Artificial Intelligence for Medicine I

Fall 2024

Lecture 1: Introduction

(Many slides adapted from CS 188, BIODS220 and other resources on the web)

Welcome
Artificial Intelligence for Medicine I

- This course is Artificial Intelligence for Medicine I
- What you will learn:
  1. Understand the basic principles of AI and machine learning.
  2. Program in Python
  3. Use Python libraries for AI model development.
  4. Develop skills for AI model development for a given medical field.
  5. Analyze the ethical, legal, and social implications of AI in healthcare

Tentative Course Outline for AI for Medicine I

WEEK	TOPIC
Week 1	What is AI? History and evolution of AI in medicine
Week 2	Programming in Python - Elementary Programming
Week 3	Programming in Python - Mathematical Functions, Strings, and Objects
Week 4	Programming in Python - Decision Statements
Week 5	Programming in Python - Loops
Week 6	Programming in Python - Functions
Week 7	Programming in Python - Multidimensional Lists, Tuples, Sets, and Dictionaries
Week 8	Programming in Python - Data Libraries/Structures for AI
Week 9	Introduction to Basic AI Models and Process for Medicine
Week 10	Data Storage, Loading, and Prepocessing/Wrangling
Week 11	Exploratory Data Analysis and Data Visualization
Week 12	Model Development
Week 13	Model Evaluation and Refinement
Week 14	Ethical, Legal, and Social Implications of AI in Medicine

Artificial Intelligence for Medicine II

- Next semester, we have a continuation of this course: Artificial Intelligence for Medicine II
- What you will learn in AI for Medicine II:
  1. Understand the fundamental concepts of AI and machine learning.
  2. Analyze the role of AI in various medical fields, including diagnostics, imaging, personalized medicine, and drug discovery.
  3. Learn and apply AI models to solve specific problems in medicine.
  4. Gain hands-on experience with AI tools and platforms through practical exercises and projects.
  5. Stay informed about the latest advancements, research studies, and trends in AI and healthcare.

Tentative Course Outline for AI for Medicine II

WEEK	TOPIC
Week 1	Introduction to AI Methods and their Applications in Medicine
Week 2	Machine Learning Basics
Week 3	Data Collection and Preprocessing
Week 4	Supervised Learning
Week 5	Unsupervised Learning
Week 6	Model Evaluation and Performance Metrics
Week 7	Deep Learning in Medicine
Week 8	Medical Imaging and AI
Week 9	Natural Language Processing (NLP) in Healthcare
Week 10	AI in Diagnostics and Disease Prediction
Week 11	AI in Personalized Medicine, Treatment Planning, Drug Discovery
Week 12	AI in Medical Robotics and Genomics
Week 13	Challenges and Limitations of AI in Medicine, and Future Trends
Week 14	Course Review and Project Presentations

Sci-Fi AI?

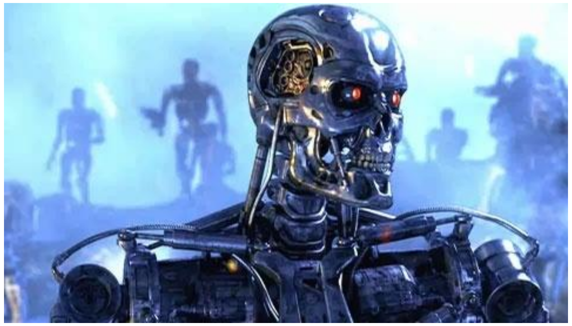

Machines that Can Speak (cont.)

- C3PO
  in Star Wars

- KITT
  in Knight Rider

But Still a Sci-Fi…

- "However, as 2001 approached it became clear that *2001*'s predictions in computer technology were far fetched. Natural language, lip reading, planning and plain common sense in computers were still the stuff of science fiction."
  HAL 9000 – Wikipedia

One who does not dream big cannot achieve big things.

Artificial general intelligence (AGI)

- *Today, we talk about* **Artificial general intelligence** (**AGI**)
- **Artificial general intelligence** (**AGI**) is a type of artificial intelligence (AI) that matches or surpasses human cognitive capabilities across a wide range of cognitive tasks. This contrasts with narrow AI, which is limited to specific tasks.[1][2] AGI is considered one of the definitions of strong AI.[3]
- **Artificial general intelligence (AGI)** is a field of theoretical AI research that attempts to create software with human-like intelligence and the ability to self-teach. The aim is for the software to be able to perform tasks that it is not necessarily trained or developed for.

- **AGI** may be comparable to, match, differ from, or even appear alien-like relative to human intelligence, encompassing a spectrum of possible cognitive architectures and capabilities that includes the spectrum of human-level intelligence.[4][5][6]

- Creating AGI is a primary goal of AI research and of companies such as OpenAI[7] and Meta.[8] A 2020 survey identified 72 active AGI R&D projects spread across 37 countries.[9]

https://en.wikipedia.org/wiki/Artificial_general_intelligence

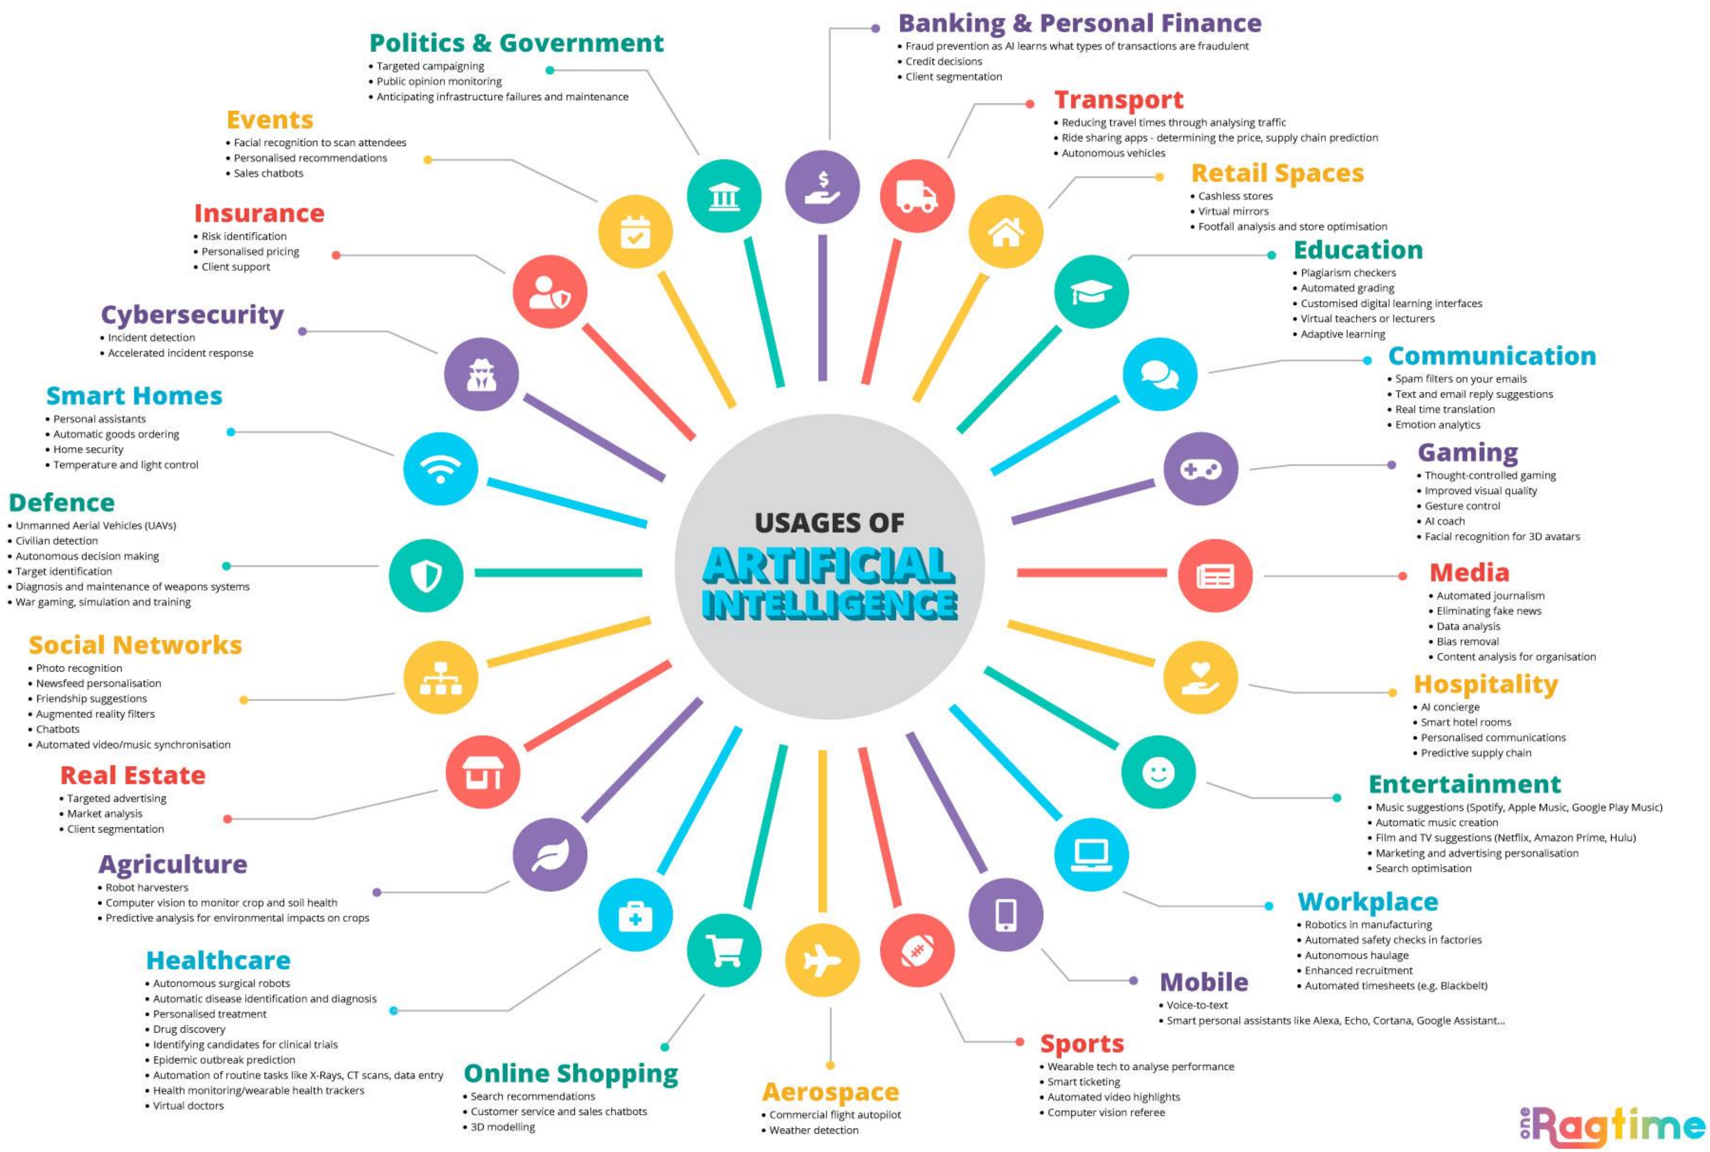

https://img.oneragtime.com/uploads/2018/08/infographic_wheel_v2.jpg

Large Language Models

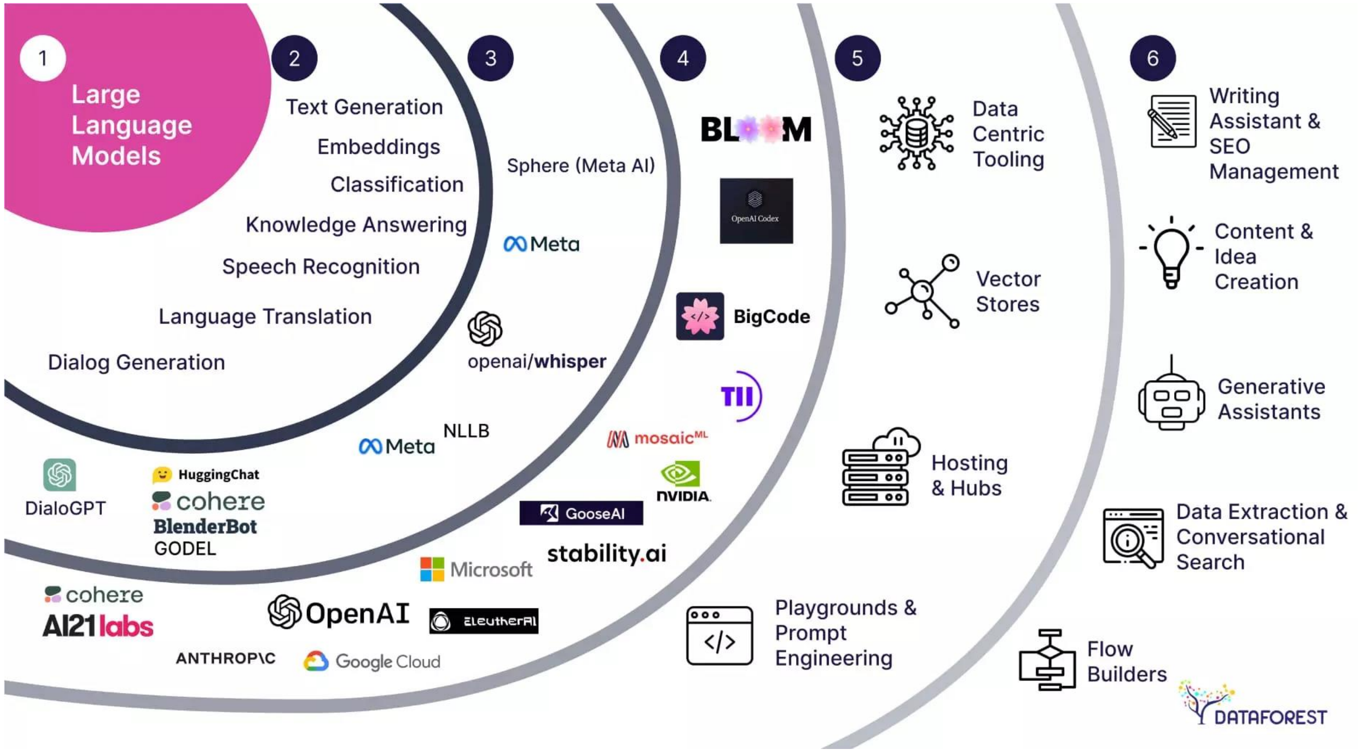

1 — Available Large Language Models

2 — General Use-Cases

3 — Specific Implementations

4 — Models

5 — Foundation Tooling

6 — End User UIs

https://dataforest.ai/blog/large-language-models-advanced-communication

The Major History of NLP / LLM

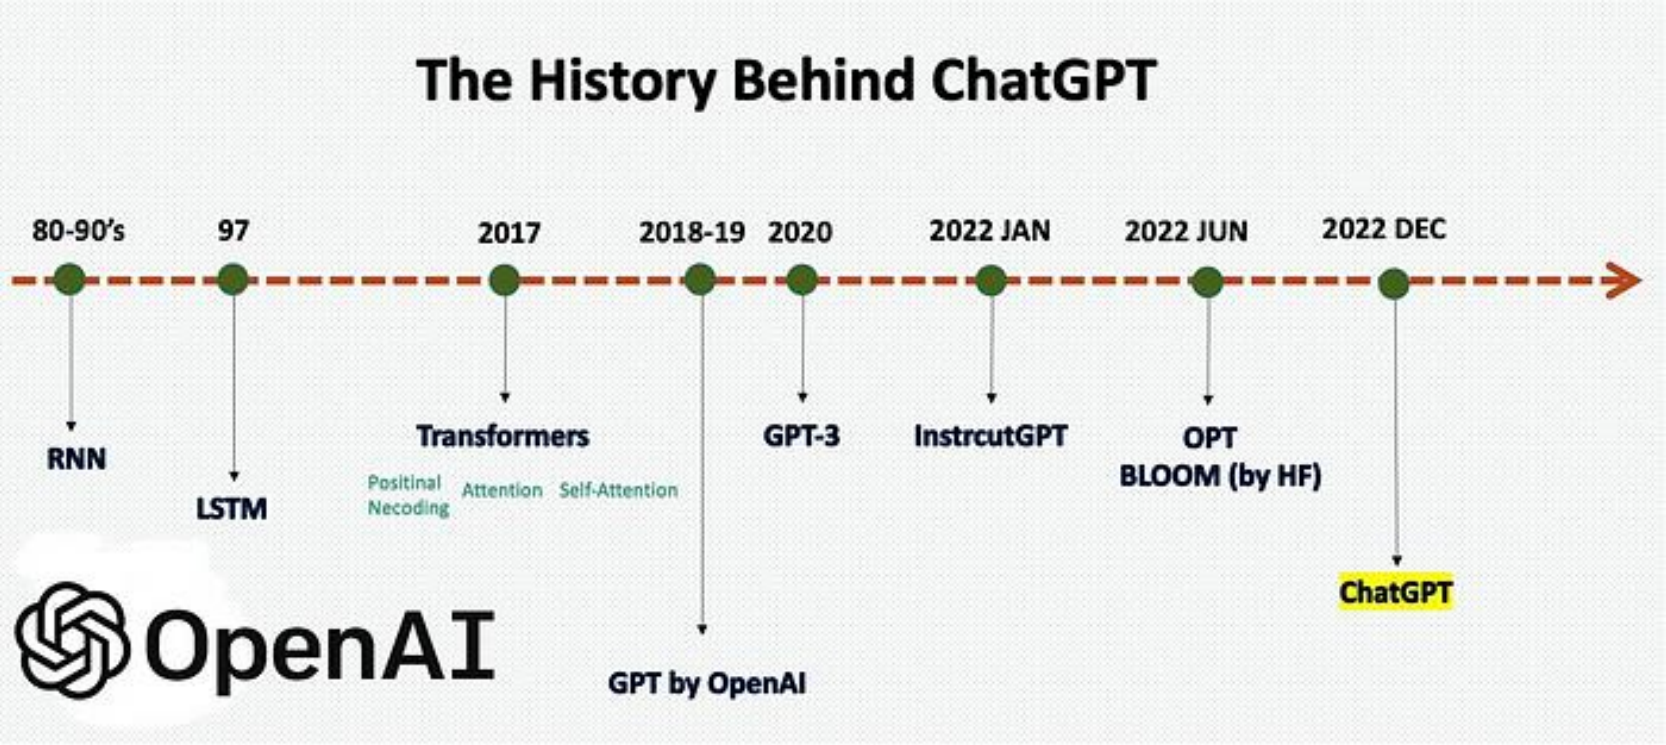

https://medium.com/the-ai-analytics-corner/demystifying-chatgpt-part-01-the-history-of-llm-nlp-53d6084758e8

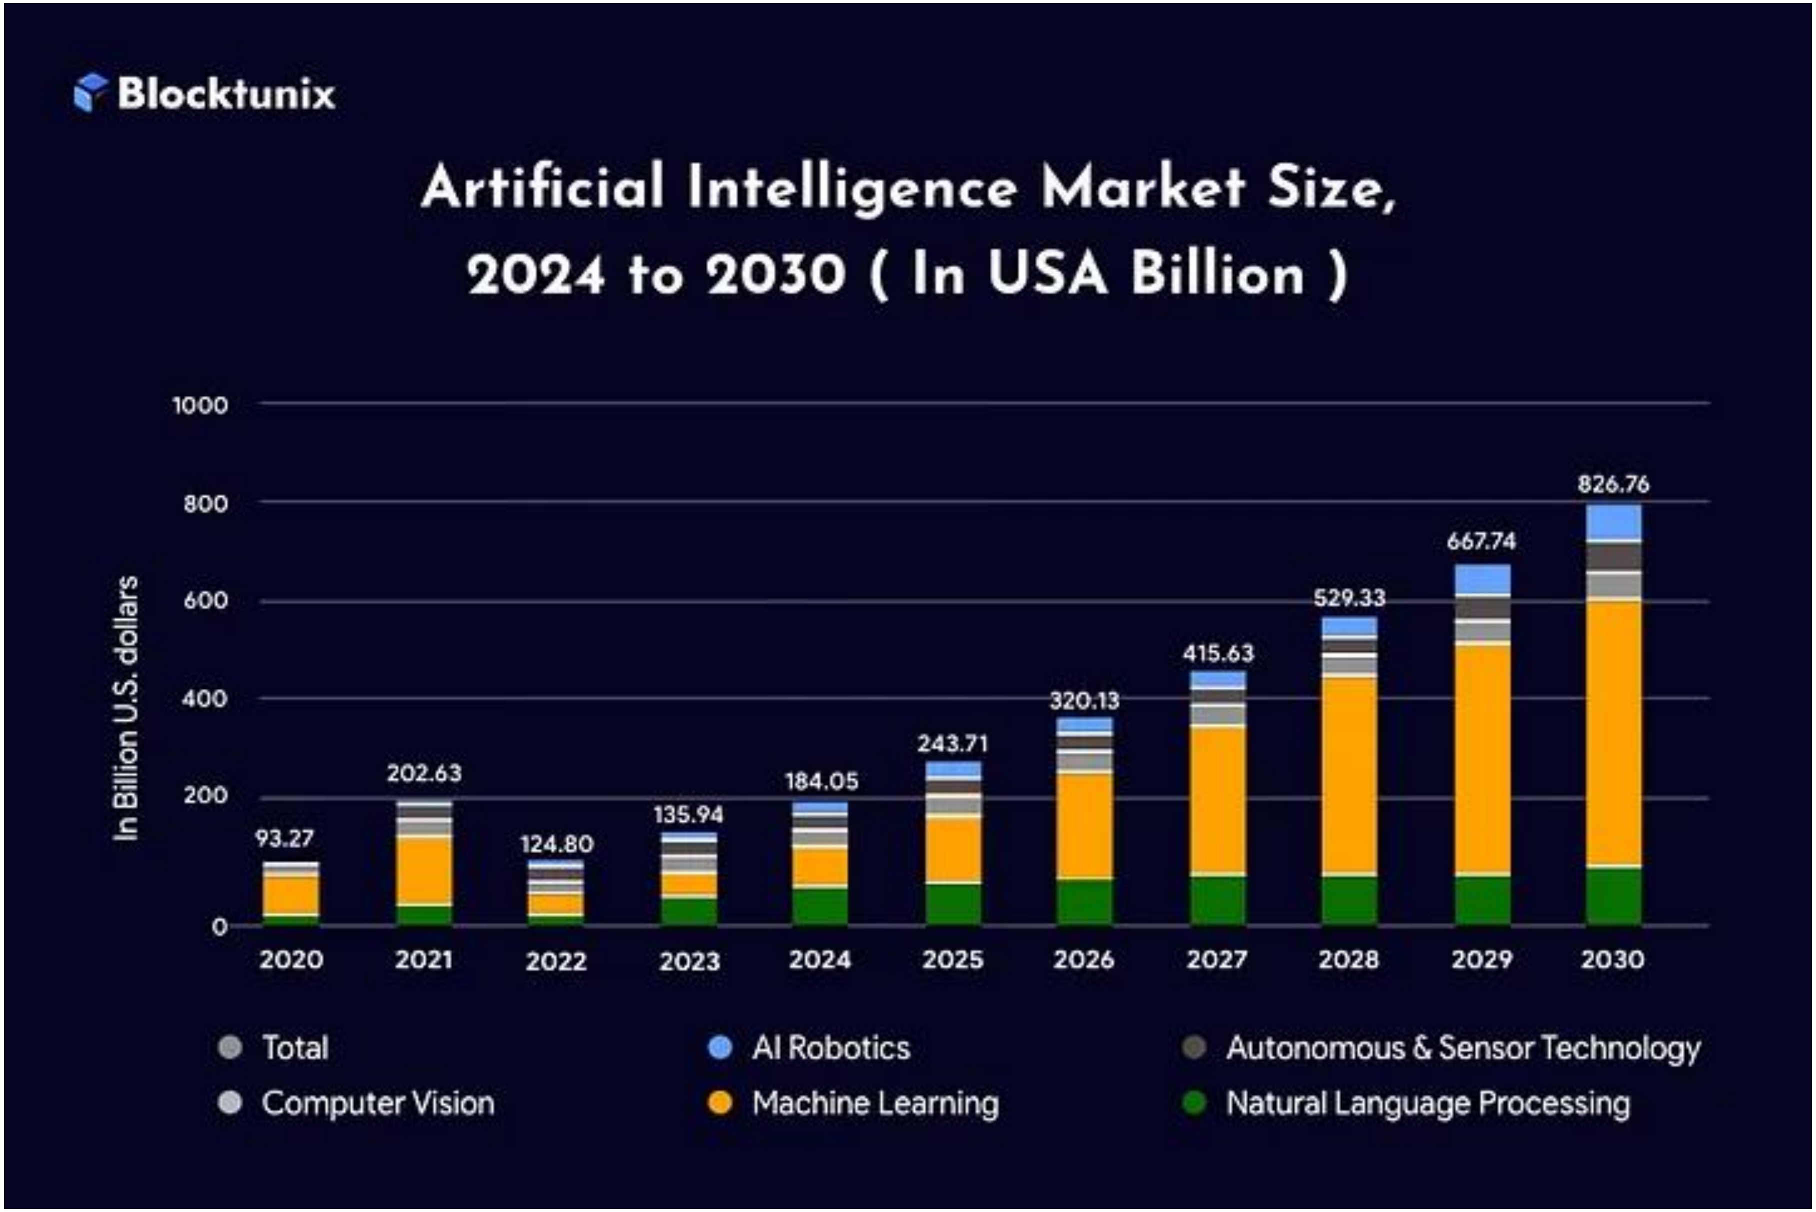

https://medium.com/@Blocktunix/key-technologies-in-ai-app-development-9a796826d107

AI in medicine and healthcare: a rapidly exploding field

Hospitals Roll Out AI Systems to Keep Patients From Dying of Sepsis

Septic shock kills 50 percent of people who are affected—Sepsis Watch could save their lives

By **Eliza Strickland**

In hospitals, doctors and nurses keep vigilant watch over patients'

Google, Verily using AI to screen for diabetic retinopathy in India

The machine learning algorithm can also help with screening for diabetic macular edema, a boon for patients in a country where physicians are in short supply.

By **Mike Miliard** | February 26, 2019 | 03:17 PM

(Photo: Google)

Source: https://biods220.stanford.edu/

AI in medicine and healthcare: a rapidly exploding field

Apple's future healthcare market moves will rely heavily on AI analysis

By Malcolm Owen
Monday, September 16, 2019, 09:03 am PT (12:03 pm ET)

Apple's moves in the healthcare market could involve the tracking of user data for further analysis by artificial intelligence and billing model based on cost-savings, with analysts pointing out areas of the consumer health industry Apple could easily advance by building upon its already-released technology and services.

SCIENCE \ BUSINESS \ TECH

Amazon is buying 'membership-based' healthcare provider One Medical for $3.9 billion

19

One Medical's Netflix-for-primary-care is a $199 subscription to a modern doctor's office

By Richard Lawler | @rjcc | Jul 21, 2022, 9:40am EDT | 19 comments

SHARE

Google to Store and Analyze Millions of Health Records

The tech company's deal with Ascension is part of a push to use artificial intelligence to aid health services.

MICROSOFT \ SCIENCE \ TECH

Microsoft Healthcare is a new effort to push doctors to the cloud

Microsoft wants to be a big part of the cloud and AI healthcare race

By Tom Warren | @tomwarren | Jun 27, 2018, 6:50am EDT

Source: https://biods220.stanford.edu/

AI in medicine and healthcare: a rapidly exploding field

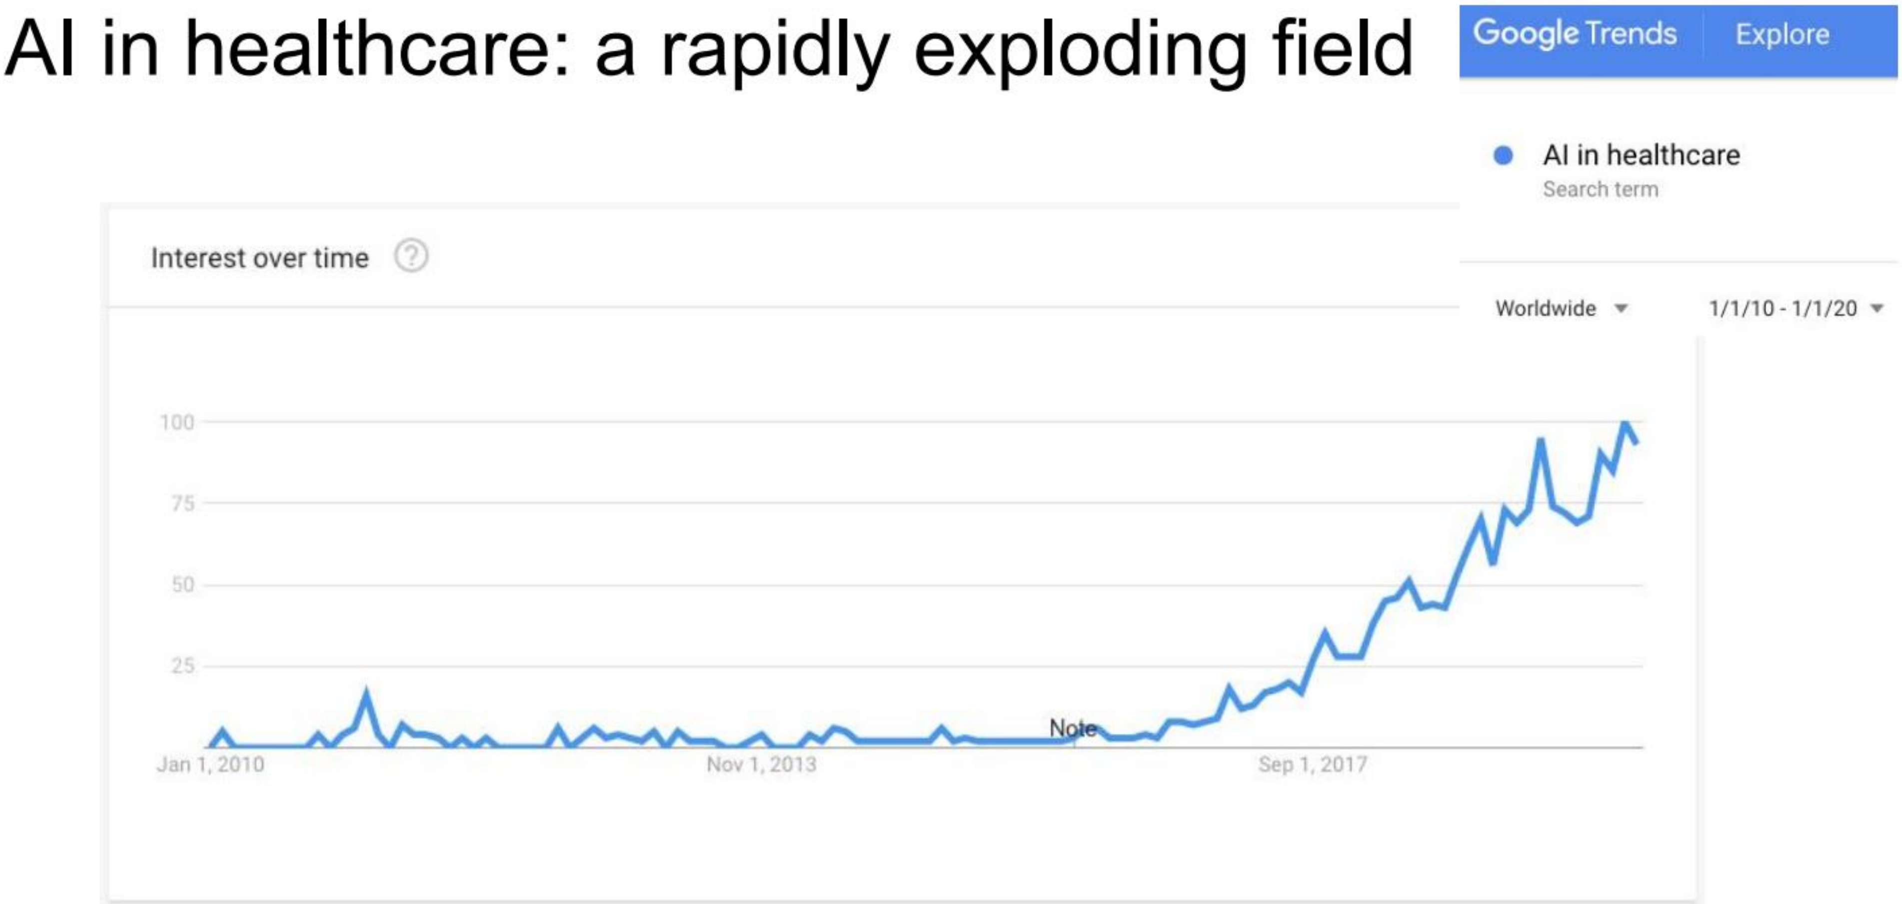

Source: https://biods220.stanford.edu/

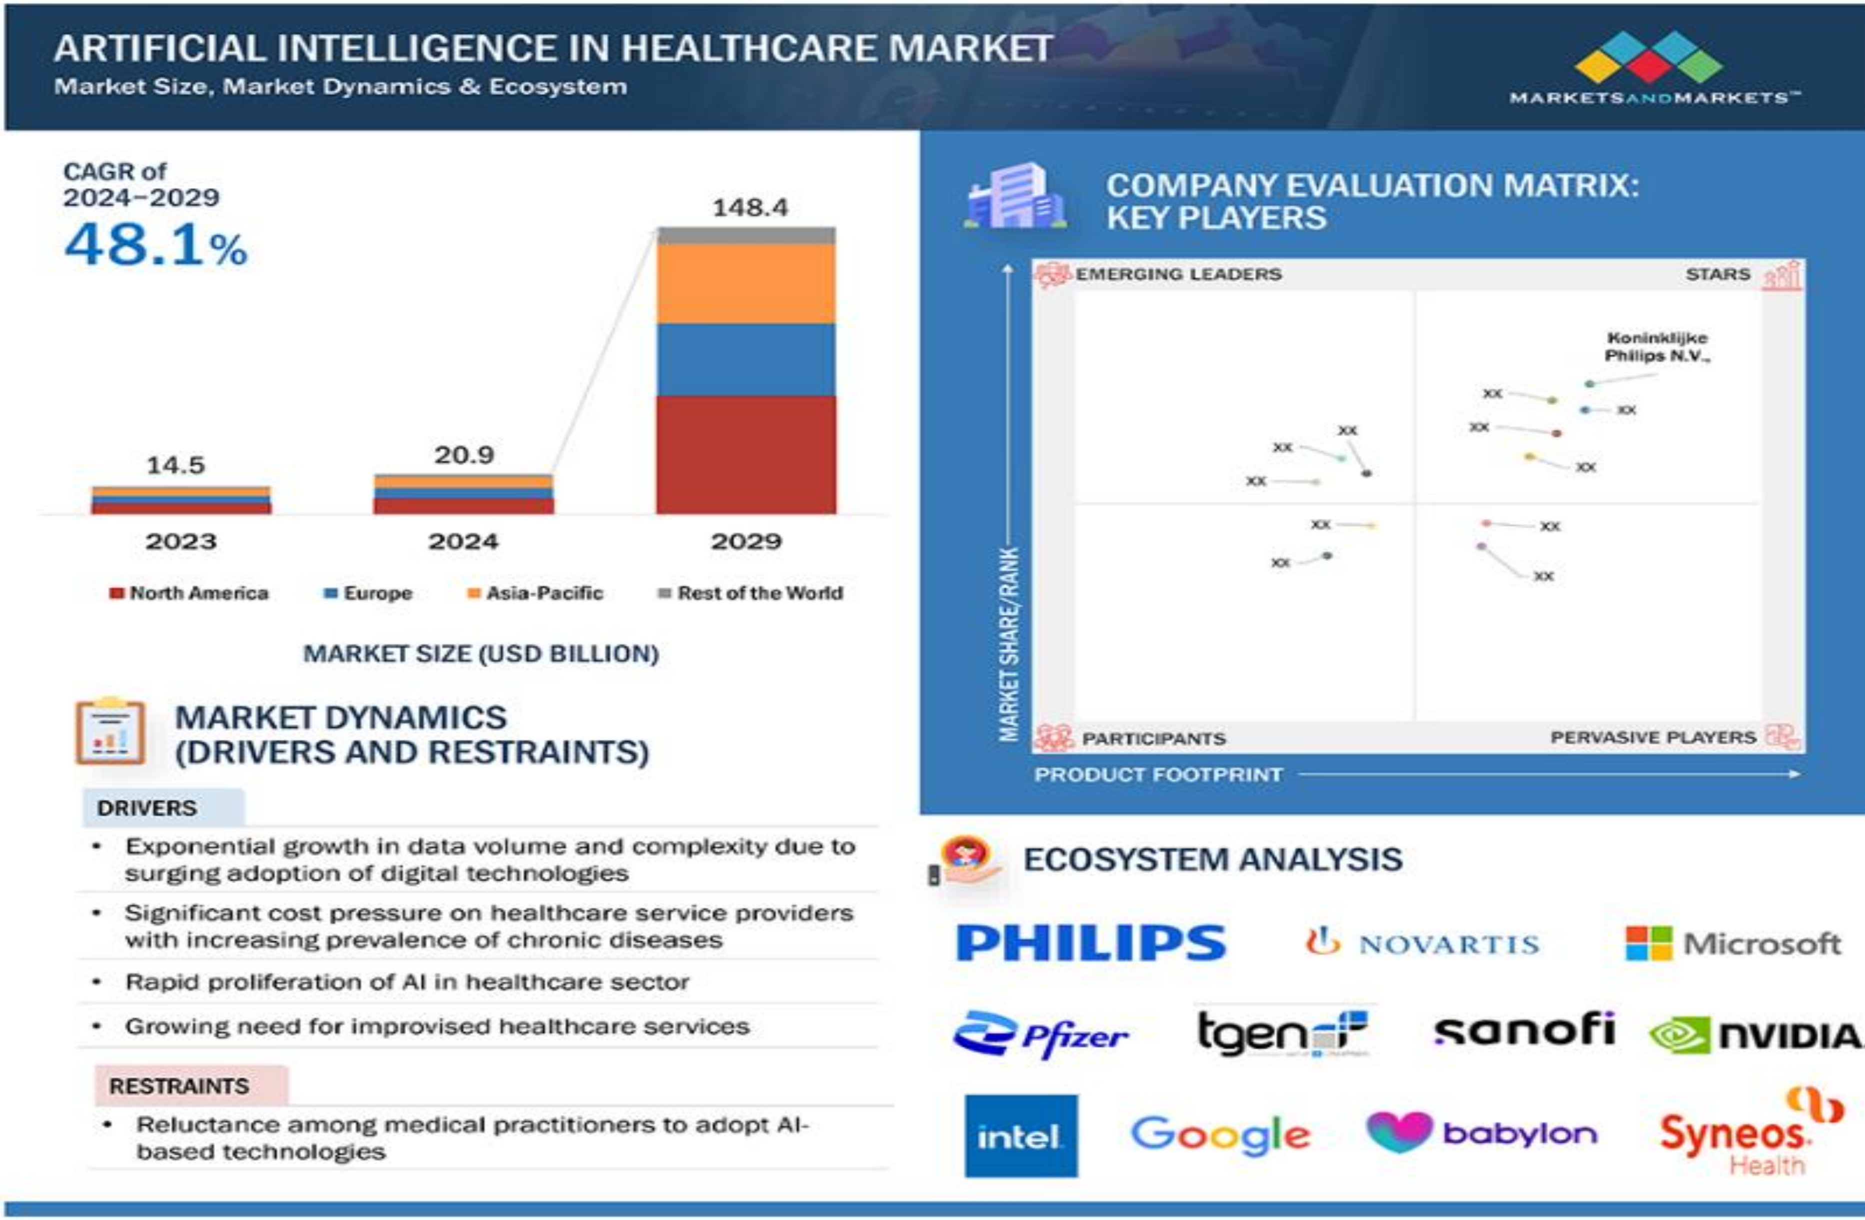

Artificial Intelligence (AI) in Healthcare Market Statistics Forecast to 2029

Artificial Intelligence (AI) in Healthcare Market - Global Forecast to 2029

- The global AI in Healthcare market size was valued at USD 20.9 billion in 2= 024 and is estimated to reach **USD 148.4 billion by 2029.**
- The growth of AI in the healthcare market is driven by the generation of large and complex healthcare datasets, the pressing need to reduce healthcare costs, improving computing power and declining hardware costs, and the rising number of partnerships and collaborations among different domains in the healthcare sector, and growing need for improvised healthcare services due to imbalance between healthcare workforce and patients.

Source: https://www.marketsandmarkets.com/Market-Reports/artificial-intelligence-healthcare-market-54679303.html?gad_source=1&gclid=Cj0KCQjwu-63BhC9ARIsAMMTLXRjxj3SvMxw2q9uyj3_cuUfpOnwWW_n8J9InV4EiVP45UcAOnBtxG8aAhhcEALw_wcB

ARTIFICIAL INTELLIGENCE IN HEALTHCARE MARKET

Market Segmentation & Geographical Spread

MARKETSANDMARKETS™

BY TECHNOLOGY

- Machine Learning
  - Deep Learning
  - Supervised Learning
  - Unsupervised Learning
  - Reinforcement Learning
  - Others
- Natural Language Processing
  - IVR
  - OCR
  - Pattern and Image Recognition
  - Auto Coding
  - Classification and Categorization
  - Text Analytics
  - Speech Analytics
- Context-aware Computing
  - Device Context
  - User Context
  - Physical Context
- Computer Vision

BY APPLICATION

- Patient Data & Risk Analysis
- Medical Imaging & Diagnostics
- Precision Medicines
- Drug Discovery
- Lifestyle Management & Remote Patient Monitoring
- Virtual Assistants
- Wearables
- Inpatient Care & Hospital Management
- Research
- Emergency Rooms & Surgeries
- Mental Health
- Healthcare Assistance Robots

BY OFFERING

- Hardware
  - Processor
  - Memory
  - Network
- Software
  - AI Platform
  - AI Solution
- Services
  - Deployment & Integration
  - Support & Maintenance

BY END USER

- Healthcare Payers
- Hospitals & Healthcare Providers
- Pharmaceuticals & Biotechnology Companies
- Other End Users

Large language models encode clinical knowledge

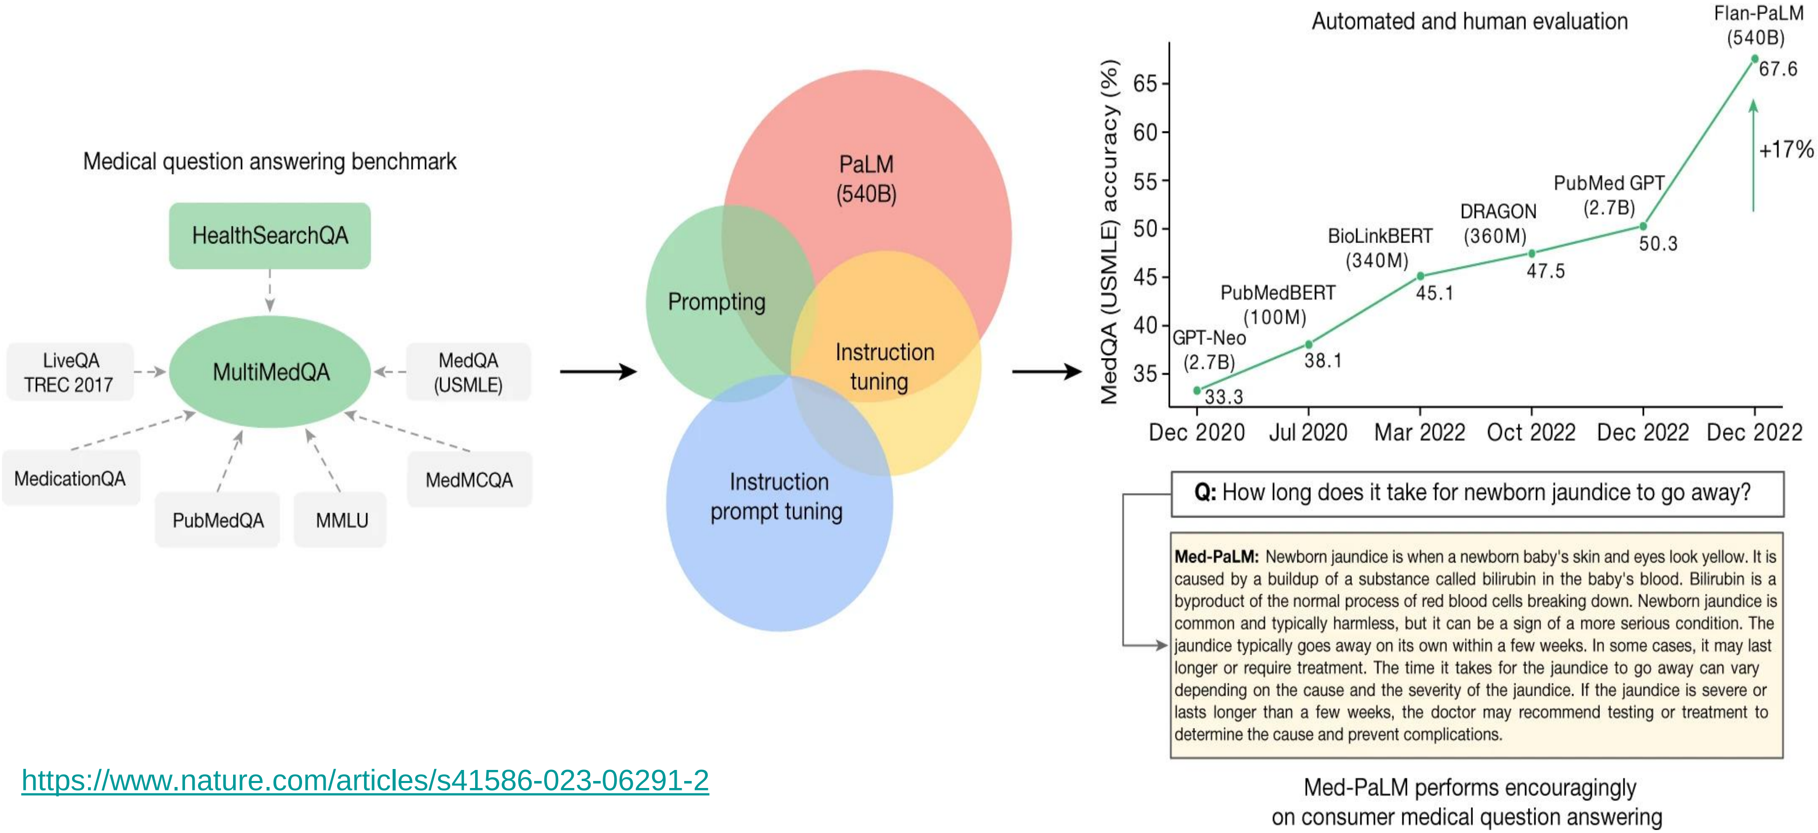

https://www.nature.com/articles/s41586-023-06291-2

We curate MultiMedQA, a benchmark for answering medical questions spanning medical exam, medical research and consumer medical questions. We evaluate PaLM and its instructed-tuned variant, Flan-PaLM, on MultiMedQA. Using a combination of prompting strategies, Flan-PaLM exceeds state-of-the-art performance on MedQA (US Medical Licensing Examination (USMLE)), MedMCQA, PubMedQA and MMLU clinical topics. In particular, it improves over the previous state of the art on MedQA (USMLE) by over 17%. We next propose instruction prompt tuning to further align Flan-PaLM to the medical domain, producing Med-PaLM. Med-PaLM's answers to consumer medical questions compare favourably with answers given by clinicians under our human evaluation framework, demonstrating the effectiveness of instruction prompt tuning.

A (Short) History of AI

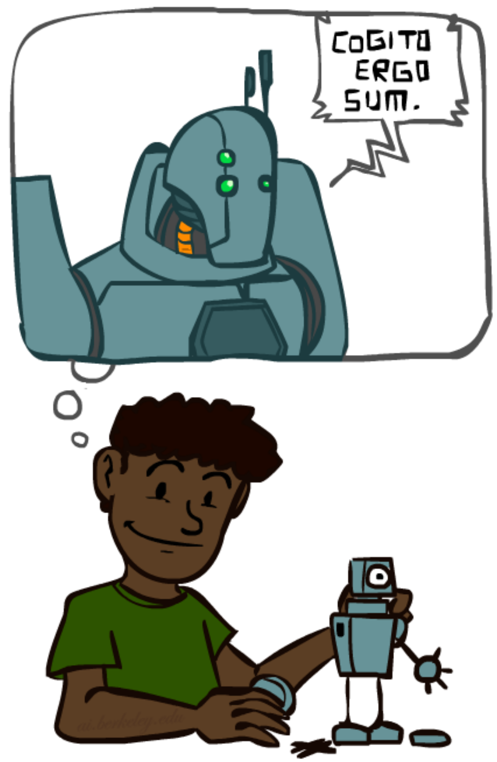

Pre-AI developments

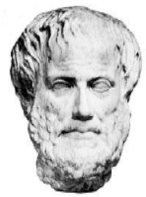

Philosophy: **intelligence** can be achieved via mechanical computation (e.g., Aristotle)

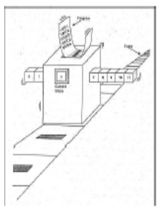

Church-Turing thesis (1930s): any computable function is **computable** by a Turing machine

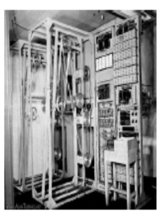

Real computers (1940s): actual **hardware** to do it: Heath Robinson, Z-3, ABC/ENIAC

- While AI is a relatively young field, one can trace back some of its roots back to Aristotle, who formulated a system of syllogisms that capture the reasoning process: how one can mechanically apply syllogisms to derive new conclusions.
- Alan Turing, who laid the conceptual foundations of computer science, developed the Turing machine, an abstract model of computation, which, based on the Church-Turing thesis, can implement any computable function.
- In the 1940s, devices that could actually carry out these computations started emerging.
- So perhaps one might be able to capture intelligent behavior via a computer. But how do we define success?

Syllogism, in logic, a valid deductive argument having two premises and a conclusion

Test for Intelligence – Turing Test

- Alan Turing (1950) proposed a test of a machine's capability to perform human-like conversation.

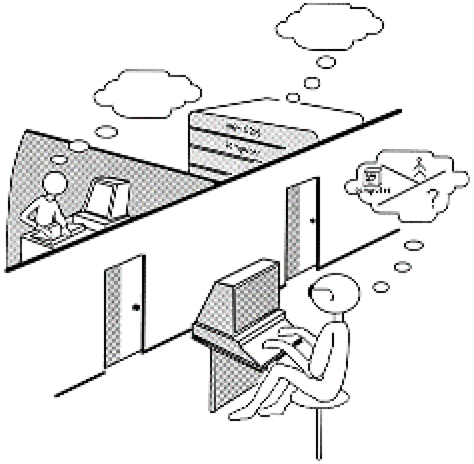

  - A human judge engages in a natural language conversation with two other parties, one a human and the other a machine; if the judge cannot reliably tell which is which, then the machine is said to pass the test.

Birth of AI (1956)

- AI's official birth: Dartmouth conference , 1956

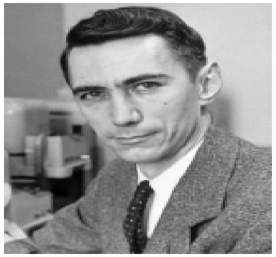

"An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. ***We think that a significant advance can be made if we work on it together for a summer."***

John McCarthy and Claude Shannon
Dartmouth Workshop Proposal

The term "Artificial Intelligence" adopted

A (Short) History of AI

- 1940-1950: Early days
  - 1943: McCulloch & Pitts: Boolean circuit model of brain
  - 1950: Turing's "Computing Machinery and Intelligence"
- 1950—70: Excitement: Look, Ma, no hands!
  - 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
  - 1956: Dartmouth meeting: "Artificial Intelligence" adopted
  - 1965: Robinson's complete algorithm for logical reasoning
- 1970—90: Knowledge-based approaches
  - 1969—79: Early development of knowledge-based systems
  - 1980—88: Expert systems industry booms
  - 1988—93: Expert systems industry busts: "AI Winter"
- 1990—: Statistical approaches
  - Resurgence of probability, focus on uncertainty
  - General increase in technical depth
  - Agents and learning systems… "AI Spring"?
- 2000—: Where are we now?

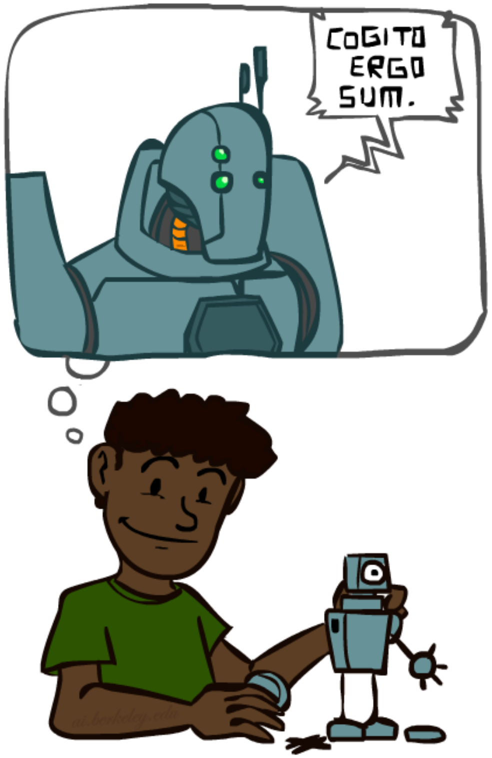

Early progress in the late 50s and 60s

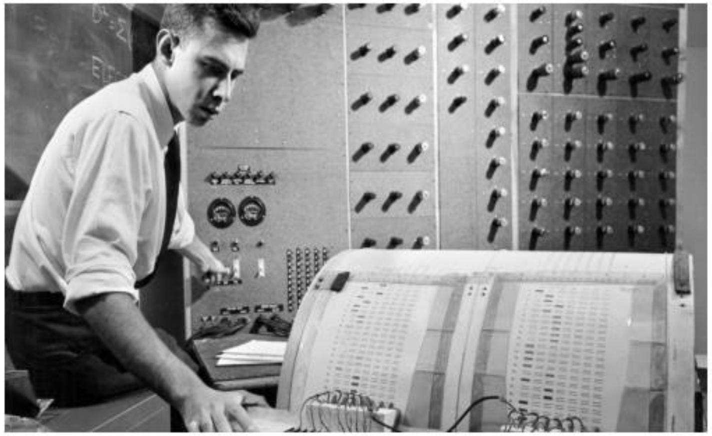

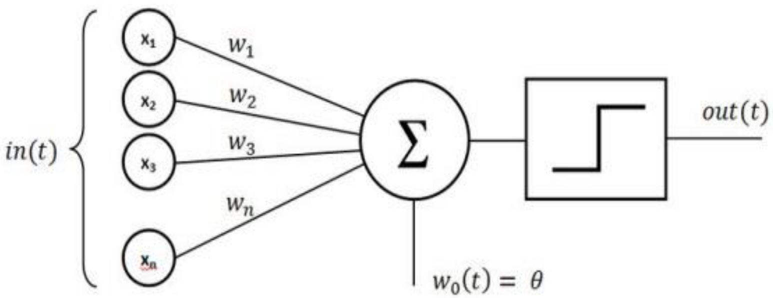

Perceptron model: Rosenblatt, 1958

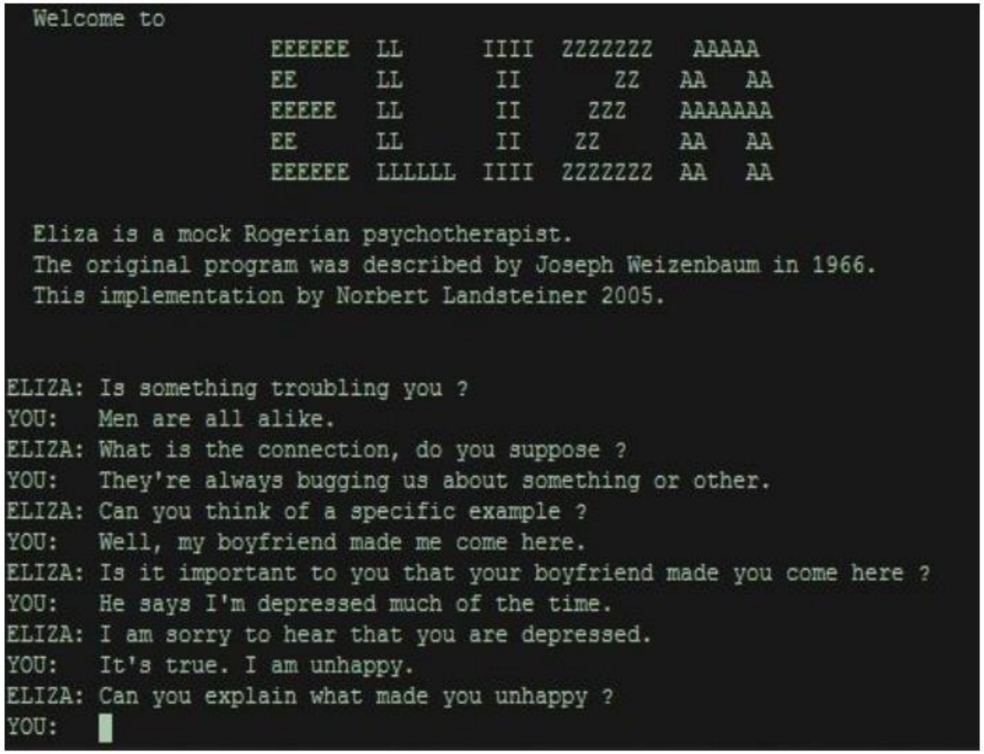

ELIZA chatbot: Weizenbaum, 1966

1970—90: Knowledge-based approaches

#DEFINE MOLFORM C 12 H 14 O
MOLECULAR FORMULA DEFINED

#DEFINE SUBSTRUCTURE Z

[Z is the structure required by constraints C2 and C3.]

...

CONSTRAINT: **SUBSTRUCTURE CH0 EXACTLY 2**
[C7: we must end up with exactly two quaternary carbons.]
CONSTRAINT: **RING 3 NONE** [C9]
CONSTRAINT: **RING 4 NONE** [C9]
CONSTRAINT:

#.#.#.#..#..#..#..#..#..#..#..#.#.#.#..#.#..

#.#..#.#..#..#...#..#..#.#.#.#..

47 STRUCTURES WERE OBTAINED

#DRAW ATNAMED (5 6)

[Comment: The following is a selection of final structures 5, 6.]

#5:

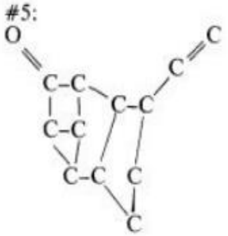

Expert systems, 1970s and 80s. Feigenbaum, etc.

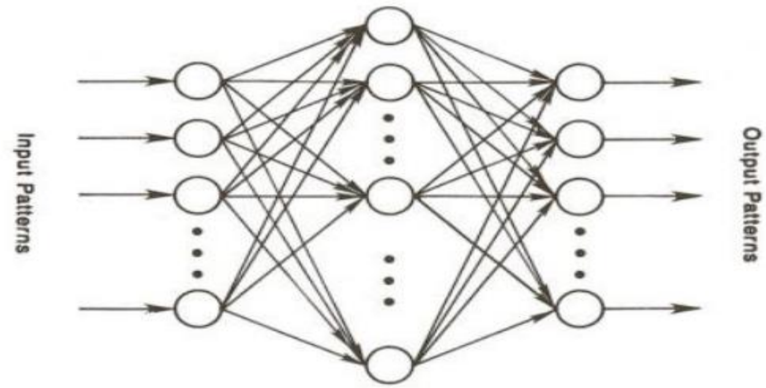

To get the correct generalization of the delta rule, we must set

$$\Delta_p w_{ji} \propto -\frac{\partial E_p}{\partial w_{ji}},$$

where E is the same sum-squared error function defined earlier. As in the standard delta rule it is again useful to see this derivative as resulting from the product of two parts: one part reflecting the change in error as a function of the change in the net input to the unit and one part representing the effect of changing a particular weight on the net input. Thus we can write

$$\frac{\partial E_p}{\partial w_{ji}} = \frac{\partial E_p}{\partial net_{pj}} \frac{\partial net_{pj}}{\partial w_{ji}}. \quad (9)$$

By Equation 7 we see that the second factor is

$$\frac{\partial net_{pj}}{\partial w_{ji}} = \frac{\partial}{\partial w_{ji}} \sum_k w_{jk} o_{pk} = o_{pi}. \quad (10)$$

Backpropagation. Rumelhart, 1986.

First appearances of modern neural networks

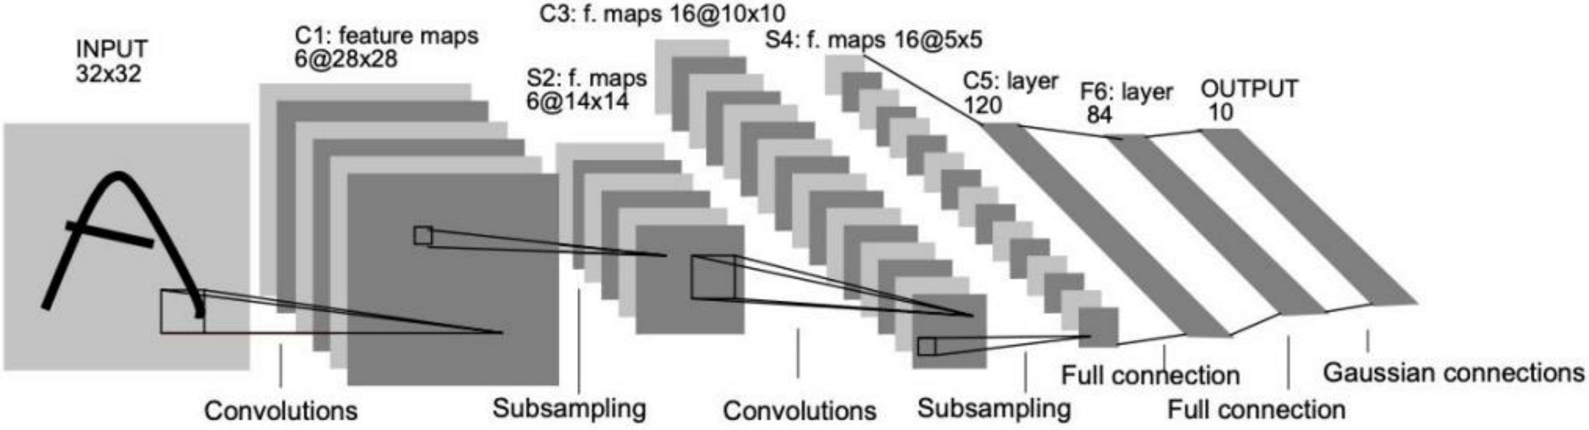

LeCun, 1990s.

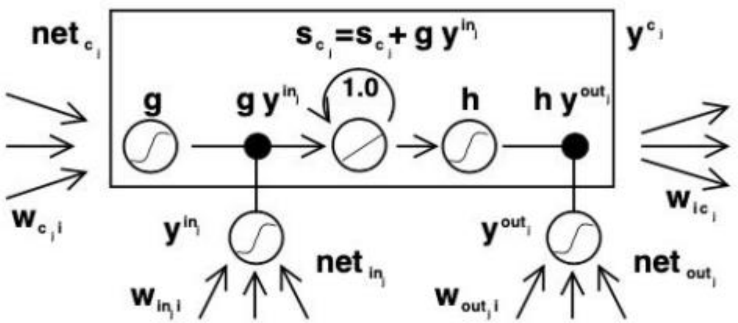

Figure 1: *Architecture of memory cell c_j (the box) and its gate units in_j, out_j. The self-recurrent connection (with weight 1.0) indicates feedback with a delay of 1 time step. It builds the basis of the "constant error carrousel" CEC. The gate units open and close access to CEC. See text and appendix A.1 for details.*

Schmidhuber, 1997.

2012: Deep learning breakthrough

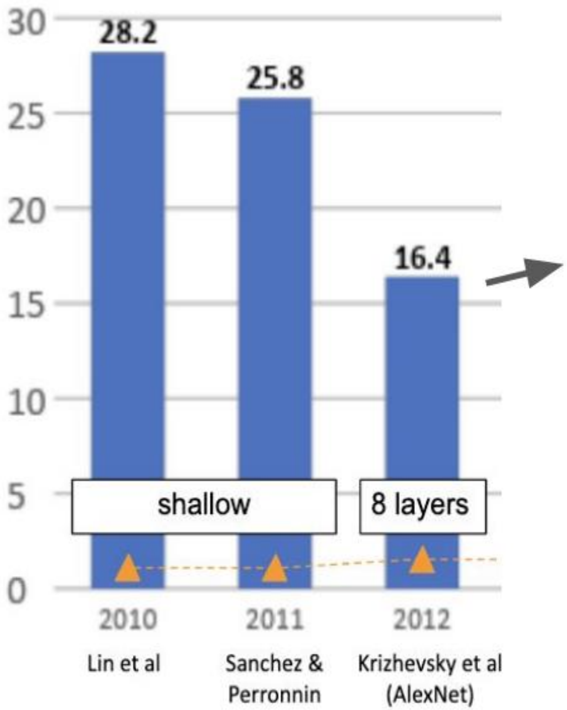

ImageNet Visual Recognition Challenge results.

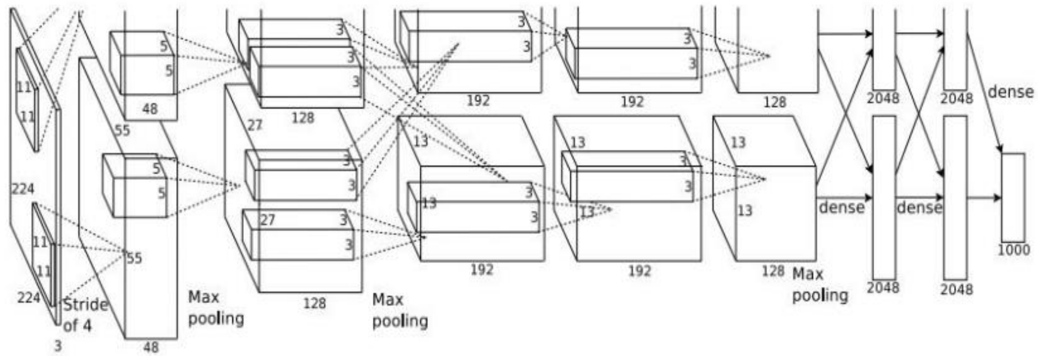

Krizhevsky et al. 2012. 8-layer “AlexNet”.

Key ingredients of deep learning

Algorithms

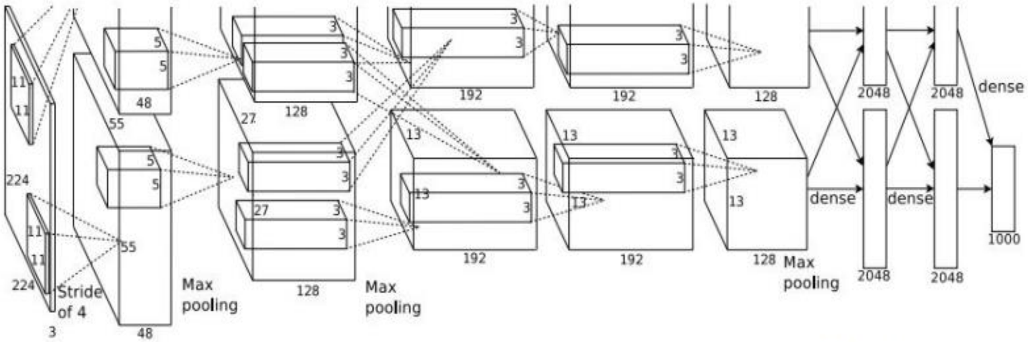

Compute

Data

ImageNet is an image database organized according to the WordNet hierarchy (currently only the nouns), in which each node of the hierarchy is depicted by hundreds and thousands of images. The project has been instrumental in advancing computer vision and deep learning research.

2015: Very deep convnets and challenging vision tasks

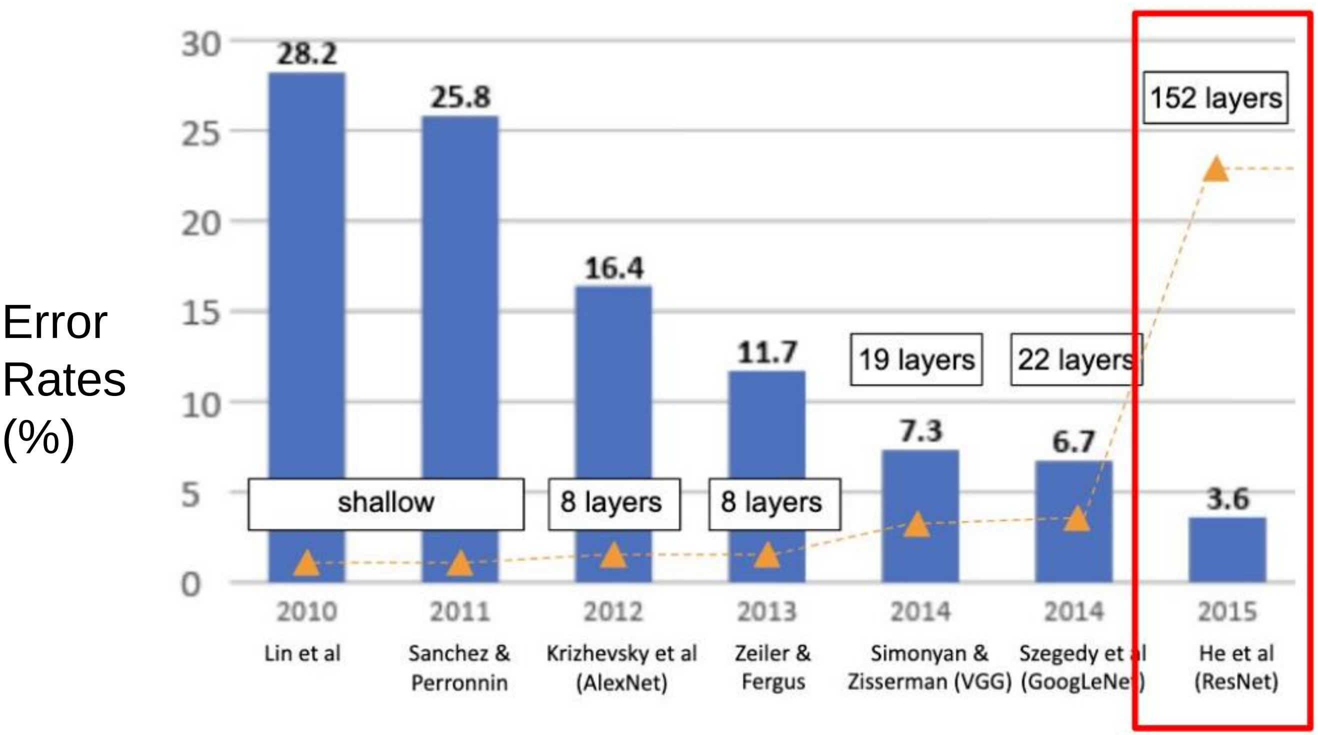

He et al. 2015. ResNet.

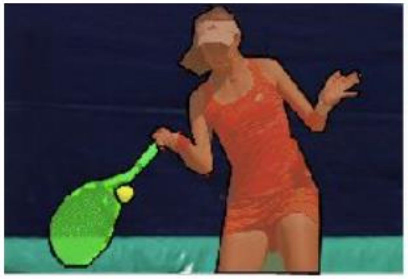

2018: Breakthroughs in deep learning for natural language processing (sequences)

Transformer architectures and pre-training -> fine-tuning. State-of-the-art on 11 NLP benchmarks.

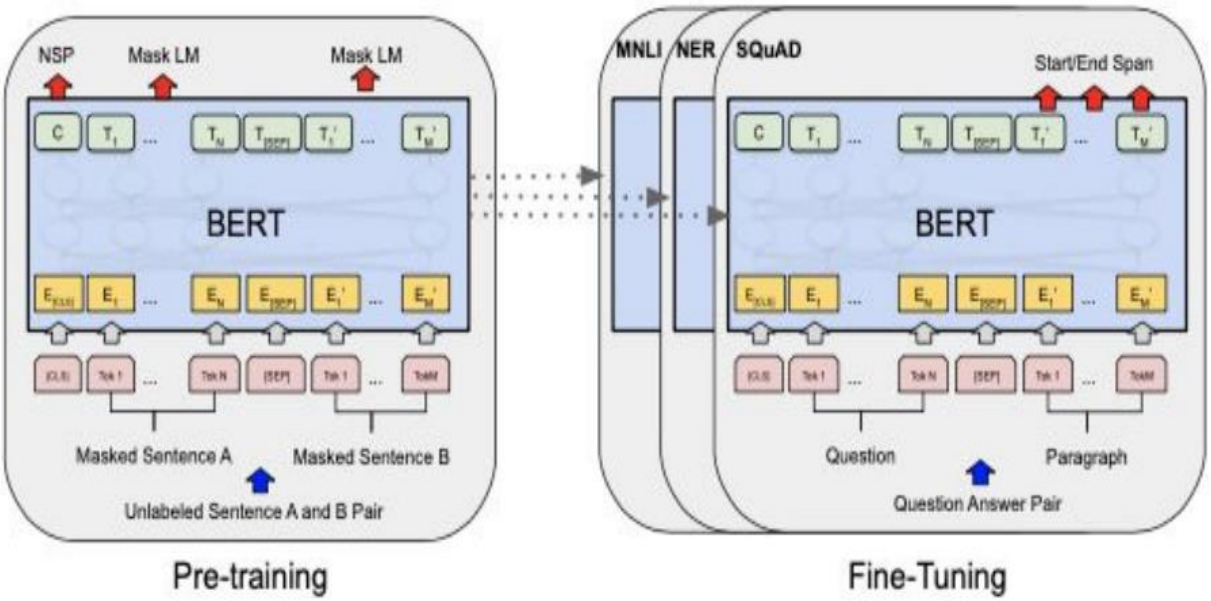

Devlin et al. 2018. BERT.

Machine Translation

- The automatic translation of texts between languages is one of the oldest non-numerical applications in Computer Science.
- In the past 15 years or so, MT has gone from a niche academic curiosity to a robust commercial industry.

Trained on text data, neural machine translation is quite good!

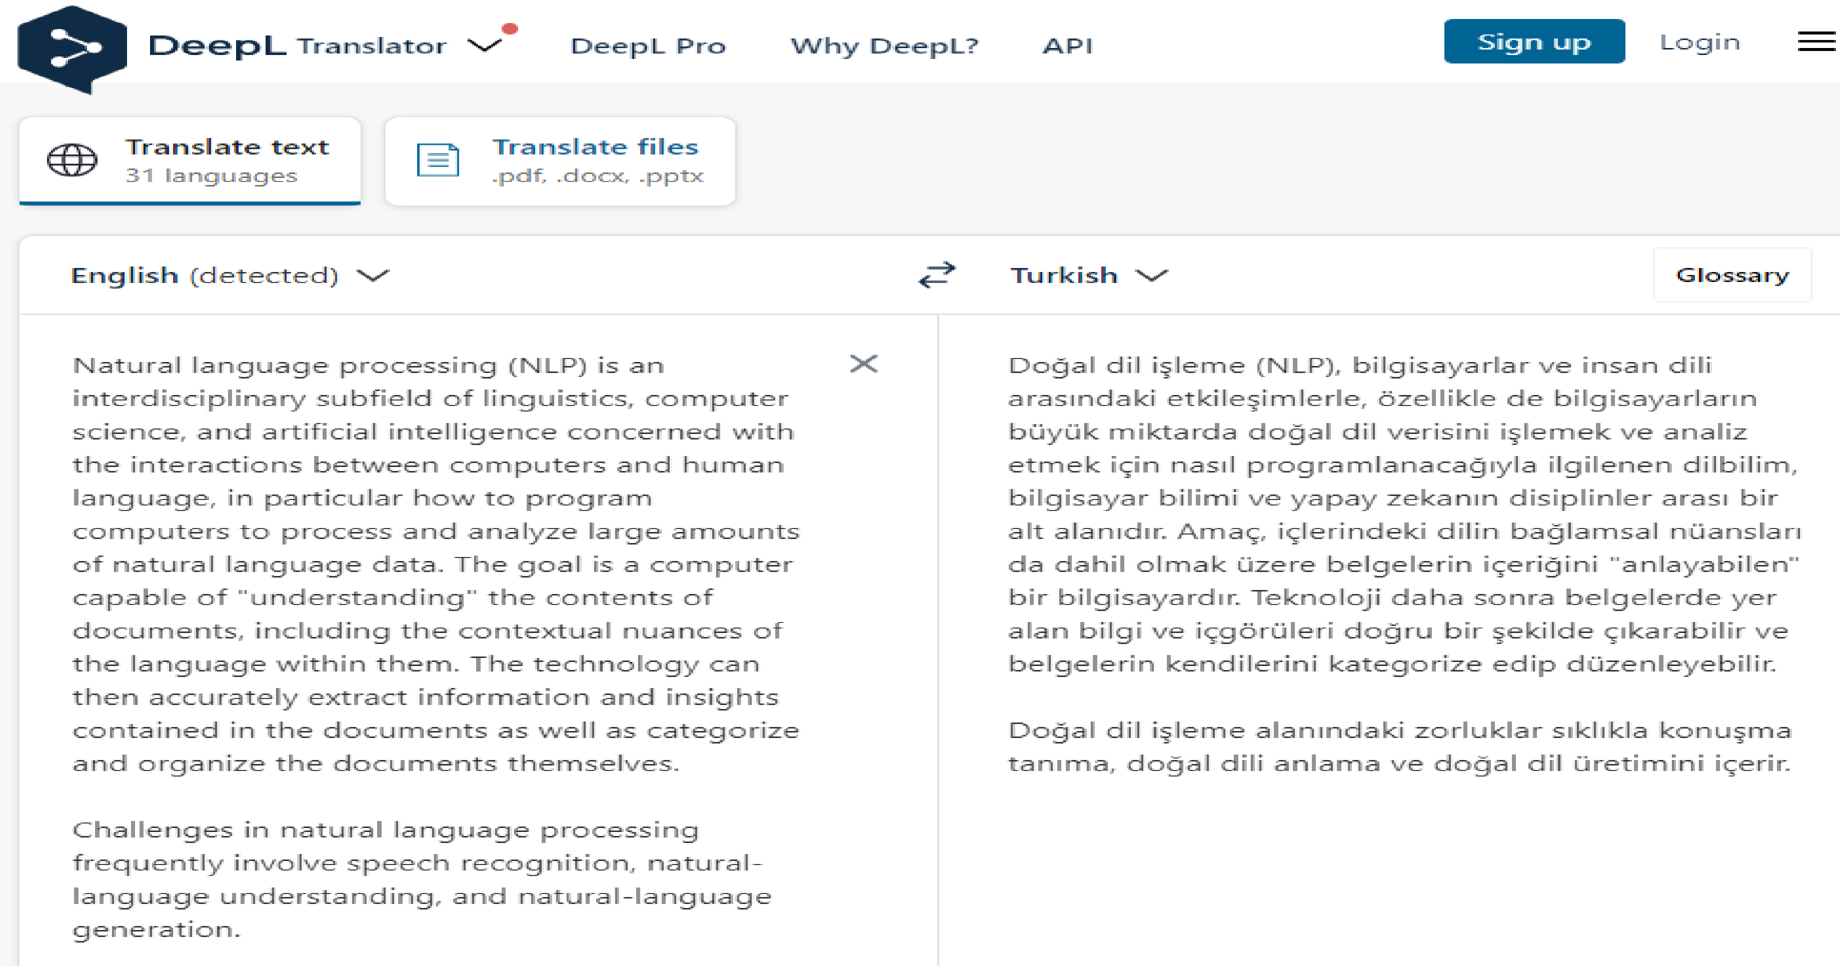

2020: Very large scale text and image generation models

2020: Very large scale text and image generation models

OpenAI models for text generation (left), text-to-image generation (right-top), and zero-shot classification tasks (right-bottom)

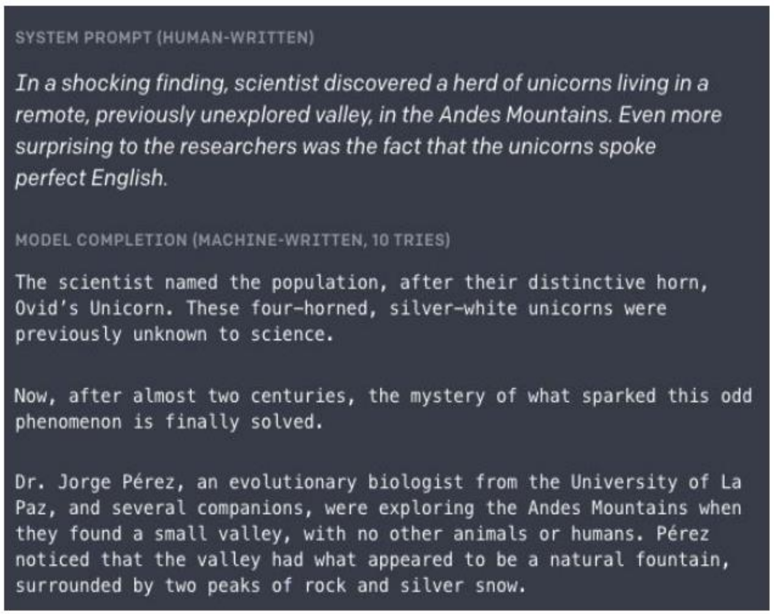

SYSTEM PROMPT (HUMAN-WRITTEN)

In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

MODEL COMPLETION (MACHINE-WRITTEN, 10 TRIES)

The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.

Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.

Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.

GPT-3 (figure from GPT-2). Brown et al. 2020.

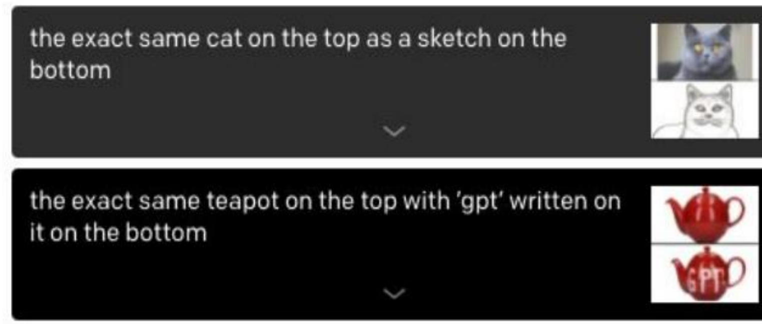

DALL-E. Ramesh et al. 2021.

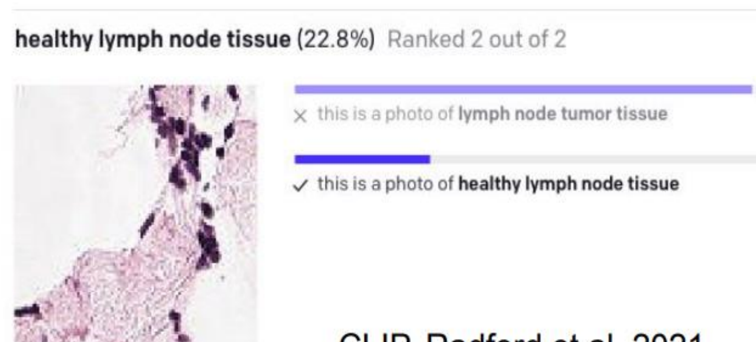

CLIP. Radford et al. 2021.

Deep learning for healthcare: the rise of medical data

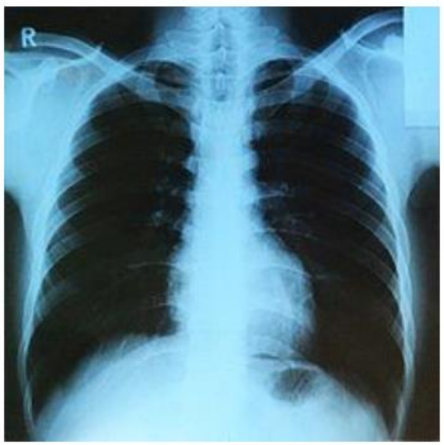

X-rays (invented 1895).

CT (invented 1972).

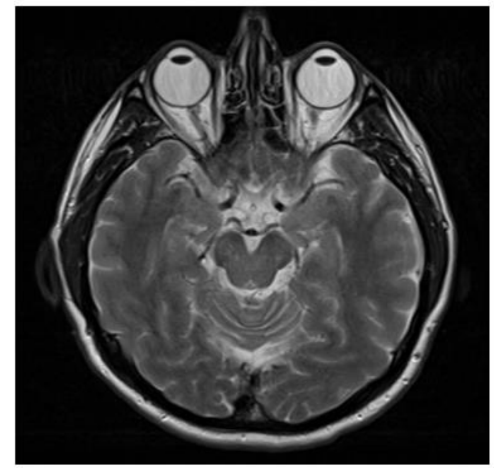

MRI (invented 1977).

Q: What are other examples of medical data?

Electronic health records -- making patient data available

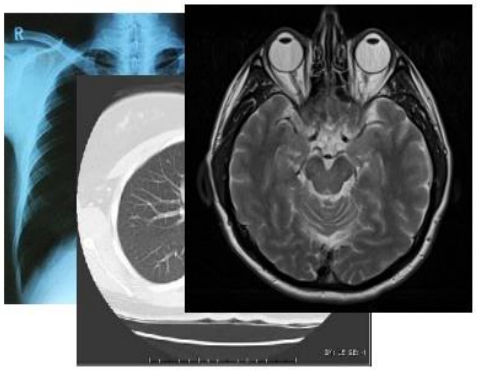

Imaging data

Patient measurements

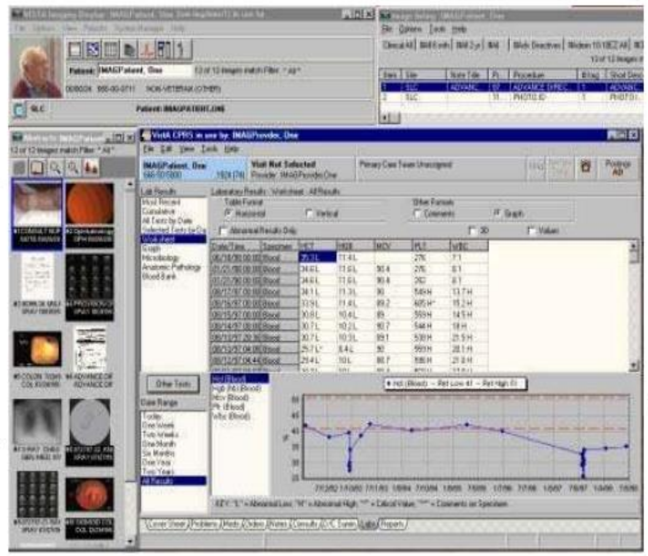

1960s: invention
1980s: increased effort
2009: 51% adoption, HITECH Act
2017: 98% adoption

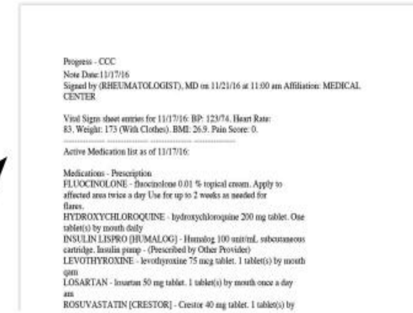

Progress - CCC
Note Date:11/17/16
Signed by (RHEUMATOLOGIST), MD on 11/21/16 at 11:00 am Affiliation: MEDICAL CENTER

Vital Signs sheet entries for 11/17/16: BP: 123/74. Heart Rate: 83. Weight: 173 (With Clothes). BMI: 26.9. Pain Score: 0.

Active Medication list as of 11/17/16:

Medications - Prescription
FLUOCINOLONE - fluocinolone 0.01 % topical cream. Apply to affected area twice a day Use for up to 2 weeks as needed for flares.
HYDROXYCHLOROQUINE - hydroxychloroquine 200 mg tablet. One tablet(s) by mouth daily
INSULIN LISPRO [HUMALOG] - Humalog 100 unit/mL subcutaneous cartridge. Insulin pump - (Prescribed by Other Provider)
LEVOTHYROXINE - levothyroxine 75 mcg tablet. 1 tablet(s) by mouth qam
LOSARTAN - losartan 50 mg tablet. 1 tablet(s) by mouth once a day am
ROSUVASTATIN [CRESTOR] - Crestor 40 mg tablet. 1 tablet(s) by

Clinical notes

Test	Value	Reference
Hemoglobin	8.0	8.5–11.0 mmol/L
C-reactive protein	279	<5 mg/L
Red blood cell count	3.86	4.3–6.0 × 10^{12}/L
White blood cell count	27.1	4.0–10.0 × 10^{9}/L
Thrombocytes	462	150–400 × 10^{9}/L
Glucose	12.9	4.0–7.8 mmol/L
Sodium	127	135–145 mmol/L
Potassium	4.2	3.5–5.0 mmol/L
Creatinine	40	50–110 μmol/L
Estimated glomerular filtration rate	>90	>60 ml/min
Ureum	3.2	2.5–7.5 mmol/L
Lactate dehydrogenase	166	<250 U/L
Aspartaat aminotransferase	14	<40 U/L
Alanine aminotransferase	13	<50 U/L
Alkaline phosphatase	127	<120 U/L
Gamma-glutamyl transferase	96	<50 U/L

Lab results

Genomics data

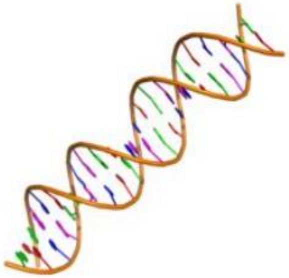

1953 - Watson and Crick discover double helix structures of DNA

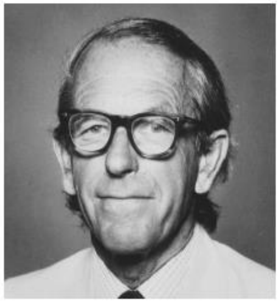

1977 - Fred Sanger sequences first full genome of a virus

1990 - 2003: Human Genome Project sequences full human genome

2003: ENCODE project launched to identify and characterize genes in human genome

2008 - 2015: 1000 Genomes Project International effort to study human genetic variation

2006 - present: UK Biobank Project Genetic data and intended 30 years of health follow-up for 500k individuals in the UK

Wearables and other sensor data

First iPhone: 2007

Fitbit: 2009

Apple Watch: 2014

AI in healthcare: biomedical image interpretation

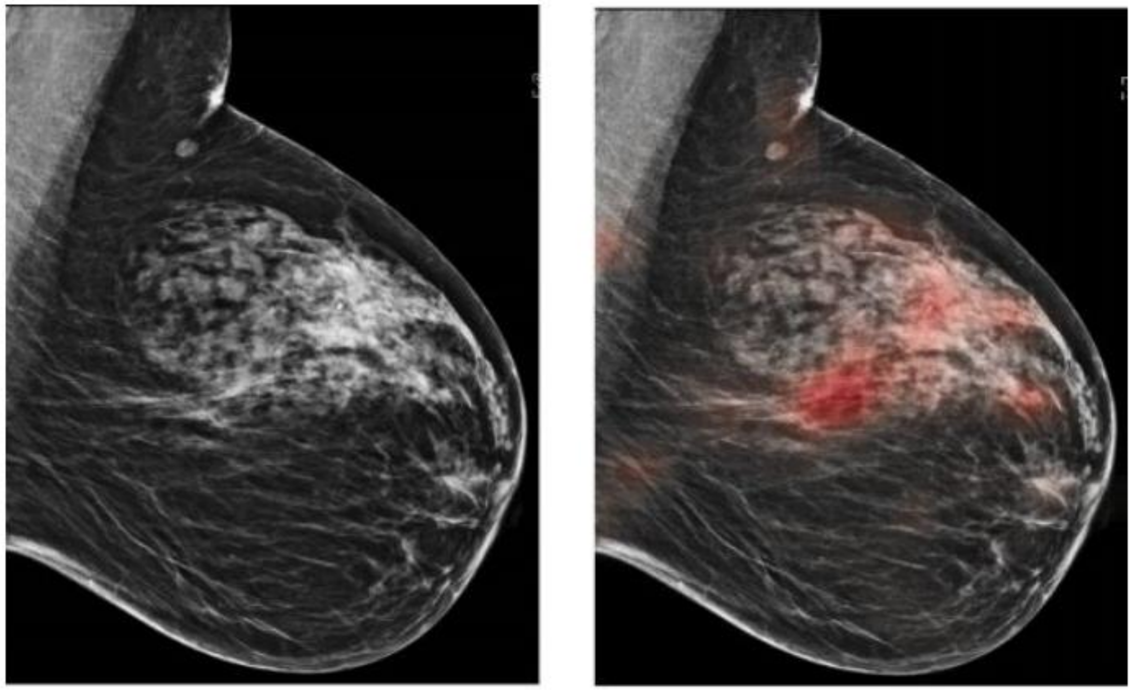

Wu et al. 2019

Cancer-associated adipocytes: key players in breast cancer progression

AI in healthcare: genomic analysis

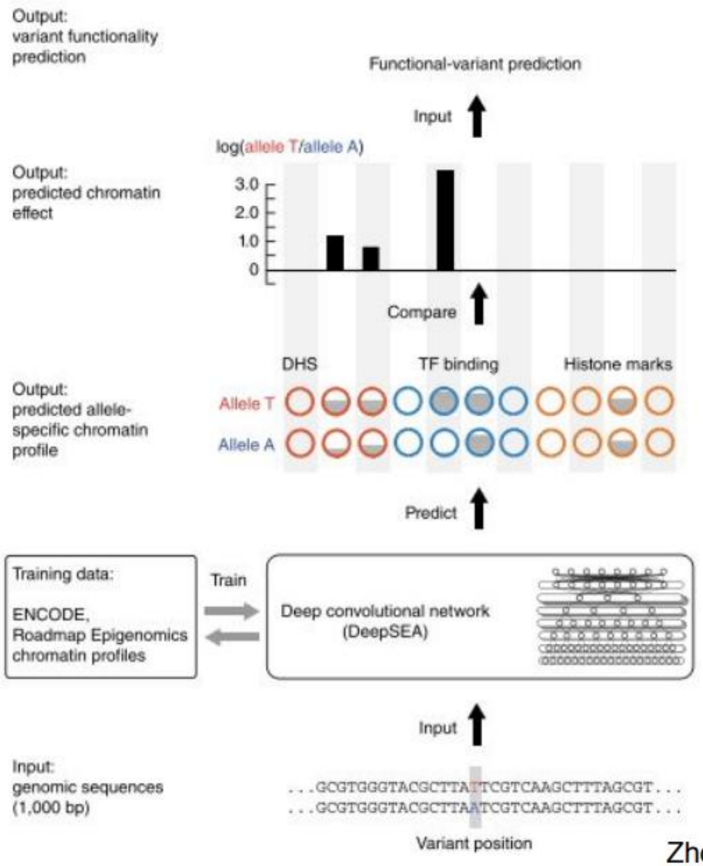

Zhou et al. 2015

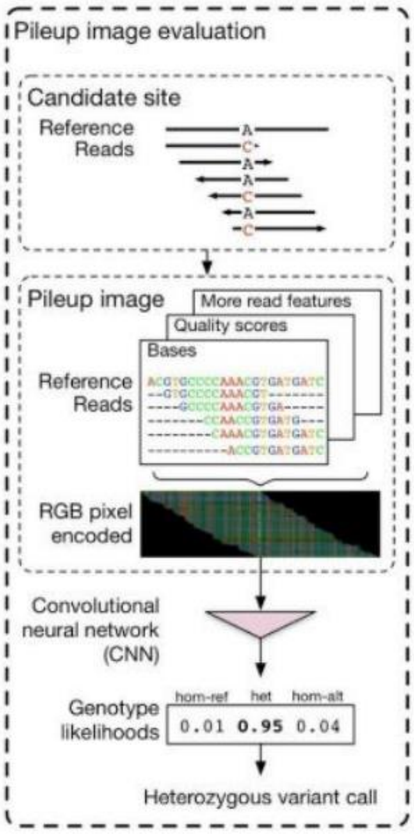

Poplin et al. 2016

AI in healthcare: drug discovery and drug interaction prediction

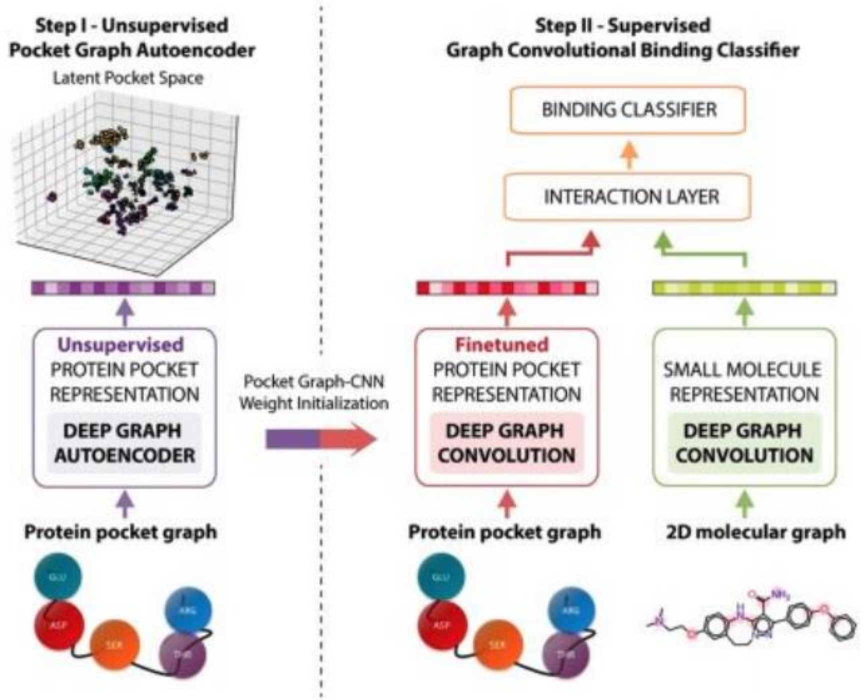

Torng et al. 2019

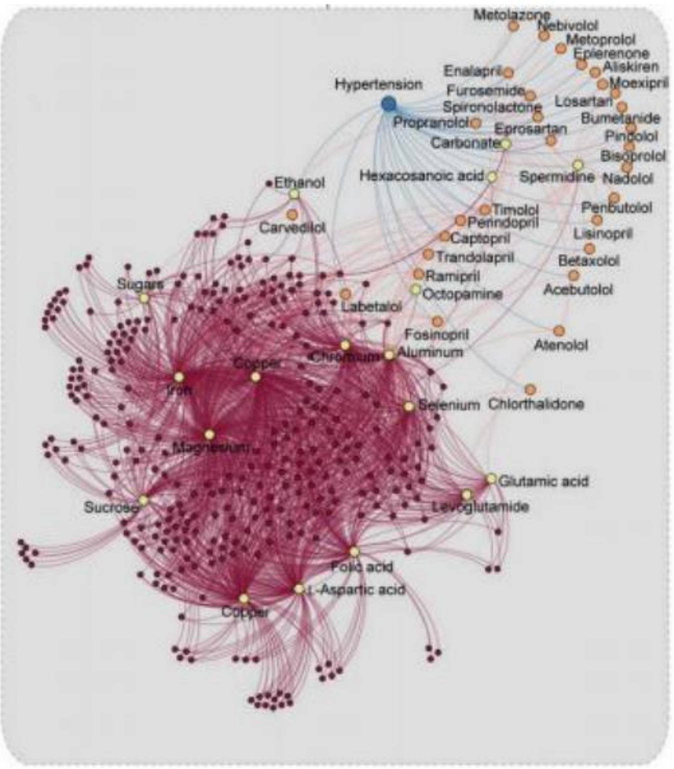

Ryu et al. 2018

Uncertainty and AI / human collaboration

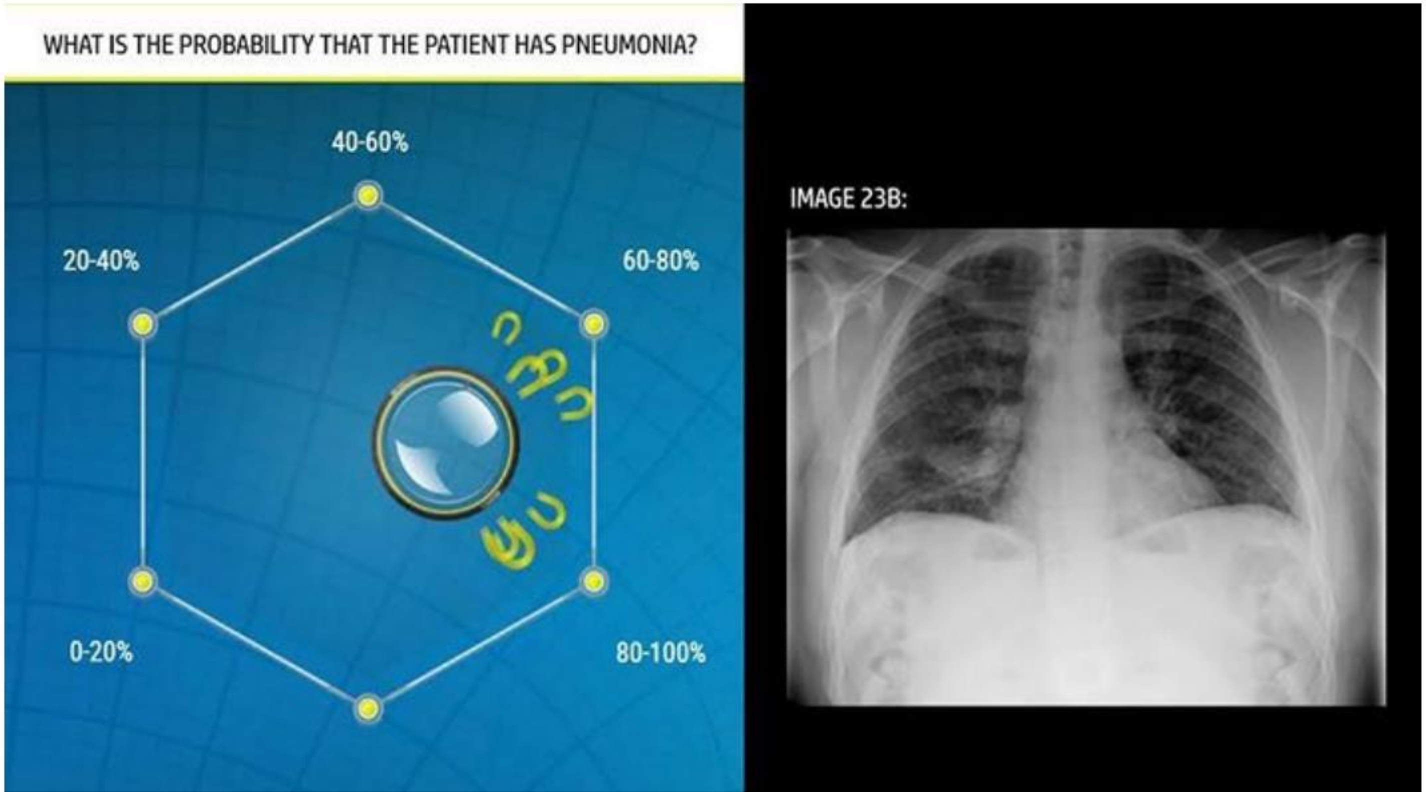

Rosenberg et al. 2018

Privacy and security

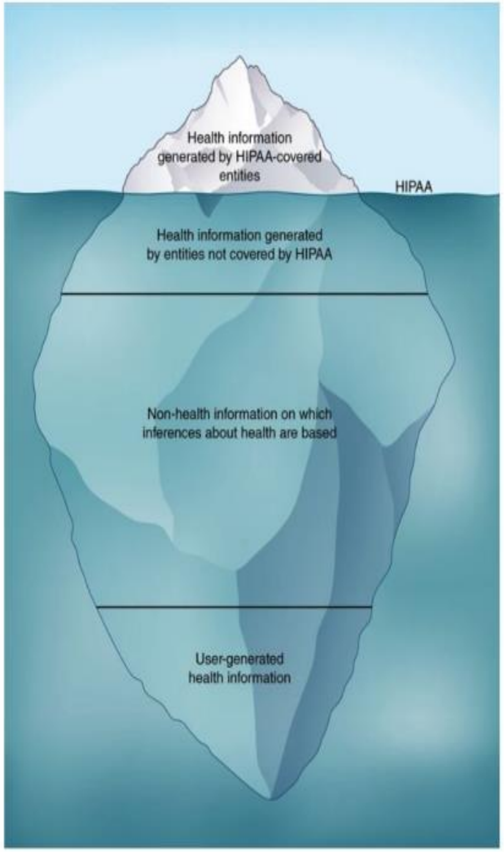

Price et al. 2019

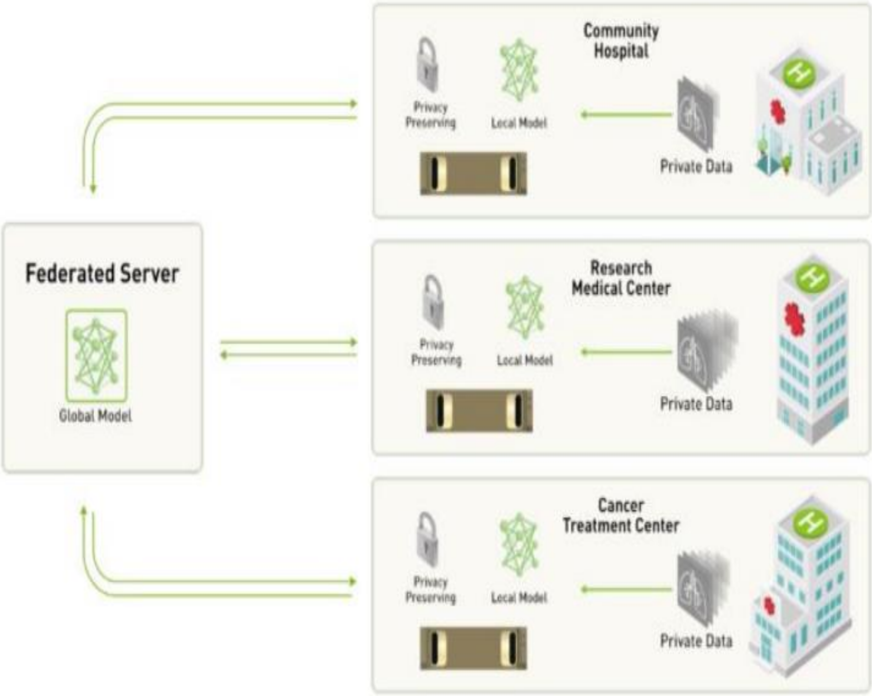

Figure: https://news.developer.nvidia.com/first-privacy-preserving-federated-learning-system/

Health Insurance Portability and Accountability Act (HIPAA)

HIPAA protects all “individually identifiable health information” held or transmitted by a covered entity or its business associate, in any form or media, whether electronic, paper, or oral.

DS vs AI vs ML vs DL – Difference

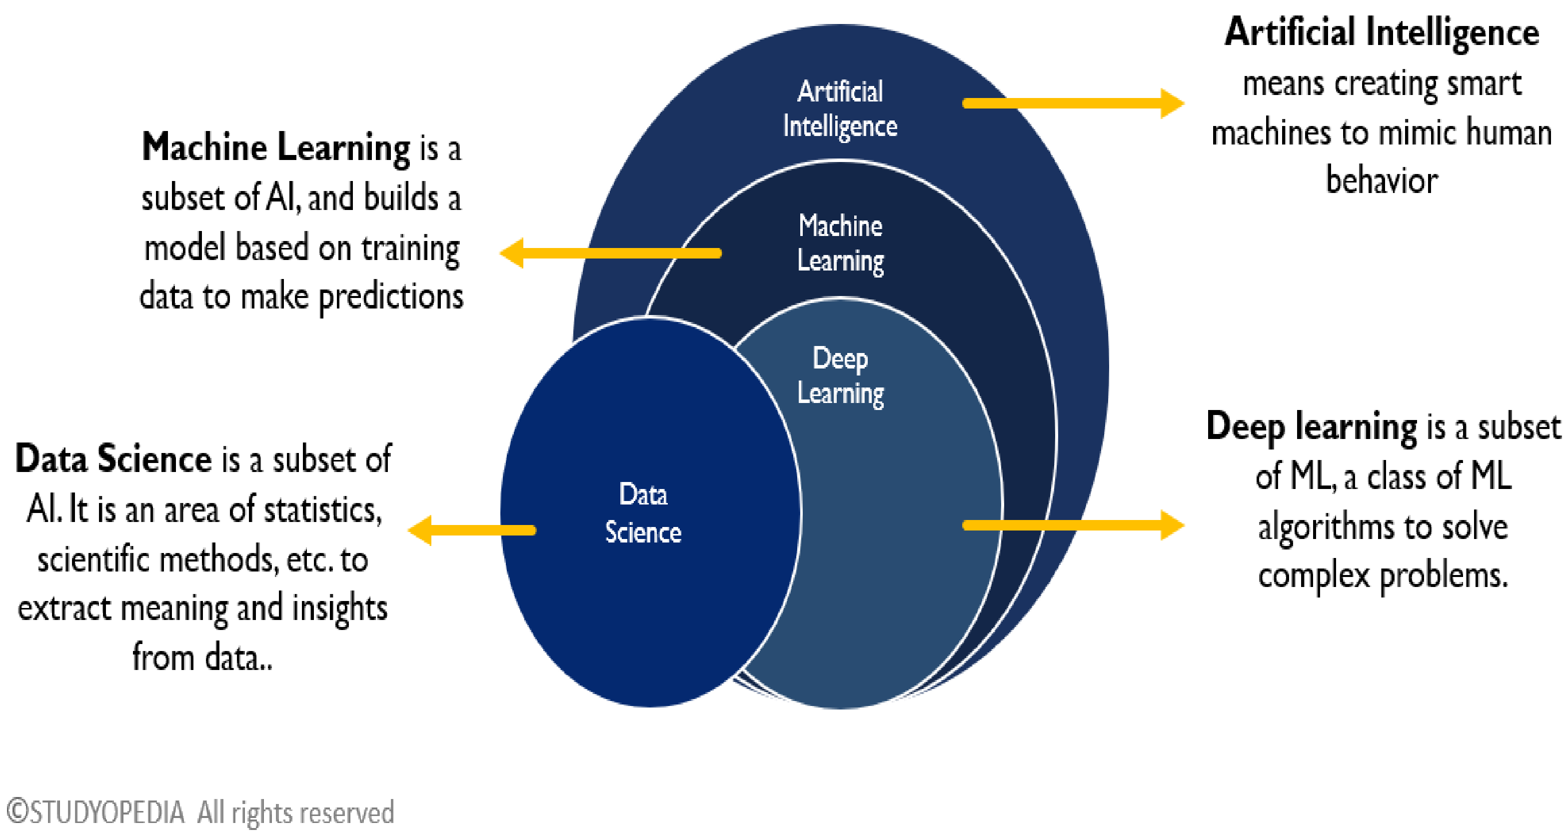

Machine Learning

Machine learning is a subset of artificial intelligence (AI) that focuses on developing algorithms and statistical models that enable computers to learn from and make predictions or decisions based on data. Here are some key points:

1. **Learning from Data**: Machine learning algorithms improve their performance as they are exposed to more data. They identify patterns and relationships within the data to make informed decisions.
2. **Types of Machine Learning**:
   - **Supervised Learning**: The algorithm is trained on labeled data, meaning the input comes with the correct output. Examples include classification and regression tasks.
   - **Unsupervised Learning**: The algorithm works with unlabeled data and tries to find hidden patterns or intrinsic structures. Examples include clustering and association tasks.
   - **Reinforcement Learning**: The algorithm learns by interacting with an environment, receiving rewards or penalties based on its actions, and aims to maximize the cumulative reward.
3. **Applications**: Machine learning is used in various fields such as healthcare (predicting diseases), finance (fraud detection), marketing (customer segmentation), and many more.

Learning from Data

Medical Insurance Price Prediction using Machine Learning

	age	sex	bmi	children	smoker	region	charges
0	19	female	27.900	0	yes	southwest	16884.92400
1	18	male	33.770	1	no	southeast	1725.55230
2	28	male	33.000	3	no	southeast	4449.46200
3	33	male	22.705	0	no	northwest	21984.47061
4	32	male	28.880	0	no	northwest	3866.85520
...	...	...	...	...	...	...	...
1333	50	male	30.970	3	no	northwest	10600.54830
1334	18	female	31.920	0	no	northeast	2205.98080
1335	18	female	36.850	0	no	southeast	1629.83350
1336	21	female	25.800	0	no	southwest	2007.94500
1337	61	female	29.070	0	yes	northwest	29141.36030

1338 rows × 7 columns

This dataset contains 1338 data points with 6 independent features and 1 target feature(charges).

https://www.geeksforgeeks.org/medical-insurance-price-prediction-using-machine-learning-python/

Scatter plot of the charges paid v/s age and BMI respectively

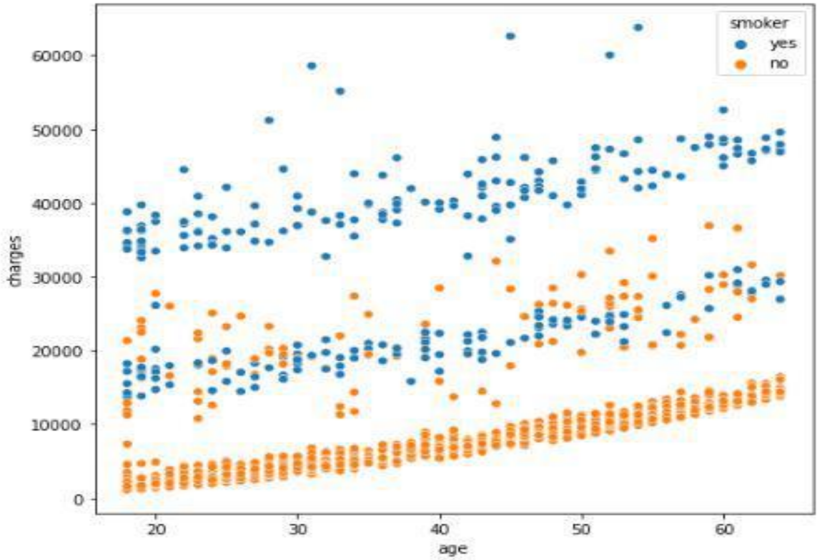

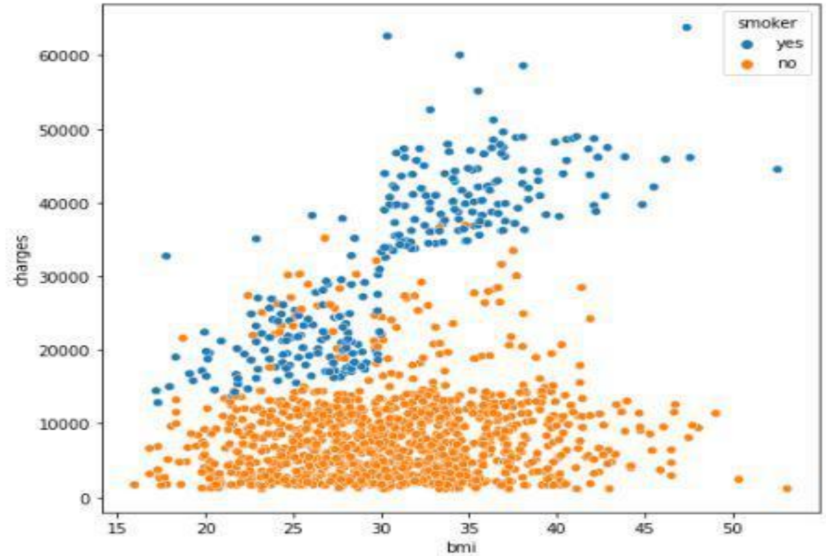

- **Exploratory Data Analysis (EDA**) : EDA is an approach to analyzing the data using visual techniques. It is used to discover trends, and patterns, or to check assumptions with the help of statistical summaries and graphical representations. While performing the EDA of this dataset we will try to look at what is the relation between the independent features that is how one affects the other.
- **Model Development** : There are many ML models available that find and learn relationship between features and target variable.
- Then, we **use these models** in our applications
  (to predict for insurance price detection in our example.)

Machine Learning Models

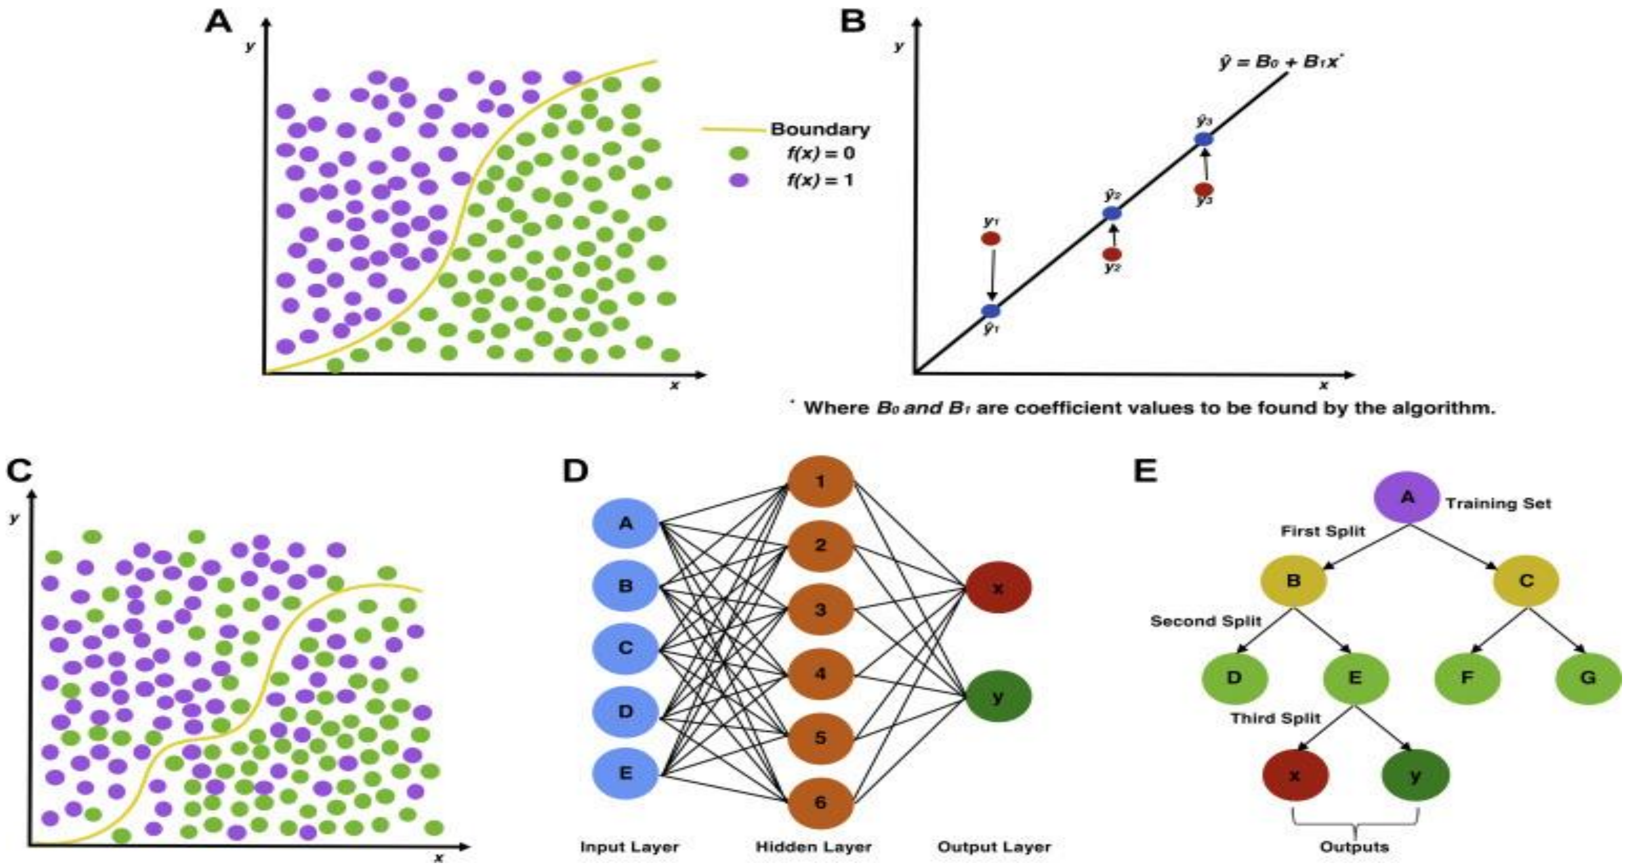

- Graphical representation of traditional statistical approaches to regression, with logistic (A) and linear regression (B) on the top row.
- The bottom row demonstrates machine learning models graphically, with support vector machine (C), artificial neural network (D), and decision tree (E) approaches.

Neural network model

- TensorFlow Playground is a powerful educational resource that enables users to gain practical experience in building and training neural networks through interactive experimentation with different architectures, activation functions, and datasets.

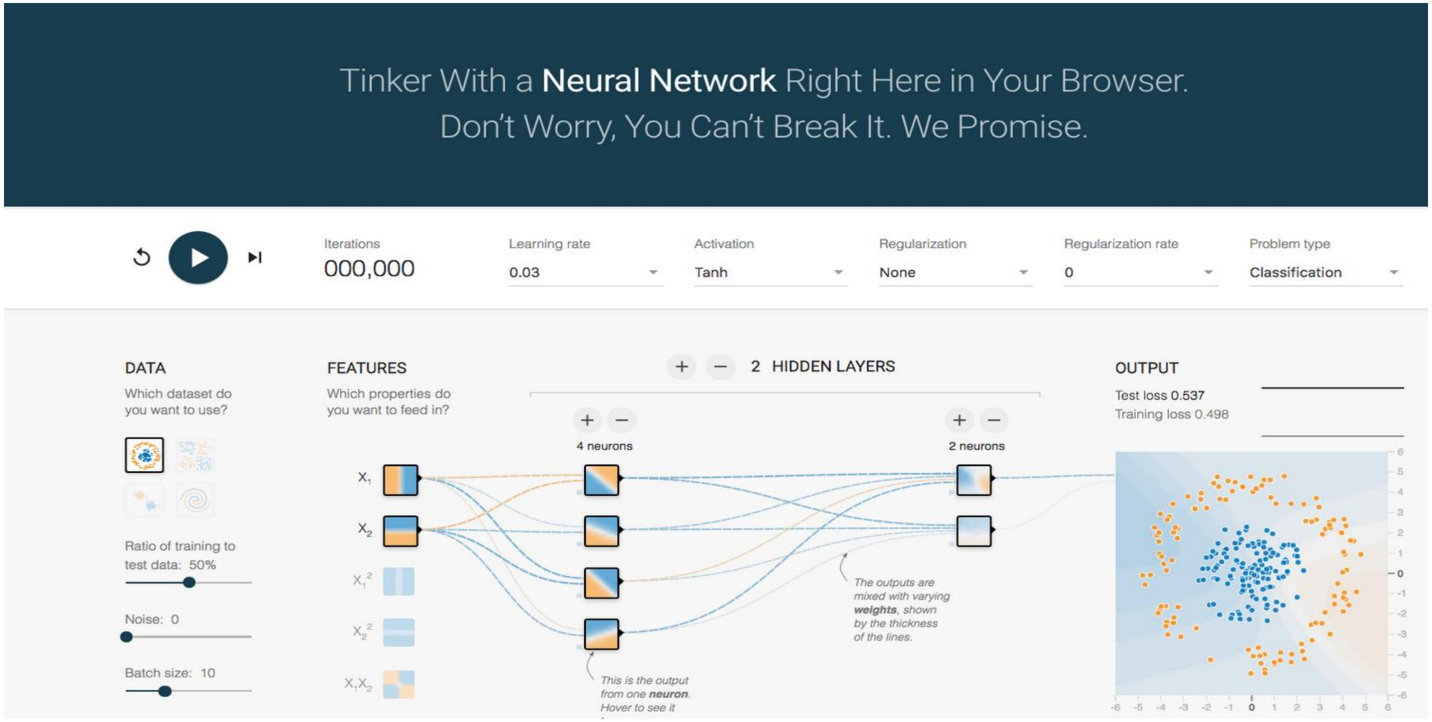

https://playground.tensorflow.org

Machine Learning vs Deep Learning

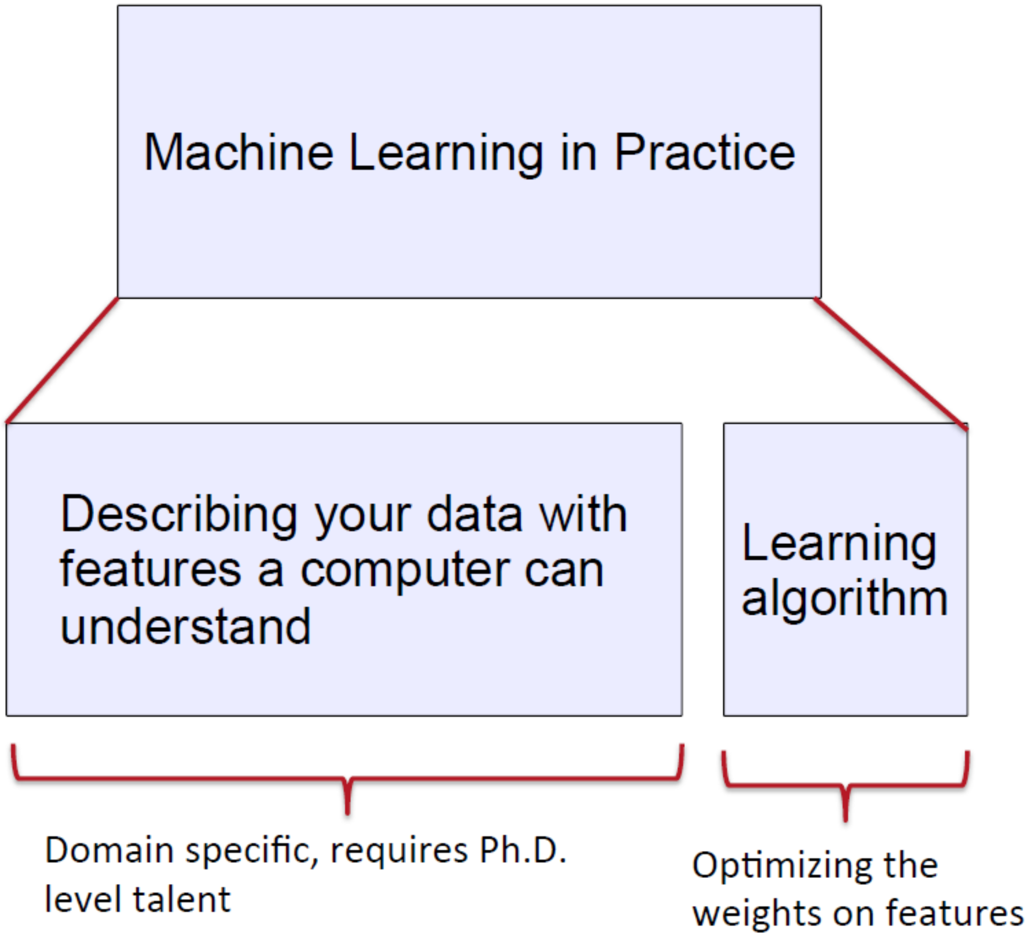

What's Deep Learning (DL)?

- Representation learning attempts to automatically learn good features or representations

- Deep learning algorithms attempt to learn (multiple levels of) representation and an output

- From "raw" inputs **x** (e.g. words)

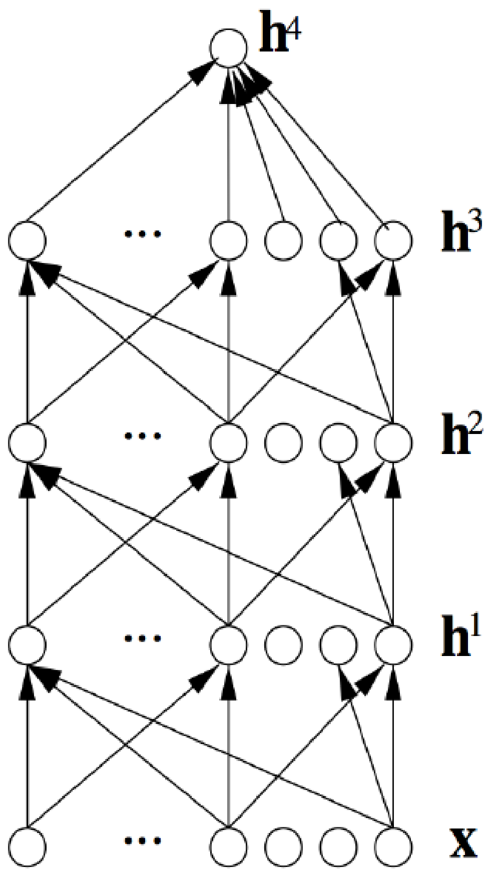

Source : cs224d

Reasons for Exploring Deep Learning

- Manually designed features are often over-specified, incomplete and take a long time to design and validate
- **Learned Features** are easy to adapt, fast to learn

- Deep learning provides a very flexible, (almost?) universal, learnable framework for **representing** world, visual and linguistic information.

- Deep learning can learn **unsupervised** (from raw text) and **supervised** (with specific labels like positive/negative)

Reasons for Exploring Deep Learning

- In 2006 **deep** learning techniques started outperforming other machine learning techniques. Why now?

- DL techniques benefit more from a lot of data
- Faster machines and multicore CPU/GPU help DL
- New models, algorithms, ideas

→ **Improved performance** (first in speech and vision, then NLP)

Machine Learning Steps

Typical steps involved in a machine learning process/project:

1. **Problem Definition**: Clearly define the problem you want to solve. Identify the target variable (the output you want to predict) and understand the business objectives.
2. **Data Collection**: Gather data that includes both input features and the target variable. Ensure the data is relevant and sufficient for training a model.
3. **Data Preprocessing**: Clean and prepare the data. This involves handling missing values, removing duplicates, normalizing or standardizing features, and encoding categorical variables.
4. **Exploratory Data Analysis (EDA)**: Analyze the data to understand its structure, distribution, and relationships. Use visualization tools and statistical methods to gain insights and identify patterns.
5. **Feature Selection and Engineering**: Select the most relevant features for the model. Create new features if necessary to improve model performance. This step can significantly impact the accuracy of the model.
6. **Data Splitting**: Split the data into training and testing sets. Typically, 70-80% of the data is used for training, and 20-30% is used for testing.
7. **Model Selection**: Choose the appropriate supervised learning algorithm based on the problem type (e.g., linear regression, decision trees, support vector machines, neural networks). Consider the nature of the data and the business requirements.

If you need more details on any of these steps, feel free to ask!

Machine Learning Steps

8. **Model Training**: Train the model using the training data. This involves fitting the model to the data and learning the relationships between the input features and the target variable.
9. **Model Evaluation**: Evaluate the model's performance using the testing data. Use metrics such as accuracy, precision, recall, F1 score, and ROC-AUC to assess the model's effectiveness.
10. **Model Tuning**: Optimize the model by adjusting hyperparameters and using techniques like cross-validation to improve its performance. This step helps in finding the best model configuration.
11. **Model Deployment**: Deploy the trained model into a production environment where it can be used to make predictions on new data. Ensure the deployment process is robust and scalable.
12. **Monitoring and Maintenance**: Continuously monitor the model's performance in the production environment. Update and retrain the model as needed to maintain its accuracy and relevance over time.

Tentative Course Outline for AI for Medicine I

WEEK	TOPIC
Week 1	What is AI? History and evolution of AI in medicine
Week 2	Programming in Python - Elementary Programming
Week 3	Programming in Python - Mathematical Functions, Strings, and Objects
Week 4	Programming in Python - Decision Statements
Week 5	Programming in Python - Loops
Week 6	Programming in Python - Functions
Week 7	Programming in Python - Multidimensional Lists, Tuples, Sets, and Dictionaries
Week 8	Programming in Python - Data Libraries/Structures for AI
Week 9	Introduction to Basic AI Models and Process for Medicine
Week 10	Data Storage, Loading, and Prepocessing/Wrangling
Week 11	Exploratory Data Analysis and Data Visualization
Week 12	Model Development
Week 13	Model Evaluation and Refinement
Week 14	Ethical, Legal, and Social Implications of AI in Medicine