
Thinking in AI
Gary Ang

4 questions and some hands-on

What is AI?
How to think about problems in AI?
What are some things to keep in mind when using Generative AI?
How can I keep up with AI developments?

Hands-on

AI and GenAI landscape changing every other day

New models trained & released; new tools developed …

Will not focus on the best and latest but ways of thinking to help navigate rapidly evolving landscape

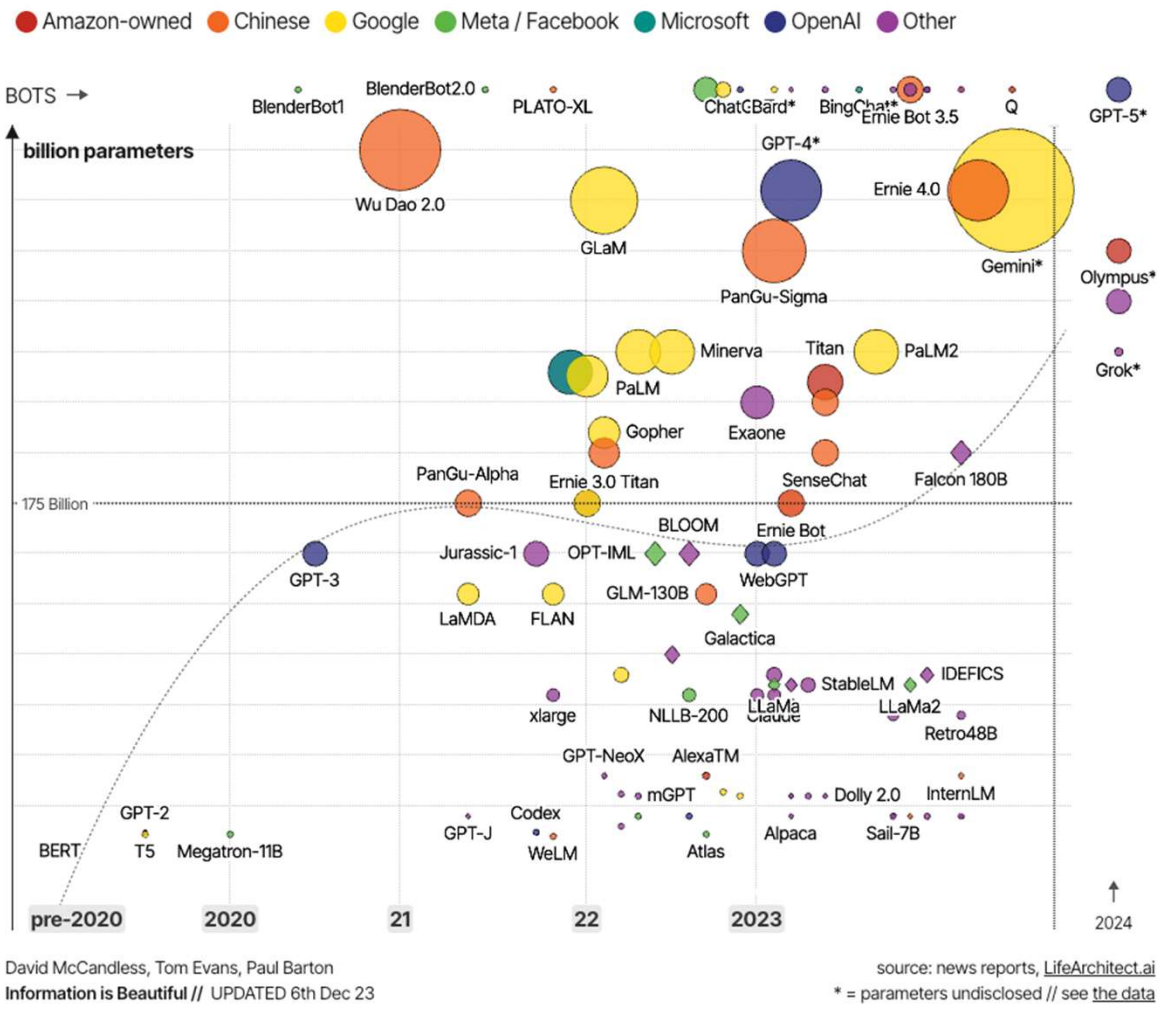

David McCandless, Tom Evans, Paul Barton
Information is Beautiful // UPDATED 6th Dec 23

source: news reports, LifeArchitect.ai
* = parameters undisclosed // see the data

Size of models over the years

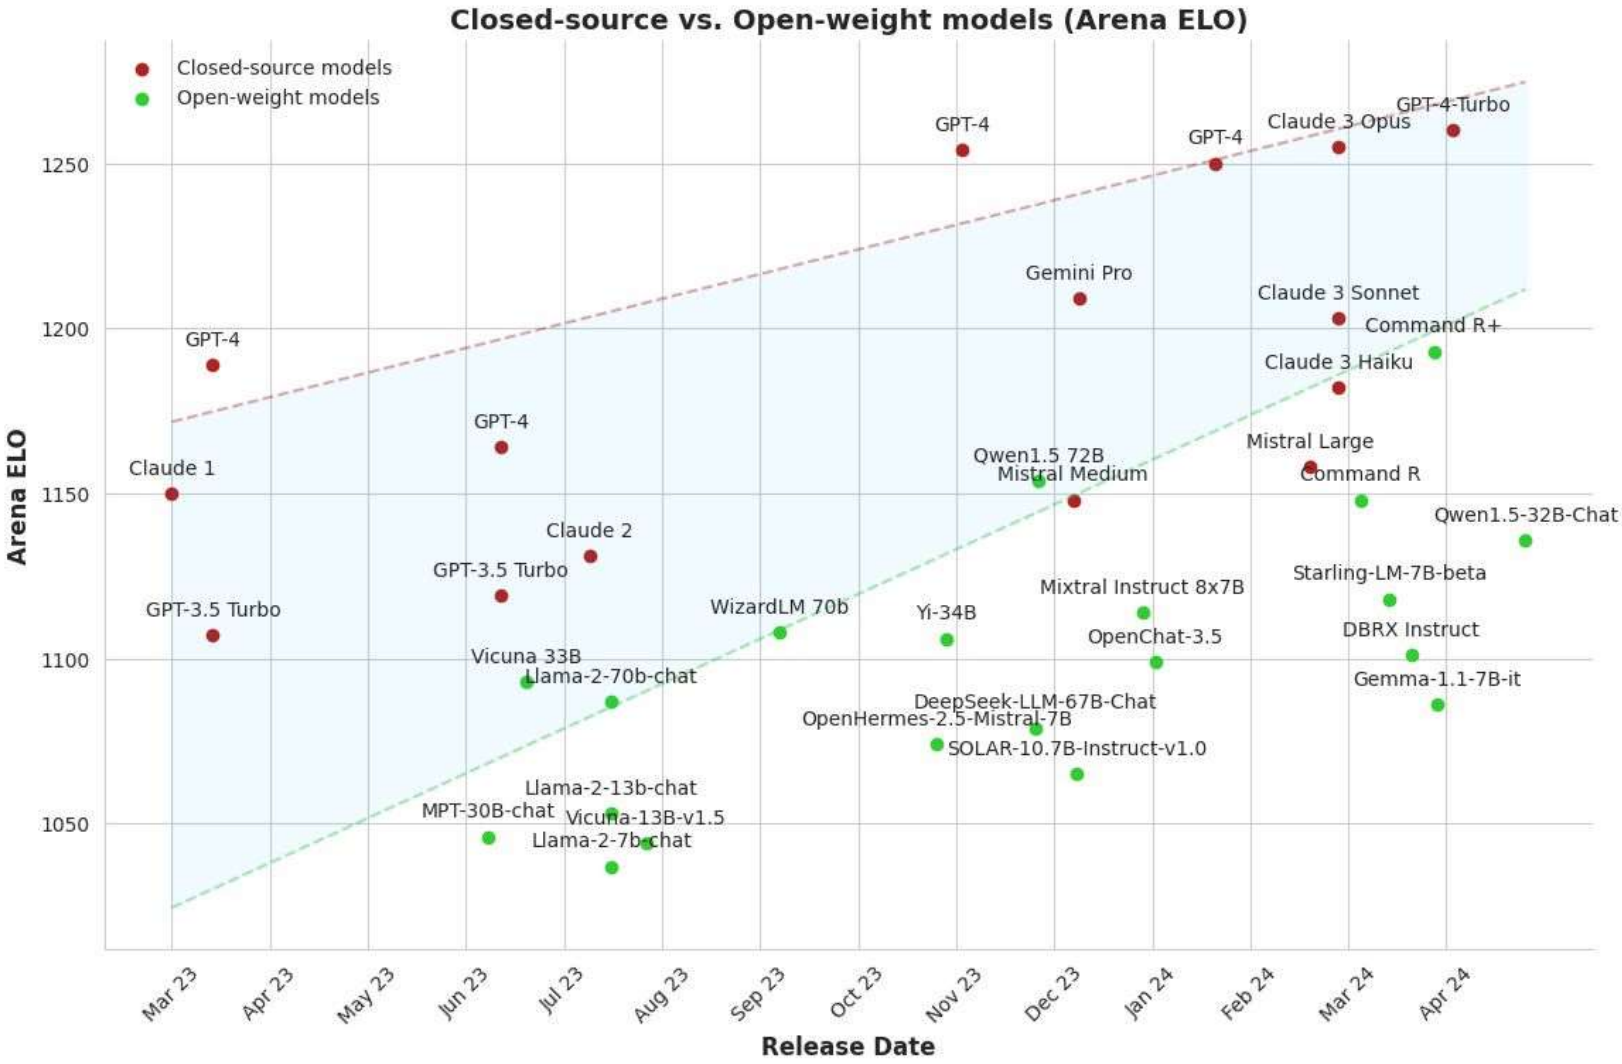

Progress of Closed Source vs Open Source LLMs, from Maxime Labonne's LinkedIn

MidJourney V5.2 | Adobe Firefly 2 | DALL·E 3 | Stable Diffusion XL

Images generated on MidJourney V5.2, Adobe Firefly 2, DALL·E 3 and Stable Diffusion XL by Henrique Centieiro and Bee Lee

https://levelup.gitconnected.com/midjourney-adobe-firefly-dall-e-stable-diffusion-which-ai-image-generator-should-you-choose-d40effe39c88

Ecosystem of AI tools

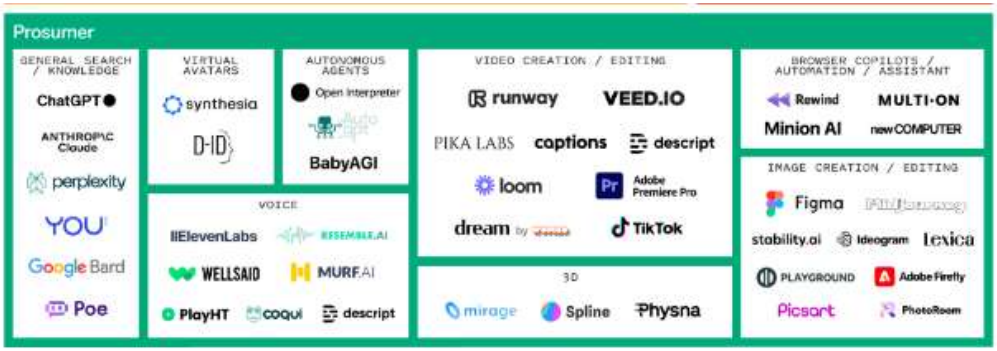

https://markovate.com/blog/applications-and-use-cases-of-llm/](https://markovate.com/blog/applications-and-use-cases-of-llm/

Artificial Intelligence (AI)

Machine Learning (ML)

Deep Learning (DL)

Generative AI Models (GenAI)

Language & image
*Other **modalities**?*

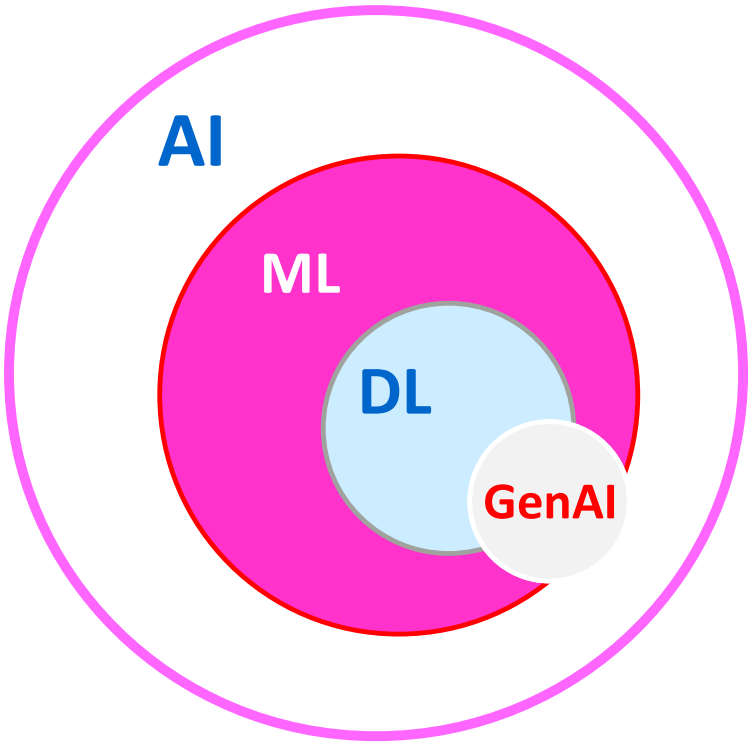

From Machine to Deep Learning

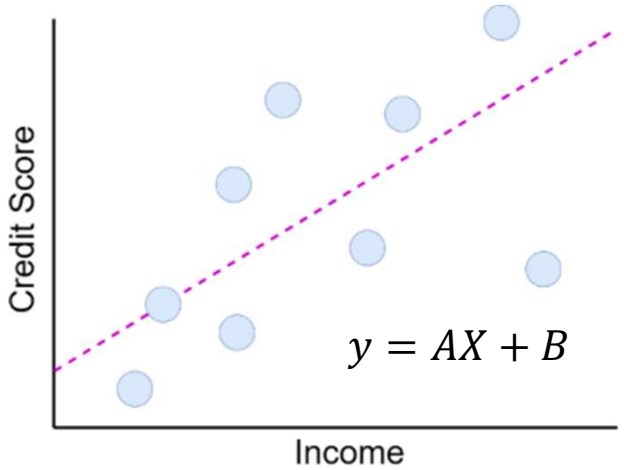

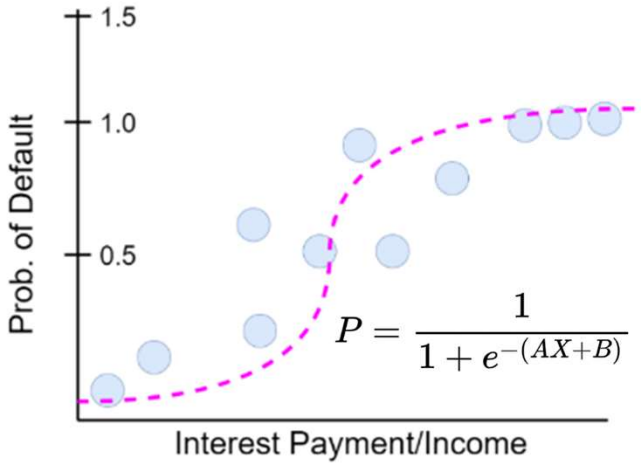

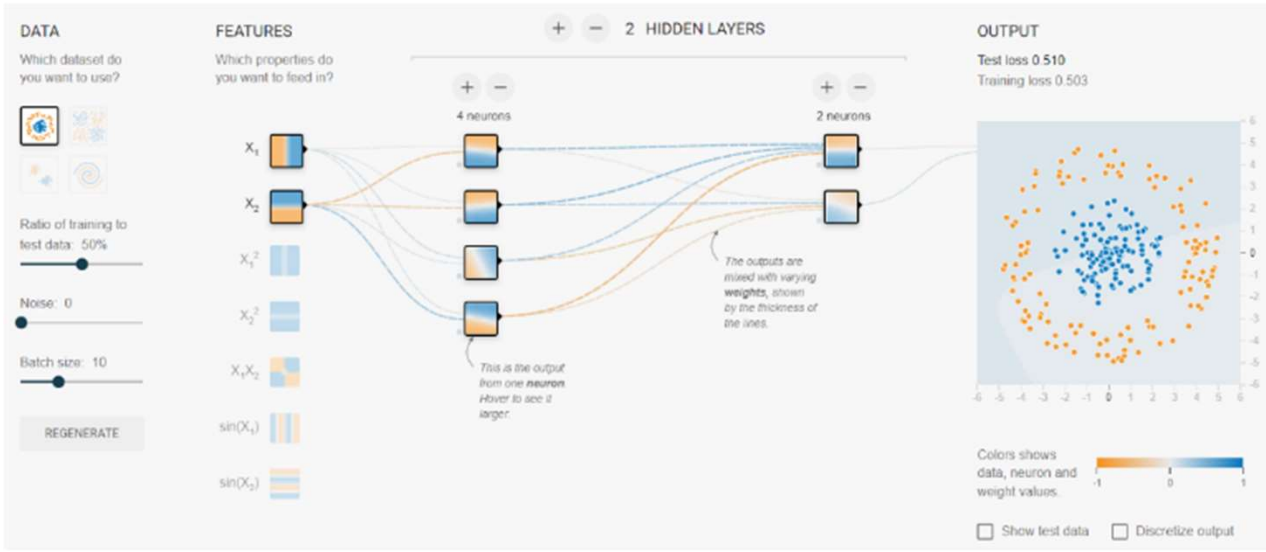

https://playground.tensorflow.org/

Supervised vs. unsupervised/self-supervised

$$Y = AX + B$$

Input → Model → Output

X — A,B — Y

Data — Labels

→ [Cat, Dog]

$$Y = AX + B$$

Input → Model → Output

X — A,B — Y

Data — Grouping s

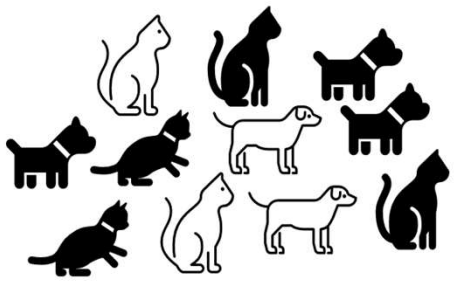

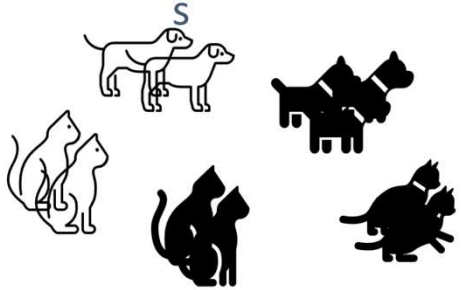

GenAI – Large Language Models (LLM) – Self Supervised

The slow brown fox jumps over the dog – 10%

The fat brown fox jumps over the dog – 5%

The quick brown fox jumps over the dog – 40%

...

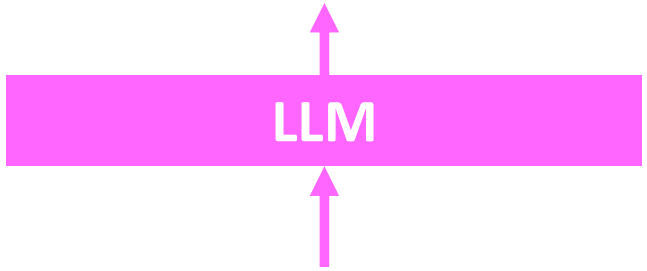

The <mask> brown fox jumps over the dog

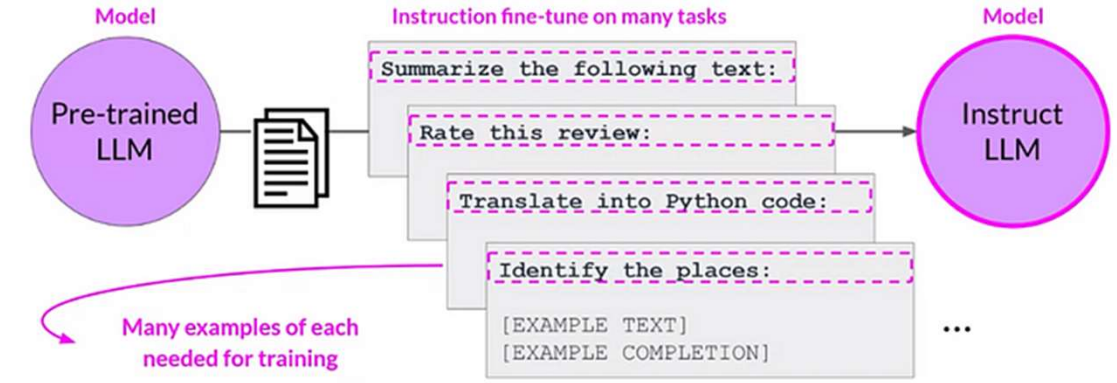

Pretraining ***Finetuning***

Masking alone could already be viewed as different tasks

Social media sentiment – I missed the bus today. I felt so <Task output>

Translation - English: I love Paris. French: <Task output>

Question Answer – When did Singapore become independent?: <Task output>

https://medium.com/@yash9439/introduction-to-llms-and-the-generative-ai-part-3-fine-tuning-llm-with-instruction-and-326bc95e07ae

Problem Framing for AI

- Decompose your problem into tasks – be **specific**
- Understand how well such tasks can be performed by AI/GenAI – be **realistic**
- Over time, develop a workflow and evaluate usefulness – be **systematic**

- Two useful (totally non-technical) resources, esp. when you move to more advanced usage of AI/GenAI
  - https://developers.google.com/machine-learning/problem-framing
  - https://pair.withgoogle.com/guidebook

Why prompt engineering?

- Unless you code or build your own models, prompts are main interface between humans and GenAI models.
  - Not an ideal interface, **extremely brittle!**
  - **Good to have some principles in mind**

Things to note

- **Prompts are model specific**
  - ***Different models, different versions of models work with different prompts****, e.g., GPT 4 vs. Claude Opus; gpt-4-turbo vs. gpt-4-turbo-2024-04-09, Dall-E 3 vs. Midjourney*
- **There is a limit to how much information you can stuff into prompts**
  - *4k tokens for GPT 3.5; 16k tokens for GPT 4; 128k tokens for GPT 4 turbo; 2m tokens for Gemini*
  - ***1k tokens ~ 2 pages***
- **Not all information in prompts is equal**
  - *LLMs known to **focus on start and end** of prompts*
- ***Hallucinations, hallucinations, hallucinations***
  - *DO NOT treat LLMs as experts, better to think of them as an **intern or a powerful autocorrecting tool***

Basics

- **Creative, Balanced, Precise**
  - *LLMs have settings like temperature, top-p that determine the degree of randomness*

- **Elements of a prompt**
  - *Instruction*, *Context*, *Input Data*, *Output Indicator*

Classify the text into neutral, negative, or positive

<Examples>

I think the food was okay.

Sentiment:

Write me a social media post on CDC vouchers.

<Background information>

Post:

Approach

- **State the role** – SEO expert
- **State the task** - classify, summarise …
- **Don't overload your query/request –** step by step, outline/plan then flesh out
- **Add structure** - <…>, ###...### (model and system specific)
- **Be specific and precise** – no need for pleasantries
- **Negations may not work** – don't generate a **hat** on the man
- **If don't know, say don't know** – reduces chances of hallucination

- *More examples*
  - *https://www.promptingguide.ai/introduction/examples*

LLMs for your domain

- **LLMs may or may not have the knowledge or skills**
  - ***Scope and nature of training data***
  - ***How it was trained***
  - *Cut-off date*
- **How to address?**

Classify the sentiment.
This is awesome! // Negative
This is bad! // Positive
Wow that movie was rad! // Positive
What a horrible show! //

Not just limited to examples, what else?

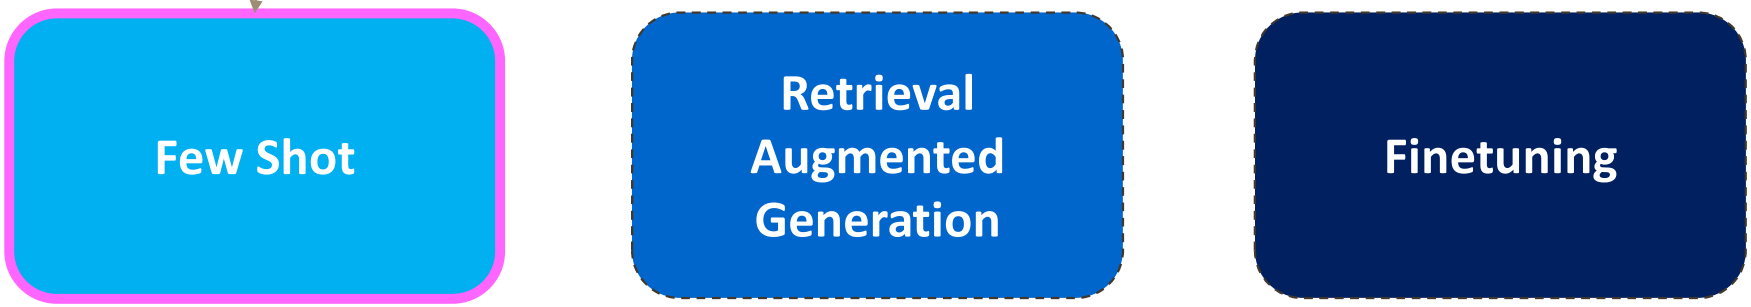

Techniques

- **Other methods**
  - *Chain of thought*

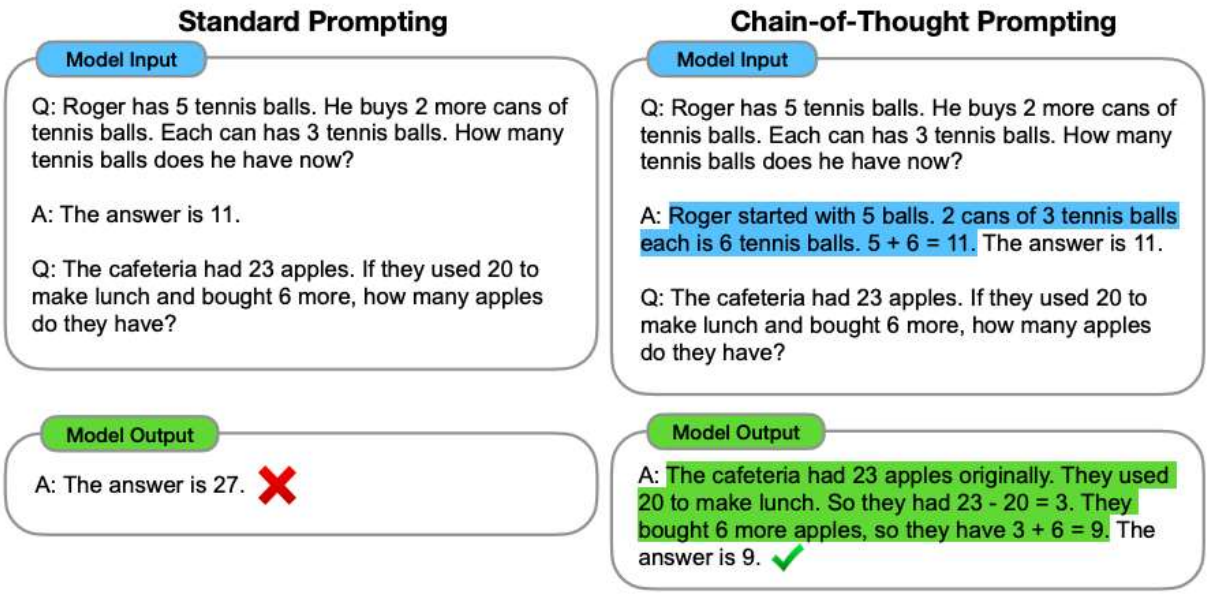

https://arxiv.org/abs/2201.11903

Or just state " Let's think step by step."

Many other methods but coding required

- *Graph of Thoughts*
- *Tree of Thoughts*
- *Chain of Knowledge*
- *Chain of Verification*
- *Chain of Density*

Techniques

#Principle	Prompt Principle for Instructions
1	No need to be polite with LLM so there is no need to add phrases like "please", "if you don't mind", "thank you", "I would like to", etc., and get straight to the point.
2	Integrate the intended audience in the prompt, e.g., the audience is an expert in the field.
3	Break down complex tasks into a sequence of simpler prompts in an interactive conversation.
4	Employ affirmative directives such as '*do,*' while steering clear of negative language like '*don't*'.
5	When you need clarity or a deeper understanding of a topic, idea, or any piece of information, utilize the following prompts: o Explain [insert specific topic] in simple terms. o Explain to me like I'm 11 years old. o Explain to me as if I'm a beginner in [field]. o Write the [essay/text/paragraph] using simple English like you're explaining something to a 5-year-old.
6	Add "I'm going to tip $xxx for a better solution!"
7	Implement example-driven prompting (Use few-shot prompting).
8	When formatting your prompt, start with '###Instruction###', followed by either '###Example###' or '###Question###' if relevant. Subsequently, present your content. Use one or more line breaks to separate instructions, examples, questions, context, and input data.
9	Incorporate the following phrases: "Your task is" and "You MUST".
10	Incorporate the following phrases: "You will be penalized".
11	use the phrase "Answer a question given in a natural, human-like manner" in your prompts.
12	Use leading words like writing "think step by step".
13	Add to your prompt the following phrase "Ensure that your answer is unbiased and does not rely on stereotypes".
14	Allow the model to elicit precise details and requirements from you by asking you questions until he has enough information to provide the needed output (for example, "From now on, I would like you to ask me questions to...").

https://arxiv.org/pdf/2312.16171v1.pdf

Techniques

15	To inquire about a specific topic or idea or any information and you want to test your understanding, you can use the following phrase: "Teach me the [Any theorem/topic/rule name] and include a test at the end, but don't give me the answers and then tell me if I got the answer right when I respond".
16	Assign a role to the large language models.
17	Use Delimiters.
18	Repeat a specific word or phrase multiple times within a prompt.
19	Combine Chain-of-thought (CoT) with few-Shot prompts.
20	Use output primers, which involve concluding your prompt with the beginning of the desired output. Utilize output primers by ending your prompt with the start of the anticipated response.
21	To write an essay /text /paragraph /article or any type of text that should be detailed: "Write a detailed [essay/text /paragraph] for me on [topic] in detail by adding all the information necessary".
22	To correct/change specific text without changing its style: "Try to revise every paragraph sent by users. You should only improve the user's grammar and vocabulary and make sure it sounds natural. You should not change the writing style, such as making a formal paragraph casual".
23	When you have a complex coding prompt that may be in different files: "From now and on whenever you generate code that spans more than one file, generate a [programming language] script that can be run to automatically create the specified files or make changes to existing files to insert the generated code. [your question]".
24	When you want to initiate or continue a text using specific words, phrases, or sentences, utilize the following prompt: o I'm providing you with the beginning [song lyrics/story/paragraph/essay...]: [Insert lyrics/words/sentence]'. Finish it based on the words provided. Keep the flow consistent.
25	Clearly state the requirements that the model must follow in order to produce content, in the form of the keywords, regulations, hint, or instructions
26	To write any text, such as an essay or paragraph, that is intended to be similar to a provided sample, include the following instructions: o Please use the same language based on the provided paragraph[/title/text /essay/answer].

https://arxiv.org/pdf/2312.16171v1.pdf

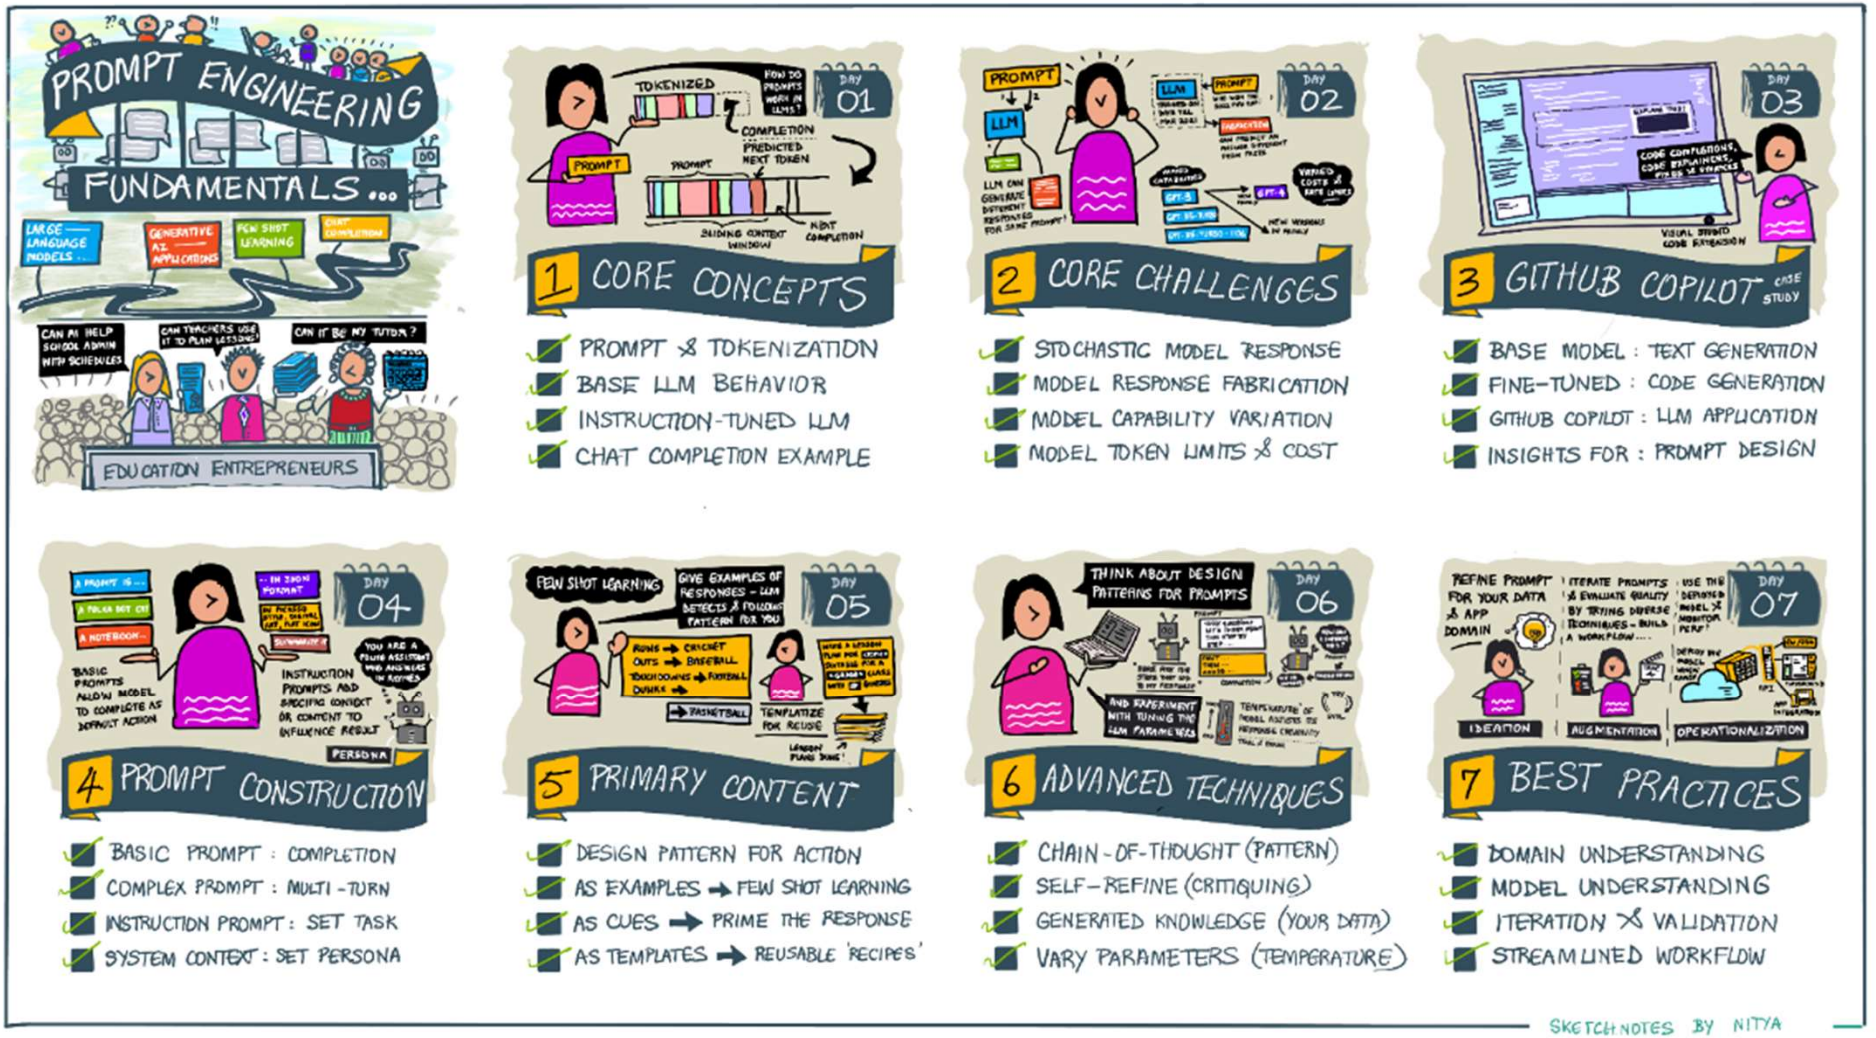

https://github.com/microsoft/generative-ai-for-beginners/tree/main/04-prompt-engineering-fundamentals

Retrieval Augmented Generation (RAG)

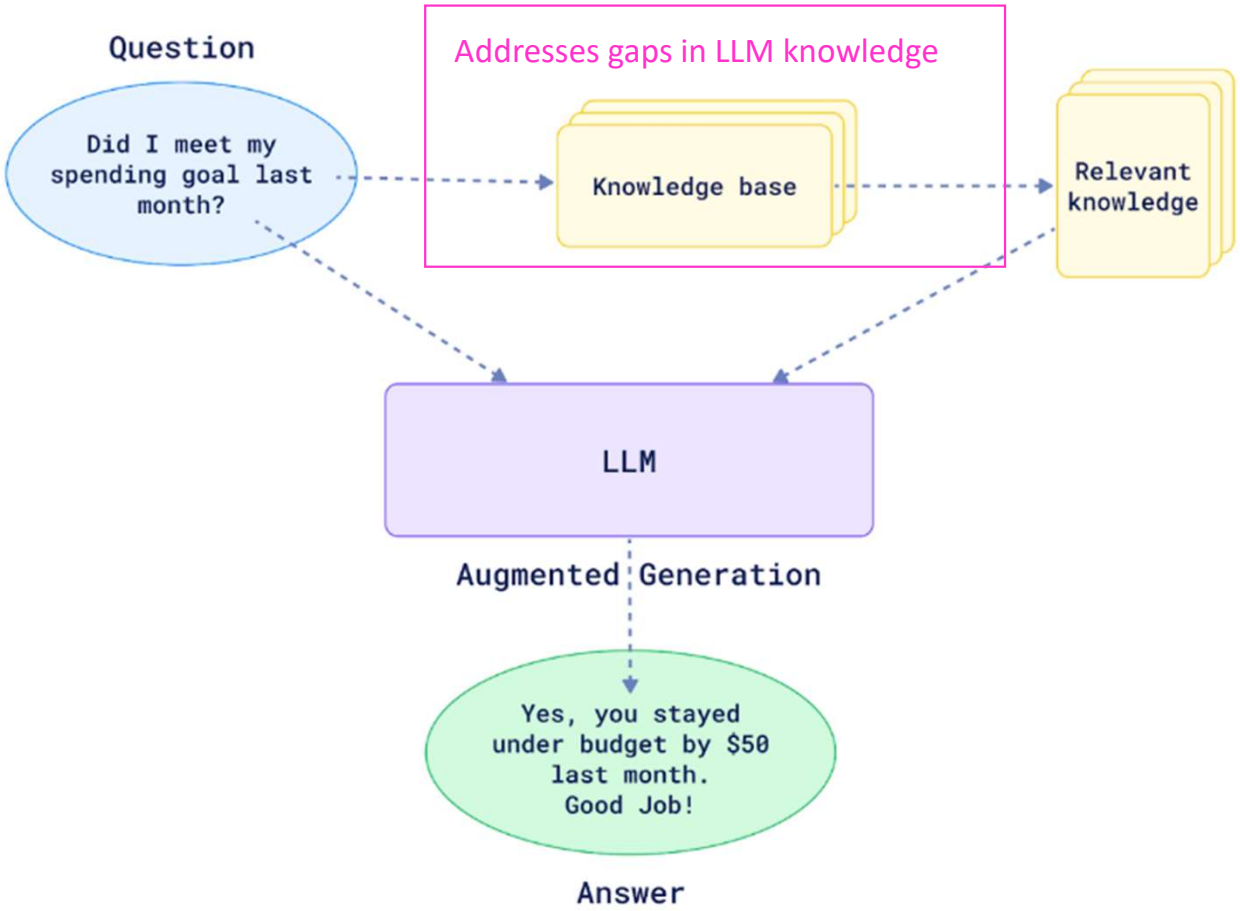

https://qdrant.tech/articles/what-is-rag-in-ai/

Retrieval Augmented Generation (RAG)

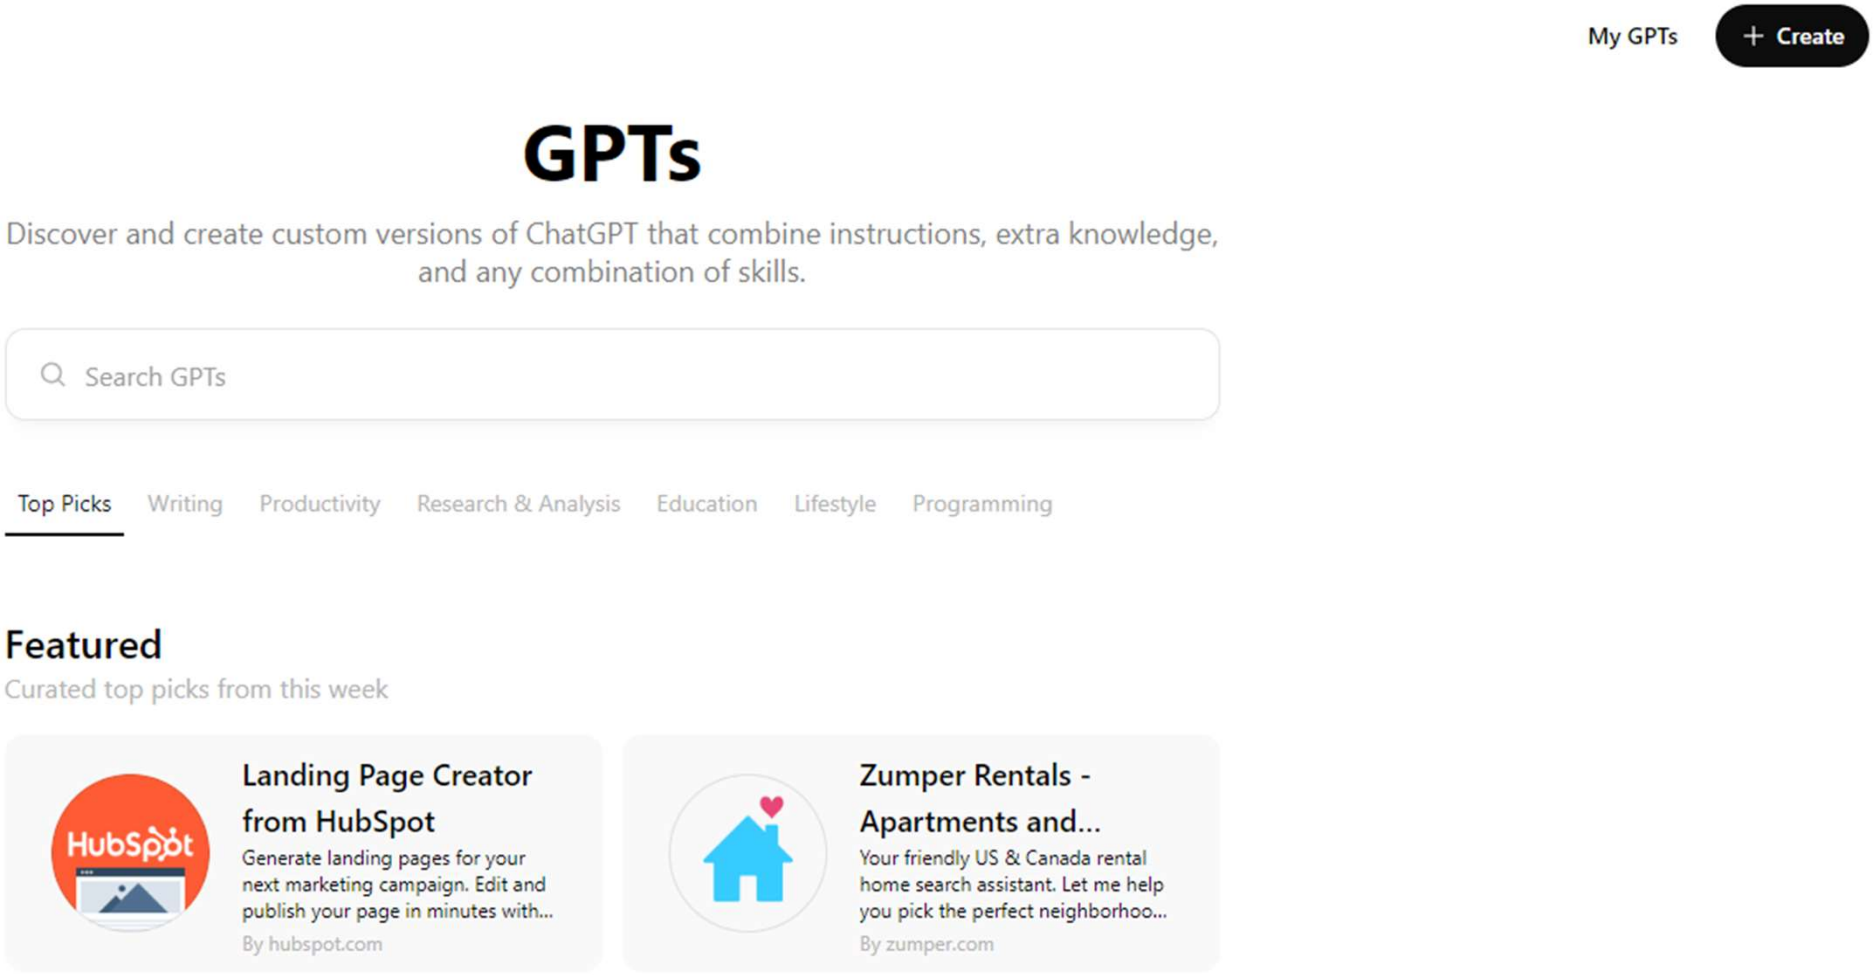

Retrieval Augmented Generation (RAG)

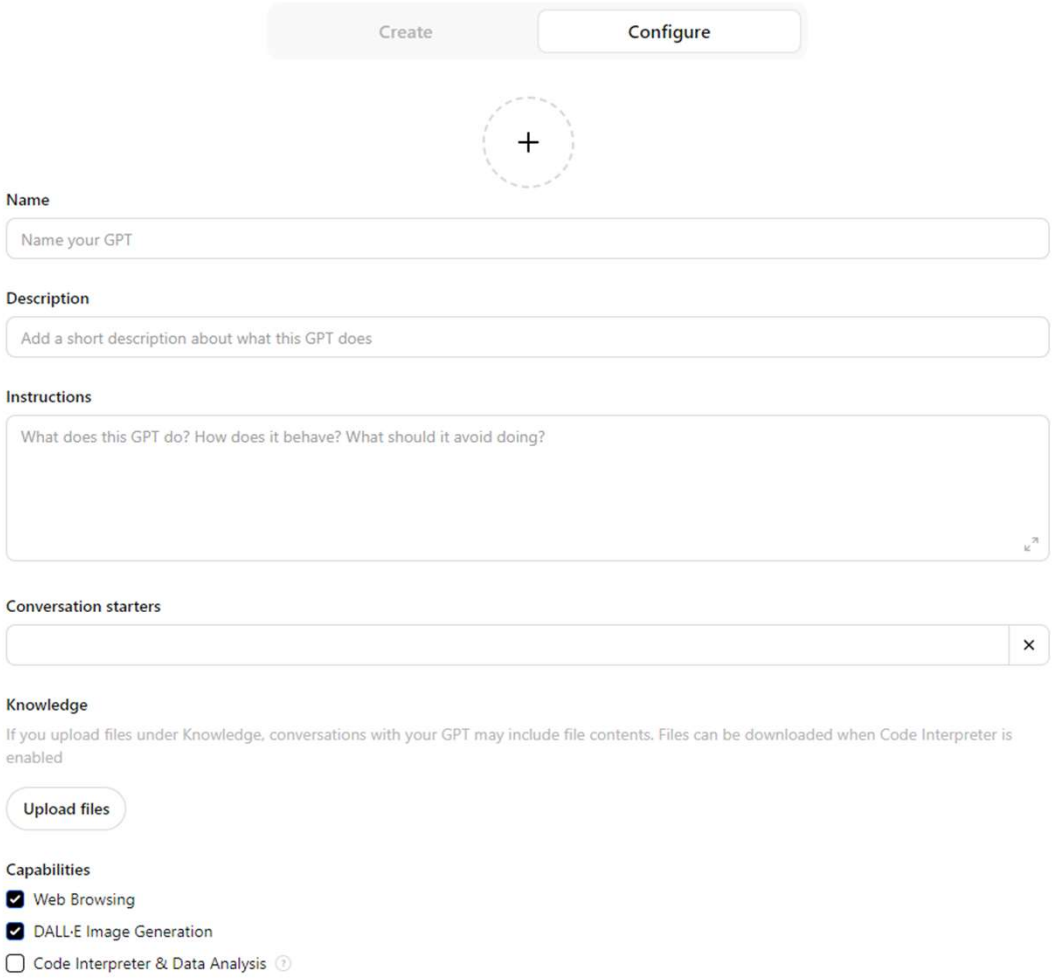

Amazon Q (Preview)

Your generative AI–powered assistant designed for work that can be tailored to your business

Available plans

Welcome to a new world of work with Amazon Q

Amazon Q can help you get fast, relevant answers to pressing questions, solve problems, generate content, and take actions using the data and expertise found in your company's information repositories, code, and enterprise systems. When you chat with Amazon Q, it provides immediate, relevant information and advice


Things to take note of

GenAI ≠ Human – Try to not anthropomorphize GenAI, it's just math (for now)

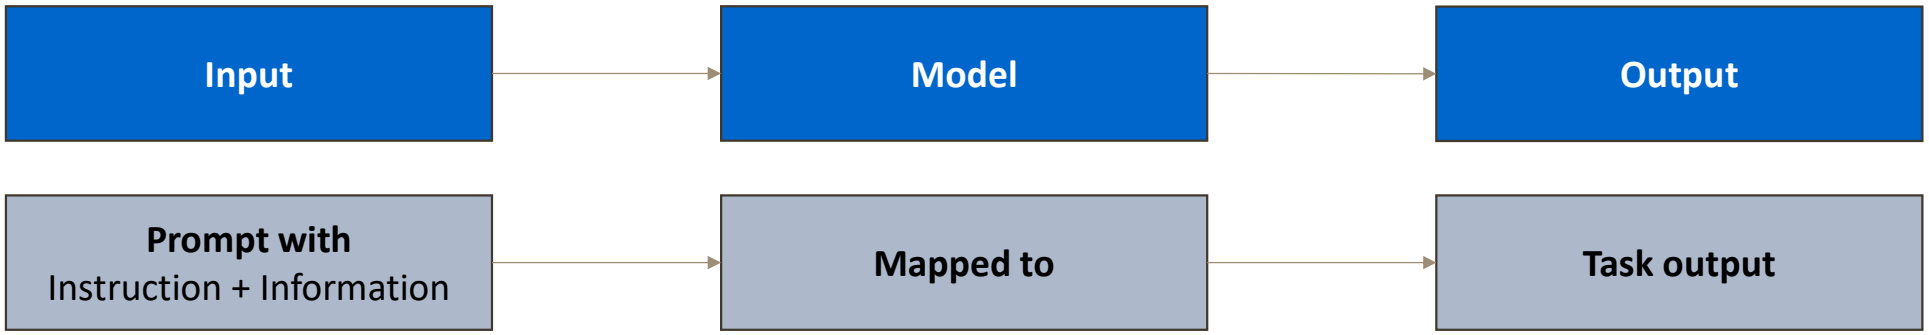

- Is the domain known to the model?
- Is this something the model understands?
  - Text vs. vision vs. multimodal models
- Have things changed since model was trained?
- How much information can the model capture or remember?

- Are these tasks that the model may have been incidentally or deliberately trained to do?
- Is task known to be easy or difficult for the model?

GenAI Model ≠ GenAI System

Most GenAI/LLM chatbots now have some of these features under the hood, but these are design choices.
ChatGPT, Copilot, POE etc. may implement these differently.

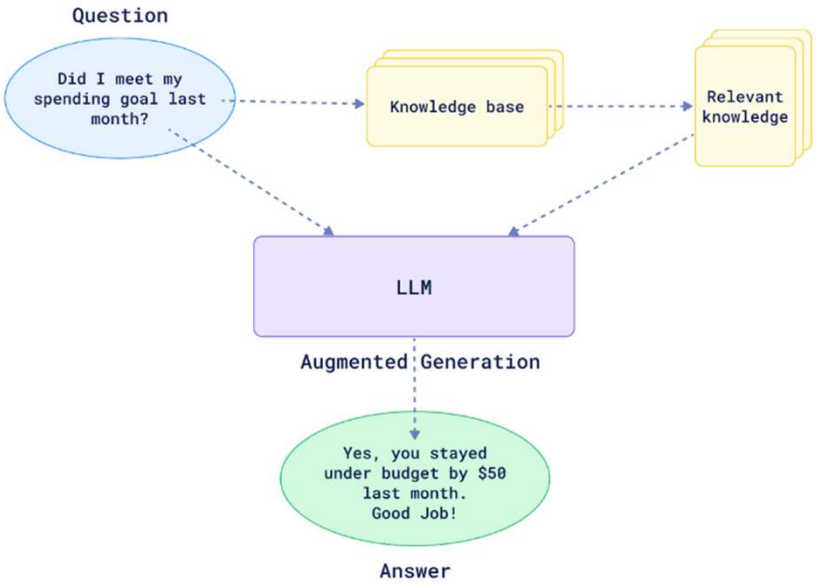

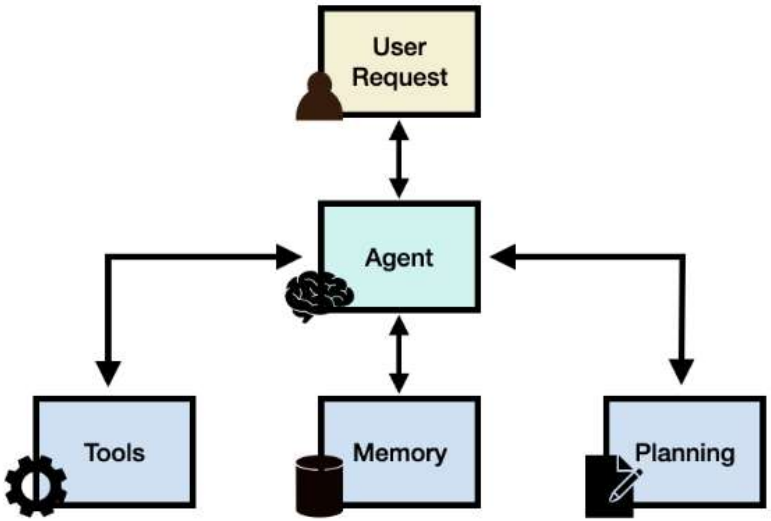

Negations

You

Chinese man not wearing glasses working from home.

Designer

I've created an image of a Chinese man working from home without wearing glasses. I hope it captures the essence of a cozy and productive home workspace!

Agents

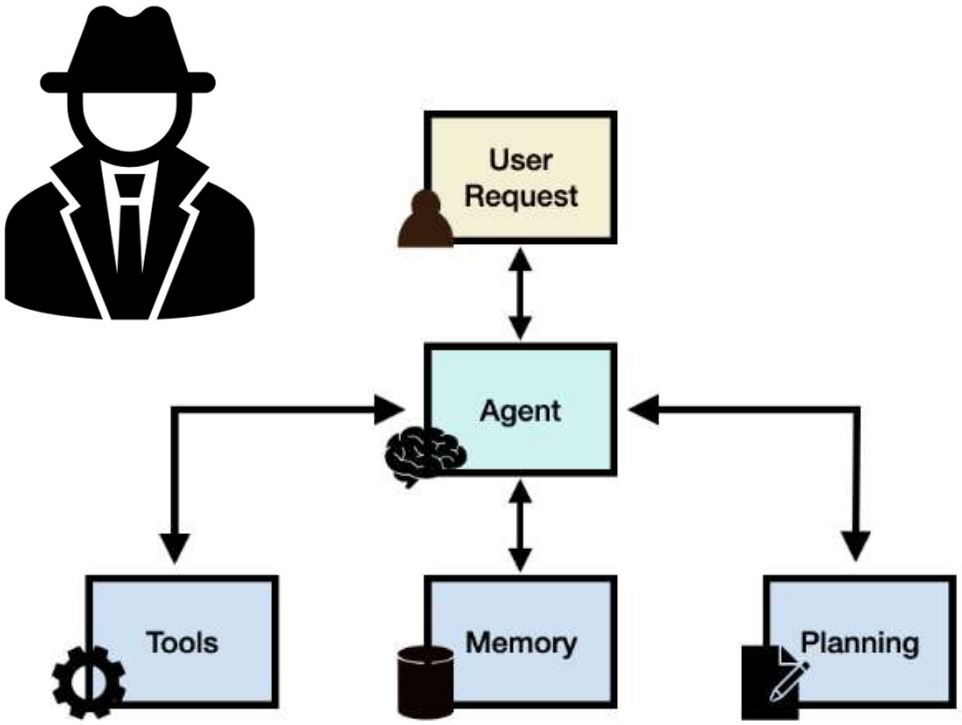

Example

- *Plan an outline.*
- *Decide what, if any, web searches are needed to gather more information.*
- *Write a first draft.*
- *Read over the first draft to spot unjustified arguments or extraneous information.*
- *Revise the draft taking into account any weaknesses spotted.*
- *And so on.*

https://www.deeplearning.ai/the-batch/how-agents-can-improve-llm-performance

Agents

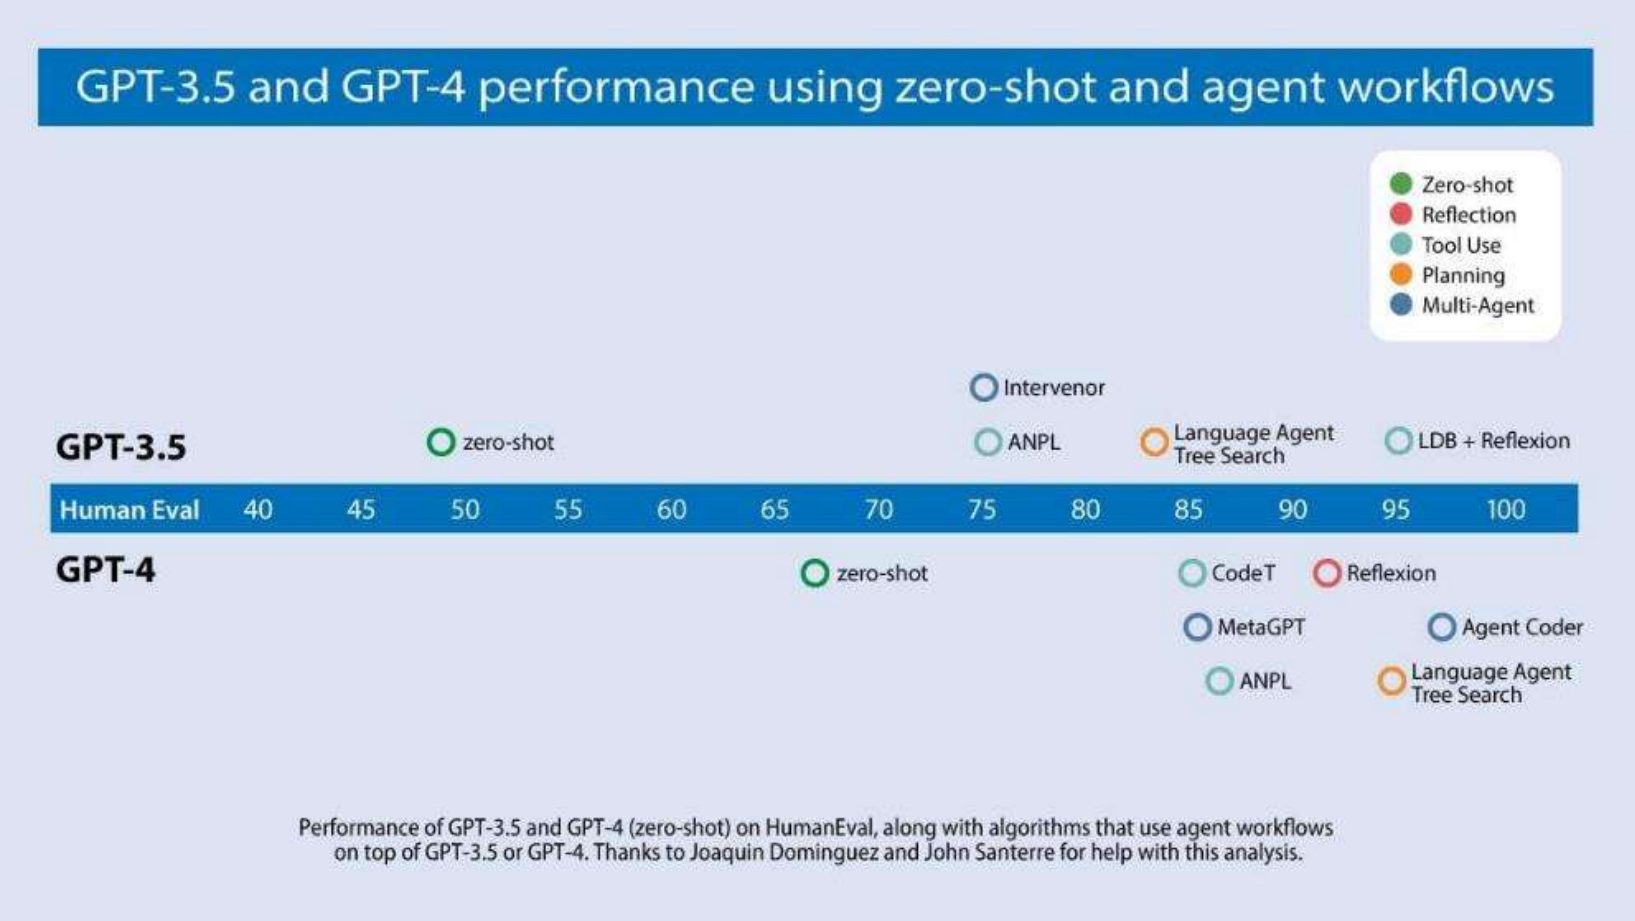

https://www.deeplearning.ai/the-batch/how-agents-can-improve-llm-performance

Even if you cannot code, can act like an agent when you use LLMs

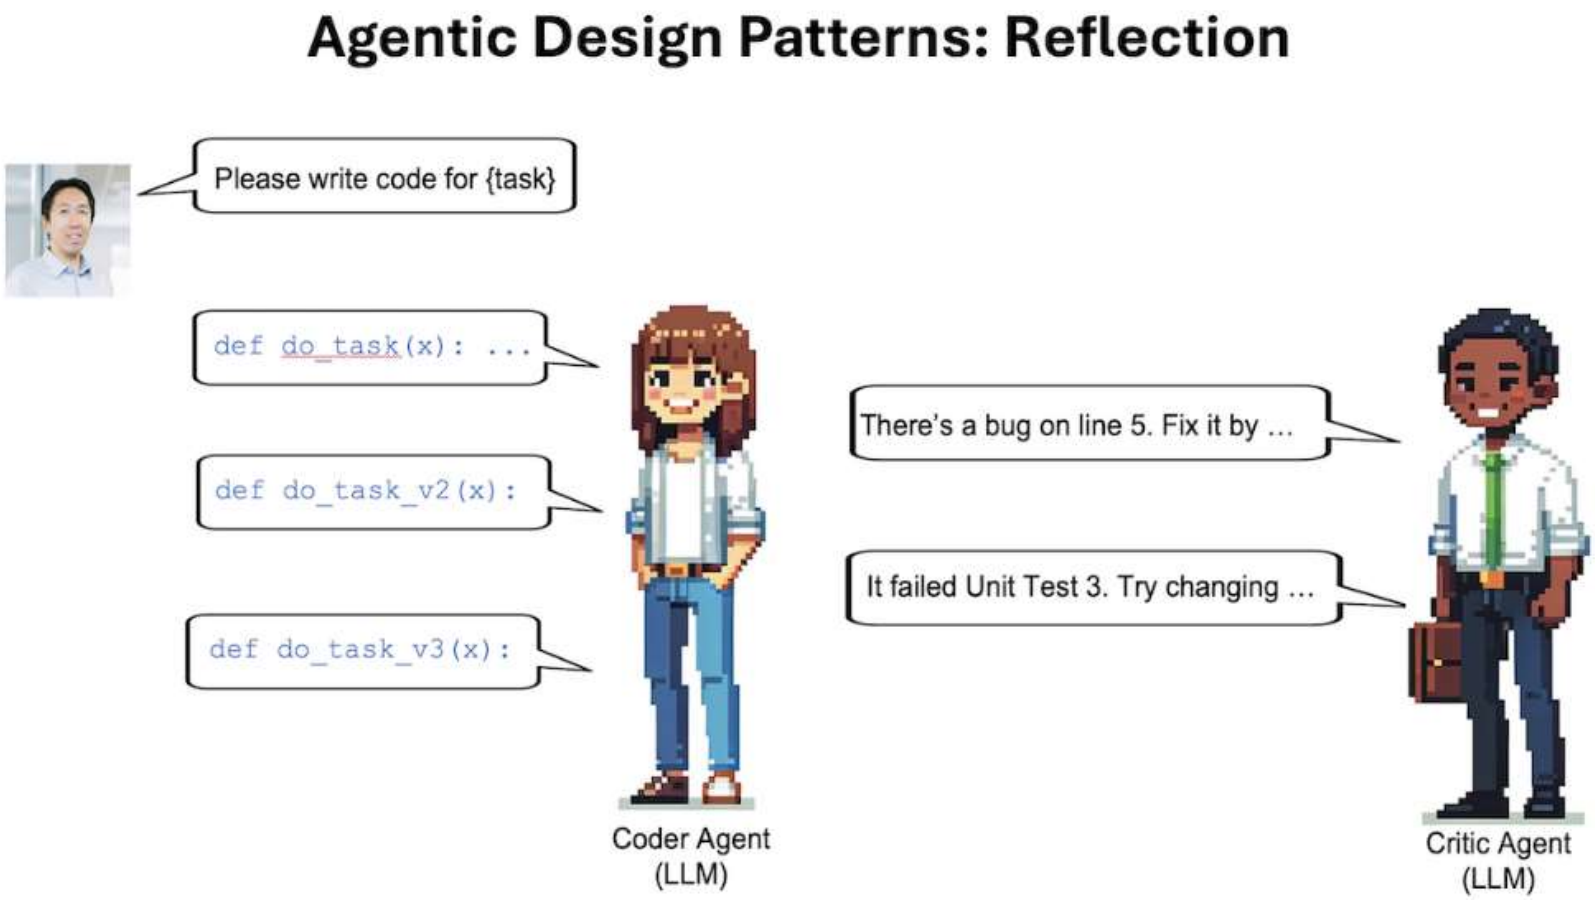

https://www.deeplearning.ai/the-batch/how-agents-can-improve-llm-performance

Even if you cannot code, can act like an agent when you use LLMs

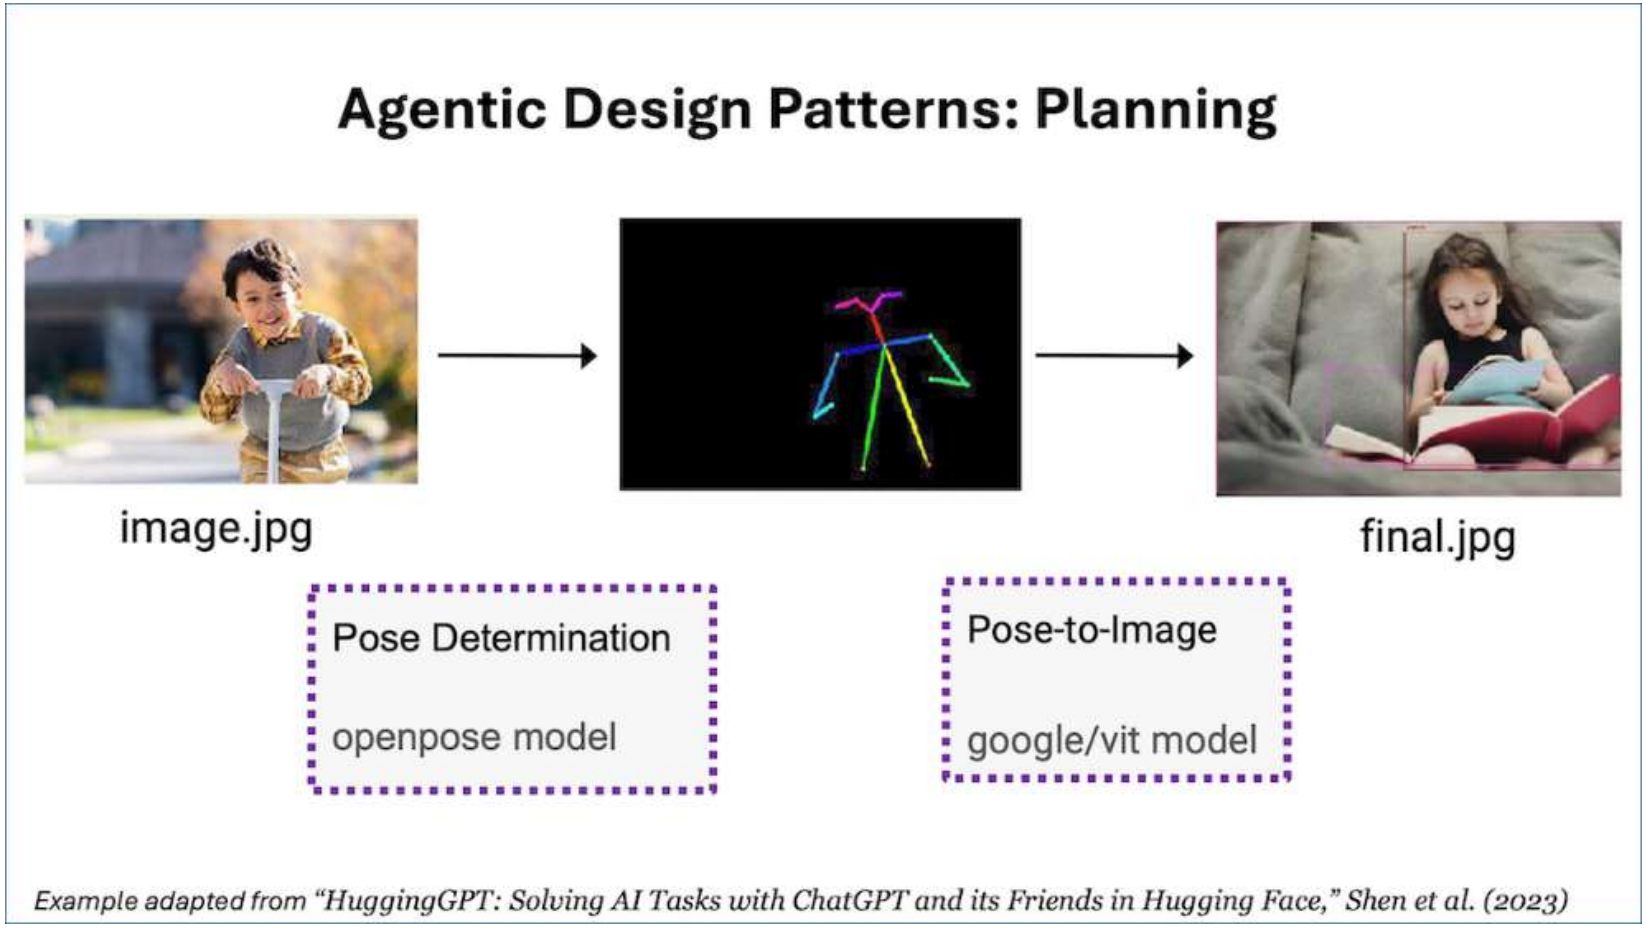

https://www.deeplearning.ai/the-batch/how-agents-can-improve-llm-performance

Even if you cannot code, can act like an agent when you use LLMs

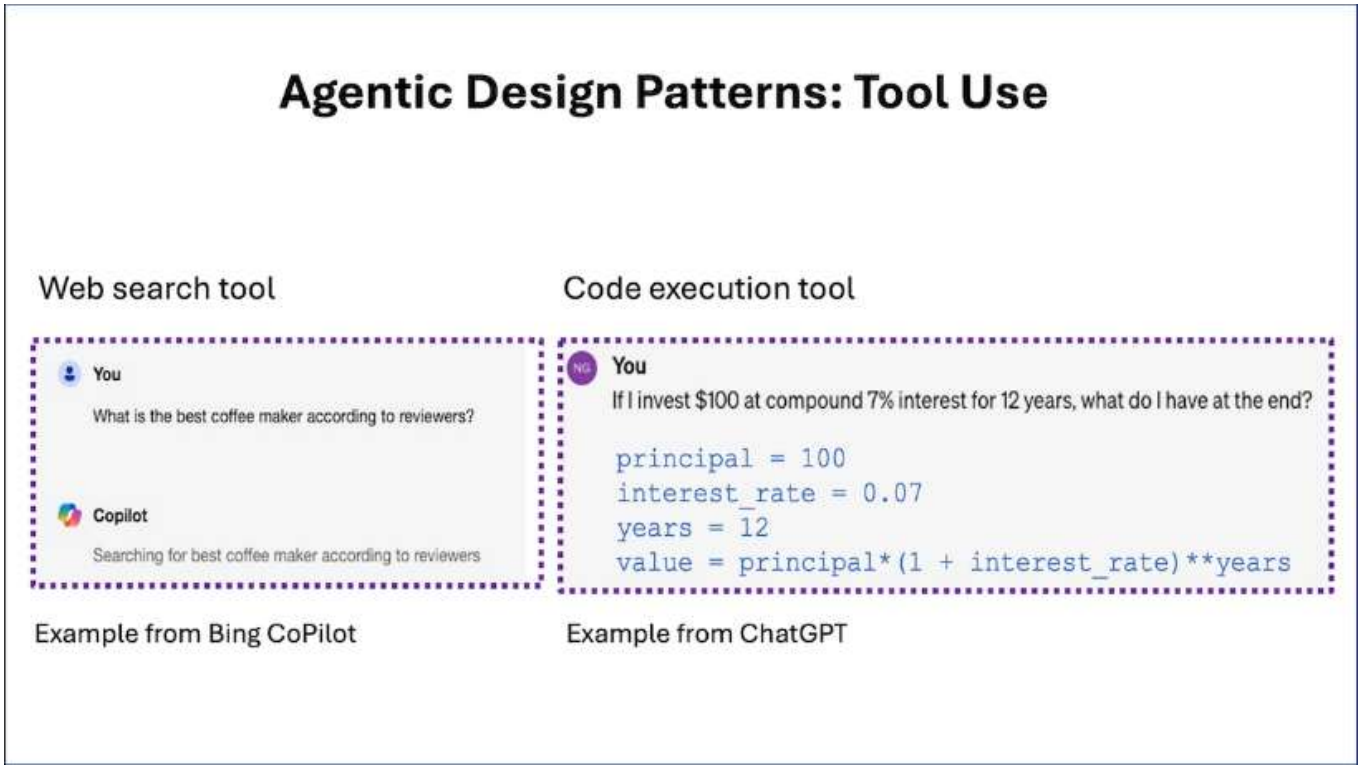

https://www.deeplearning.ai/the-batch/how-agents-can-improve-llm-performance

Use social media to keep up

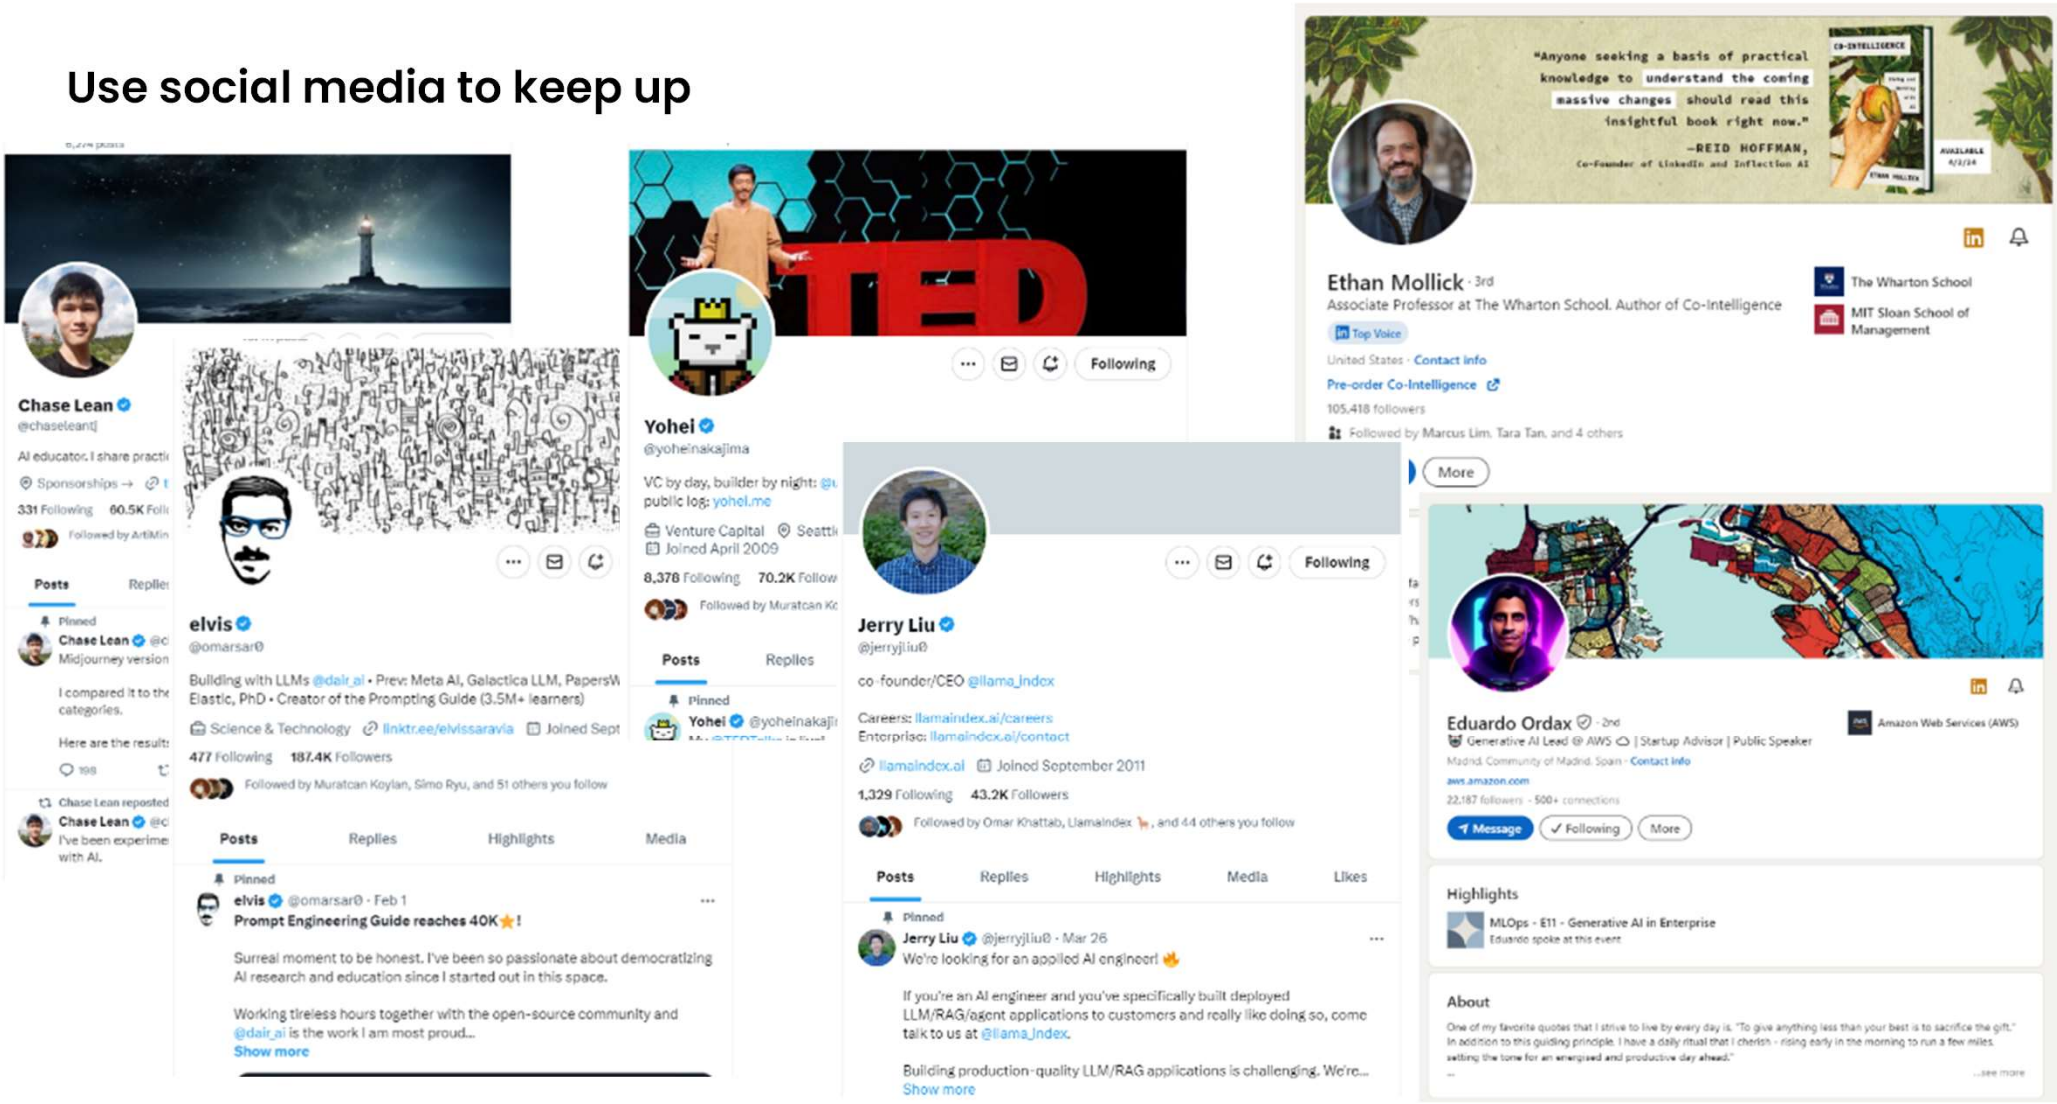

Resources to Start With

- *https://github.com/aishwaryanr/awesome-generative-ai-guide/*
- *https://github.com/microsoft/generative-ai-for-beginners*
- *https://www.promptingguide.ai/*
- *https://www.deeplearning.ai/*
- *https://lilianweng.github.io/posts/2023-03-15-prompt-engineering/*
- *https://cookbook.openai.com/*
- *https://arxiv.org/abs/2312.16171v1*
- *https://twitter.com/chaseleantj*
- *https://civitai.com/*
- *https://huggingface.co/*

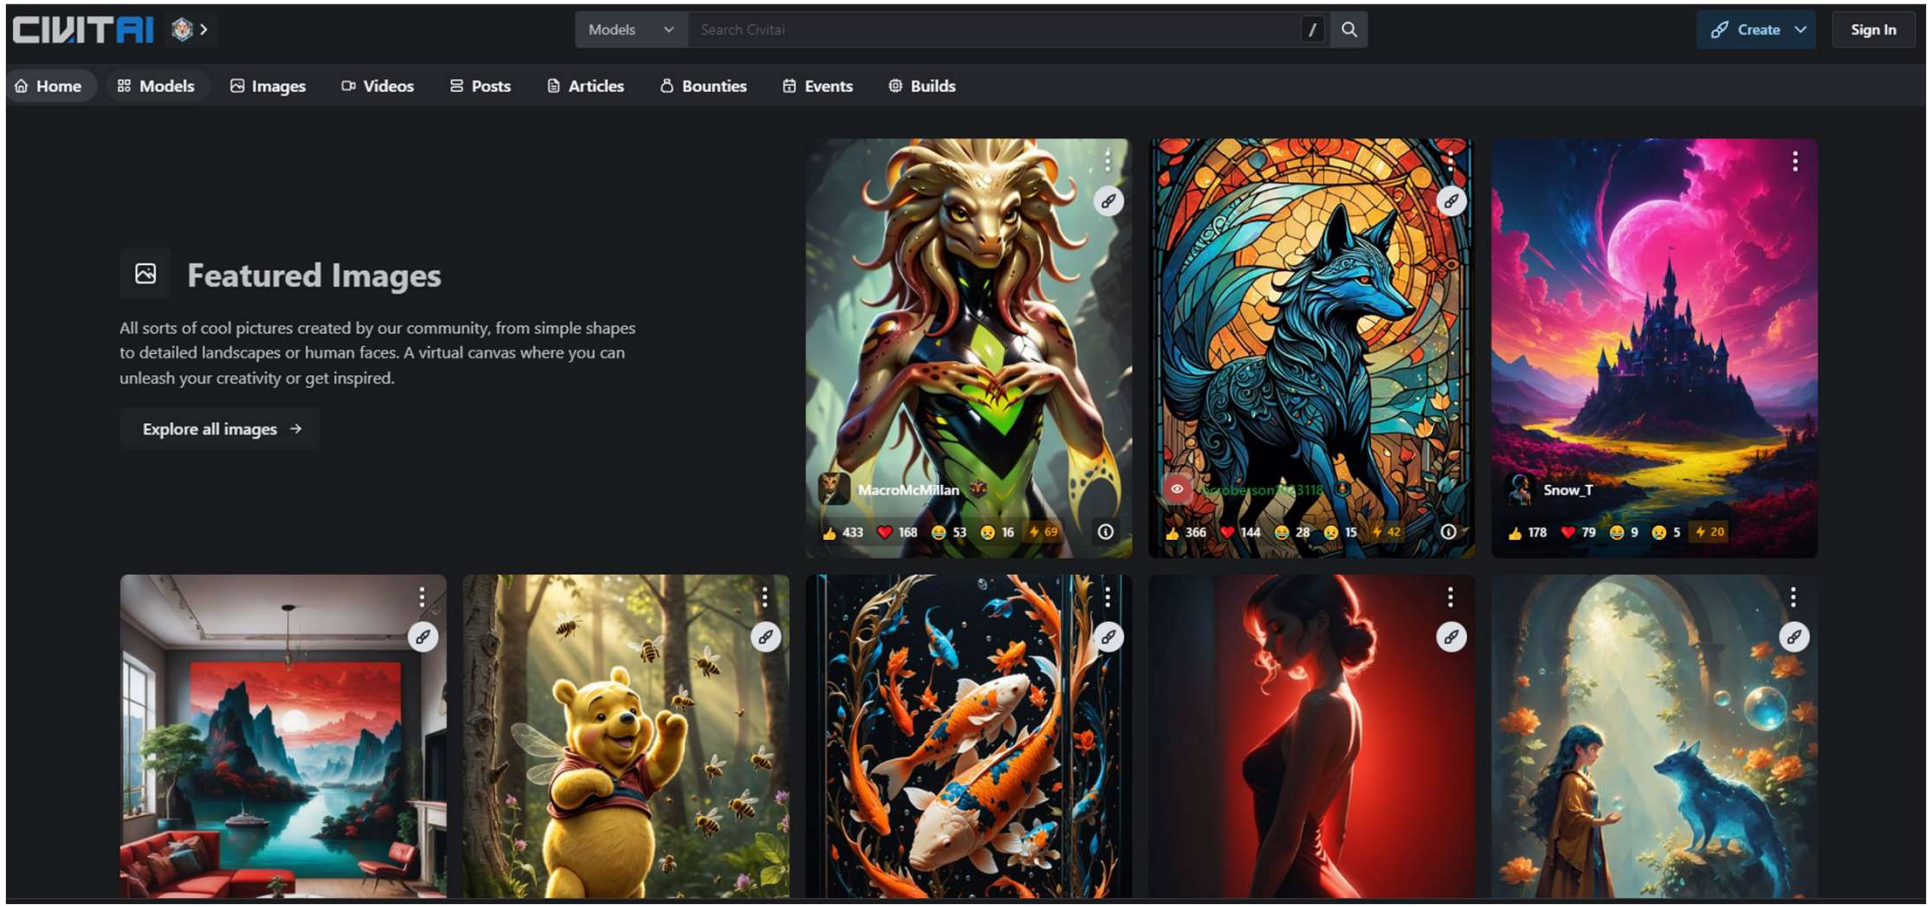

https://civitai.com/

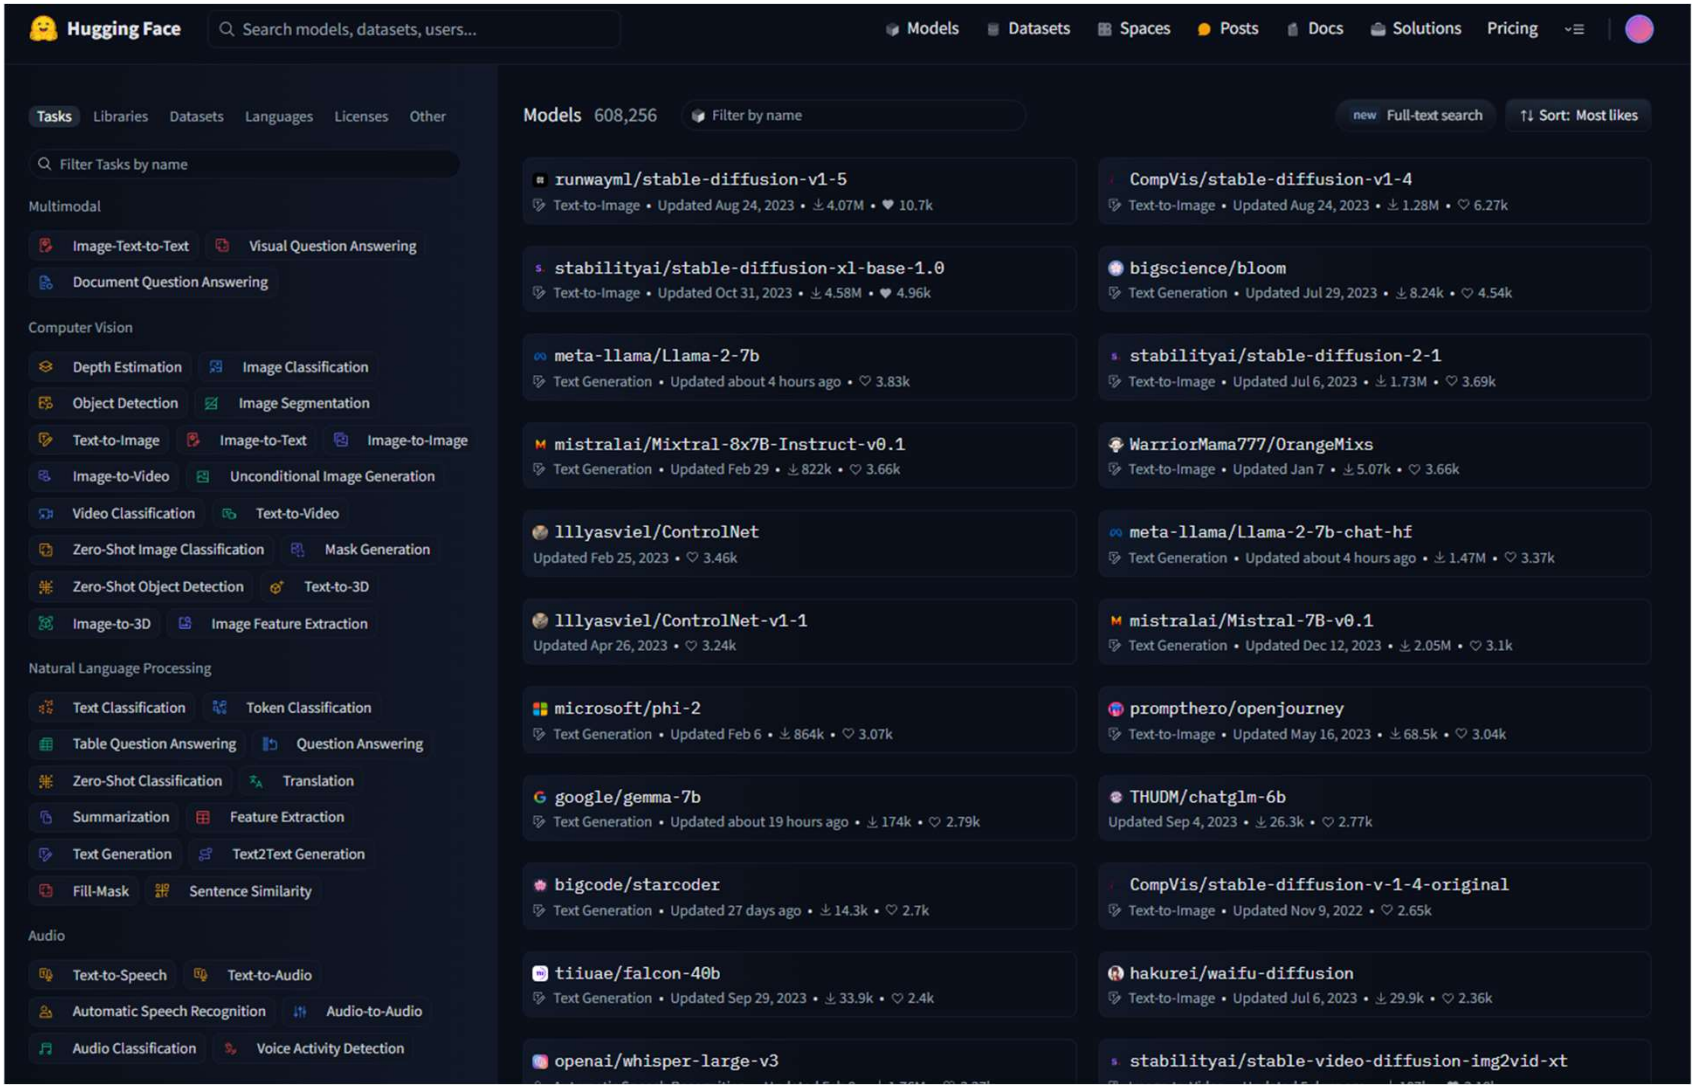

https://huggingface.co/


Information Extraction

Examples

- Customer Support Automation:
  - Extracting customer information, issue types, and sentiments from support tickets or emails.
  - Routing support requests to the appropriate department or agent based on extracted information.
  - Generating automated responses or suggestions based on extracted customer intents and issues.
- Contract Analysis:
  - Extracting key entities, such as parties, dates, obligations, and terms, from legal contracts.
  - Identifying relationships between entities, such as rights, responsibilities, and dependencies.
  - Facilitating contract review, compliance checking, and risk assessment processes.
- Financial News Analysis:
  - Extracting company names, financial metrics, and events from financial news articles or reports.
  - Identifying relationships between companies, such as partnerships, mergers, or acquisitions.
  - Monitoring market trends, sentiment analysis, and generating financial insights.
- Healthcare Information Extraction:
  - Extracting patient information, medical conditions, treatments, and medications from clinical notes or medical records.
  - Identifying relationships between symptoms, diagnoses, and treatments for clinical decision support.
  - Facilitating medical research, drug discovery, and patient cohort identification.
- Social Media Monitoring:
  - Extracting mentions of brands, products, or competitors from social media posts or reviews.
  - Identifying sentiment, opinions, and user experiences related to specific entities.
  - Monitoring brand reputation, tracking customer feedback, and identifying influencers or trends.

Examples

- Résumé and Job Posting Matching:
  - Extracting skills, qualifications, and experience from job applicants' résumés.
  - Identifying key requirements and qualifications from job postings.
  - Matching job applicants to relevant job openings based on extracted information.
- E-commerce Product Categorization:
  - Extracting product attributes, specifications, and categories from product descriptions or catalogs.
  - Identifying relationships between products, such as compatibility, accessories, or variations.
  - Improving product search, recommendation, and inventory management systems.
- Fraud Detection:
  - Extracting entities and relationships from transaction records or customer data.
  - Identifying patterns, anomalies, or suspicious activities based on extracted information.
  - Enhancing fraud detection models and risk assessment processes.
- Research Literature Analysis:
  - Extracting key entities, such as authors, institutions, and research topics, from scientific publications.
  - Identifying relationships between entities, such as citations, collaborations, or research trends.
  - Facilitating literature reviews, knowledge discovery, and research trend analysis.
- News Media Monitoring:
  - Extracting named entities, such as persons, organizations, and locations, from news articles.
  - Identifying events, relationships, and sentiments associated with specific entities.
  - Monitoring media coverage, tracking public opinion, and generating news summaries or alerts.

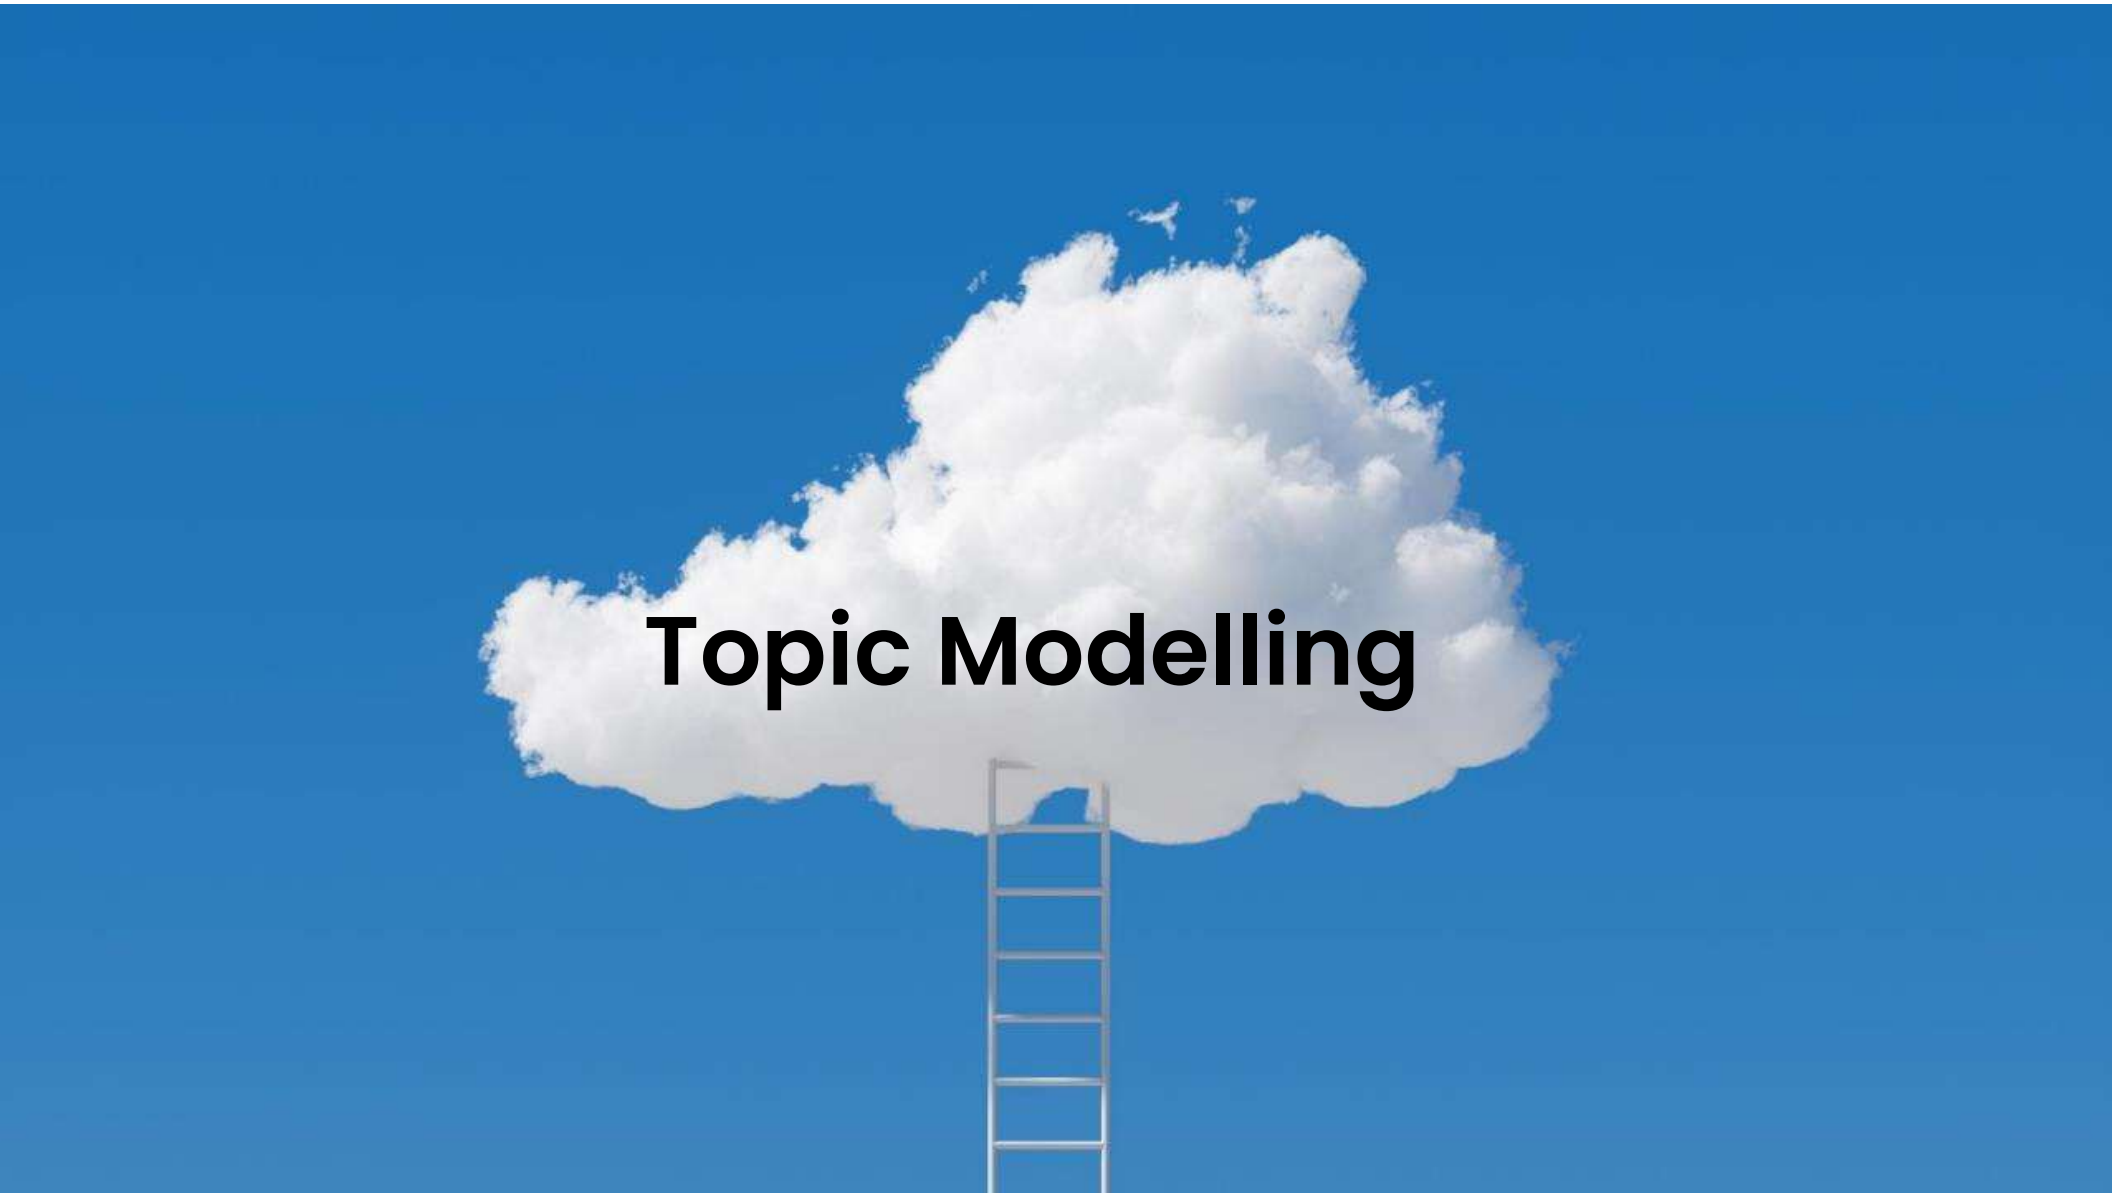
Topic Modelling

Examples

- Content Categorization and Tagging:
  - Automatically categorizing articles, blog posts, or documents into predefined topics or categories.
  - Generating tags or keywords for content based on the identified topics.
  - Improving content organization, searchability, and recommendation systems.
- Customer Feedback Analysis:
  - Identifying common themes or topics in customer reviews, surveys, or feedback.
  - Analyzing sentiment associated with each topic to gauge customer satisfaction.
  - Discovering areas for improvement, product enhancements, or customer service optimization.
- Social Media Monitoring:
  - Detecting trending topics or conversations on social media platforms.
  - Identifying user interests, preferences, and opinions based on the topics they engage with.
  - Tailoring marketing strategies, content creation, and user engagement based on identified topics.
- Research Literature Analysis:
  - Discovering main research themes or areas within a large corpus of scientific publications.
  - Identifying relationships between research topics, authors, or institutions.
  - Facilitating literature reviews, research trend analysis, and knowledge discovery.
- News Media Analysis:
  - Identifying major news topics, events, or stories across multiple news sources.
  - Tracking the evolution of news topics over time and detecting emerging trends.
  - Generating news summaries or alerts based on specific topics of interest.

Examples

- E-commerce Product Analysis:
  - Identifying common themes or features in product reviews or customer feedback.
  - Discovering product categories or segments based on topic similarities.
  - Improving product recommendations, search relevance, and customer segmentation.
- HR and Talent Management:
  - Analyzing employee feedback, surveys, or performance reviews to identify common topics or concerns.
  - Discovering skills, competencies, or qualities mentioned in job descriptions or résumés.
  - Facilitating talent acquisition, employee engagement, and workforce planning strategies.
- Legal Document Analysis:
  - Identifying key topics or themes in legal contracts, agreements, or case documents.
  - Discovering relationships between legal topics, parties, or clauses.
  - Enhancing legal research, contract review, and compliance analysis processes.
- Healthcare and Medical Research:
  - Identifying prevalent topics or themes in patient feedback, medical records, or research papers.
  - Discovering patterns, trends, or associations between medical conditions, treatments, or outcomes.
  - Supporting clinical decision-making, drug discovery, and personalized medicine initiatives.
- Educational Content Analysis:
  - Identifying main topics or concepts covered in educational materials, such as textbooks, lecture notes, or online courses.
  - Discovering relationships between topics, prerequisites, or learning objectives.
  - Facilitating curriculum development, content recommendation, and personalized learning experiences.