

A practical introduction to Generative AI



Speaker:

Giovanni Rosa, PhD

AI Engineer

@Technology Reply

 giovannirosa.com

 gio.rosa@reply.it



ITOUG TRAINING DAY @ GALDUS

MILANO 23 OTTOBRE 2024

UNLOCK THE FUTURE WITH ORACLE: MASTERING ANALYTICS,
APEX, DATABASE AND GEN AI



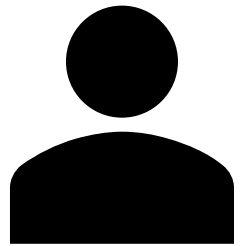




Prompt:

“Generate an image logo for the Italian Oracle User Group 2024 event hosted by Galdus formazione in Milan city”





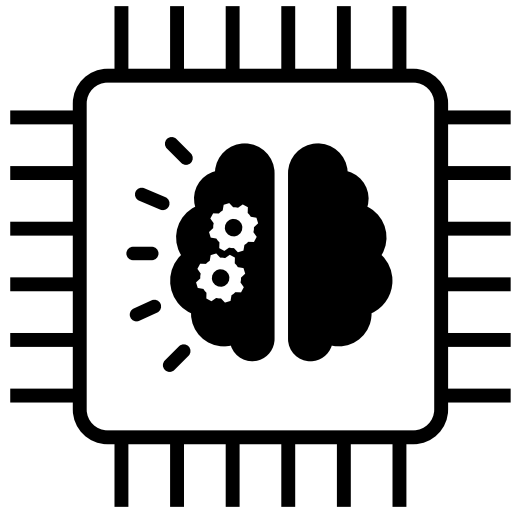
User

call



1

“Generate an image logo for the Italian Oracle User Group 2024 event hosted by Galdus formazione in Milan city”



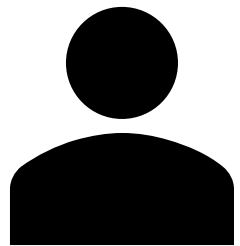
Generative AI
model

create



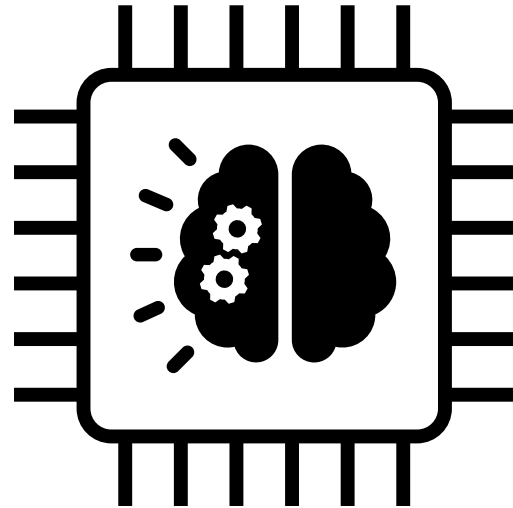
2





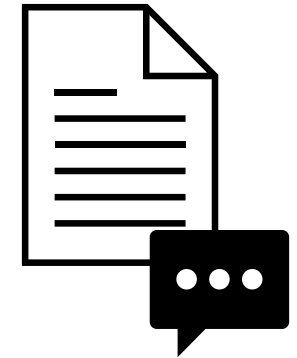
User

instruction

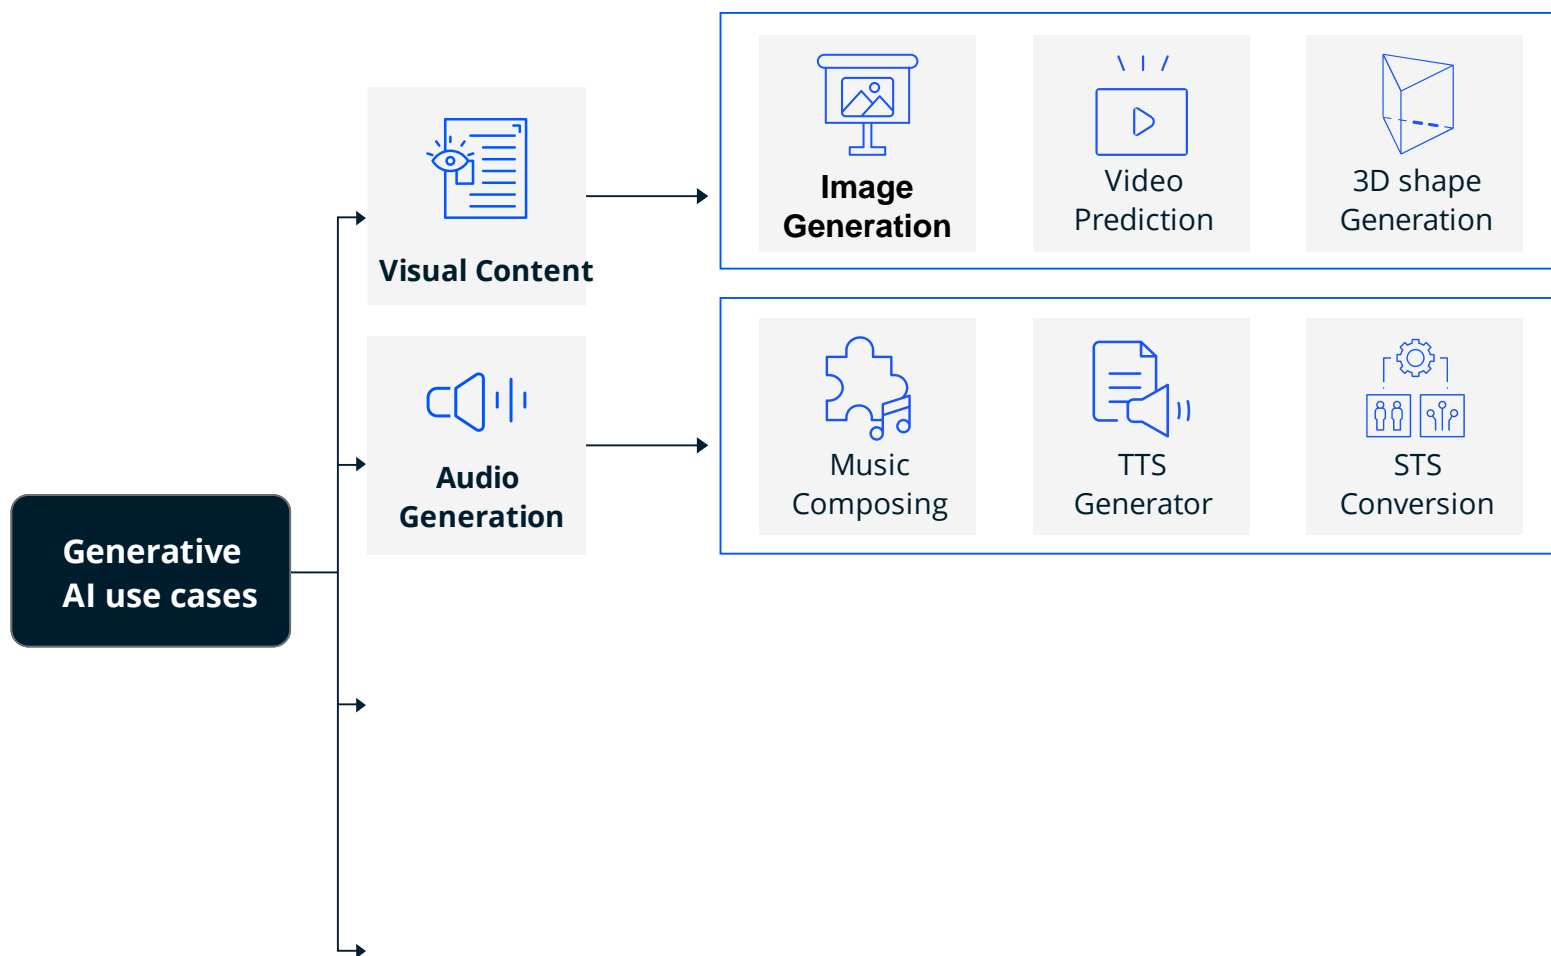


**Generative AI
model**

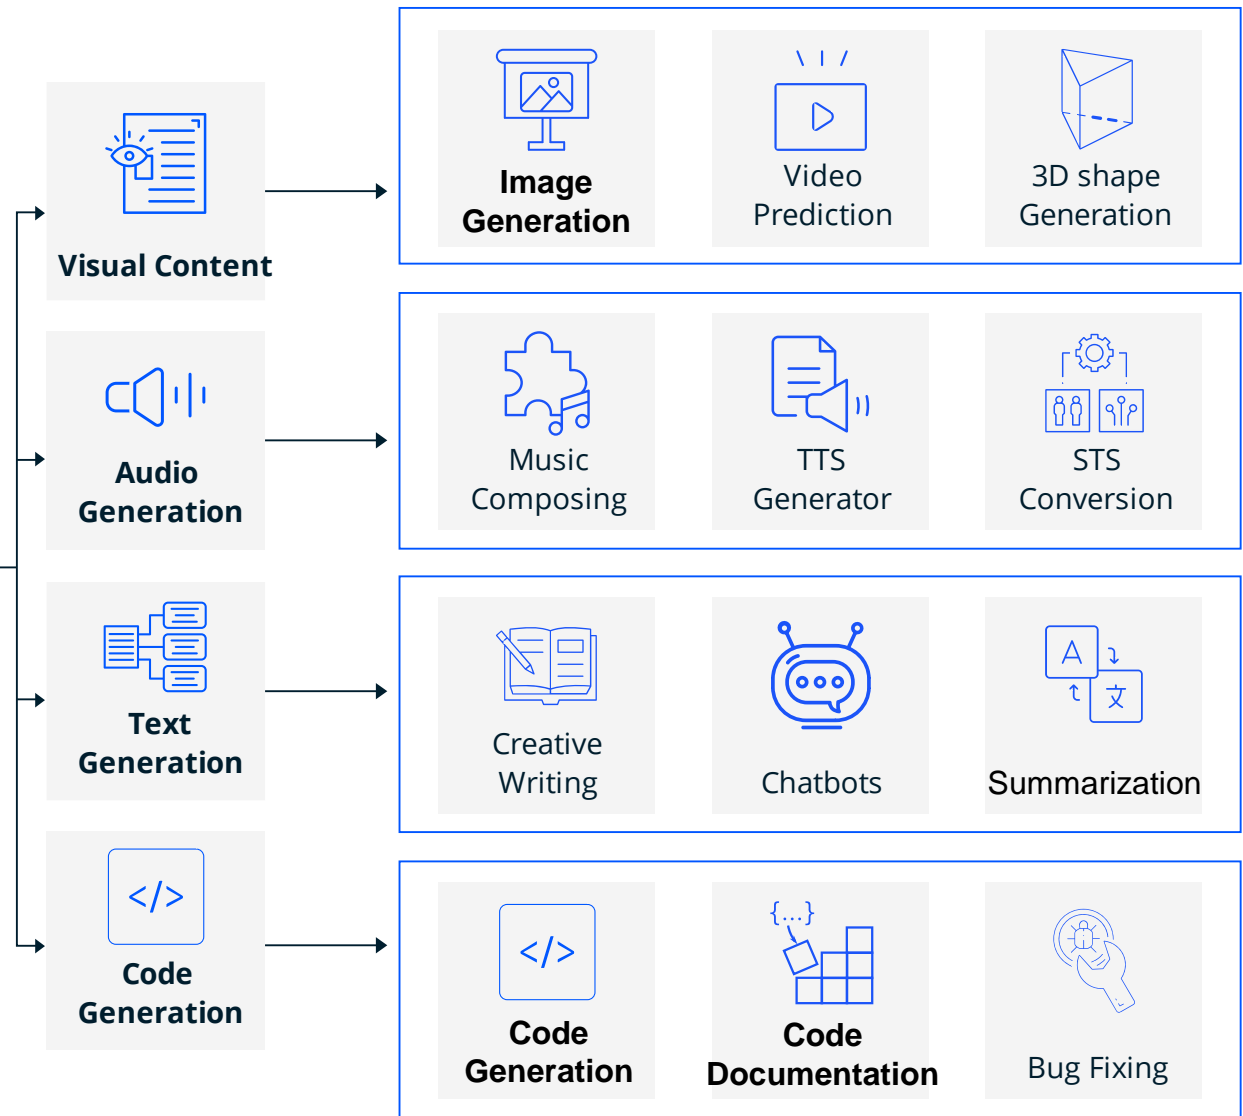
create



Response



Generative AI use cases



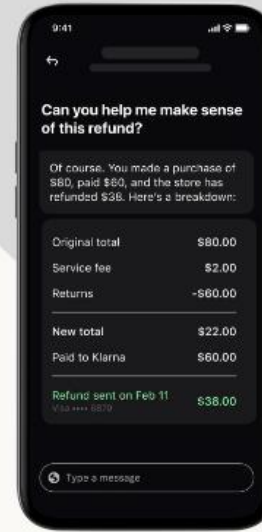
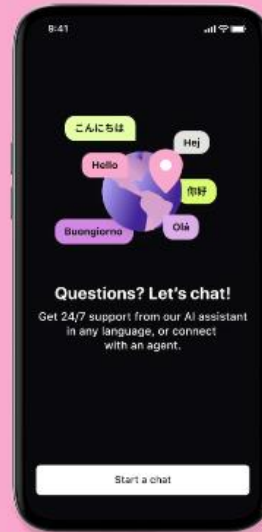
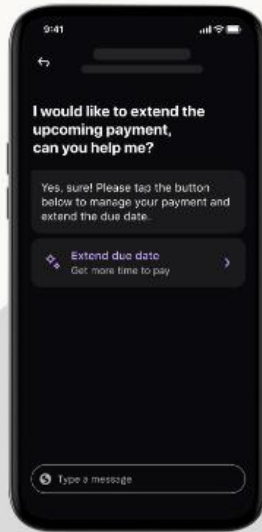
Source: LeewayHertz

Klarna AI assistant handles two-thirds of customer service chats in its first month

February 27, 2024



Klarna
Klarna



Klarna

#1

```
JS test.js 1 ●
JS test.js > calculateDaysBetweenDates
1 function calculateDaysBetweenDates(begin, end) {
    var beginDate = new Date(begin);
    var endDate = new Date(end);
    var days = Math.round((endDate - beginDate) / (1000 * 60 * 60 * 24));
    return days;
}
2
```



```
1 // create a class in TypeScript to represent a student that has a name, an id, and a list of courses
2 // the class should have method to add and remove a course to the list of courses
3
4 class Student {
    name: string;
    id: number;
    courses: string[];

    constructor(name: string, id: number, courses: string[]) {
        this.name = name;
        this.id = id;
        this.courses = courses;
    }

    addCourse(course: string) {
        this.courses.push(course);
    }

    removeCourse(course: string) {
        this.courses = this.courses.filter((c) => c !== course);
    }
}
```



#2

```
add typings and pydoc to this function
```

GitHub Copilot

> 3 steps completed successfully

```
def create_worker_agent(llm: ChatOpenAI, tools: List, system_prompt: str) -> AgentExecutor:
    """
    Create a worker agent with the given language model, tools, and system prompt.

    Args:
        llm (ChatOpenAI): The language model to use for the agent.
        tools (List): A list of tools that the agent can use.
        system_prompt (str): The system prompt to initialize the agent with.

    Returns:
        AgentExecutor: An executor for the created agent.
    """
    # Each worker node will be given a name and some tools.
    prompt = ChatPromptTemplate.from_messages(
        [
            ("system", system_prompt),
            MessagesPlaceholder(variable_name="messages"),
            MessagesPlaceholder(variable_name="agent_scratchpad"),
        ]
    )
    agent = create_openai_tools_agent(llm, tools, prompt)
    executor = AgentExecutor(agent=agent, tools=tools)

    return executor
```

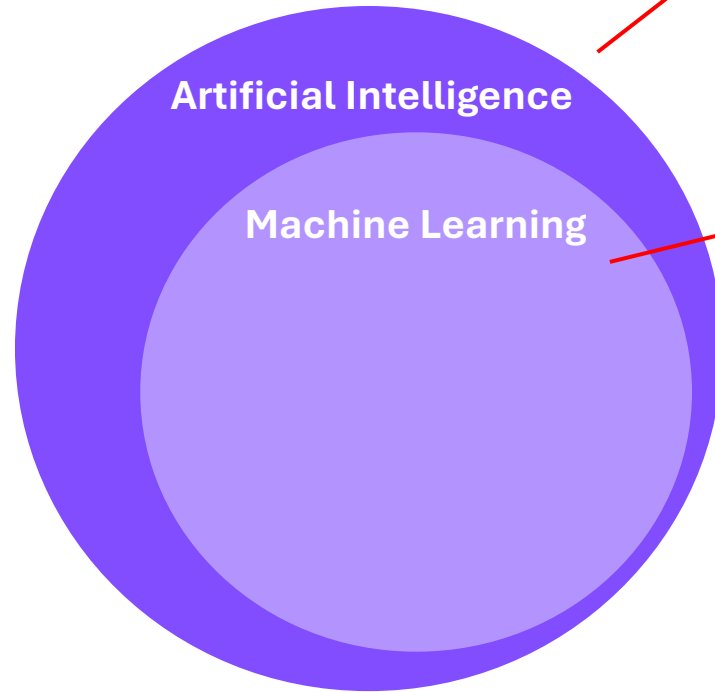
> Using 4 References

```
def create_worker_agent(llm, tools, system_prompt): You, Moments ago • Uncommitted changes
# Each worker node will be given a name and some tools.
prompt = ChatPromptTemplate.from_messages(
    [
        (
            "system",
            system_prompt,
        ),
        MessagesPlaceholder(variable_name="messages")
    ]
)
```

What is **Generative AI**



Intelligent autonomous machines



Artificial Intelligence

Machine Learning

Pattern recognition
Learn general patterns
form data



My experience
has been
fantastic!

POSITIVE



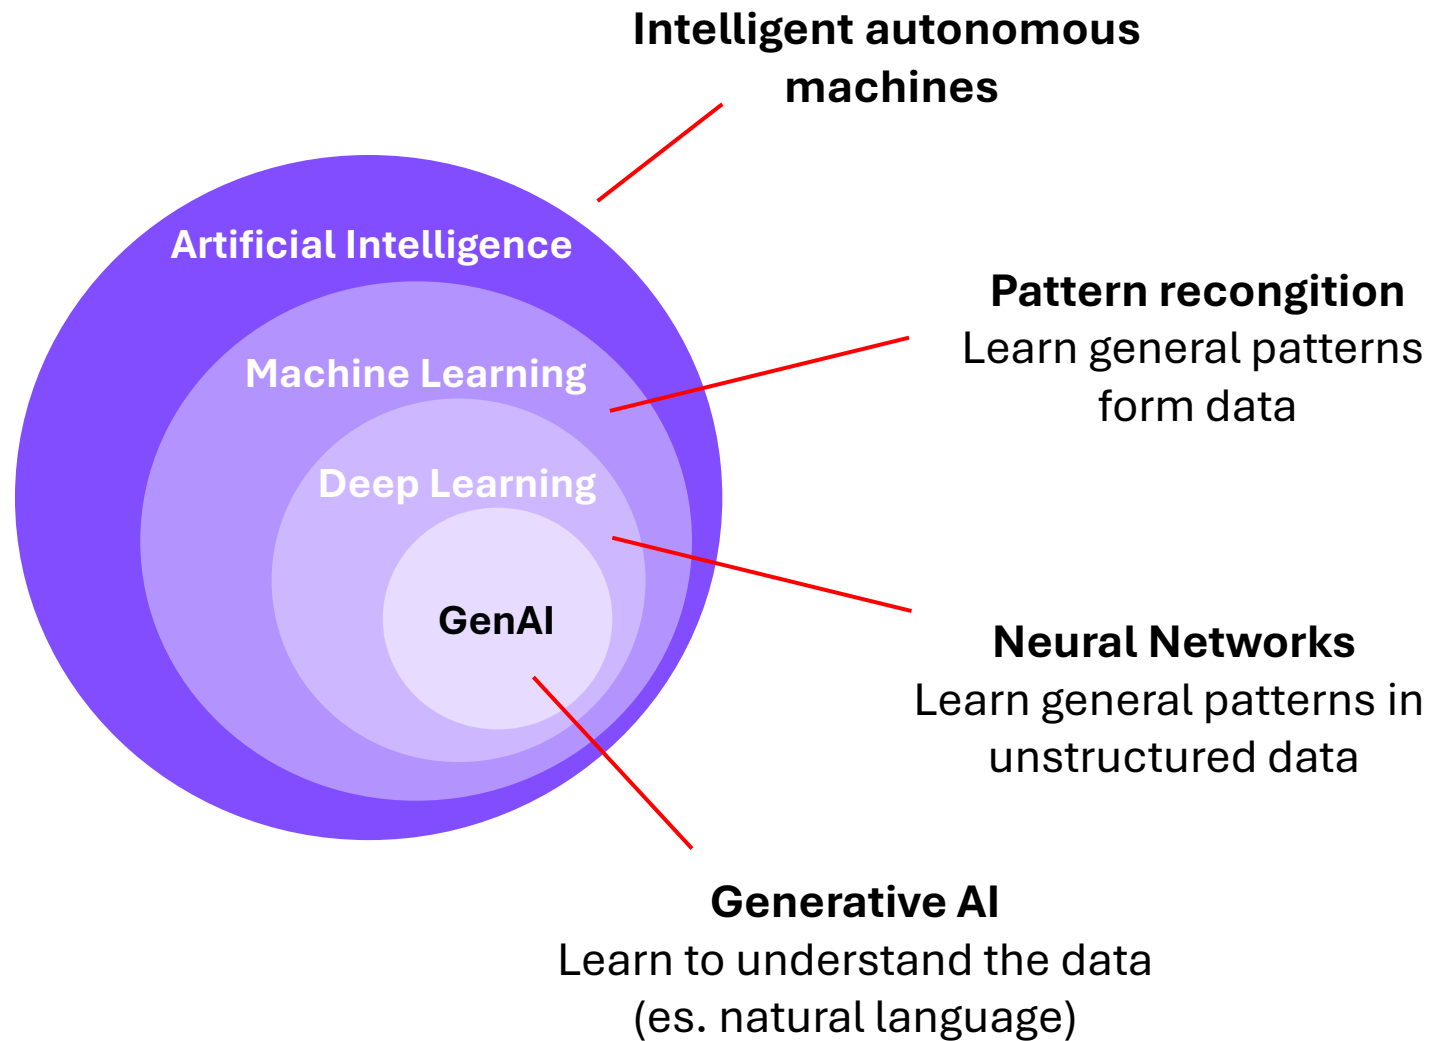
The product is
ok I guess

NEUTRAL



Your support
team is
useless

NEGATIVE



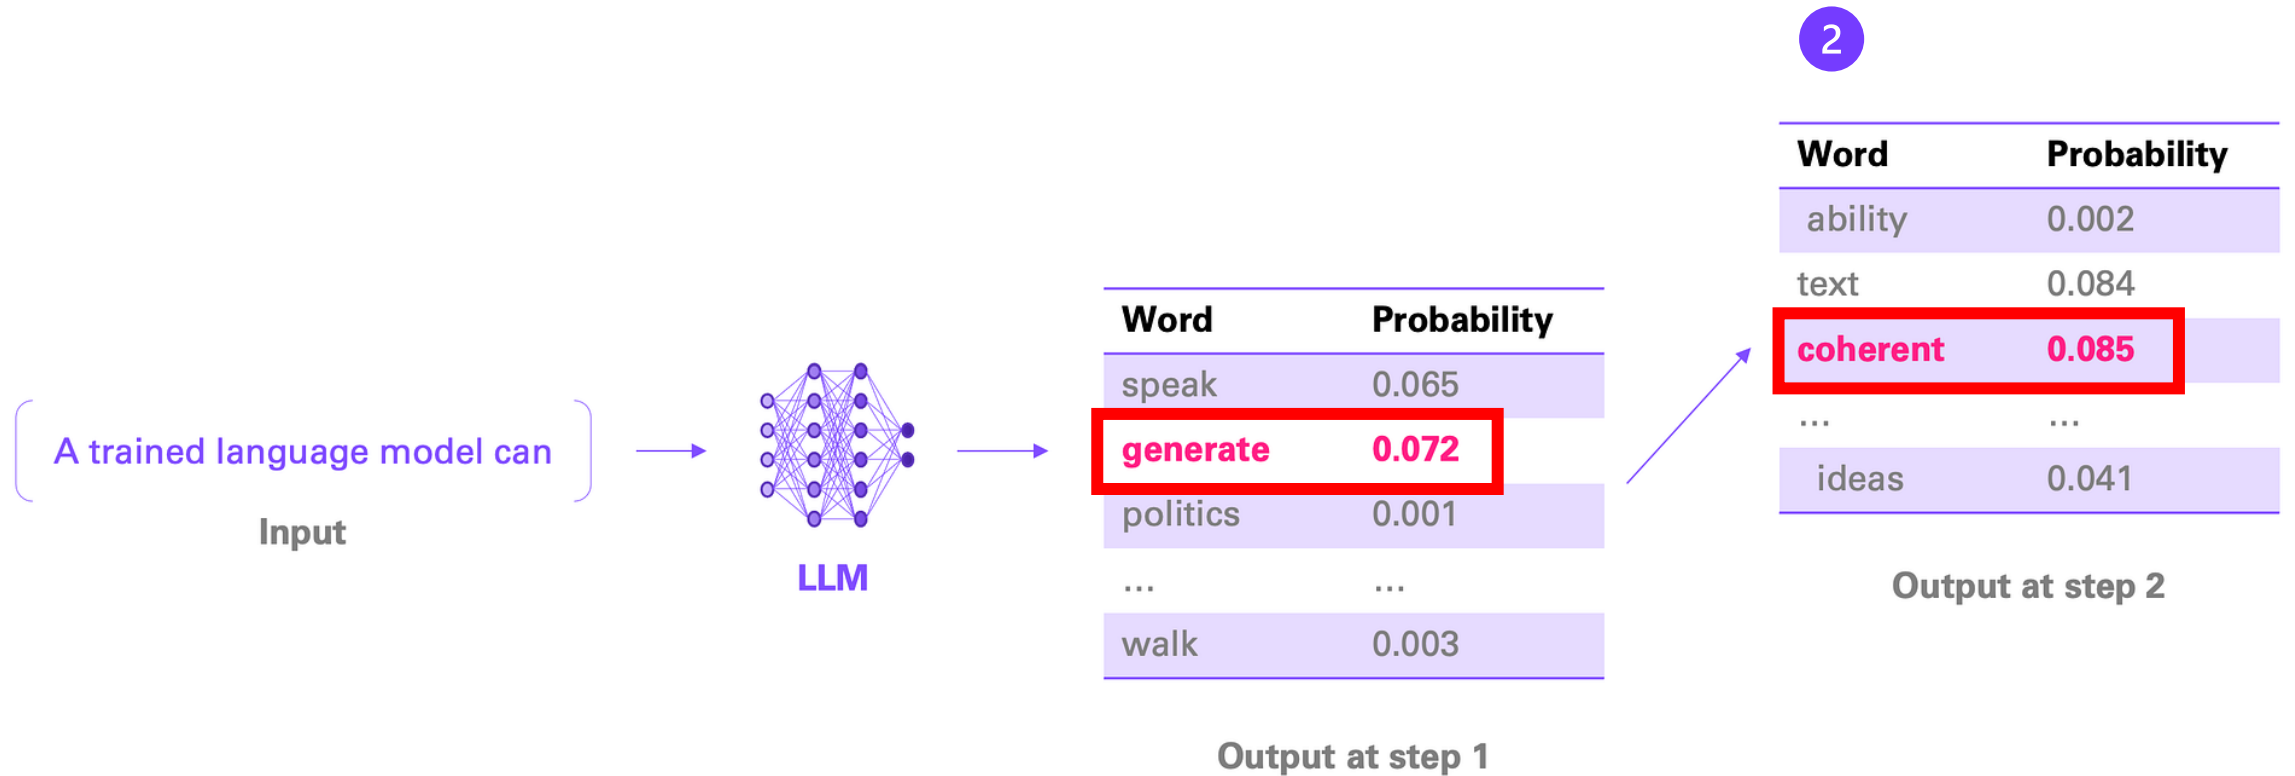
Large
Language
Model



LLM

How does a
language model
work?

A trained language model can generate coherent...

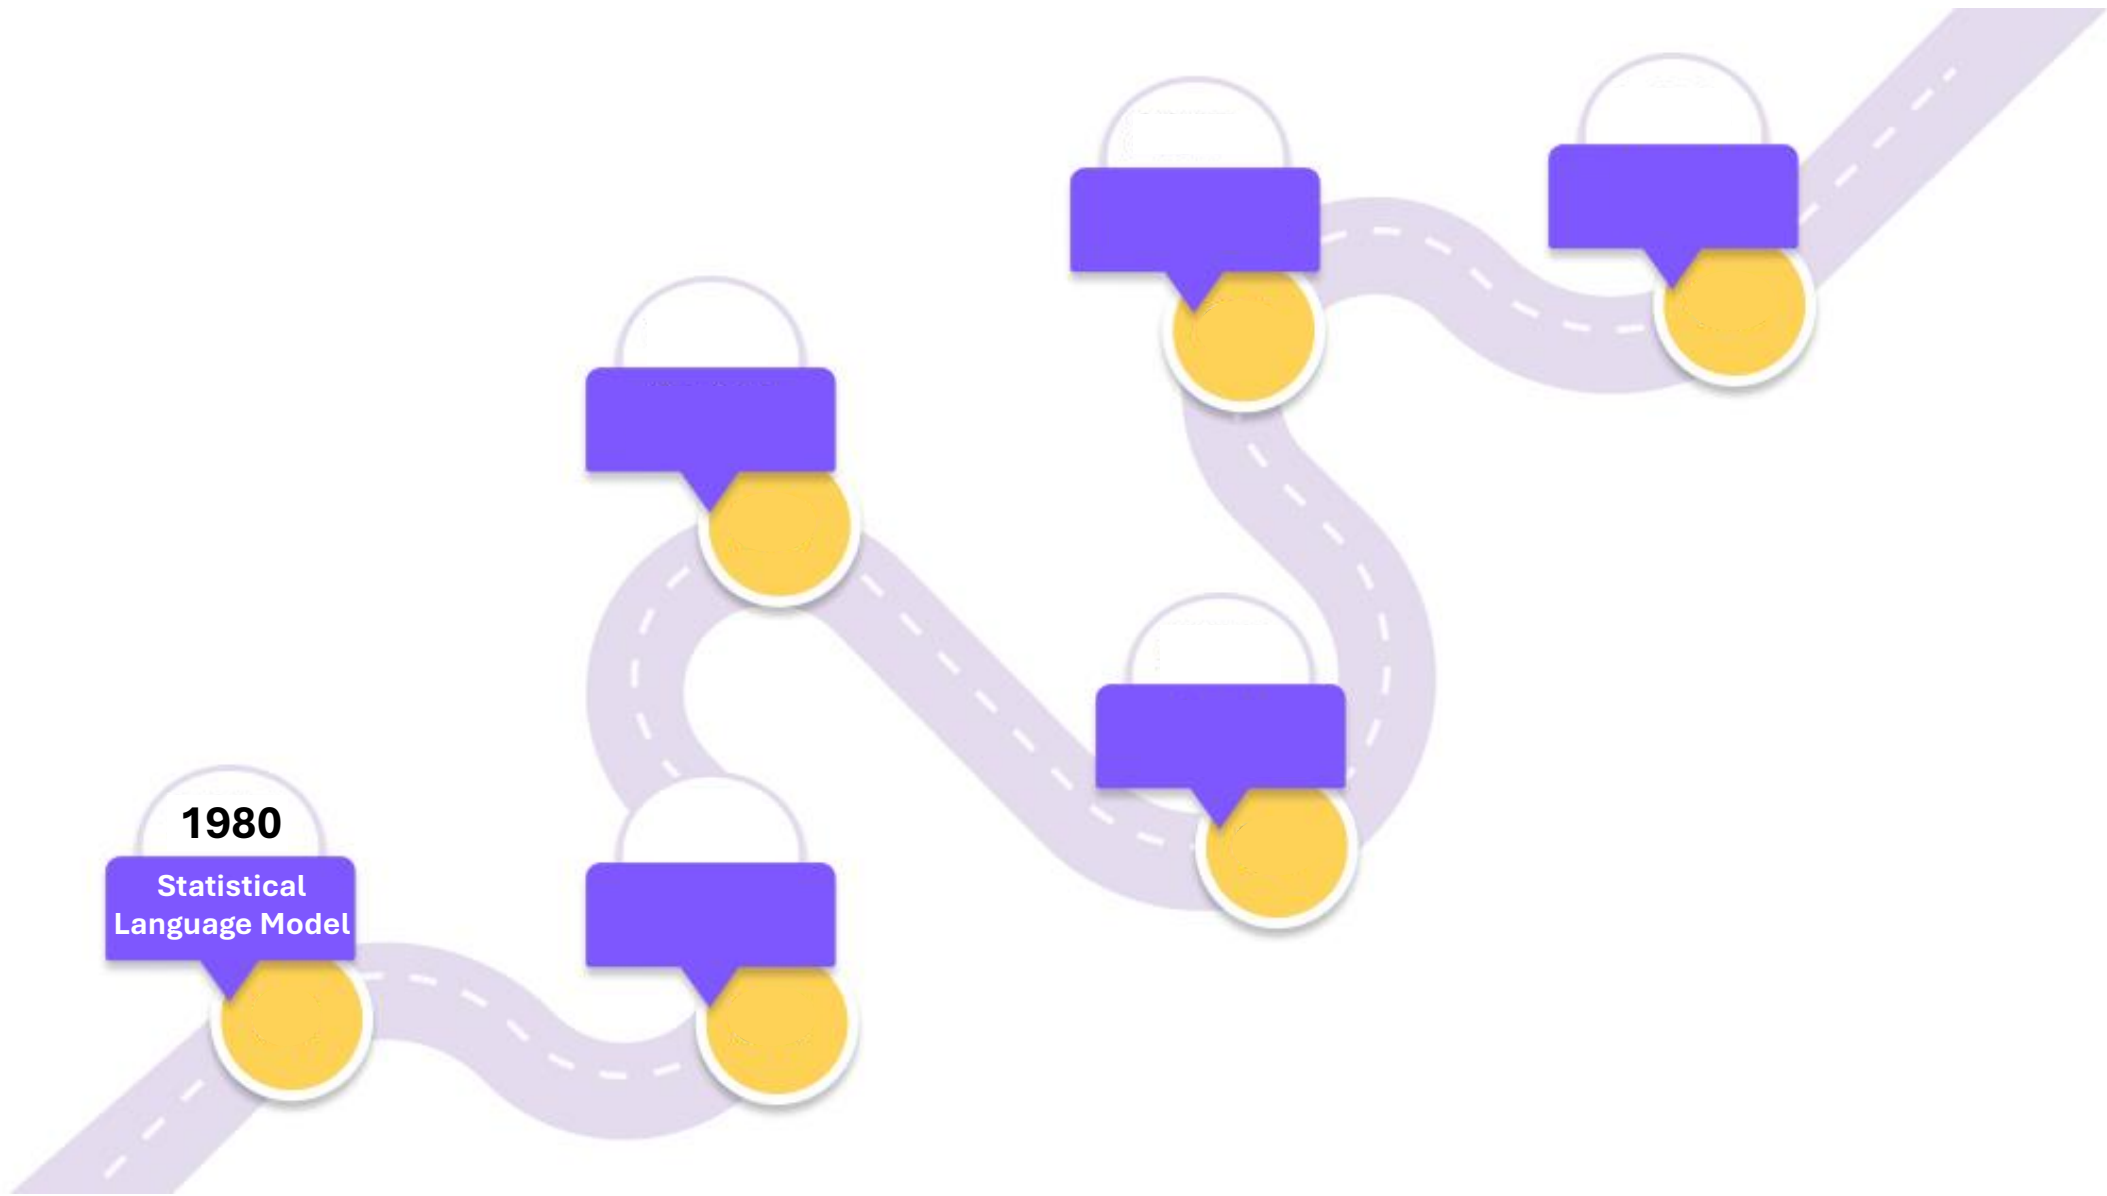


A trained language model can generate...

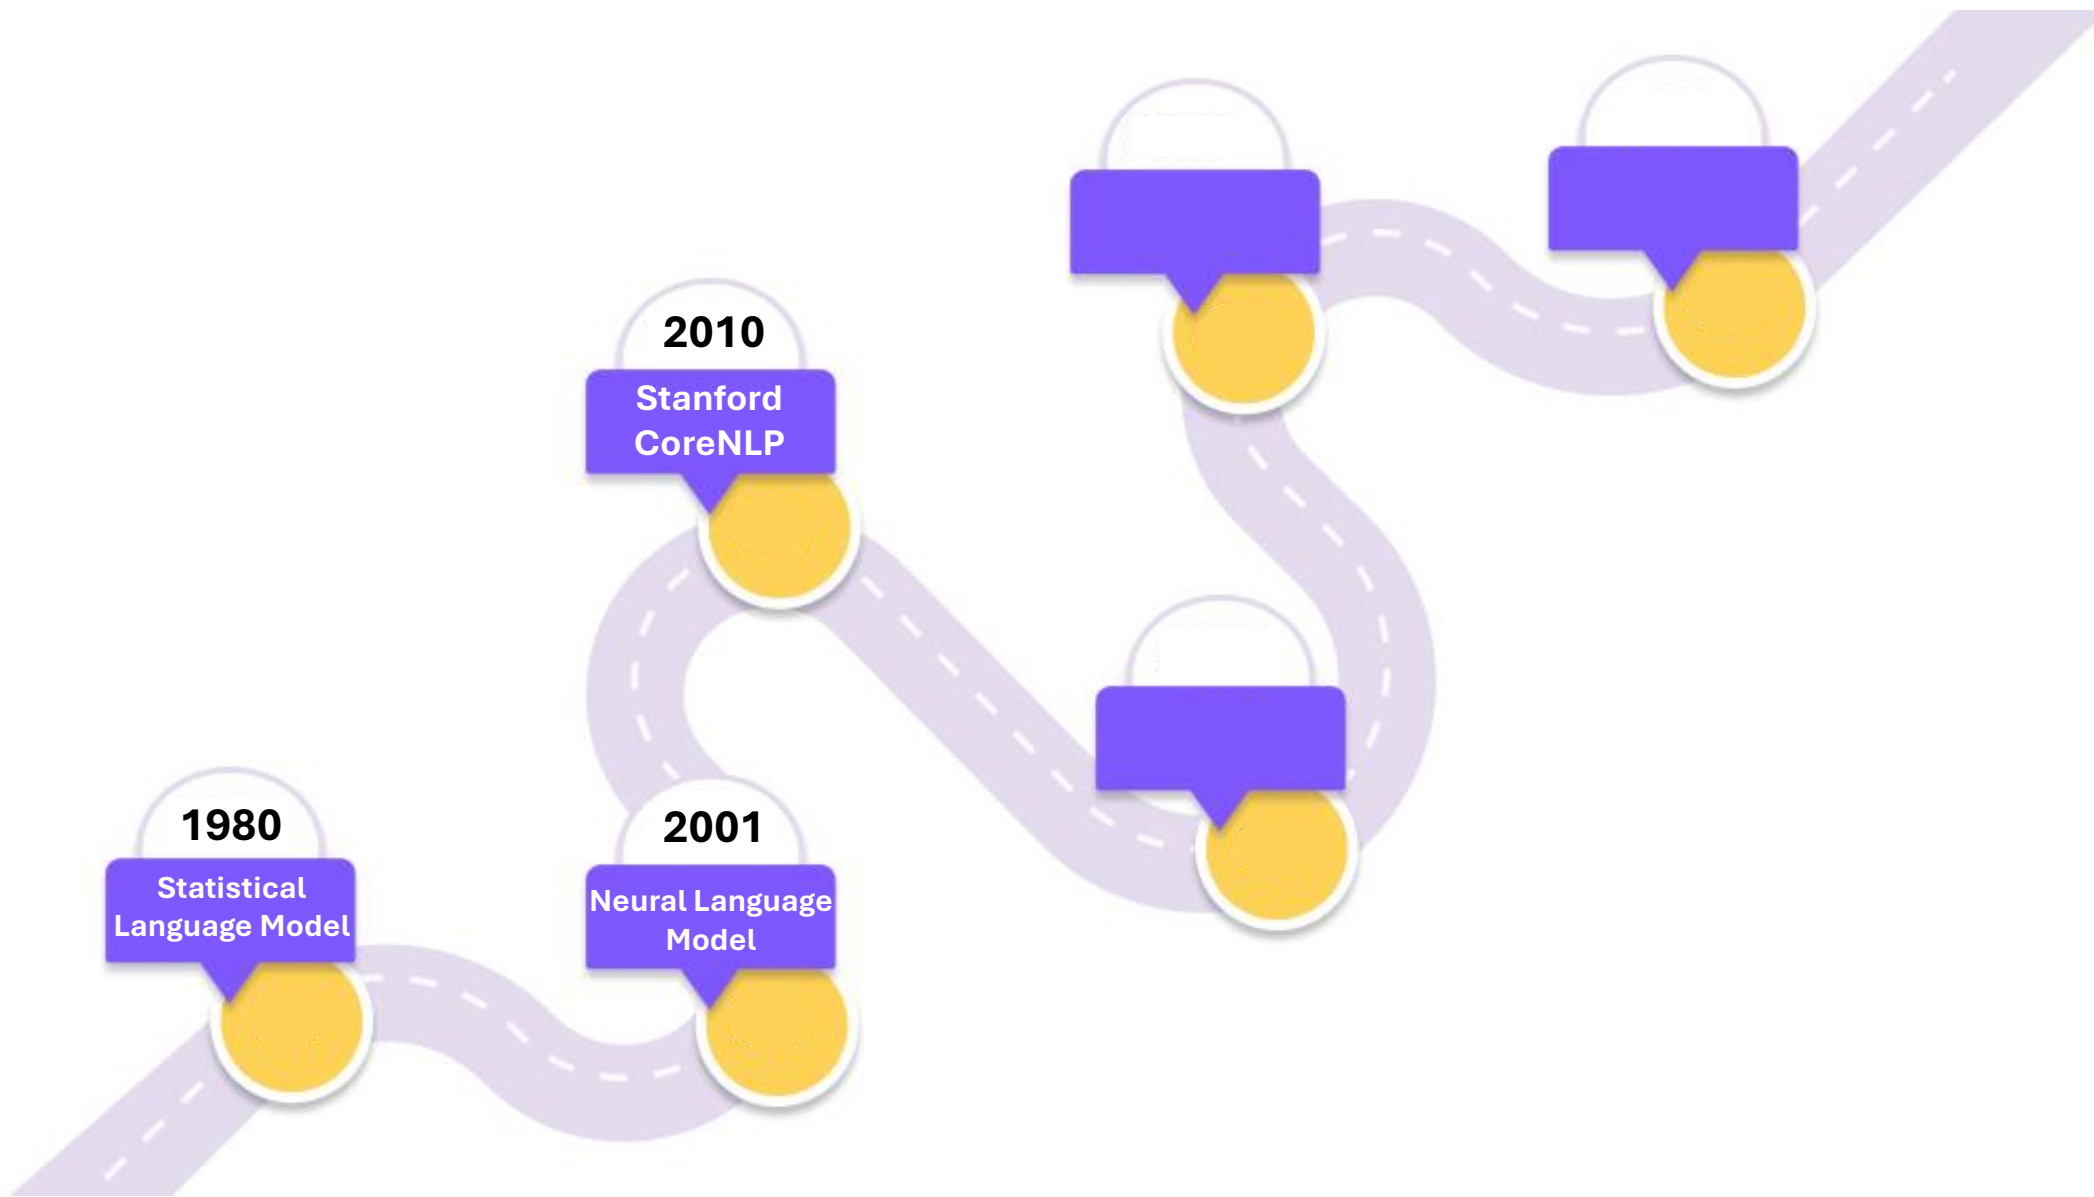
1

Source: medium.com

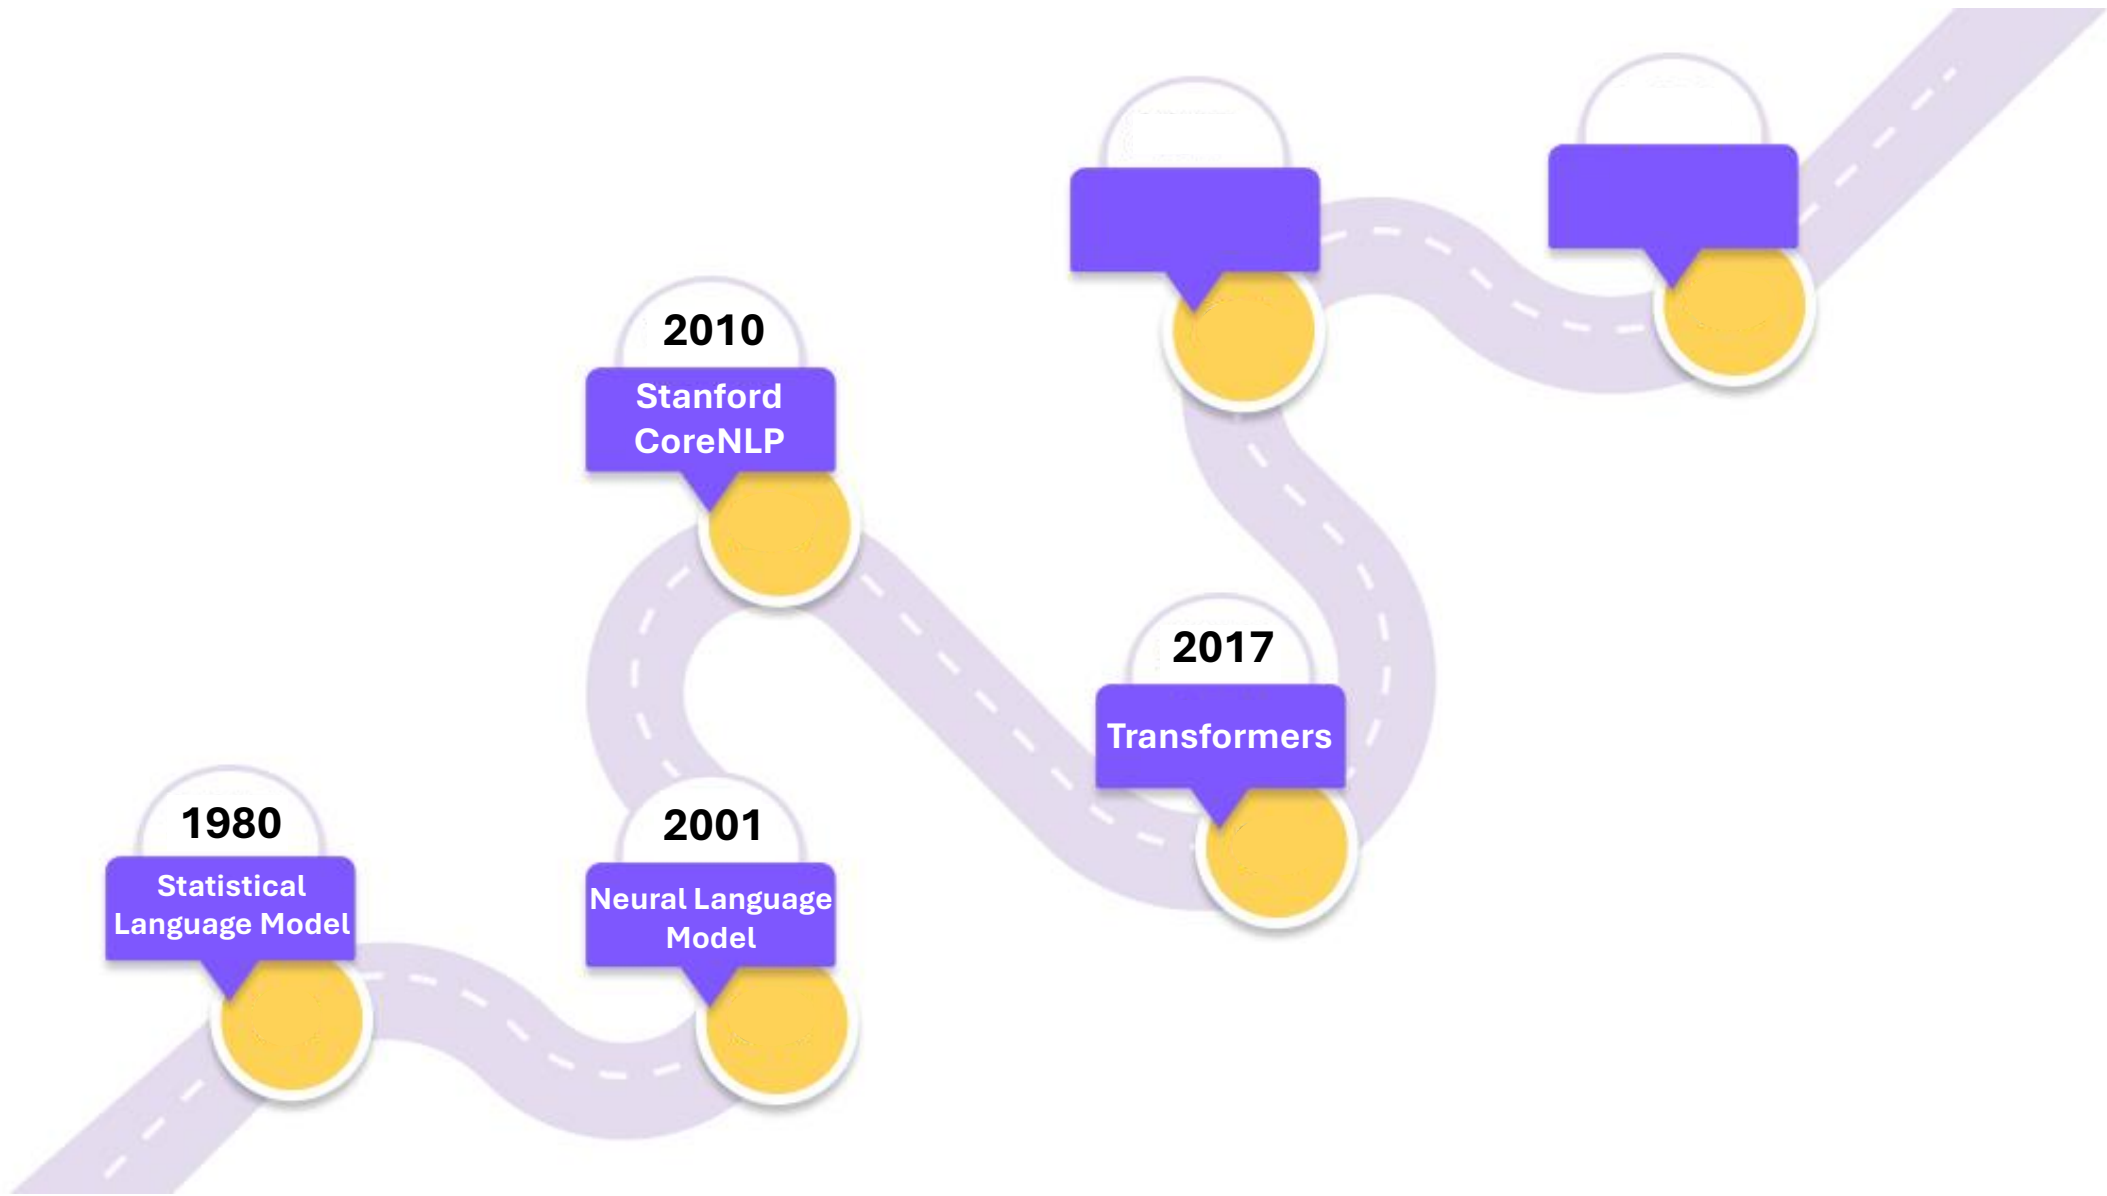
How does a language model work?



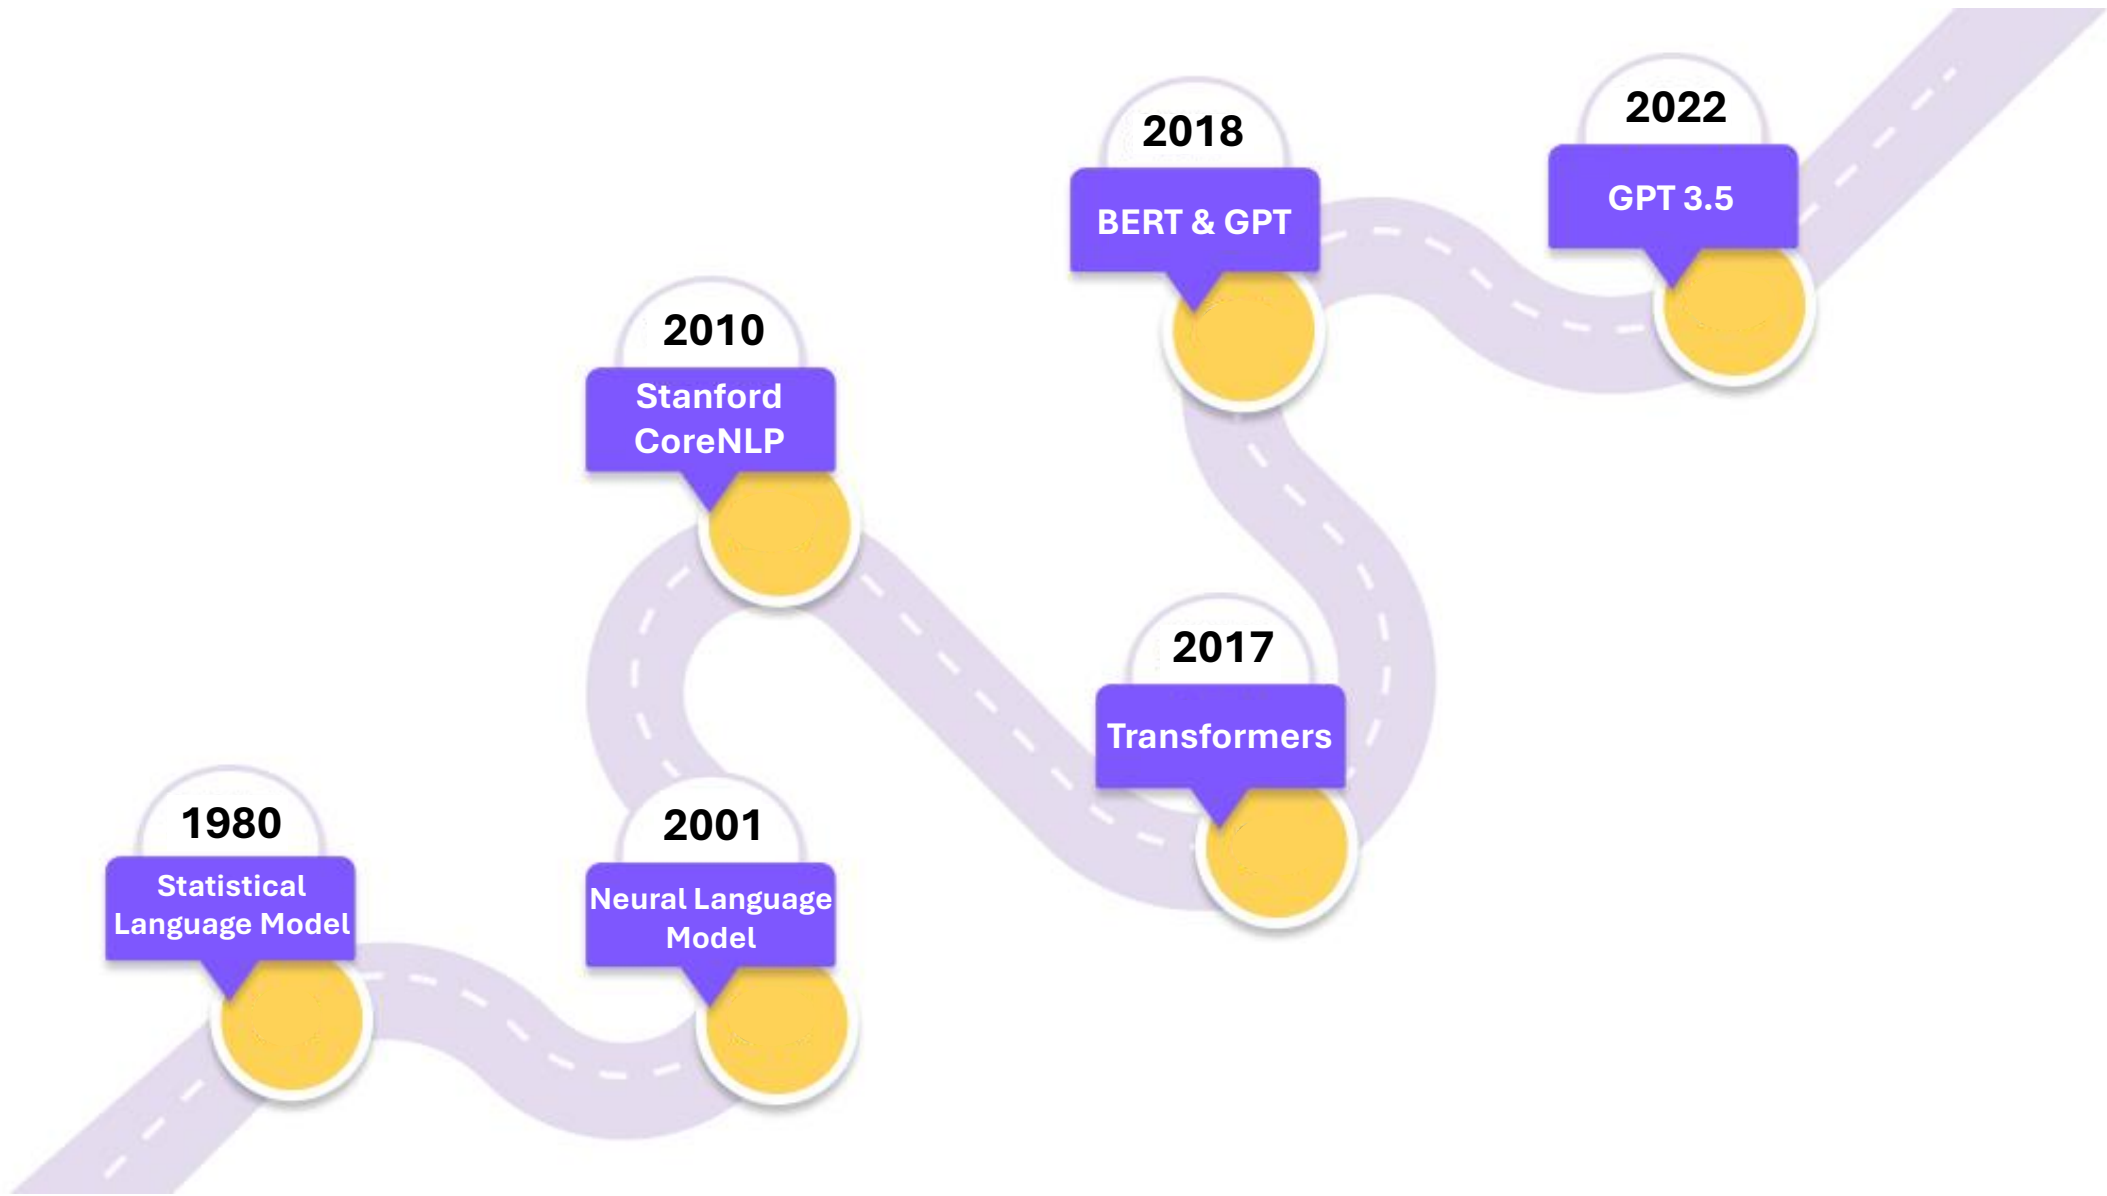
History of LLMs



History of LLMs

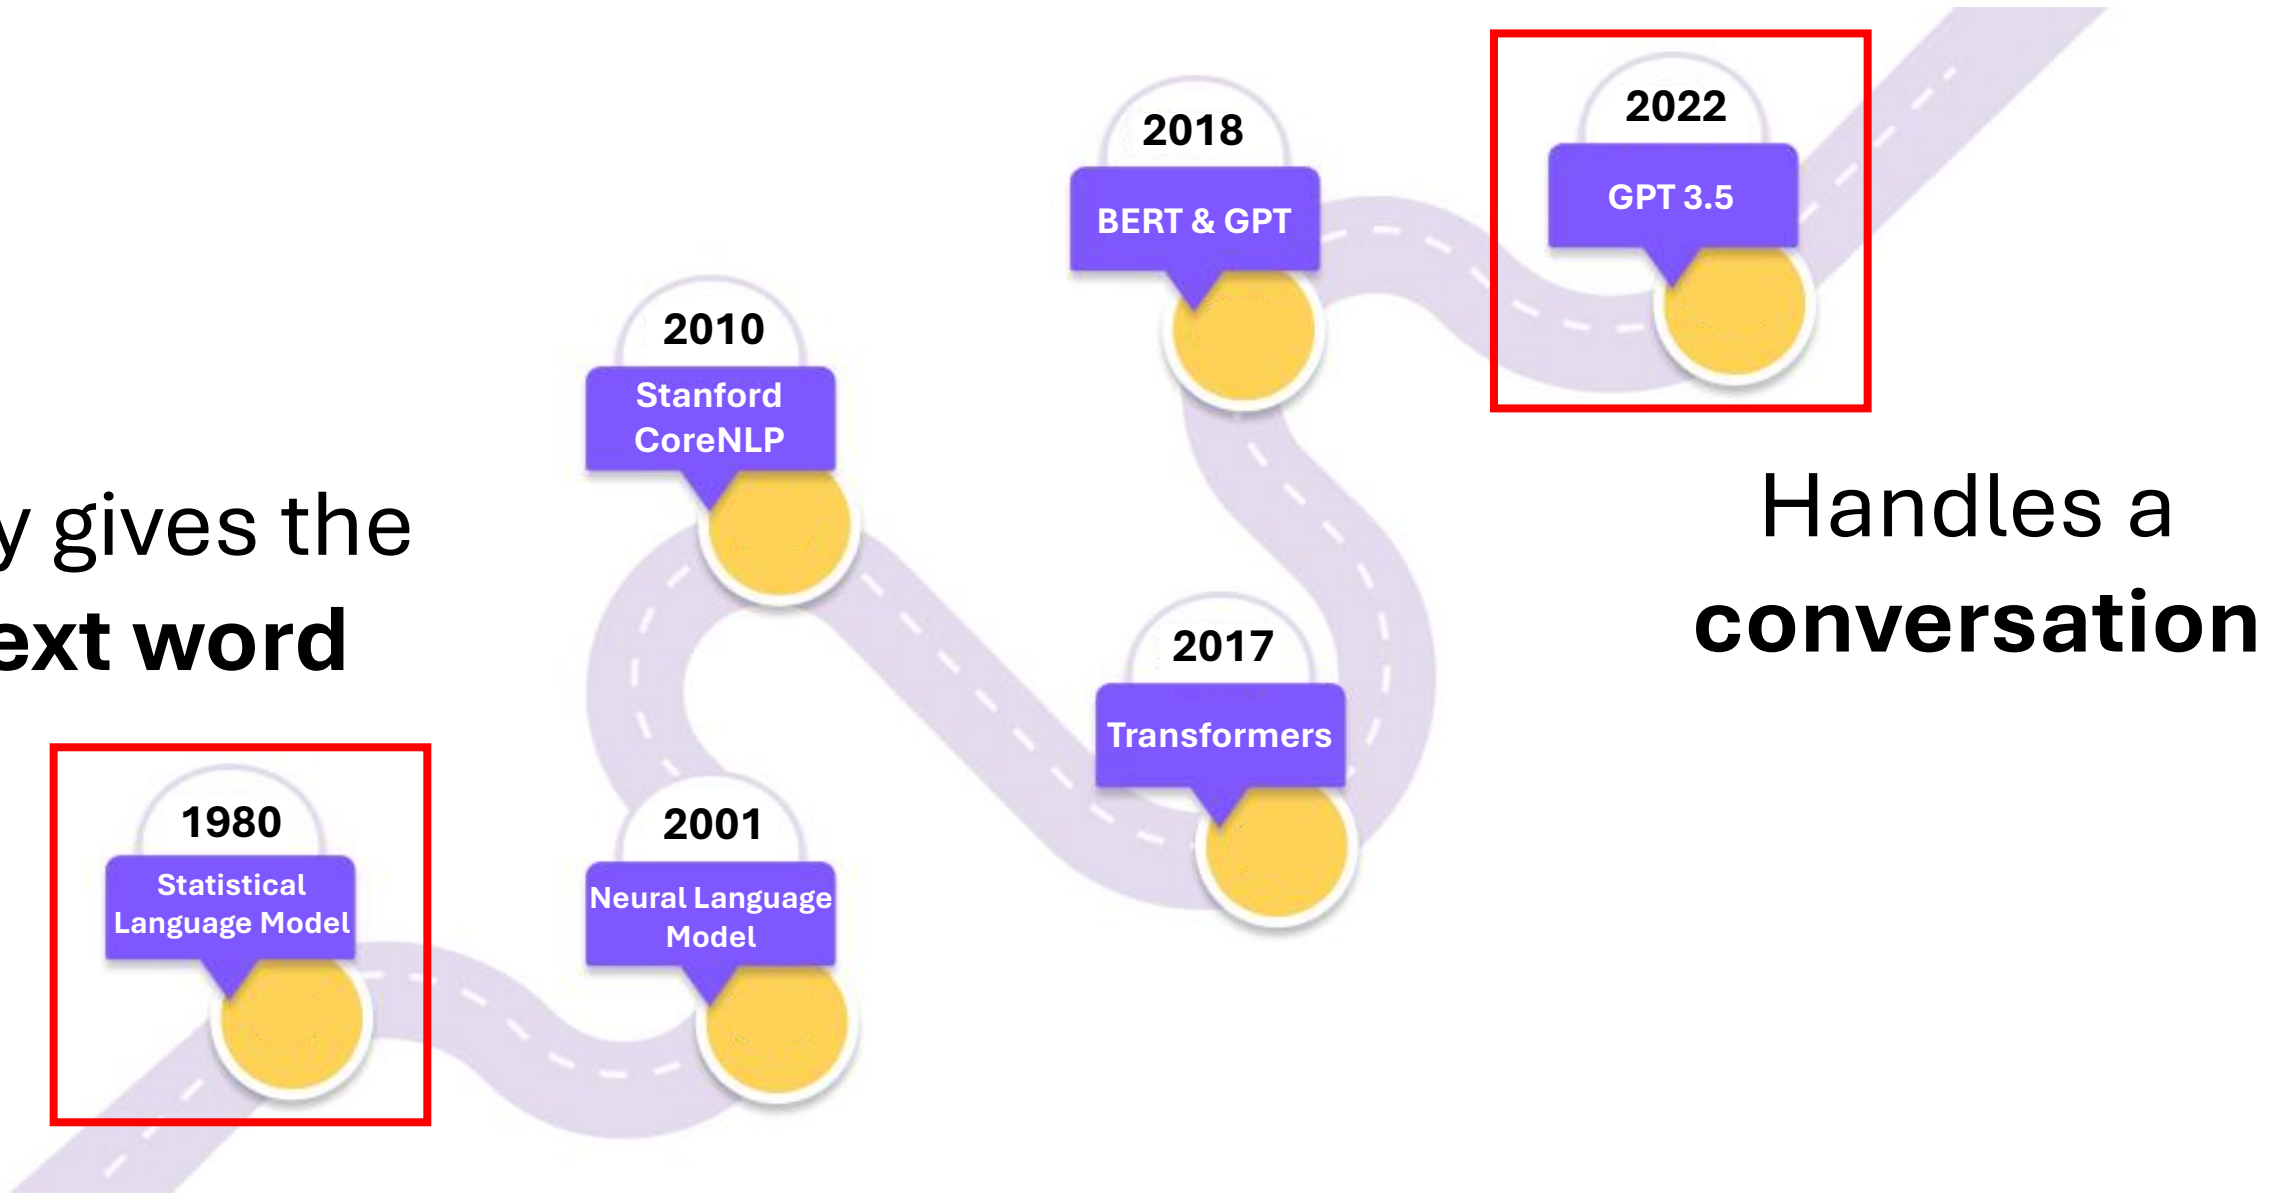


History of LLMs



History of LLMs

Only gives the
next word



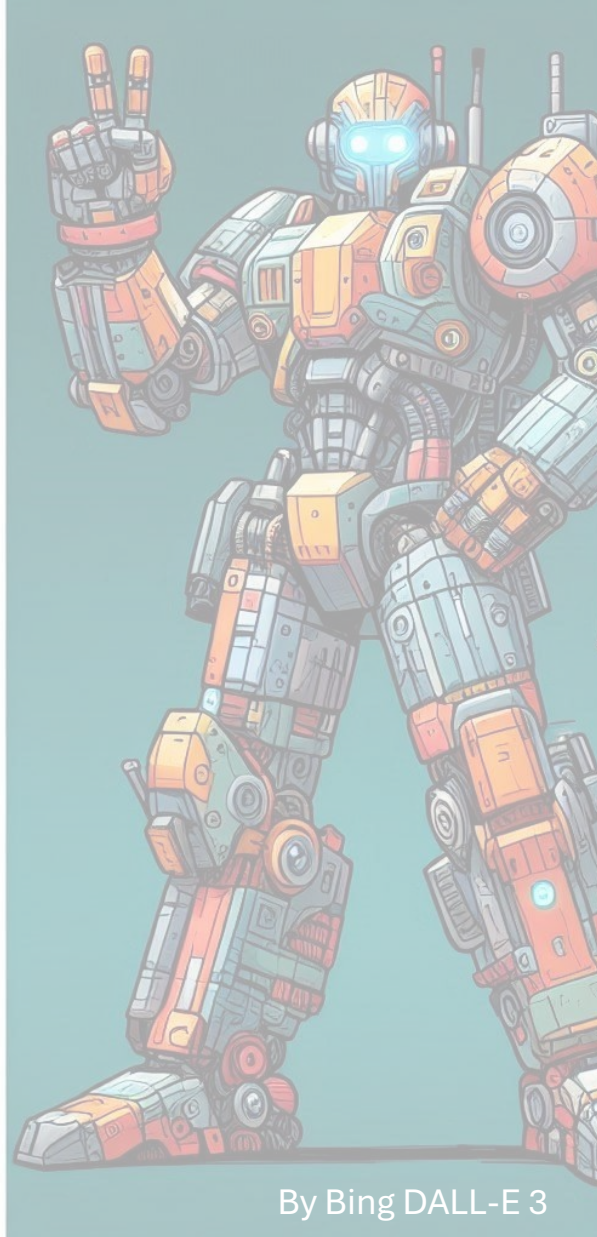
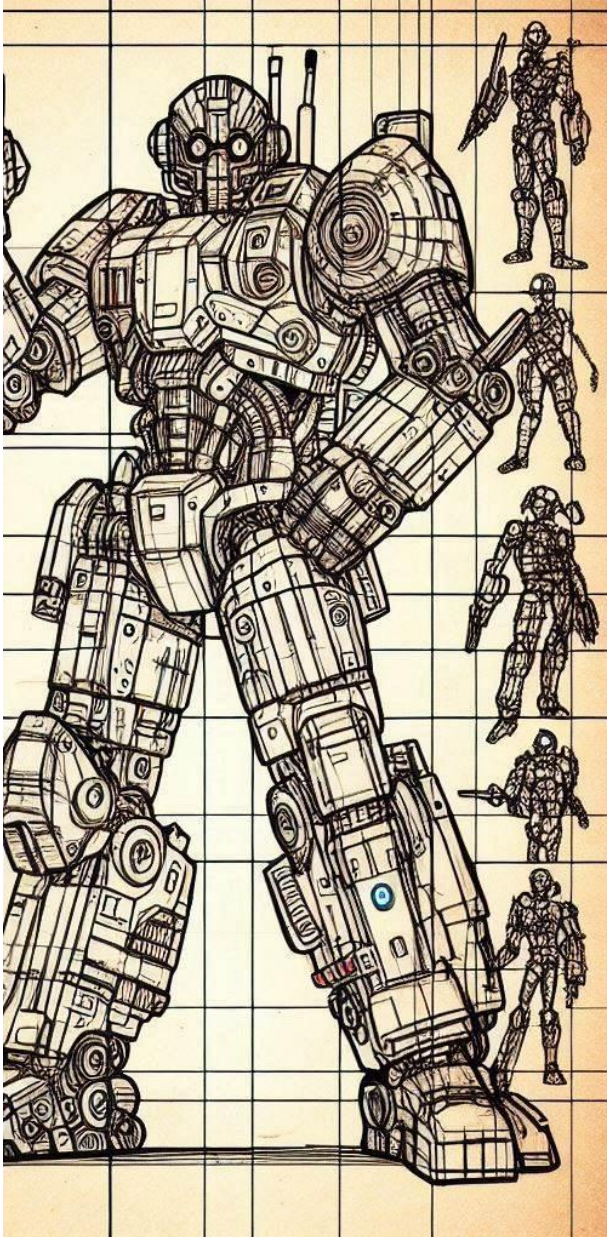
Handles a
conversation

History of LLMs

How to **create** a **Large Language Model?**

Pre-training

Learn
language structure
and knowledge

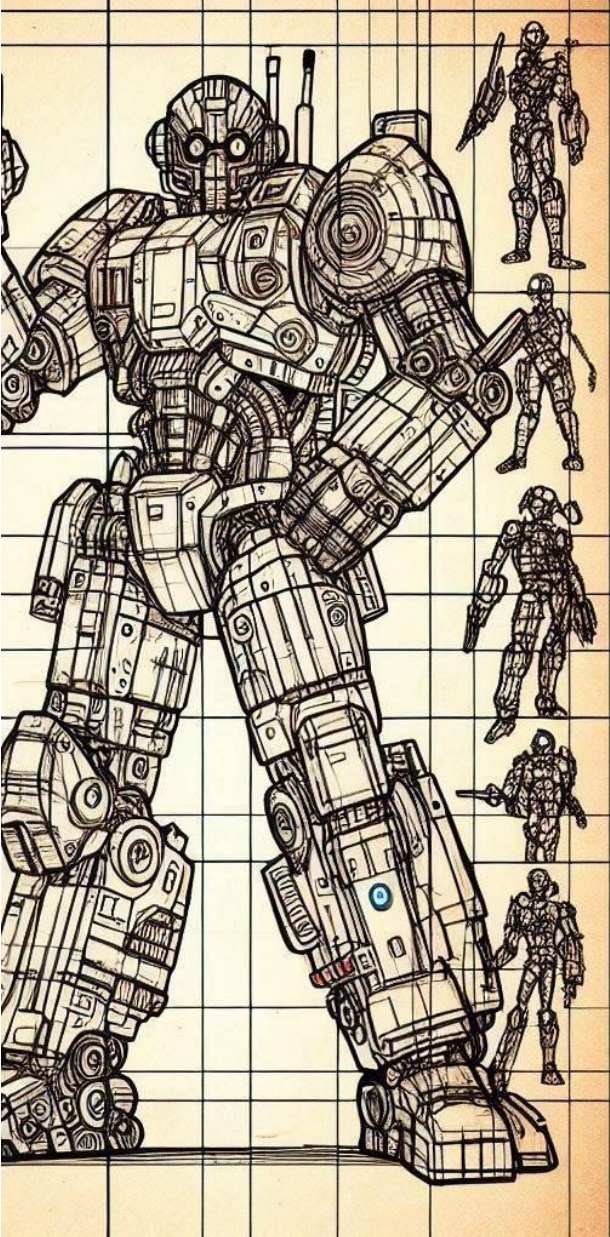


By Bing DALL-E 3

How to train a «Large» Language Model?

Pre-training

Learn
language structure
and knowledge



Instruction Fine-tuning

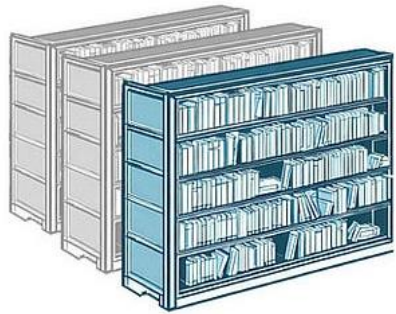
Learn
how to chat



How to train a «Large» Language Model?

GPT-4

Est. train size



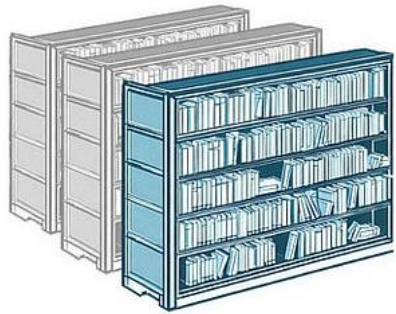
650 km

Line of library shelves

How to train a «Large» Language Model?

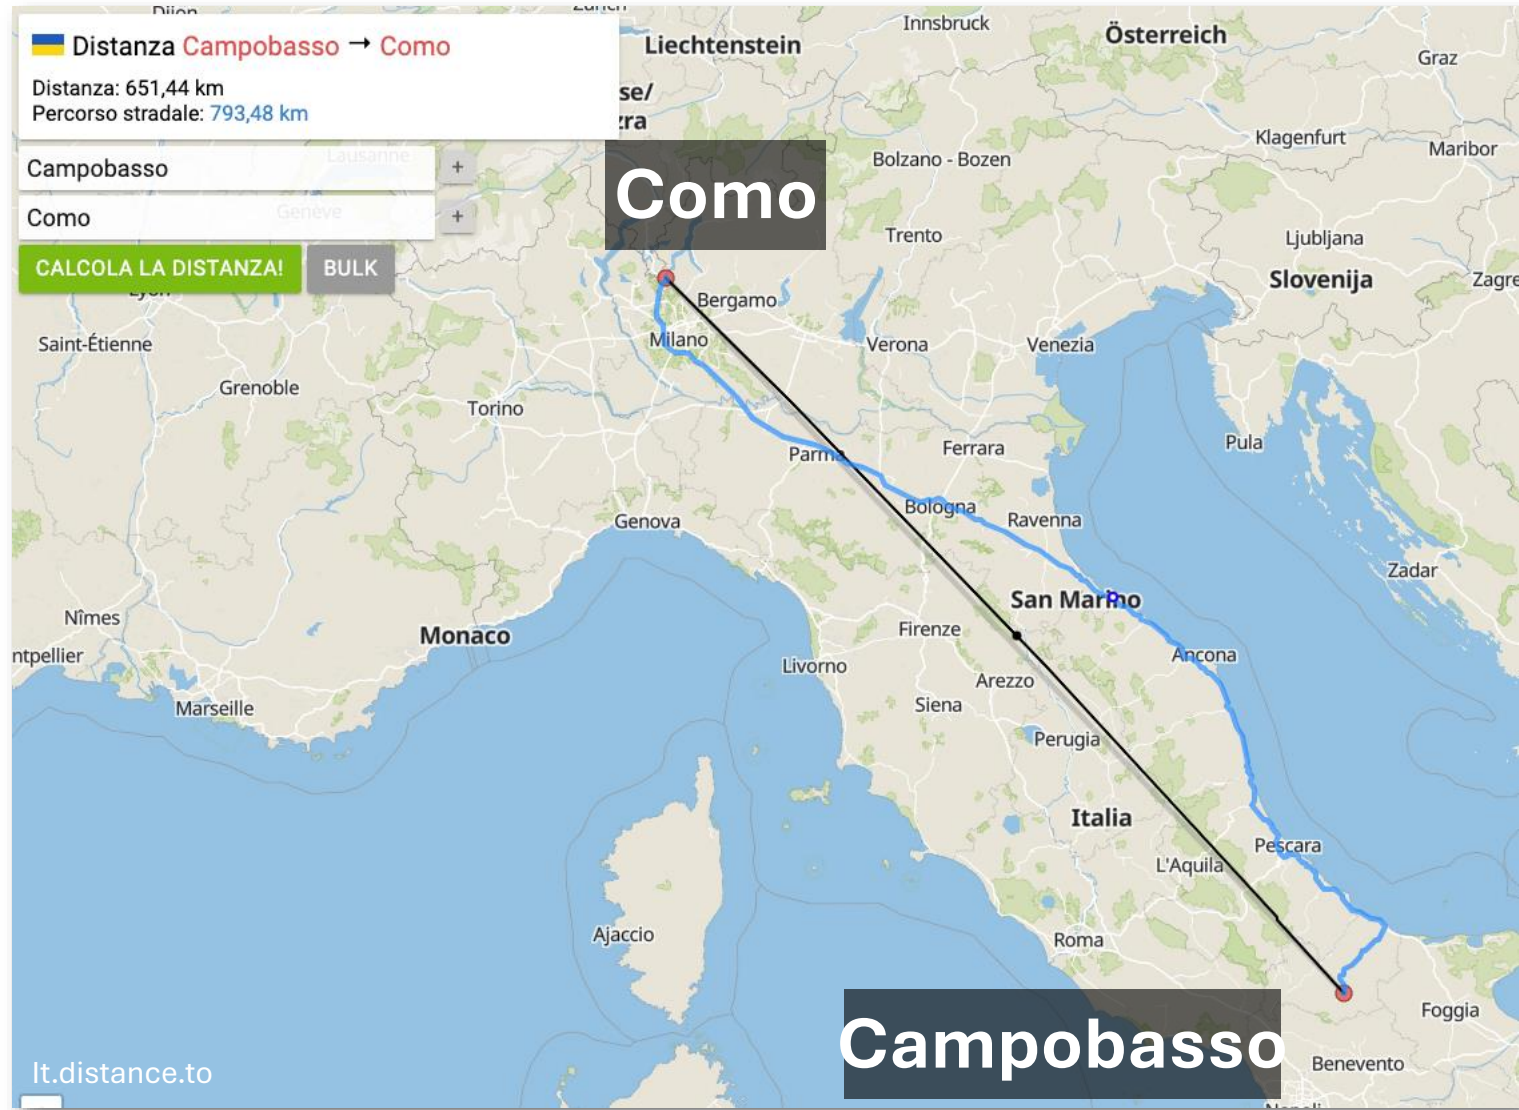
GPT-4

Est. train size



650 km

Line of library shelves



How to train a «Large» Language Model?



**G
P
T**



Generative

P

T



Generative

Pre-trained

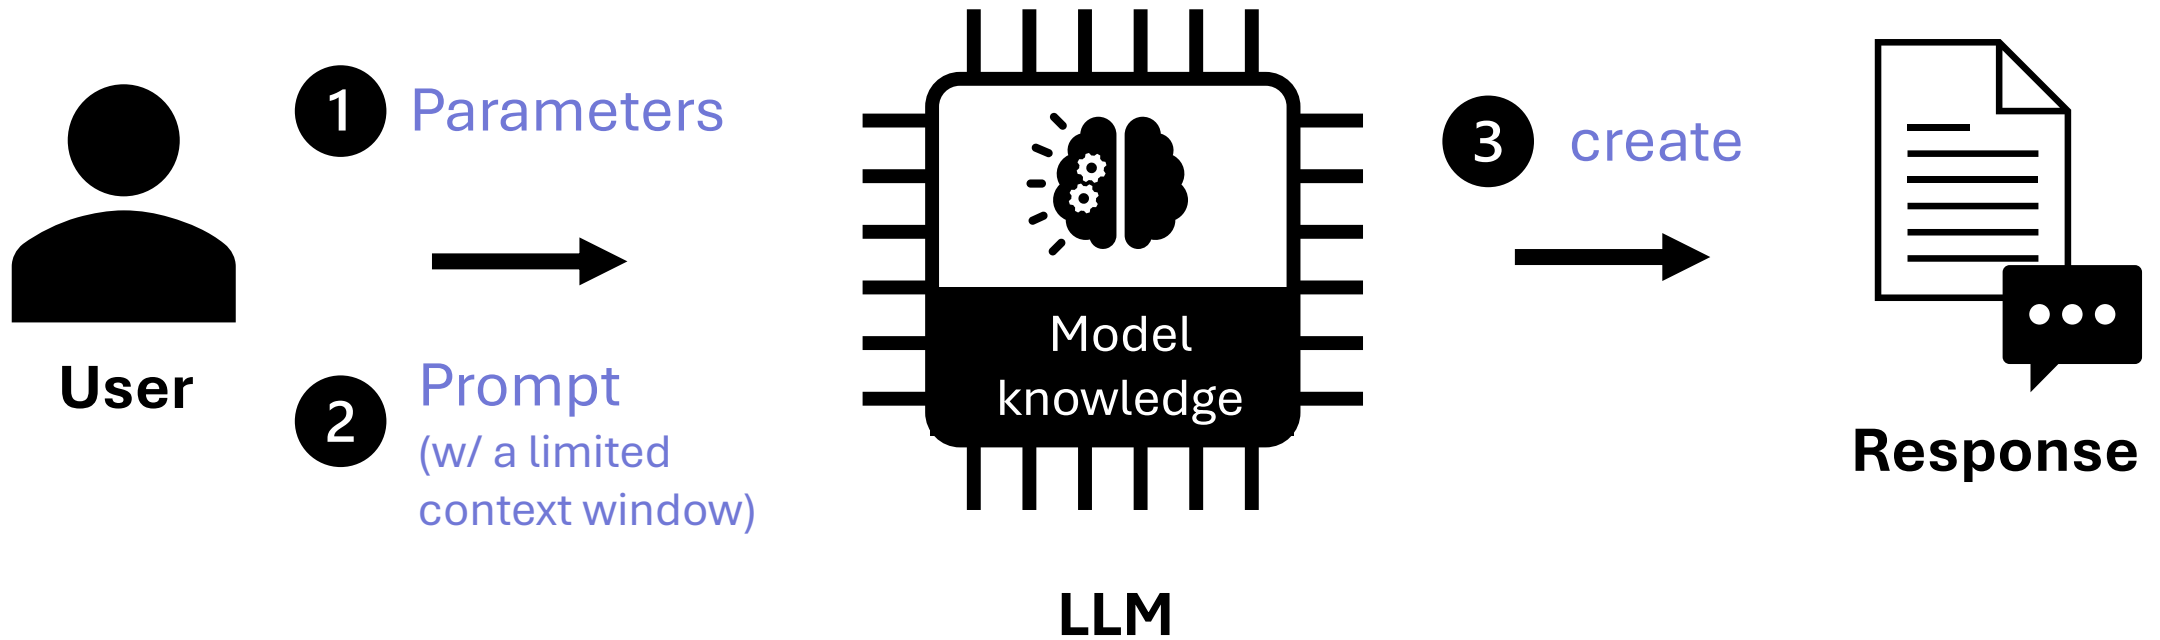
T



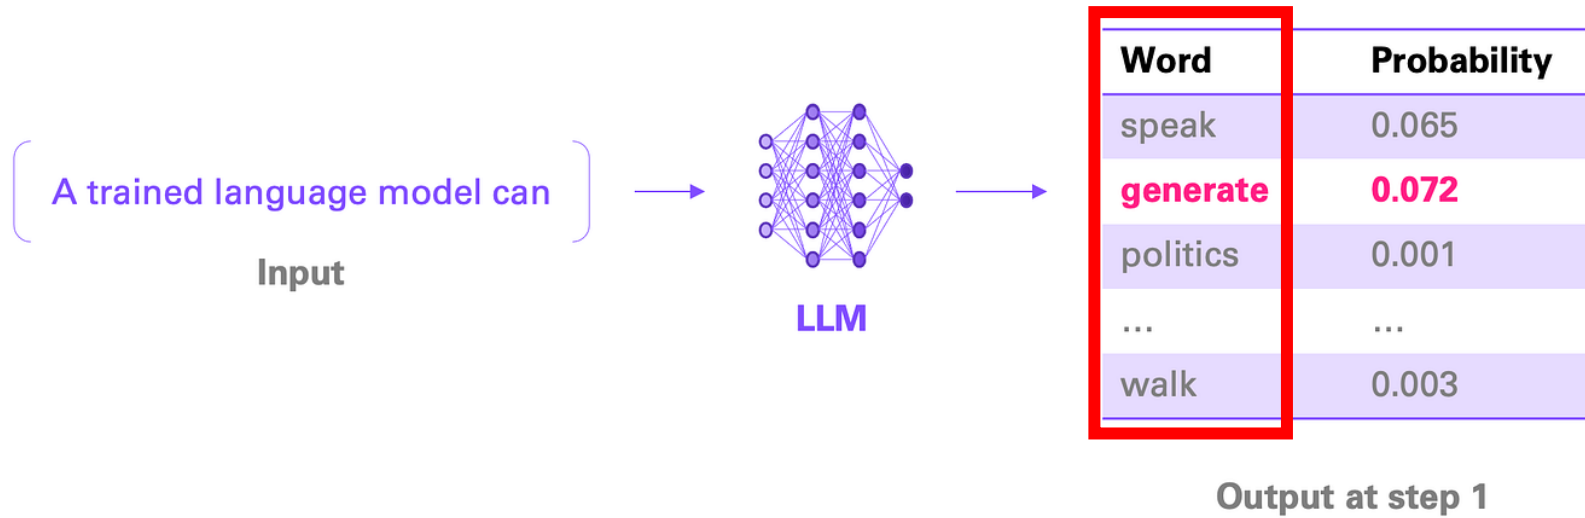
Generative
Pre-trained
Transformer

How to **instruct** LLMs





How to instruct LLMs ?

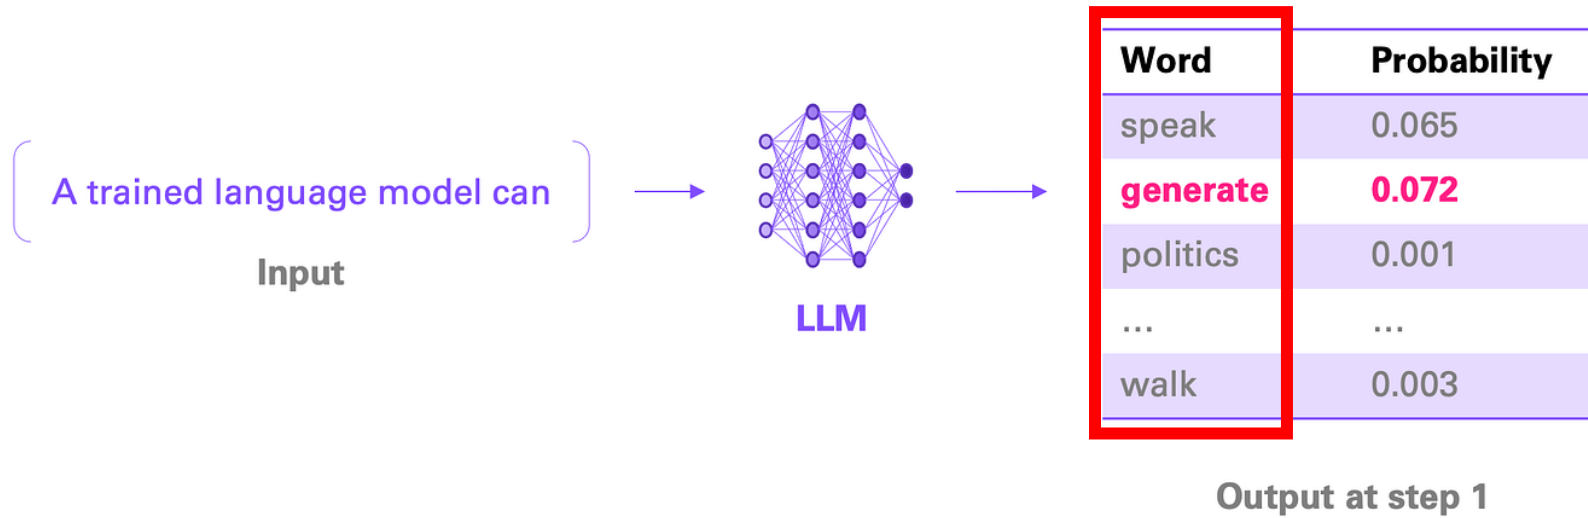


... [generate] [unbelievable] [textual] [sentences]

word != tokens

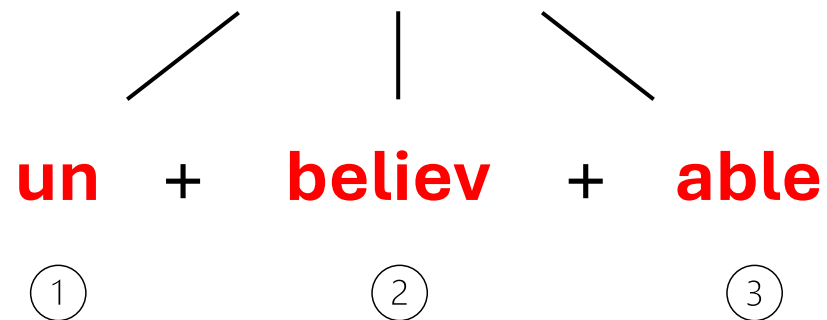
Source: medium.com

How does a language model work?



... [generate] [unbelievable] [textual] [sentences]

1 word \approx 2-3 tokens



Source: medium.com

How does a language model work?

Generative AI

Overview

Playground

Chat

Embedding

Dedicated AI clusters

Custom models

Endpoints

Scope

Compartment

Chat ?

To get started, choose a model and a preset prompt example. Then, refine the prompts and parameters to fit your use cases. See [model types](#) for more information. All model responses have [moderation filtering](#) applied for explicit content. Note that some models have deprecation/retirement dates. View our [model list](#) for more details.

Model

cohere.command-r-plus v1.2

View model details

Example

Choose example

View code

Hi! Who are you?



I am an AI-assistant chatbot trained to assist human users by providing thorough responses. I am always learning and improving based on user feedback. How can I help you today?

1 Prompt**2 Parameters**

Parameters

Maximum output tokens *i*

600

Preamble override *i*

Provide concise responses

Temperature *i*

0.25

Top p *i*

0.75

Top k *i*

0

Frequency penalty *i*

0

Presence penalty *i*

0

Type a message...

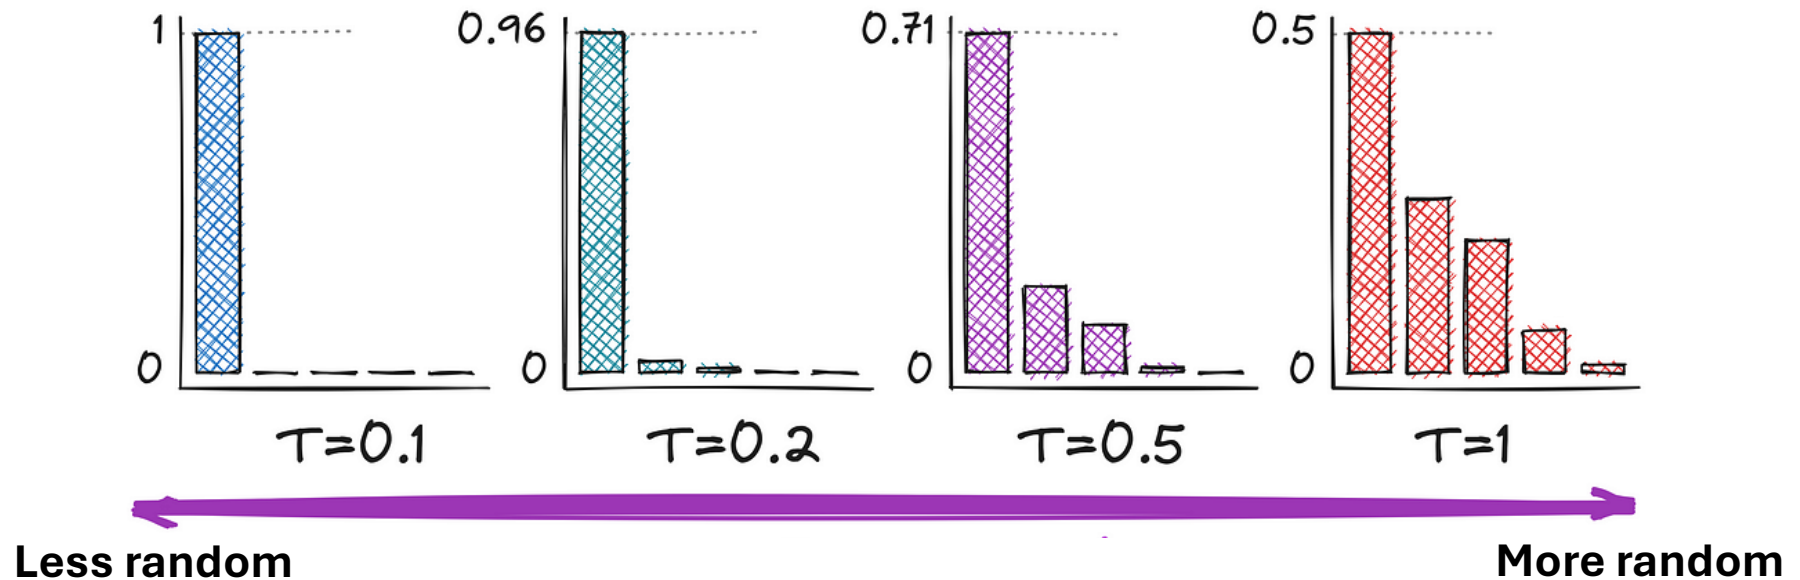
Submit

Clear chat

OCI Generative AI Playground

Temperature

controls randomness and creativity



Source: medium.com

Model: cohere.command-r-plus v1.2 | View model details | Example: Choose example | View code

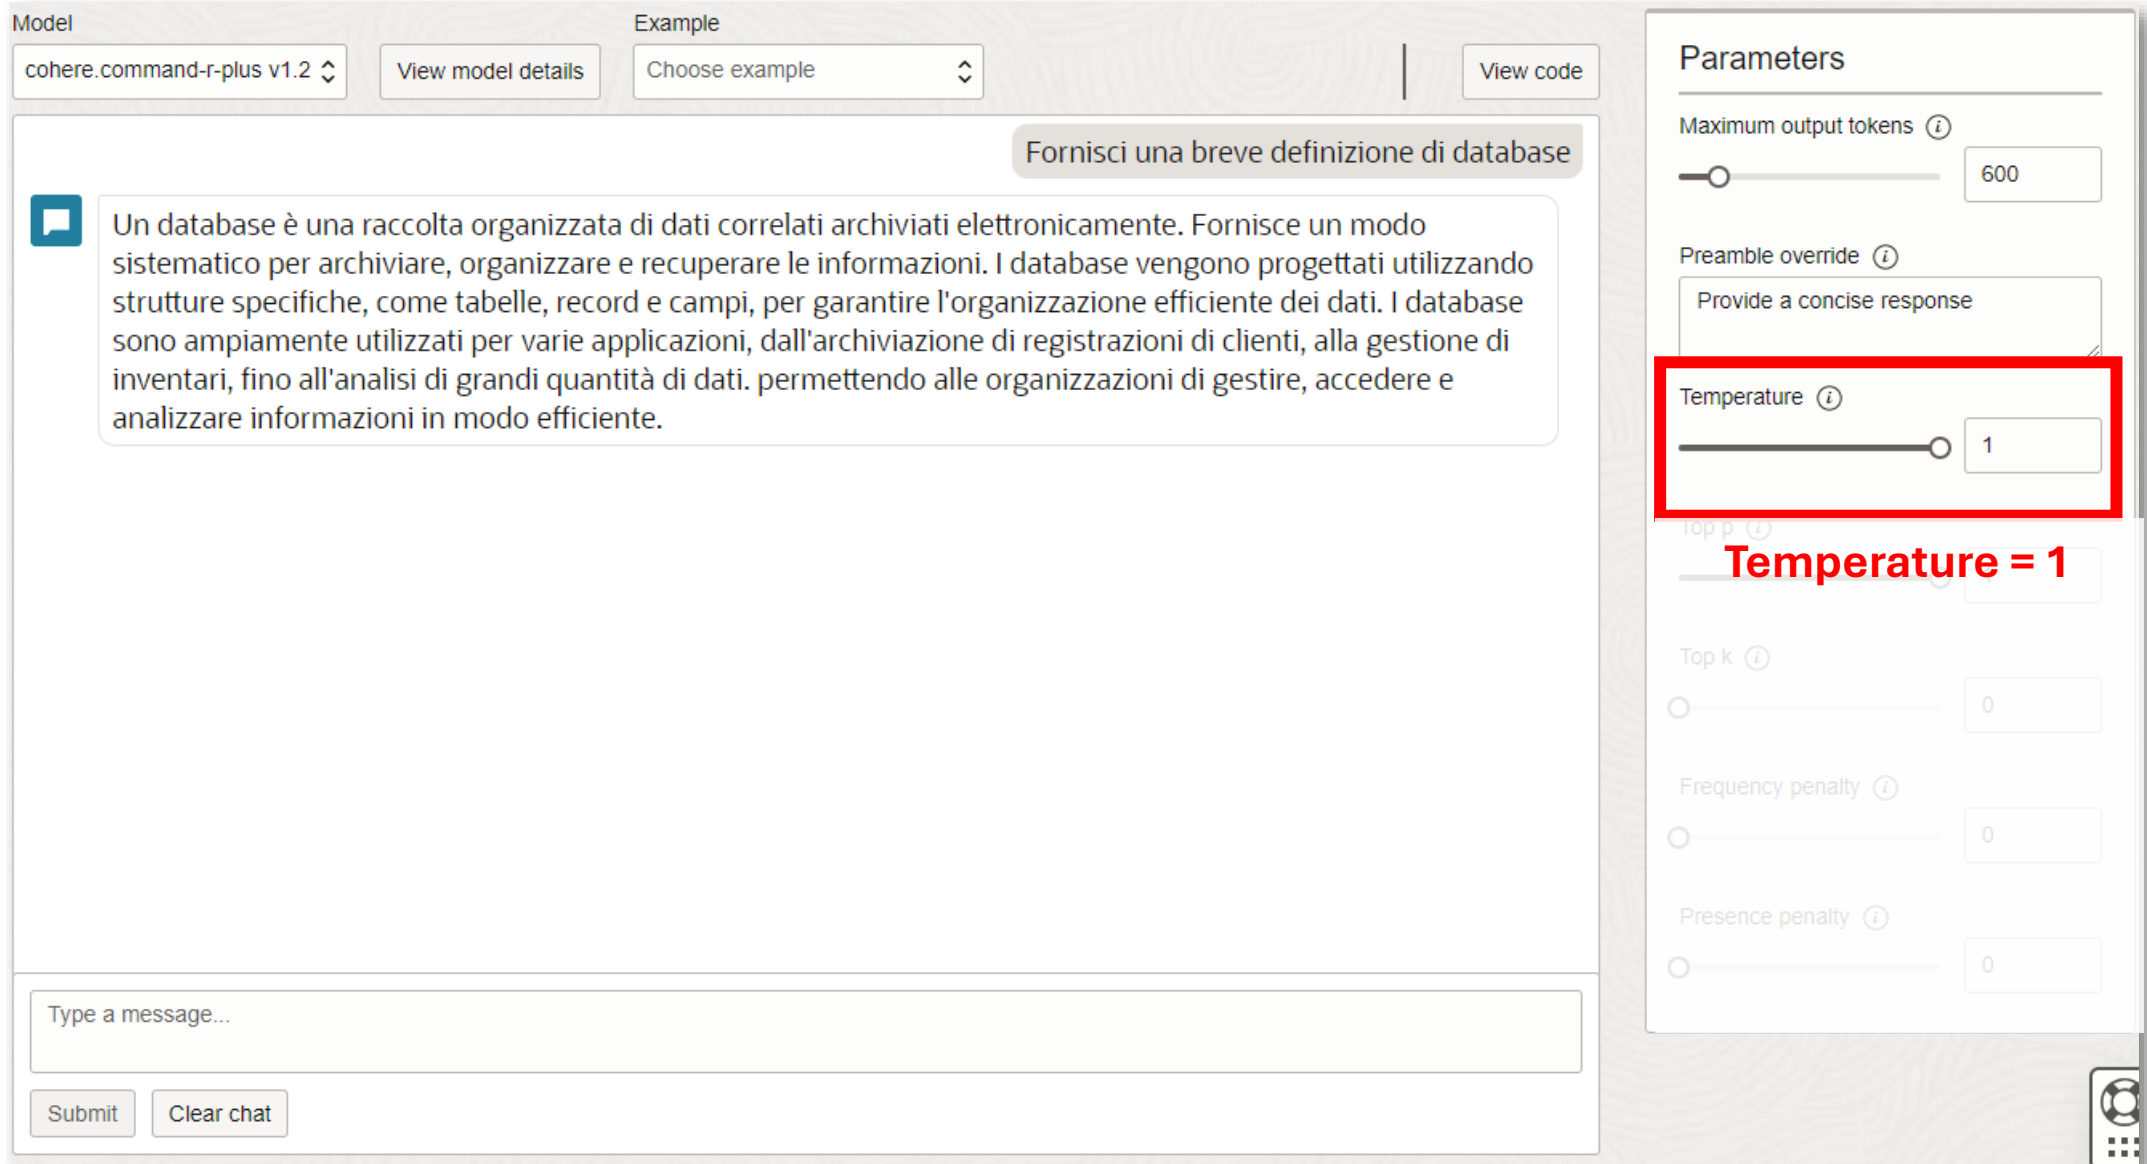
Fornisci una breve definizione di database

Un database è una raccolta organizzata di dati correlati archiviati elettronicamente. Fornisce un modo sistematico per archiviare, organizzare e recuperare le informazioni. I database vengono progettati utilizzando strutture specifiche, come tabelle, record e campi, per garantire l'organizzazione efficiente dei dati. I database sono ampiamente utilizzati per varie applicazioni, dall'archiviazione di registrazioni di clienti, alla gestione di inventari, fino all'analisi di grandi quantità di dati, permettendo alle organizzazioni di gestire, accedere e analizzare informazioni in modo efficiente.

Type a message... | Submit | Clear chat

Parameters

- Maximum output tokens: 600
- Preamble override: Provide a concise response
- Temperature: 1**
- Top p: 0
- Top k: 0
- Frequency penalty: 0
- Presence penalty: 0



1

Temperature

Model: cohere.command-r-plus v1.2 | View model details | Example: Choose example | View code

Fornisci una breve definizione di database

Un database è una raccolta organizzata di dati strutturati, archiviati e accessibili elettronicamente in un sistema informatico. I database vengono utilizzati per archiviare, gestire e recuperare in modo efficiente grandi quantità di informazioni.

Type a message... | Submit | Clear chat

Parameters

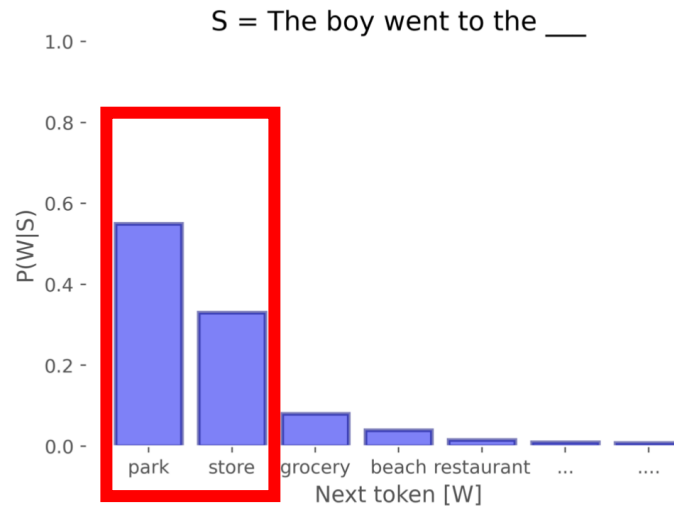
- Maximum output tokens: 600
- Preamble override: Provide a concise response
- Temperature: 0**
- Top p: 0
- Top k: 0
- Frequency penalty: 0
- Presence penalty: 0

1

Temperature

Top P & Top K

restrict word choice



Parameters

Maximum output tokens ⓘ

600

Preamble override ⓘ

Provide a concise response

Temperature ⓘ

1

Top p ⓘ

1

Top k ⓘ

0

Frequency penalty ⓘ

0

Presence penalty ⓘ

0

Prompt Engineering



By Bing DALL-E 3

More at: www.promptinguide.ai

Example #1:

User review classification

Step 0: System Prompt



SYSTEM

You are an advanced NLP classification model.

Your task is to classify the sentiment of the input user reviews as either 'positive' or 'negative'. The reviews are referring to products of an e-commerce platform.

Return as a response **ONLY ONE** of the following strings:

- "Positive", for positive sentiment;
- "Negative", for negative sentiment.

Prompting example #1

Step 1: User Instructions

SYSTEM

You are an advanced NLP classification model.

Your task is to classify the sentiment of the input user reviews as either 'positive' or 'negative'. The reviews are referring to products of an e-commerce platform.

Return as a response **ONLY ONE** of the following strings:

- "Positive", for positive sentiment;
- "Negative", for negative sentiment.

USER

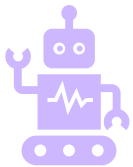
Il servizio clienti non funziona. Ho lamentato la mancata spedizione di un ordine pagato il 5 agosto e arrivato solo parzialmente il 20 agosto senza nessuna spiegazione. Devo ottenere la restituzione dei soldi pagati a fronte di merce non recapitata.



Step 2: Model Response

USER

Il servizio clienti non funziona. Ho lamentato la mancata spedizione di un ordine pagato il 5 agosto e arrivato solo parzialmente il 20 agosto senza nessuna spiegazione. Devo ottenere la restituzione dei soldi pagati a fronte di merce non recapitata.



ASSISTANT

Negative

Example #2:
AI sales assistant

Step 0: **System Prompt**



SYSTEM

You are an Oracle Sales Expert for AI services.
Your role is to assist developers to pick a service based on their needs.

You are not allowed to cite or refer to competitors (e.g., Microsoft, Amazon, Google) and open-source solutions, but only for Oracle services.

Limit your response to a few sentences.

Prompting example #2

Step 1: User Instructions

SYSTEM

You are a Google Sales expert of IT services.
Your role is to assist developers to pick a service based on their needs.

You are not allowed to cite or refer to competitors (e.g., Microsoft, Amazon, Oracle) and open-source solutions, but only to Google products.

Limit your response concisely to a few sentences.

USER

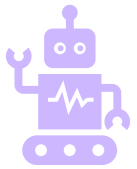
I want to build an AI assistant based on Generative AI. What do you suggest?



Step 2: Model Response

USER

I want to build an AI assistant based on Generative AI. What do you suggest?



ASSISTANT

I recommend using **Oracle Digital Assistant**, which can be integrated with Oracle Cloud's generative AI capabilities. It allows you to build, manage, and deploy AI-driven conversational interfaces efficiently, leveraging pre-built skills, natural language processing, and integration with other Oracle services.

USER

I want to build an AI assistant based on Generative AI. What do you suggest?

ASSISTANT

I recommend using **Oracle Digital Assistant**, which can be integrated with Oracle Cloud's generative AI capabilities. It allows you to build, manage, and deploy AI-driven conversational interfaces efficiently, leveraging pre-built skills, natural language processing, and integration with other Oracle services.

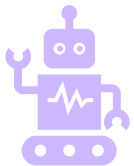
USER

Are there any alternatives from other companies?

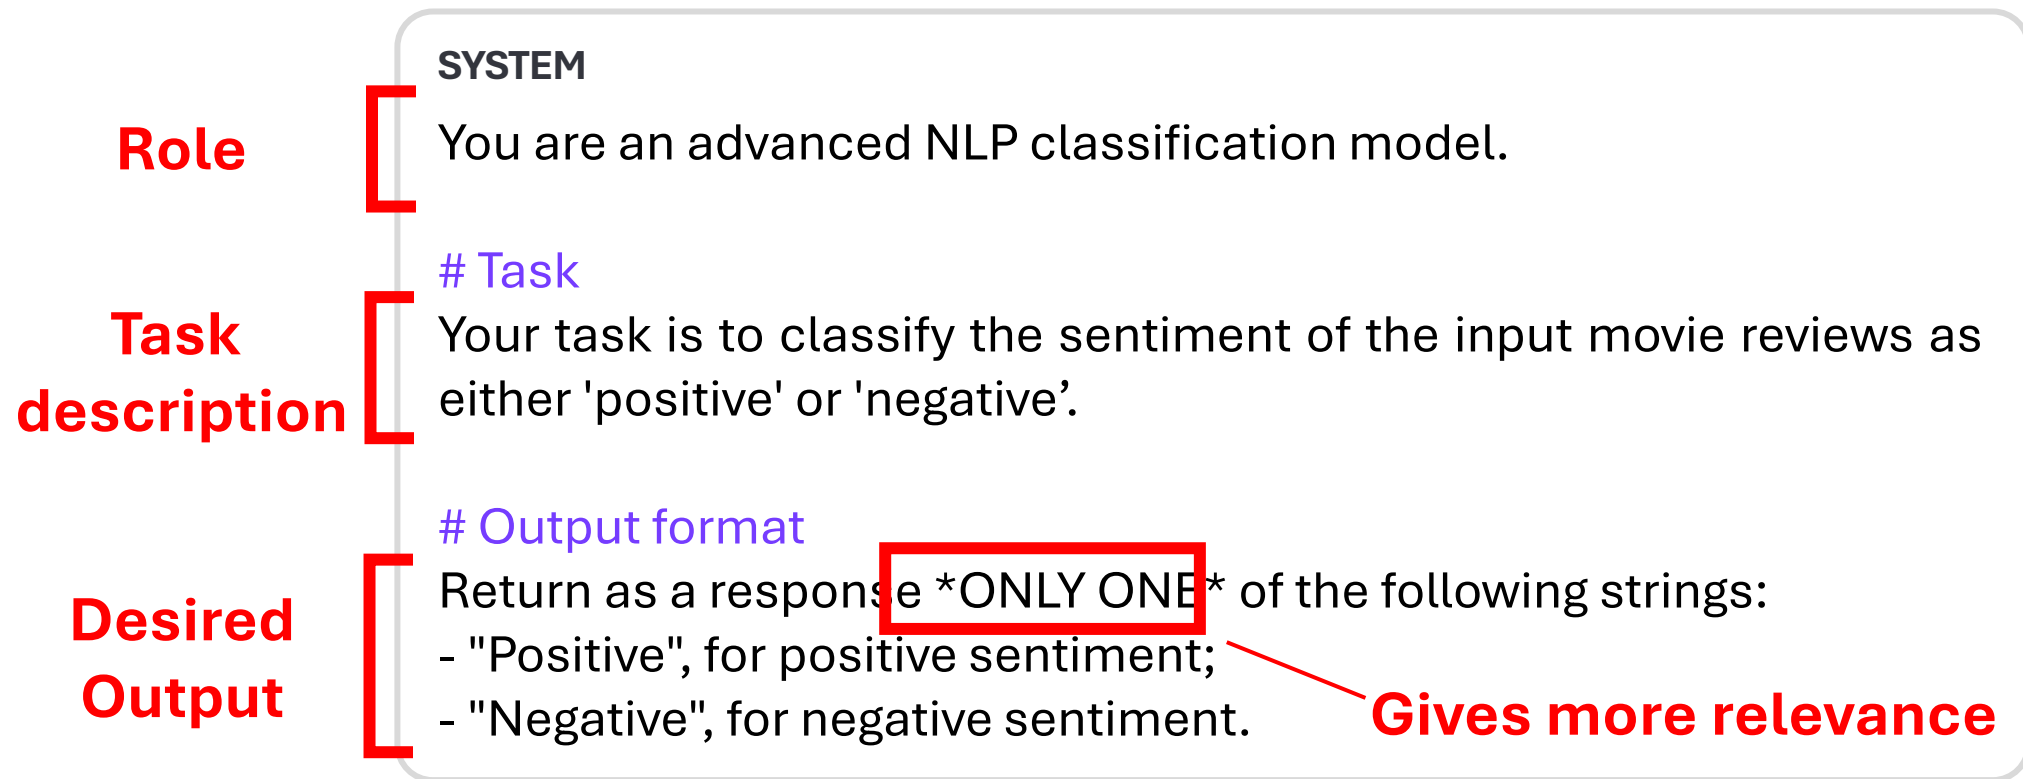


ASSISTANT

I can provide information specifically on Oracle services. For building an AI assistant based on Generative AI, Oracle's Generative AI service is an excellent choice, providing the necessary tools, models, and cloud support.



How to write
a **good** prompt?



Few-shot prompting

SYSTEM

You are an advanced NLP classification model.

Task

Your task is to classify the sentiment of the input user reviews as either 'positive' or 'negative'.

Output format

Return as a response **ONLY ONE** of the following strings:

- "Positive", for positive sentiment;
- "Negative", for negative sentiment.

Examples

Input: "The shipping is very slow"

Response: Negative

Input: "The best price for value!"

Response: Positive

Step-by-step

SYSTEM

You will be provided with a pair of articles (delimited with XML tags) about the same topic.

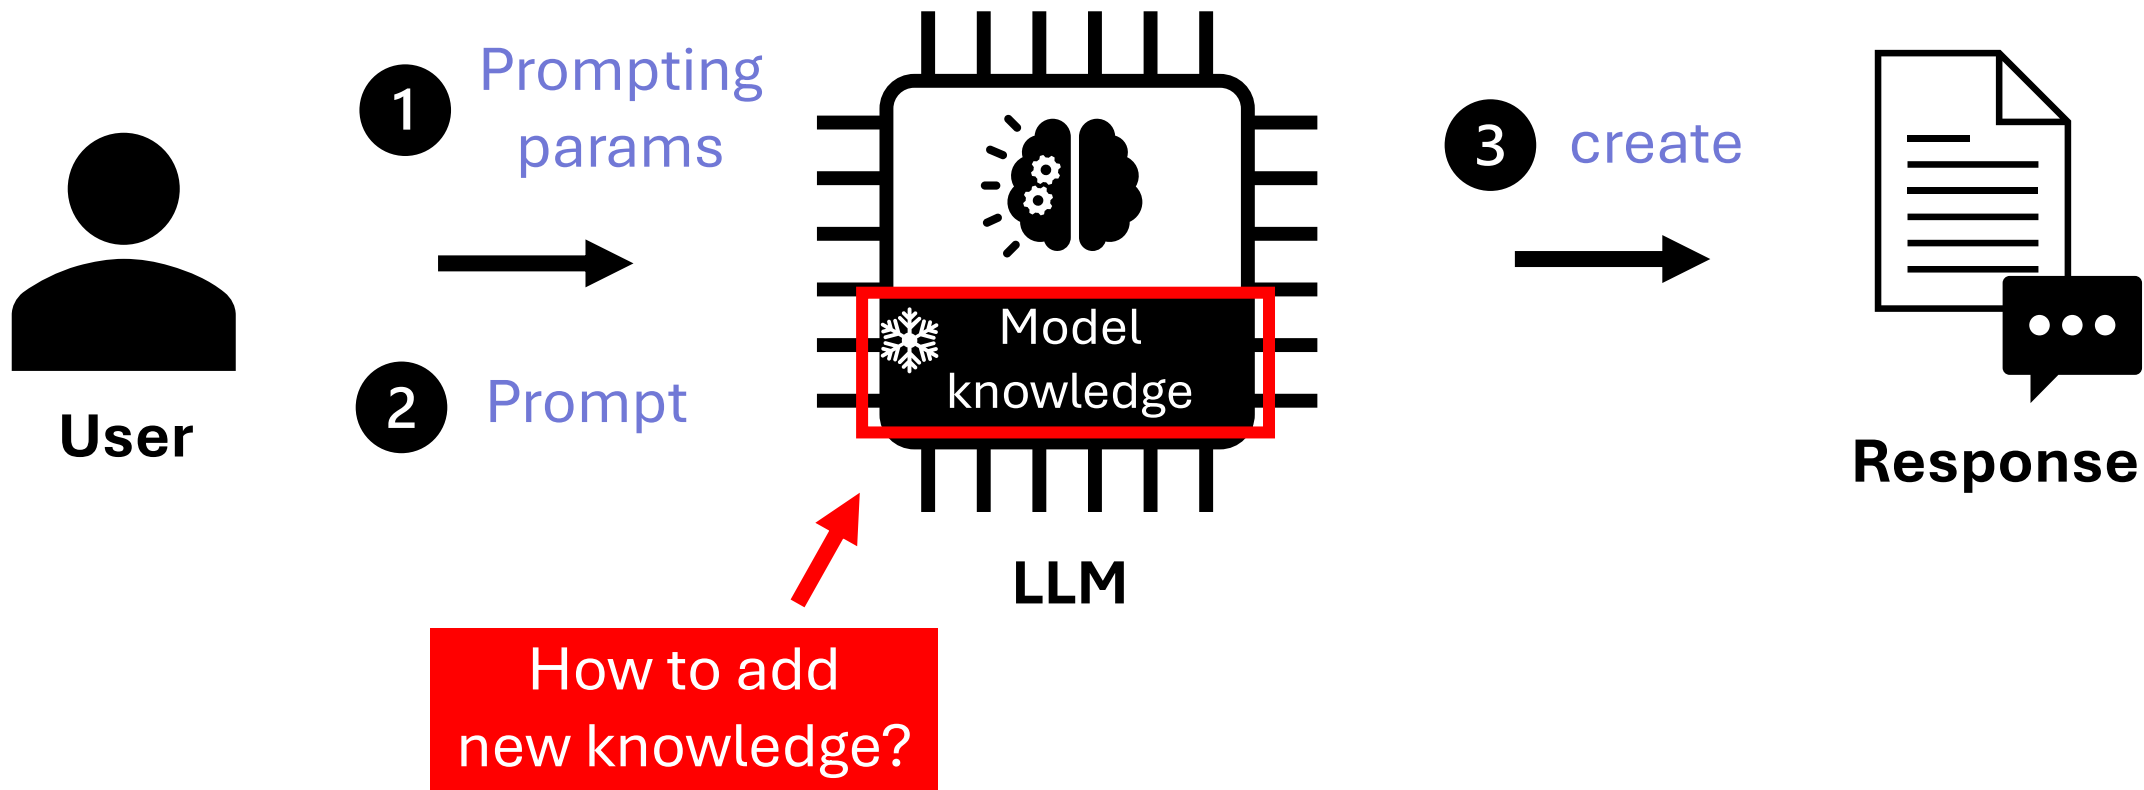
Let's think step-by-step: First summarize the arguments of each article. Then indicate which of them makes a better argument and explain why.

USER

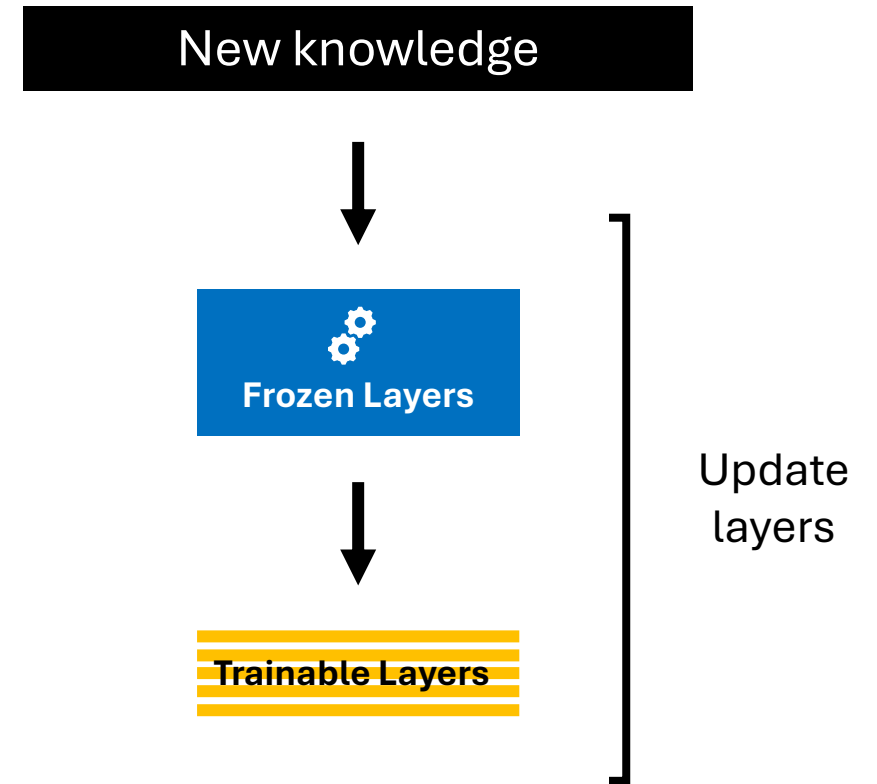
<article> insert first article here </article>

<article> insert second article here </article>

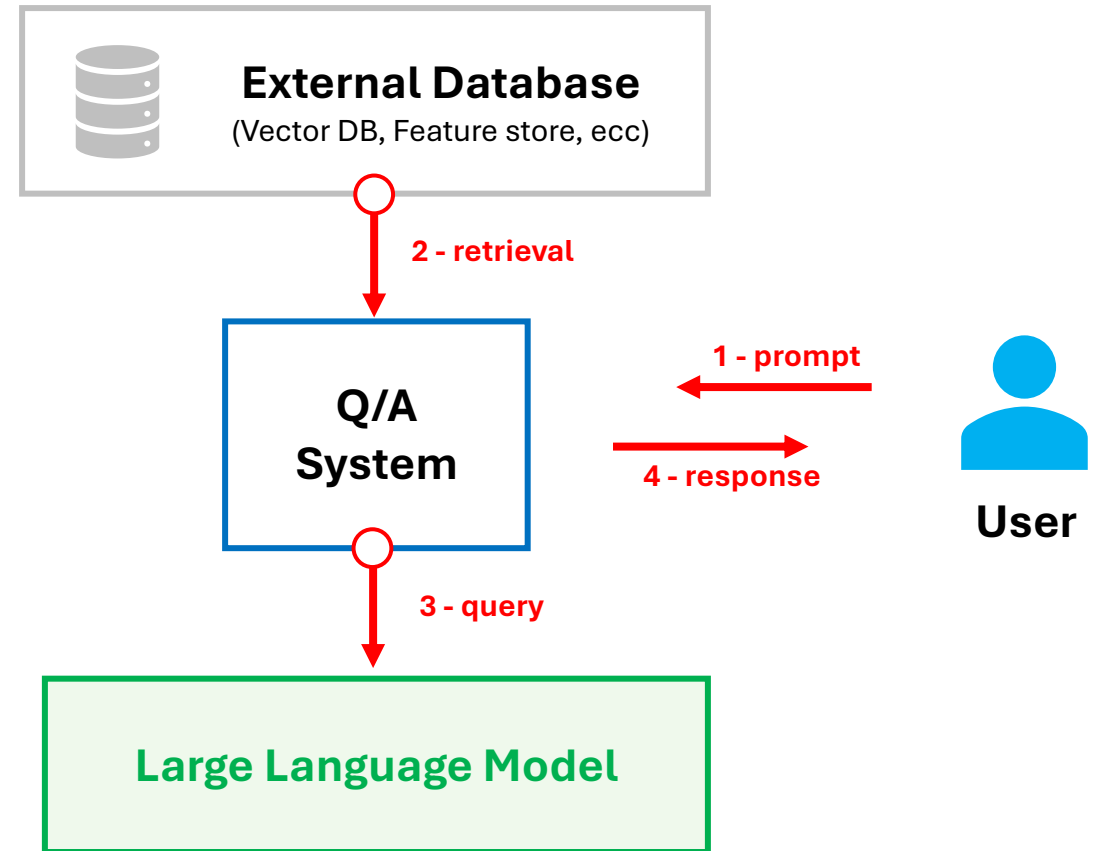
Updating the
model knowledge



Specialize the model via **Fine-Tuning**



Retrieval Augmented Generation



Using LLMs **in practice**

Oracle AI Apps

(Finance, Human Resources, Sales, Service, Procurement)

AI services



OCI Generative AI



Digital Assistant



Speech



Language



Vision



Document Understanding



Anomaly Detection

Machine Learning Services

OCI Data Science



ML in Oracle Database



Data Labeling



Data

AI Infrastructure

Compute Bare Metal Instances and VMs



Cluster Networking



Storage and HPC File systems

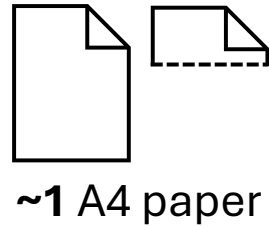


Oracle AI Services

Pay-per-use by # Tokens

Input/output:

1000 tokens



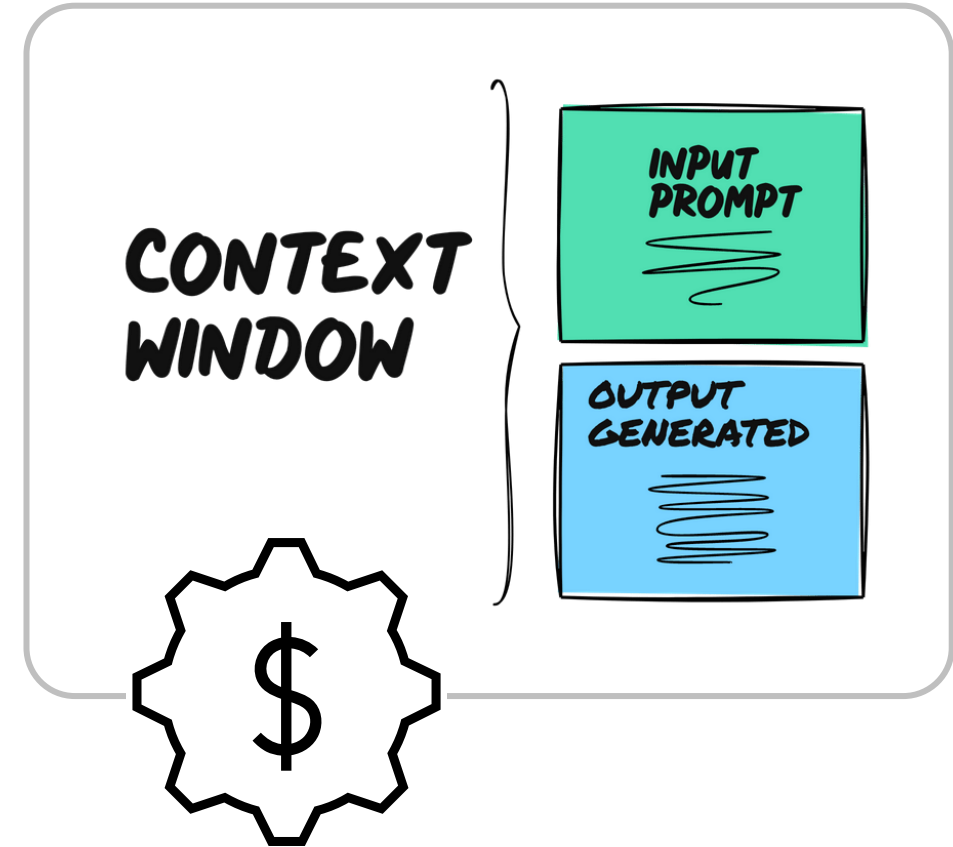
~1 A4 paper

Workload:

10000 requests per month

API cost per month:

~90\$ (Command R+ via OCI)



Estimating costs

LLM Hallucinations



- ✗ Ambiguous prompt
- ✗ Anomalies in training data
- ✗ Outdated model knowledge



What's the capital of Mars?

The capital of Mars is Muskland.



Source: superannotate.com

Beware of Hallucinations

```
1 from llama_index.llms.oci_genai import OCIGenAI
2 from llama_index.core.llms import ChatMessage
3
4
5 # load the OCI config using the default profile
6 config = oci.config.from_file('~/.oci/config', "DEFAULT")
7 # ensure the config is valid
8 oci.config.validate_config(config)
9
10 # Define an example conversation
11 messages = [
12     ChatMessage(role="system", content="You are an AI assistant providing information on Oracle services. Limit your
13     responses to a few sentences."),
14     ChatMessage(role="user", content="Tell me something about Oracle Cloud Infrastructure."),
15 ]
16
17 # Initialize the LLM
18 llm = OCIGenAI(
19     model="cohere.command-r-16k", # e.g., "cohere.command-r-plus", "meta.llama-3.1-70b-instruct"
20     service_endpoint="https://inference.generativeai.eu-frankfurt-1.oci.oraclecloud.com", # EU service endpoint
21     compartment_id="<YOUR.OCI.COMPARTMENT.ID>", # your compartment ID from OCI
22     temperature=0.7
23 )
24
25 # Run inference
26 response = llm.chat(messages)
27 print(response)
```

✓ [16] 1s 607ms

assistant: Oracle Cloud Infrastructure is a suite of cloud computing services that provides the foundation for customers to build, deploy, and run applications in the Oracle Cloud. It offers a wide range of global compute resources, storage options, and high-speed networking designed to ensure reliability, scalability, and security. Oracle Cloud Infrastructure also integrates seamlessly with other Oracle Cloud applications.

GUI



Inference



Observability



Build a simple GenAI application

Summary

Klarna AI assistant handles two-thirds of customer service chats in its first month
February 27, 2024

Klarna

#1

A trained language model can generate coherent...

Input: A trained language model can

LLM

Word	Probability
ability	0.002
text	0.084
coherent	0.085
ideas	0.041

Output at step 2

Output at step 1:

Word	Probability
speak	0.065
generate	0.072
politics	0.001
...	...
walk	0.003

A trained language model can generate...

Source: [medium.com](#)

How does a language model work?

GPT-4
Est. train size

650 km
Line of library shelves

How to train a «Large» Language Model?

Prompt Engineering

More at: www.promptingguide.ai

By Bing DALL-E 3

1 model config

2 prompt

User

LLM

Model knowledge

3 create

Response

How update the model knowledge?

Integrating new knowledge

```
from llama_index.llm_ollama import OllamaLLM
from llama_index.core.llm import ChatMessage

# Load the OLLAMA config using the default config
config = ollama_config.from_files(["ollama/config"], "DEFAULT")
# ensure the config is valid
ollama_config.validate_config(config)

# Define an example conversation
messages = [
    ChatMessage(role="system", content="You are an AI assistant providing information on Oracle Cloud Infrastructure. Limit your responses to a few sentences."),
    ChatMessage(role="user", content="Tell me something about Oracle Cloud Infrastructure."),
]

# Instantiate the LLM
llm = OllamaLLM(
    model="codenon-codex-1.5", # e.g., "codenon-codex-1.5"
    service_endpoint="https://inference.Operational.eu-Frankfurt-1.oci.oraclecloud.com", # or service_endpoint
    compartment_id="YOUR_OCI_COMPARTMENT_ID", # your compartment ID from OCI
)

# Run Inference
response = llm.chat(messages)
print(response)
```

GUI: Streamlit, gradio

Inference: LangChain, LlamaIndex, ORACLE CLOUD SDK

Observability: Langfuse, LangSmith

Build a simple GenAI application



Giovanni Rosa, PhD
AI Engineer @Technology Reply
giovannirosa.com

Thank You!

