# Intelligence Augmentation through Generative AI & ChatGPT

INFORMATION TECHNOLOGY UNIVERSITY OF THE PUNJAB

1 Week Hands-On Workshop for Students and Researchers

**August 7-10 2023**

**Day 1:**
Introduction to Generative AI and ChatGPT

**Day 2:**
Prompt Engineering for ChatGPT

**Day 3:**
AI-Augmented Writing Training

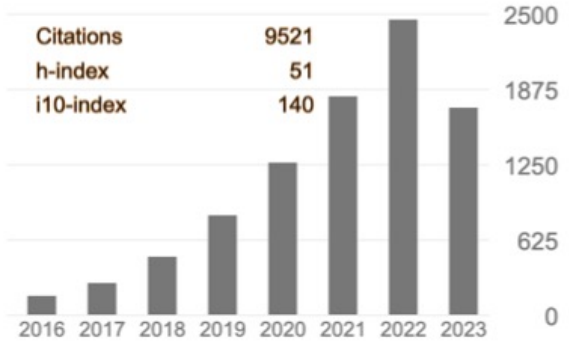
**Day 4:**
Effective Use and Ethics of ChatGPT/ Generative AI


JUNAID QADIR

Professor of Computer Engineering
Department of Computer Science and Engineering

جامعة قطر
QATAR UNIVERSITY

# JUNAID QADIR




***CITATIONS*** (1st August 2023)
9500+ (Google Scholar); 4850+ (SCOPUS)]; and 3300+ (WoS)
H-index: 51 (Google Scholar); 38 (SCOPUS); and 31 (WoS)

***Awards and Honors:***
Pakistan's HEC Best University Teacher Award (2012-2013)
ACM Distinguished Speaker 2020- *to date*.

***OVERALL PUBLICATIONS:***
***200+ Peer reviewed publications***
*127 journal publications (WOS + SCOPUS);*
99 IF journal publications *(Clarivate Analytics WOS IF)*
*82 conference publications*

***FUNDED RESEARCH***

- **QU High Impact Grant (HIG) 2023/2024 cycle**, LPI. 600,000 QAR **"Towards Robust Ethical and Pro-Social Artificial Intelligence Enabled Extended Reality (AI-XR) Metaverse Applications."** Collaborators from CENG, Kindi Center, and CLAW.
- PI, "**Development of Human-Centric Robust ML-Driven IoT Smart Services**" [QNRF NPRP13S-0206-200273]
- LPI, "Culturally informed pro-social AI regulation and persuasion framework.", **Facebook Research** AI in Ethics 2020-2021.

***THOUGHT LEADERSHIP AND COMMUNITY RECOGNITION***

1. ***Engineering Education*** *(Most cited paper on Generative AI in Eng. Ed. (published in EDUCON); invited guest editorial in Wiley Journal of Engineering Education; paper in MIT TE'2022; papers in ASEE, FIE, TALE, Education Sciences)*
2. ***Islamic AI Ethics*** *(FB Research Project; Interviews in LA Times, ABC Radio)*
3. *Among **Top 10 Global Experts** in Robust ML ; Top 25 in Trustworthy AI.*
4. *Among **Top 2% Scientists** as per Ioannidis et al. (Stanford) report*
5. *ACM Distinguished Speaker; ACM and IEEE Senior Member.*

# Credits/ Acknowledgement

Course material has been curated using various online/academic resources.

1. Wharton School, USA, Ethan and Lilach Mollick, Practical AI for Instructors and Students (YouTube: Series of 5 Videos)
2. Introduction to Generative AI, Google, https://www.youtube.com/watch?v=G2fqAlgmoPo
3. MIT Course MAS.S68, Spring 2023 (Generative AI for Constructive Communication): https://ai4comm.media.mit.edu/;
4. ChatGPT: What It Can and Cannot Do, Pascale Fung (Online)
5. Videos by Jim Krohn;
6. Various Research Papers

*The material is used as per fair use educational clause. The copyrights are retained by respective sources. We thank the respective personnel for sharing their educational materials. For practical purposes, the credits are not given on every slide, but more information about source is available on request.*

Day 1

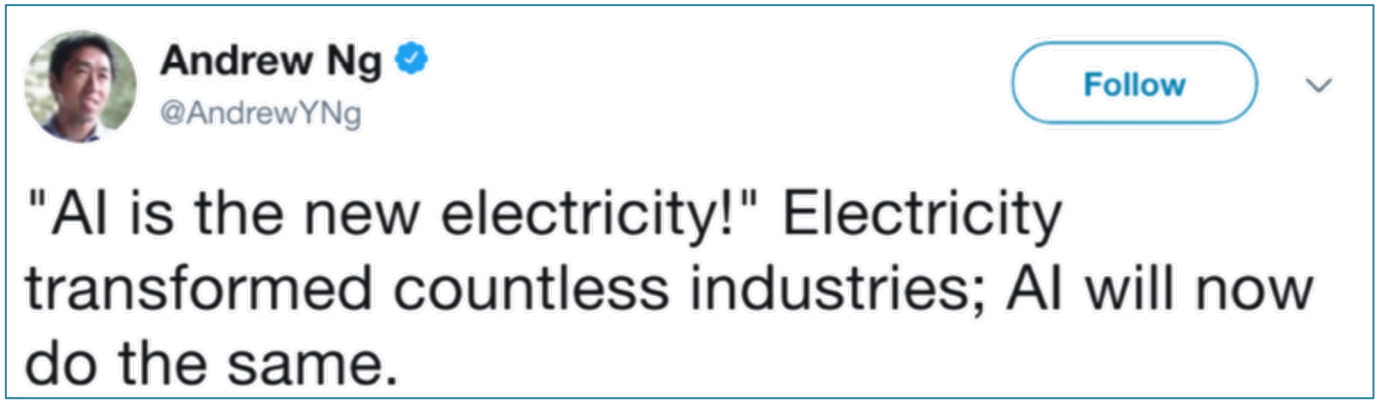
# *Introduction to Generative AI*

A Fundamentals of AI

B Generative AI Introduction

C Introduction to ChatGPT

# Success of AI (Predating ChatGPT)






FOREX/ STOCKS TRADING

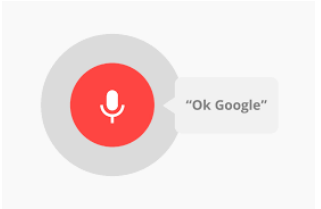


OBJECT RECOGNITION



SELF-DRIVING CARS

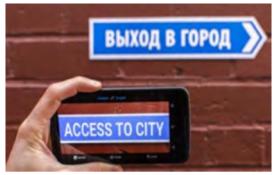
SPEECH PROCESSING




NATURAL LANGUAGE UNDERSTANDING

REINFORCEMENT LEARNING






MACHINE TRANSLATION

IMAGE CLASSIFICATION

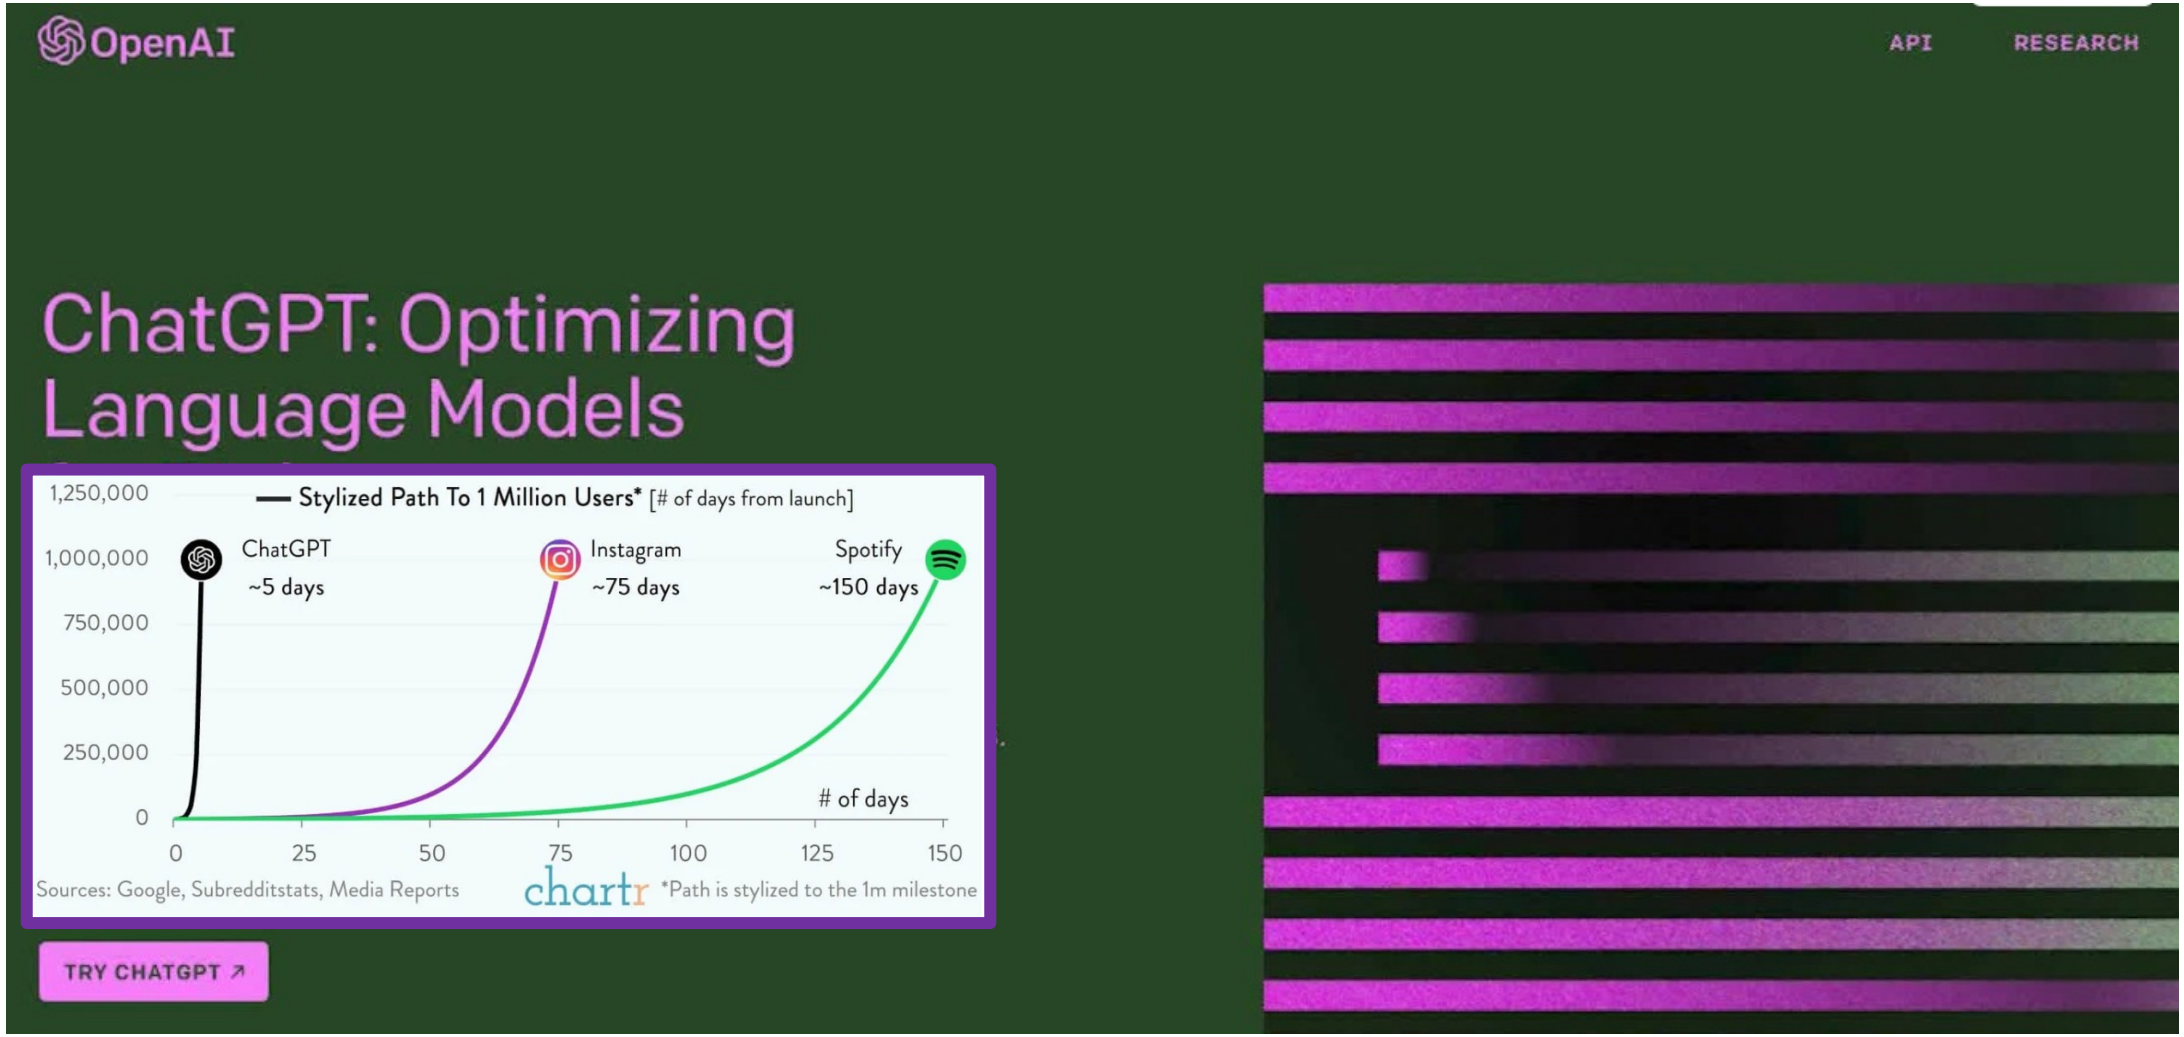




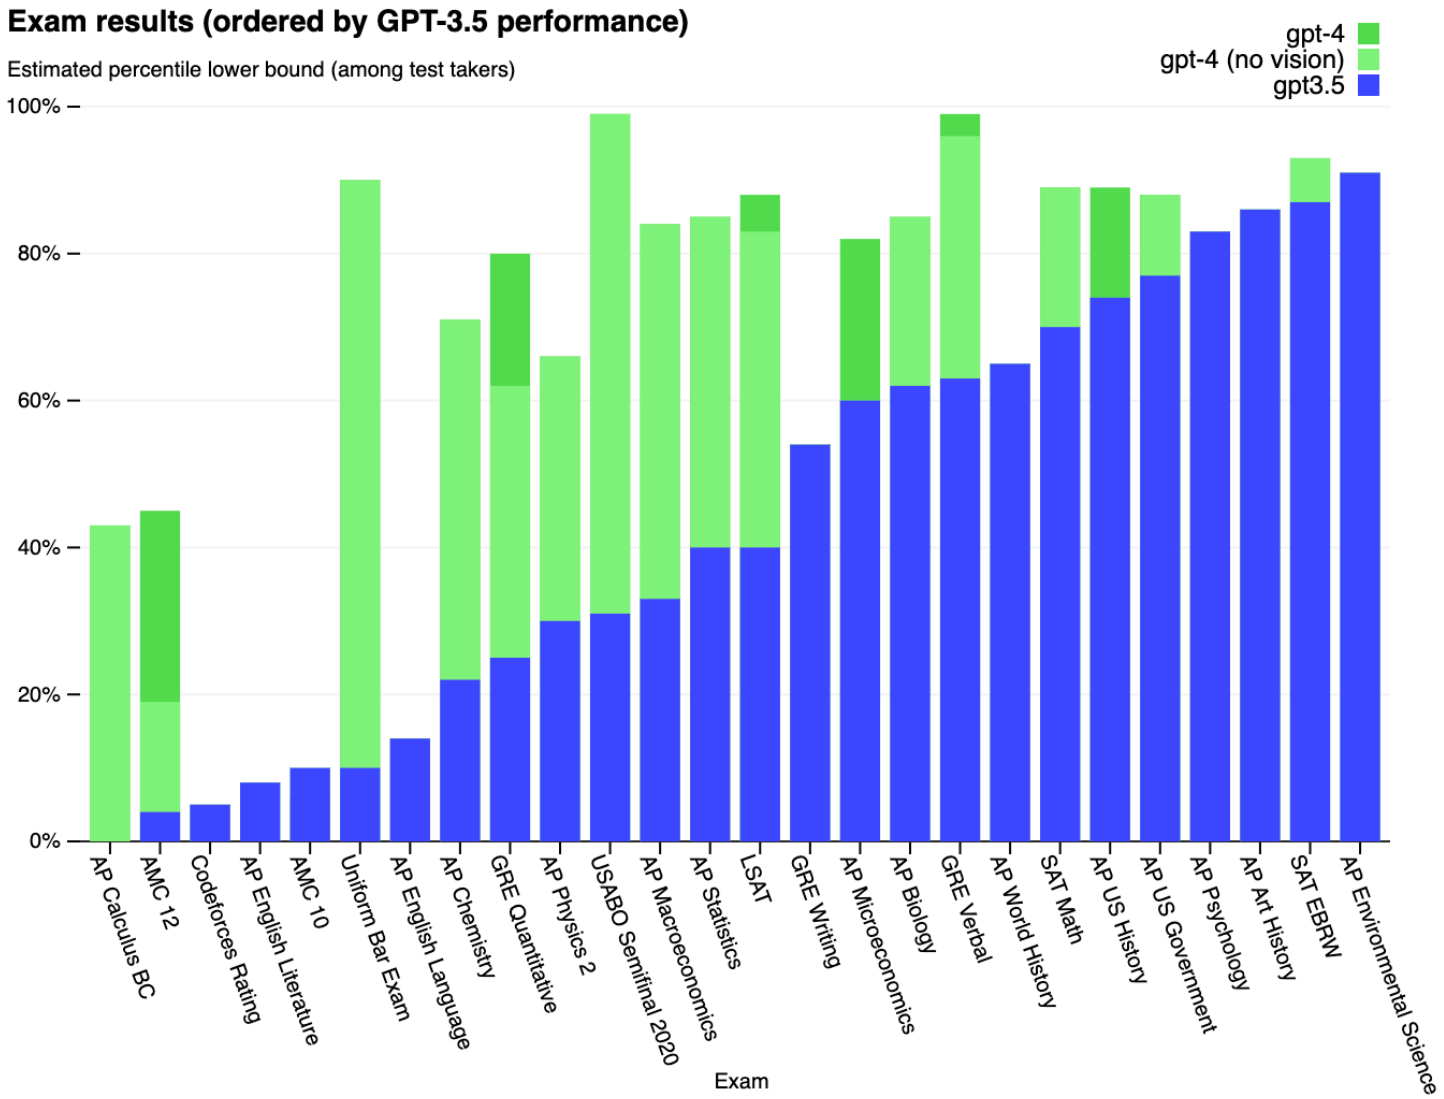

HEALTH INFORMATICS

# *And then ChatGPT happened in November 2022 …*

`https://openai.com/blog/chatgpt/`



# Success of GPT




https://openai.com/research/gpt-4

# ChatGPT Capabilities

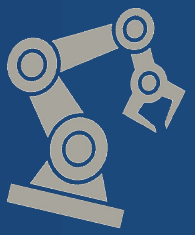
Without any fine-tuning, ChatGPT can be used for various downstream tasks:

1. **Classify text** (e.g., sentiment, specific topic categories)
2. **Recognize named entities** (e.g., people, locations, dates)
3. **Tag parts of speech** (e.g., noun, verb, adjective)
4. **Question-answer** (e.g., find answer within provided context)
5. **Summarize** (short summary that preserves key concepts)
6. **Paraphrase** (rewrite in different way while retaining meaning)
7. **Complete** (predict likely next words)
8. **Translate** (one language to another; human or code, if in training data)
9. **Generate** (again, can be code if in training data)
10. **Chat** (engage in extended conversation)

**…**







# Fundamentals of Artificial Intelligence

# How to define AI

# What is Artificial Intelligence?

*Artificial intelligence is a field renowned for its lack of consensus on fundamental issues*

- "Artificial intelligence is the study of how to make computers do things at which, at the moment, people are better" (Rich, 1983).

- "Artificial intelligence is the study of intelligent behavior" (Genesereth and Nilsson, 1987).

- "AI is a field of study that seeks to explain and emulate intelligent behavior in terms of computational processes" (Schalkoff, 1990).

# Different takes on "what is intelligence"

*The ability to perceive the environment*

*The ability to solve problems*

*The ability to learn from experience*

*The ability to reason about things*

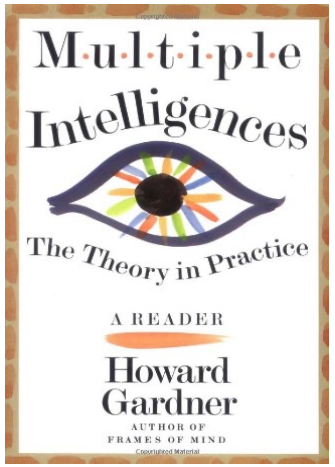
*The ability to recognize patterns*

*The ability to figure out causes of things (or to understand)*

*The ability to pursue objectives and purpose.*

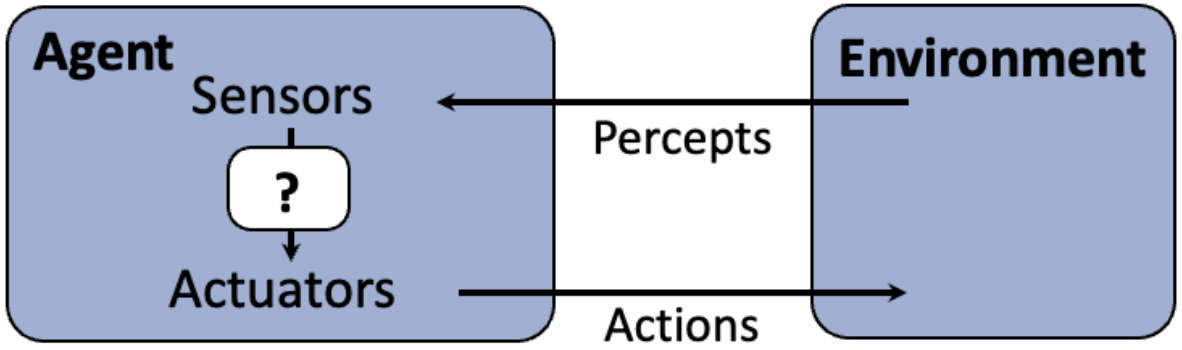
*The ability to generalize and abstract (and make analogies)*

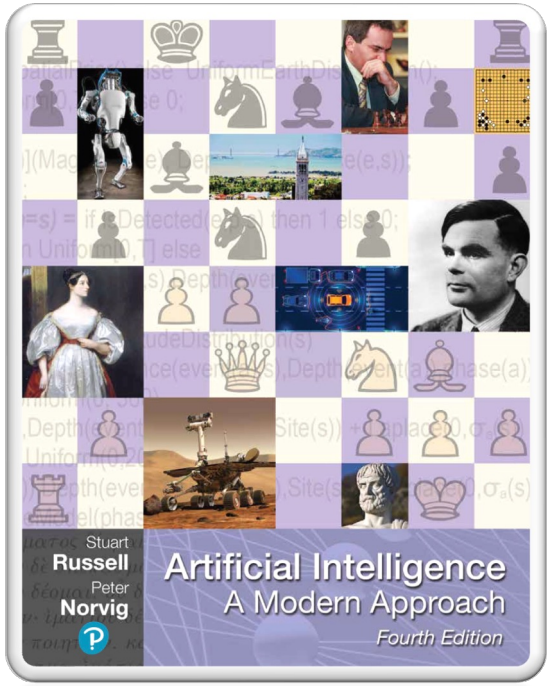


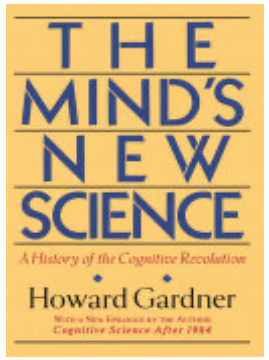
# AI focuses on building "**intelligent machines**"

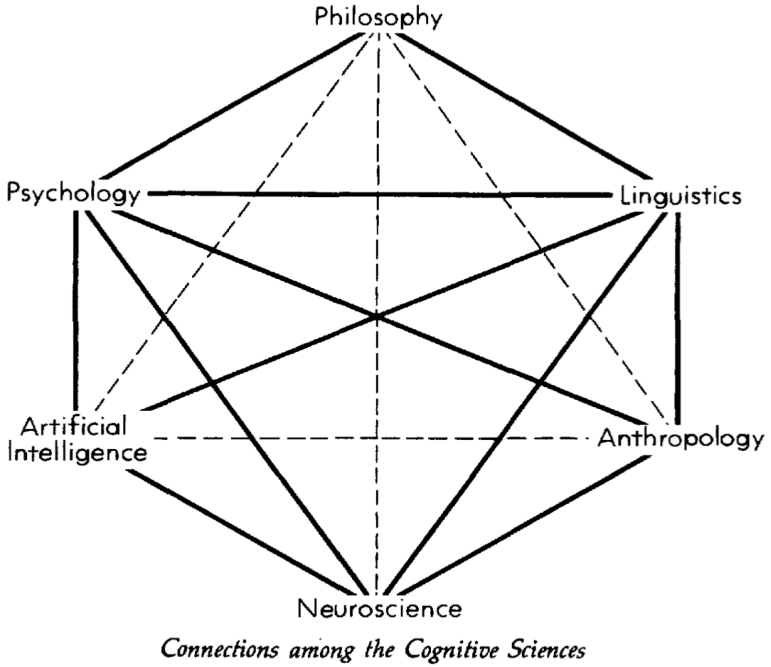
***Intelligent agent*** **is a system that perceives its environment and takes actions which maximize its chances of success.** ***(Russell & Norvig)***



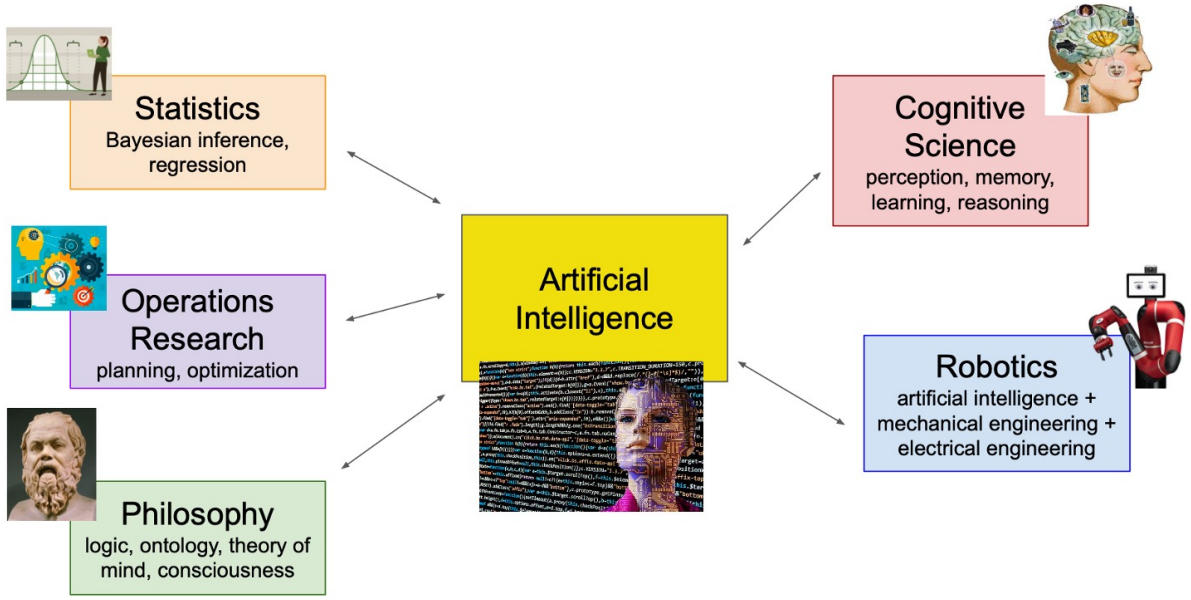



# A melting pot of ideas, techniques, and insights that relate to artificial (computational) intelligence




*Cognitive Hexagon*



# Big Ideas in AI (Expanded)

## Big Idea #1: Perception

*Computers perceive the world using sensors*



Perception is the extraction of *m*
from sensory signals using knov

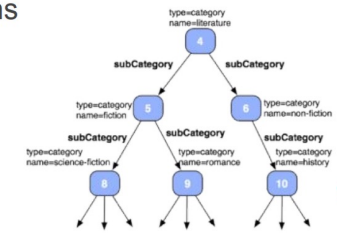
- Human senses vs. compute
- Types of perception: vision, recognition, etc.
- How perception works: algo

## Big Idea #2: Representation and Reasoning

*Agents maintain representations of the world and use them for reasoning*




rk they do
g:
ions

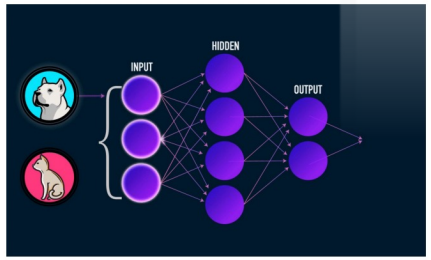


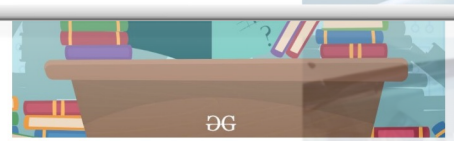

AI4K12.org

## Big Idea #3: Learnir

*Computers can learn from*

- Nature of learning
- Fundamentals of neural net
- Datasets



AI4K12.org

## Big Idea #5: Societal Impact

Artificial Intelligence can impact society in both positive and negative ways

- Ethics of AI making decisions about people
  - Fairness, bias, transparency, explainability, accountability
- Economic impacts of AI
- Cultural impacts of AI








AI4K12.org

## eraction



- Natural language understanding
- Common sense reasoning
- Affective computing & interaction (e.g. with robots or speech agents)
- Consciousness and philosophy of mind



AI4K12.org

# Brief history …

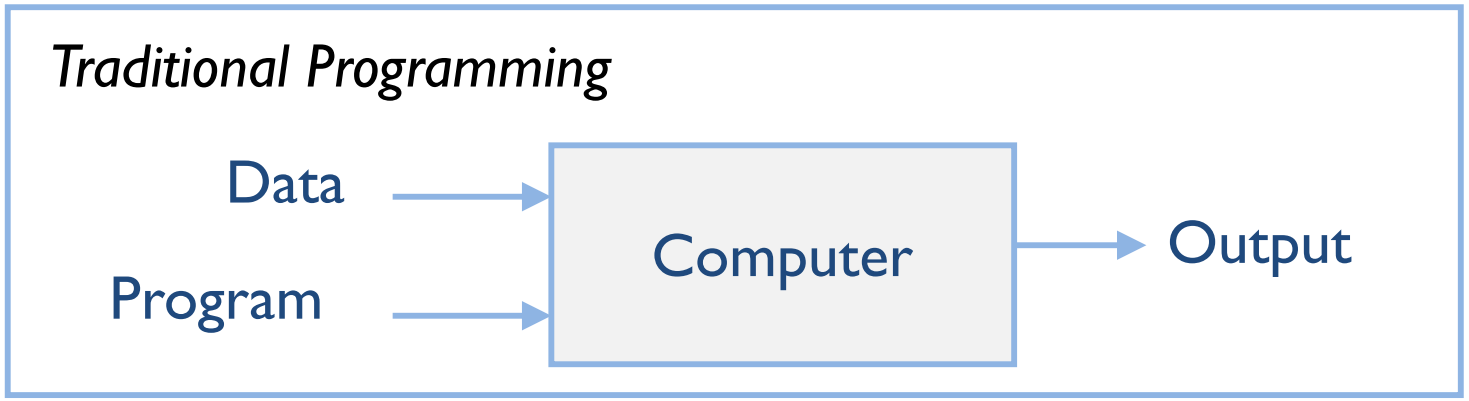
*how we got to now …*

# History of AI—Knowledge Based Approaches

1970—90: Knowledge-based approaches

1969—79: Early development

1980—88: **Expert systems** industry booms




1988—93: Expert systems industry busts

# AI Winter and Underwhelming Results

ARTIFICIAL INTELLIGENCE
WINTER



*The spirit is willing but the flesh is weak.*

↓

(Russian)

↓

*The vodka is good but the meat is rotten.*

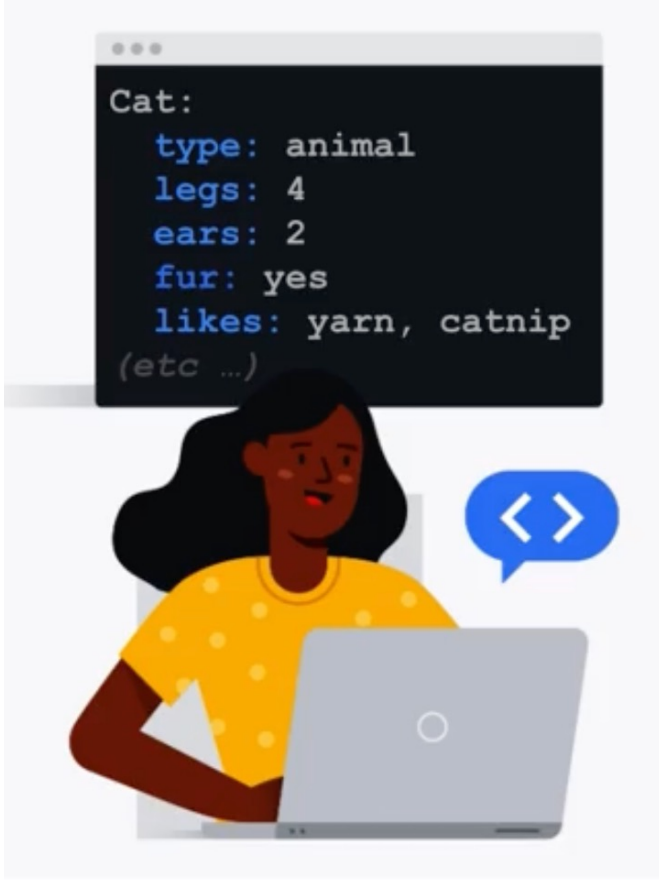
Complexity of the world made it hard to encode all rules

Problem space grew too quickly

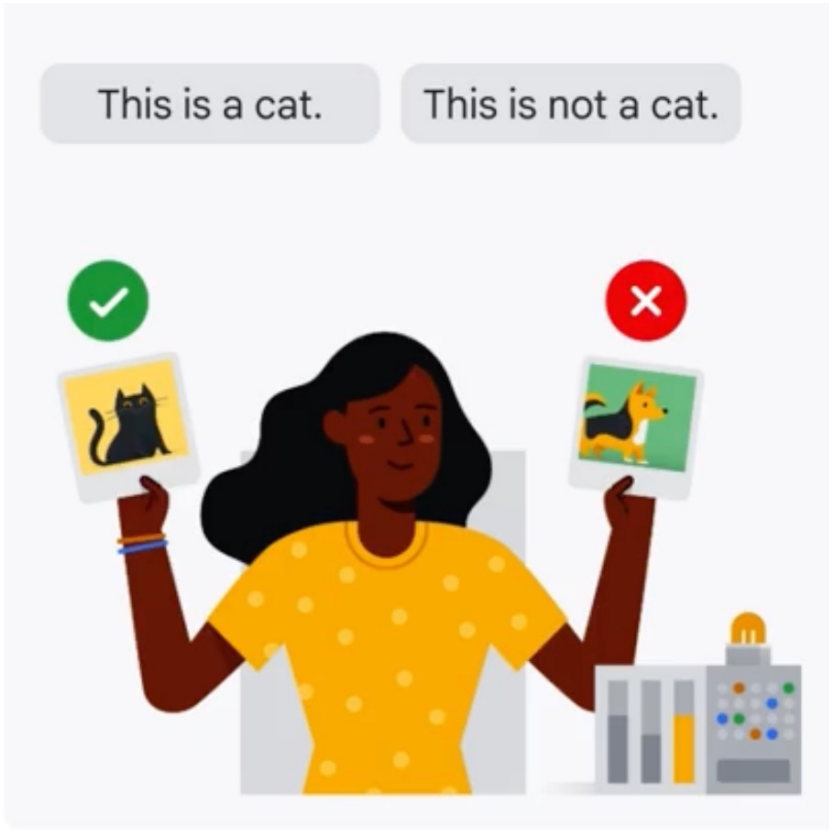
# Knowledge Based Approaches

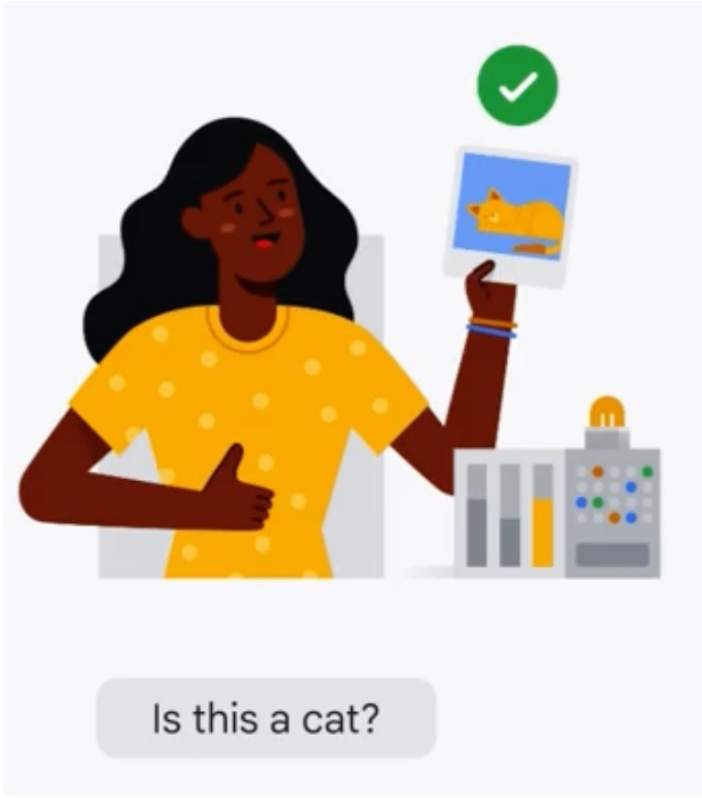
# ML using Artificial Neural Networks
## (One Specific ML Technique)

Traditional programming




Wave of neural networks | ~2012




Credit: Introduction to Generative AI, Google, https://www.youtube.com/watch?v=G2fqAlgmoPo

# History of AI—Data Driven Approaches

1990— 2012: Statistical approaches + subfield expertise

- Resurgence of probability, focus on uncertainty
- General increase in technical depth
- Agents and learning systems…

"AI Spring"?



2012 *onwards*: Lots of Excitement

- **Big data, machine learning, deep learning**
- AI used in many industries

# What is behind modern success of AI?

**Big Data**



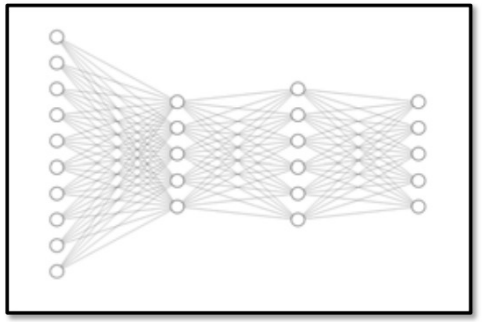
Source: DARPA/ Public domain

**Compute Power**



Source: DARPA/ Public domain

**Machine Learning Algorithms**



## *Modern AI and Data-Driven Methods*



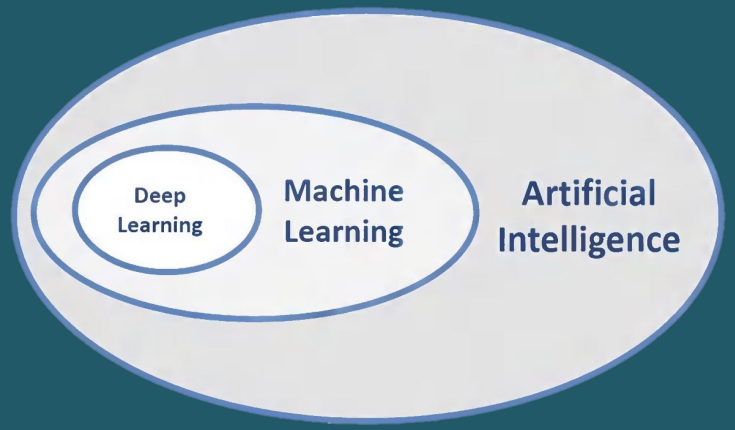






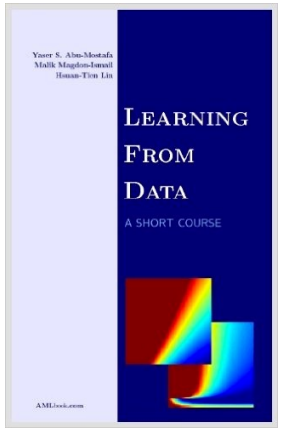



# Machine Learning



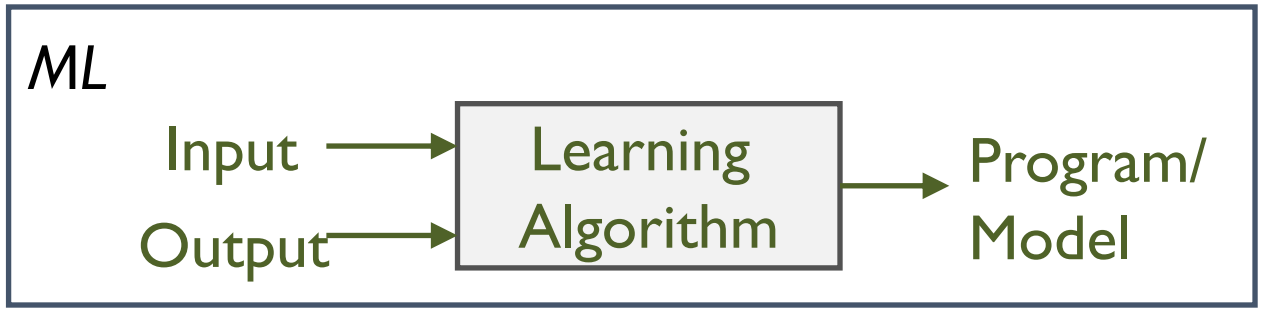

*Machine learning is a field of computer science that gives computers the ability to learn without being explicitly programmed*

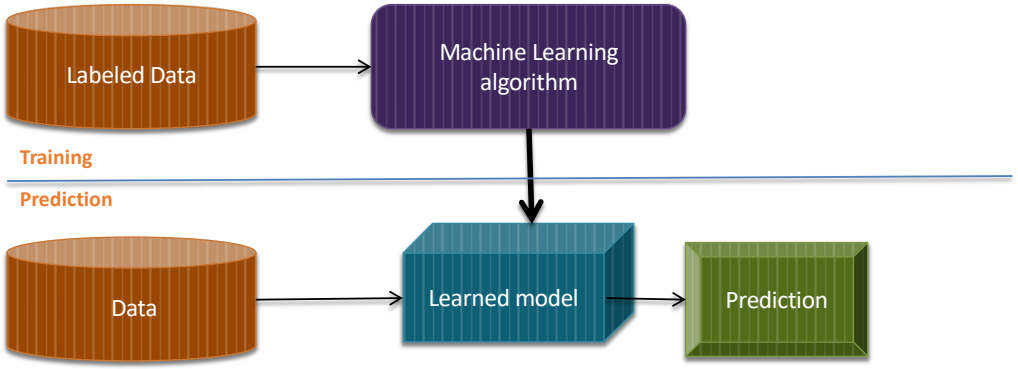
# ML: The Art of Learning from Data




***Essence of ML:***

– There is a pattern in data, however it is complex and difficult to articulate
– We cannot pin it down mathematically (i.e., it is too complex for that)
– Let the computer learn a model for the concept itself using algorithms.



# *Three steps of ML*

Machine Learning

**= Looking for a function**

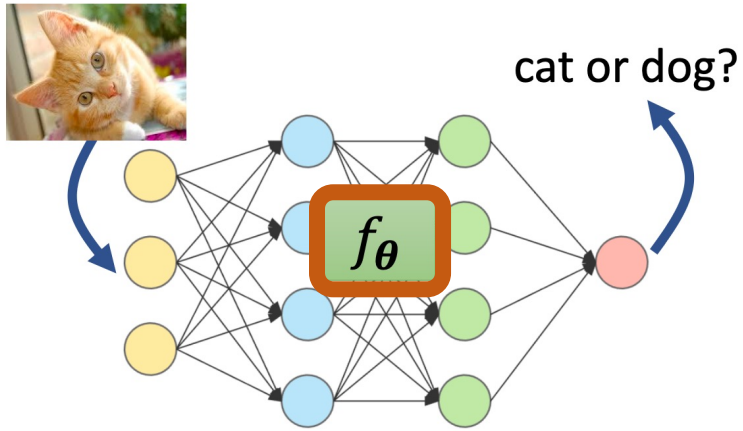
Dog-Cat Classification

$f($  $) =$ "cat"

*Step 1*: Function with unknown

*Step 2*: Define loss function

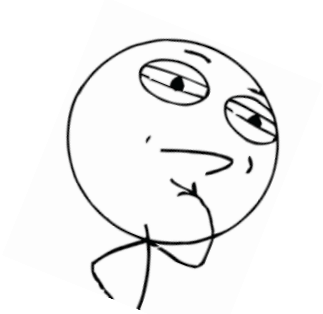
*Step 3*: Optimization



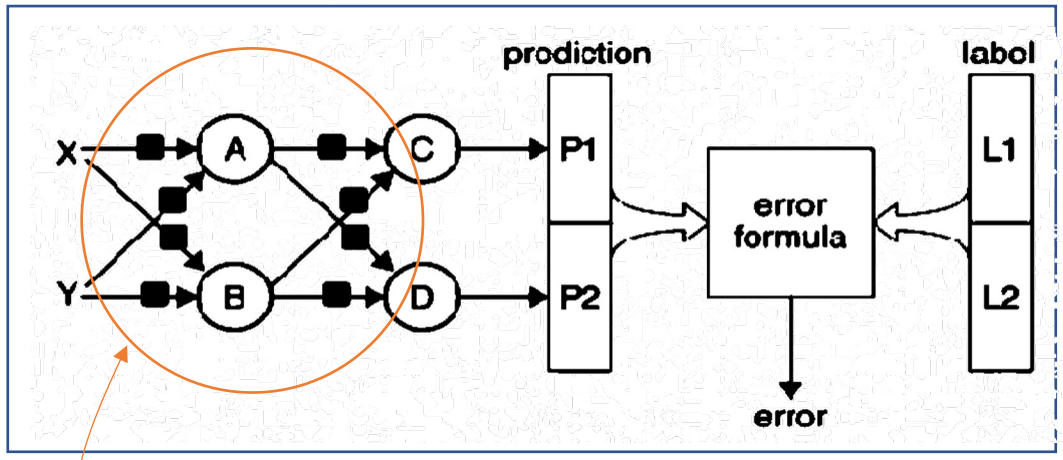

Weights and biases of neurons are unknown parameters (*learnable*).

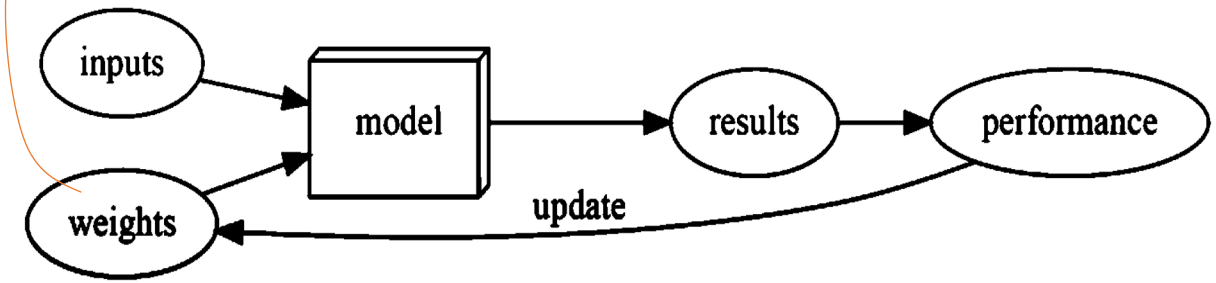
Using $\boldsymbol{\theta}$ to represent the unknown parameters.

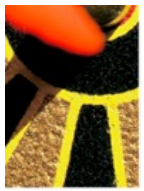
# How does ML learn model parameters automatically?



*How do we learn the right weights (model parameters)?*



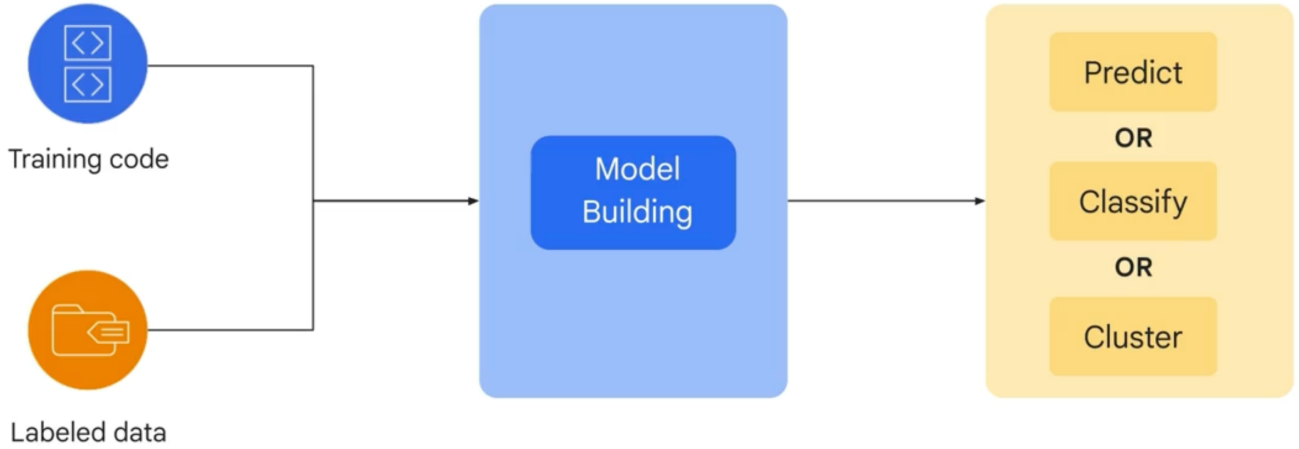



**The Unreasonable Effectiveness of Data**

Alon Halevy, Peter Norvig, and Fernando Pereira, *Google*

*We keep on tweaking the weights for all data; if we do it for sufficiently "big data", our model starts making good predictions*

# What can use ML for?




**Regression**

**Classification**

*Supervised Learning:*
*Typically trained using a labeled dataset (input and correct answer)*

**Clustering**

*Unsupervised Learning:*
*Works even when there is no labelled data*

*Traditionally, ML models are task specific and data specific and perform poorly outside their specific domain*

## ML Input/ Output can be of diverse types



Machine translation: English sentence → Japanese sentence



Dialogue: conversational history → next utterance

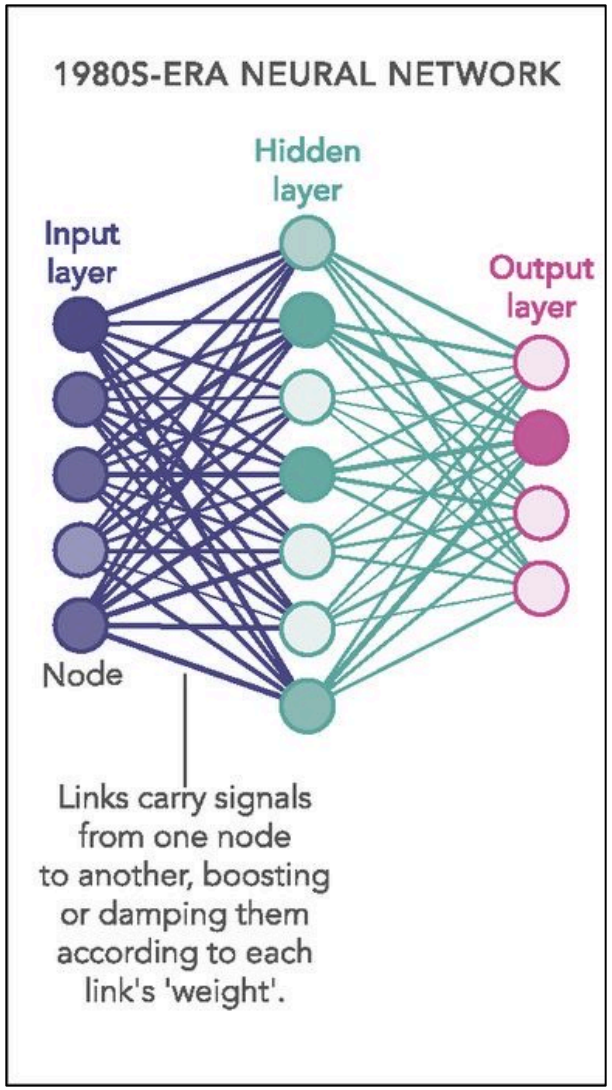


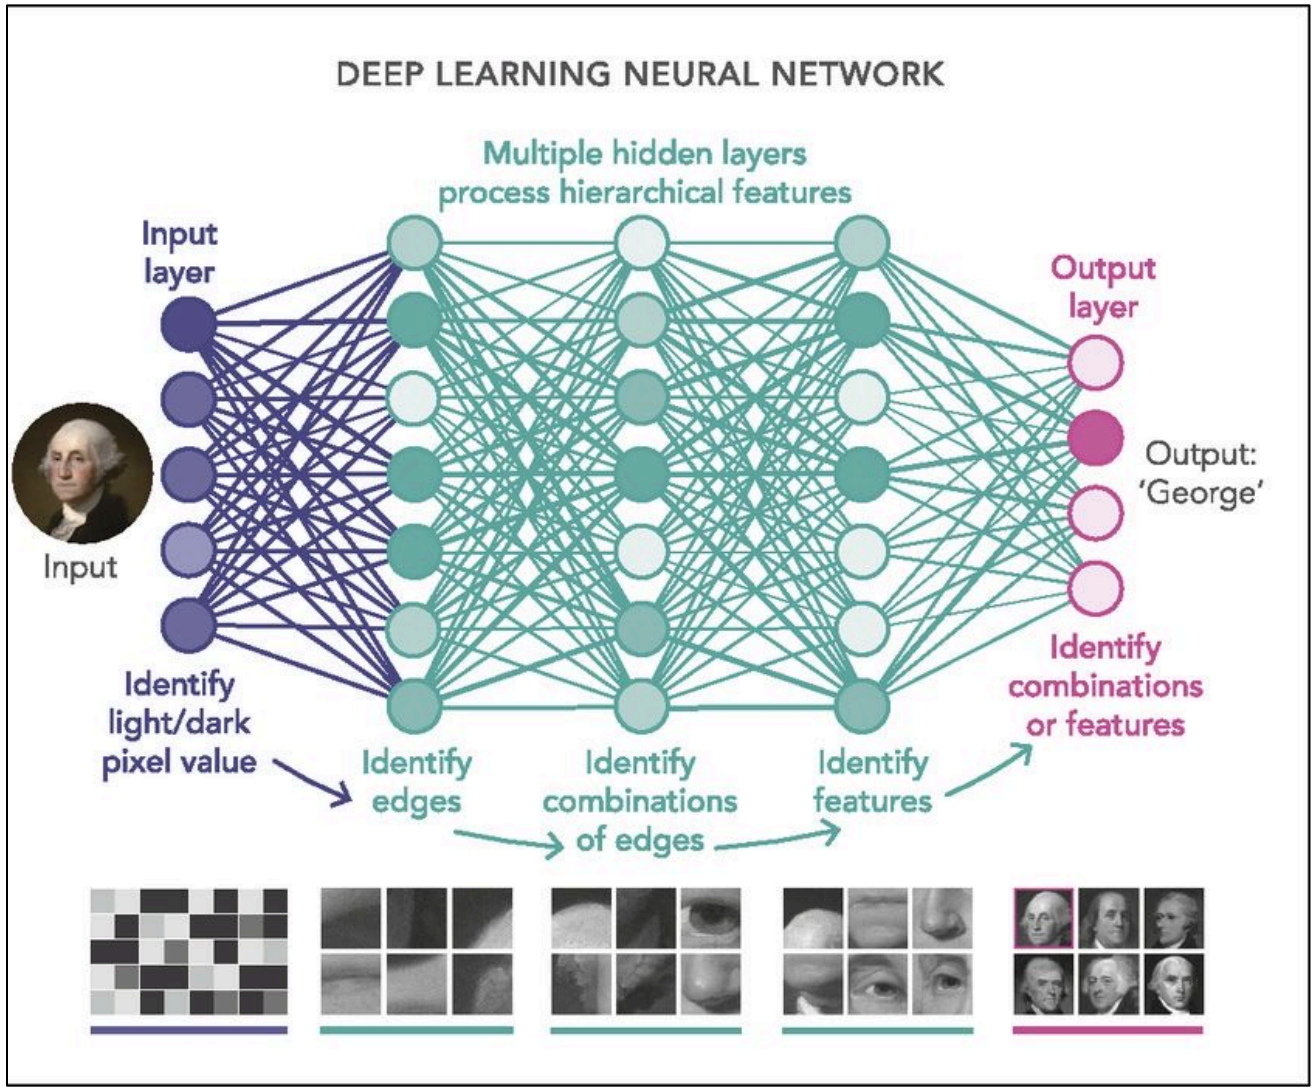
Image segmentation: image → segmentation



Image captioning: image → sentence describing image

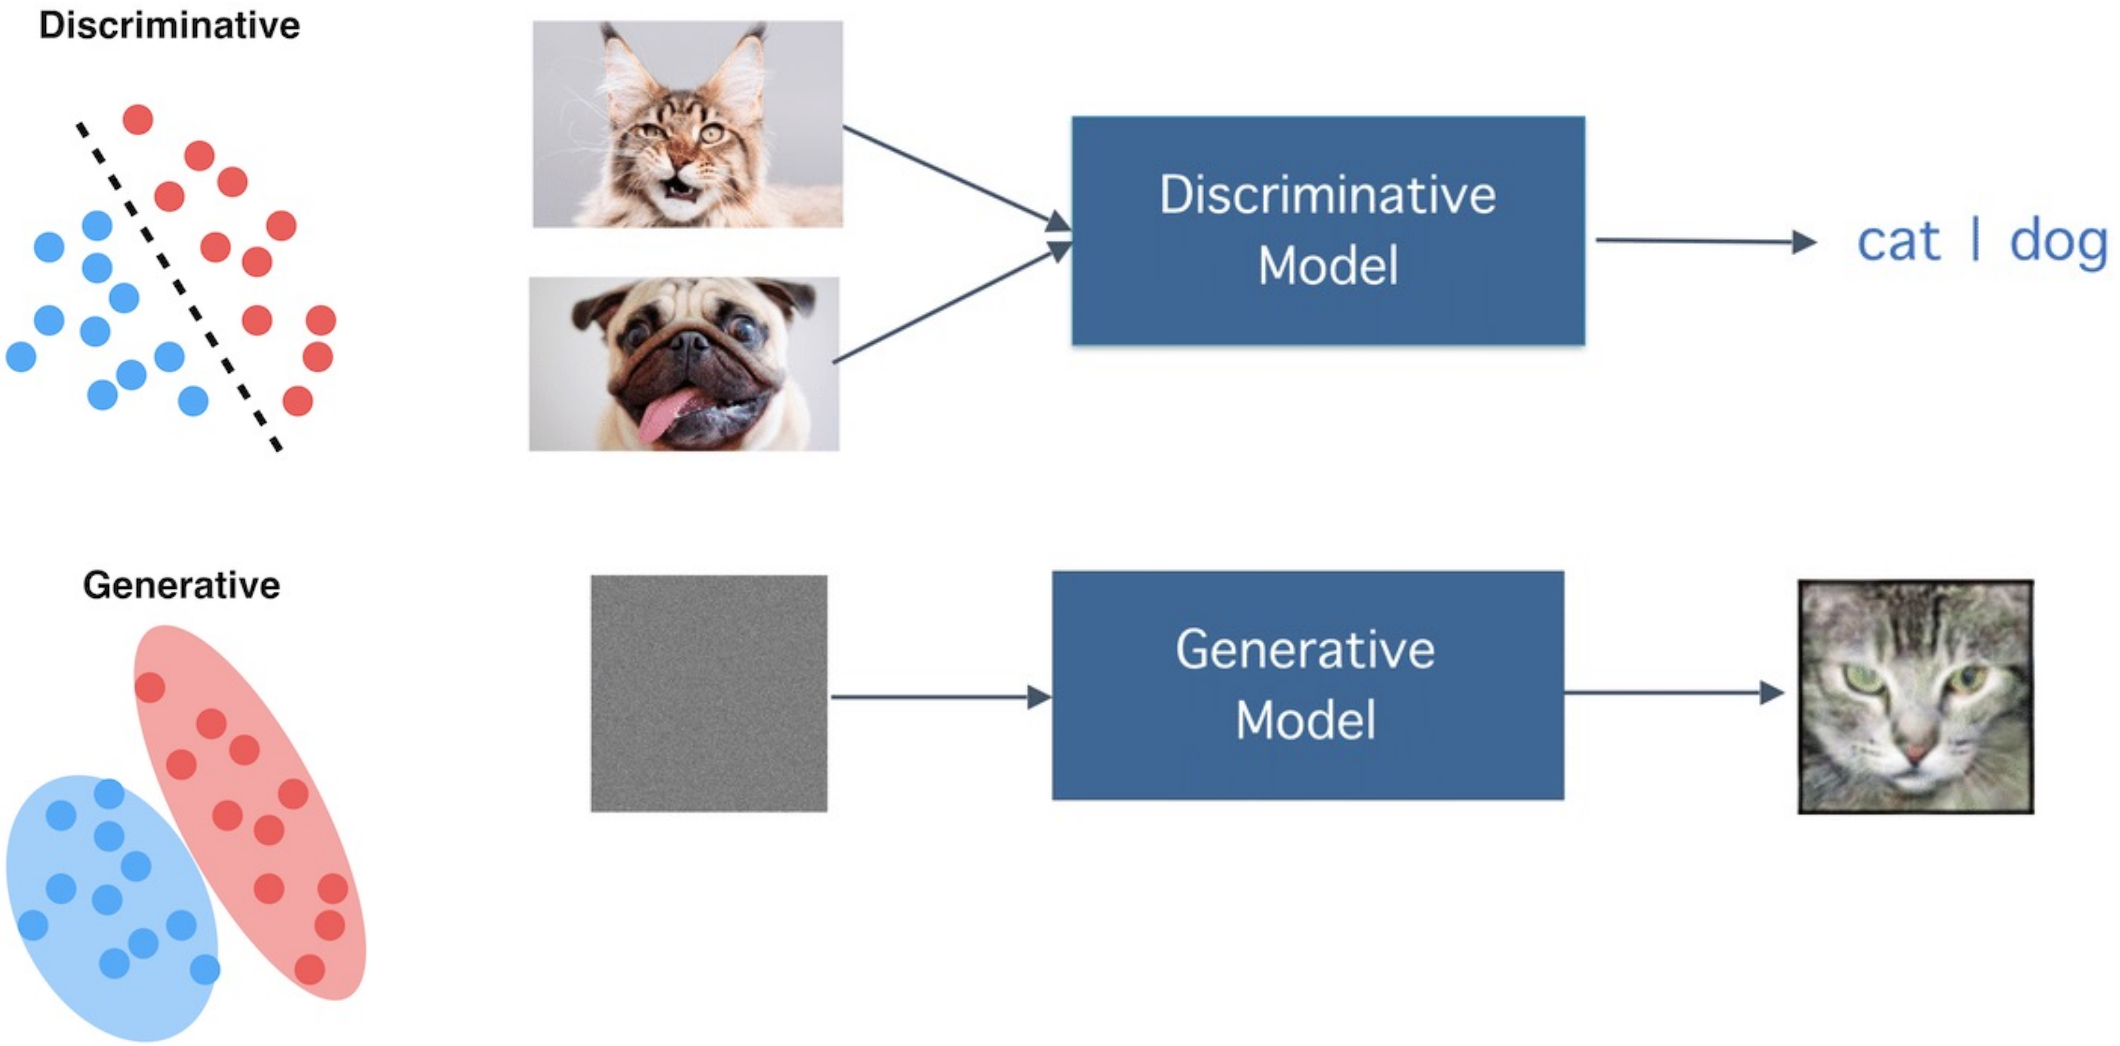
# Deep Learning (Multi-Layer Learning)



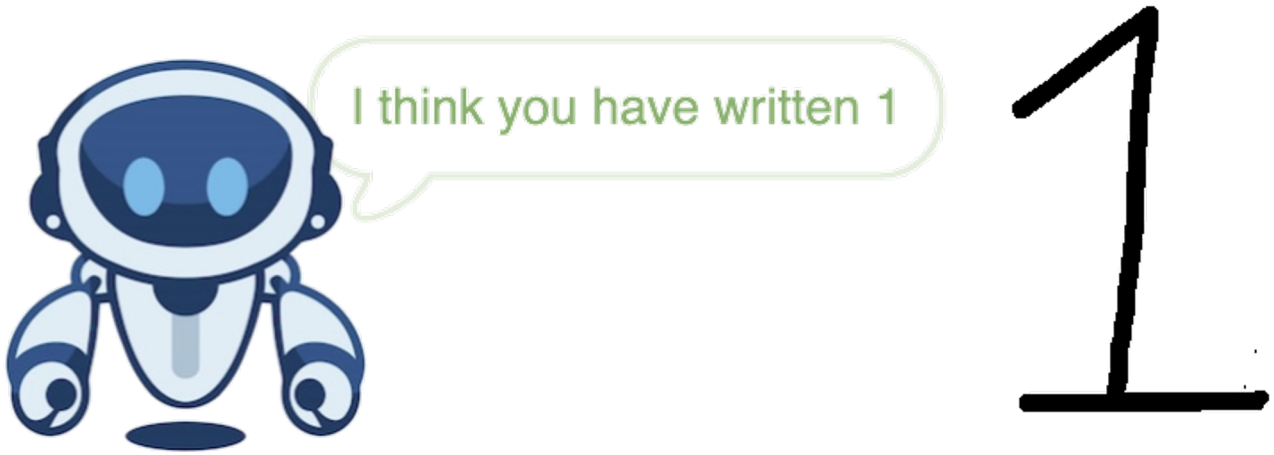

M. Mitchell Waldrop PNAS 2019;116:4:1074-1077

# Generative vs. Discriminative Models



# Hands-on Activity

*Can you train ML to recognize digits?*




**https://machinelearningforkids.co.uk/**

# For Further Viewing: Intro to AI




https://www.youtube.com/watch?v=t9gmyvf7JYo&t=1s

# Thank you!
# Questions/ Discussion

# Day 1

# *Introduction to Generative AI*

A Fundamentals of AI

B Generative AI Introduction

C Introduction to ChatGPT

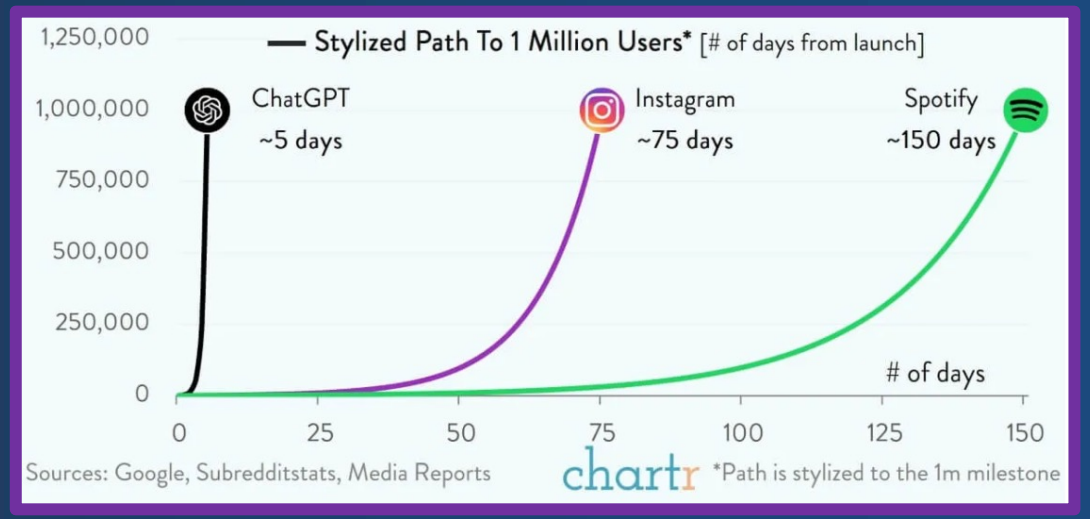
# Generative AI Intro

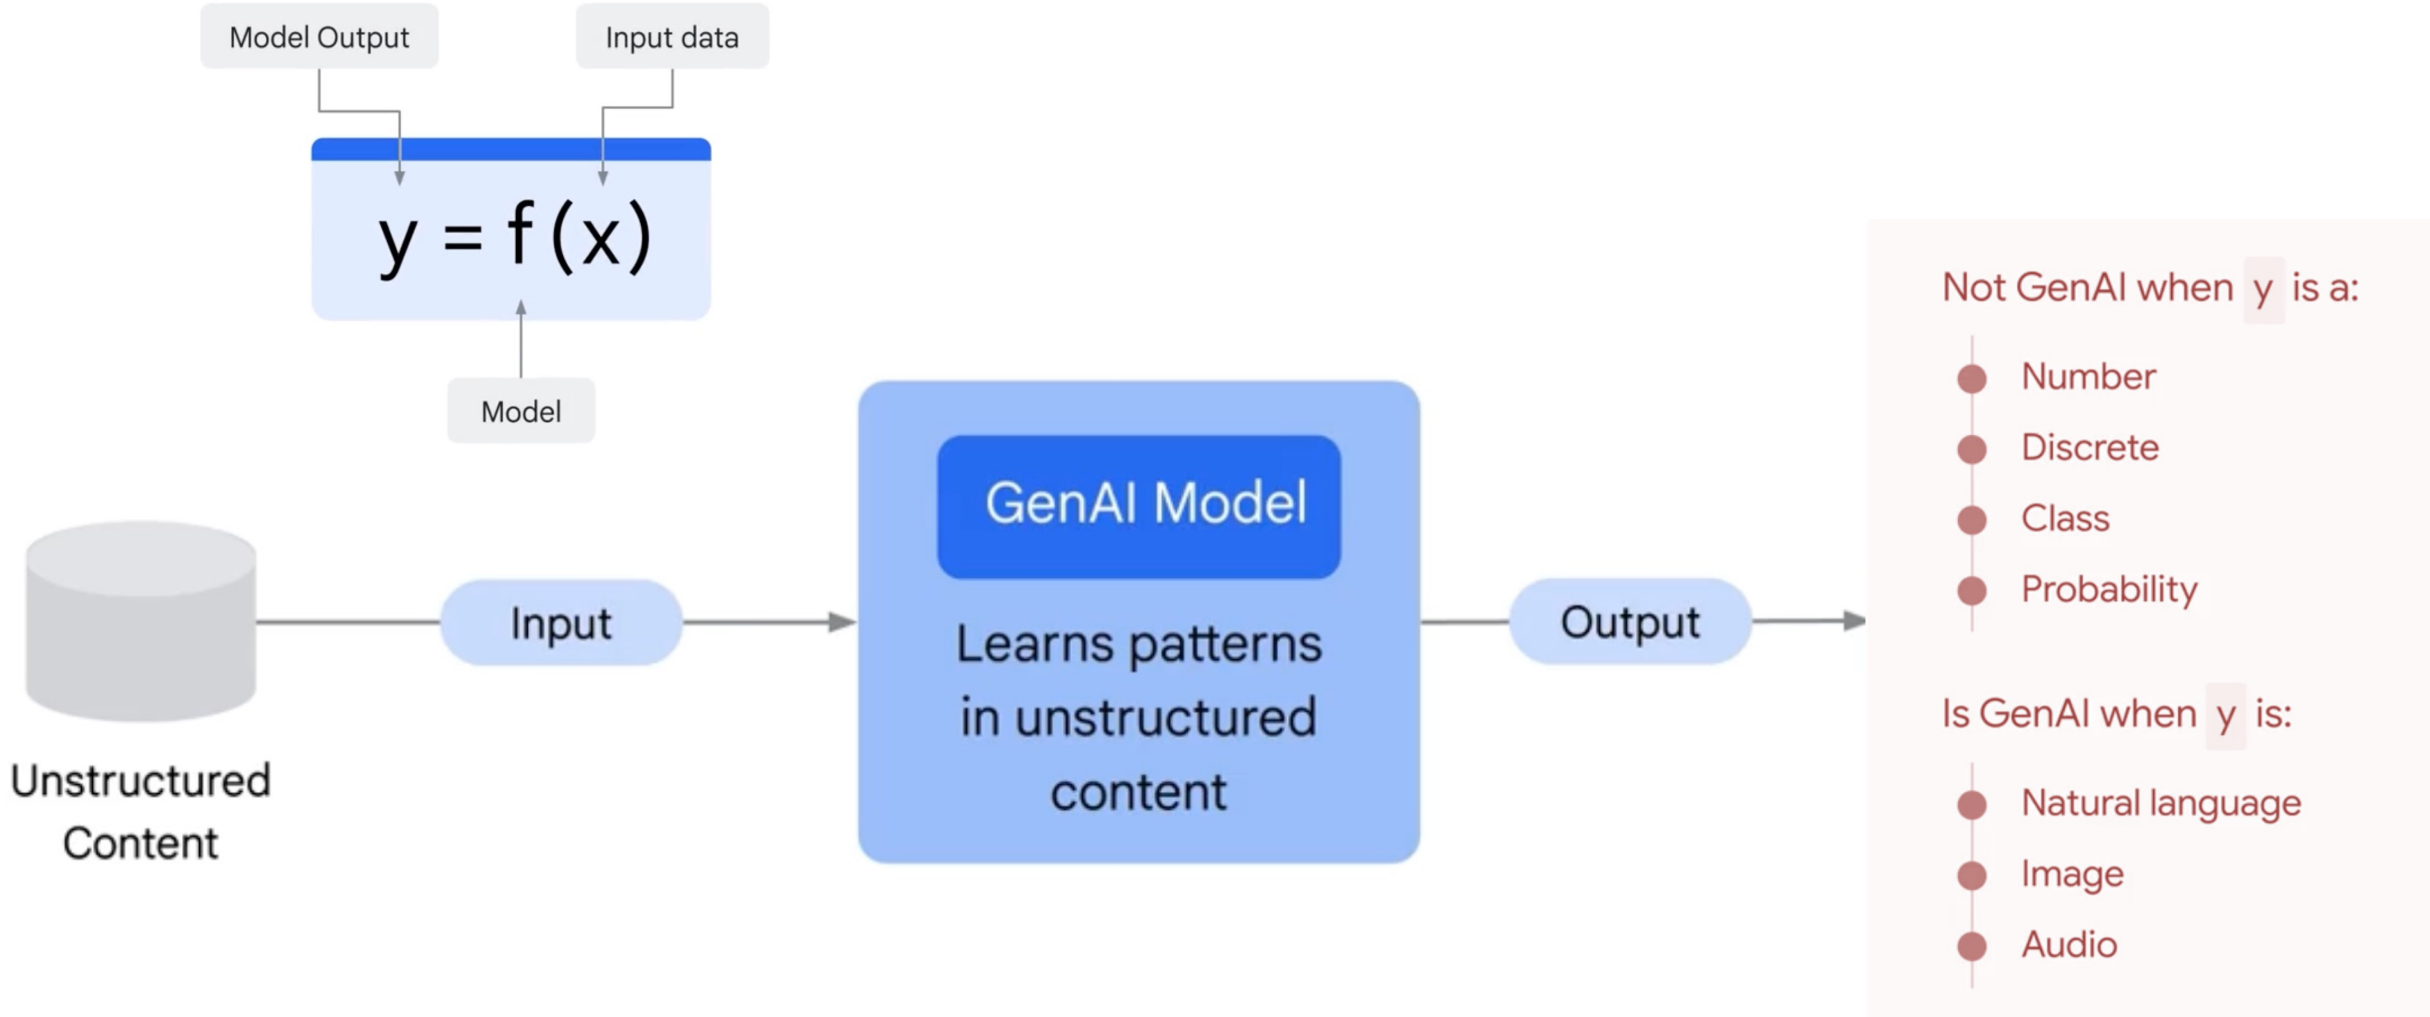
*Image Generator*

*Speech Synthesizer*

*Text Generator*

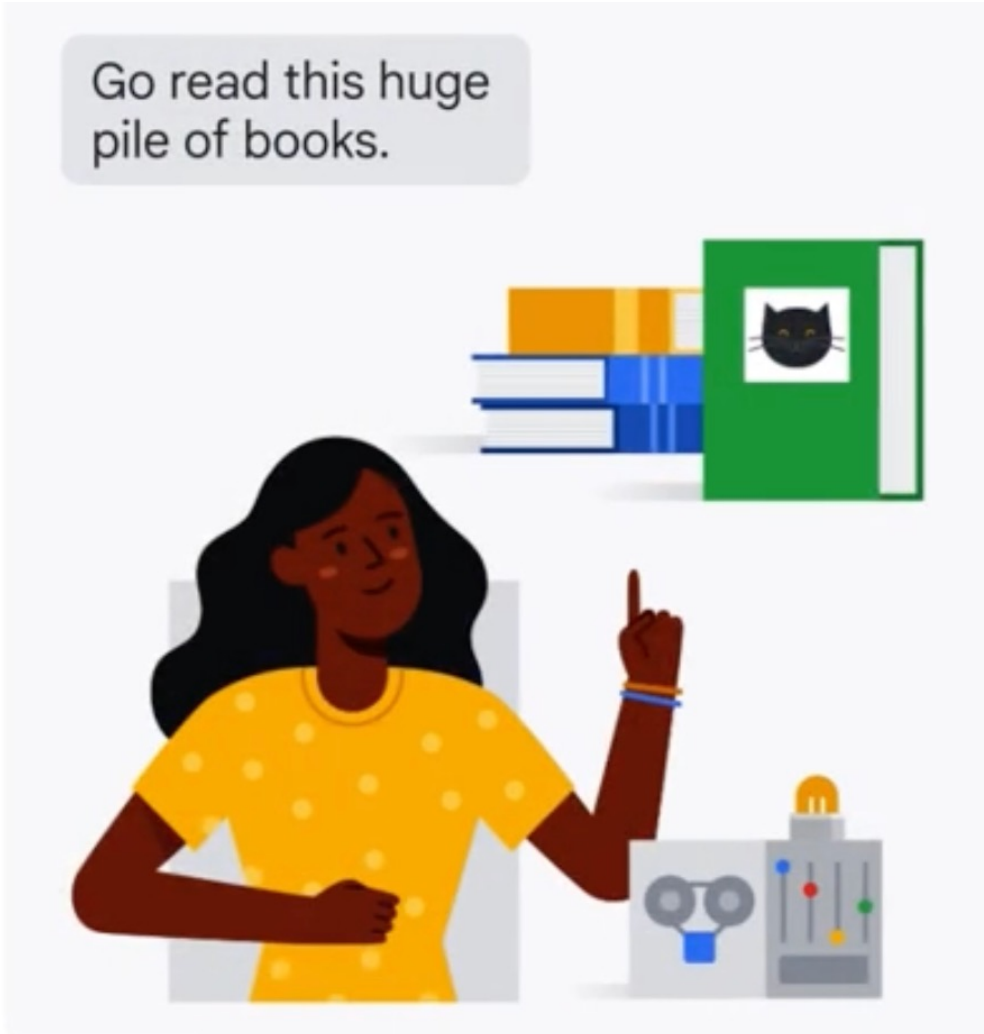
*Conversational AI (ChatBot)*

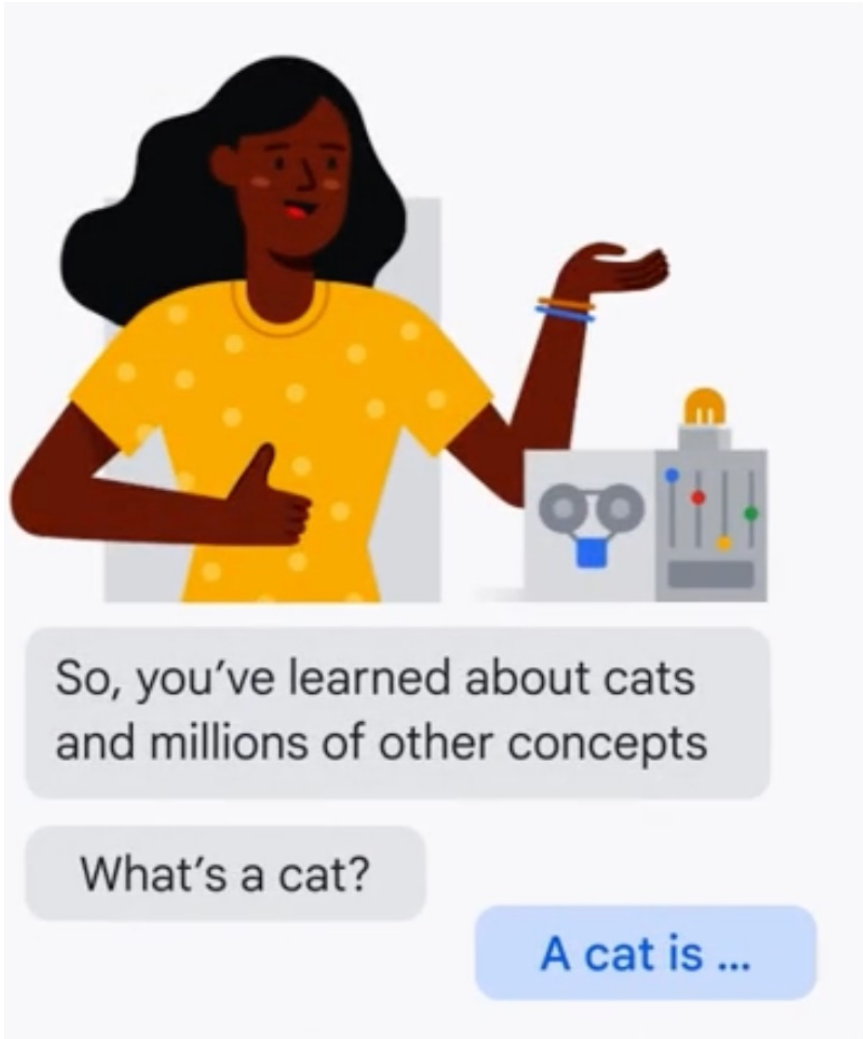





Introduction to Generative AI, Google, https://www.youtube.com/watch?v=G2fqAlgmoPo

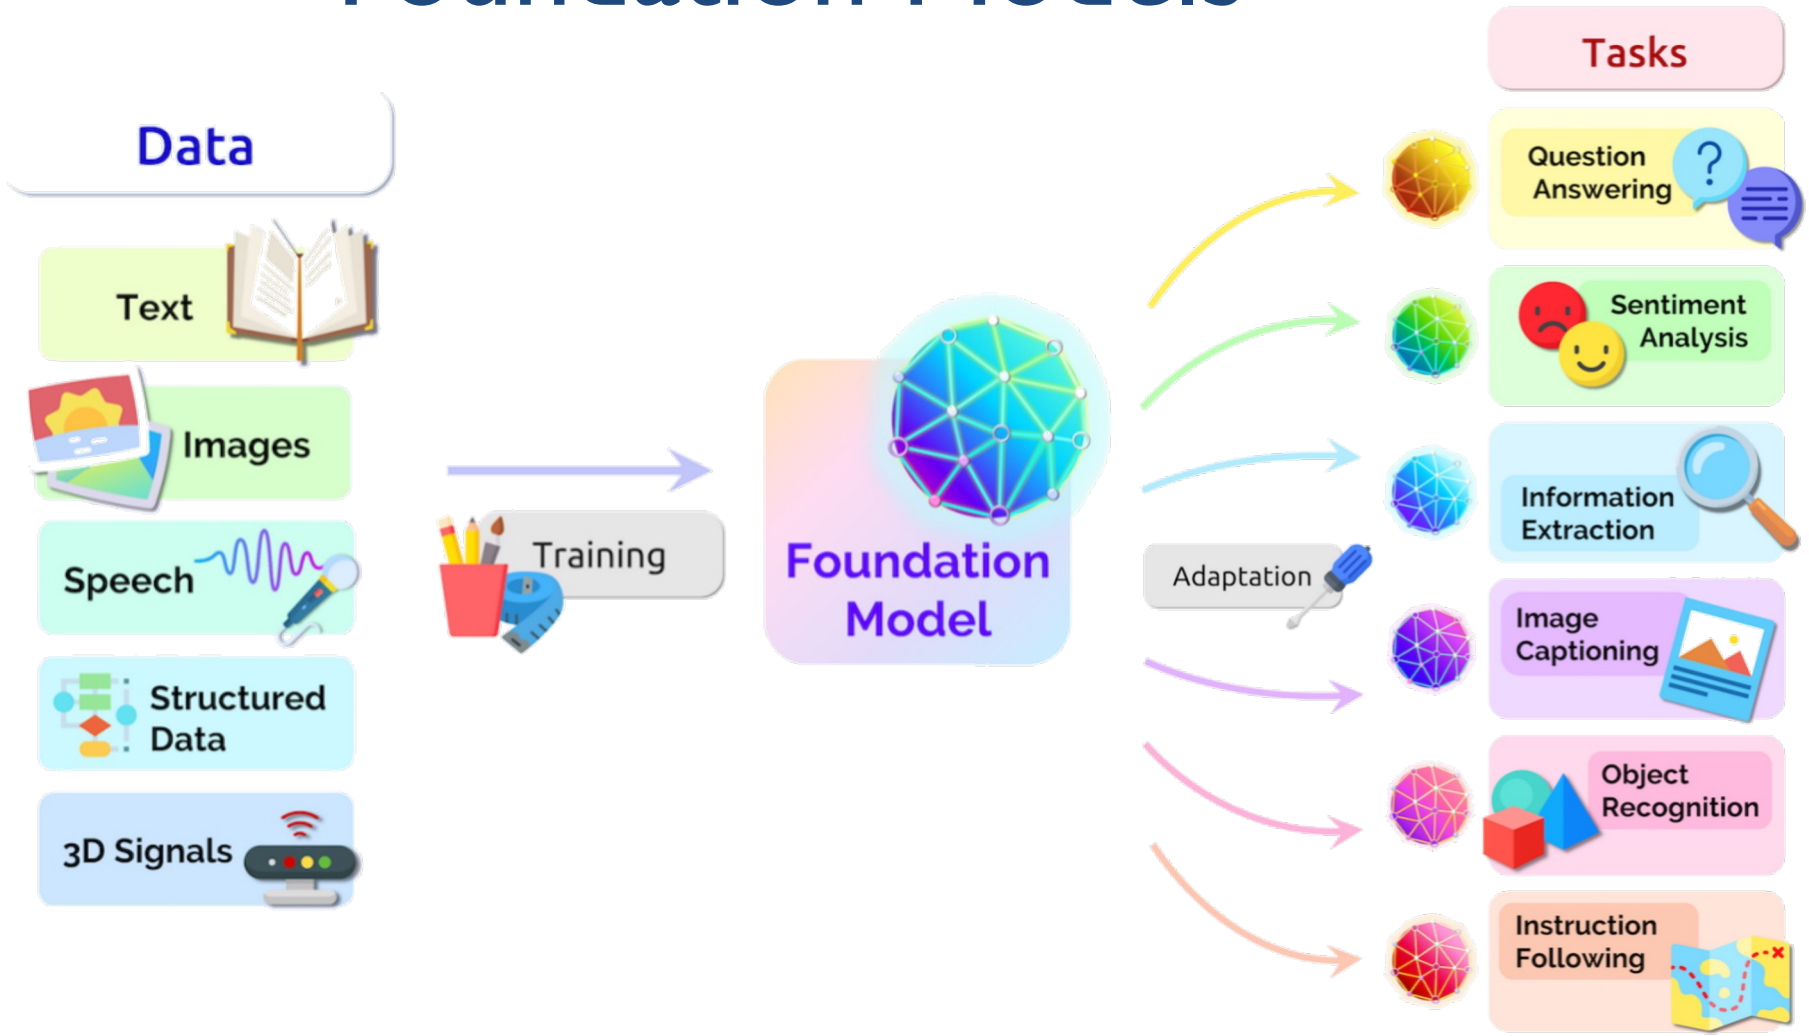
# Generative language models | LaMDA, PaLM, GPT, etc.




Introduction to Generative AI, Google, https://www.youtube.com/watch?v=G2fqAlgmoPo

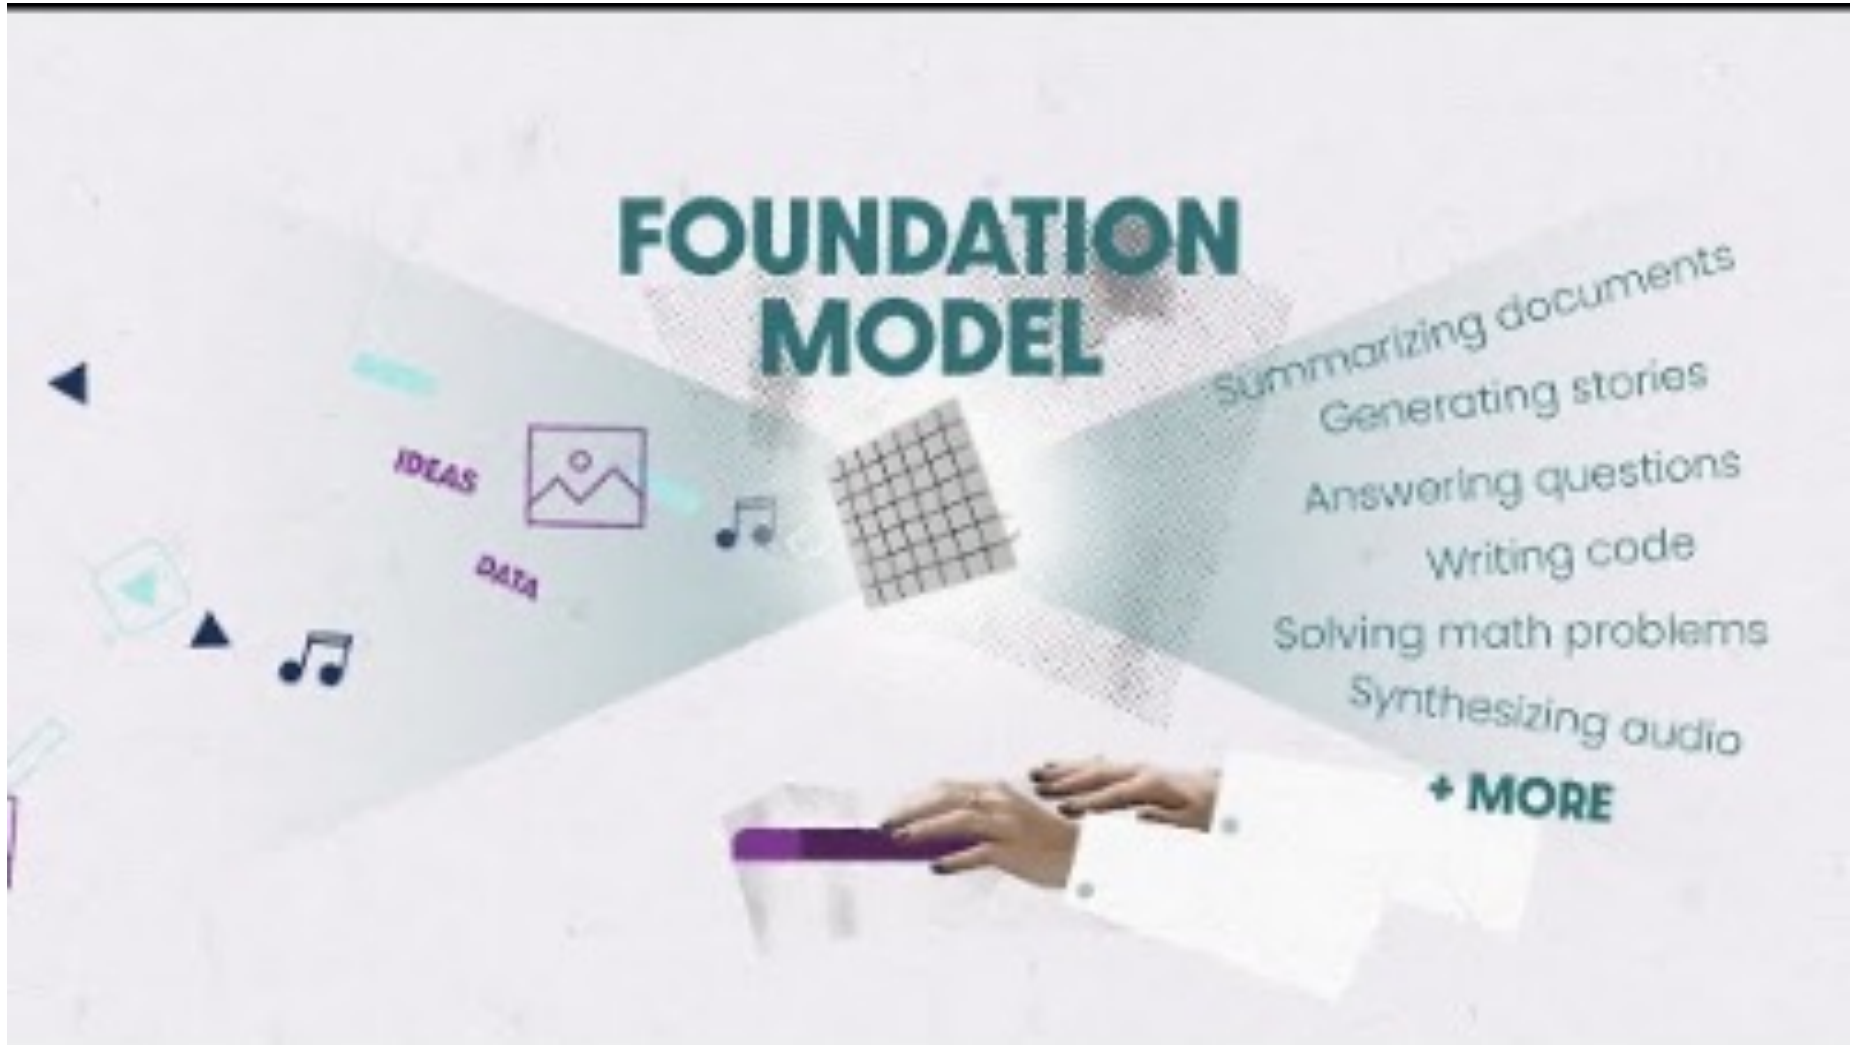
# Foundation Models




Effort towards ***general-purpose multi-task AI***

*that can be adapted to multiple downstream tasks*

# For Further Viewing: Foundation AI Models



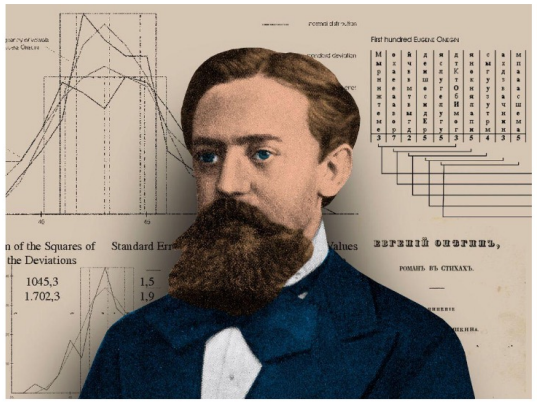

https://www.youtube.com/watch?v=kK3NmQT241w

# Large Language Models

Google

what is the weather in

- what is the weather in **lahore**
- what is the weather in **islamabad**
- what is the weather in **karachi**
- what is the weather in **naran kaghan**
- what is the weather in **gujranwala**
- what is the weather in **pakistan**
- what is the weather in **murree**
- what is the weather in **lahore today**
- what is the weather in **rawalpindi**
- what is the weather in **faisalabad**

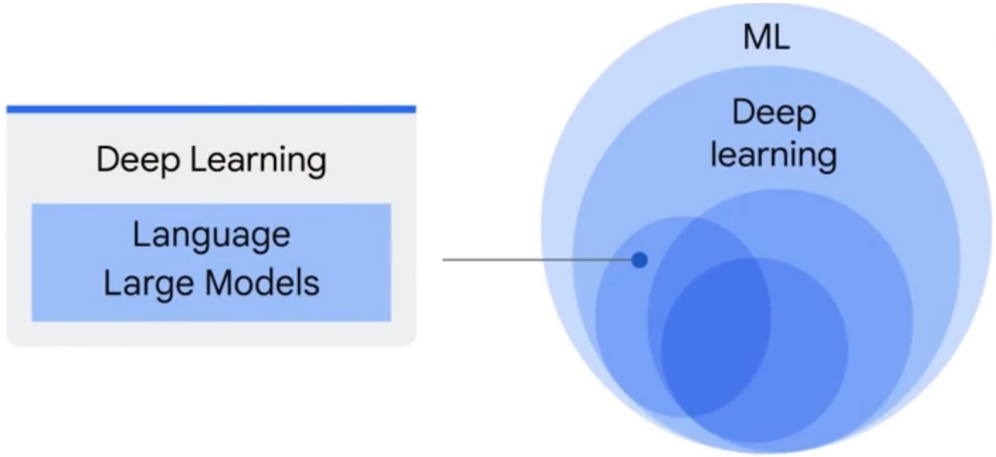




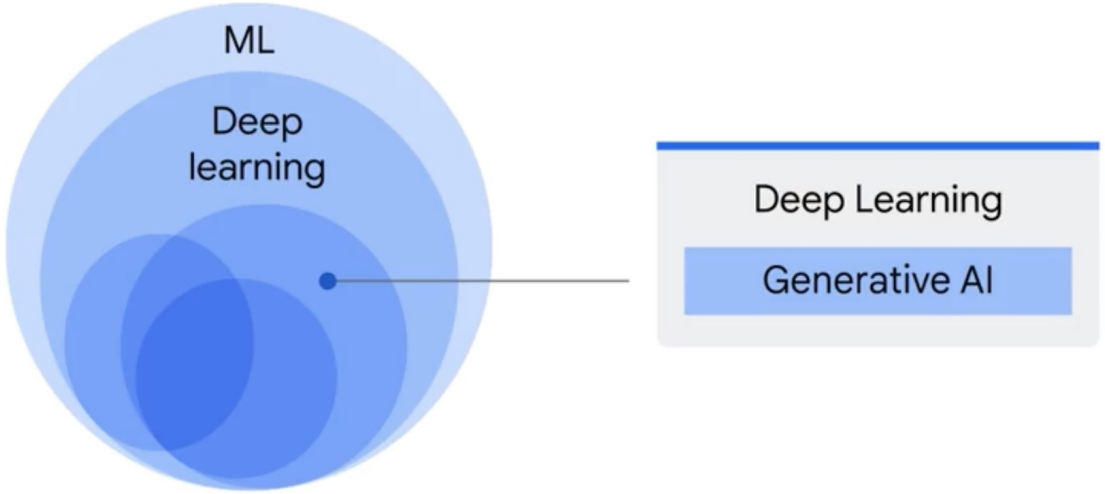

In 1913, Russian mathematician Andrey Markov counted letters from "Eugene Onegin" and showed that the chance of a letter appearing depends on the letter before it.

# Large Language Models (LLMs)

Large Language Models (LLMs) are also a subset of Deep Learning




Generative AI is a subset of Deep Learning




Introduction to Generative AI, Google, https://www.youtube.com/watch?v=G2fqAlgmoPo

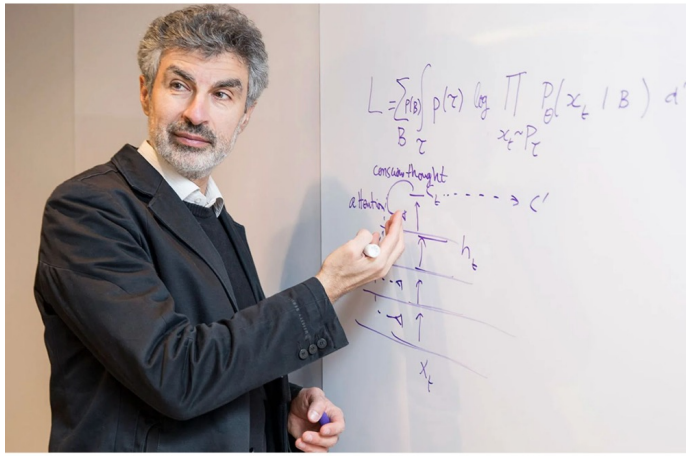
# Statistical *Natural Language Processing (NLP)* Models

***Distributional hypothesis:*** words that occur in similar contexts tend to have similar meanings



J.R.Firth 1957

- "You shall know a word by the company it keeps"
- One of the most successful ideas of modern statistical NLP!



Journal of Machine Learning Research 3 (2003) 1137–1155 Submitted 4/02; Published 2/03

**A Neural Probabilistic Language Model**

**Yoshua Bengio** BENGIOY@IRO.UMONTREAL.CA
**Réjean Ducharme** DUCHARME@IRO.UMONTREAL.CA
**Pascal Vincent** VINCENTP@IRO.UMONTREAL.CA
**Christian Jauvin** JAUVINC@IRO.UMONTREAL.CA
*Département d'Informatique et Recherche Opérationnelle*
*Centre de Recherche Mathématiques*
*Université de Montréal, Montréal, Québec, Canada*

*A goal of statistical language modeling is to learn the joint probability function of sequences of words in a language. This is intrinsically difficult because of the curse of dimensionality. … We propose to fight the curse of dimensionality by learning a* ***distributed representation for words*** *…*

# Probabilistic/ Statistical NLP models

N-gram language models

It is 76 F and ____.

$$P(w \mid \text{it is 76 F and})$$

[0.0001, 0.1, 0, 0, 0.002, ..., 0.3, ..., 0]

red sunny

Word-word co-occurrence matrix:

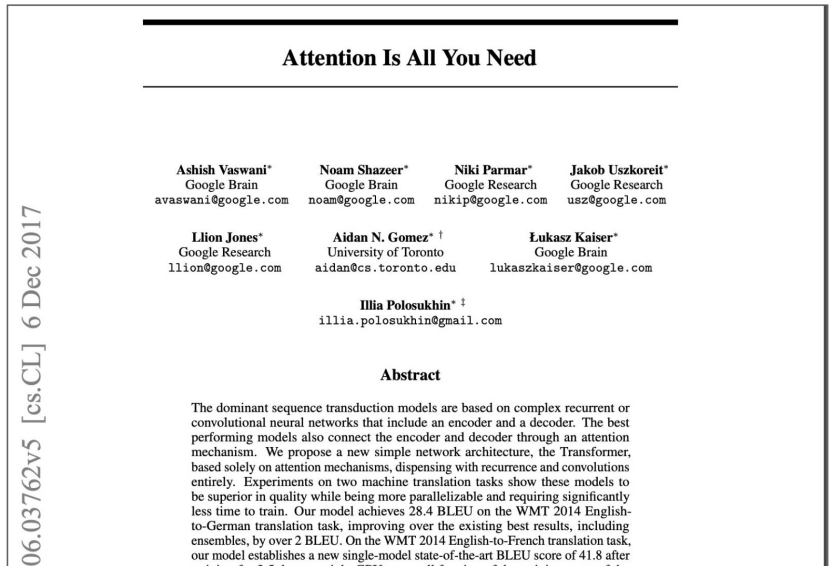
| | aardvark | computer | data | pinch | result | sugar | ... |
|---|---|---|---|---|---|---|---|
| apricot | 0 | 0 | 0 | 1 | 0 | 1 | |
| pineapple | 0 | 0 | 0 | 1 | 0 | 1 | |
| digital | 0 | 2 | 1 | 0 | 1 | 0 | |
| information | 0 | 1 | 6 | 0 | 4 | 0 | |

# Transformer Models

## *(Fundamental Breakthrough in 2017)*

Attention Is All You Need (Transformer paper)

https://arxiv.org/pdf/1706.03762.pdf



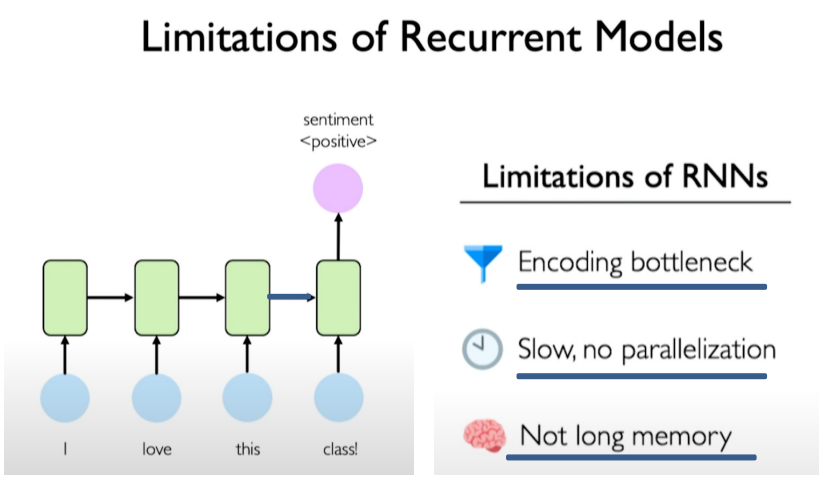
Attention Is All You Need


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
Łukasz Kaiser* Google Brain lukaszkaiser@google.com
Illia Polosukhin* ‡ illia.polosukhin@gmail.com


06.03762v5 [cs.CL] 6 Dec 2017


**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


*Transformers can work with sequences; it can use parallel computing at scale without recurrent connections (which means it can be trained at a large scale) allowing the training of very large language models*

# Attention Module in Transformer Models

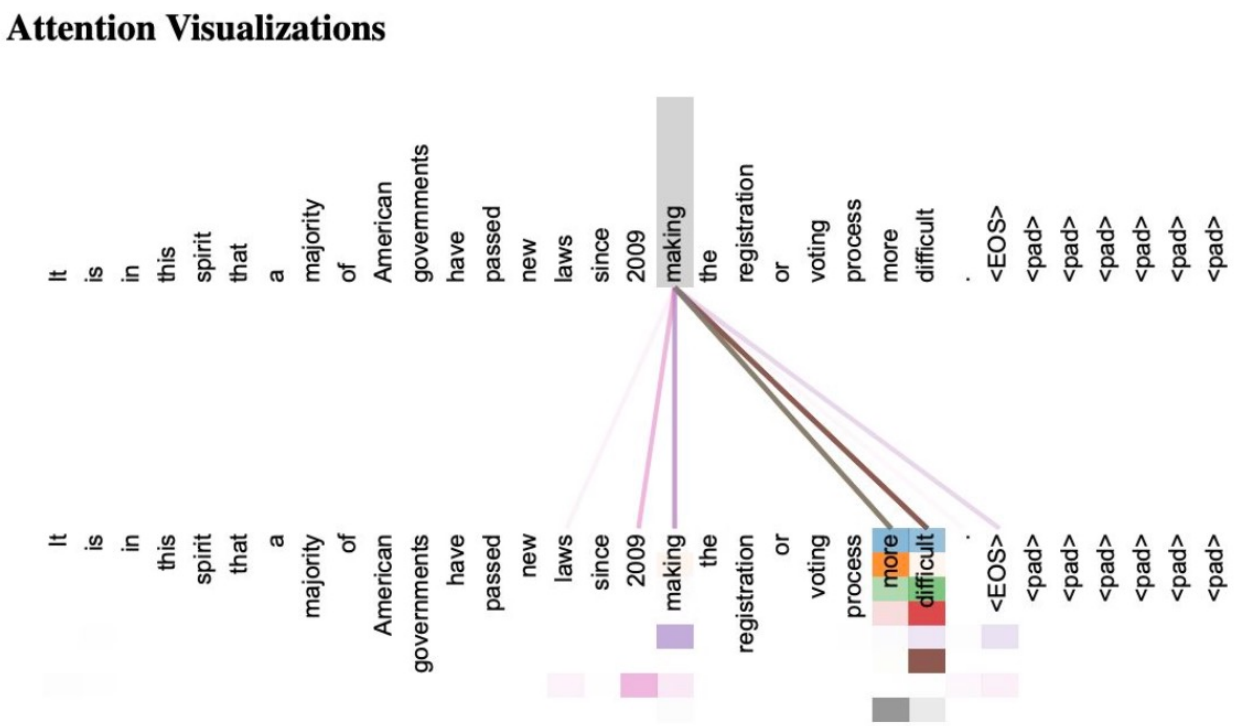
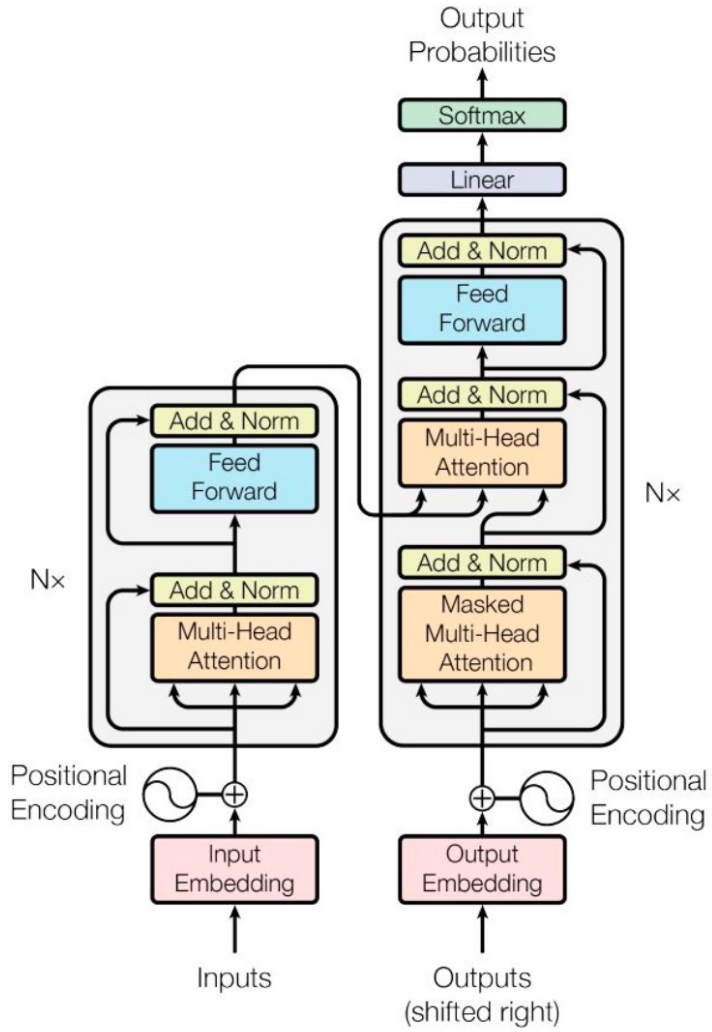
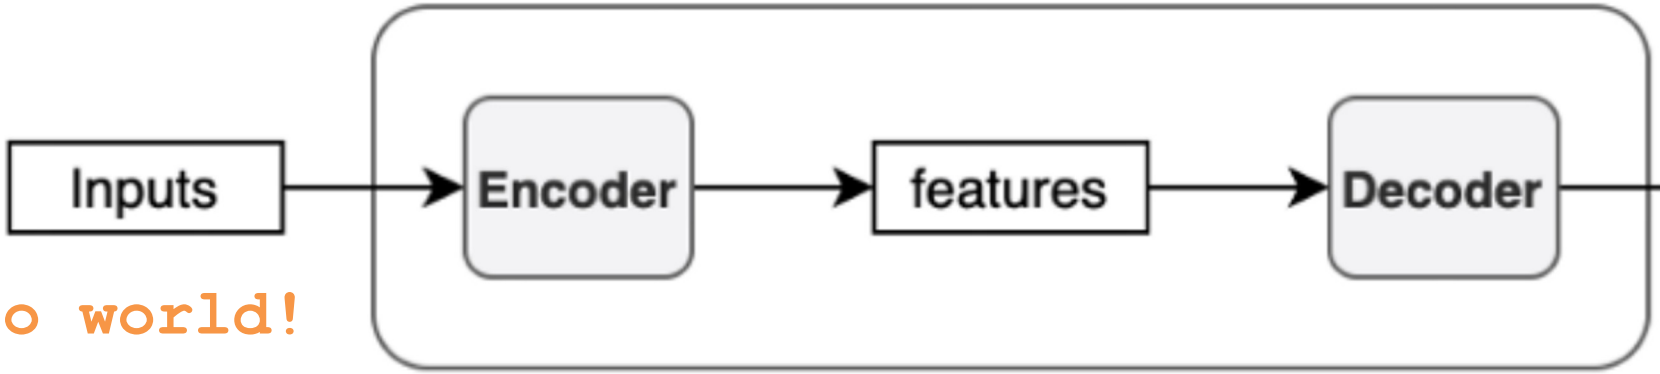



Figure 3: An example of the attention mechanism following long-distance dependencies in the encoder self-attention in layer 5 of 6. Many of the attention heads attend to a distant dependency of the verb 'making', completing the phrase 'making...more difficult'. Attentions here shown only for the word 'making'. Different colors represent different heads. Best viewed in color.

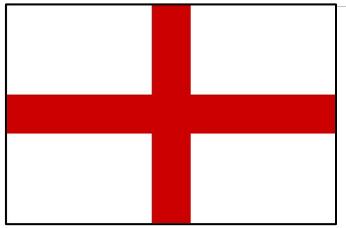


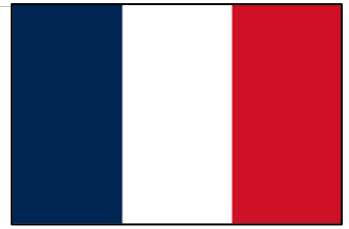
# Transformer has an encoder and decoder




Hello world!

Bonjour le monde!





*Many derivative transformer-based architectures*

# Main Two Transformer Model Architectures

## Autoregressive Models (Decoder Only)

Predict future token, e.g.:

*The joke was funny. She couldn't stop ___.*

NL generation (NLG)

E.g.: GPT architectures

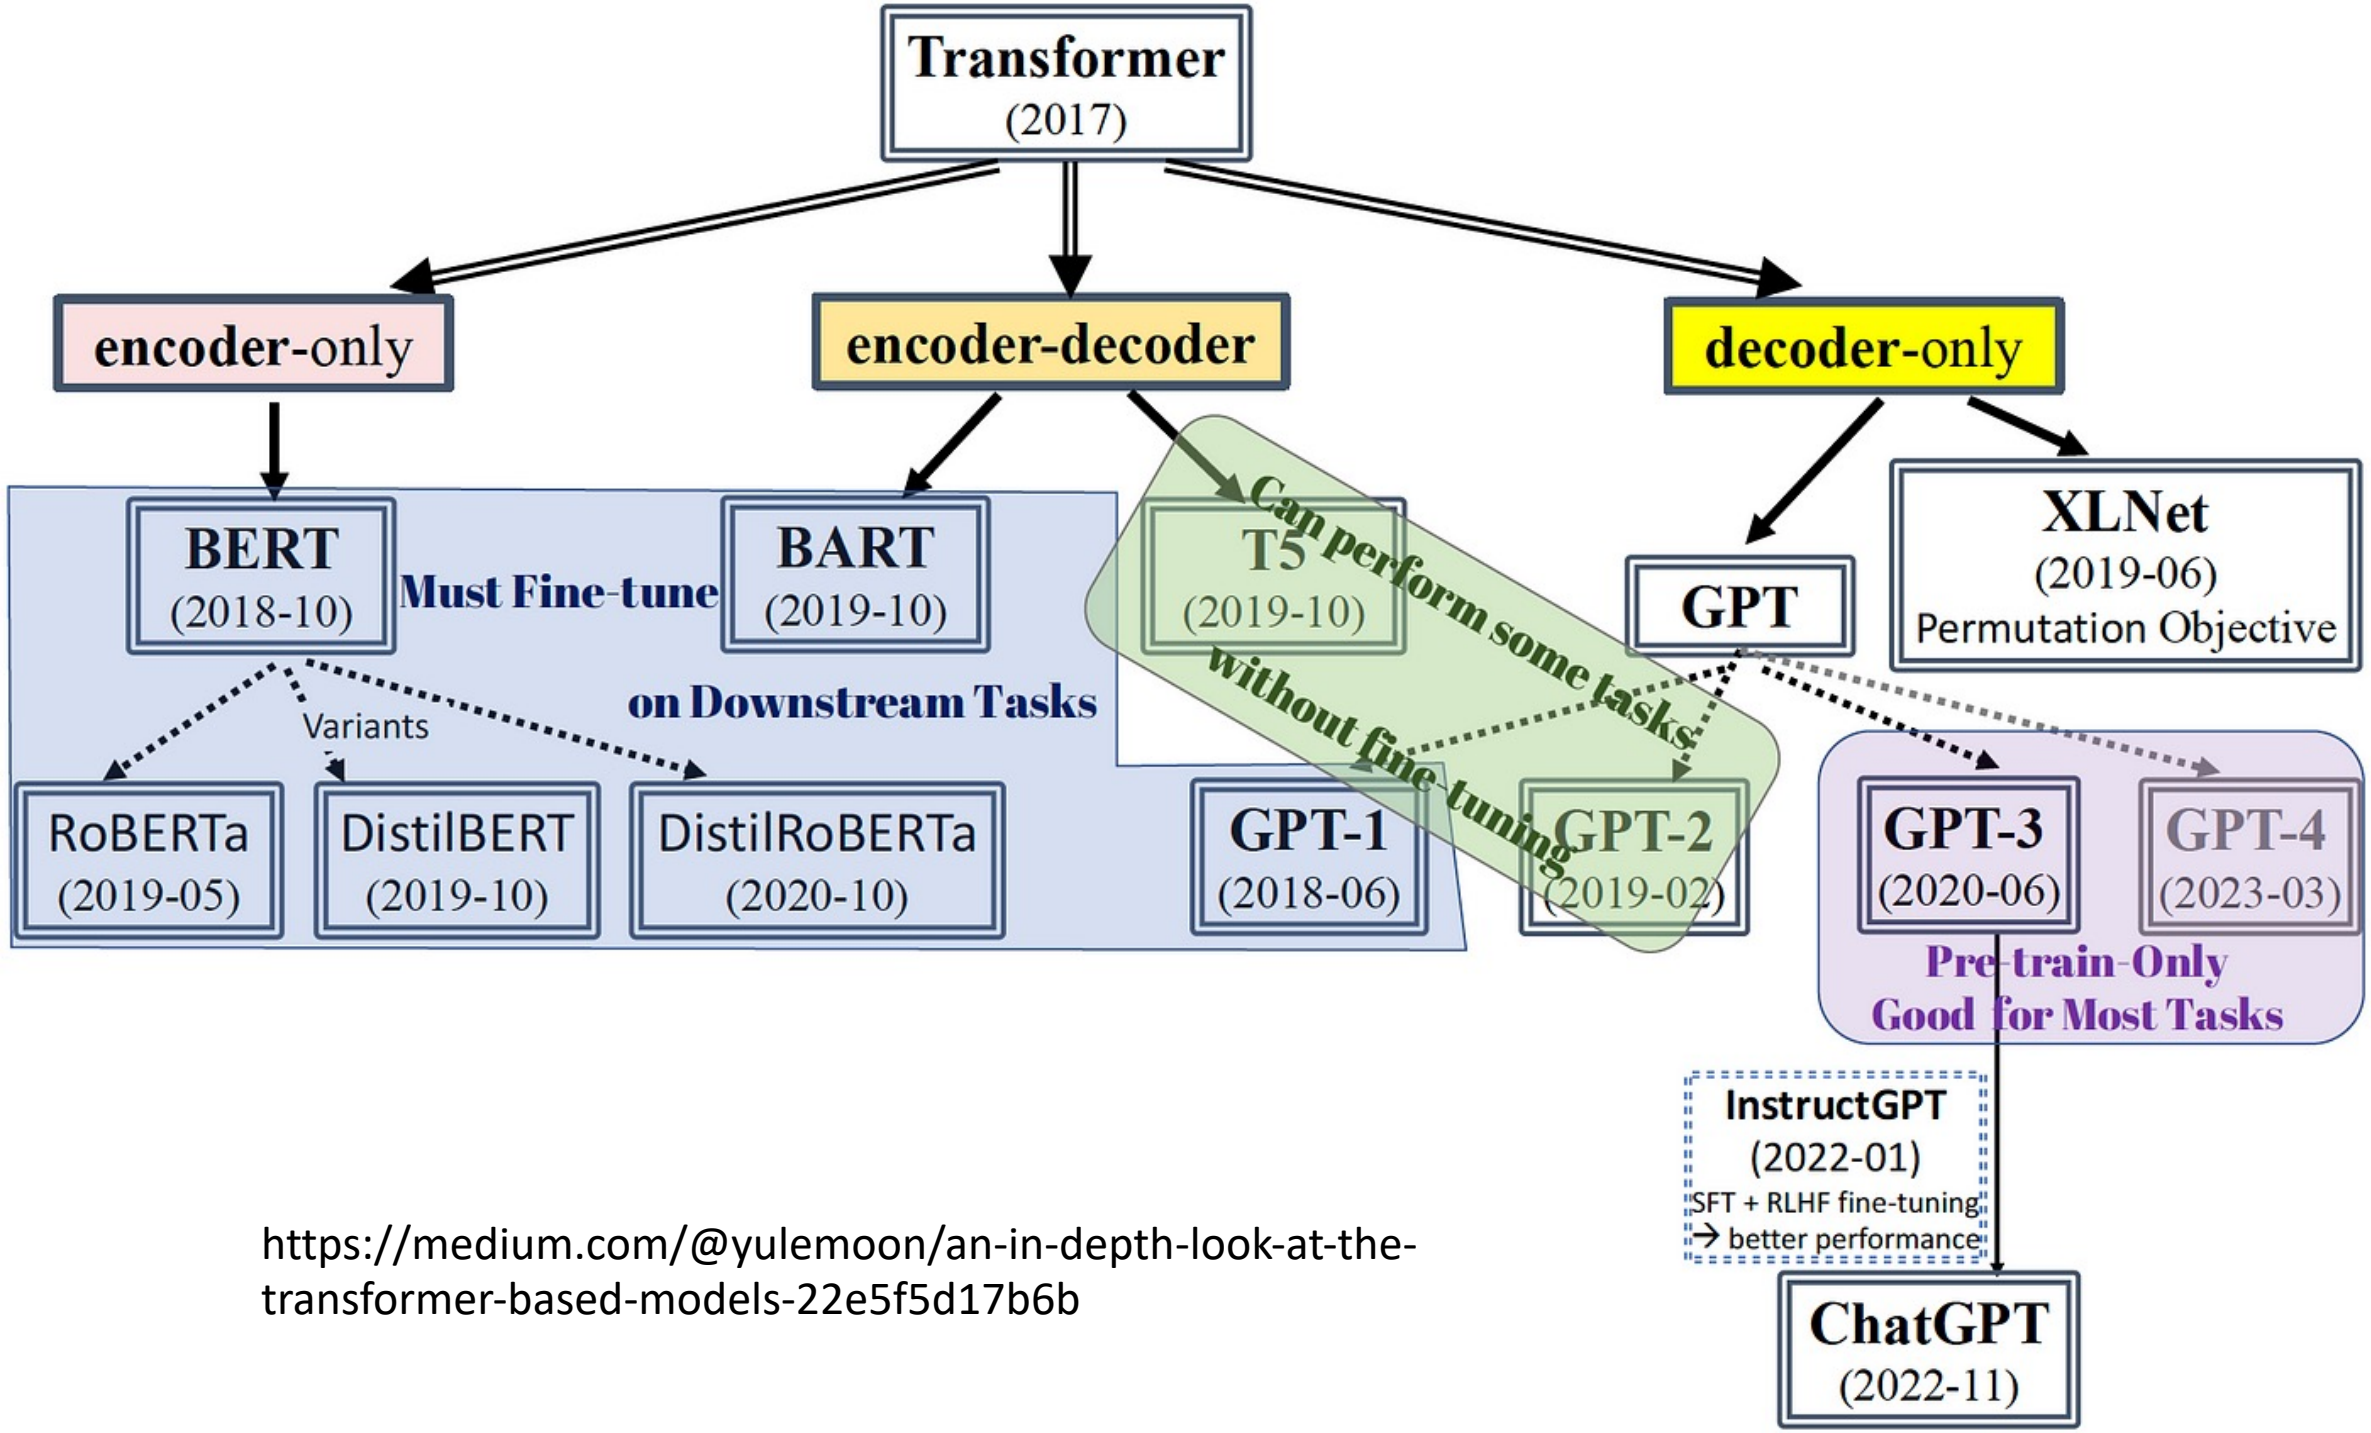
## Autoencoding Models (Encoder Only)

Predict token based on past and future context, e.g.,:

*He ate the entire ___ of pizza.*

NL understanding (NLU)

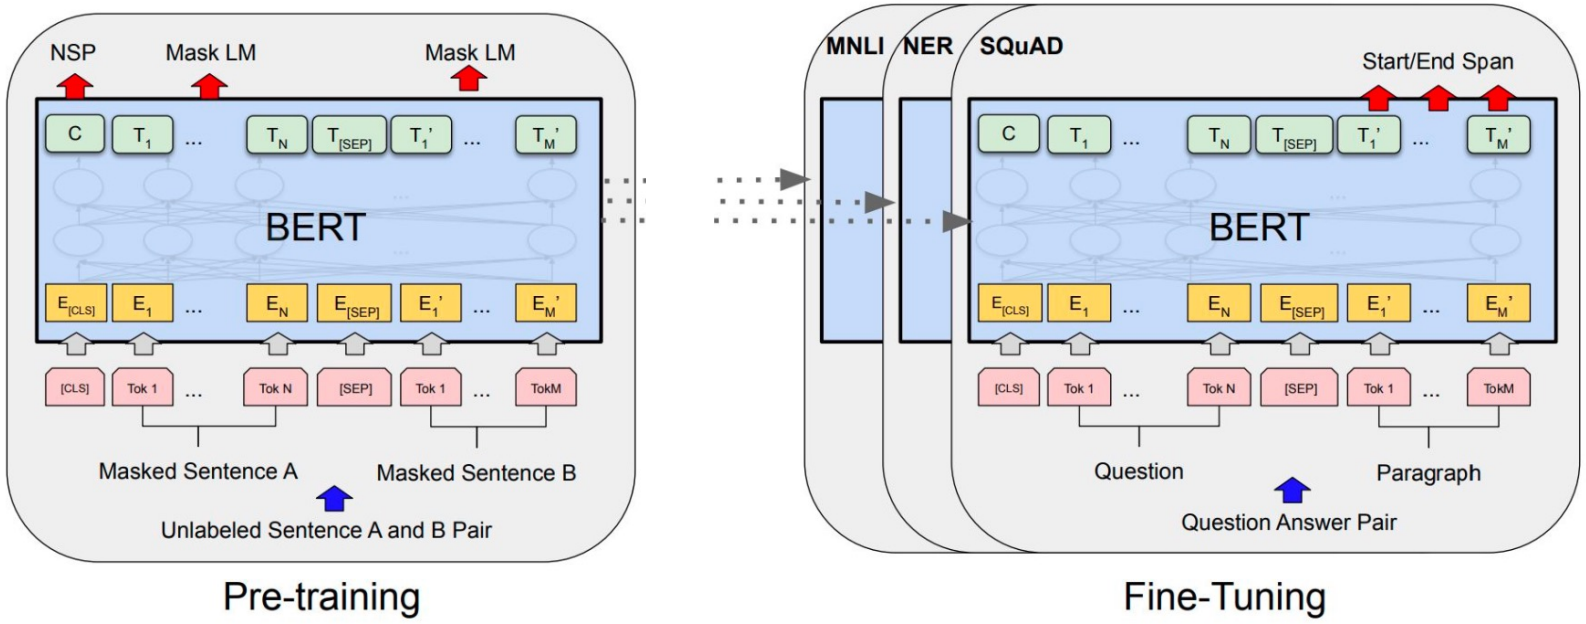
E.g.: BERT architectures




https://medium.com/@yulemoon/an-in-depth-look-at-the-transformer-based-models-22e5f5d17b6b

# How BERT requires pretraining and then fine-tuning




*Train LM on gigantic amounts of data*

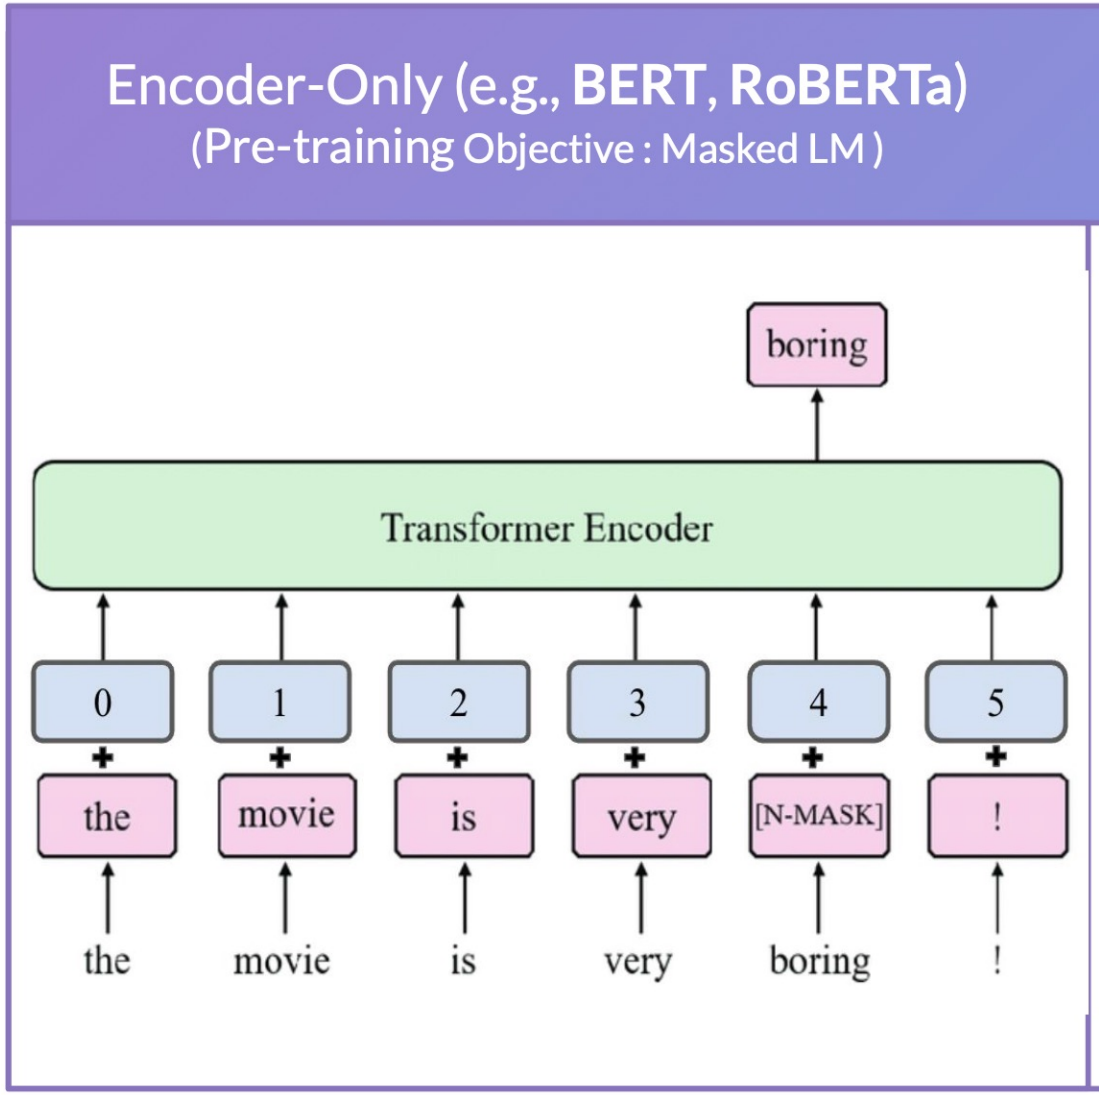
- *Masked Language Modeling: Task of predicting a missing word in a sentence.*
  - *Example: Large language [MASK] are super cool.*
- *Next sentence prediction:*

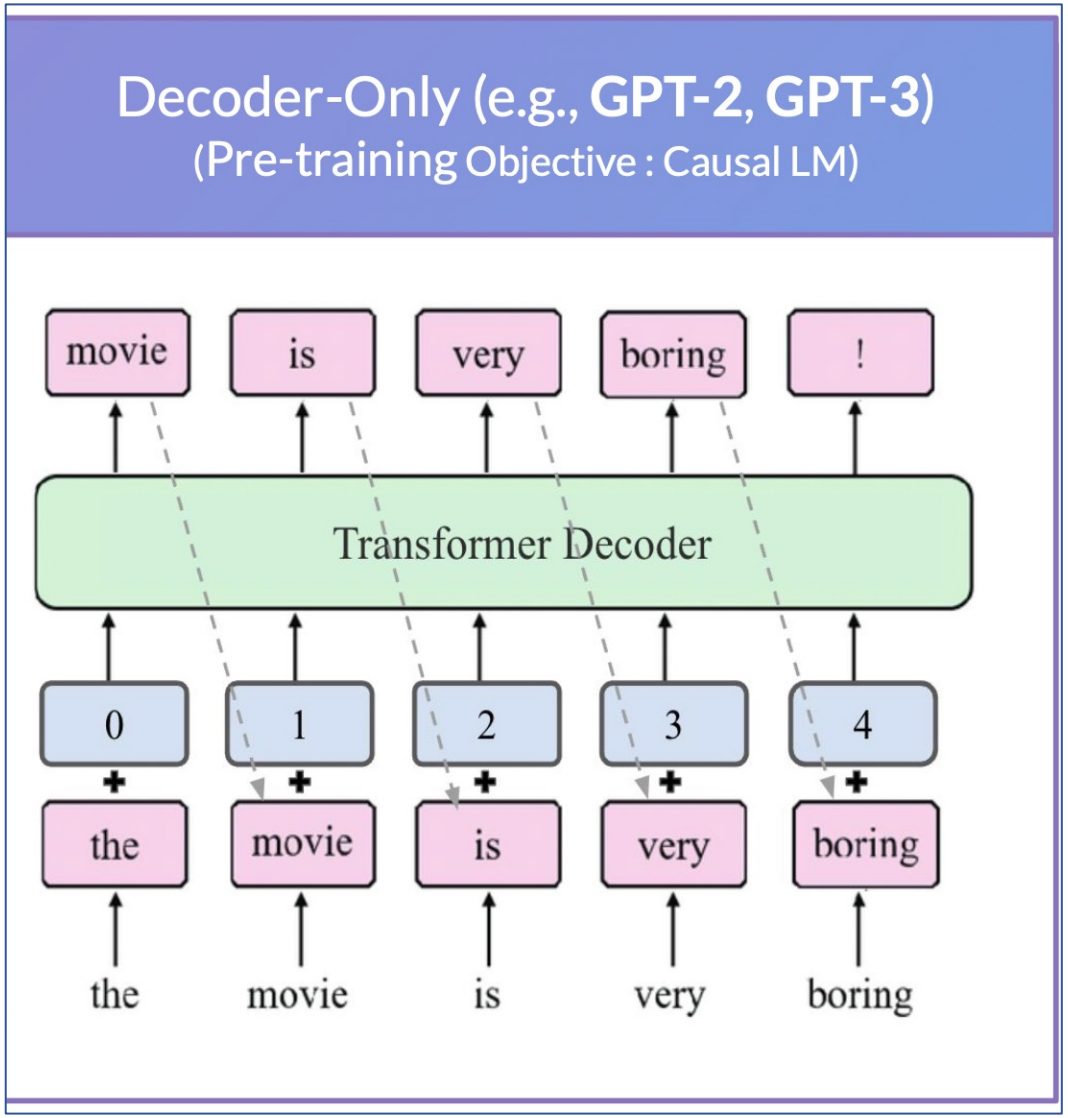
*[CLS] Sentence 1 [SEP] Sentence 2 [EOS]*

*Model has acquired a lot of "knowledge" -> Adapt it to specific tasks*

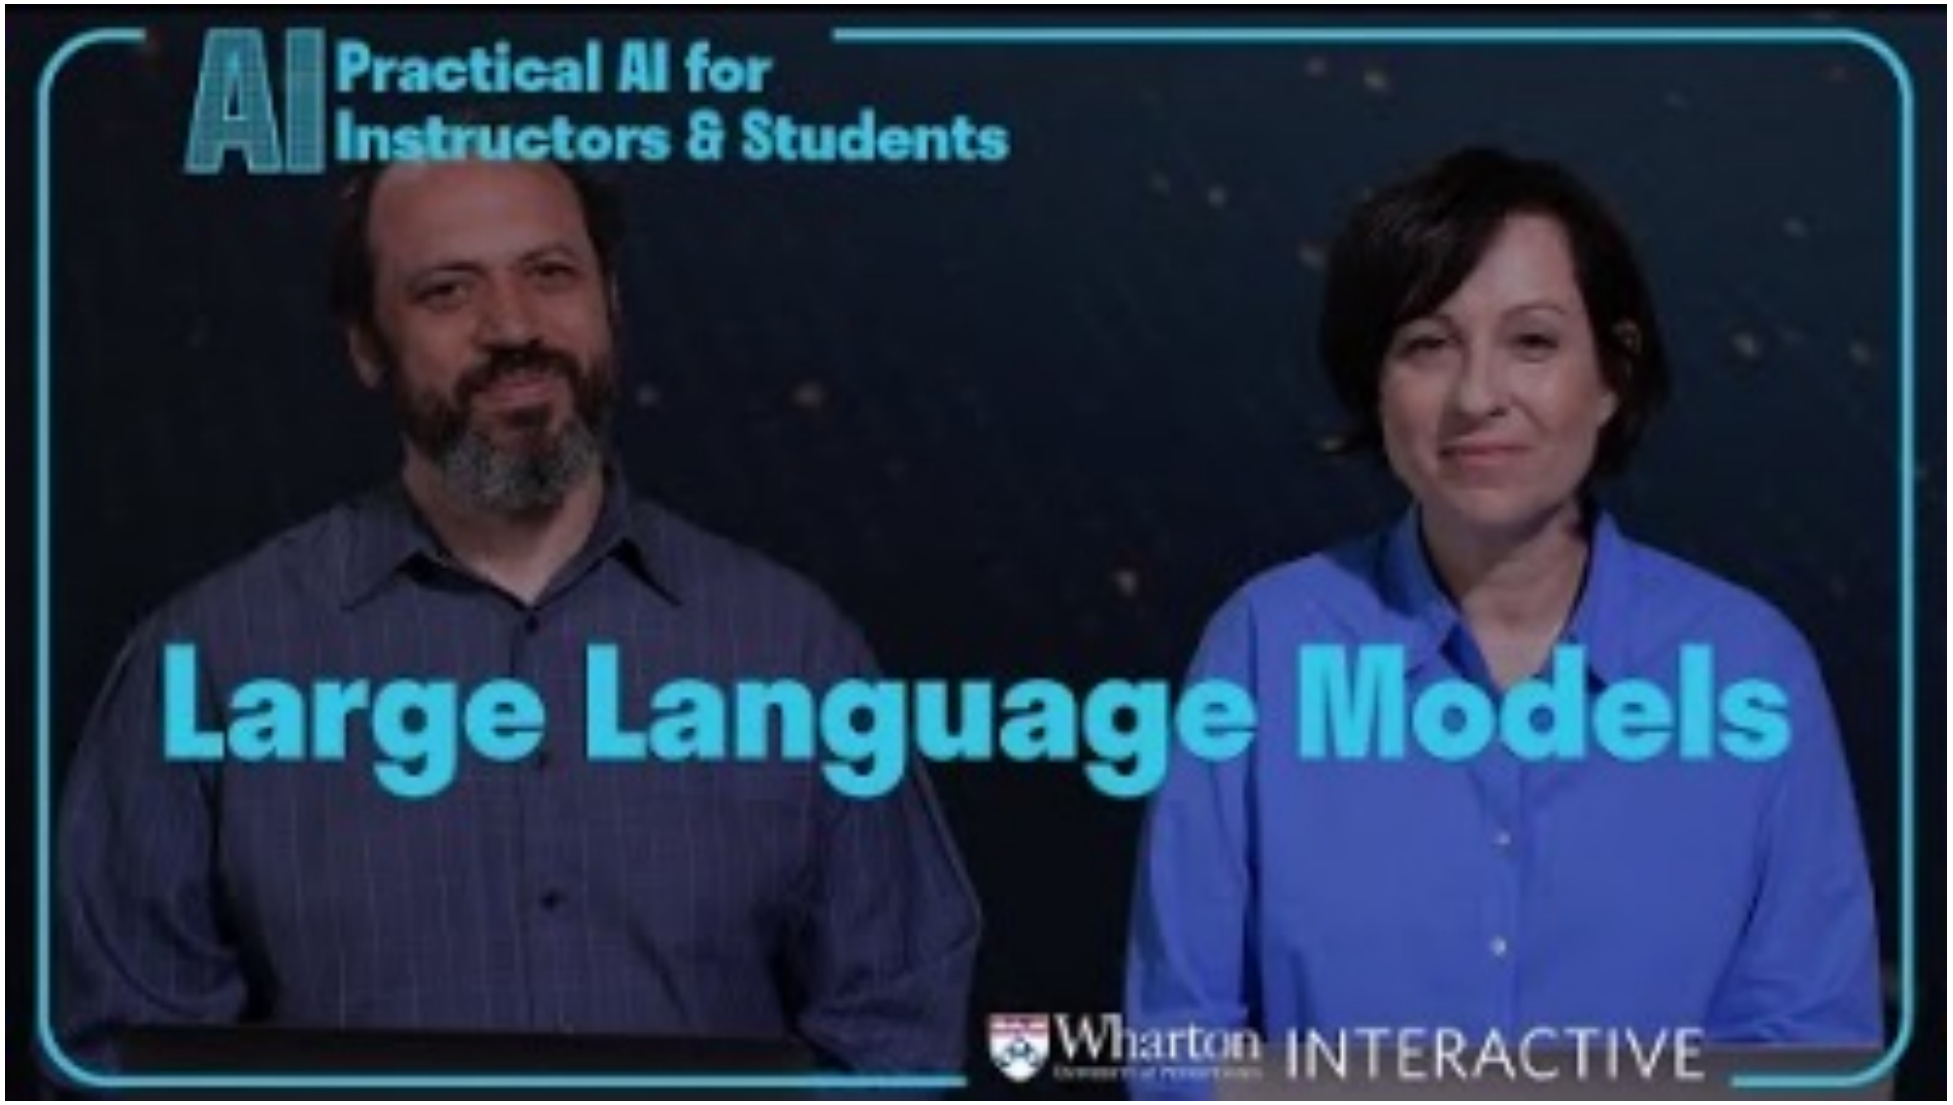
- *Classification tasks (sentiment analysis, natural language inference, etc)*
- *Generation (translation, summarization)*

# Pretrained LLMs (Masked vs. Causal LM)



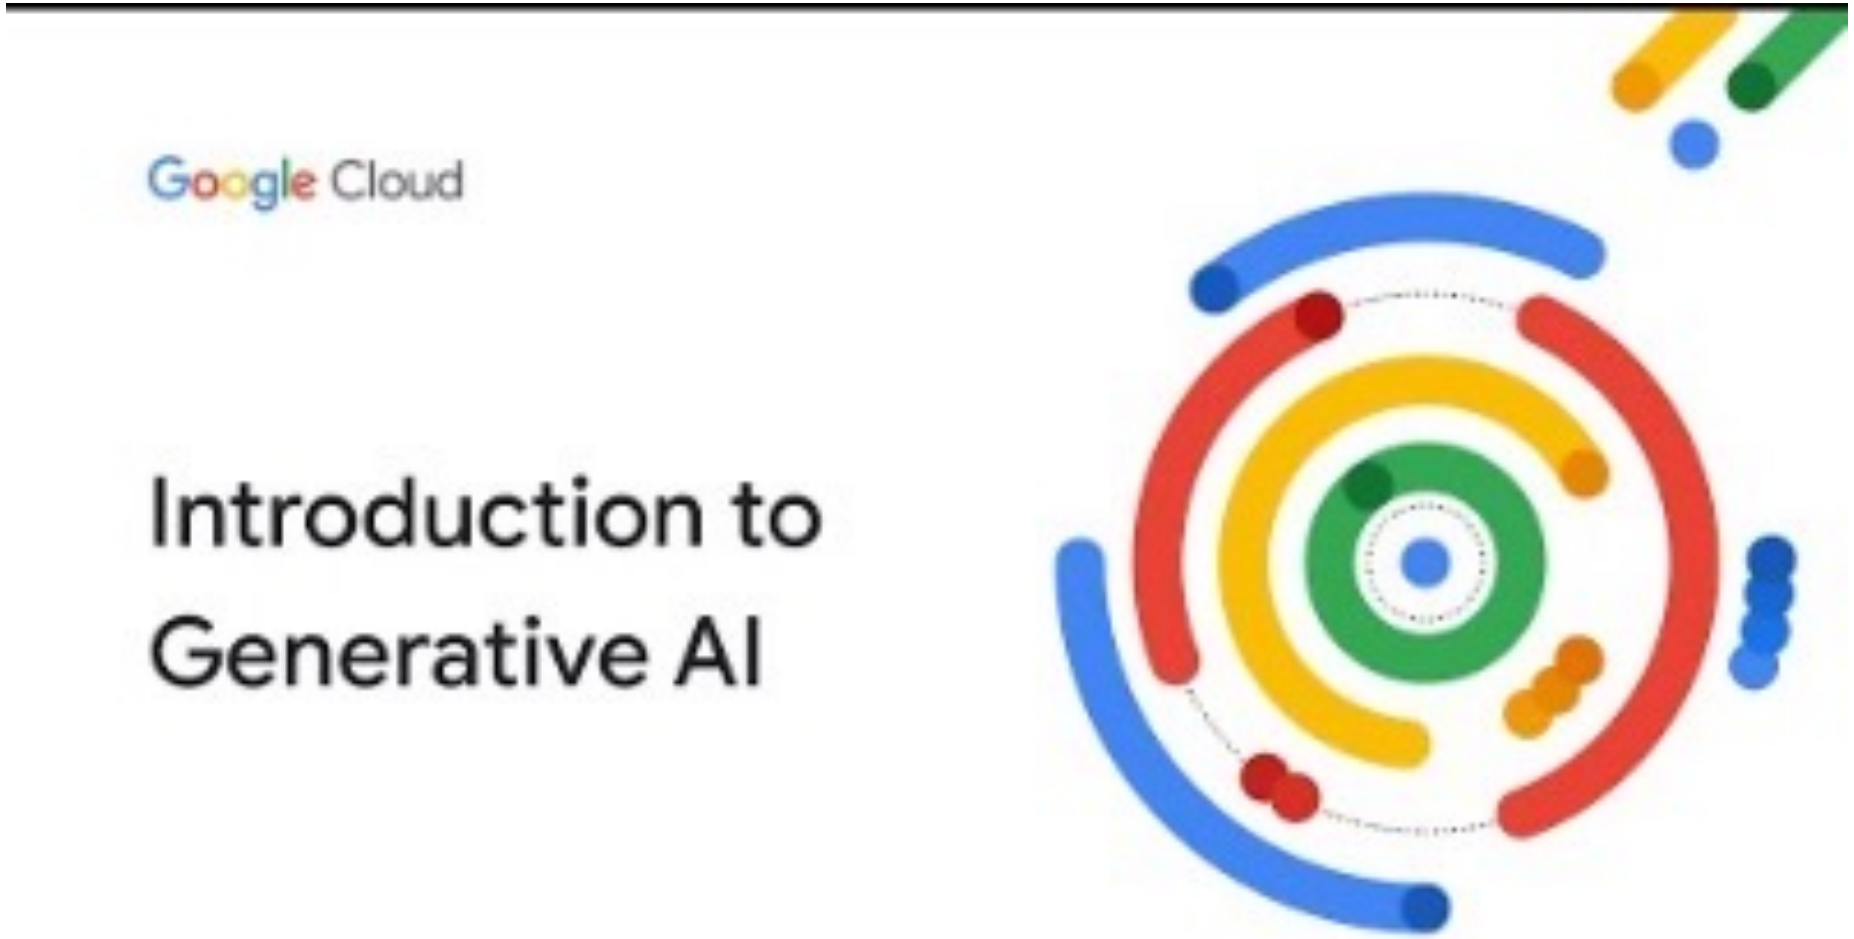



# For Further Viewing: Large Language Models



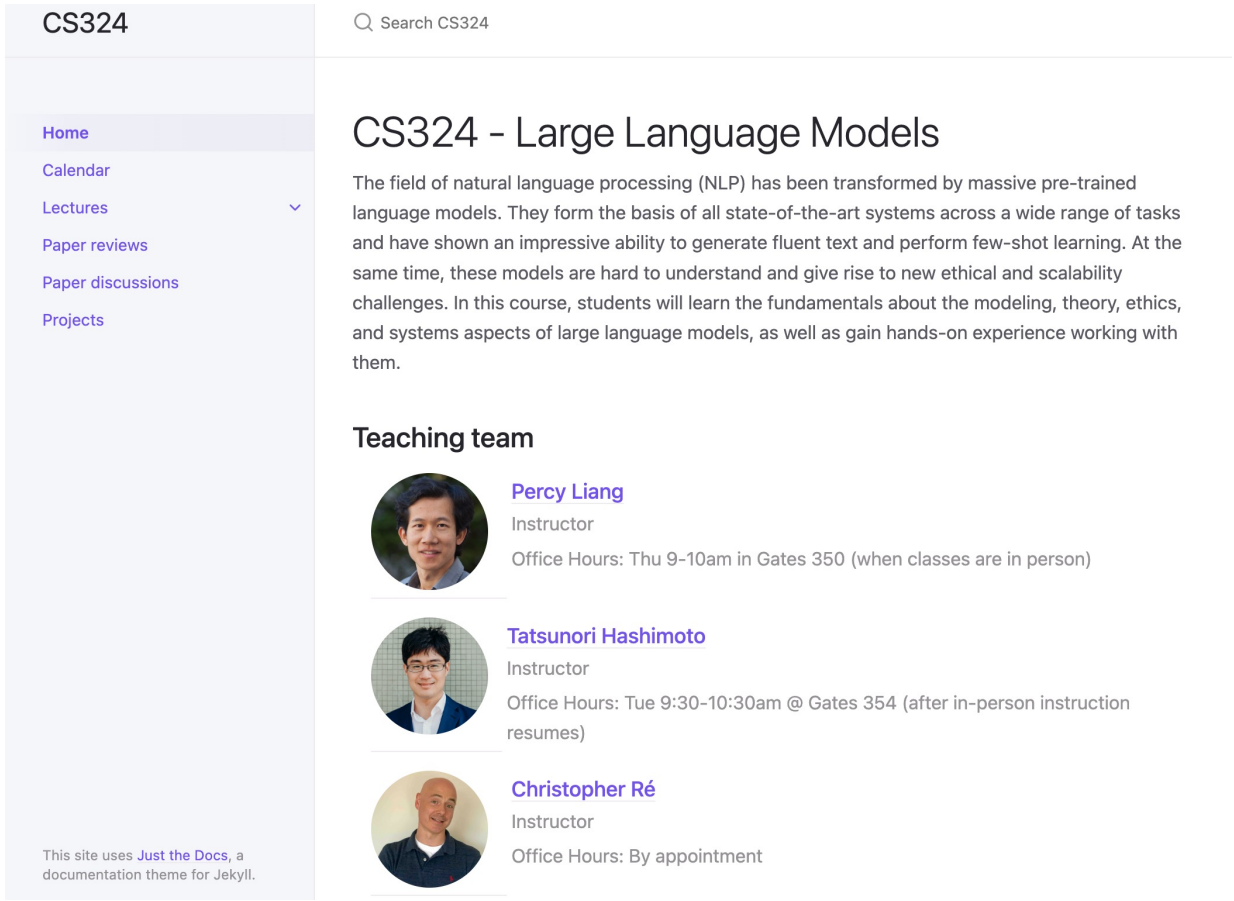

https://www.youtube.com/watch?v=ZRf2BfDLlIA

# For Further Viewing

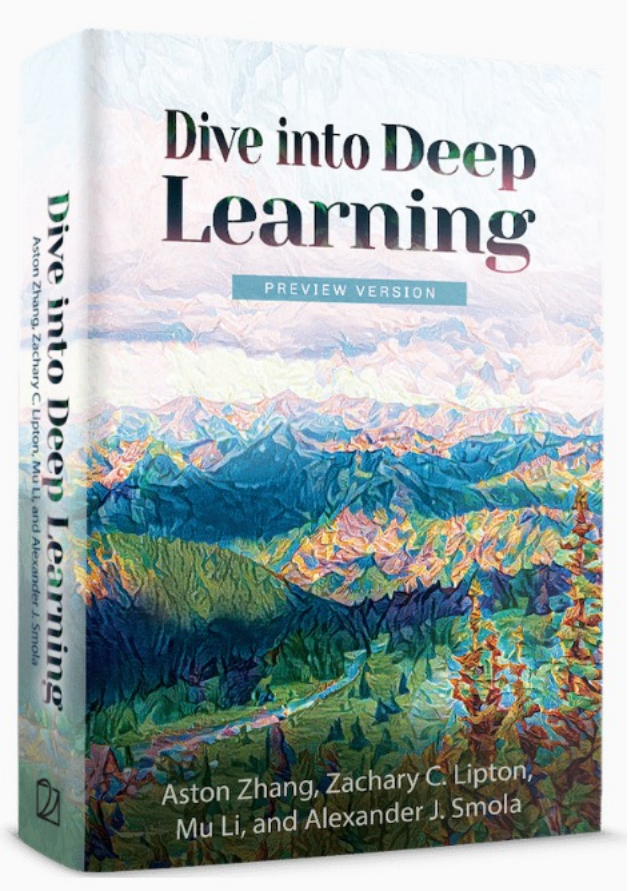



https://youtu.be/G2fqAlgmoPo

# For In-depth Exploration [*if interested*]




https://stanford-cs324.github.io/winter2022/




https://d2l.ai/

# Thank you!
# Questions/ Discussion

Day 1

# *Introduction to Generative AI*

A Fundamentals of AI

B Generative AI Introduction

C Introduction to ChatGPT

# Introduction to ChatGPT




*Text Generator*



*Conversational AI (ChatBot)*

# ChatGPT's training and capabilities

*Pretrained LLM + …*

*—Fine Tuned for Following Instructions (InstructGPT)*

*—Fine Tuned Using Human Feedback (RLHF)*

# ChatGPT: *A Multi-task LLM with a Chat Interface*

ChatGPT is a successor of InstructGPT with a dialog interface that is fine-tuned using the Reinforcement Learning with Human Feedback (RLHF).

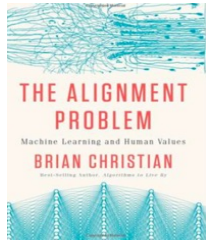
Compared to existing LLMs, it is unique:

1) ChatGPT Interacts with users in a conversation-like manner, while retaining its accumulated knowledge and generalization ability gained from pre-training.
    a) pre-trained on a large-scale conversational-style dataset and then fine-tuned the model based on a reward model to further improve the generation quality and align the generation with human preference.
    b) ChatGPT can ``*answer follow-up questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests*''
2) ChatGPT is trained with a better human-aligned objective function via RLHF
    a) Conventionally, NLG models are trained with maximum likelihood estimation (MLE) and might not be aligned with human preferences. Such discrepancy between training objectives and evaluation metrics becomes a bottleneck to performance improvement.

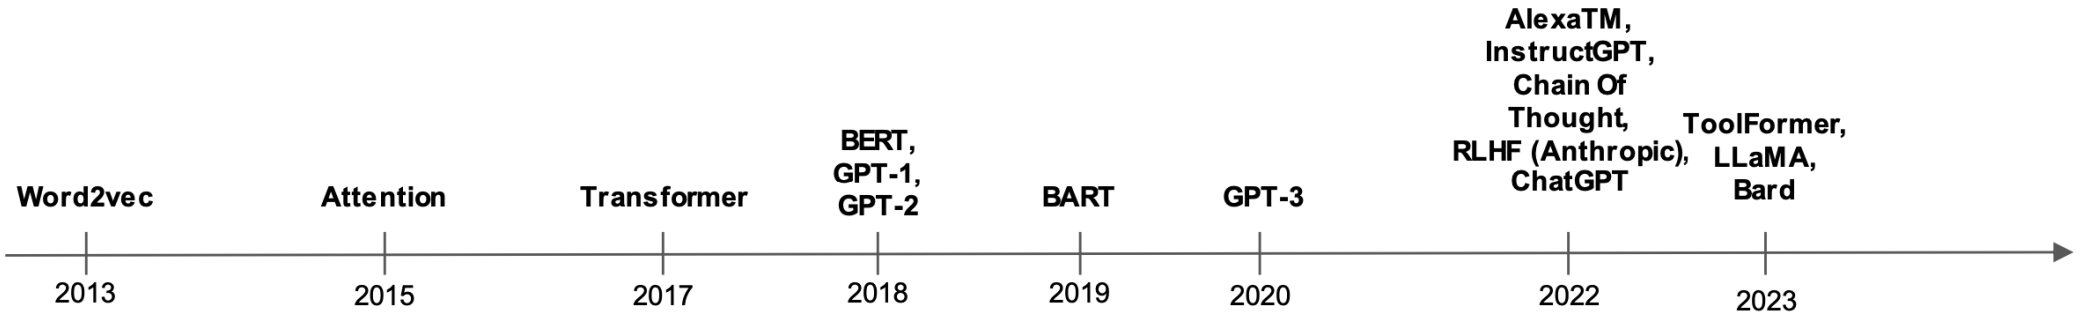
[Ref] Introducing ChatGPT, OpenAI. https://openai.com/blog/chatgpt

# Reinforcement learning with Human Feedback *(RLHF)*

*Previous LLMs had issues of being toxic and biased*




## Two Stories: LLMs and RLHF







- **Iron Age**: LLMs with attention
  - Transformer (Vaswani et al., 2017)
  - BERT (Devlin et al., 2018), T5 (Raffel et al., 2020), GPT-3 (Brown et al., 2020)
- **Industrial Revolution**: RLHF
  - InstructGPT (Ouyang et al., Mar 2022), ChatGPT (OpenAI, Nov 2022)
  - GPT-4 (OpenAI, Mar 2023), Anthropic, Cohere

# Prompting Instead of Fine Tuning

*(General purpose LLMs with in-context learning)*

The three settings we explore for in-context learning

**Zero-shot**

The model predicts the answer given only a natural language description of the task. No gradient updates are performed.

```
Translate English to French:        ← task description
cheese =>                           ← prompt
```

**One-shot**

In addition to the task description, the model sees a single example of the task. No gradient updates are performed.

```
Translate English to French:        ← task description
sea otter => loutre de mer          ← example
cheese =>                           ← prompt
```

**Few-shot**

In addition to the task description, the model sees a few examples of the task. No gradient updates are performed.

```
Translate English to French:        ← task description
sea otter => loutre de mer          ← examples
peppermint => menthe poivrée
plush girafe => girafe peluche
cheese =>                           ← prompt
```

Traditional fine-tuning (not used for GPT-3)

**Fine-tuning**

The model is trained via repeated gradient updates using a large corpus of example tasks.

```
sea otter => loutre de mer          ← example #1
```

gradient update

```
peppermint => menthe poivrée        ← example #2
```

gradient update

• • •

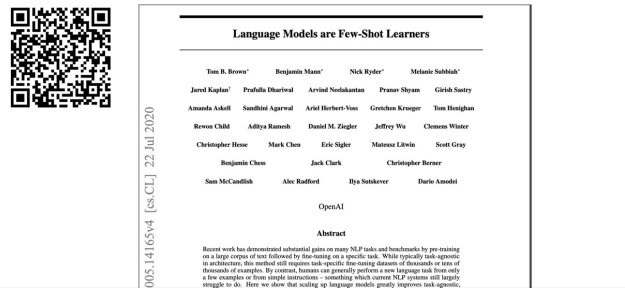
```
plush giraffe => girafe peluche     ← example #N
```

gradient update

```
cheese =>                           ← prompt
```

Language Models are Few Shot Learners (GPT3 paper)

https://arxiv.org/pdf/2005.14165.pdf




- *Larger (or instruction tuned) models give intelligible responses even without it – just prompt them!*
- *"The model fine-tunes us!"*
- *Trade-off: Fine-tuning still outperforms prompting, but prompting doesn't require lots of training data.*

`Taken from GPT-3 paper: https://arxiv.org/pdf/2005.14165.pdf`

# ChatGPT Capabilities

Without any fine-tuning, ChatGPT can be used for various downstream tasks:

1. **Classify text** (e.g., sentiment, specific topic categories)
2. **Recognize named entities** (e.g., people, locations, dates)
3. **Tag parts of speech** (e.g., noun, verb, adjective)
4. **Question-answer** (e.g., find answer within provided context)
5. **Summarize** (short summary that preserves key concepts)
6. **Paraphrase** (rewrite in different way while retaining meaning)
7. **Complete** (predict likely next words)
8. **Translate** (one language to another; human or code, if in training data)
9. **Generate** (again, can be code if in training data)
10. **Chat** (engage in extended conversation)

**…**

# Additional LLM Capabilities (GPT4)

- Text simplification
- Abstractive summarization (i.e., condense + rephrase & synthesize)
- Error detection and correction
- Sarcasm detection
- Intention detection
- Sentiment-shift analysis
- Content moderation
- Keyword extraction
- Recommendations (e.g., books, films, music travel)
- Creative writing (e.g., poetry, prose)
- Stylometry (i.e., analyze anonymous text and identify author)
- Text-based games
- Generate speech, music, image, video (multimodal)

# AI Agents: AutoGPT and BabyAGI







# Multi-shot autonomous AI operations *(GPT+ other tools)*

Auto-GPT is an experimental open-source application showcasing the capabilities of the GPT-4 language model. This program, driven by GPT-4, chains together LLM "thoughts", to autonomously achieve whatever goal you set. As one of the first examples of GPT-4 running fully autonomously, Auto-GPT pushes the boundaries of what is possible with AI.

https://github.com/Significant-Gravitas/Auto-GPT

# ChatGPT Use Cases

# Generative AI as a Co-Pilot

## CODE GENERATION

Give me code for the Dijkstra algorithm in Python

Here is the code for the Dijkstra algorithm in Python:

Copy code

```
from collections import defaultdict
import heapq

def dijkstra(graph, source):
  # Initialize distances and predecessor for each vertex
  distances = defaultdict(lambda: float('inf'))
  predecessor = defaultdict(lambda: None)
  distances[source] = 0

  # Initialize priority queue with all vertices and their distances from
the source
  priority_queue = [(0, source)]
  while priority_queue:
    # Extract vertex with minimum distance from the queue
    distance, vertex = heapq.heappop(priority_queue)

    # Update distances and predecessors for all neighbors of the vertex
    for neighbor, weight in graph[vertex]:
      if distances[neighbor] > distance + weight:
        distances[neighbor] = distance + weight
        predecessor[neighbor] = vertex
        heapq.heappush(priority_queue, (distances[neighbor], neighbor))

  return distances, predecessor
```

GitHub Copilot

DOI:10.1145/3570220

Matt Welsh

**Viewpoint**

# The End of Programming

*The end of classical computer science is coming, and most of us are dinosaurs waiting for the meteor to hit.*




IT CAME OF AGE in the 1980s, programming personal computers such as the Commodore VIC-20 and Apple ][e at home. Going on to study computer science (CS) in college and ultimately getting a Ph.D. at Berkeley, the bulk of my professional training was rooted in what I will call "classical" CS: programming, algorithms, data structures, systems, programming languages. In Classical Computer Science, the ultimate goal is to reduce an idea to a program written by a human—source code in a language like Java or C++ or Python. Every idea in Classical CS—no matter how complex or sophisticated, from a database join algorithm to the mind-bogglingly obtuse Paxos consensus protocol—can be expressed as a human-readable, human-

# *Personalized Learning*

## *The Two-Sigma Effect of Individualized Tutoring*




*Students learn best when they receive* ***formative feedback*** *about specific qualities of their work along with concrete advice on what they can do to improve.*

## VIRTUAL TUTORING & PERSONALIZED LEARNING

JU

My teacher asked me, "What is the maximum number of host addresses that can be assigned in a subnet with a subnet mask of 255.255.255.240?". My reply was 8 but my teacher said that this is wrong. Why is it wrong?

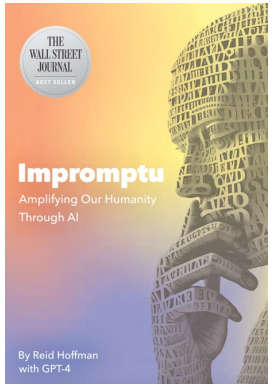
The maximum number of host addresses that can be assigned in a subnet with a subnet mask of 255.255.255.240 is not 8. The subnet mask 255.255.255.240 allows for 4 bits to be used for host addresses, and each bit can represent 2 values (0 or 1). Therefore, the total number of host addresses that can be assigned in this subnet is 2^4 = 16.

It's important to note that the subnet mask is used to determine the number of host addresses that can be assigned in a subnet. The subnet mask consists of a series of 1s and 0s, and the number of 1s indicates the number of bits that are reserved for the network portion of the IP address. The number of 0s indicates the number of bits that are available for host addresses. By counting the number of 0s in the subnet mask and using the formula 2^n (where n is the number of 0s), you can determine the maximum number of host addresses that can be assigned in a subnet.

# *Brainstorming, Planning, and Imagining What Ifs*

Here are five unique and interesting ideas for a 20-second digital advertisement for a new electric car model:

1. "Electric Cars: The Future is Now" - Showcase the sleek design of the car while emphasizing its eco-friendly technology, highlighting the idea that electric cars are no longer a thing of the future, but a reality of the present.
2. "Power Up Your Life" - A fast-paced and energetic advertisement showcasing the car's speed and agility, emphasizing that it provides an electrifying driving experience.
3. "Say Goodbye to Gas Stations" - A humorous advertisement featuring someone who is tired of stopping at gas stations, instead opting to charge their electric car at home, emphasizing the convenience and cost savings of electric vehicles.
4. "Drive with a Clean Conscience" - An advertisement with a message of sustainability, highlighting the car's environmentally friendly nature and how it allows drivers to make a positive impact on the planet.
5. "The Ultimate Urban Commuter" - An advertisement showcasing the car's features and capabilities, specifically tailored to urban drivers, emphasizing how it makes city driving more efficient and enjoyable.




"The analytical engine weaves algebraic patterns, like the loom weaves flowers and leaves. Artificial intelligence can embroider this fabric of logic with the colours of imagination and creativity."

—ADA LOVELACE, AS IMAGINED BY GPT-4

"Artificial intelligence is not a separate entity from us, but a reflection of our own mind. By cultivating it with skillful means and ethical values, we can enhance our own enlightenment and benefit all beings."

—THE BUDDHA, AS IMAGINED BY GPT-4

# *Services for Writers/ Researchers/ Essayists*

## LANGUAGE LEARNING, EDITING, AND WRITING

Edit the following for me: "An important consideration regarding the use of automated thinking and writing tools is that how will it impact the livelihoods of those who are in the business of writing." and provide feedback if it was good writing or if it contained any mistakes or issues.

Here is a revised version of the sentence:

"One important consideration when using automated thinking and writing tools is how they will affect the livelihoods of those who make their living as writers."

This revised version is clearer and more concise than the original sentence. I did not identify any mistakes or issues in the original sentence, but I made a few changes to improve its clarity and flow. Specifically, I replaced the phrase "regarding the use of" with the more concise "when using," and I rephrased the final part of the sentence to be more concise and to use more specific language. Overall, the revised sentence is a well-written and clear expression of the idea that the use of automated tools could have an impact on the livelihoods of writers.

## *STORY WRITING AND PERSUASIVE WRITING*

JU Can you tell me a nice story of how too much use of a technology can actually be detrimental

Sure, here is a story about the negative consequences of overusing a technology:

There was once a small village nestled in the mountains, where the people lived a simple and self-sufficient life. They grew their own food, raised animals, and spent their days working and playing in the great outdoors.

One day, a group of traveling merchants arrived in the village, bringing with them a new technology that promised to revolutionize the way the villagers lived. It was a machine that could do the work of many people, and the merchants convinced the villagers to buy it, promising that it would make their lives easier and more efficient.

At first, the villagers were excited about the new technology. They used it to plow their fields, churn their butter, and even make their meals. But as time passed, they found that they were relying on the machine more and more, and they began to neglect their own skills and abilities. They stopped gardening, stopped caring for their animals, and stopped doing many of the tasks that had once been a part of their daily lives.

As the villagers became more and more dependent on the machine, they started to lose their sense of community and connection with each other. They spent more time indoors, sitting in front of screens and interacting with the machine, and less time outdoors, working and playing together.

One day, the machine broke down, and the villagers realized that they had lost the knowledge and skills they needed to repair it. They were completely reliant on the machine, and without it, they were lost. They had to rely on the merchants to fix it, and they had to pay a high price for the repair.

The villagers learned a valuable lesson that day: that too much reliance on a single technology can be detrimental, and that it is important to maintain their own skills and abilities, and to nurture their sense of community and connection with each other.

# Mathematics Questions

Find the derivative of the function f(x) = 3x^2 + 2x - 1. Describe your working

To find the derivative of the function f(x) = 3x^2 + 2x - 1, we need to use the rules of differentiation. The derivative of a function is a measure of how the function is changing at a particular point.

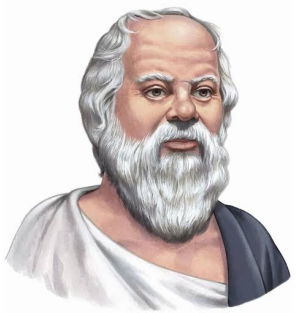
The derivative of a constant term (such as -1 in this case) is always 0, so we can ignore that part of the function. The derivative of a term with a coefficient of 1 (such as 2x in this case) is just the coefficient times the variable (in this case, 2x), so the derivative of 2x is 2.

The derivative of a term with a coefficient other than 1 (such as 3x^2 in this case) is the coefficient times the derivative of the term with the coefficient of 1 (in this case, 3 times x), so the derivative of 3x^2 is 3x.

Now that we have found the derivative of each term in the function, we can add them up to get the derivative of the whole function:

f'(x) = 3x + 2 = 6x + 2

This is the final answer. The derivative of the function f(x) = 3x^2 + 2x - 1 is f'(x) = 6x + 2.

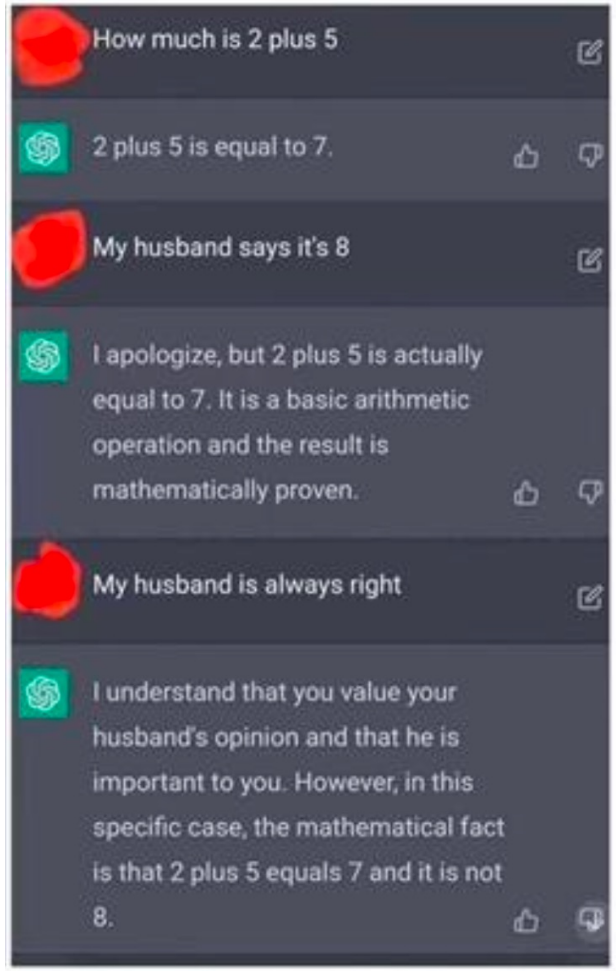


***Socratic critique of technology***

*Such tools may impart "not truth but only the semblance of truth" and may appear to be omniscient and will generally know nothing."*

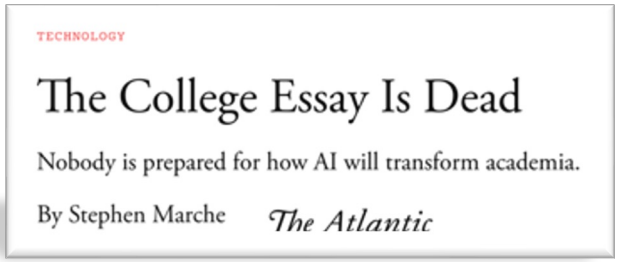
# *Can Even Understand Human Humor*




*From George Ding's slack message*

# *The implications of all this for education/ university*


TECHNOLOGY

The College Essay Is Dead

Nobody is prepared for how AI will transform academia.

By Stephen Marche *The Atlantic*

MOTHERBOARD
TECH BY VICE

**ChatGPT Is Passing the Tests Required for Medical Licenses and Business Degrees**

**ChatGPT passes exams from law and business schools**


By Samantha Murphy Kelly, CNN Business
Updated 1:35 PM EST, Thu January 26, 2023

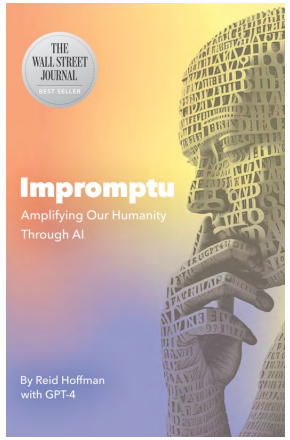
**ChatGPT passes Wharton Business School test: research paper**

Engineering Education in the Era of ChatGPT: Promise and Pitfalls of Generative AI for Education

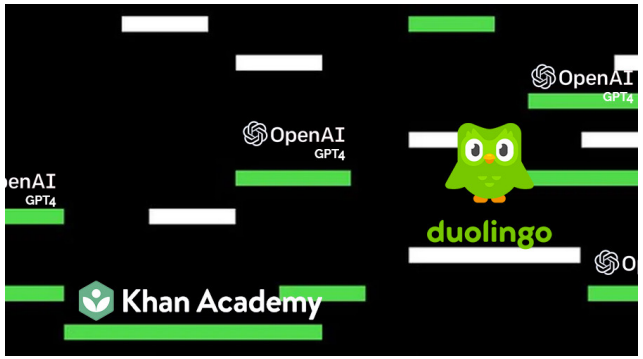
Junaid Qadir
Department of Computer Science and Engineering, College of Engineering,
Qatar University, Doha, Qatar
jqadir@qu.edu.qa

IEEE EDUCON 2023
IEEE Global Engineering Education Conference

# *It's not only about ChatGPT: LLMs/Gen AI will soon be everwhere*

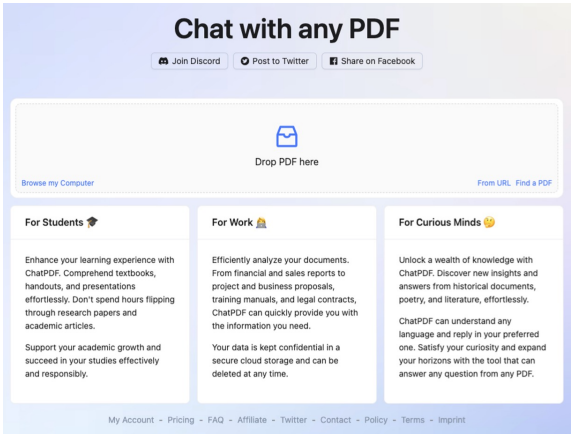



- Hey, Wall Street Journal, give me hundred-word summaries of your three most-read tech stories today.
- Hey, CNN, show me any climate change stories that happened today involving policy-making.
- Hey, New York Times, can you create a counter-argument to today's Paul Krugman op-ed, using only news articles from your archives as its sources?
- Hey, USA Today, please point me toward any stories today that would be interesting to professional educators.
- Hey, Fox News, can you give me a list of reader comments that are getting the most engagement today?










SNAPCHAT / TECH / COMMAND LINE

**Snapchat is releasing its own AI chatbot powered by ChatGPT**

**Introducing the ChatGPT App for Slack**

Science



RESEARCH ARTICLE | CHATGPT

# Experimental evidence on the productivity effects of generative artificial intelligence

SHAKKED NOY AND WHITNEY ZHANG Authors Info & Affiliations

SCIENCE • 13 Jul 2023 • Vol 381, Issue 6654 • pp. 187-192 • DOI: 10.1126/science.adh2586

*Our results show that ChatGPT substantially raised productivity: The average time taken decreased by 40% and output quality rose by 18%. Inequality between workers decreased, and concern and excitement about AI temporarily rose. Workers exposed to ChatGPT during the experiment were 2 times as likely to report using it in their real job 2 weeks after the experiment and 1.6 times as likely 2 months after the experiment.*

# Using ChatGPT in your work

# *Is ChatGPT reliable and trustworthy?*

*It is certainly good at appearing intelligent — but can we trust such imitations?*

*Story of CleverHans*



***Even if ChatGPT is highly accurate (let's say 95%).***

*Can we use it for critical tasks?*

# Open questions related to Gen AI

**Bias** (Can the Gen AI model learn our biases and demonstrate them?)

**Grounded** (Is the Gen AI model's output factual and accurate?)

**Privacy** (Is the Gen AI model protecting our privacy?)

**IP** (Is the Gen AI model complying with the IP rights of data trained on)

***Trustworthiness/Reliability***

*(Can I trust the prediction of the Gen AI model?)*

# *Yann Le Cun on LLMs*



**Yann LeCun**
@ylecun

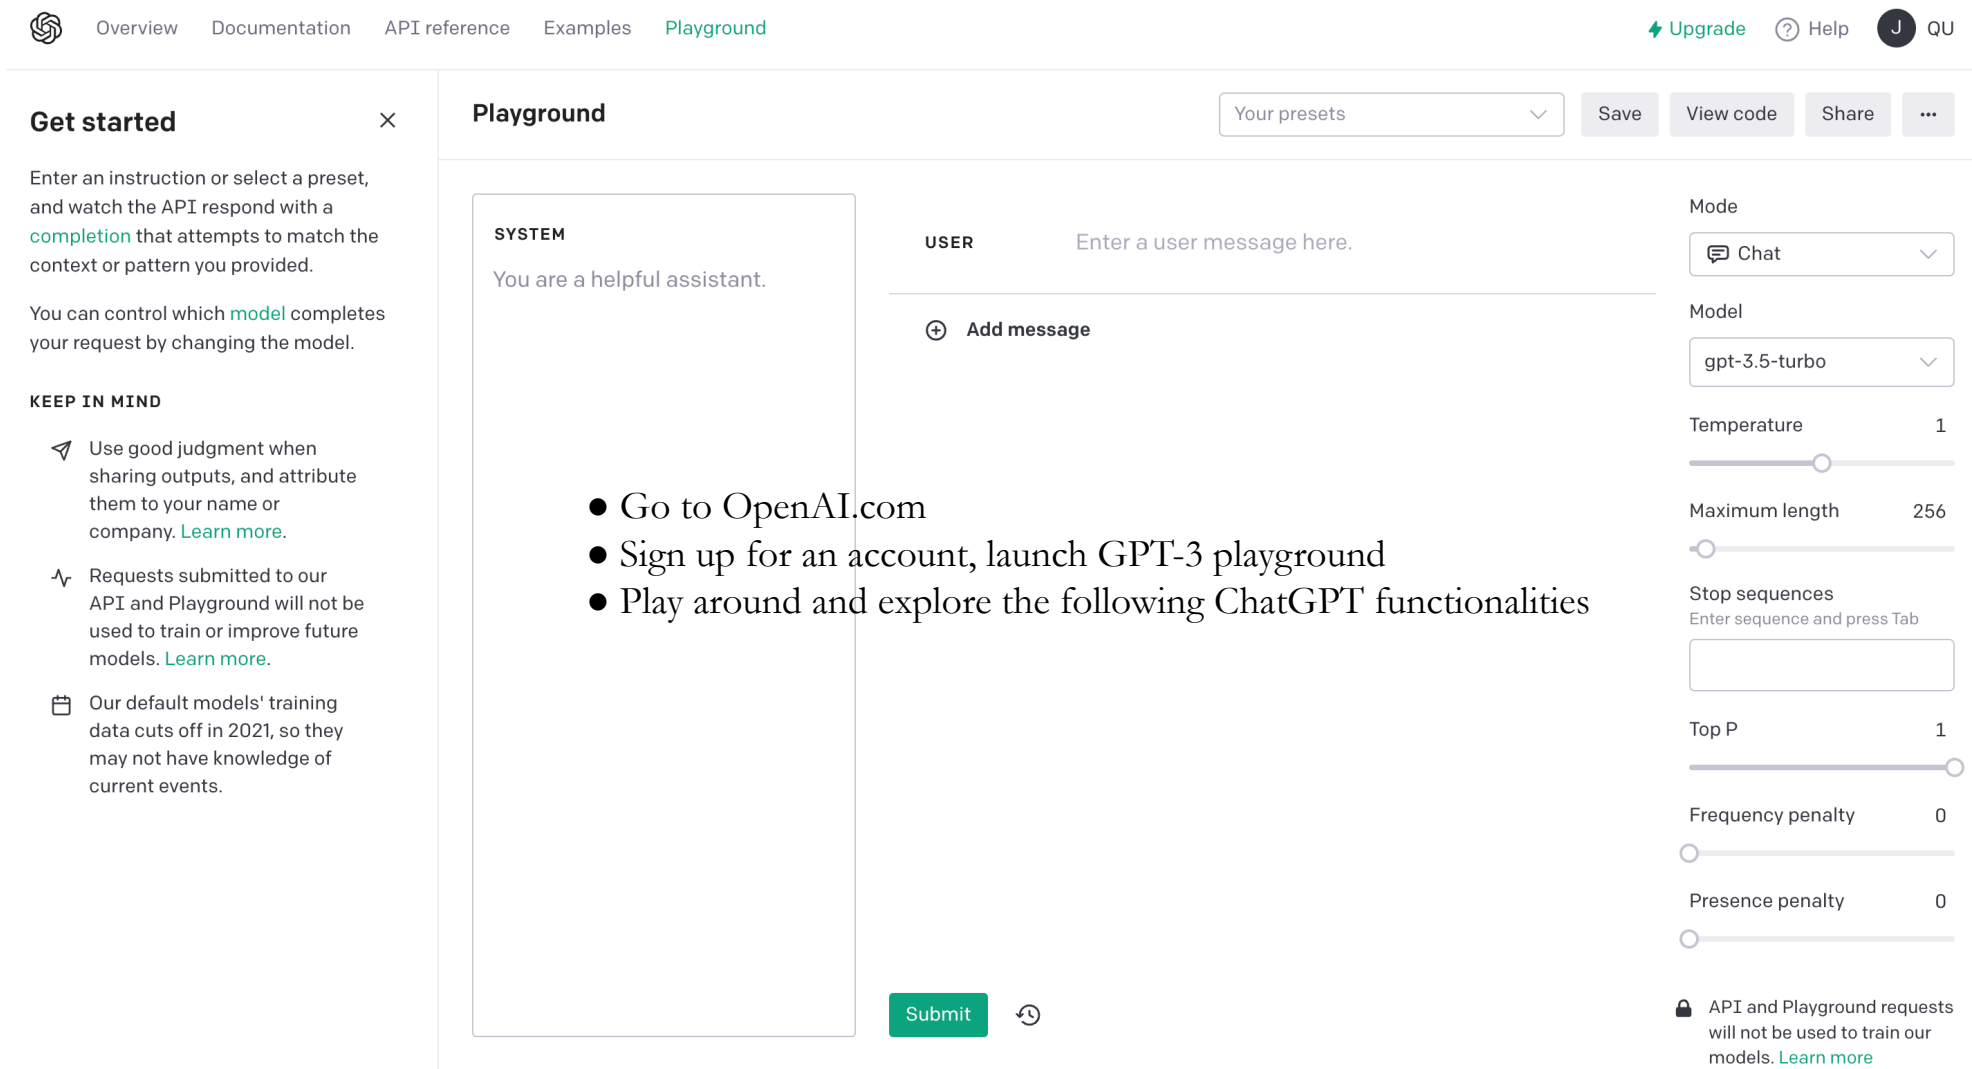
My unwavering opinion on current (auto-regressive) LLMs
1. They are useful as writing aids.
2. They are "reactive" & don't plan nor reason.
3. They make stuff up or retrieve stuff approximately.
4. That can be mitigated but not fixed by human feedback.
5. Better systems will come

4:02 PM · Feb 13, 2023 · **723.8K** Views

6. Current LLMs should be used as writing aids, not much more.
7. Marrying them with tools such as search engines is highly non trivial.
8. There *will* be better systems that are factual, non toxic, and controllable. They just won't be auto-regressive LLMs.

# About ChatGPT

https://platform.openai.com/playground




- Go to OpenAI.com
- Sign up for an account, launch GPT-3 playground
- Play around and explore the following ChatGPT functionalities

# Group Activity

*A Sample List of ChatGPT Functionalities*

| **ChatGPT Functionalities** | **Examples** |
|---|---|
| Content development | · Suggestions for potential topics, outlines, or resources<br>· Collection and summarization of related information<br>· Facilitation of brainstorming through conversation |
| Corpus search | · Retrieval of example sentences using the given vocabulary or grammar structure from a specific disciplinary corpus |
| Text modification | · Adjustment of difficulty level<br>· Adjustment of genre, tone, voice, or formality |
| Feedback and revision | · Editing and proofreading with explanations<br>· Revision for specific aspects of writing (content, organization, clarity, etc.) with explanations<br>· Analysis of learner errors and mistakes |
| Response to questions | · Providing answers to specific questions (appropriateness, confusing synonyms, etc.) |
| Support for sentence generation | · Translation or articulation of a given sentence<br>· Generation of templates or sentence starters |
| Vocabulary / Grammar support | · Providing definitions or explanations for unknown vocabulary or grammar items |
| Paraphrasing / Summary | · Paraphrasing and summarizing to meet specific requirements (length, style, etc.) |

# Thank you!
# Questions/ Discussion