Enhancing AI Search with Machine Learning

Suman Debnath

Developer Advocate, Data and Machine Learning
Amazon Web Services

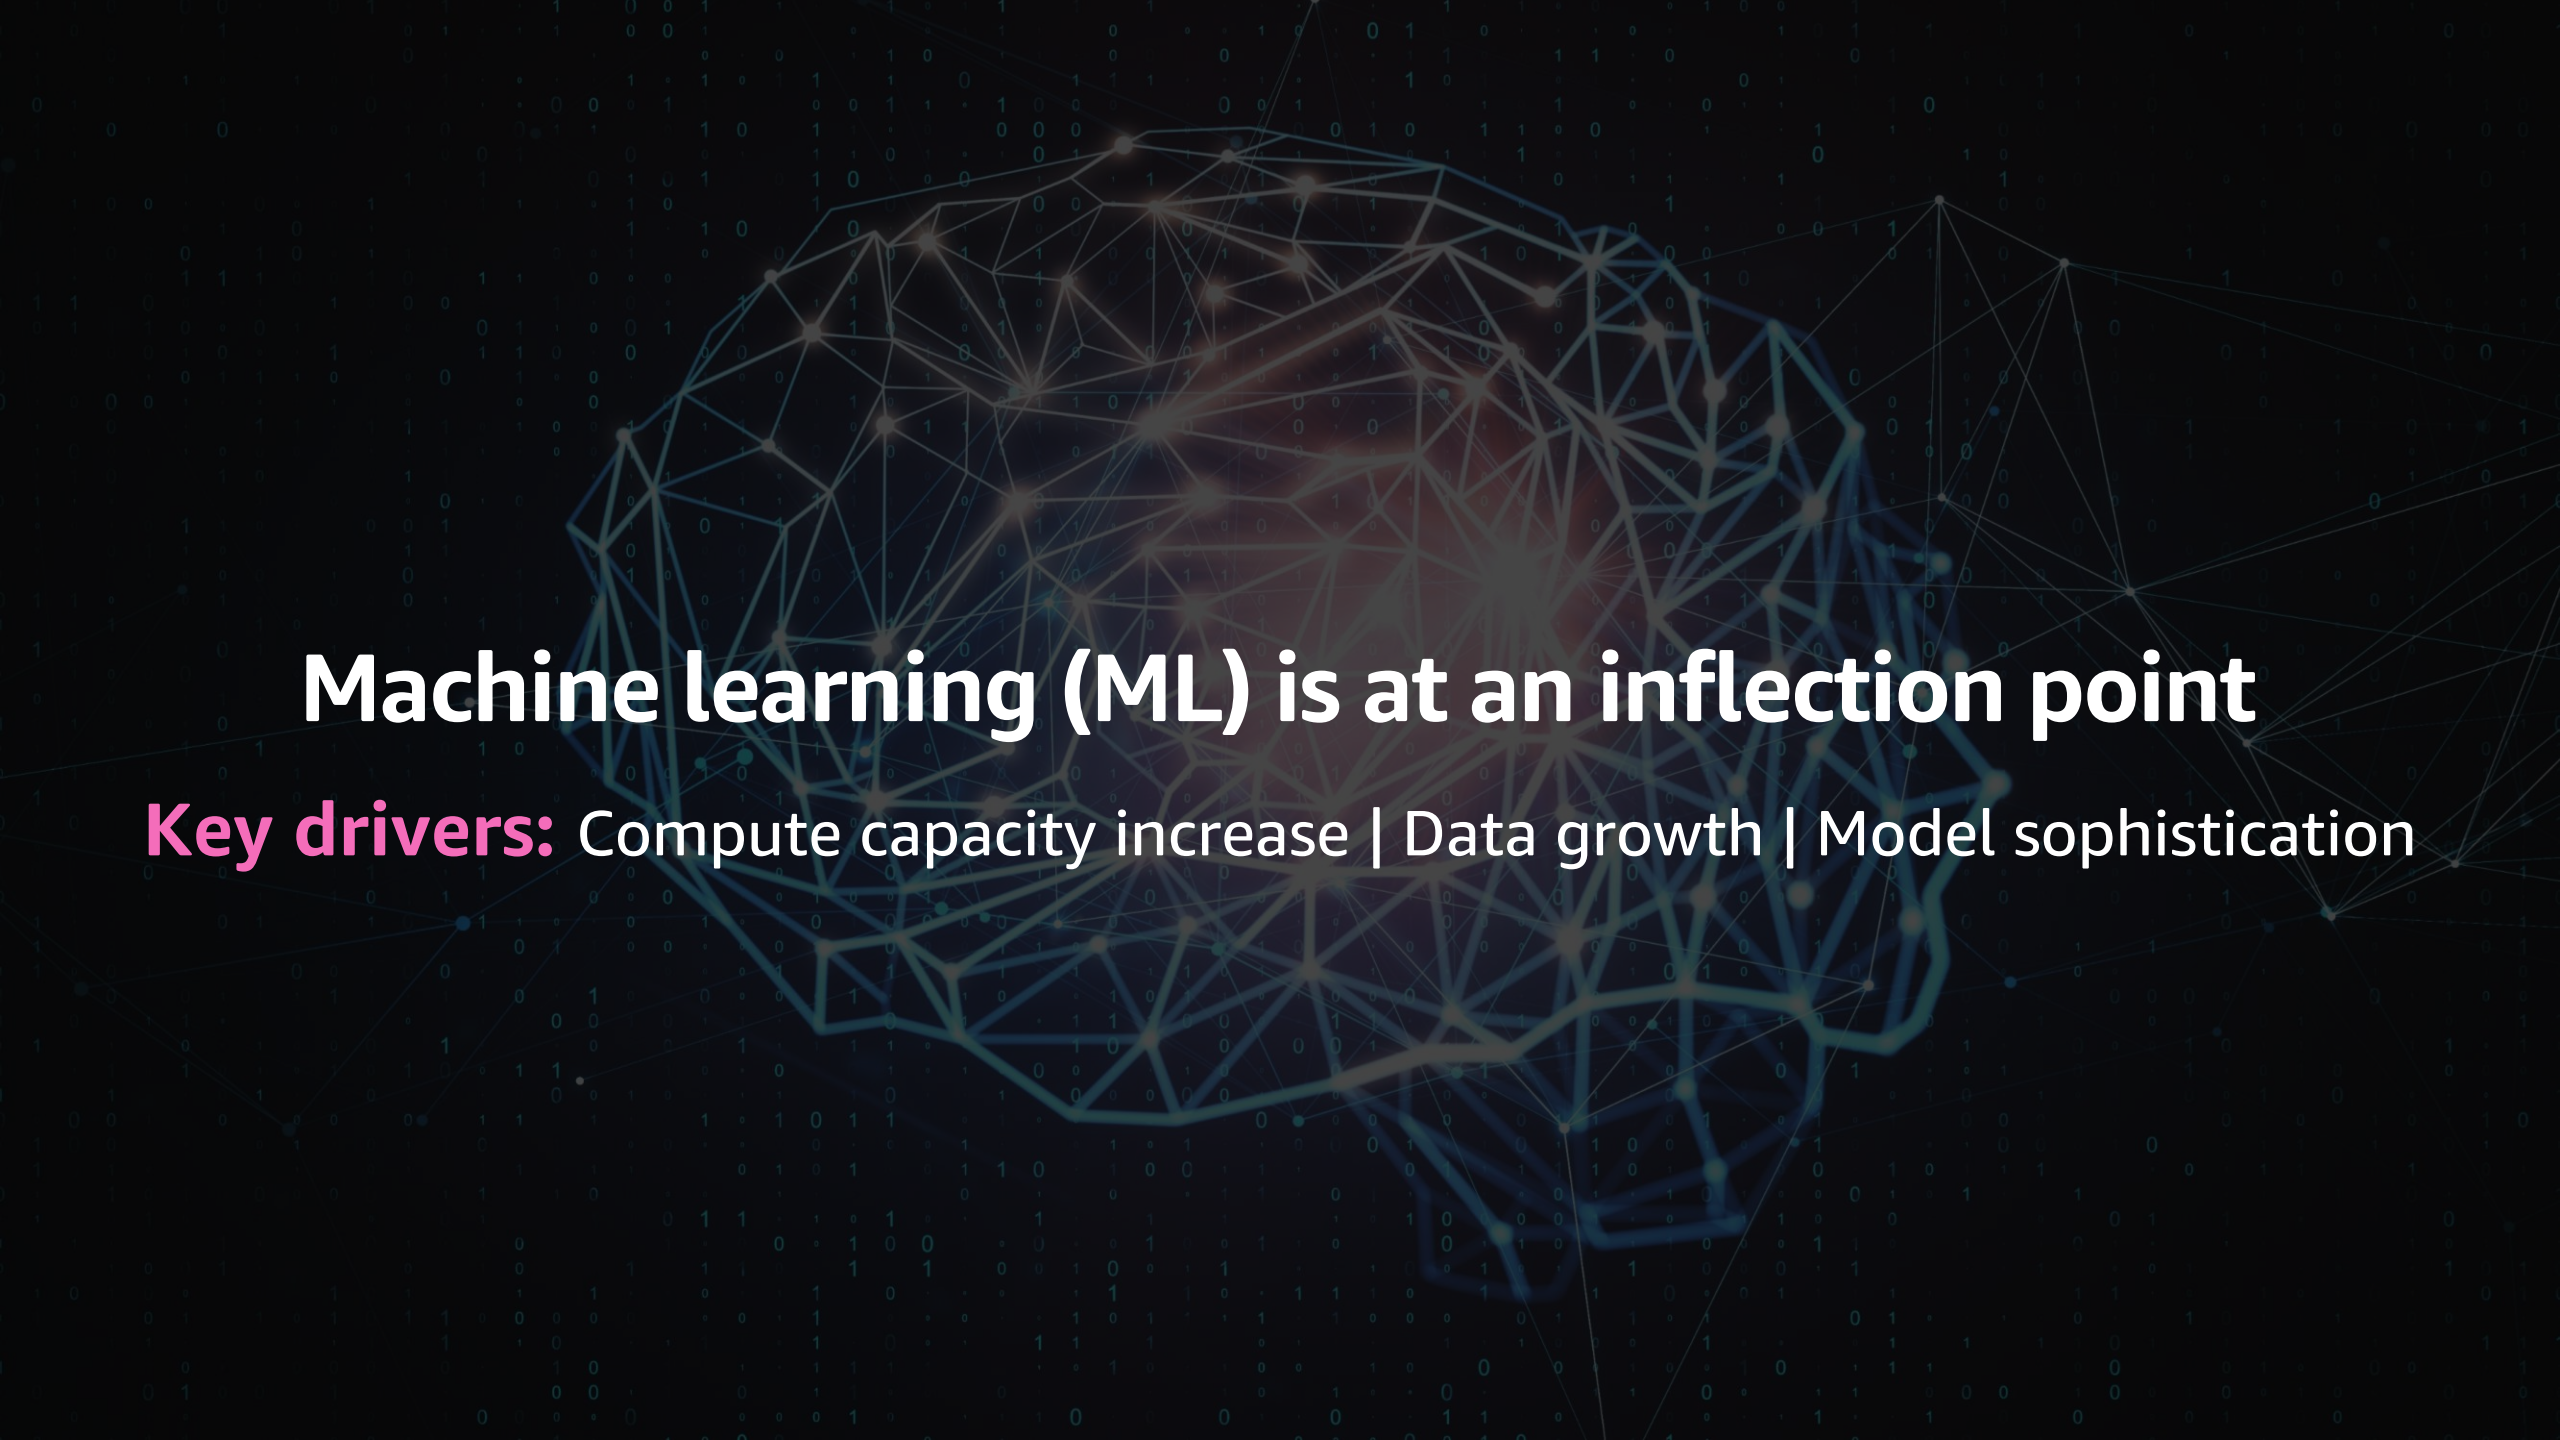

Machine learning (ML) is at an inflection point

Key drivers: Compute capacity increase | Data growth | Model sophistication

Question: What is generative artificial intelligence (AI)?

- Creates new content and ideas, including conversations, stories, images, videos, and music
- Powered by large models that are pre-trained on vast corpora of data and commonly referred to as foundation models (FMs)

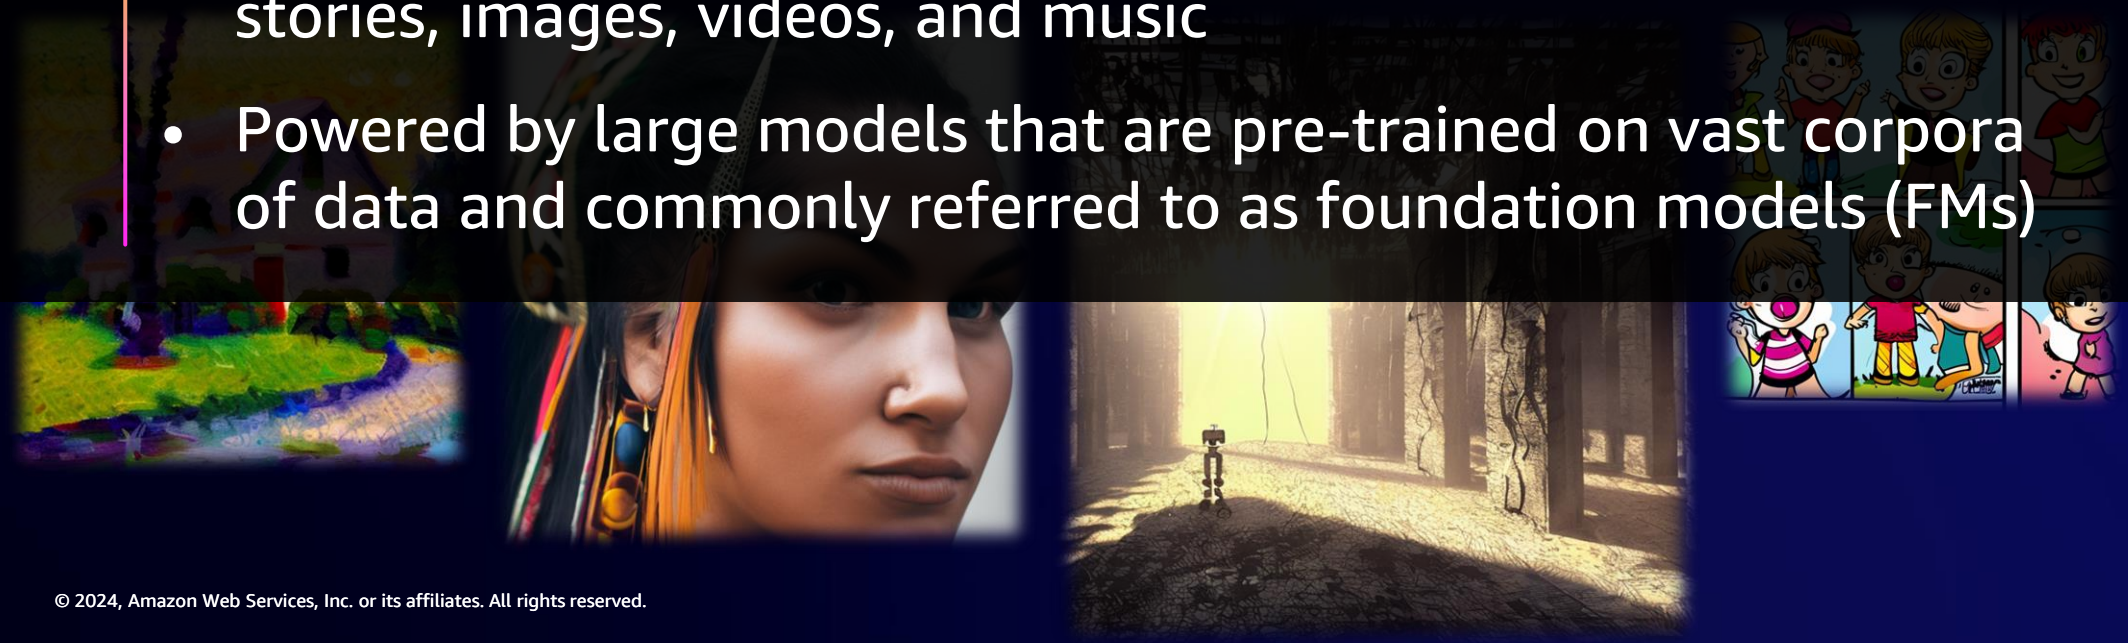

How foundation models differ from other ML models ?

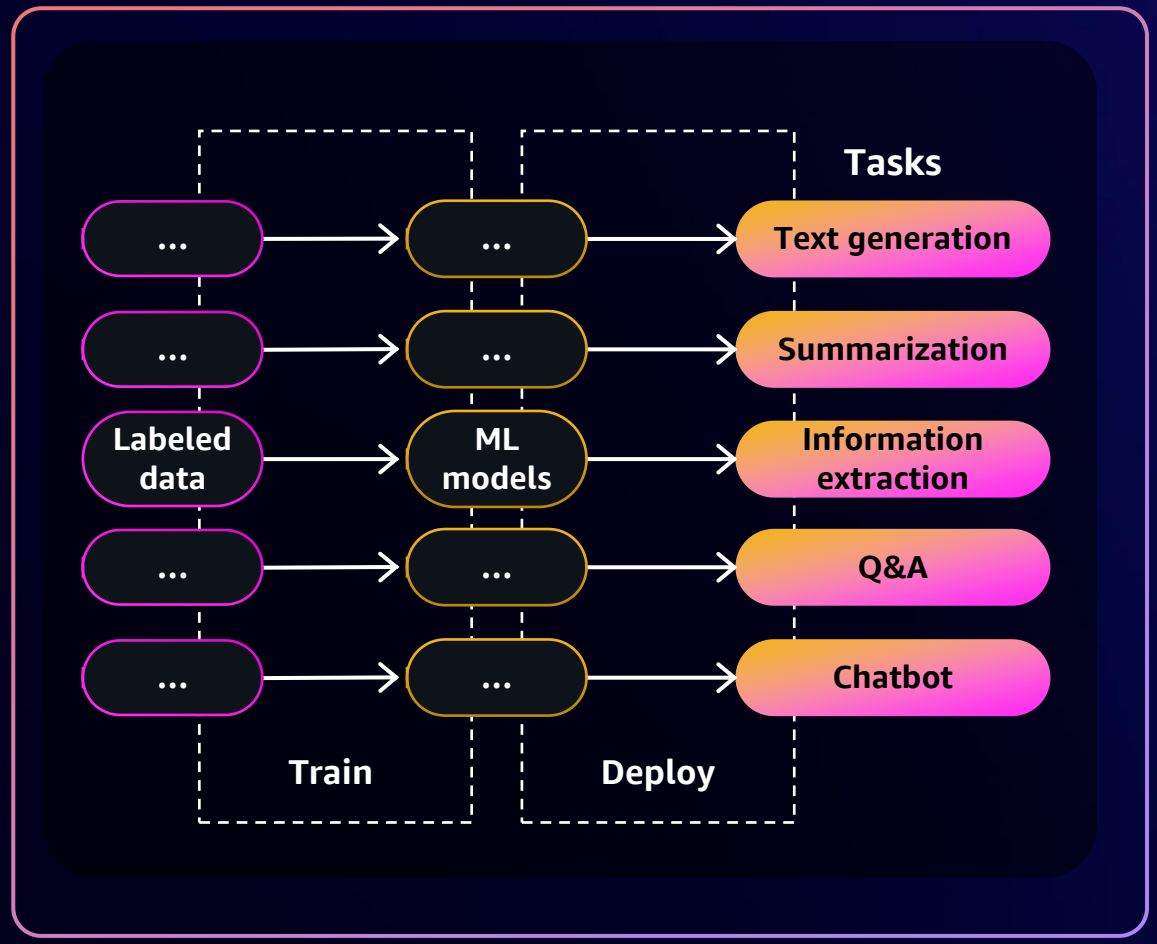

Traditional ML models

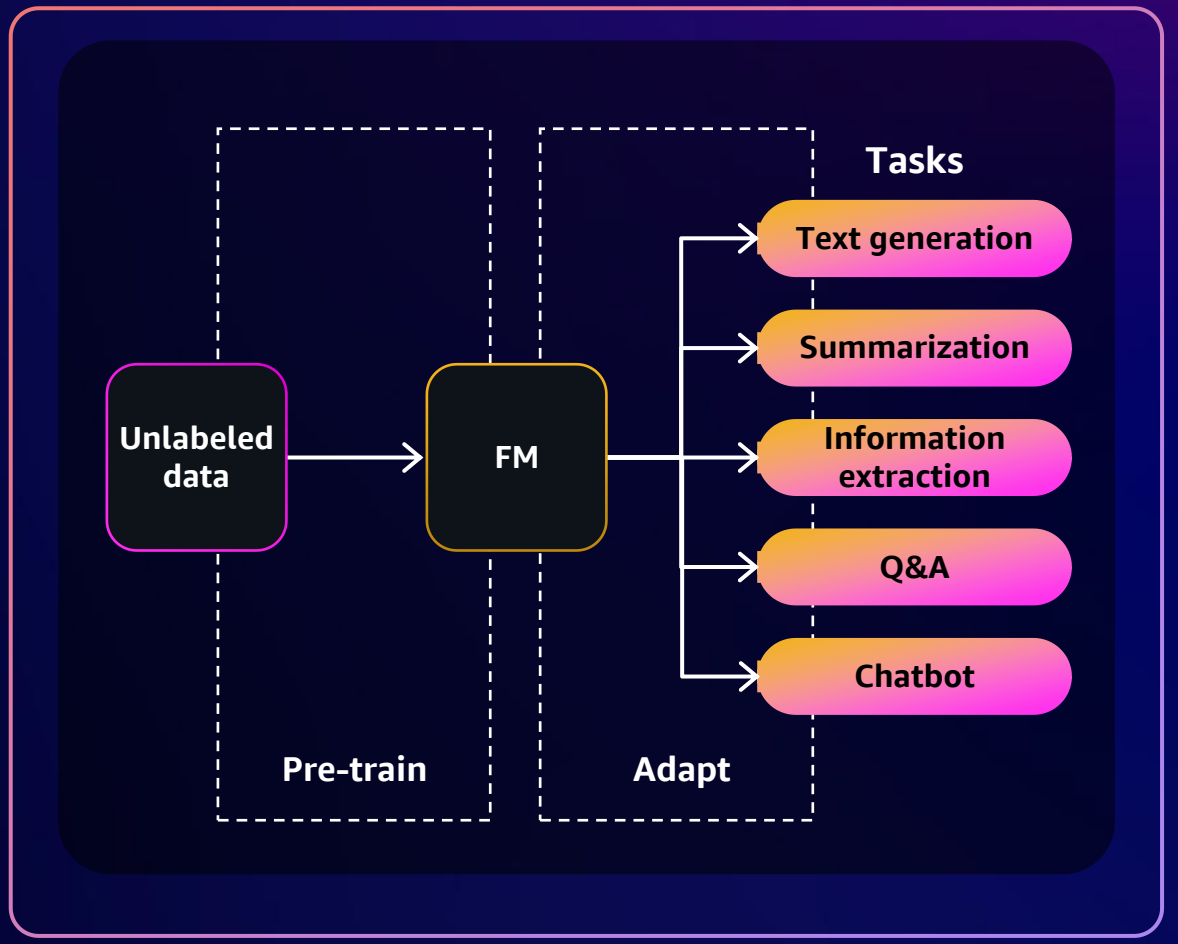

Foundation models

What are inputs & outputs of foundation models ?

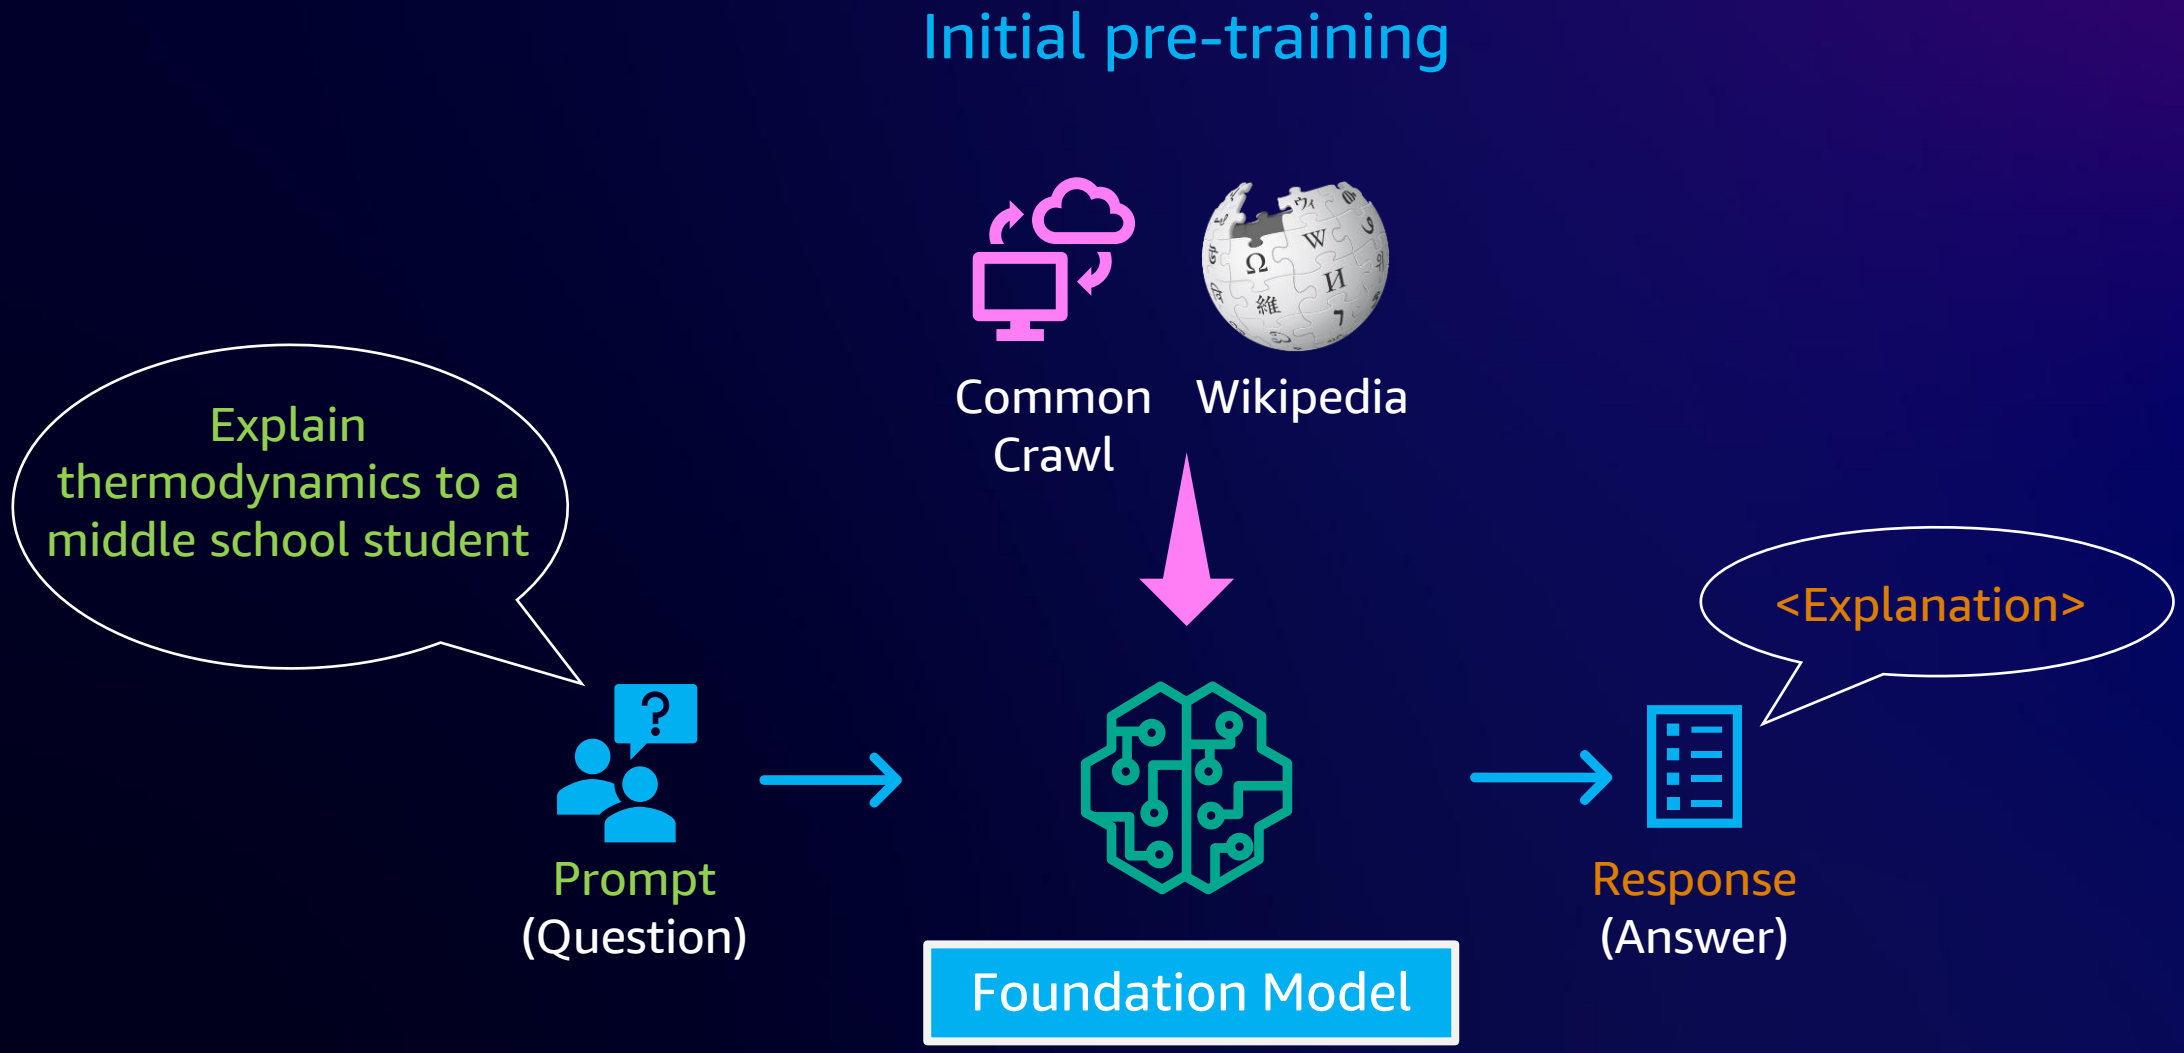

Let’s say we want to know…

- Which are the products that got the best reviews on XYZ platform in last 15 days ?

Hallucination

- Who won the India vs Afghanistan 2024 T20 championship ?

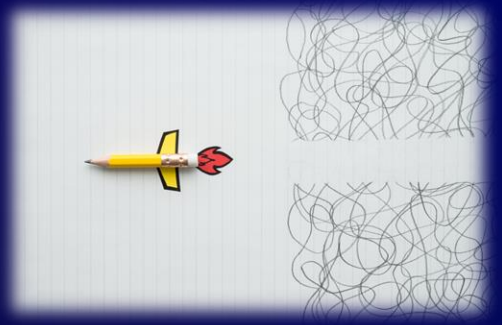

Knowledge Cutoff

Large Language Model Limitations


CUSTOMIZE
Macro
Shift

How can we customize a foundation model ?

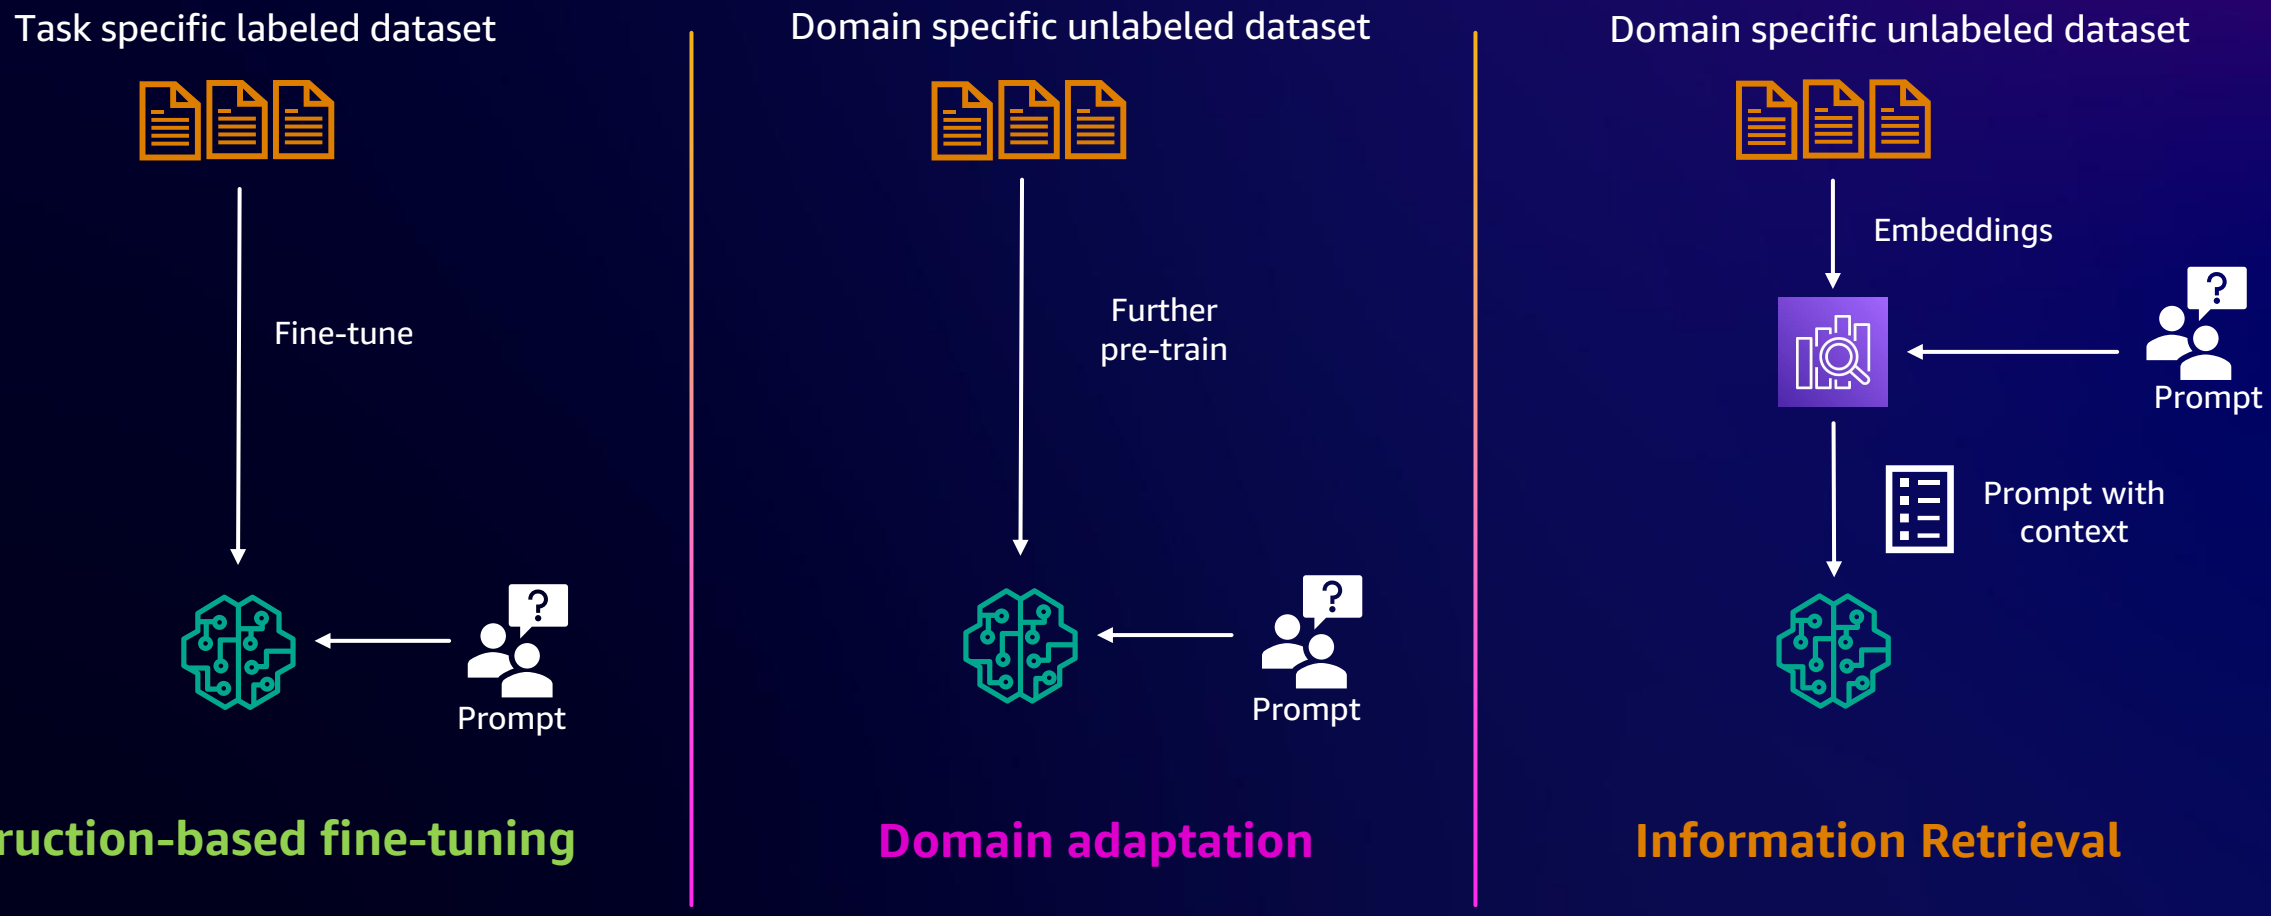

Vector store

RAW Data

Documents

Images

Audio

Machine Learning Model (Embedding)

Amazon Bedrock

Dense Vector Encodings

0.3, 2.1, 0, 0.9, 1.0,...

1.3, 8.1, 0, 4.6, 3.6,...

7.3, 1.1, 0, 2.9, 1.0,...

Vector Embedding Space

Vector Database

Dev ready and Operationalized

Vector store

Dev ready and Operationalized

Query :

+

Context :

LangChain

Vector Database

LangChain

Response

Retrieval-Augmented Generation (RAG)

Visual Search

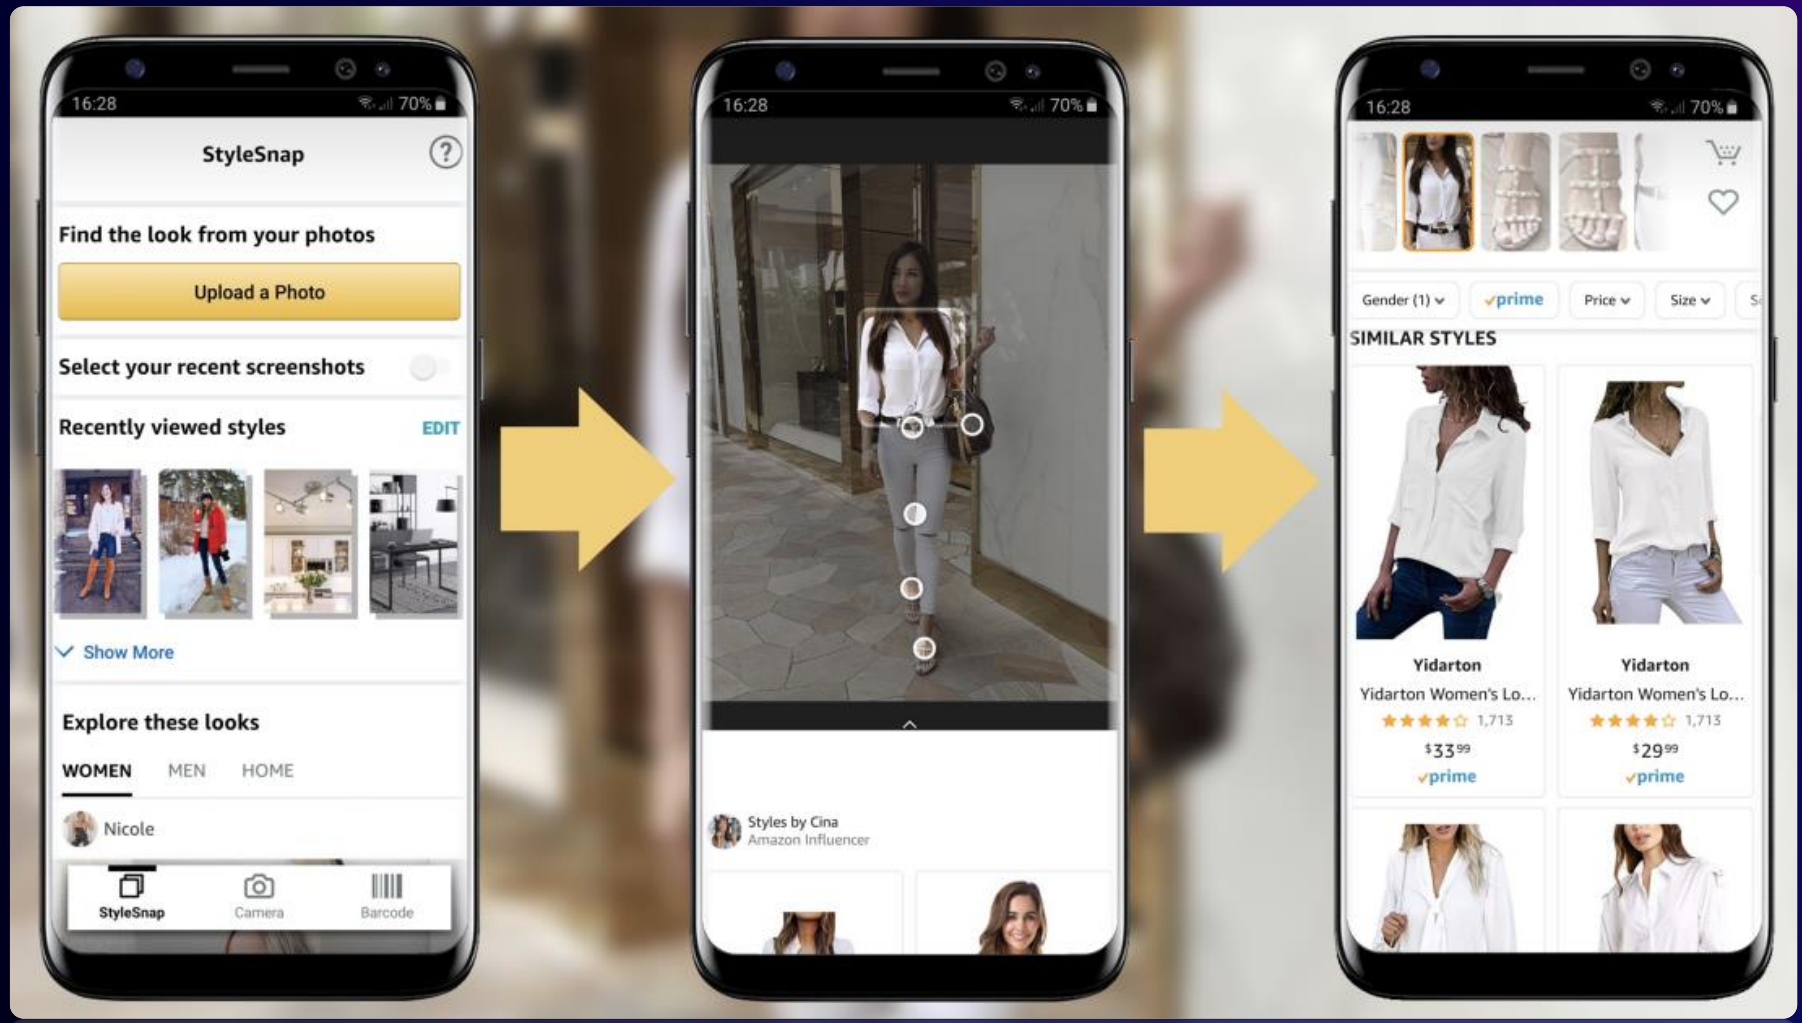

Let’s build it : Resume Screening

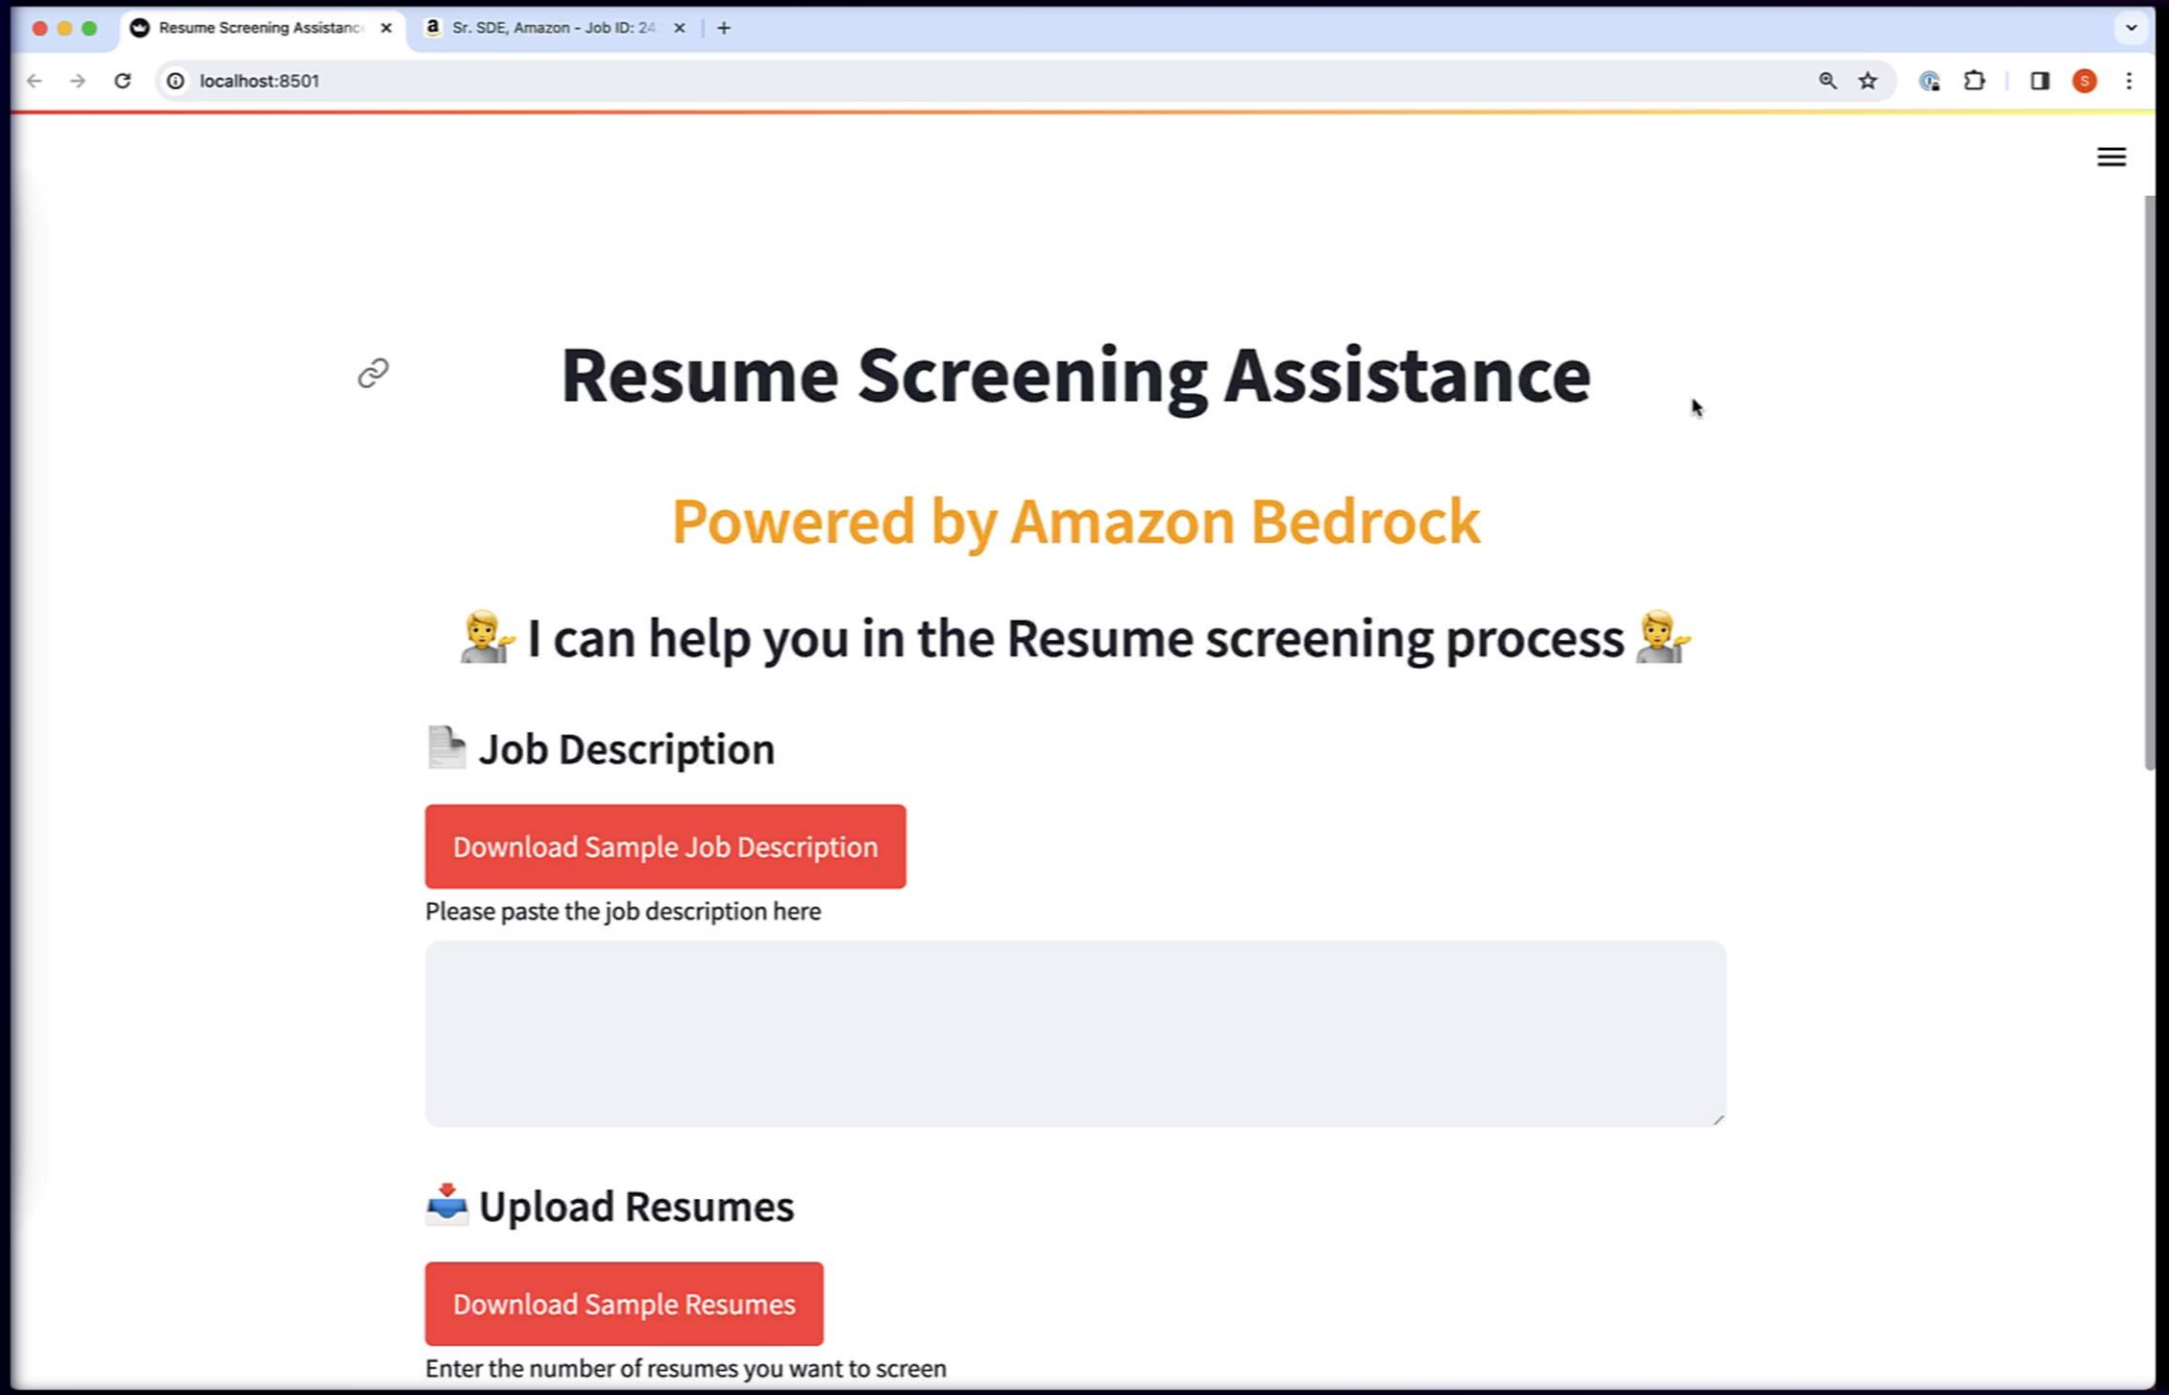
Resume Screening Assistance
Sr. SDE, Amazon - Job ID: 24
localhost:8501
Resume Screening Assistance
Powered by Amazon Bedrock
I can help you in the Resume screening process
Job Description
Download Sample Job Description
Please paste the job description here
Upload Resumes
Download Sample Resumes
Enter the number of resumes you want to screen

Let's build it : Resume Screening

Input

Job Description

No. of resumes you want ?

```
relavant_docs = vectorstore.similarity_search_with_score(job_description,
                                                         resume_count)
```

Amazon Bedrock

```
from langchain.llms.bedrock import Bedrock
from langchain.chains.summarize import load_summarize_chain

llm = Bedrock()
chain = load_summarize_chain(llm,..)

summary = chain.run(relavant_docs)
```

Bunch of Resumes

Amazon Bedrock
(Embedding)

0.3, 2.1, 0, 0.9, 1.0,...

1.3, 8.1, 0, 4.6, 3.6,...

7.3, 1.1, 0, 2.9, 1.0,...

Dense Vector Encodings

Vector Database

```
from langchain.vectorstores.pgvector import PGVector
vectorstore = PGVector.from_documents(...)
```

https://bit.ly/3SmHjbk

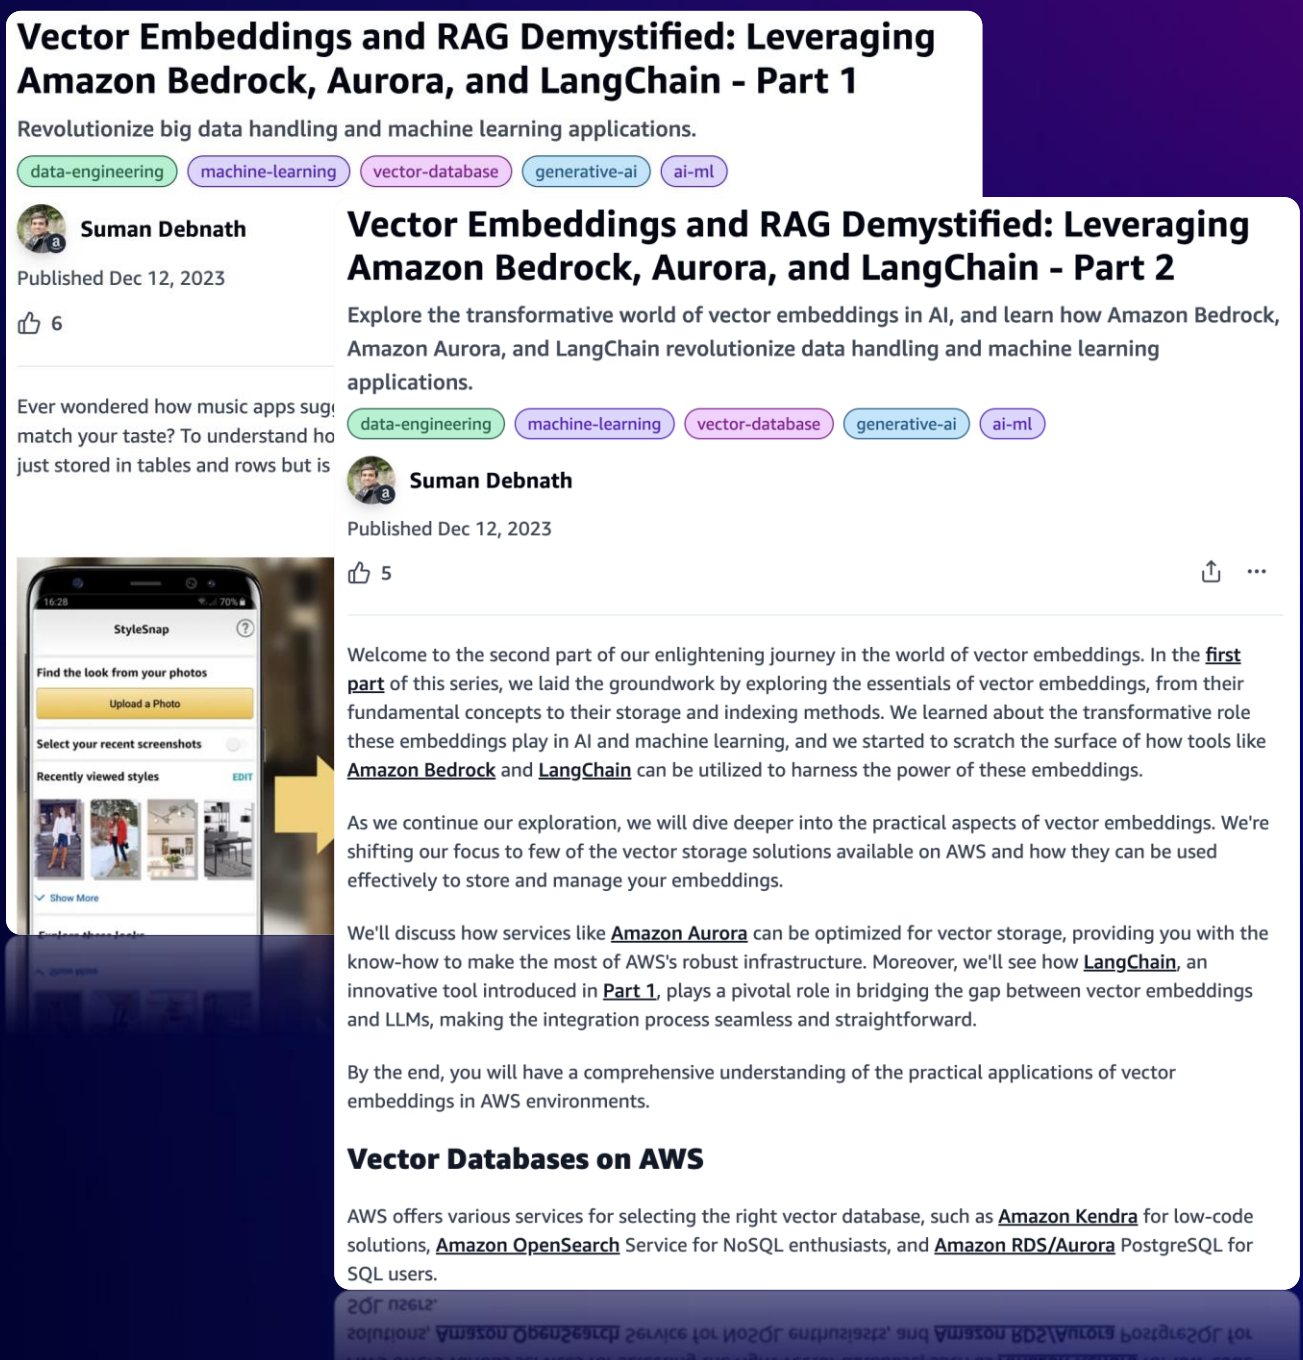

https://bit.ly/3Q3amy0

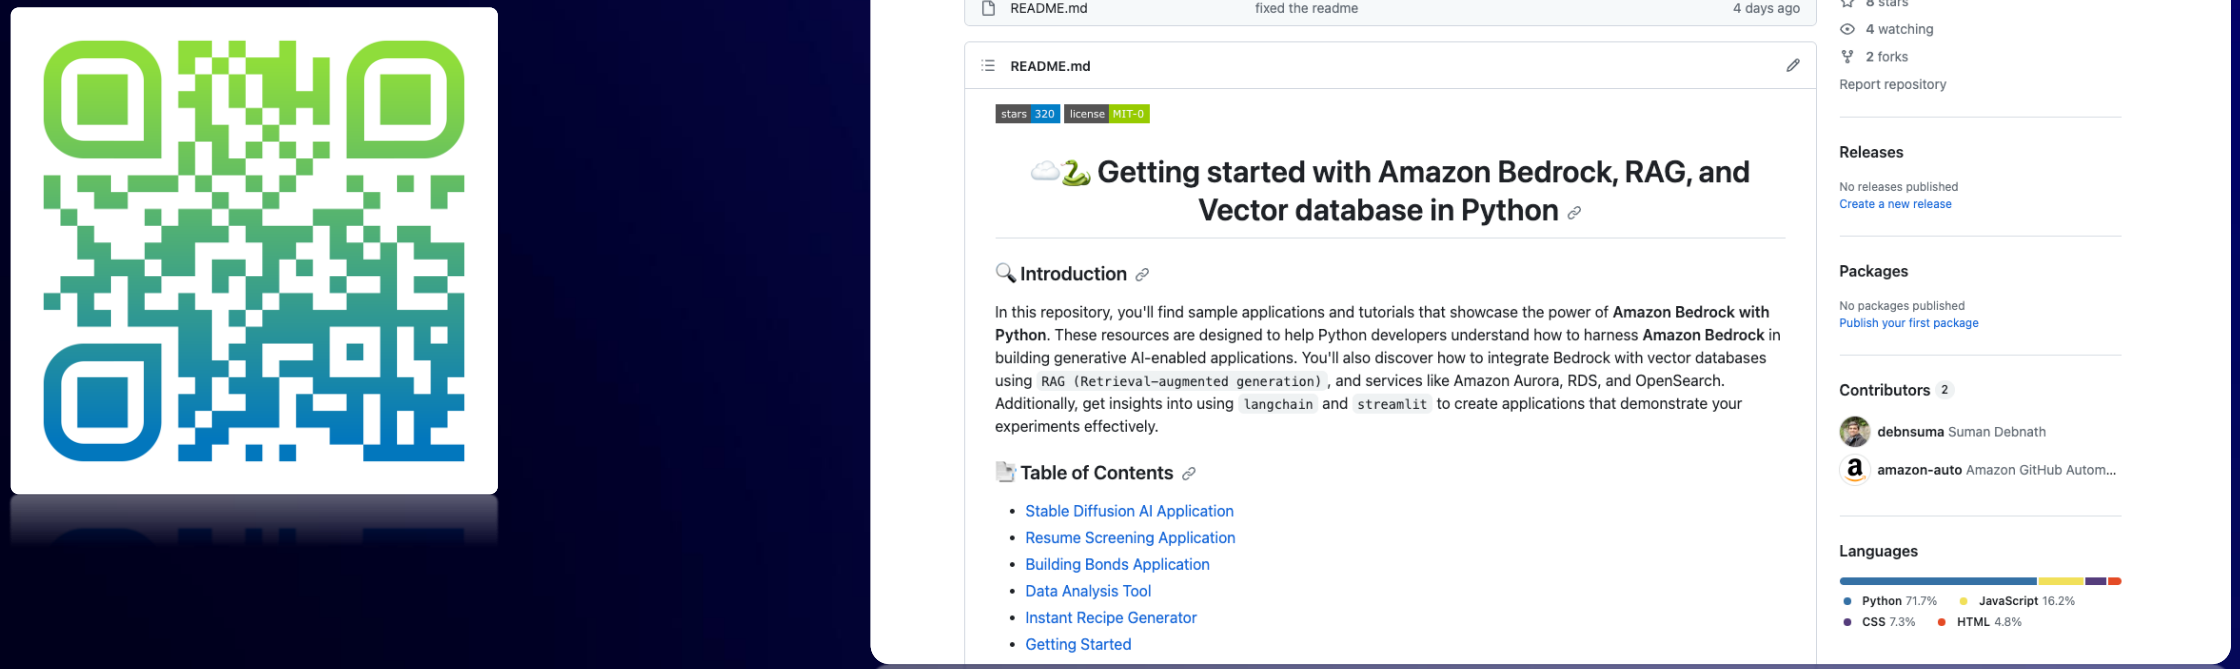

GENERATIVE AI STACK

APPLICATIONS THAT LEVERAGE FMs

- Amazon Q
- Amazon Q in Amazon QuickSight
- Amazon Q in Amazon Connect
- Amazon CodeWhisperer

TOOLS TO BUILD WITH LLMs & OTHER FMs

Amazon Bedrock

Guardrails | Agents | Customization Capabilities

SageMaker

INFRASTRUCTURE FOR FM TRAINING & INFERENCE

- GPUs
- Trainium
- Inferentia
- SageMaker
- UltraClusters
- EFA
- EC2 Capacity Blocks
- Nitro
- Neuron

Suman Debnath

linkedin.com/in/suman-d

debnsuma@amazon.com