An introduction to AI

…and why you might avoid that term

Brendan Clarke, NHS Education for Scotland

01/07/2024

KIND Learning Network

Welcome

- this session is 🌶️: for beginners
- we'll get going properly at 13.05
- if you can't access the chat, you might need to join our Teams channel: tinyurl.com/kindnetwork
- you can find session materials at tinyurl.com/kindtrp

KIND Learning Network
NHS Education for Scotland

The KIND network

- a social learning space for staff working with **k**nowledge, **i**nformation, and **d**ata across health, social care, and housing in Scotland
- we offer social support, free training, mentoring, community events, …
- Teams channel / mailing list

KIND Learning Network

Motive

- There's a *lot* of hype about AI at the moment (see this graph)
- Underneath the hype, there's a lot of genuinely exciting stuff going on too
- That exciting stuff is likely to have some impact on health and care work
- But the timing and nature of that impact is unclear

Three questions for you

KIND Learning Network
NHS Education for Scotland

What does AI mean to you?

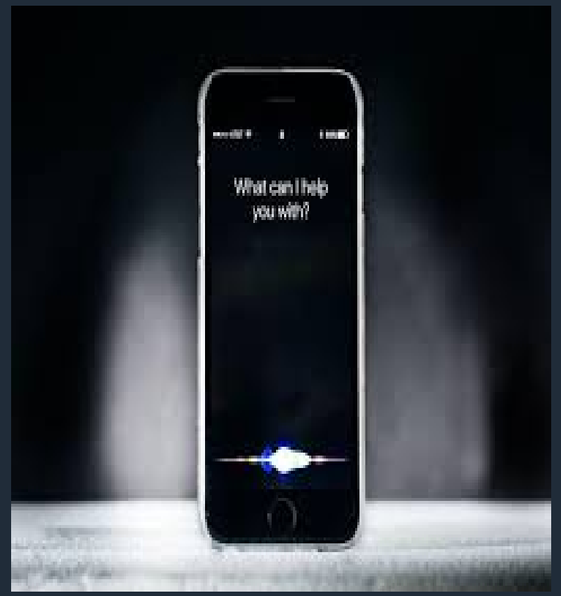

KIND
Learning Network
NHS Education for Scotland

Is AI…

- Over-hyped?
- Somewhere in between?
- Neglected?
- Other / don’t know

KIND Learning Network
NHS Education for Scotland

Do submarines swim?

About this talk

- AI is hard
  - lots of different technologies
  - lots of new words
  - lots of promises and implications
- So let’s start with a thought experiment

The Chinese room

Searle (1980)

> “Suppose that I’m locked in a room and given a large batch of Chinese writing. Suppose furthermore (as is indeed the case) that I know no Chinese, either written or spoken, and that I’m not even confident that I could recognize Chinese writing”

However, he is supplied with a set of intelligible rules for manipulating these Chinese symbols

“火” is the opposite of “水”

“六” is more than “四”

Question

Does this poor bloke locked in a room understand the Chinese symbols?

Now suppose that we start asking him questions (in English):

Is “六” more than “四”?

If so, respond with “是”. Otherwise respond “不”

Question

- Is understanding the same thing as being able to produce output in response to input?
- Searle (1980) - this is the difference between strong and weak AI

Back to nice safe words

- we usually don’t worry too much about what words like intelligence, understanding, etc really mean
- for most purposes, understanding something, and doing that thing, pretty well overlap
- AI, unfortunately, is an exception
- big difference between producing output and understanding here

Why does this matter?

- Because the current conversation around AI does violence to our usual understanding of basic terms (like intelligence)
  - We need to do a bit of re-interpreting…
  - …particularly because AI can do the input-output part *really* well
- (side effect) The Chinese Room is an excellent way of understanding what's going on inside some of the current tech

What are we talking about

- AI = big umbrella term, problematic
  - understanding?
- Let’s stick to some narrower concepts
  - *Algorithms* = rule-based ways of producing sensible output
  - *Expert systems* = more sophisticated expertise-based production of output
  - *Machine learning* = umbrella term for non-expertise-based production of output
  - *Large Language Models* = sub-species of machine learning

So what's an algorithm?

KIND Learning Network
NHS Education for Scotland

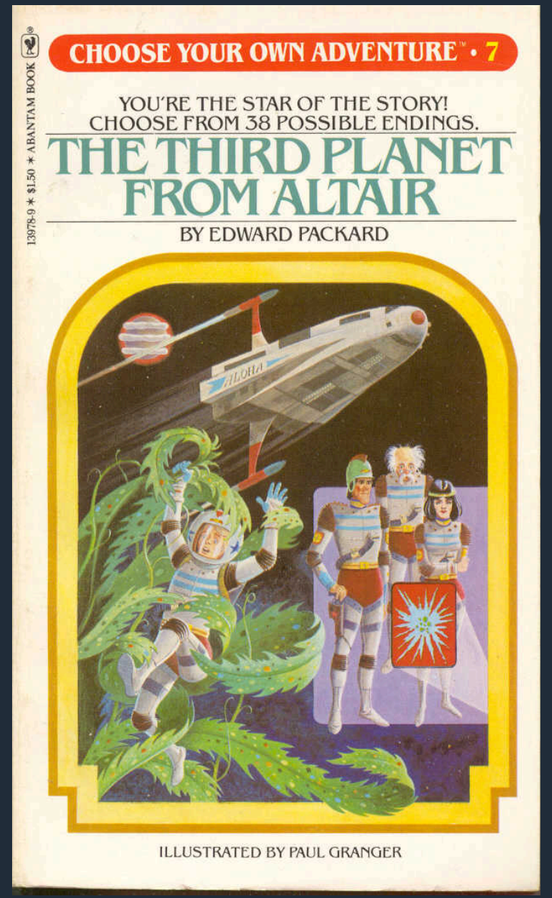

(Packard 1979)

- Algorithm = rule (roughly)
  - if something happens, do something
- made from expert input and evidence

An example algorithm

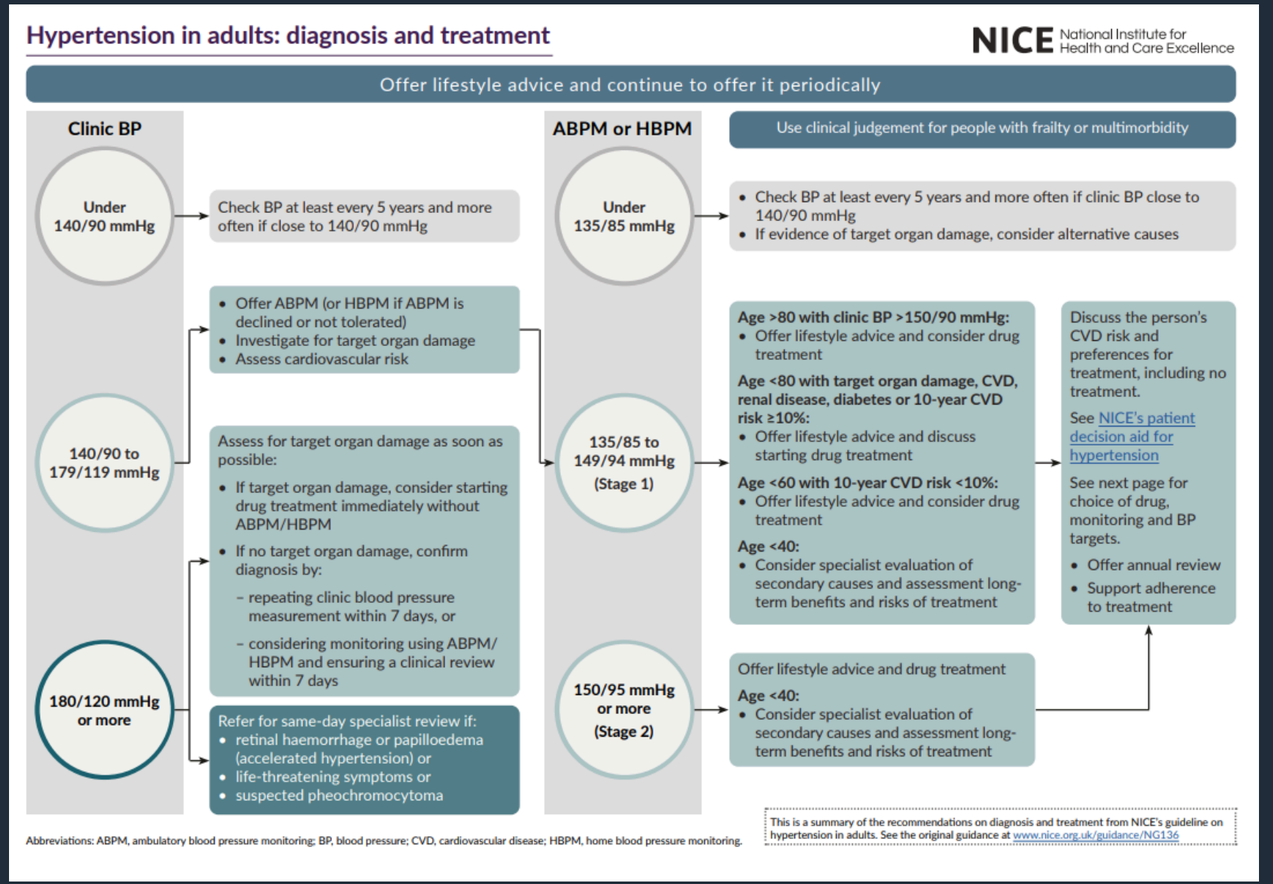

KIND Learning Network
NHS Education for Scotland

Related expertise-based tools

"See also…" references in indexes, library catalogues, wikipedia

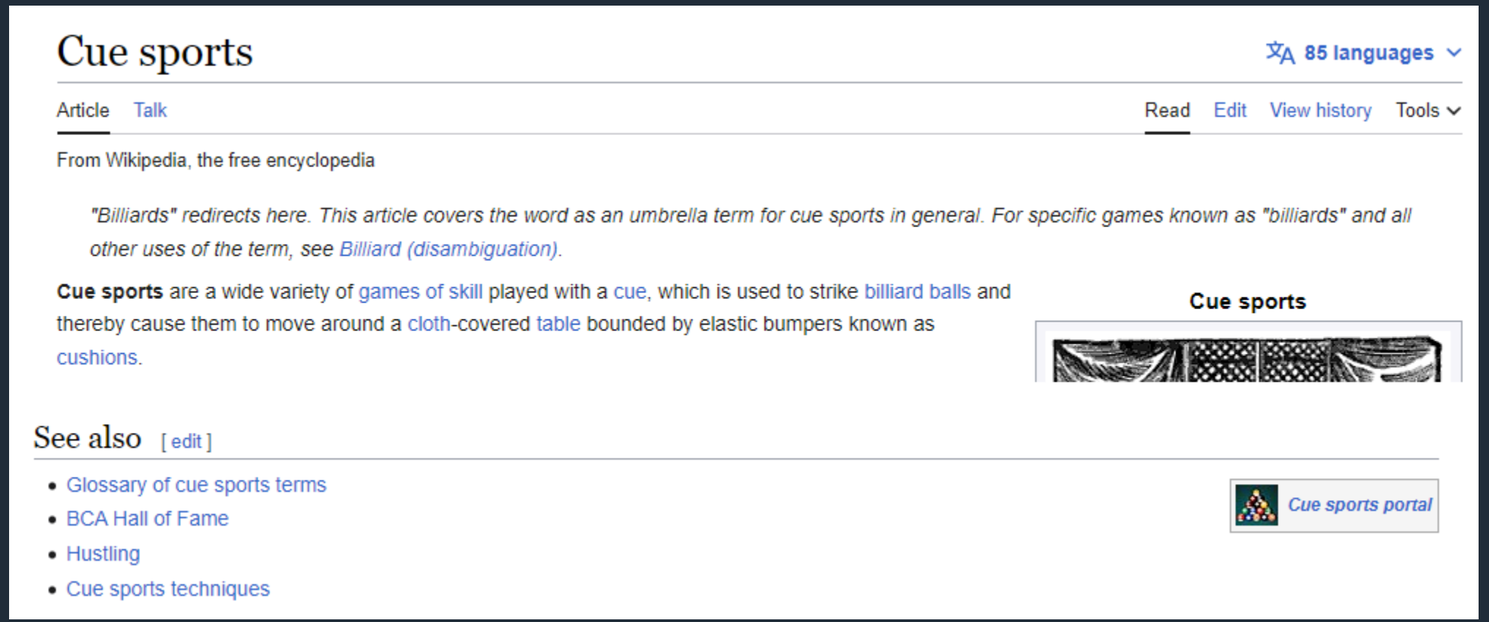

Cue sports

85 languages

Article Talk Read Edit View history Tools

From Wikipedia, the free encyclopedia

"Billiards" redirects here. This article covers the word as an umbrella term for cue sports in general. For specific games known as "billiards" and all other uses of the term, see Billiard (disambiguation).

Cue sports are a wide variety of games of skill played with a cue, which is used to strike billiard balls and thereby cause them to move around a cloth-covered table bounded by elastic bumpers known as cushions.

Cue sports

See also [edit]

- Glossary of cue sports terms
- BCA Hall of Fame
- Hustling
- Cue sports techniques

Cue sports portal

. . .

KIND Learning Network

NHS Education for Scotland

How about something more complicated?

KIND Learning Network
NHS Education for Scotland

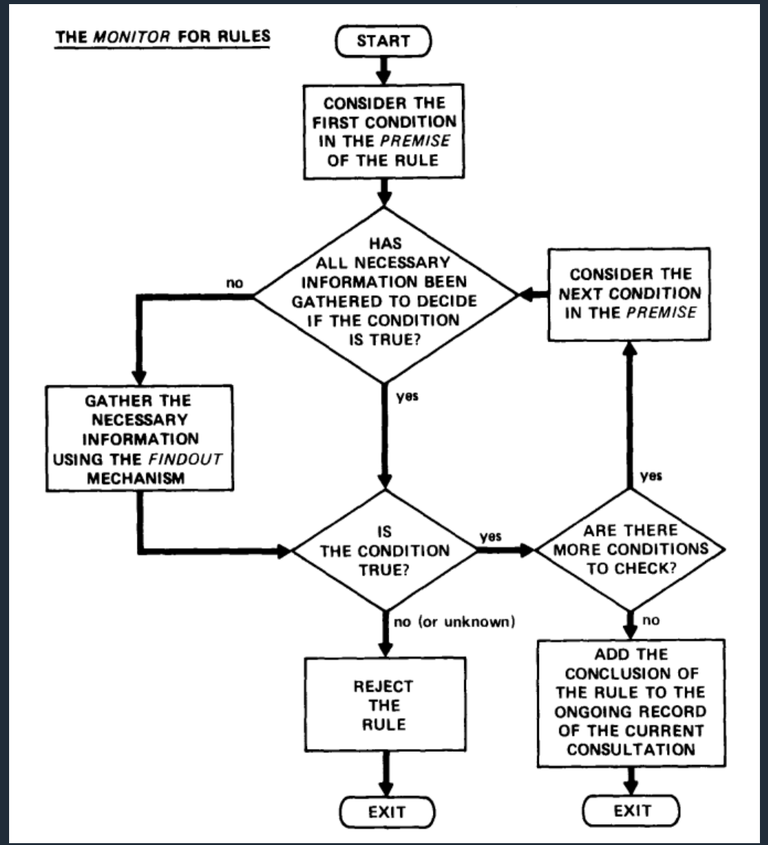

- one problem with algorithms: how to handle conflicting information?
- An expert system - MYCIN (Shortliffe and Buchanan 1975)
  - designed to identify bacterial infections and suitable Rx
  - 600 rules, supplied by experts
  - asks users a series of clinical questions
  - combines the answers using a (fairly simple) inference system
  - able to manage some conflicting information - unlike simpler algorithms

Machine learning

- A next step: can we provide learning rules to a system, and let it figure out the details for itself?

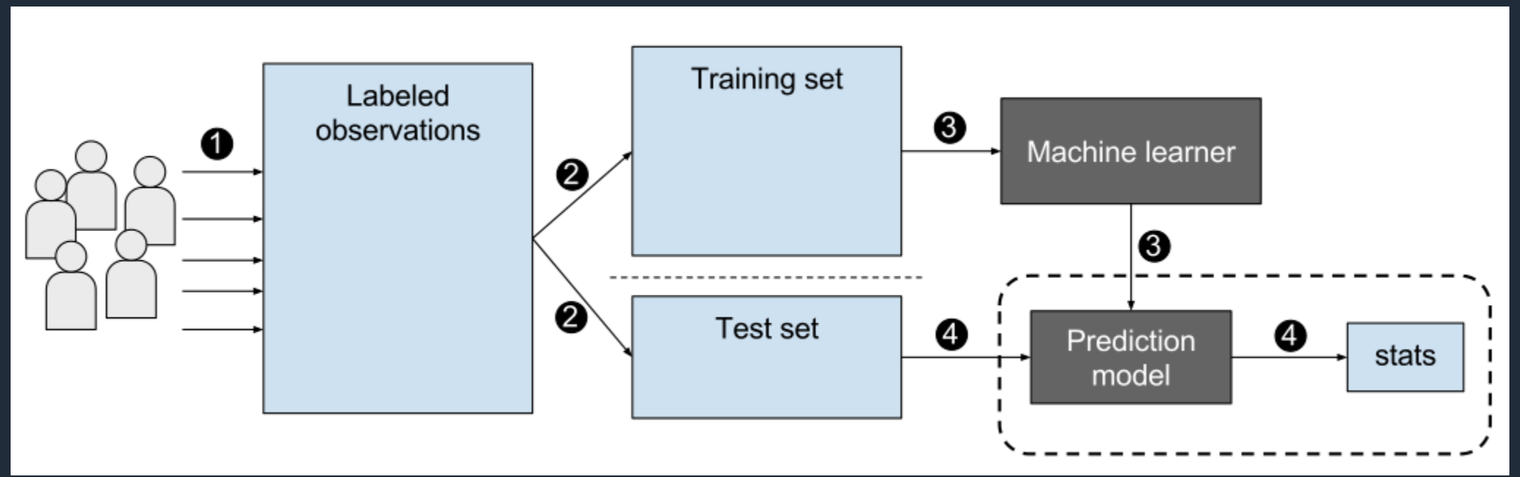

https://commons.wikimedia.org/wiki/File:Supervised_machine_learning_in_a_nutshell.svg

This is supervised learning

- supervision = labelled observations used for training and testing
- Lots of health examples with promising results:
  - diabetic retinopathy (Mookiah et al. 2013)
  - ECG (Aziz, Ahmed, and Alouini 2021)
  - fractures, melanoma, …

A dataset downside

KIND Learning Network
NHS Education for Scotland

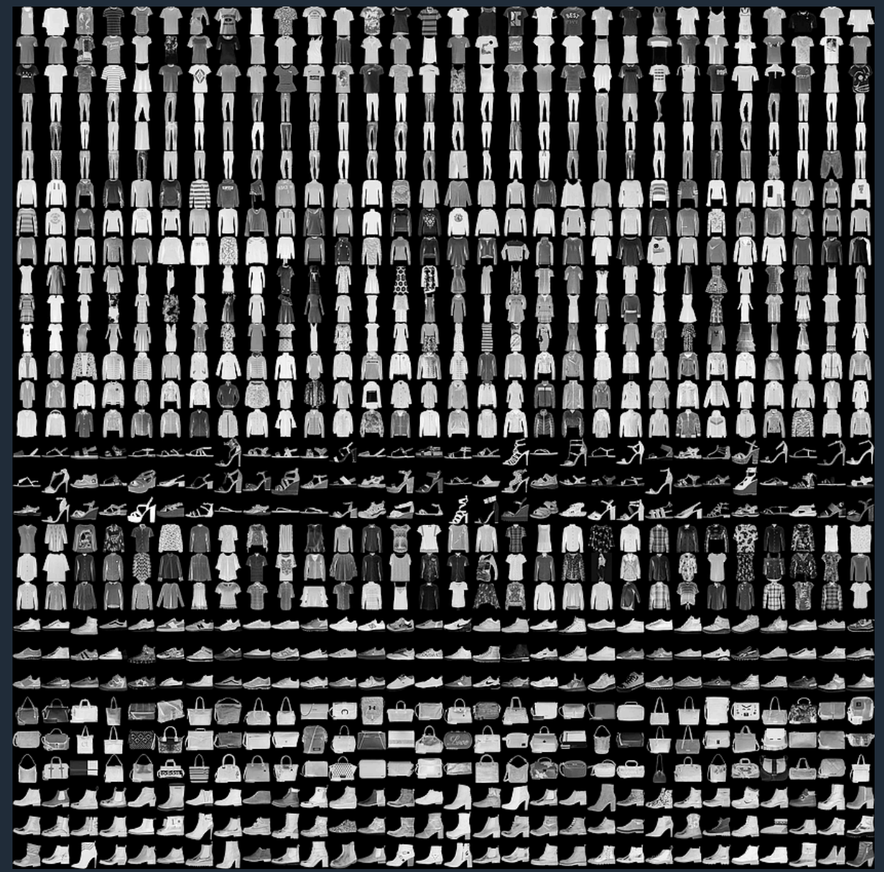

Fashion-MNIST dataset

Producing labelled datasets is hard:

- generally must be very large
- generally requires expert classification
- must be done with great accuracy
  - scale bar problem (Winkler et al. 2021)
- so dataset labelling is wildly expensive and thankless
  - Is there a way of doing something similar without spending millions classifying everything in the world by hand?

Unsupervised learning

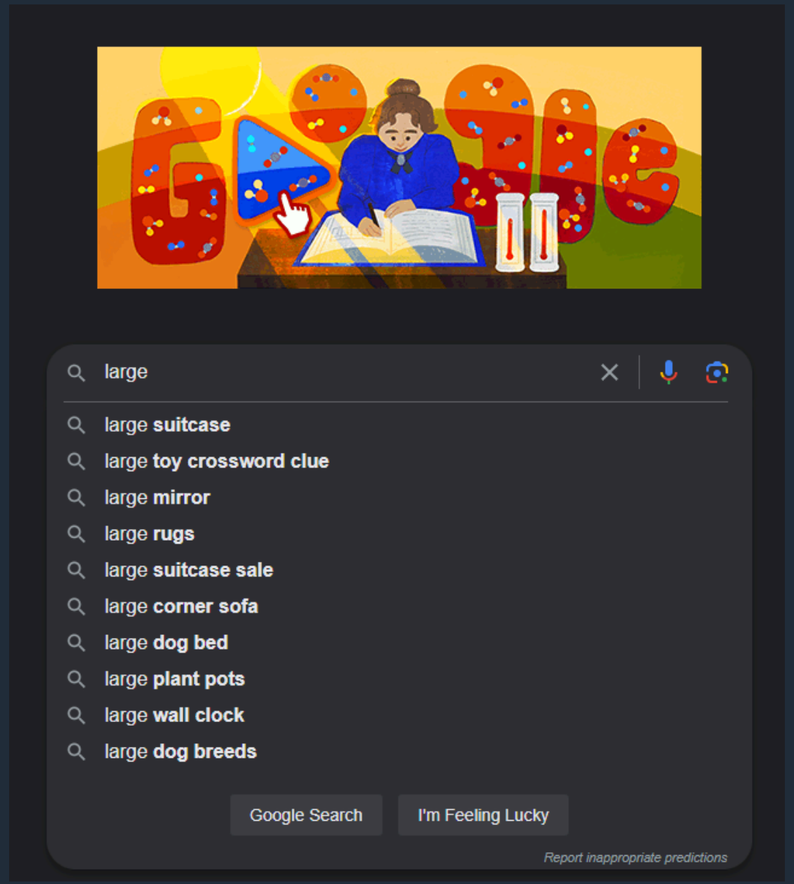

Unsupervised learning

Where autocomplete predictions come from

Autocomplete predictions reflect real searches that have been done on Google. To determine what predictions to show, our systems look for common queries that match what someone starts to enter into the search box but also consider:

- The language of the query
- The location a query is coming from
- Trending interest in a query
- Your past searches

These factors allow autocomplete to show the most helpful predictions that are unique to a particular location or time, such as for breaking news events.

In addition to full search predictions, Autocomplete may also predict individual words and phrases that are based on both real searches as well as word patterns found across the web.

Unsupervised learning

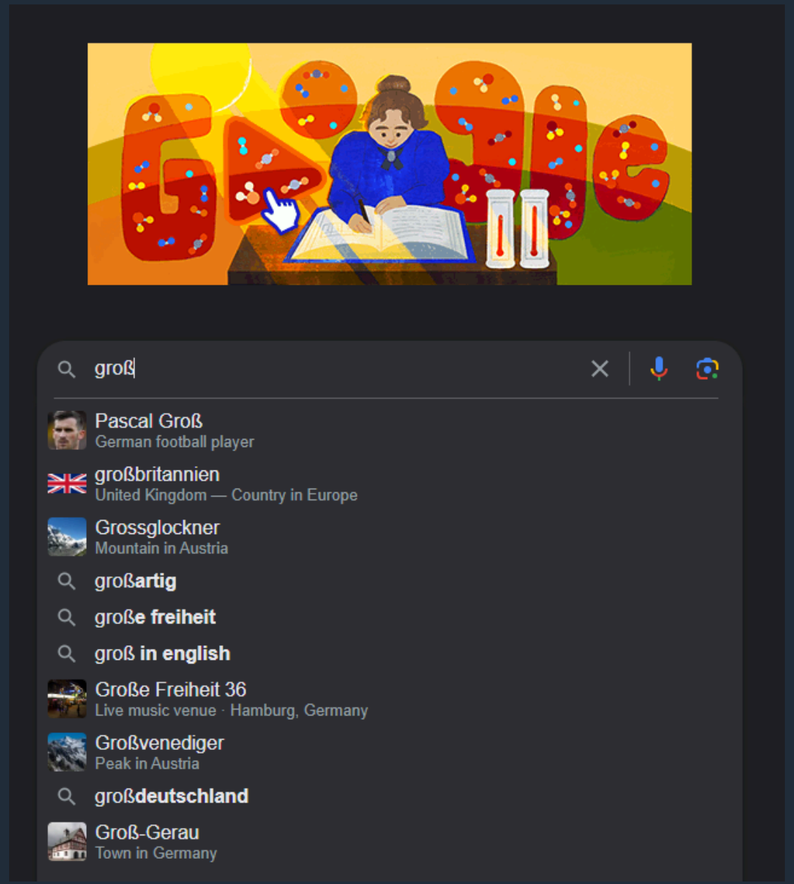

KIND Learning Network

NHS Education for Scotland

Unsupervised learning

KIND Learning Network
NHS Education for Scotland

- No-one is writing a list of possible searches starting with “Large…”
- Nor are they classifying searches into likely/unlikely, then training a model
- Instead, the model is looking at data (searches, language, location, trends) and calculating probabilities
  - 2011 blog post
  - 2020 PR piece
  - 2020 build your own in JS
- The terminology gets confusing again at this point:
  - some describe this as *deep learning*
  - better to call this a *language model*

Large language models

What if we were more ambitious with the scope of our language model?

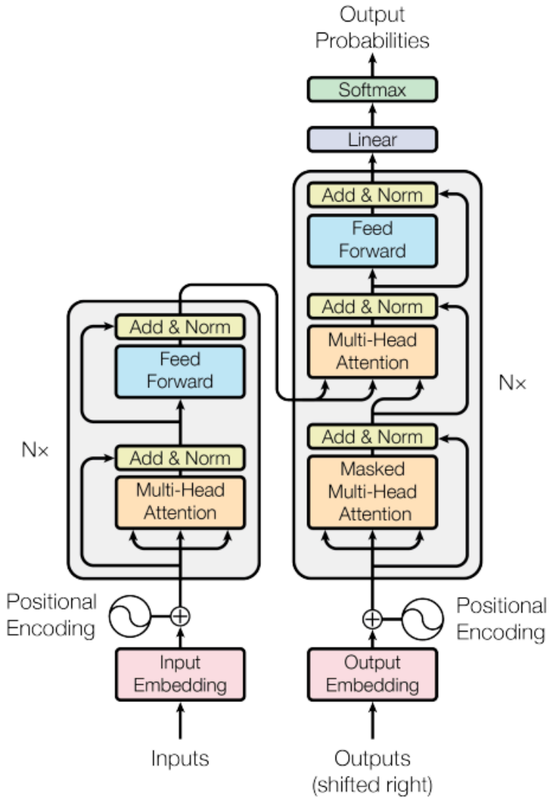

Figure 1: The Transformer - model architecture.

- Find masses of language data
  - chatGPT uses basically the whole web before September 2021
- Build a model capable of finding patterns in that data
  - Attention model used in chatGPT (Vaswani et al. 2017)
- Allow the model to calculate probabilities based on those patterns
  - lots of work going on at present allowing models to improve in response to feedback etc

Large language models

KIND Learning Network
NHS Education for Scotland

- superb at generating appropriate text, code, images, music…
- but production vs understanding
  - e.g. hallucinations, phantom functions…
- training is extremely computationally expensive
  - questions about inequality and regulatory moating
    - no-one but FAANG-sized companies can afford to do this
  - training is also surprisingly manual
- queries about ethics and attribution
  - your web content, my model, my paycheque
  - big serious worries about bias in some kinds of output

Punchline

- On balance, while there's hype here, there's also lots of substance and interest
- LLMs have become *much* better at producing plausible output, across a *greatly* expanded area
- A strength: fantastic ways for those with expertise to work faster
- A danger: LLMs are great at producing truth-like output. Good enough so that some will be tempted to use them to extend their apparent expertise…
- But big serious legal and ethical trouble ahead - we're not good at dealing with distributed responsibility

Feedback

Feedback link

Please give us one minute of your time. We add feedback comments to our training pages, because we think this is the most useful resource for people looking for specific training that suits their needs

KIND Learning Network
NHS Education for Scotland

Further reading

Aziz, Saira, Sajid Ahmed, and Mohamed-Slim Alouini. 2021. “ECG-Based Machine-Learning Algorithms for Heartbeat Classification.” *Scientific Reports* 11 (1). https://doi.org/10.1038/s41598-021-97118-5.

Mookiah, Muthu Rama Krishnan, U. Rajendra Acharya, Chua Kuang Chua, Choo Min Lim, E. Y. K. Ng, and Augustinus Laude. 2013. “Computer-Aided Diagnosis of Diabetic Retinopathy: A Review.” *Computers in Biology and Medicine* 43 (12): 2136–55. https://doi.org/10.1016/j.compbiomed.2013.10.007.

Searle, John R. 1980. “Minds, Brains, and Programs.” *Behavioral and Brain Sciences* 3 (3): 417–24. https://doi.org/10.1017/s0140525x00005756.

Shortliffe, Edward H., and Bruce G. Buchanan. 1975. “A Model of Inexact Reasoning in Medicine.” *Mathematical Biosciences* 23 (3-4): 351–79. https://doi.org/10.1016/0025-5564(75)90047-4.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. “Attention Is All You Need.” https://doi.org/10.48550/ARXIV.1706.03762.

Winkler, Julia K., Katharina Sies, Christine Fink, Ferdinand Toberer, Alexander Enk, Mohamed S. Abassi, Tobias Fuchs, and Holger A. Haenssle. 2021. “Association Between Different Scale Bars in Dermoscopic Images and Diagnostic Performance of a Market-Approved Deep Learning Convolutional Neural Network for Melanoma Recognition.” *European Journal of Cancer* 145 (March): 146–54. https://doi.org/10.1016/j.ejca.2020.12.010.