Large Language Models with Financial Analysis

MAFS5440 2024Fall

Large language models use transformer models and are trained using massive datasets — hence, large. This enables them to recognize, translate, predict, or generate text or other content.

What LLM can do for financial analysis?

1. Interpreting Data
2. Modeling Predict
3. Automated Reporting
4. Market Analysis
5. Risk Assessment

Financial Statement Analysis with Large Language Models

Alex G. Kim[1]   Maximilian Muhn[2]   Valeri V. Nikolaev[3]

This draft: May 20, 2024

GPT-4 can outperform professional financial analysts in predicting future earnings changes, generating useful narrative insights, and resulting in superior trading strategies with higher Sharpe ratios and alphas, thereby suggesting a potential central role for LLMs in financial decision-making.

https://arxiv.org/pdf/2407.17866

BloombergGPT
- trained an LLM using a mixture of finance data and general-purpose data, which took about 53 days, at a cost of around **$3M**)
- Improving existing financial NLP tasks, such as sentiment analysis, named entity recognition, news classification, and question answering.

Open Financial LLMs. Can be fine-tuned swiftly to incorporate new data (the cost falls significantly, less than **$300 per fine-tuning**).

- FinGPT: Instruction Tuning Benchmark for Open-Source Large Language Models in Financial Datasets
- FinGPT: Democratizing Internet-scale Data for Financial Large Language Models

PIXIU(貔貅)

A Large Language Model, Instruction Data and Benchmark for Finance

https://arxiv.org/abs/2306.05443

HKUST's InvestLM: Hong Kong's First Open-Source Large Language Model for Financial Generative AI Applications
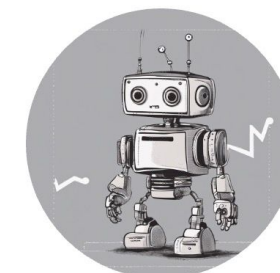
2023-11-29

InvestLM: A Large Language Model for Investment using Financial Domain Instruction Tuning

Yi Yang ; Yixuan Tang , Kar Yan Tam

Hong Kong University of Science and Technology
{imyiyang,yixuantang,kytam}@ust.hk

InvestLM

Example 1
Financial News Analysis with Llama Index, ChatGPT, and Streamlit

LlamaIndex: Swiftly start LLM apps, create (retrieval-augmented generation) RAG systems in minutes, access advanced data tools for customization.

A free and open-source framework to rapidly build and share beautiful machine learning and data science web apps. It is a Python-based library specifically designed for machine learning engineers.

What we will build?

- Single Stock Outlooks
- Competitor Analysis

Financial Analyst

Financial Reports

What type of report do you want?

Single Stock Outlook

Stock Symbol

goog

` Title: Alphabet (GOOG) Stock Outlook for 2023-2027

Introduction: Based on the analysis from Wall Street analysts over the past three months, Alphabet (NASDAQ: GOOG) stock has garnered positive reviews. With 1 out of 8 analysts giving it a bullish rating and the remaining 7 being somewhat bullish, it seems that the market expects Alphabet to perform well over the next few years. The following report will outline the outlook for GOOG stock from 2023 to 2027, including potential risks and headwinds.

Average Price Target: The average price target for Alphabet stock in the next 12 months, according to the 8 analysts, is 125.75, $with a high of$ $155.00 and a low of $116.00. This average price target has increased by 1.31% over the past month.

Positive Factors:

1. Continued Dominance in Online Advertising: Alphabet, through its subsidiary Google, largely dominates the online advertising market, which is expected to grow further in the coming years. As more businesses shift their advertising budgets to online platforms, Alphabet stands to benefit from this transition.

2. Expansion into Cloud Computing and Artificial Intelligence: Alphabet is investing heavily in cloud computing and AI technologies, which could further strengthen its position in these high-growth markets. The company's Google Cloud Platform is being increasingly adopted by businesses, and its AI offerings continue to innovate and adapt to new market needs.

Financial Analyst

Financial Reports

What type of report do you want?

Competitor Analysis

Stock Symbol 1

goog

Stock Symbol 2

msft

` Title: Competition Between Alphabet's (GOOG) and Microsoft's (MSFT) Stocks Intensifies as AI Developments Advance

Introduction: Alphabet (GOOG) and Microsoft (MSFT) stocks have been under the spotlight as emerging competition in the search engine space, driven by advancements in artificial intelligence (AI) technologies. Recent reports of Samsung considering replacing Google Search with Microsoft's Bing across its devices has further intensified the rivalry between these tech giants.

Samsung's Potential Switch to Bing: While Alphabet's Google has enjoyed a stronghold in the search engine space, a recent New York Times report suggests that tech giant Samsung may replace Google Search with Microsoft Bing as the default search engine across its devices. This has led to a decline in Alphabet's stocks by 2% as around

Overall Framework
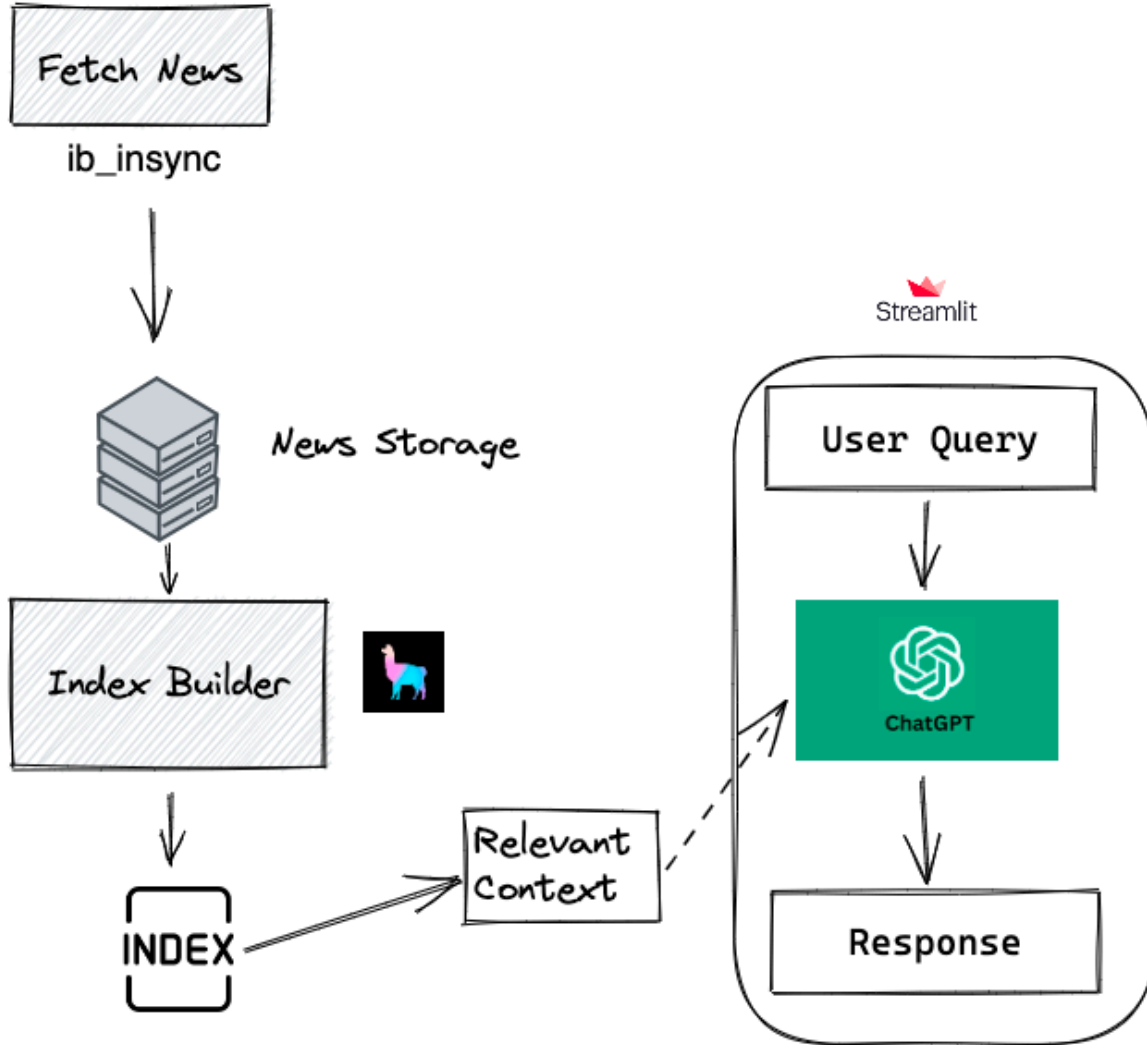
1. Few line codes
2. Very clear infrastructure
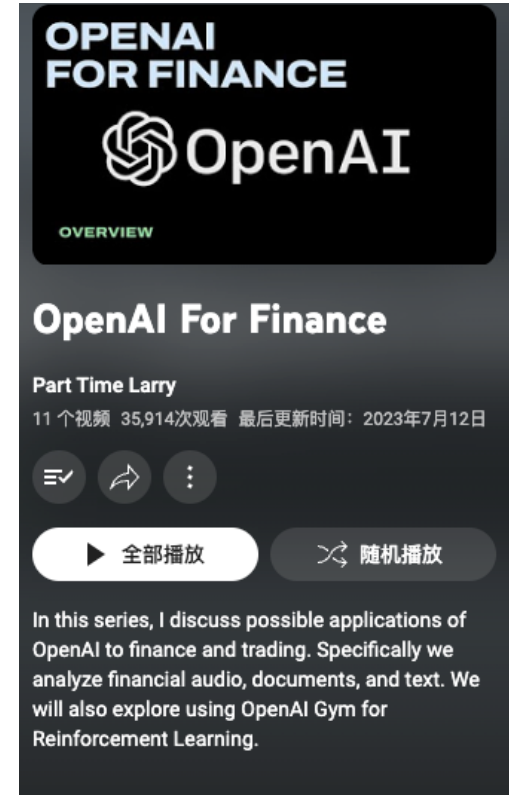
[Courses on Youtube](#)

Code Example

https://github.com/hackingthemarkets/financial-news-llama-index

hkust azure-openai-api-service
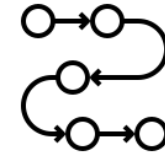
AZURE OPENAI API SERVICE

Available To

Students, Staff

Service Fee

You will receive a monthly free credit of HKD 8.00, which will be automatically recharged on the first day of each month at 8:00am and shared with the usage of HKUST ChatGPT platform. For additional usage, it is available for top-up credit at your own cost.

Get the OpenAI API Key step by step

Supported OpenAI API Models and Versions

- API Models

OpenAI Model name	Azure Deployment name
gpt-35-turbo	gpt-35-turbo
gpt-35-turbo-16k	gpt-35-turbo-16k
gpt-4	gpt-4
gpt-4-32k	gpt-4-32k
text-embedding-ada-002	text-embedding-ada-002

- API Versions
  - 2022-12-01
  - 2023-03-15-preview
  - 2023-05-15

Models	Context	Prompt (Per 1,000 tokens)	Completion (Per 1,000 tokens)
GPT-3.5-Turbo	4K	$0.0015	$0.002
GPT-3.5-Turbo	16K	$0.003	$0.004
GPT-4	8K	$0.03	$0.06
GPT-4	32K	$0.06	$0.12

Secondary key	XXXXXXXXXXXXXXXXXXXXXXXXX XXXXX	Show | Regenerate			
Name	My OpenAI	Rename	Azure OpenAI (personal)	Active	Cancel
Started on	07/25/2023				
Primary key	XXXXXXXXXXXXXXXXXXXXXXXXX XXXXX	Show | Regenerate			
Secondary key	XXXXXXXXXXXXXXXXXXXXXXXXX XXXXX	Show | Regenerate			

Other way to get API tokens

https://api.bianxie.ai/

	可用性 ⇅	模型名称 模糊搜索	计费类型 ⇅	倍率 ❓	模型价格
☐	✅	gpt-3.5-turbo	按量计费	模型: 0.25 补全: 3 分组: 1	提示 $0.5 / 1M tokens 补全 $1.5 / 1M tokens
☐	✅	gpt-3.5-turbo-0125	按量计费	模型: 0.25 补全: 3 分组: 1	提示 $0.5 / 1M tokens 补全 $1.5 / 1M tokens
☐	✅	gpt-3.5-turbo-0301	按量计费	模型: 0.75 补全: 1.333 分组: 1	提示 $1.5 / 1M tokens 补全 $2 / 1M tokens
☐	✅	gpt-3.5-turbo-0613	按量计费	模型: 0.75 补全: 1.333 分组: 1	提示 $1.5 / 1M tokens 补全 $2 / 1M tokens
☐	✅	gpt-3.5-turbo-1106	按量计费	模型: 0.75 补全: 2 分组: 1	提示 $1.5 / 1M tokens 补全 $3 / 1M tokens
☐	✅	gpt-3.5-turbo-16k	按量计费	模型: 1.5 补全: 1.333 分组: 1	提示 $3 / 1M tokens 补全 $4 / 1M tokens
☐	✅	gpt-3.5-turbo-16k-0613	按量计费	模型: 1.5 补全: 1.333 分组: 1	提示 $3 / 1M tokens 补全 $4 / 1M tokens
☐	✅	gpt-3.5-turbo-instruct	按量计费	模型: 0.75 补全: 1.333 分组: 1	提示 $1.5 / 1M tokens 补全 $2 / 1M tokens

AI/ML API   Product ⌄   Enterprise   Pricing   Resources ⌄   Contact us

One API
200+ AI Models
Uptime 99%

Integrate AI features via our secure API

https://aimlapi.com/

Example 2
Local-deployed FinGPT-Forecaster

What is FinGPT-Forecaster?

1. An AI tool analyzing market news and optional financial data for a specific company.
2. Offers a stock price movement forecast for the upcoming week along with summarized analysis.
3. Fine-tuned on **Llama-3-7b-chat-hf** with **LoRA** using past year's DOW30 market data.
4. Demonstrates adaptability and generalization capabilities across various stock symbols.

 GitHub link   Huggingface Demo

FinGPT: Open-Source Financial Large Language Models

downloads `2k` downloads/week `70` python `3.6` pypi `v0.0.1` license `MIT`

FINGPT

Let us not expect Wall Street to open-source LLMs or open APIs, due to FinTech institutes' internal regulations and policies.

Build by

AI4Finance

An open-source organization focused on advancing AI within the finance sector

FinGPT-Forecaster

FinGPT-Forecaster takes random market news and optional basic financials related to the specified company from the past few weeks as input and responds with the company's **positive developments** and **potential concerns**. Then it gives out a **prediction** of stock price movement for the coming week and its **analysis** summary.

This model is finetuned on Llama2-7b-chat-hf with LoRA on the past year's DOW30 market data. Inference in this demo uses fp16 and **welcomes any ticker symbol**. Company profile & Market news & Basic financials & Stock prices are retrieved using **yfinance & finnhub**. This is just a demo showing what this model is capable of. Results inferred from randomly chosen news can be strongly biased. For more detailed and customized implementation, refer to our FinGPT project: https://github.com/AI4Finance-Foundation/FinGPT

Disclaimer: Nothing herein is financial advice, and NOT a recommendation to trade real money. Please use common sense and always first consult a professional before trading or investing.

Ticker
Companys from Dow-30 are recommended

> AAPL

Date
Date from which the prediction is made, use format yyyy-mm-dd

> 2023-11-08

n_weeks
                                                                      3
Information of the past n weeks will be utilized, choose between 1 and 4

▬▬▬▬▬▬▬▬▬▬▬▬▬▬●▬▬▬▬▬▬▬

If checked, the latest quarterly reported basic financials of the company is taken into account.

☐ Use Latest Basic Financials

| Clear | **Submit** |

Information

[Company Introduction]:

Apple Inc is a leading entity in the Technology sector. Incorporated and publicly traded since 1980-12-12, the company has established its reputation as one of the key players in the market. As of today, Apple Inc has a market capitalization of 2827801.47 in USD, with 15634.23 shares outstanding.

Apple Inc operates primarily in the US, trading under the ticker AAPL on the NASDAQ NMS - GLOBAL MARKET. As a dominant force in the Technology space, the company continues to innovate and drive progress within the industry.

From 2023-10-18 to 2023-10-25, AAPL's stock price decreased from 175.84 to 171.10. Company news during this period are listed below:

[Headline]: Apple's iPhone 15 sold at deep discount by China's e-commerce giants
[Summary]: Top Chinese e-commerce platforms are offering discounts on Apple's latest iPhone 15, with JD.com cutting over $200 off the retail price.

[Headline]: Foxconn Hit by Chinese Investigations. Why the Apple Supplier Is in China's Sights.
[Summary]: Foxconn, also known as Hon Hai Precision Industry, was subjected recently to searches by Chinese authorities, according to China's state-run Global Times newspaper.

[Headline]: Apple (AAPL) Shares Pulled Back In Q3
[Summary]: Aristotle Atlantic Partners, LLC, an investment advisor, released its "Large Cap Growth Strategy" third

Llama 2

MODEL SIZE (PARAMETERS)	PRETRAINED	FINE-TUNED FOR CHAT USE CASES
7B	**Model architecture:**	**Data collection for helpfulness and safety:**
13B	**Pretraining Tokens:** 2 Trillion	**Supervised fine-tuning:** Over 100,000
70B	**Context Length:** 4096	**Human Preferences:** Over 1,000,000

Llama 2 is Meta's open source large language model (LLM). It's freely available for almost anyone to use for research and commercial purposes. **7B model could deploy on RTX-3090(24GB)**

LLama3.1 series

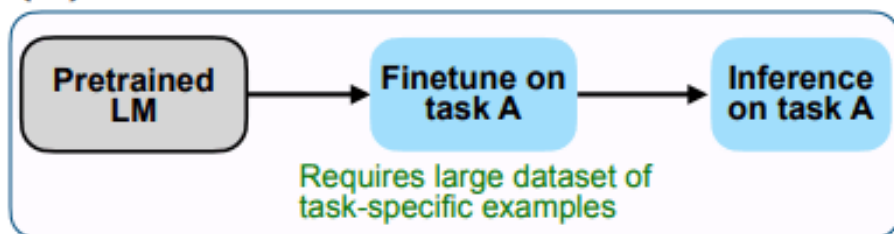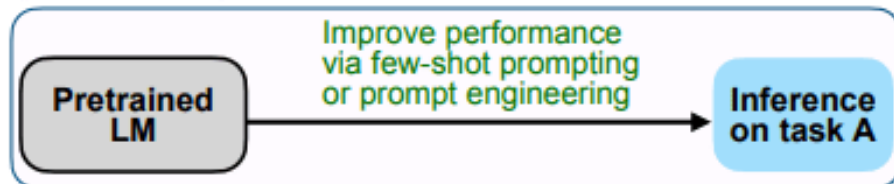| Category Benchmark	Llama 3.1 8B	Gemma 2 9B IT	Mistral 7B Instruct	Llama 3.1 70B	Mixtral 8x22B Instruct	GPT 3.5 Turbo
General MMLU (0-shot, CoT)	**73.0**	72.3 (5-shot, non-CoT)	60.5	**86.0**	79.9	69.8
MMLU PRO (5-shot, CoT)	**48.3**	-	36.9	**66.4**	56.3	49.2
IFEval	**80.4**	73.6	57.6	**87.5**	72.7	69.9
Code HumanEval (0-shot)	**72.6**	54.3	40.2	**80.5**	75.6	68.0
MBPP EvalPlus (base) (0-shot)	**72.8**	71.7	49.5	**86.0**	78.6	82.0
Math GSM8K (8-shot, CoT)	**84.5**	76.7	53.2	**95.1**	88.2	81.6
MATH (0-shot, CoT)	**51.9**	44.3	13.0	**68.0**	54.1	43.1
Reasoning ARC Challenge (0-shot)	83.4	**87.6**	74.2	**94.8**	88.7	83.7
GPQA (0-shot, CoT)	**32.8**	-	28.8	**46.7**	33.3	30.8
Tool use BFCL	**76.1**	-	60.4	84.8	-	**85.9**
Nexus	**38.5**	30.0	24.7	**56.7**	48.5	37.2
Long context ZeroSCROLLS/QuALITY	81.0	-	-	90.5	-	-
InfiniteBench/En.MC	65.1	-	-	78.2	-	-
NIH/Multi-needle	98.8	-	-	97.5	-	-
Multilingual Multilingual MGSM (0-shot)	**68.9**	53.2	29.9	**86.9**	71.1	51.4
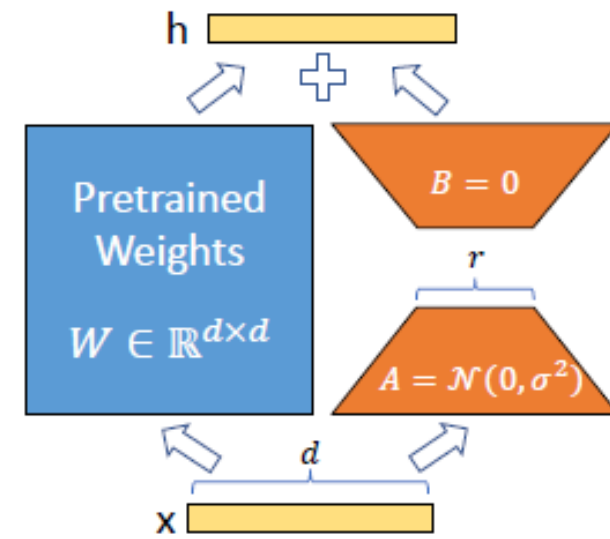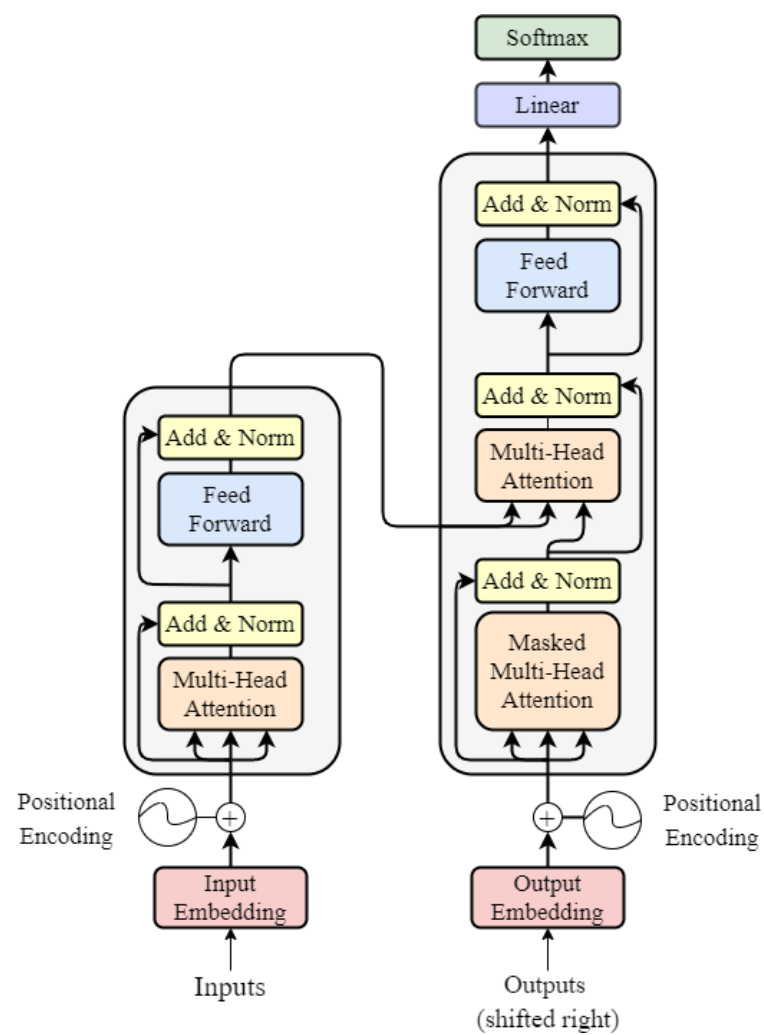
Figure 1: Our reparametriza-tion. We only train A and B.

LoRA (Low-Rank Adaptation of Large Language Models) is a popular and lightweight training technique that significantly reduces the number of trainable parameters. It works by inserting a smaller number of new weights into the model and only these are trained.
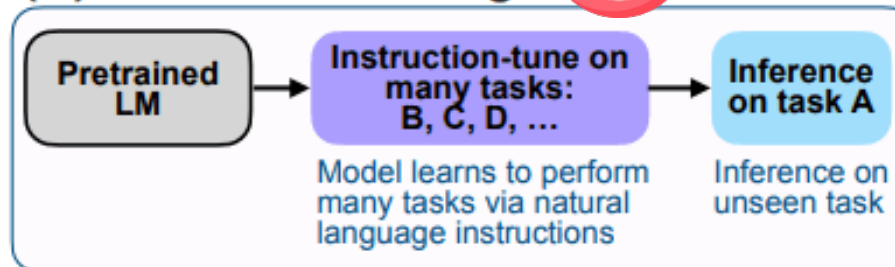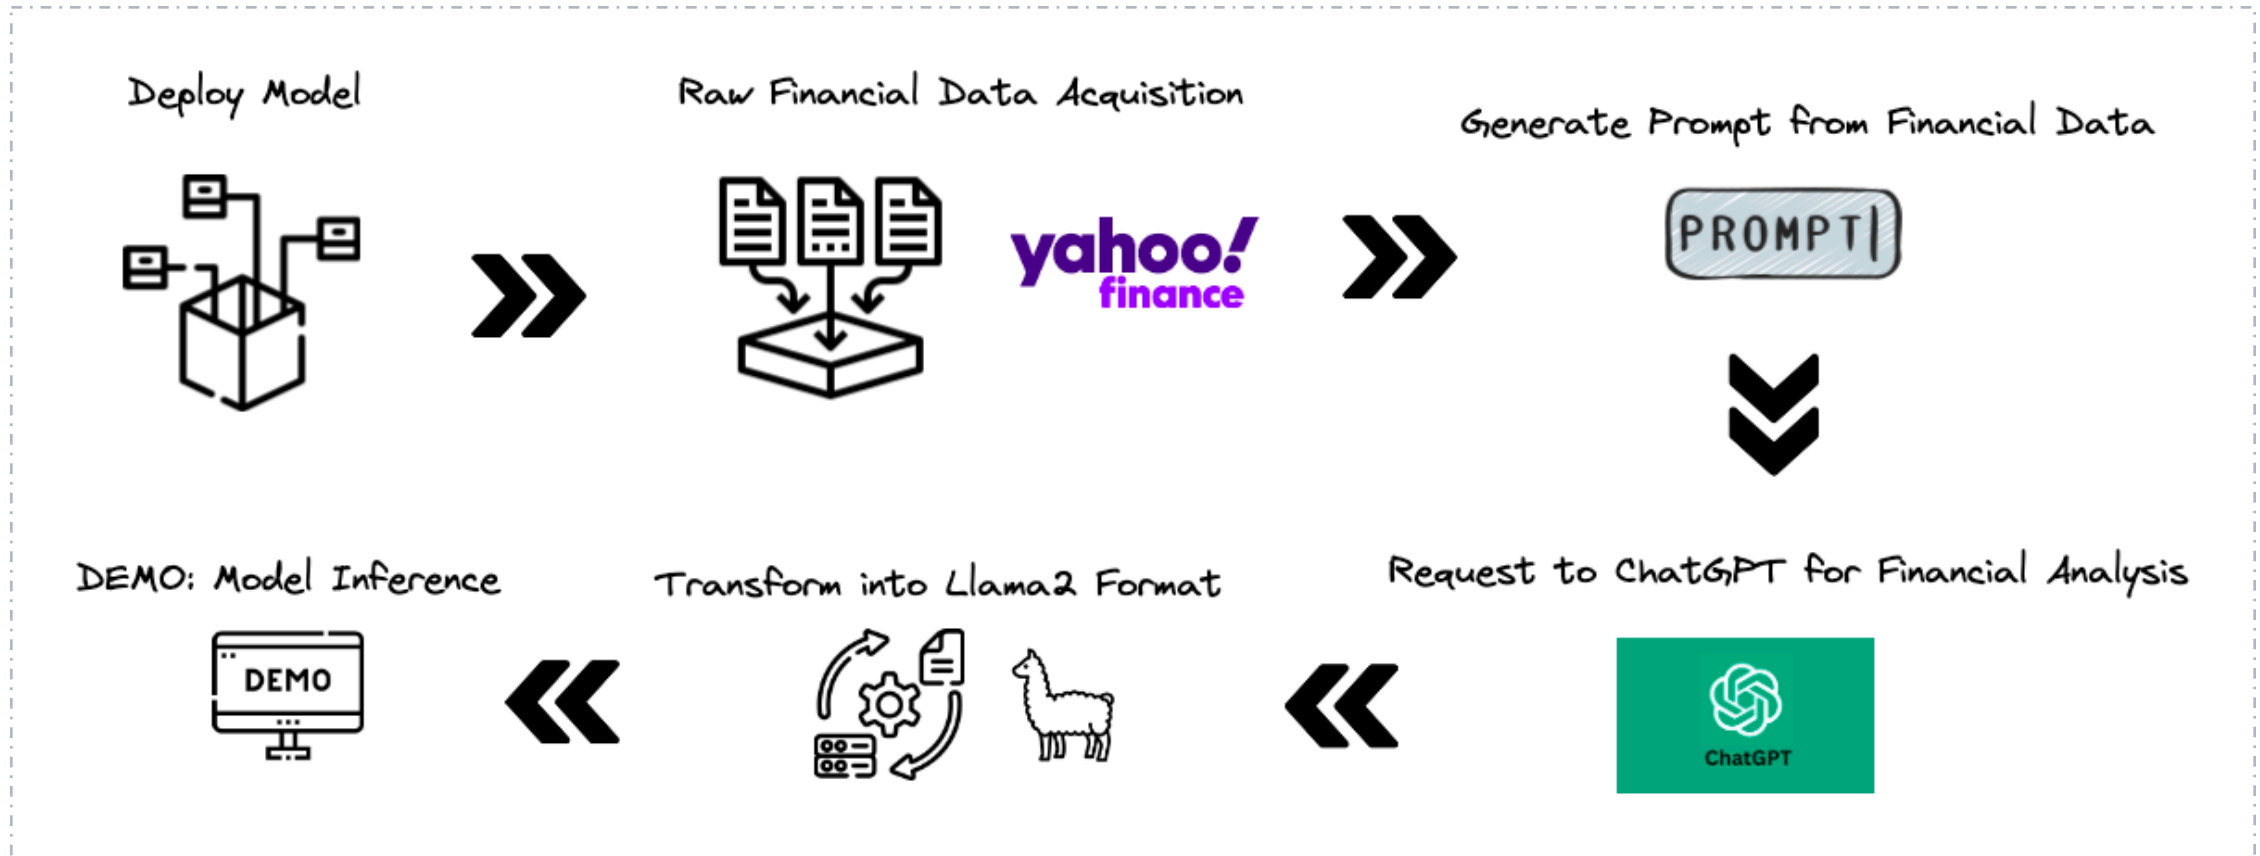
(A) Pretrain–finetune

Pretrained LM	→	Finetune on task A	→	Inference on task A

Requires large dataset of task-specific examples

(B) Prompting

Pretrained LM	→	Inference on task A

Improve performance via few-shot prompting or prompt engineering

(C) Instruction tuning

Pretrained LM	→	Instruction-tune on many tasks: B, C, D, …	→	Inference on task A

Model learns to perform many tasks via natural language instructions

Inference on unseen task

TO DO Try to deploy the **FinGPT-Forecaster** locally

Overall Framework

Deploy Model → Raw Financial Data Acquisition yahoo! finance → Generate Prompt from Financial Data PROMPT|

DEMO: Model Inference ← Transform into Llama2 Format ← Request to ChatGPT for Financial Analysis ChatGPT

References

Data Preparation Reference [CODE](#)

Demo Reference [CODE](#)

FinGPT-forecaster Huggingface Page [Huggingface](#)
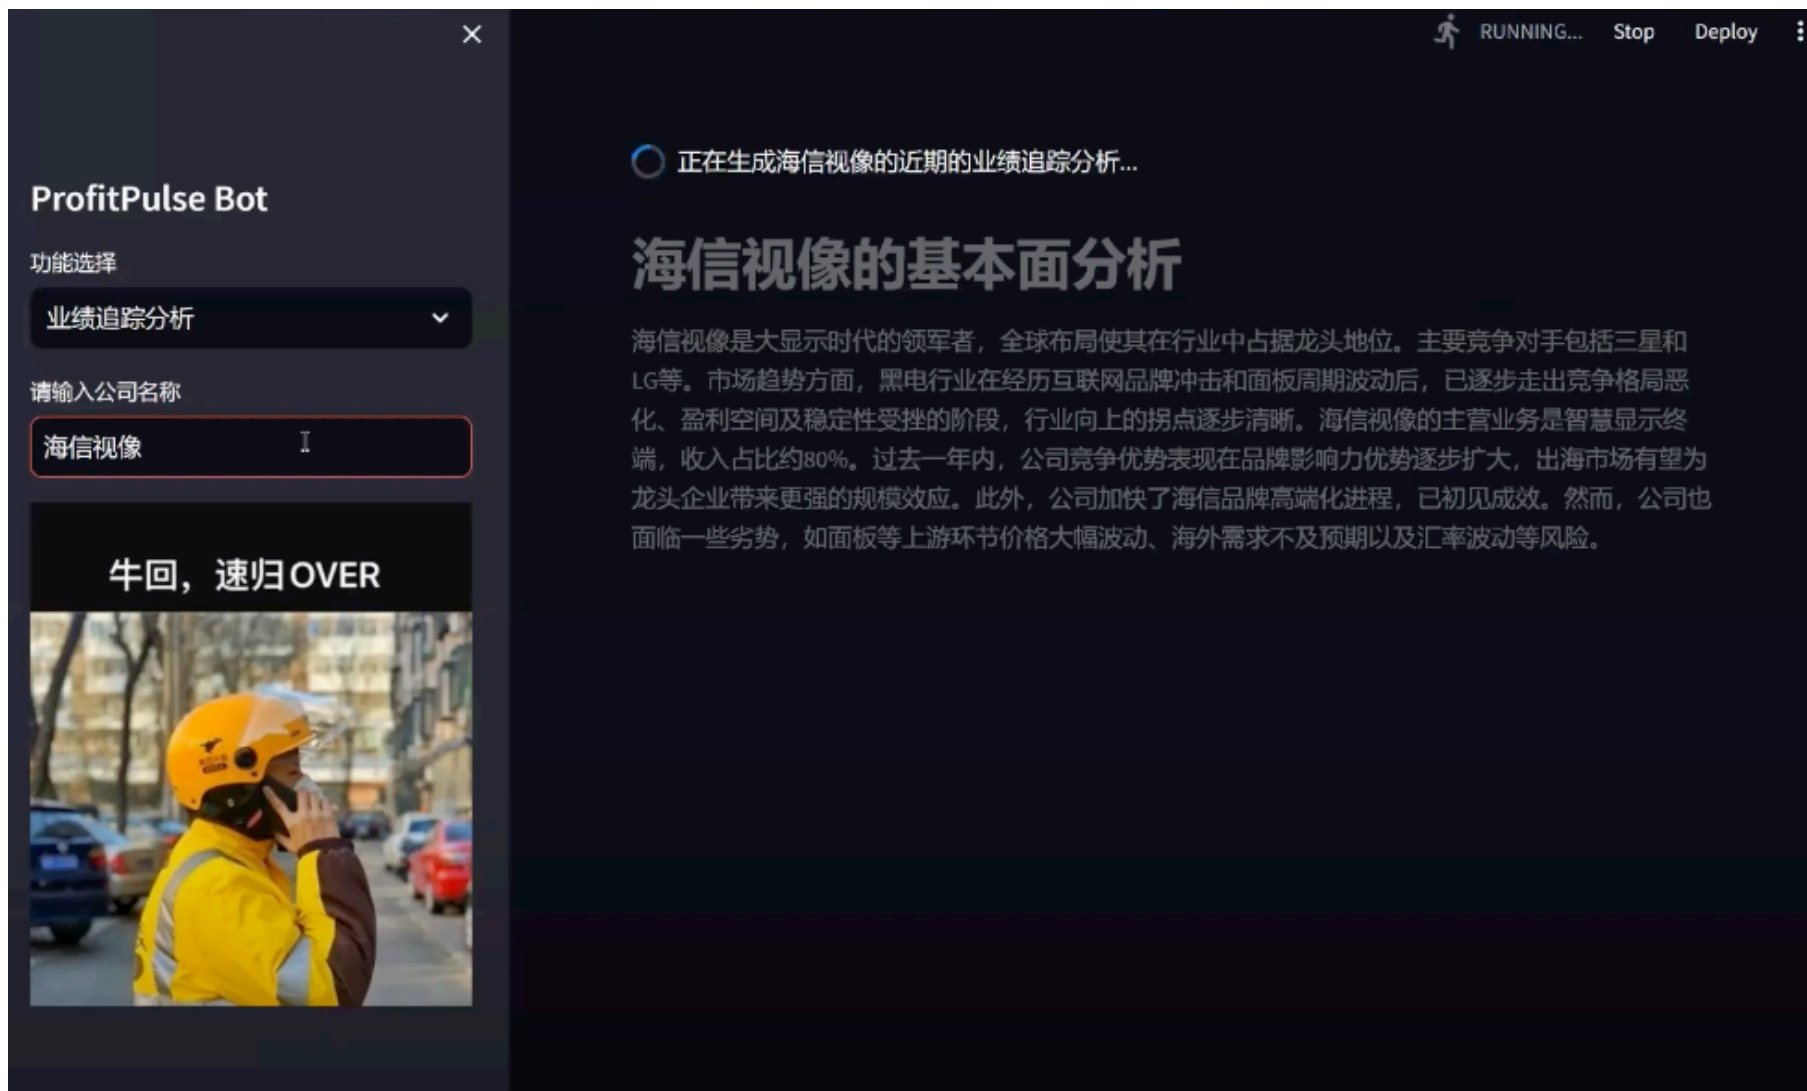
[Optional] OpenAI API Key: see previous slides

Huggingface Tutorials:
1. **Generation with LLMs** [link](#)
2. **NLP Course** [link](#)

Last year students' excellent projects presentation:

- https://www.youtube.com/watch?v=4nkZgyifR_s
- https://youtu.be/Lk8iM29FGcl?si=Pgf9oBBw4hH_YwEd