EE P 596
LLMs: From Transformers to ChatGPT

Introduction | LLM Motivation | History of LLMs

Dr. Karthik Mohan, Jan 5 2026 | Winter Quarter course | PMP, ECE, UW

Bit about Me

Bit about Me

Phd, ECE, UW

Bit about Me

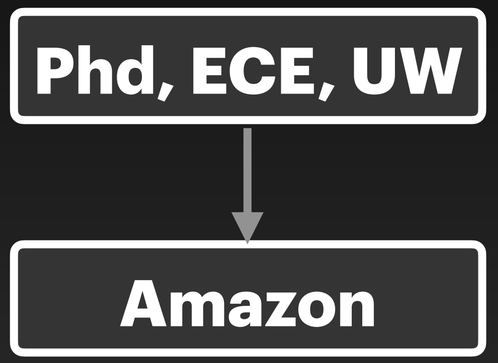

Bit about Me

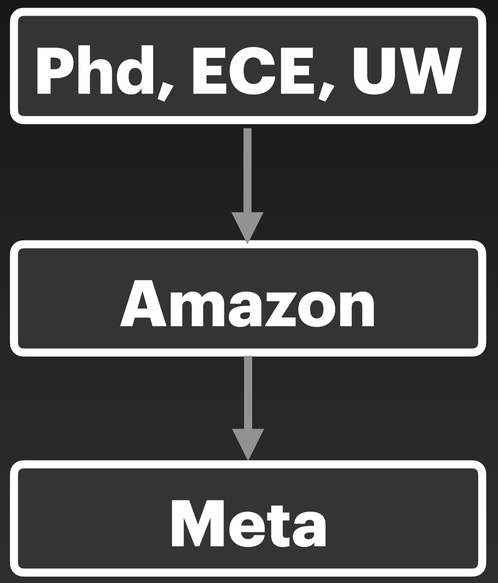

Bit about Me

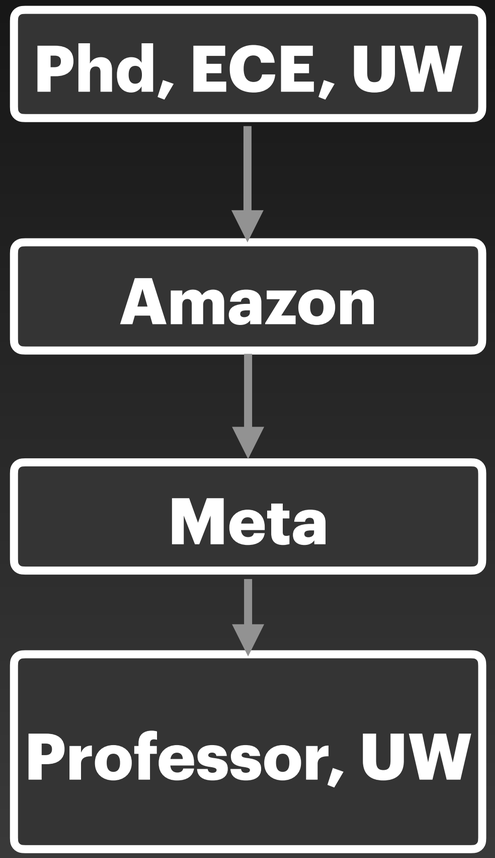

Bit about Me

Phd, ECE, UW

Amazon

Meta

Professor, UW

Staff Scientist, eBay

Teaching Support Team

Arnav Das (TA)

Urvashi Taki (TA)

Haofei Ding (Grader)

ChatGPT and LLMs are everywhere!

ChatGPT and LLMs are everywhere!

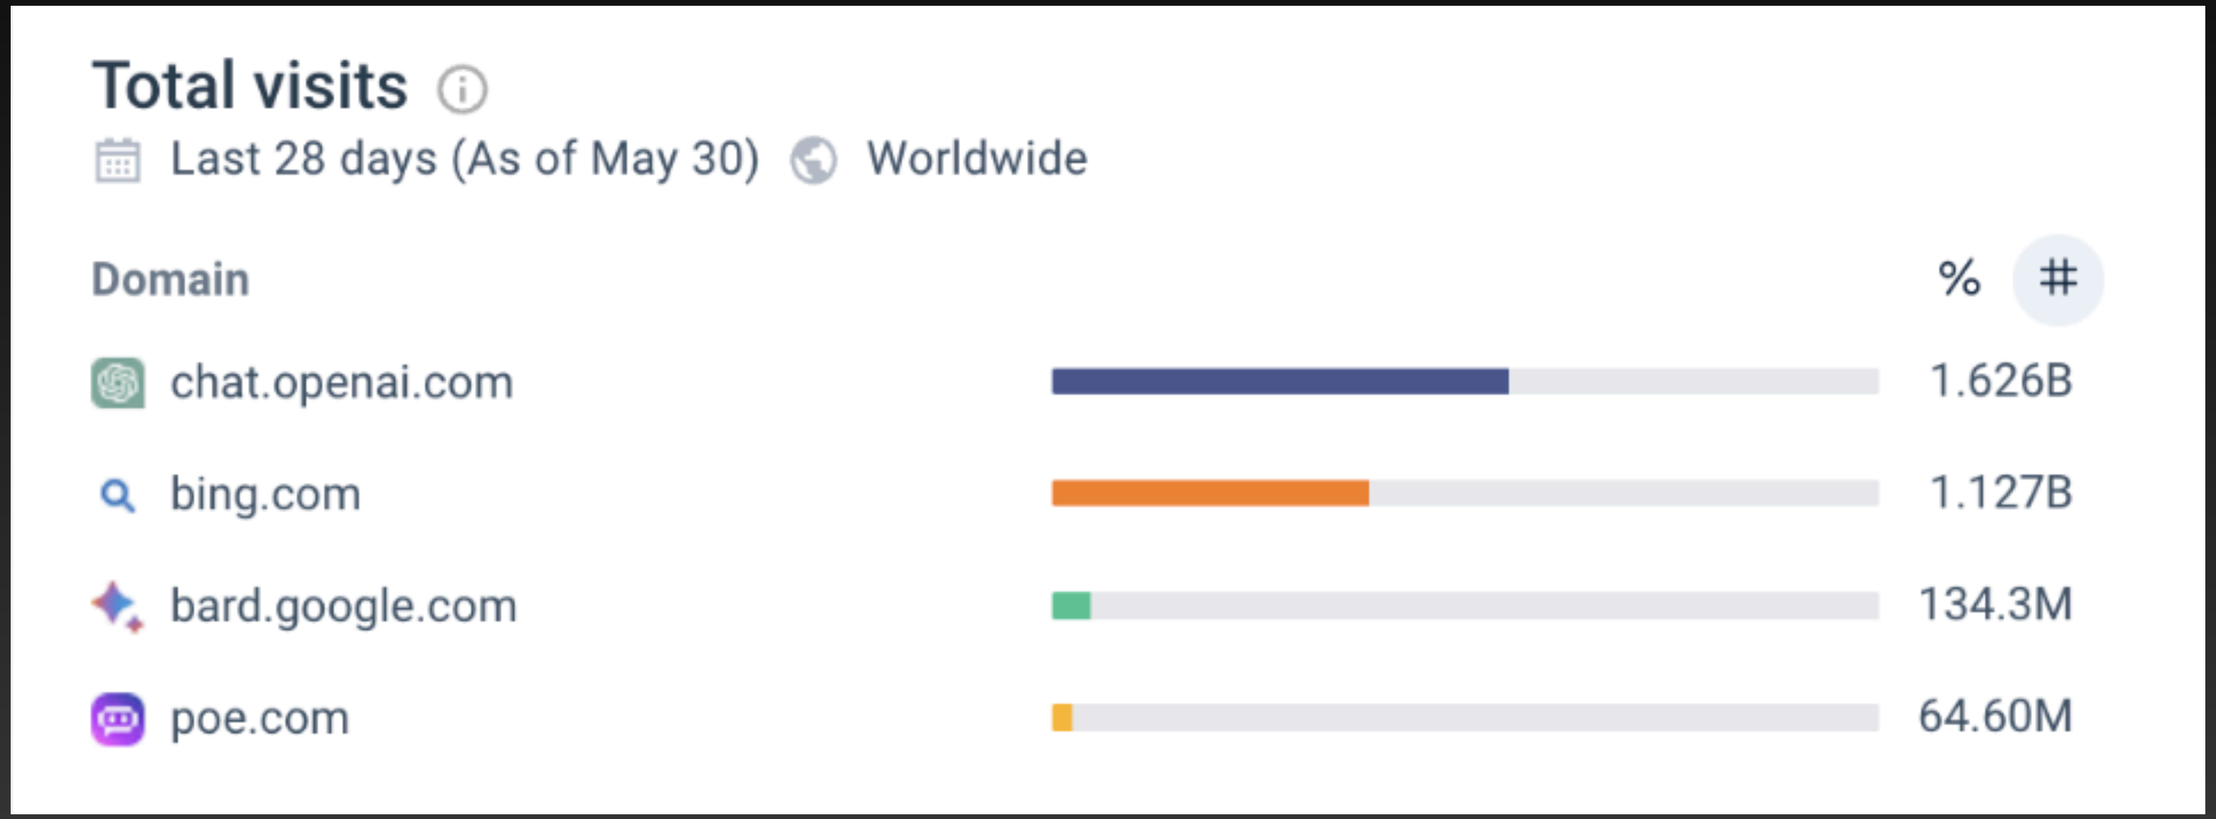

ChatGPT and LLMs are everywhere!

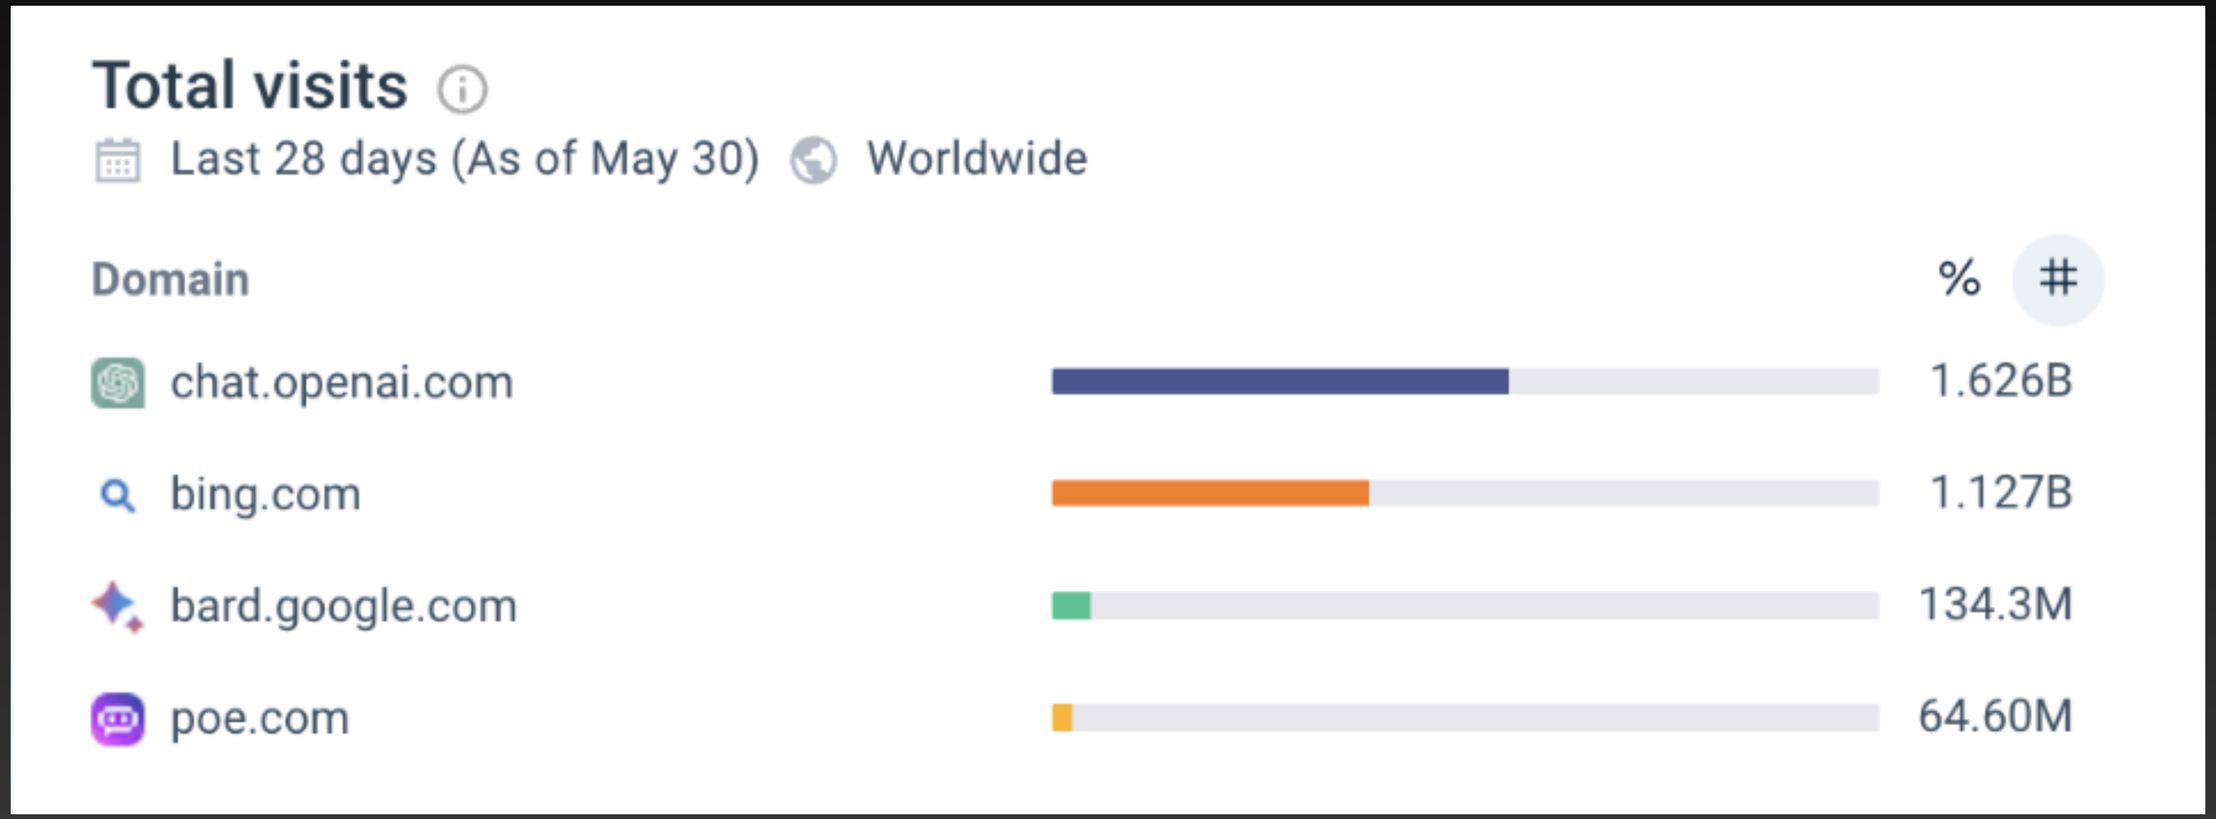

Super Intelligence

Computer scientist Yann LeCun: ‘Intelligence really is about learning’

ft.com

ChatGPT and LLMs are everywhere!

Let’s look at some examples!

ChatGPT and LLMs are everywhere!

Paraphrasing

ChatGPT and LLMs are everywhere!

Paraphrasing

Math

ChatGPT and LLMs are everywhere!

Paraphrasing

Math

Coding

ChatGPT and LLMs are everywhere!

Sentiment Annotation

Data Augmentation

LLM as a judge

ChatGPT and LLMs are everywhere!

Let’s go checkout ChatGPT live!

Business Applications

Looking for specific info?

Ask Rufus or search reviews and Q&A

Does it have a wax coating? | Is this floss good for sensitive gums? | Can this floss be used with braces? | Does it come in different flavors? | Is it easy to grip?

Customer reviews

4.2 out of 5

4,163 global ratings

5 star 67%
4 star 14%
3 star 7%
2 star 4%
1 star 8%

How customer reviews and ratings work

Review this product

Share your thoughts with other customers

Write a customer review

Customers say

Customers appreciate the dental floss's minty flavor and consider it good value for money, with one noting it's a great discount dupe for Oral-B floss. However, the floss receives mixed feedback about its effectiveness and durability - while some say it works well and doesn't break or shred, others report it shreds easily during use. Moreover, customers find the floss slippery and difficult to hold onto, and criticize its thickness, noting it's too thin for thorough cleaning.

Generated from the text of customer reviews

Select to learn more

Value for money (175) | Taste (74) | Quality (278) | Durability (181) | Effectiveness (151) | Slippage (142) | Slipperiness (135) | Thickness (131)

Reviews with images

See all photos

Responsible AI

https://www.nbcnews.com/tech/tech-news/ai-toys-gift-present-safe-kids-robot-child-miko-grok-alilo-miiloo-rcna246956

Engine behind ChatGPT

ChatGPT heavily relies on Large Language Models to power its responses to users!

How do you understand ChatGPT?

To Understand ChatGPT?

How do you understand ChatGPT?

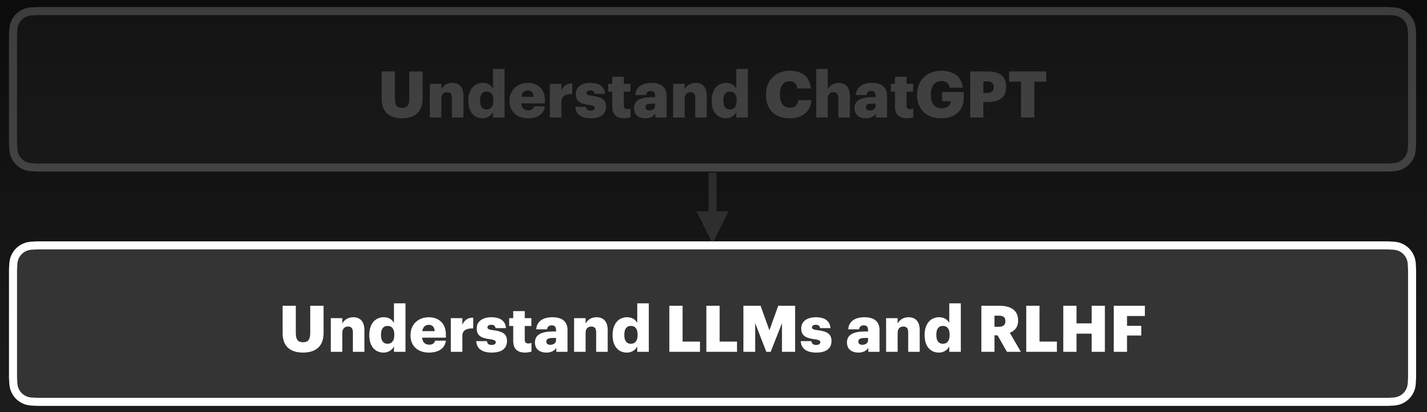

How do you understand ChatGPT?

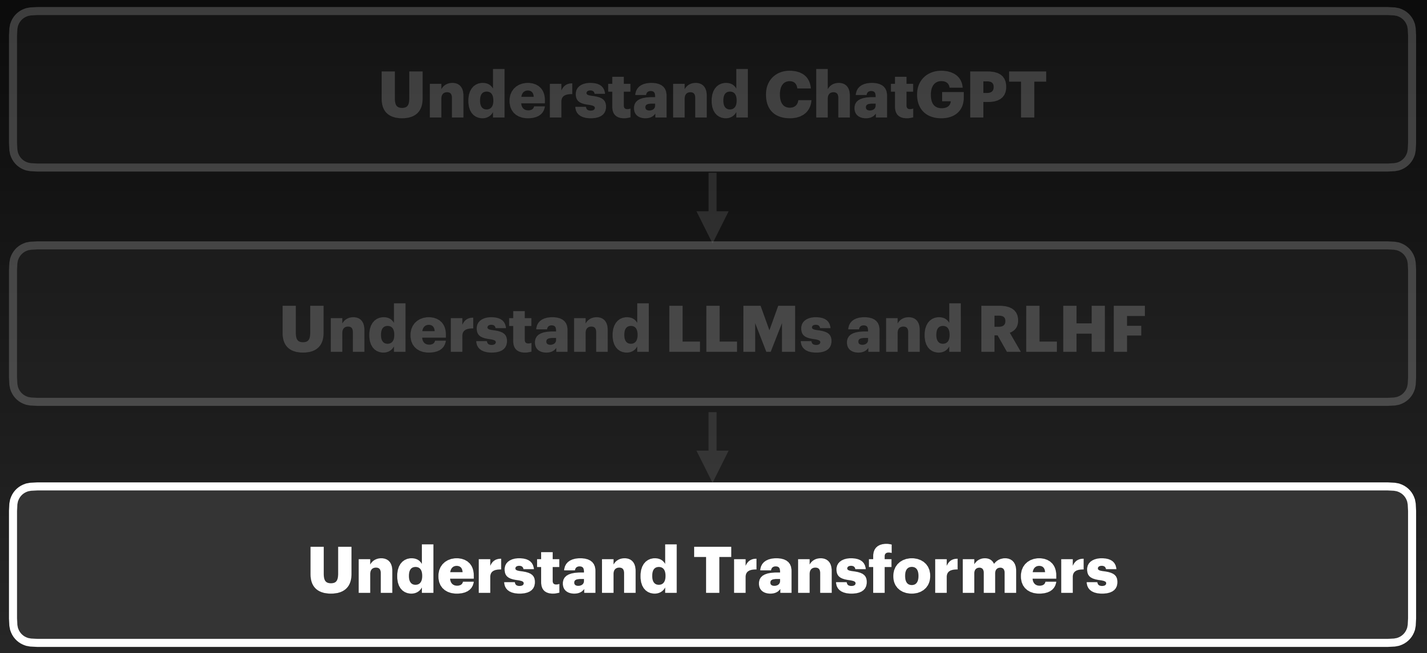

How do you understand ChatGPT?

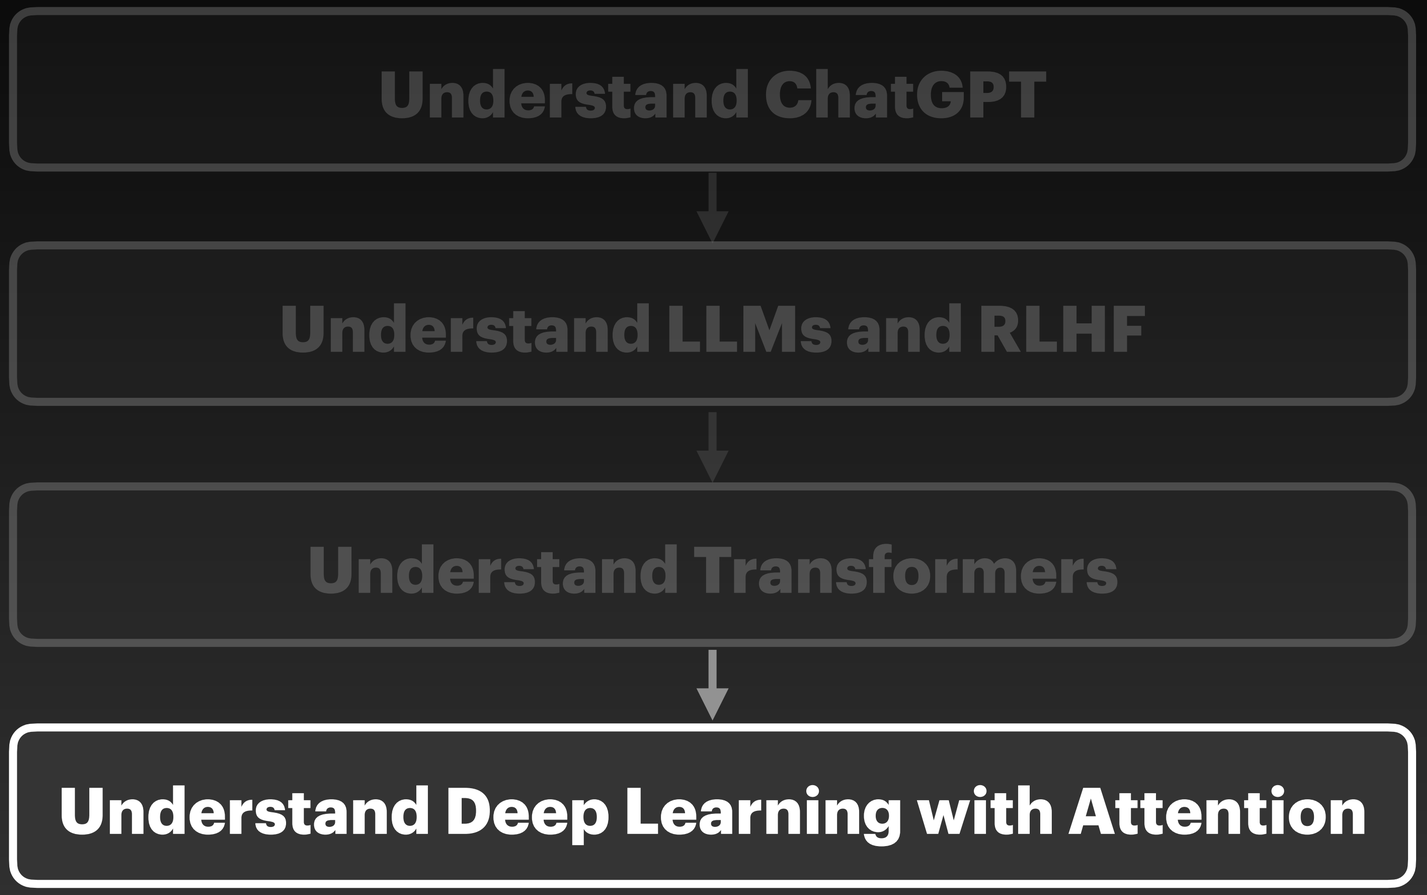

How do you understand ChatGPT?

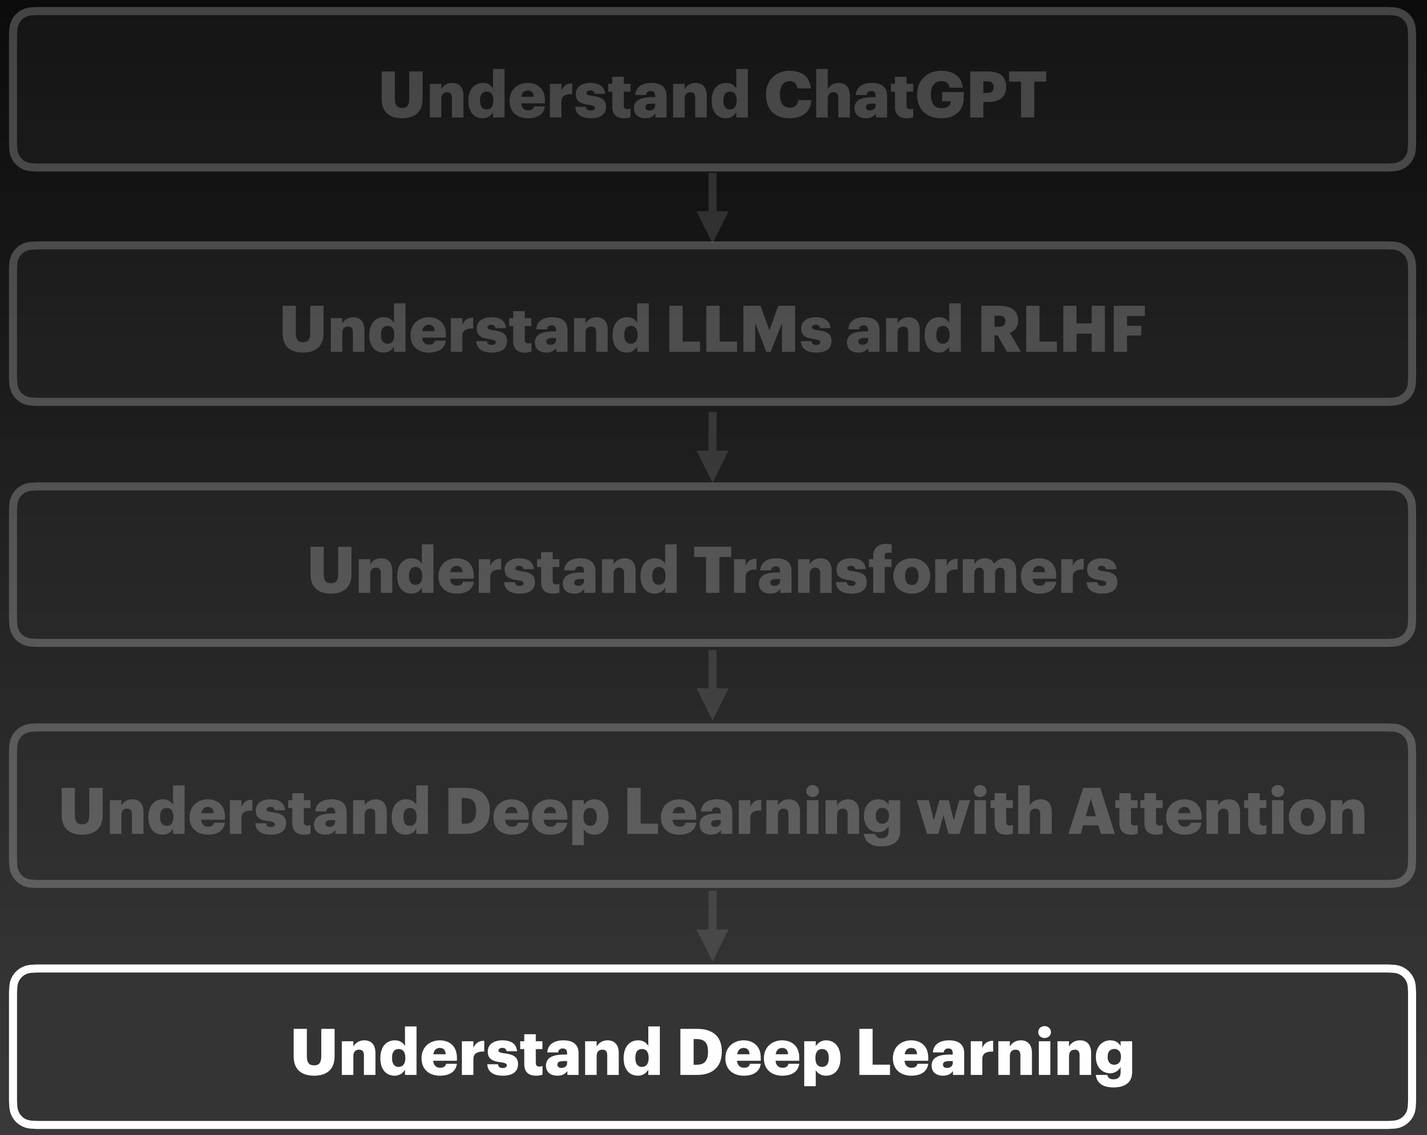

How do you understand ChatGPT?

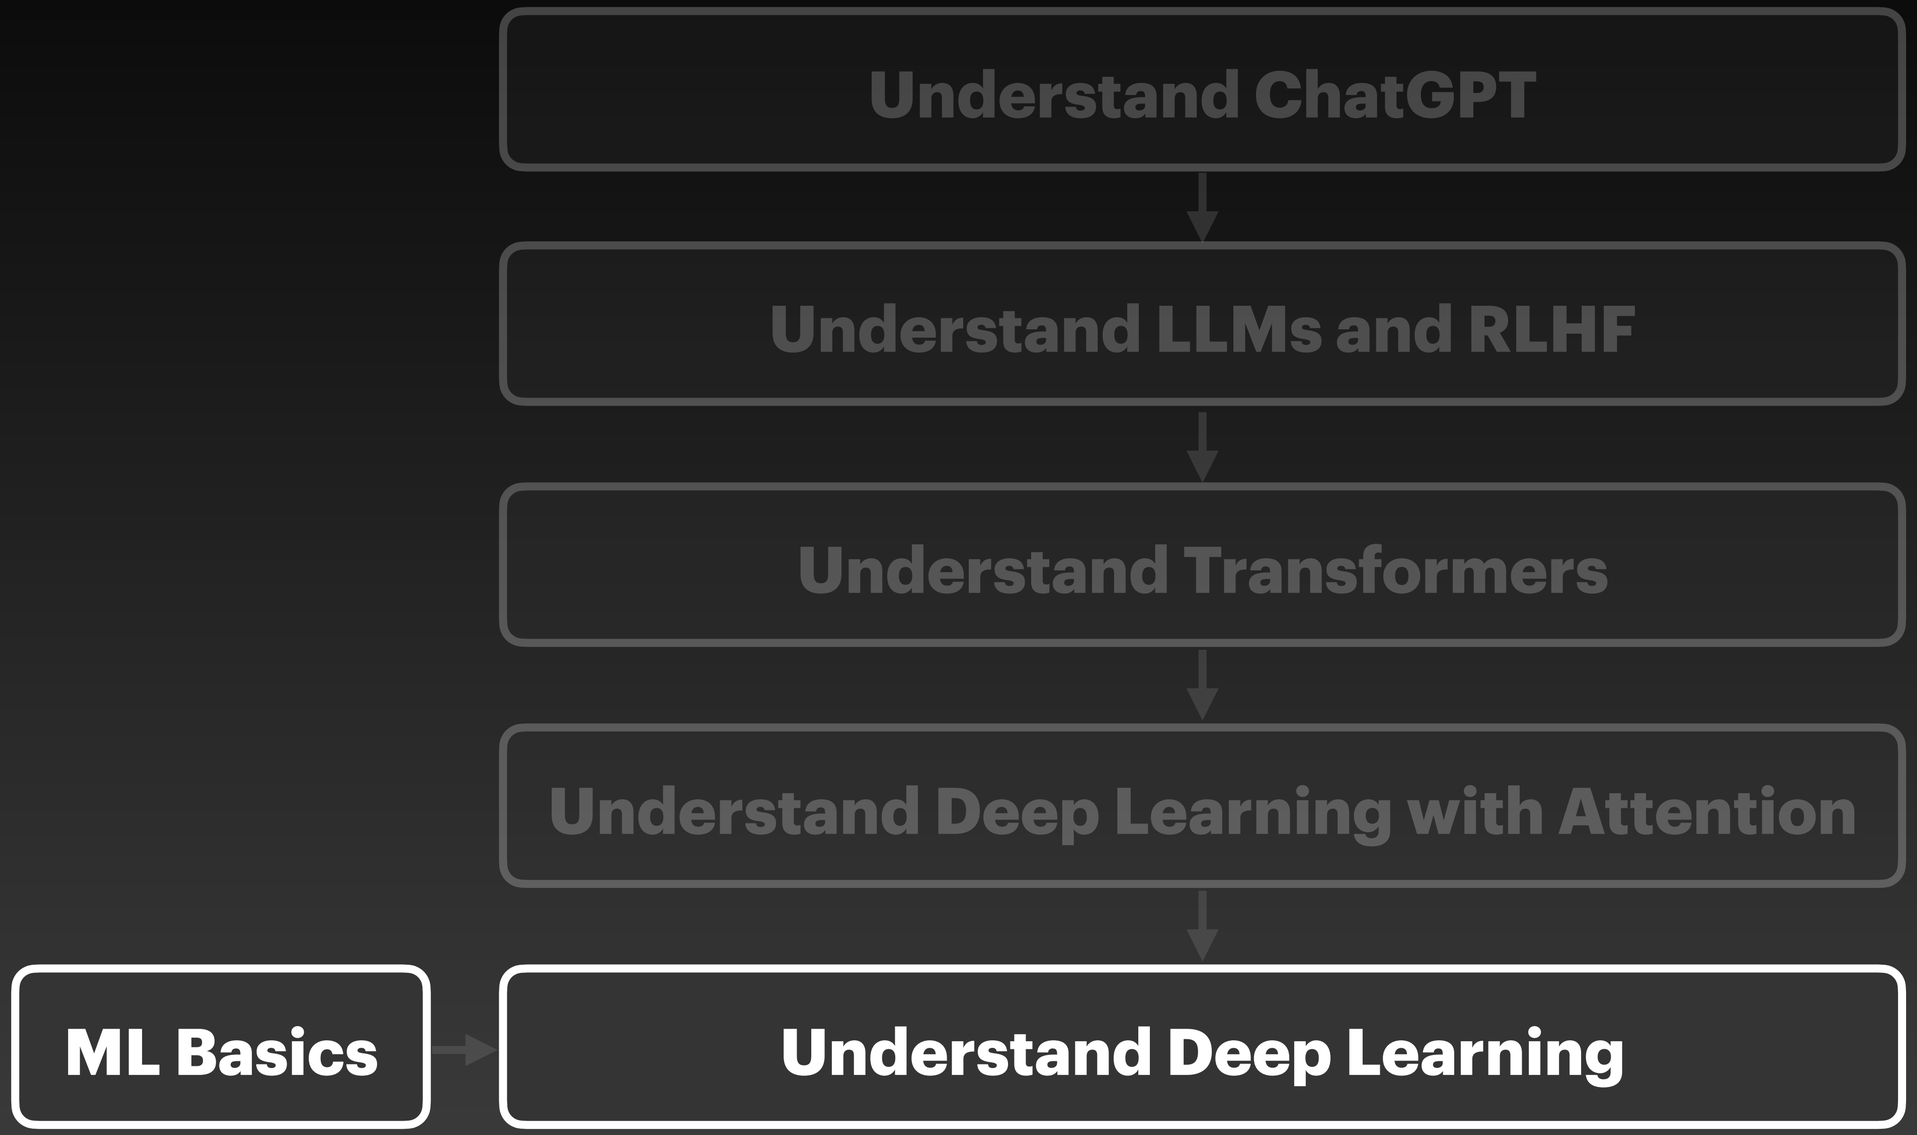

How do you understand ChatGPT?

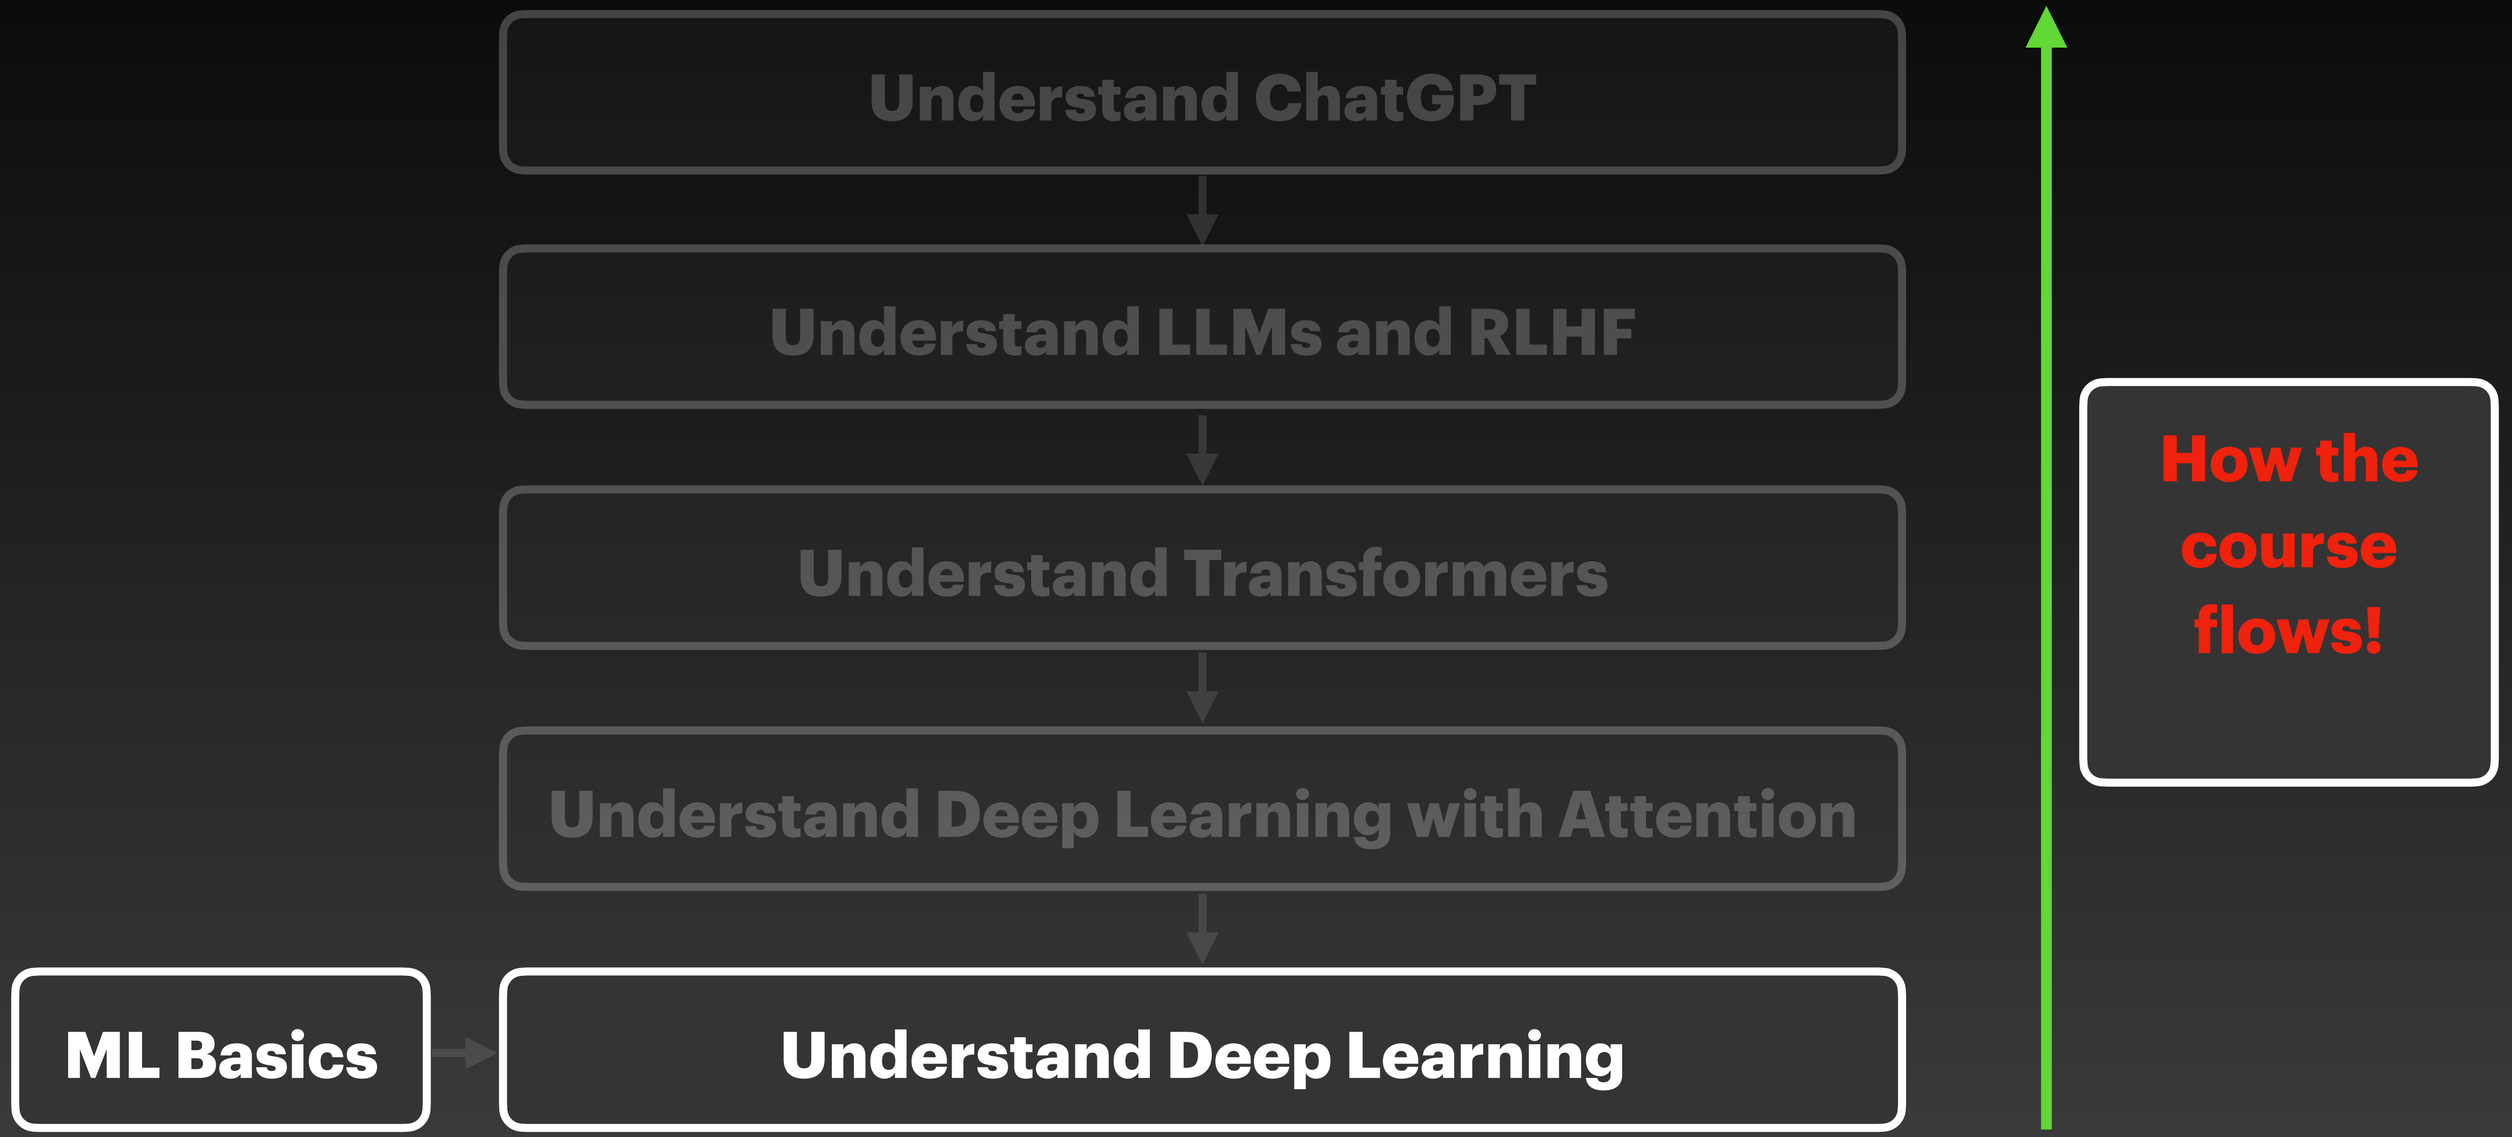

How do you understand ChatGPT?

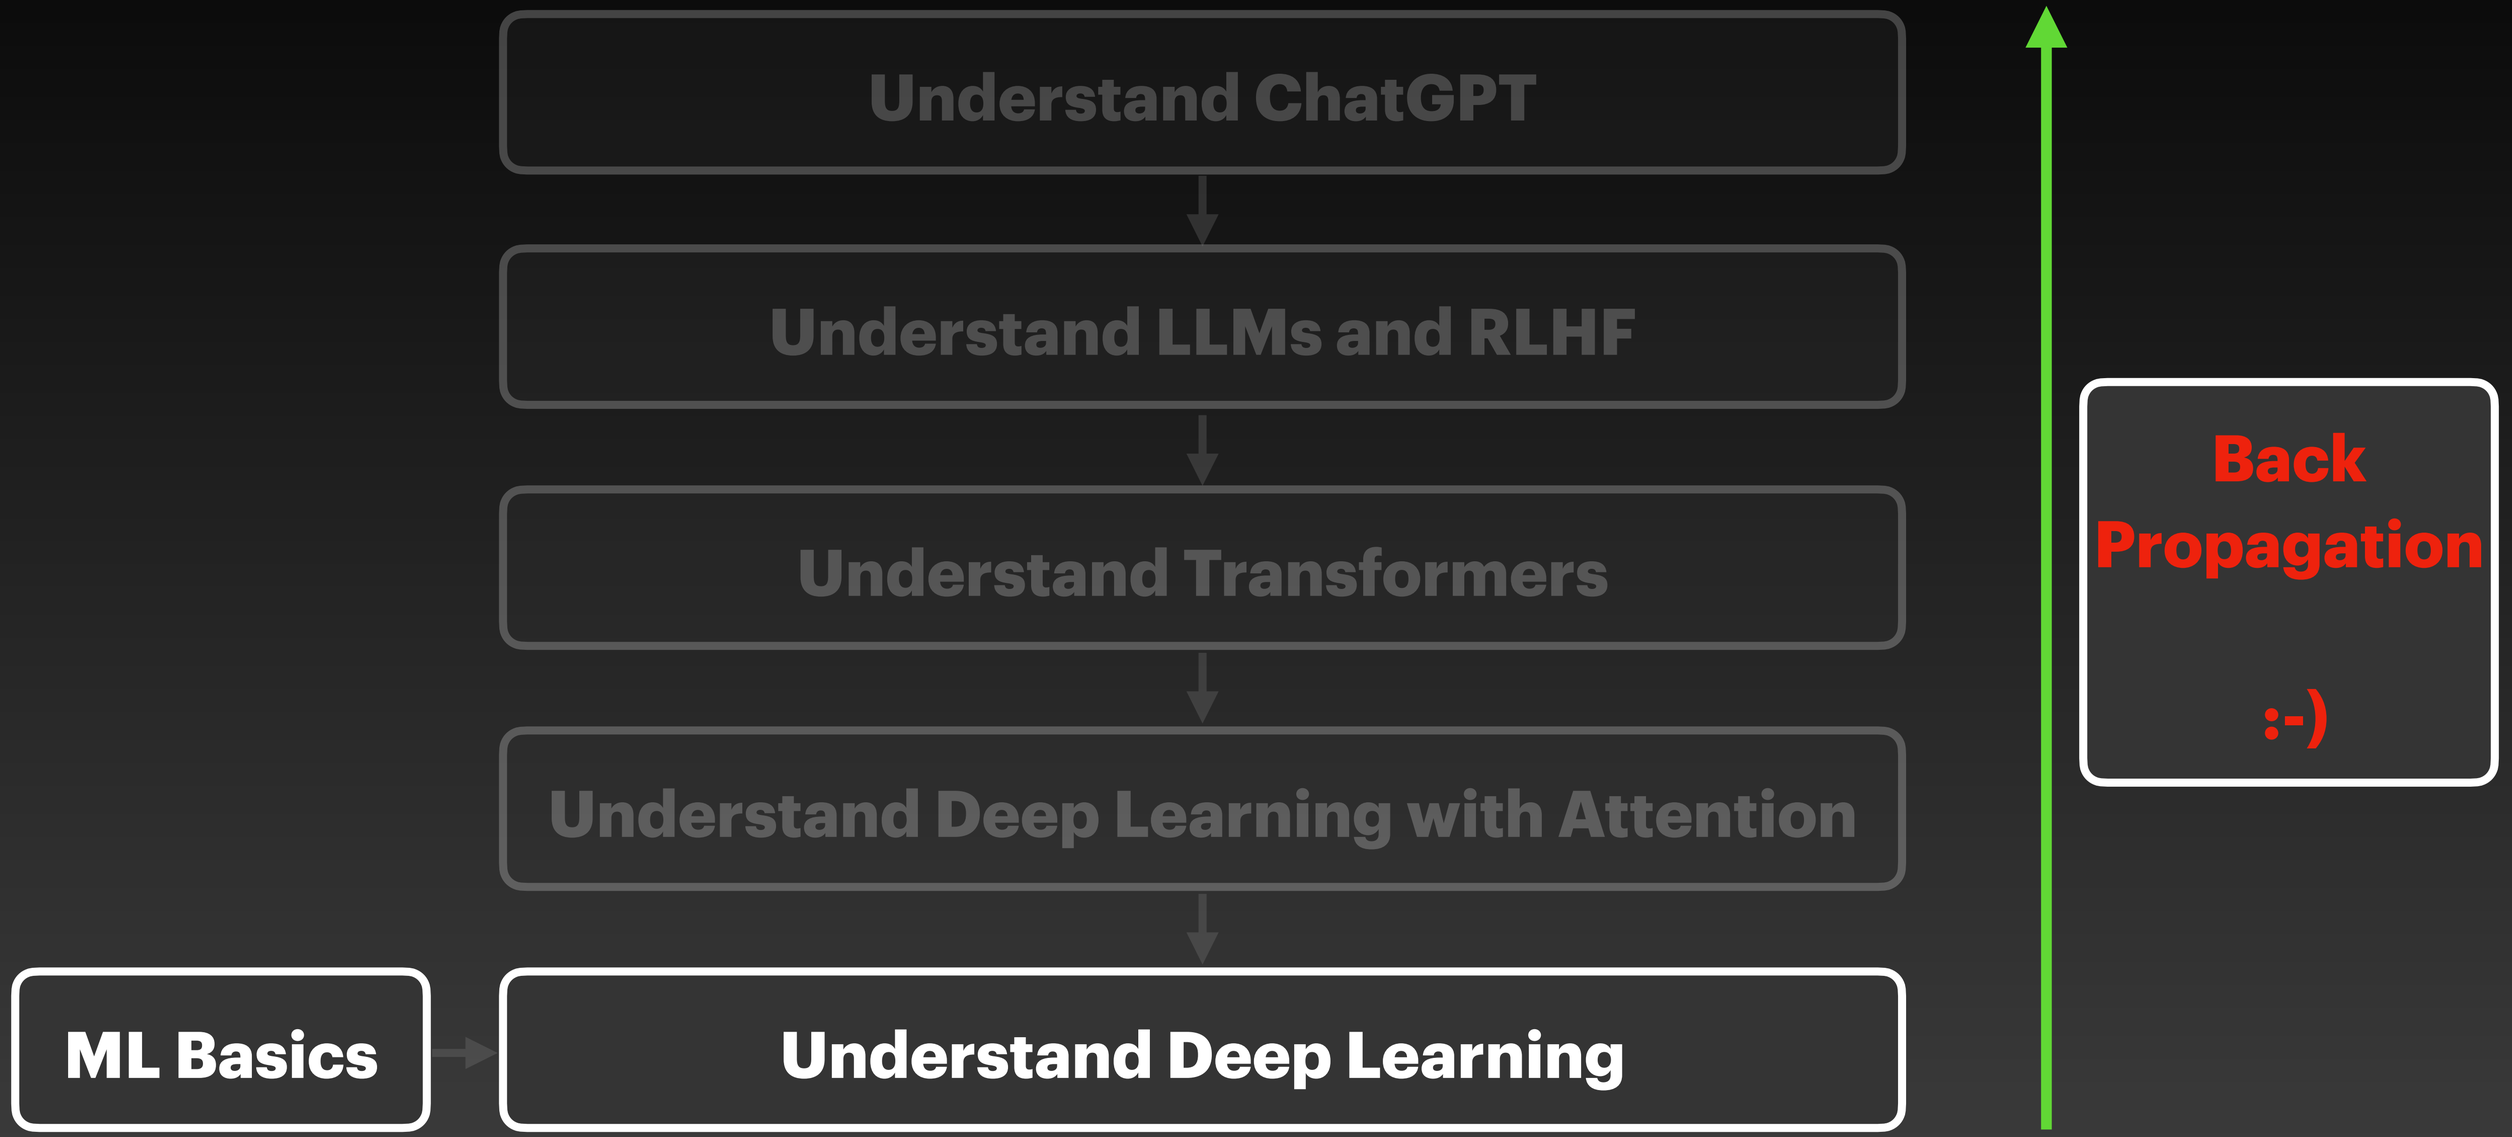

Course Outline

1. Building the foundations

- Logistics and Motivation
- ML fundamentals
- Logistic Regression
- Deep Learning

2. Transformers

- Transformers
- Discriminative and Generative
- Embeddings
- Applications

3. Generative AI

- LLMs
- GPT, GPT-2,GPT-3
- GPT 3.5, GPT 4
- Prompt Engineering
- Fine-tuning and Evaluating LLMs
- Open source vs closed LLMs

4. Miscellaneous Topics

- Auto Encoders
- Stable Diffusion
- Text to Image models
- Applications

Course Outline

1. Building the foundations

- Logistics and Motivation
- ML fundamentals
- Logistic Regression
- Deep Learning

2. Transformers

- Transformers
- Discriminative and Generative
- Embeddings
- Applications

3. Generative AI

- LLMs
- GPT, GPT-2,GPT-3
- GPT 3.5, GPT 4
- Prompt Engineering
- Fine-tuning and Evaluating LLMs
- Open source vs closed LLMs

4. Miscellaneous Topics

- Auto Encoders
- Stable Diffusion
- Text to Image models
- Applications

Course Outline

1. Building the foundations

- Logistics and Motivation
- ML fundamentals
- Logistic Regression
- Deep Learning

2. Transformers

- Transformers
- Discriminative and Generative
- Embeddings
- Applications

3. Generative AI

- LLMs
- GPT, GPT-2,GPT-3
- GPT 3.5, GPT 4
- Agentic Framework
- Prompt Engineering
- Fine-tuning and Evaluating LLMs
- Open source vs closed LLMs

4. Miscellaneous Topics

- Auto Encoders
- Stable Diffusion
- Text to Image models
- Applications

Course Outline

1. Building the foundations

- Logistics and Motivation
- ML fundamentals
- Logistic Regression
- Deep Learning

2. Transformers

- Transformers
- Discriminative and Generative
- Embeddings
- Applications

3. Generative AI

- LLMs
- GPT, GPT-2,GPT-3
- GPT 3.5, GPT 4
- Prompt Engineering
- Fine-tuning and Evaluating LLMs
- Open source vs closed LLMs

4. Miscellaneous Topics

- Auto Encoders
- Stable Diffusion
- Text to Image models
- SLMs
- Responsible AI

Course Webpage and Resources

https://bytesizeml.github.io/llm2026

(Almost) Every Class

First 60 Minutes

- Theory
- Code samples

Next 15 minutes

- In-Class Exercise

Next 35 minutes

- Theory
- In-class Coding Exercise

(Almost) Every Class

First 60 Minutes

- Theory
- Code samples

Next 15 minutes

- In-Class Exercise

Next 35 minutes

- Theory
- In-class Coding Exercise

(Almost) Every Class

First 60 Minutes

- Theory
- Code samples

Next 15 minutes

- In-Class Exercise

Next 35 minutes

- Theory
- Business Applications
- In-class Coding Exercise

What I would like you to take away!

Conceptually

- Better understanding of LLMs
- Of LLM application areas
- Of APIs
- Intuition behind LLM models
- Theory behind LLMs

Implementation

- Coding up baselines in Colab
- Comfort with APIs
- Use of Hugging Face models
- Showcasing your work on webpage
- Fine-tuning LLM models

Ideas

- Where can you apply LLMs next?
- How can you chain LLMs
- to solve a problem?

What I would like you to take away!

Conceptually

- Better understanding of LLMs
- Of LLM application areas
- Of APIs
- Intuition behind LLM models

Implementation

- Coding up baselines in Colab
- Comfort with APIs
- Use of Hugging Face models
- Showcasing your work on web-apps
- Fine-tuning LLM models

Ideas

- Where can you apply LLMs next?
- How can you chain LLMs
- to solve a problem?

What I would like you to take away!

Conceptually

- Better understanding of LLMs
- Of LLM application areas
- Of APIs
- Intuition behind LLM models

Implementation

- Coding up baselines in Colab
- Comfort with APIs
- Use of Hugging Face models
- Showcasing your work on webpage
- Fine-tuning LLM models

Ideas

- Where can you apply LLMs next?
- Grasping cutting-edge technical papers
- Being able to do basic LLM designs
- for pet projects

Survey Results

Find a AI buddy

Find a buddy and catch up on your AI goals for this class and also for your career

Discord Setup and Message

Set up your access to the class discord channel
&
send a link to an interesting article on LLMs on the channel

Practical applications of LLMs you want to talk about

Discuss in groups of 3 or 4

- What business applications or pet projects are you looking at?
- How can use LLMs for them?
- Would LLMs be good to have or actually needed for the application?

Assignments

1. In-Class Exercises (15%)

- 3-4 questions per lecture
- Respond on a poll or a form
- Will be graded for credit

2. Coding Assignments (15%)

- Two Coding Assignments
- First two weeks of class

3. Mini-projects (45%)

- 3 for this class
- Get 2 weeks to work on each
- More involved than a coding assignment
- Could include a Kaggle Contest
- Could include a web demo

4. Project Presentation (15%)

- Present on one of the mini-projects
- Presentation on M or W of finals week
- 7 minutes per team + 3 minute questions
- Methodology + working demo and results

Assignments

1. In-Class Exercises (15%)

- 3-4 questions per lecture
- Respond on a poll or a form
- Will be graded for credit

2. Coding Assignments (15%)

- Two Coding Assignments
- First two weeks of class

3. Mini-projects (45%)

- 3 for this class
- Get 2 weeks to work on each
- More involved than a coding assignment
- Could include a Kaggle Contest
- Could include a web demo

4. Project Presentation (15%)

- Present on one of the mini-projects
- Presentation on M or W of finals week
- 7 minutes per team + 3 minute questions
- Methodology + working demo and results

Assignments

1. In-Class Exercises (15%)

- 3-4 questions per lecture
- Respond on a poll or a form
- Will be graded for credit

2. Coding Assignments (15%)

- Two Coding Assignments
- First two weeks of class

3. Mini-projects (45%)

- 3 for this class
- Get 2 weeks to work on each
- More involved than a coding assignment
- Could include a Kaggle Contest
- Could include a web demo

4. Project Presentation (15%)

- Present on one of the mini-projects
- Presentation on M or W of finals week
- 7 minutes per team + 3 minute questions
- Methodology + working demo and results

Assignments

1. In-Class Exercises (15%)

- 3-4 questions per lecture
- Respond on a poll or a form
- Will be graded for credit

2. Coding Assignments (15%)

- Two Coding Assignments
- First two weeks of class

3. Mini-projects (45%)

- 3 for this class
- Get 2 weeks to work on each
- More involved than a coding assignment
- Could include a Kaggle Contest
- Could include a web demo

4. Project Presentation (15%)

- Present on one of the mini-projects
- Presentation on M or W of finals week
- 7 minutes per team + 3 minute questions
- Methodology + working demo and results

Assignments

5. Paper Presentation (10%)

- Present on a technical paper
- Present in teams of 2
- Will be graded on presentation quality
- And ability to respond to questions

Assignments Summary

Summary

- In-Class Exercises - 15%
- Coding Assignments - 15%
- Mini-Projects - 45%
- Mini-Project Presentation - 15%
- Paper Presentation - 10%

Assignment 1

- To be assigned tomorrow and due next Monday
- Focuses on set up for colab, apis + llm discovery and insights

Assignments Flow

- Assignment 1 and 2 due next and week after next
- Mini-Projects 1,2,3 follow
- Paper presentation slots will be available and presentation will happen last 15 mins of class

Course Flow

Engine vs API

Engines are different from APIs and we shouldn’t confuse the two.

Engine vs API

Engines are different from APIs and we shouldn’t confuse the two.

BERT and Llama are Engines/Foundation Models whereas ChatGPT 4o is an API

Engine vs API

Foundation Models
(Pre-Trained Models)

- BERT (Encoder only)
- GPT (Decoder only)
- Claude
- Stable Diffusion (Vision)

Chat APIs

Engine vs API

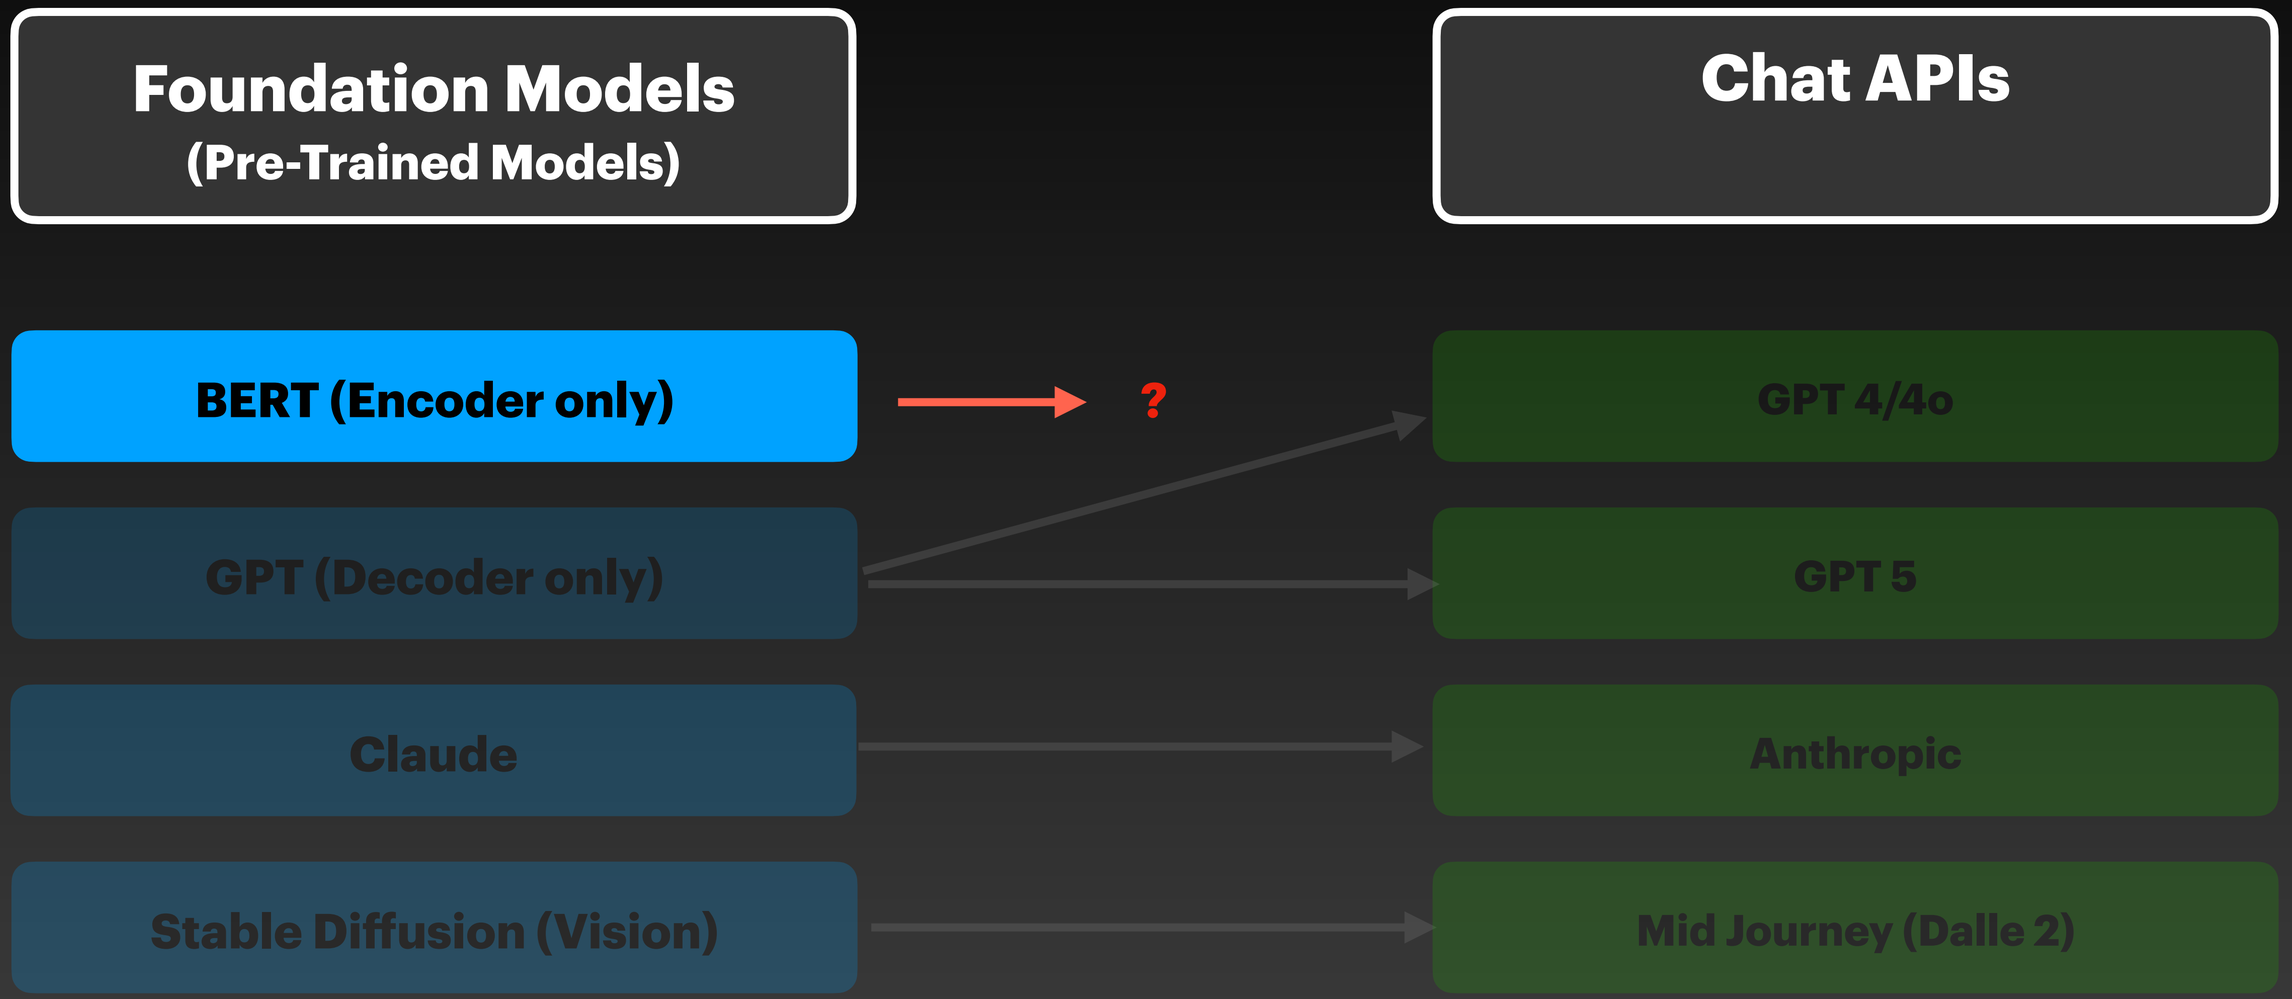

Engine vs API

Foundation Models
(Pre-Trained Models)

Chat APIs

BERT (Encoder only)

GPT (Decoder only)

Claude

Stable Diffusion (Vision)

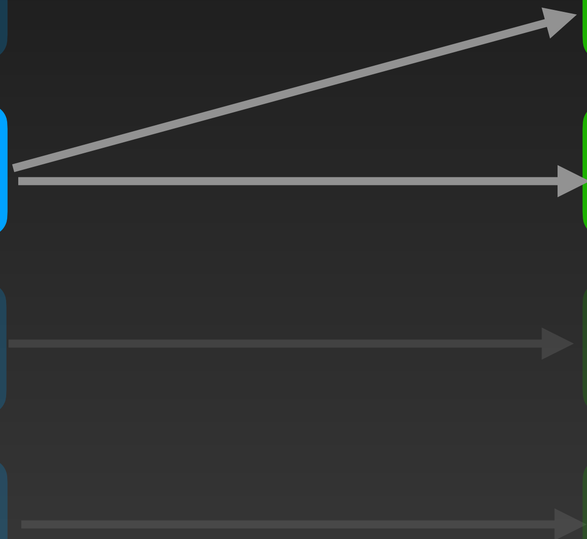

GPT 4/GPT 4o

GPT 5.2 (Thinking)

Anthropic

Mid Journey (Dalle 2)

Engine vs API

Foundation Models
(Pre-Trained Models)

Chat APIs

BERT (Encoder only)

GPT (Decoder only)

Claude

Stable Diffusion (Vision)

GPT 4/4o

GPT 5

Opus 4.5, Sonnet 4.5

Mid Journey (Dalle 2)

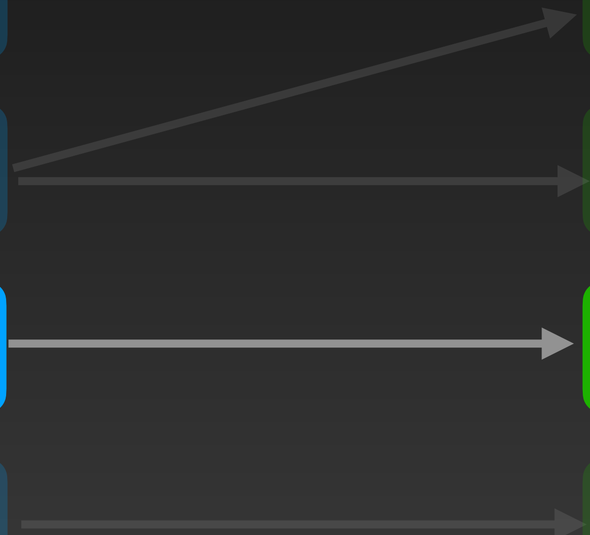

Engine vs API

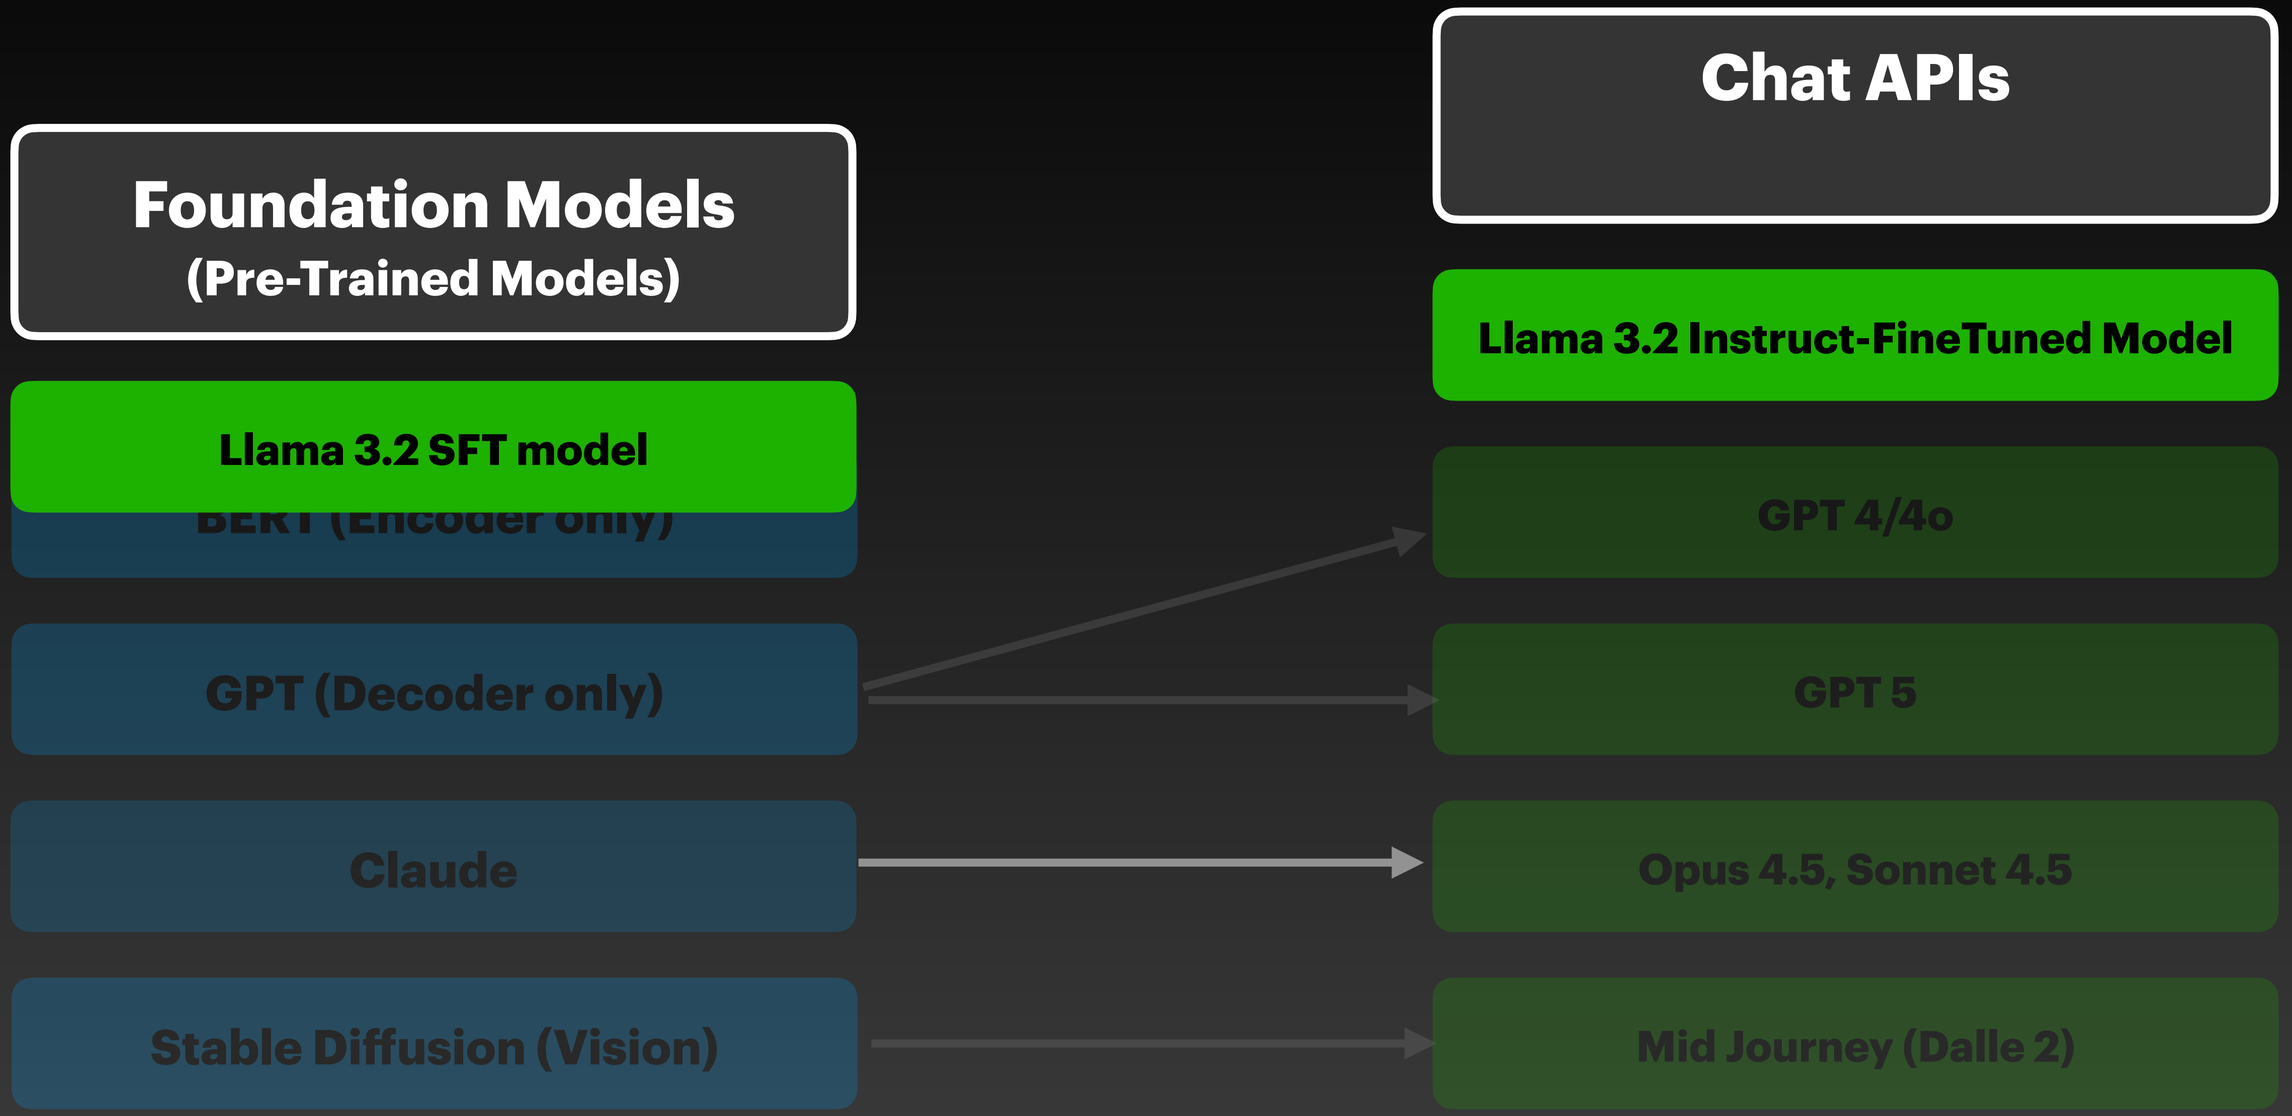

Engine vs API

Foundation Models
(Pre-Trained Models)

Llama 3.2 SFT model

BERT (Encoder only)

GPT (Decoder only

Claude

Stable Diffusion (Visi

Chat APIs

Llama 3.2 Instruct-FineTuned Model

GPT 4/4o

GPT 5

pus 4.5, Sonnet 4.5

lid Journey (Dalle 2)

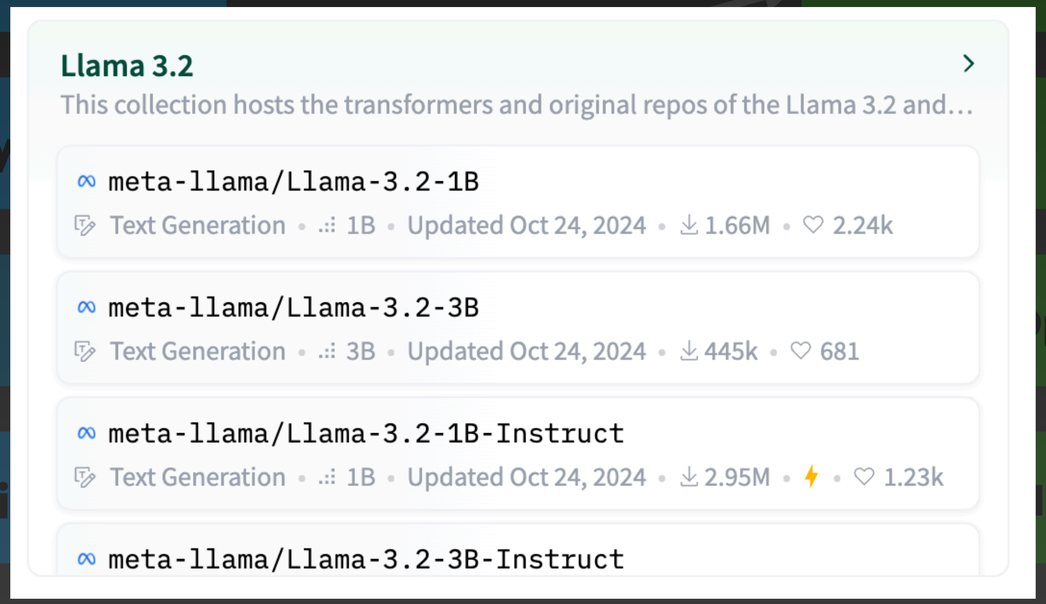

Engine vs API

Foundation Models
(Pre-Trained Models)

Chat APIs

BERT (Encoder only)

GPT (Decoder only)

Claude

Stable Diffusion (Vision)

GPT 4/GPT 4o

GPT 5

Anthropic

Mid Journey (Dalle 2)

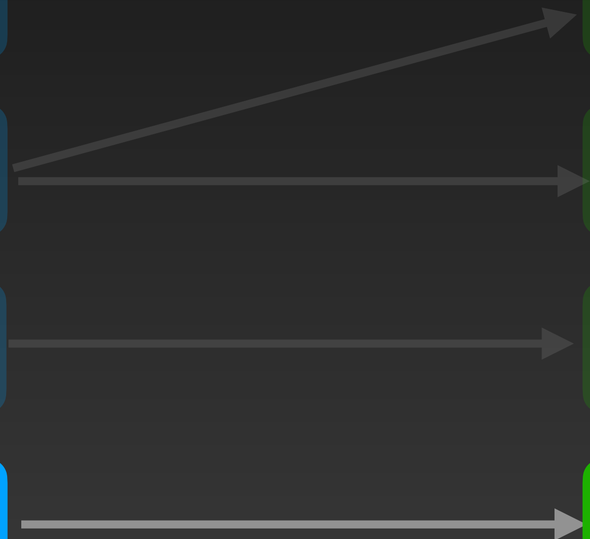

Engine vs API

Foundation Models
(Pre-Trained Models)

Chat APIs

BERT (Encoder only)

GPT (Decoder only)

Large Vision Model (LVM)
Pre-Trained Model
Foundation Model

GPT 4/ GPT 4o

GPT 5

Claude

Opus, Sonnet

Stable Diffusion (Vision)

Mid Journey (Dalle 2)

Engine vs API

Foundation Models
(Pre-Trained Models)

t APIs

LLM - Large Language Model
Pre-Trained Model
Foundation Model

BERT (Encoder only)

GPT (Decoder only)

Claude

Stable Diffusion (Vision)

GPT 3.5

GPT 4

Opus, Sonnet

Mid Journey (Dalle 2)

Engine vs API

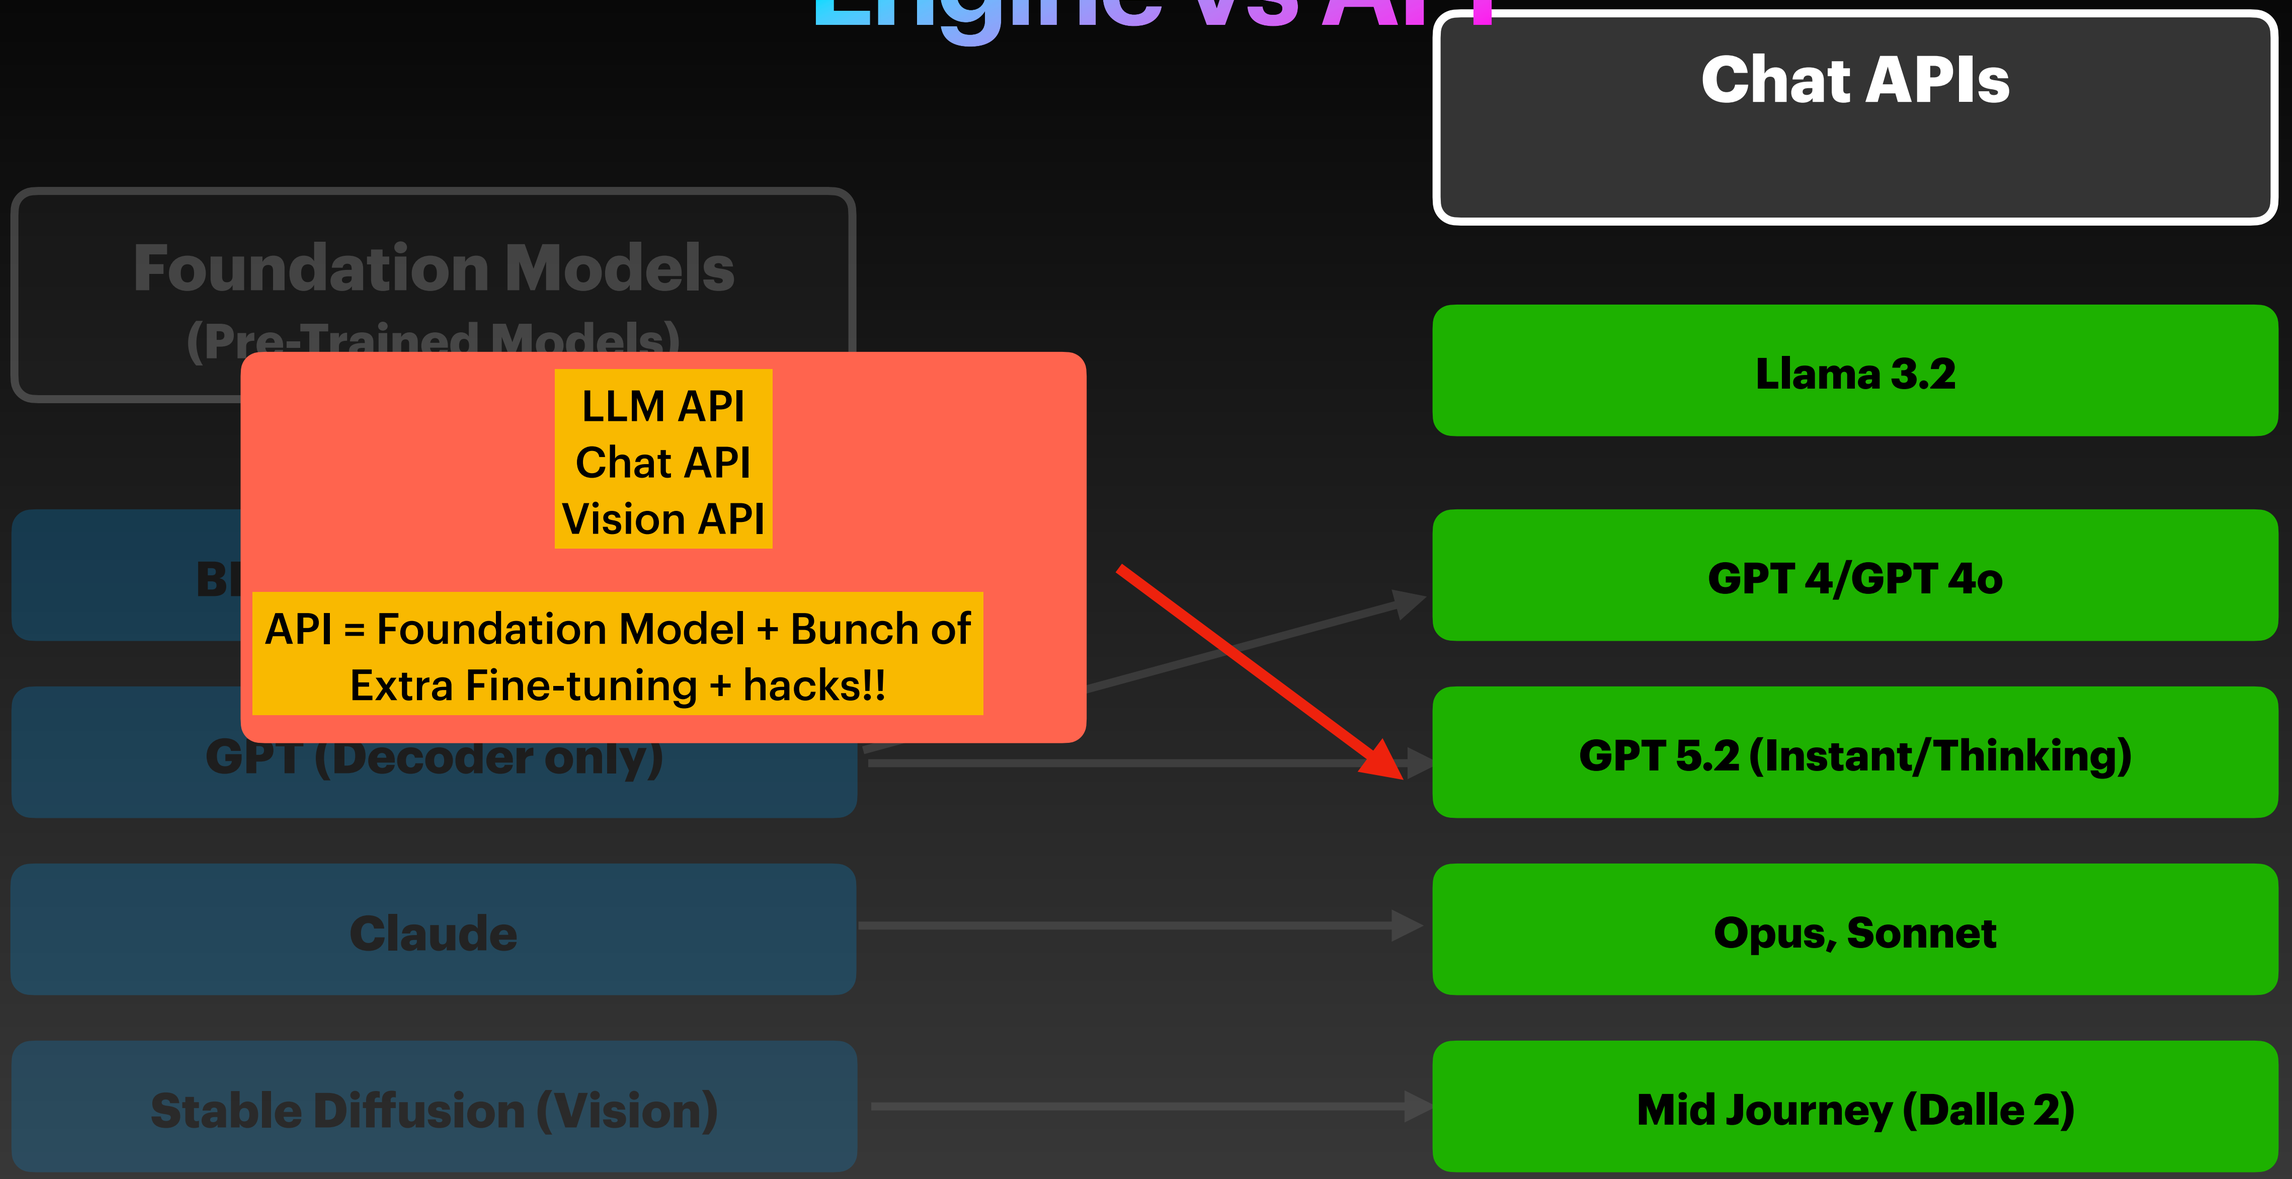

What is a Language Model?

Scientific Data-driven Model that helps machines understand language and patterns in sentence construction

What is a Language Model?

Example: I just got promoted. I am feeling so ———

What is a Language Model?

Example: I just got promoted. I am feeling so happy

What is a Language Model?

Example: I just checked my application status and it got ———. It’s frustrating!

What is a Language Model?

Example: I just checked my application status and it got rejected. It's frustrating!

What is a Large Language Model (LLM)?

LLMs are language models that are learned from massive corpuses of text, that are mined from the web. They are known to be sophisticated in understanding language and can be **generative** in nature.

History of (Large) Language Models

How did machines work with language before and how we do it now?

History of (Large) Language Models

Perceptron

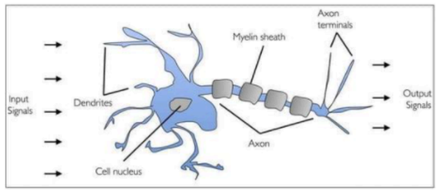

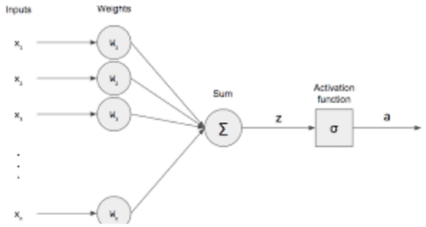

History of (Large) Language Models

Perceptron

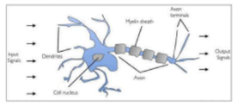

Feed Forward NN

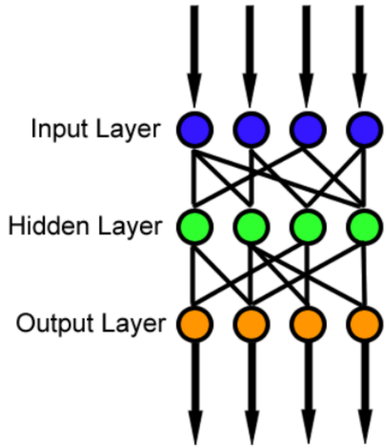

History of (Large) Language Models

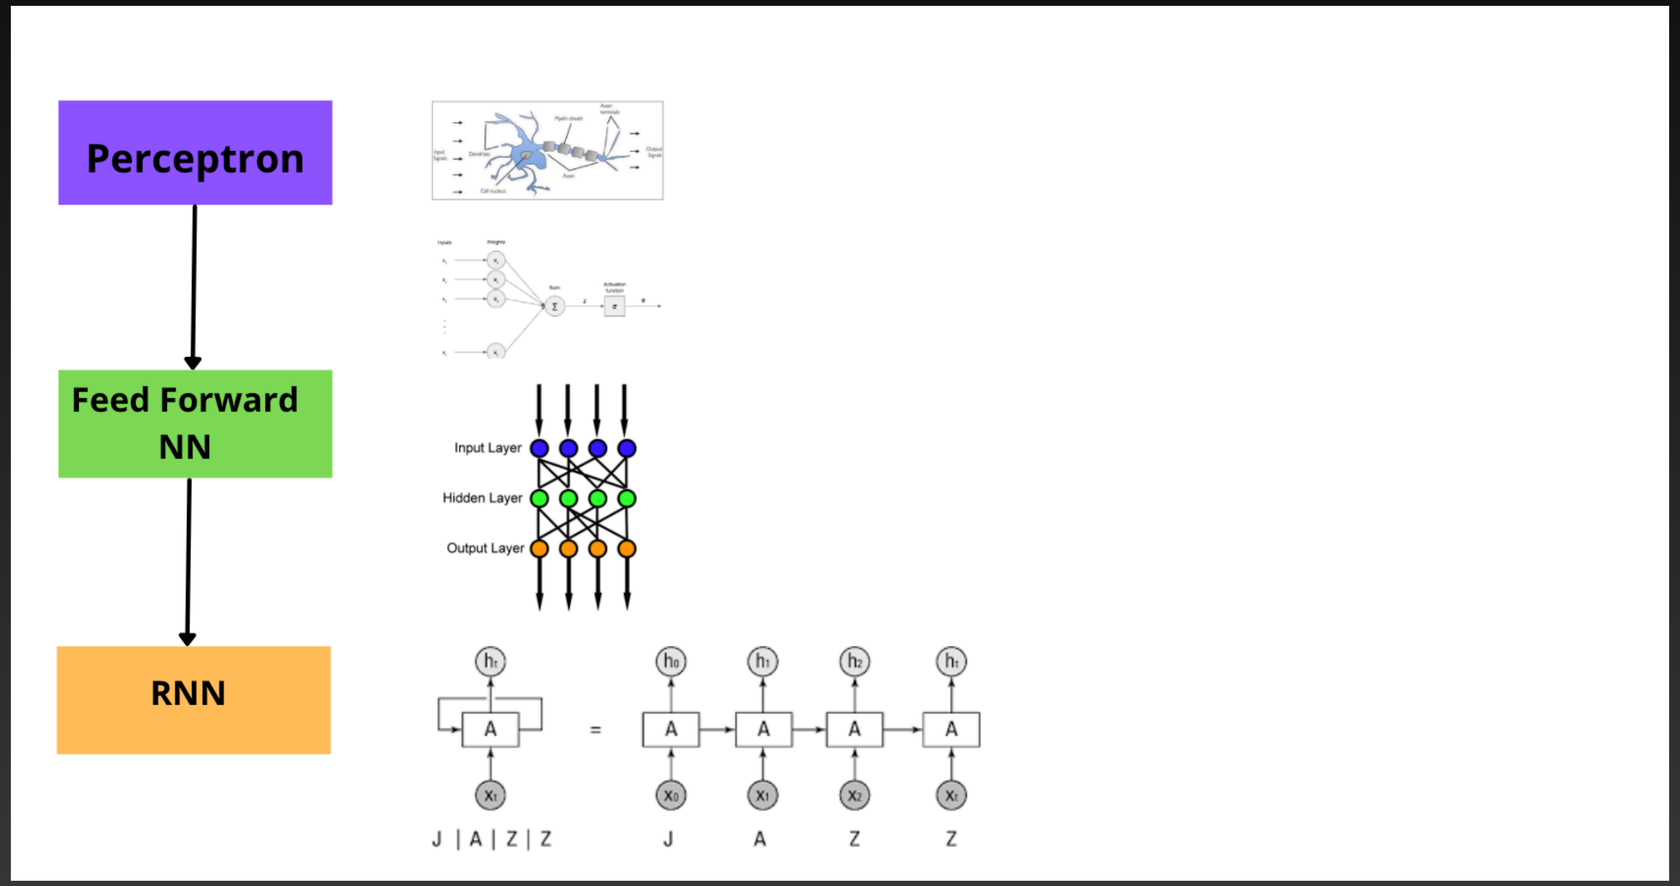

History of (Large) Language Models

RNN Issue:

I just arrived in NY. In a few days, I would like to visit the city, ———

History of (Large) Language Models

RNN Issue:

I just arrived in NY. In a few days, I would like to visit the city, NY

History of (Large) Language Models

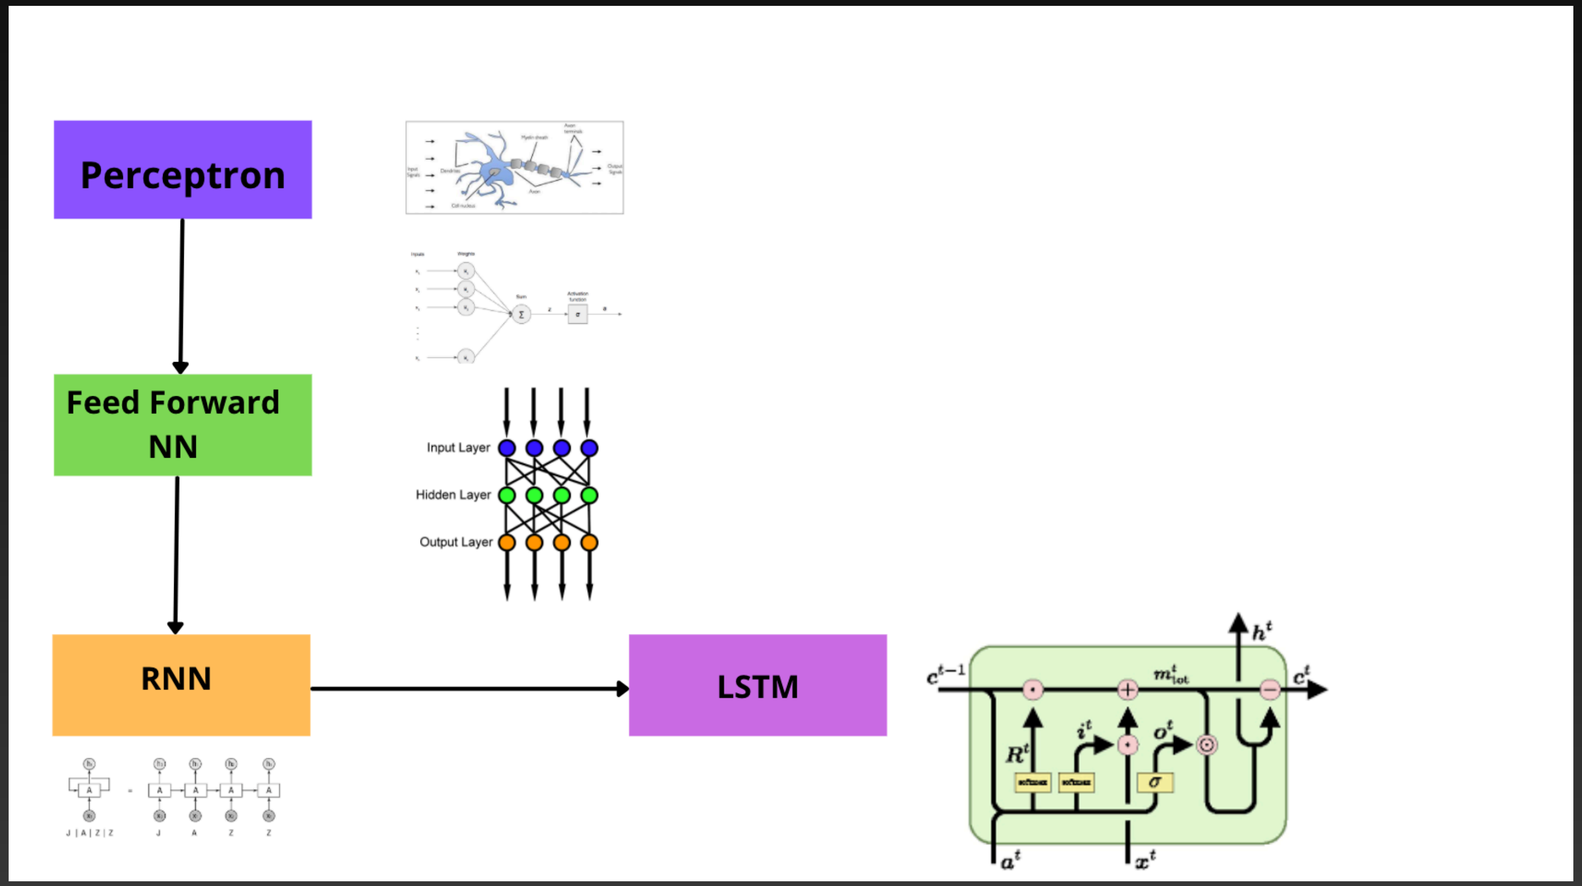

History of (Large) Language Models

LSTM

I just arrived in NY. In a few days, I would like to visit the city, ———

History of (Large) Language Models

LSTM

I just arrived in NY. In a few days, I would like to visit the city, Seattle

History of (Large) Language Models

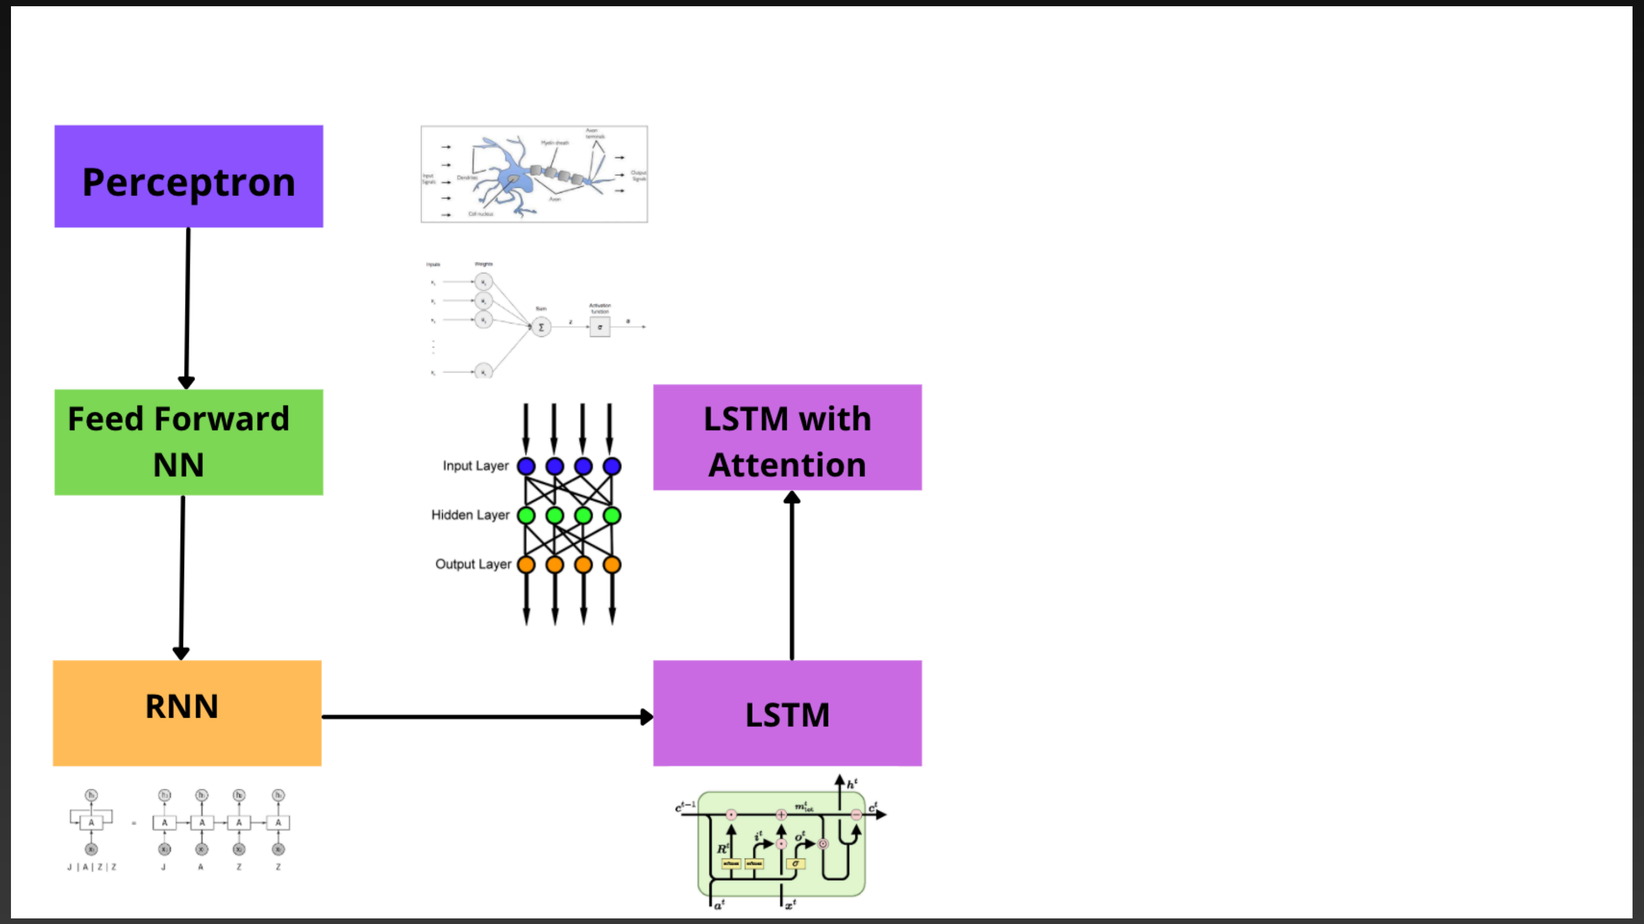

History of (Large) Language Models

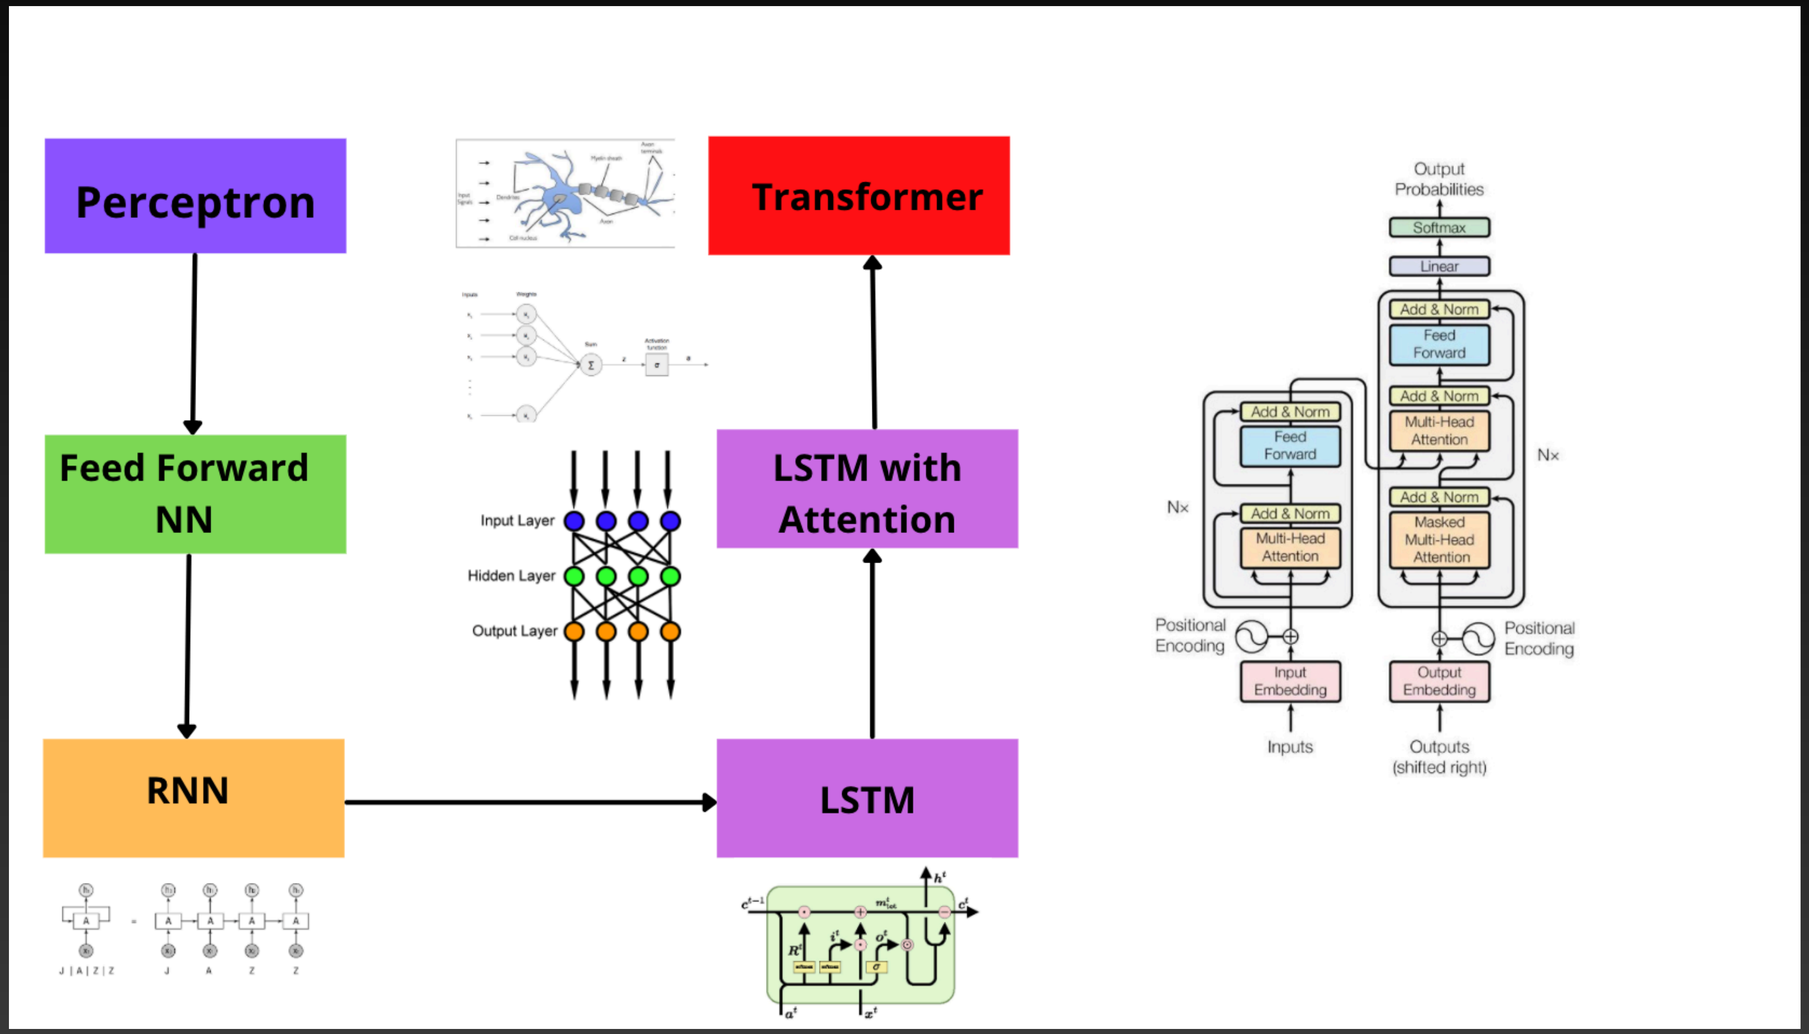

History of (Large) Language Models

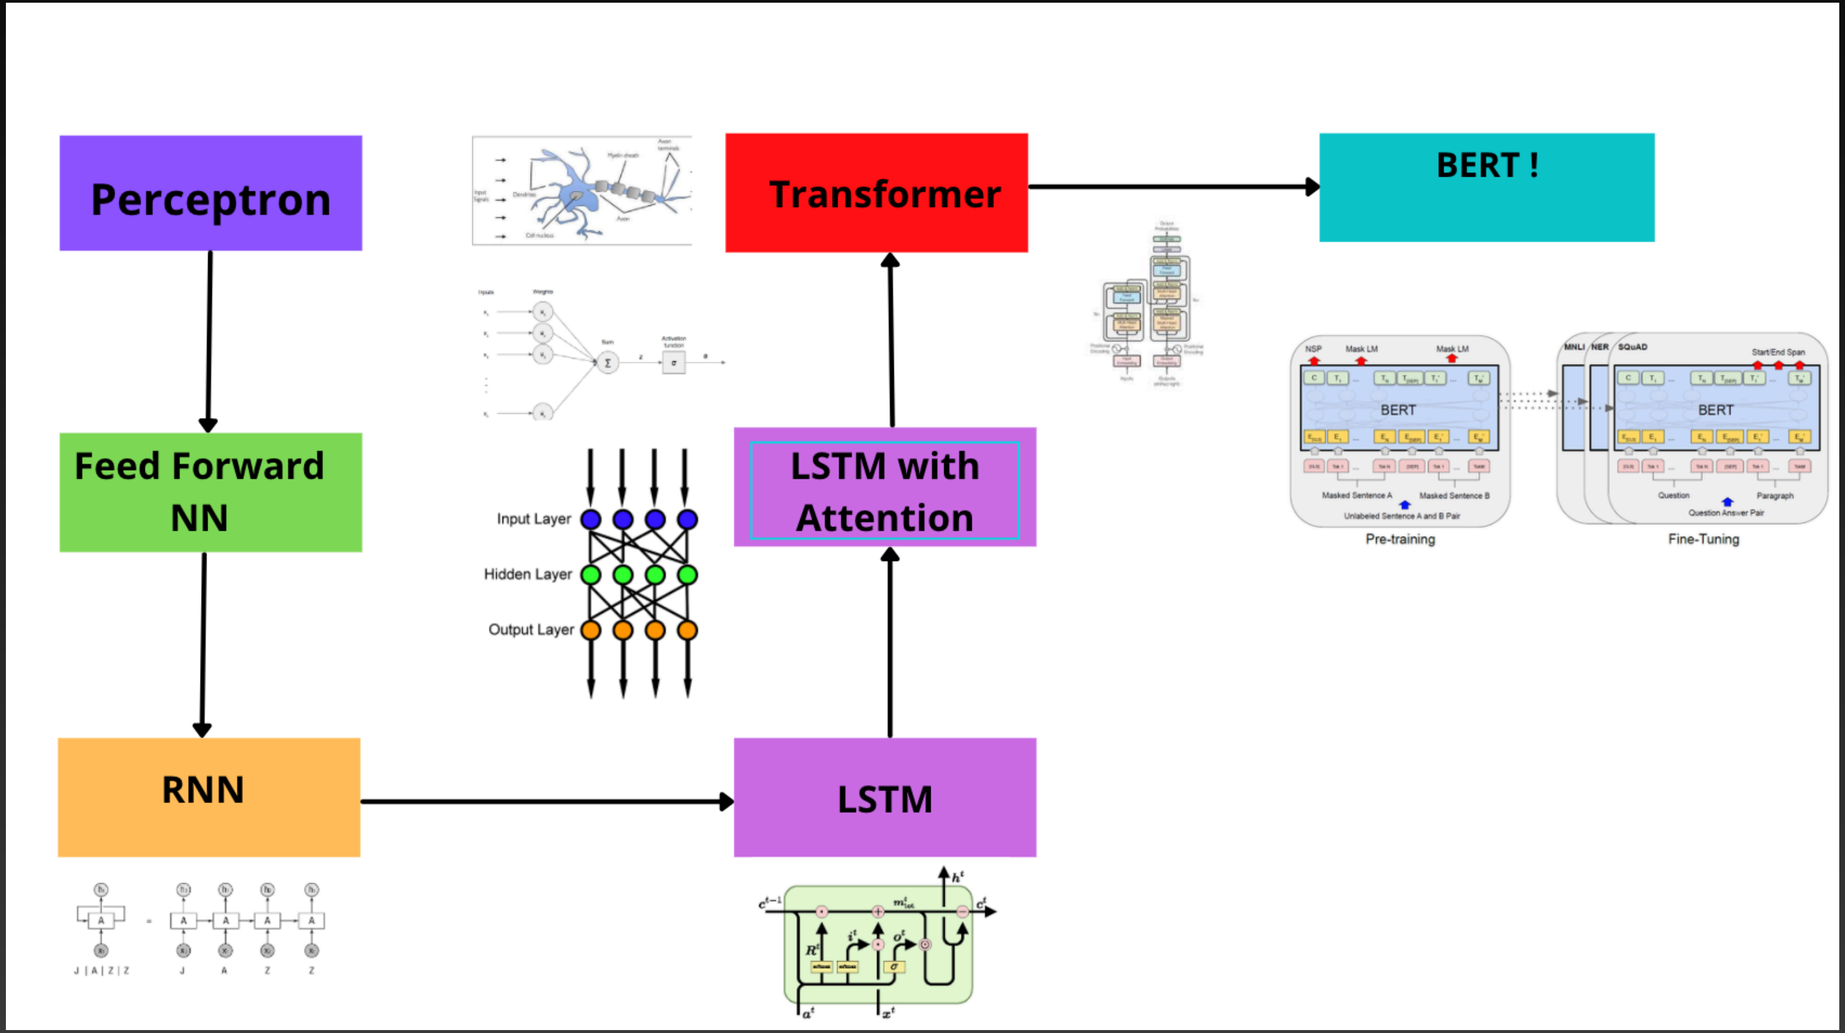

History of (Large) Language Models

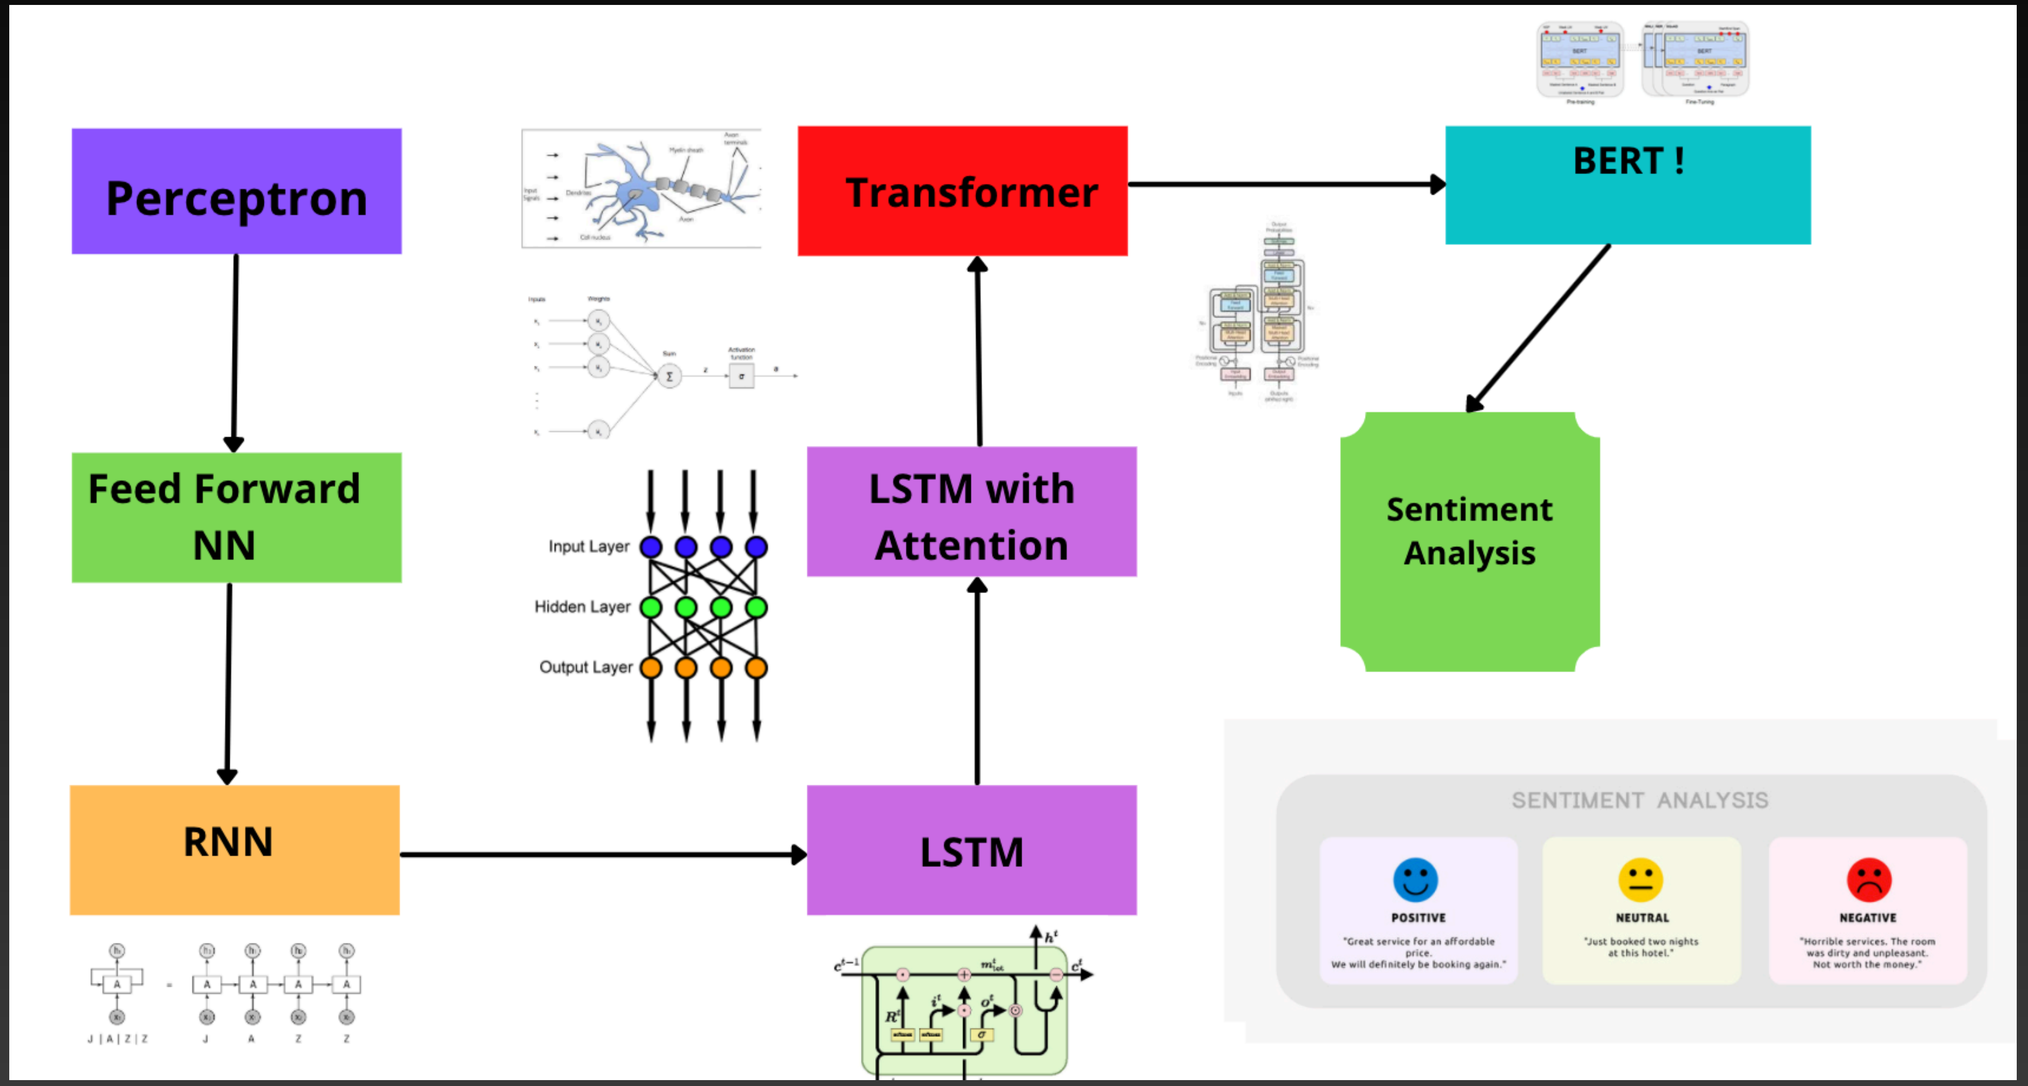

History of (Large) Language Models

Perceptron

Feed Forward NN

RNN

LSTM

LSTM with Attention

Transformer

BERT !

Extractive Summarization

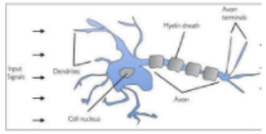

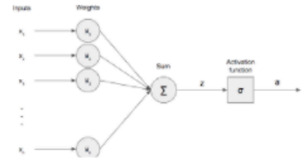

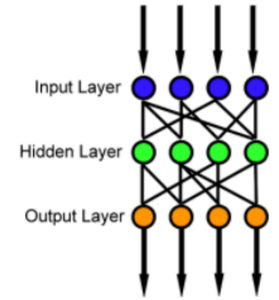

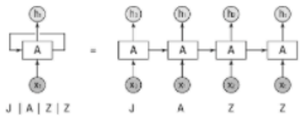

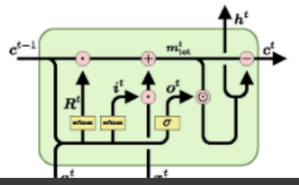

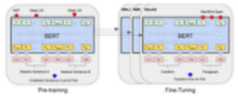

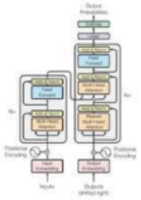

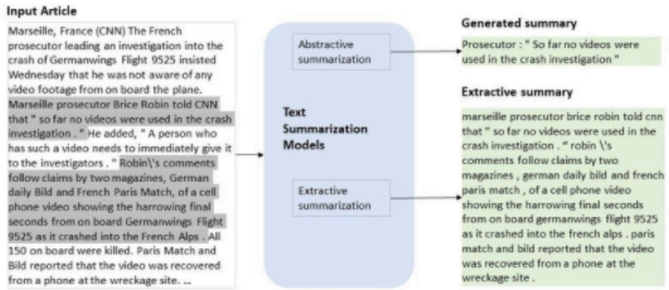

History of (Large) Language Models

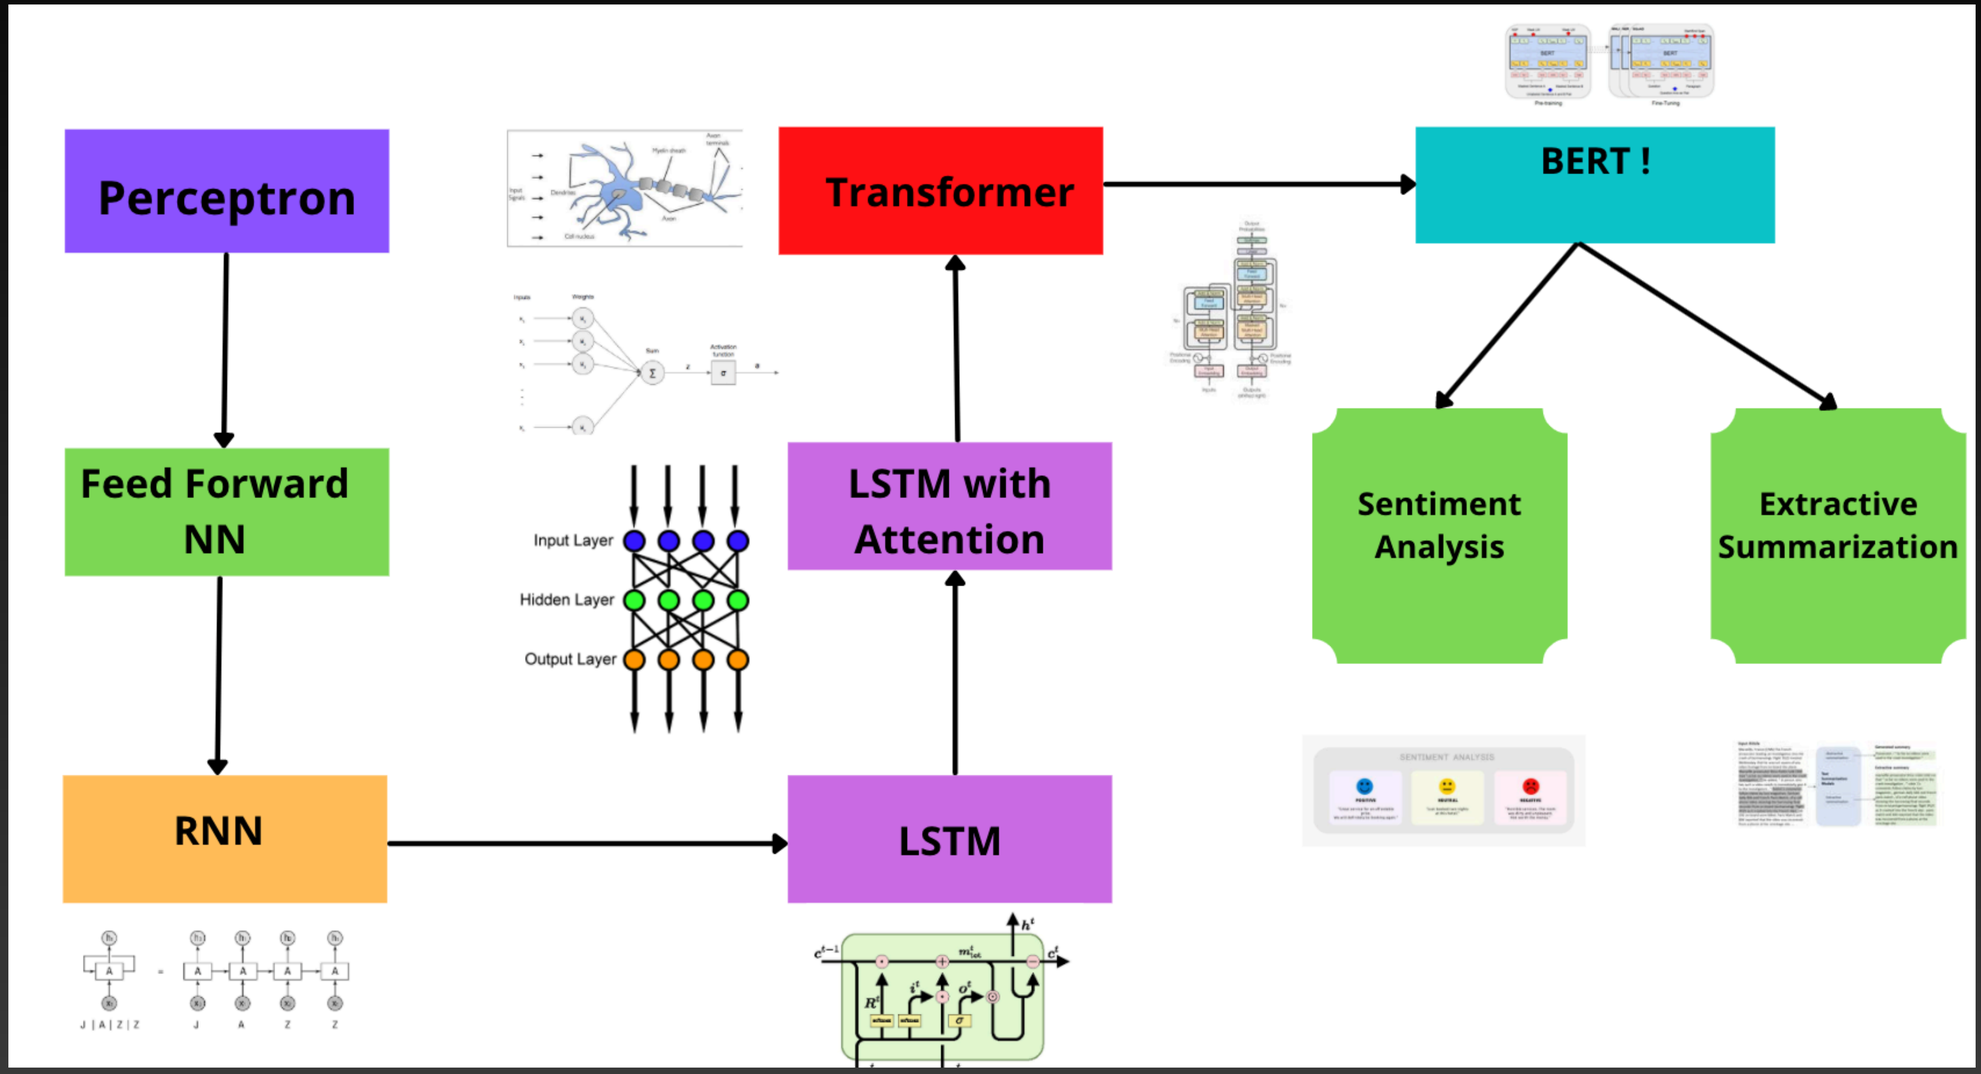

History of (Large) Language Models

GPT vs BERT

While BERT is purely about encoding and is called an encoding Transformer. GPT is purely a decoder and is called a decoding transformer.

History of (Large) Language Models

GPT-x

GPT-x (GPT, GPT-2, GPT-2.5, etc) are decoding transformers that are trained to predict the next token given the past and do a very good job at it! That's how they can generate entire paragraphs that look logical, grammatical and structured.

1 Trillion Tokens!

	RedPajama	LLaMA*
CommonCrawl	878 billion	852 billion
C4	175 billion	190 billion
Github	59 billion	100 billion
Books	26 billion	25 billion
ArXiv	28 billion	33 billion
Wikipedia	24 billion	25 billion
StackExchange	20 billion	27 billion
Total	1.2 trillion	1.25 trillion

1 Trillion Tokens requires how many books?

	RedPajama	LLaMA*
CommonCrawl	878 billion	852 billion
C4	175 billion	190 billion
Github	59 billion	100 billion
Books	26 billion	25 billion
ArXiv	28 billion	33 billion
Wikipedia	24 billion	25 billion
StackExchange	20 billion	27 billion
Total	1.2 trillion	1.25 trillion

1 Trillion Tokens requires how many books?

	RedPajama	LLaMA*
CommonCrawl	878 billion	852 billion
C4	175 billion	190 billion
Github	59 billion	100 billion
Books	26 billion	25 billion
ArXiv	28 billion	33 billion
Wikipedia	24 billion	25 billion
StackExchange	20 billion	27 billion
Total	1.2 trillion	1.25 trillion

1 Book ~ 50k Tokens

1 Trillion Tokens requires how many books?

	RedPajama	LLaMA*
CommonCrawl	878 billion	852 billion
C4	175 billion	190 billion
Github	59 billion	100 billion
Books	26 billion	25 billion
ArXiv	28 billion	33 billion
Wikipedia	24 billion	25 billion
StackExchange	20 billion	27 billion
Total	1.2 trillion	1.25 trillion

1 Book ~ 50k Tokens

15 Million Books ~ 1 Trillion Tokens

ChatGPT use cases for NLP

ChatGPT use cases for NLP

Table 1: Distribution of use case categories from our API prompt dataset.

Use-case	(%)
Generation	45.6%
Open QA	12.4%
Brainstorming	11.2%
Chat	8.4%
Rewrite	6.6%
Summarization	4.2%
Classification	3.5%
Other	3.5%
Closed QA	2.6%
Extract	1.9%

Table 2: Illustrative prompts from our API prompt dataset. These are fictional examples inspired by real usage—see more examples in Appendix A.2.1.

Use-case	Prompt
Brainstorming	List five ideas for how to regain enthusiasm for my career
Generation	Write a short story where a bear goes to the beach, makes friends with a seal, and then returns home.
Rewrite	This is the summary of a Broadway play: """ {summary} """ This is the outline of the commercial for that play: """

The distribution of prompts used to finetune InstructGPT

Dialing it back a bit...

Deep Learning Foundations