guide experimentation, explore the space of theories compatible with the data, and provide actionable and reliable models integrated with scientific workflows for autonomous discovery.

Data collection and analysis are fundamental to scientific understanding and discovery, two of the central aims in science², and quantitative

methods and emerging technologies, from physical instruments such as microscopes to research techniques such as bootstrapping, have long been used to reach these aims³. The introduction of digitization in the 1950s paved the way for the general use of computing in scientific research. The rise of data science since the 2010s has enabled AI to provide valuable guidance by identifying scientifically relevant patterns from large datasets.

Although scientific practices and procedures vary across stages of scientific research, the development of AI algorithms cuts across traditionally isolated disciplines (Fig. 1). Such algorithms can enhance the design and execution of scientific studies. They are becoming

¹Department of Engineering, University of Cambridge, Cambridge, UK. ²Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA, USA. ³Department of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA, USA. ⁴Department of Computer Science, Cornell University, Ithaca, NY, USA. ⁵Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA. ⁶Department of Computer Science, Stanford University, Stanford, CA, USA. ⁷Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA. ⁸Harvard-MIT Program in Health Sciences and Technology, Cambridge, MA, USA. ⁹Mila - Quebec AI Institute, Montreal, Quebec, Canada. ¹⁰Université de Montréal, Montreal, Quebec, Canada. ¹¹Department of Earth, Environmental and Planetary Sciences, Brown University, Providence, RI, USA. ¹²Data Science Institute, Brown University, Providence, RI, USA. ¹³NVIDIA, Santa Clara, CA, USA. ¹⁴Center for Computational Astrophysics, Flatiron Institute, New York, NY, USA. ¹⁵Department of Astrophysical Sciences, Princeton University, Princeton, NJ, USA. ¹⁶Department of Physics, Carnegie Mellon University, Pittsburgh, PA, USA. ¹⁷Department of Physics and Center for Data Science, New York University, New York, NY, USA. ¹⁸Google DeepMind, London, UK. ¹⁹Microsoft Research, Beijing, China. ²⁰Department of Biomedical Informatics, Harvard Medical School, Boston, MA, USA. ²¹Department of Systems Biology, Harvard Medical School, Boston, MA, USA. ²²Broad Institute of MIT and Harvard, Cambridge, MA, USA. ²³Deep Forest Sciences, Palo Alto, CA, USA. ²⁴BioMap, Beijing, China. ²⁵Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, United Arab Emirates. ²⁶University of Illinois at Urbana-Champaign, Champaign, IL, USA. ²⁷HEC Montréal, Montreal, Quebec, Canada. ²⁸CIFAR AI Chair, Toronto, Ontario, Canada. ²⁹Department of Computer Science and Technology, University of Cambridge, Cambridge, UK. ³⁰University of Amsterdam, Amsterdam, Netherlands. ³¹Microsoft Research Amsterdam, Amsterdam, Netherlands. ³²DP Technology, Beijing, China. ³³AI for Science Institute, Beijing, China. ³⁴Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA. ³⁵Harvard Data Science Review, Cambridge, MA, USA. ³⁶Kempner Institute for the Study of Natural and Artificial Intelligence, Harvard University, Cambridge, MA, USA. ³⁷Present address: Department of Research and Early Development, Genentech Inc., South San Francisco, CA, USA. ³⁸Present address: Department of Computer Science, Stanford University, Stanford, CA, USA. ³⁹These authors contributed equally: Hanchen Wang, Tianfan Fu, Yuanqi Du. [✉]e-mail: marinka@hms.harvard.edu

ChatGPT (GPT 3.5-4)

ChatGPT

🌐 94 languages ▾

Article [Talk](#)

[Read](#) [View source](#) [View history](#) [Tools](#) ▾

From Wikipedia, the free encyclopedia



ChatGPT, which stands for **Chat Generative Pre-trained Transformer**, is a [large language model](#)-based [chatbot](#) developed by [OpenAI](#) and launched on November 30, 2022, notable for enabling users to refine and steer a conversation towards a desired length, format, style, level of detail, and [language](#) used. Successive prompts and replies, known as [prompt engineering](#), are considered at each conversation stage as a context.^[2]

ChatGPT is built upon [GPT-3.5](#) and [GPT-4](#) —members of OpenAI's proprietary series of [generative pre-trained transformer](#) (GPT) models, based on the [transformer](#) architecture developed by [Google](#)^[3]—and it is [fine-tuned](#) for conversational applications using a combination of [supervised](#) and [reinforcement learning](#) techniques.^[4] ChatGPT was released as a freely available research preview, but due to its popularity, OpenAI now operates the service on a [freemium model](#). It allows users on its free tier to access the GPT-3.5-based version. In contrast, the more advanced GPT-4 based version and priority access to newer features are provided to paid subscribers under the commercial name "ChatGPT Plus".

By January 2023, it had become what was then the fastest-growing consumer software application in history, gaining over 100 million users and contributing to OpenAI's [valuation](#) growing to [US\\$29 billion](#).^{[5][6]} Within months, [Google](#), [Baidu](#), and [Meta](#) accelerated the development of their competing products: [Bard](#), [Ernie Bot](#), and [LLaMA](#).^[7] Microsoft launched its [Bing Chat](#) based on OpenAI's GPT-4. Some observers expressed concern over the potential of ChatGPT to displace or atrophy [human intelligence](#) and its potential to enable [plagiarism](#) or fuel [misinformation](#).^{[4][8]}

Training

ChatGPT is based on particular [GPT foundation models](#), namely [GPT-3.5](#) and [GPT-4](#), that were [fine-tuned](#) to target conversational usage.^[9] The fine-tuning process leveraged both [supervised learning](#) as well as

ChatGPT



Developer(s)	OpenAI
Initial release	November 30, 2022; 9 months ago
Stable release	August 3, 2023; 31 days ago ^[1]
Written in	Python
Engine	GPT-3.5 GPT-4
Platform	Cloud computing platforms
Type	Chatbot Large language model Generative text-to-image model Generative pre-trained transformer
License	Proprietary
Website	chat.openai.com/chat

Number of AI papers on arXiv, 2010-2019

Number of AI papers on arXiv, 2010-2019

Source: arXiv, 2019.





Fig. 1.6.

Growth of Deep Learning

'Deep Learning' is coined by Hinton et al. in their Restricted Boltzman Machine paper, *Science* 2006, not yet popular until championing ImageNet competitions.



Some Cold Water: Tesla Autopilot Misclassifies Truck as Billboard




Problem: Why? How can you trust a blackbox?

Deep Learning may be fragile in generalization against noise!


 x

“panda”

57.7% confidence

 $+ .007 \times$

 $\text{sign}(\nabla_x J(\theta, x, y))$

“nematode”

8.2% confidence


 $=$

 $x +$
 $\epsilon \text{sign}(\nabla_x J(\theta, x, y))$

“gibbon”

99.3 % confidence


[Goodfellow et al., 2014]



“black hole”

87.7% confidence

 $+ .007 \times$

 $=$


“donut”

99.3% confidence



CNN learns **texture** features, not **shapes**



(a) Texture image
81.4% **Indian elephant**
10.3% indri
8.2% black swan



(b) Content image
71.1% **tabby cat**
17.3% grey fox
3.3% Siamese cat




(c) Texture-shape cue conflict
63.9% **Indian elephant**
26.4% indri
9.6% black swan

Geirhos et al. ICLR 2019

<https://videoken.com/embed/W2HvLBMhCJQ?tocitem=46>

Lack of Causality or Interpretability

- ImageNet training learns non-semantic texture features: after random shuffling of patches, shapes information are destroyed which does not affect CNN's performance much.





Capture spurious correlations and can't do causal inference on **counterfactuals**

<https://videoken.com/embed/8UxS4ls6g1g?tocitem=2>

Leon Bottou, ICLR 2019

Example: detection of the action "giving a phone call"



(Oquab et al., CVPR 2014)
~70% correct (SOTA in 2014)



Not giving a phone call.

Giving a phone call ????

Overfitting causes **privacy leakage**

- ▶ Model inversion attack leaks privacy



Figure: Recovered (Left), Original (Right)

What's wrong with deep learning?

Ali Rahimi NIPS'17: Machine (deep) Learning has become **alchemy**.


<https://www.youtube.com/watch?v=ORHFOndEzPc>

Yann LeCun CVPR'15, invited talk: **What's wrong with deep learning?**
One important piece: **missing some theory (clarity in understanding)!**

<http://techtalks.tv/talks/whats-wrong-with-deep-learning/61639/>



Being alchemy is certainly not a shame, not wanting to work on advancing to chemistry is a shame! -- **by Eric Xing**




“ Shall we see soon an
emergence
from Alchemy to Science
in deep leaning? ”

How can we teach our students in the next generation science rather than alchemy?


Kaggle survey: Top Data Science Methods

<https://www.kaggle.com/surveys/2017>

Academic



Industry




What type of data is used at work?

<https://www.kaggle.com/surveys/2017>

Academic

What type of data is used at work?


Relational data is the most commonly reported type of data used at work for all industries except for **Academia** and the **Military and Security** industry where text data's used more.



Industry

What type of data is used at work?

Relational data is the most commonly reported type of data used at work for all industries except for **Academia** and the **Military and Security** industry where text data's used more.



All models are wrong, but some are useful ...


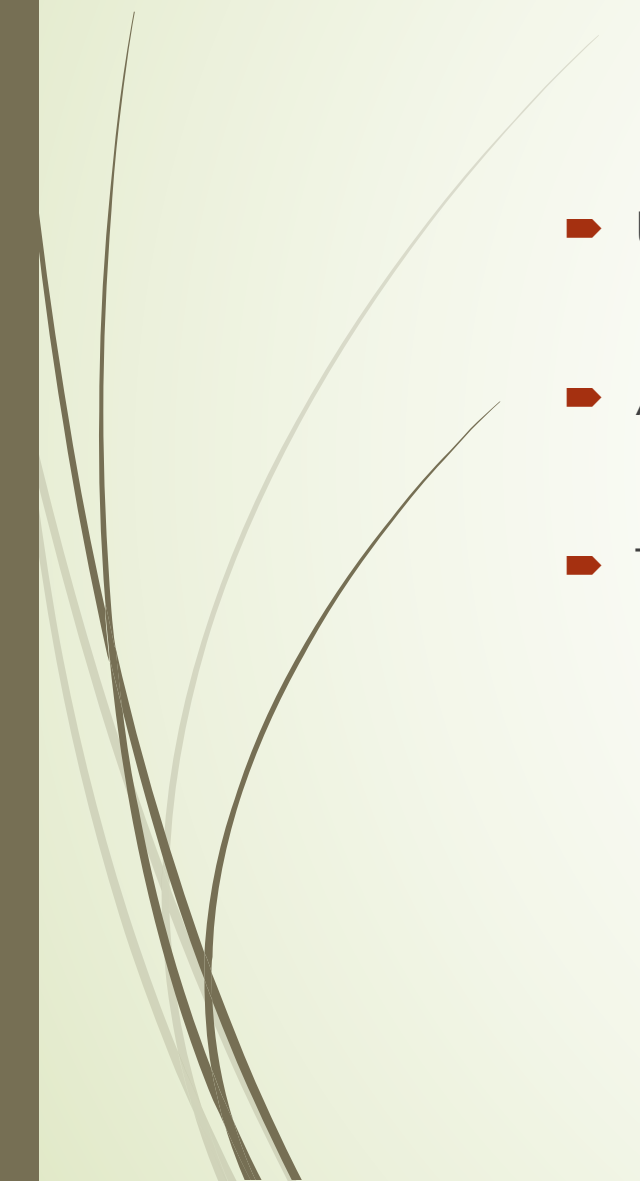


Figure 7: George Box: “Essentially, all models are wrong, but some are useful.”



In this class

- Understand its principles: statistics, optimization
 - Analyze the real world data with the methods
 - Team-work in projects
- 

Thank you!

