

FinRobot

***FinGPT Forecaster:
Leveraging Large Language Models
for Enhanced Robo-Advising in Finance***

DMO-FinTech@PAKDD2024


May 7, 2024

Hongyang (Bruce) Yang

Founder and President
AI4Finance Foundation, a 501(c)(3)
nonprofit organization
(ai4finance.org)



FINGPT



AI4Finance

Agenda

Why Open-Source FinGPT

Challenges of Handling Financial Data

FinGPT Framework

FinGPT-FinNLP: Data-Centric Design of Data Curation Pipeline

FinGPT-Benchmark: Instruction Tuning Paradigm for Financial Data

FinGPT-RAG: Retrieval Augmented Generation Framework

FinGPT-Forecaster: The Future of Robo-Advisory Services

FinGPT Project 2023 Academic Achievements

Conclusion

Why Open-Source FinGPT?

Data Perspective:

- **Witnessing a shifting trend towards democratizing Internet-scale financial data**
- **Promoting quality of financial data in a collaborative way, crucial for accurate modelling**
- **Sharing the understandings of financial data, crucial for events**

Model Perspective:

- **Directly applying general-purpose LLMs to finance may lead to sub-optimal or even contradictory results.**
- **Example: A layoff, negative to the public, may result in market going up (positive to investors)**

Wall Street cannot open-source LLMs nor open APIs, due to FinTech institutes' internal regulations and policies.

Financial Innovation through AI and Open-Source Collaboration

Our mission emphasizes the promotion of standardized practices and the development of open-source resources, benefiting both the research community and industry professionals.

Short-term Goal:

- **We aim to foster their adoption across the open-source finance ecosystem, provide ongoing maintenance and bug fixes for mature environments, and integrate more key open-source projects into our portfolio, elevating them to our standards of maturity.**

Long-term Goal:

- **Offer enhanced standardization tools for professionals engaged in applied financial tasks. This includes developing resources akin to a standardized API for financial environments, as well as broadening our suite of open-source tools beyond the foundational layers.**

Lowering the Cost of Training Domain-Specific LLMs

Cost Estimation for GPT-3 Training & BloombergGPT

- **GPT-3 (175B) Estimated Cost:** According to OneFlow, training GPT-3 once costs approximately **\$1.398 million**.
- **BloombergGPT (50B) Estimated Cost:** 512 GPUs for 53 days, 24 hours a day = 651,264 GPU hours. With \$4.1 per hour for an A100 GPU, the total cost is approximately **\$2,670,182.40**.

Scaling Down: Training a 17.5B Domain-Specific Model

- **Between \$140k - \$890k** based on the above estimated only for the GPU cost per training
- At least **one million dollar cost** to train a domain-specific LLM (GPU + Data + Manpower)

	GPT-3 (OpenAI)	Gopher (Google DeepMind)	MT-NLG (Microsoft/Nvidi a)	PaLM (Google Research)
Model Parameters	175B	280B	530B	540B
FLOPs/Token/Model Parameter			6	
TPUs/Machine			4	
Peak FLOPS/TPU			275T	
FLOPS Utilization			46.20%	
Cost/Machine/Hour(1-year reserved)			\$8.12	
Seconds/Hour			3600	
Training Cost/1000 Tokens	\$0.0047	\$0.0075	\$0.0141	\$0.0144
Train Tokens	300B	300B	270B	780B
Training Cost	\$1,398,072	\$2,236,915	\$3,810,744	\$11,216,529

Cost-Effective Strategies: Training smaller models can significantly reduce costs while maintaining efficacy.

For FinGPT: We use LoRA + open-source LLMs

- **Example:** Llama2-14B + LoRA cost about \$65.6 (One A100 16 hours)
- **Lower cost:** \$1,000,000 -> \$65.6

Challenges of Handling Financial Data

High Temporal Sensitivity:

- **Financial data are characterized by their time-sensitive nature**
- **Market-moving news provides a narrow window for investors to capture alpha signal**

High Dynamism:

- **Constant state of flux due to deluging of news, social media updates, etc.**
- **Retraining LLMs frequently is expensive and impractical**

Low Signal-to-Noise Ratio (SNR):

- **Financial data often contain a significant amount of irrelevant or noisy data**
- **Extracting valuable insights is labor-intensive**

Key Features of FinGPT

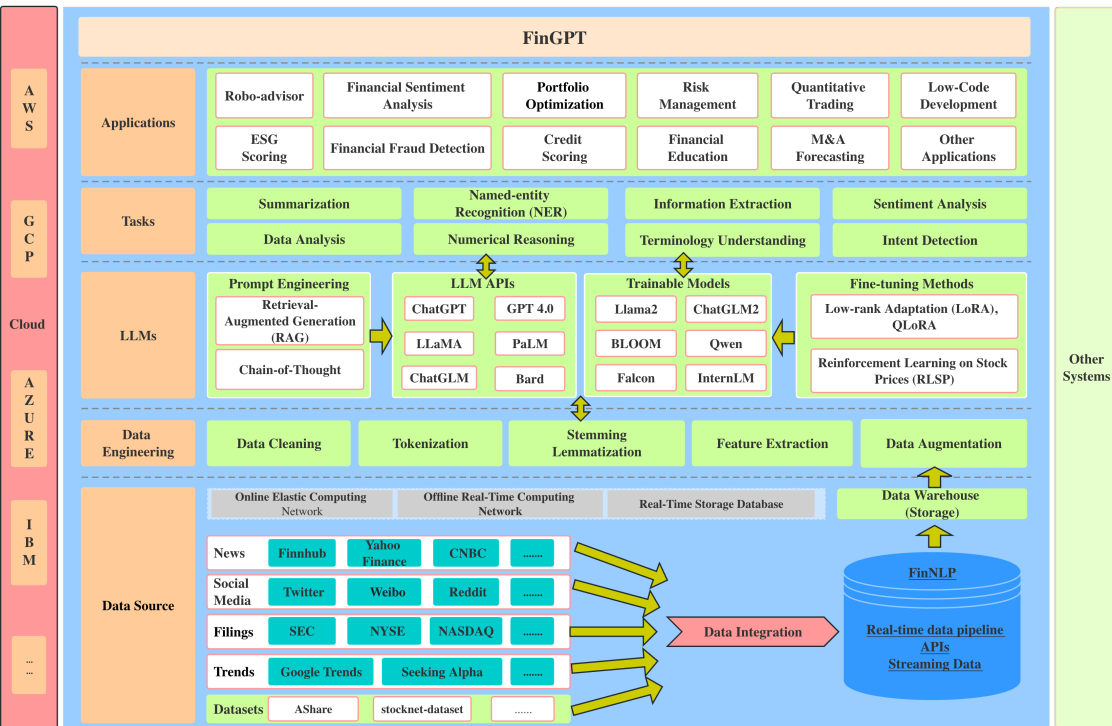
Which data to train? Democratizing Internet-scale Financial Data & Data-centric design of data curation pipeline

How to train? Instruction Tuning Paradigm & Retrieval Augmented Generation (RAG)

How to train efficiently? Low-rank Adaptation (LoRA, QLoRA).



FinGPT Framework



An End-to-End Framework

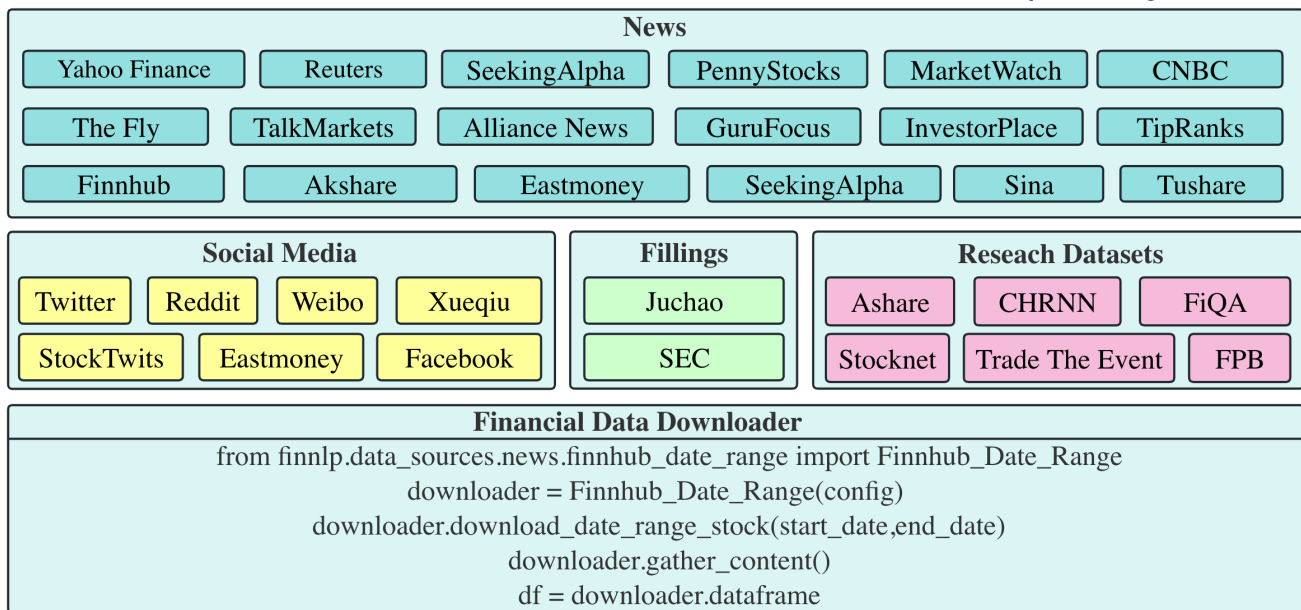
- **Applications layer:** Showcasing practical applications and demos, this layer highlights the potential capability of FinGPT in the finance sector
- **Tasks Layer:** the fundamental tasks serve as the benchmarks for performance evaluations and cross-comparisons in the realm of FinLLMs.
- **LLMs Layer:** On the LLMs layer, FinGPT focuses on a range of fine-tuning methodologies, this layer takes care of the highly dynamic nature of financial data, ensuring the model's relevance and accuracy
- **Data Engineering & Data Source Layer:** tackles the inherent challenges of high temporal sensitivity and low signal-to-noise ratio in financial data

FinGPT-FinNLP: Data-Centric Design of Data Curation Pipeline

Challenges: Diverse Data Sources, Data Quality Issues, High Time-Validity

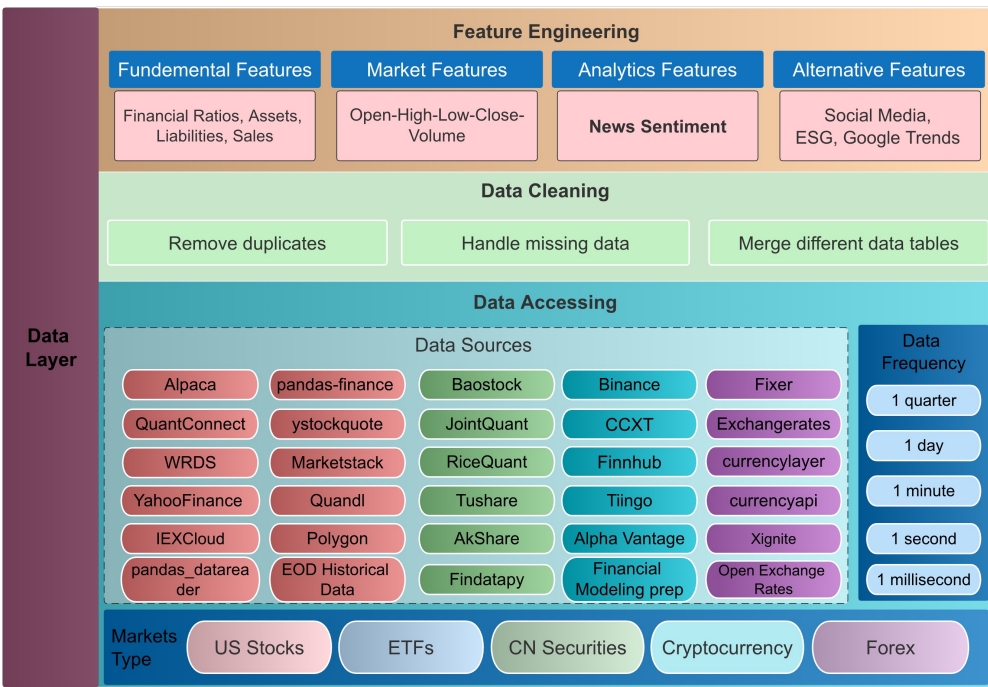
Our Solution: Data Curation Pipeline

Contributions: An open-source and data-centric framework, automating the collection and curation of real-time financial data from the Internet, with 34 financial data sources and the corresponding code.



FinGPT-FinNLP: Data-Centric Design of Data Curation Pipeline

• US



```
# Finnhub (Yahoo Finance, Reuters, SeekingAlpha, CNBC...)
from finnlp.data_sources.news.finnhub_date_range import Finnhub_Date_Range

start_date = "2023-01-01"
end_date = "2023-01-03"
config = {
    "use_proxy": "us_free", # use proxies to prvent ip blocking
    "max_retry": 5,
    "proxy_pages": 5,
    "token": "YOUR_FINNHUB_TOKEN" # Available at https://finnhub.io/dashboard
}

news_downloader = Finnhub_Date_Range(config) # init
news_downloader.download_date_range_stock(start_date,end_date) # Download headers
news_downloader.gather_content() # Download contents
df = news_downloader.dataframe
selected_columns = ["headline", "content"]
df[selected_columns].head(10)

-----

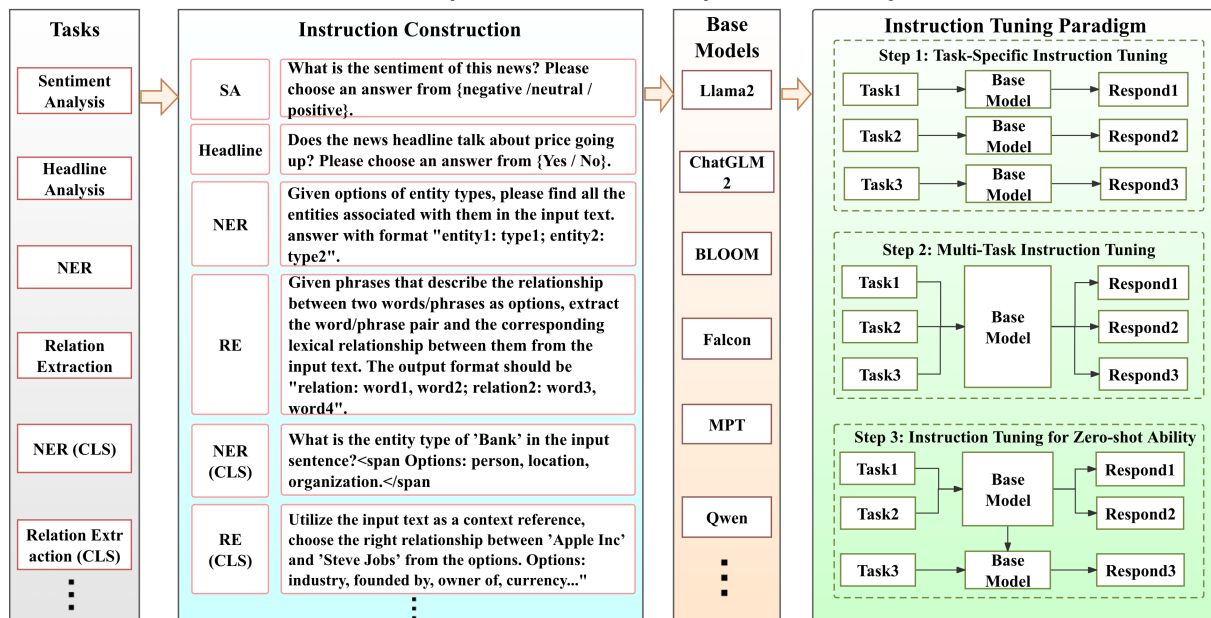
#     headline                                     content
# 0     My 26-Stock $349k Portfolio Gets A Nice Petrob... Home\nInvesting Strategy\nPortfc
# 1     Apple's Market Cap Slides Below $2 Trillion fo... Error
# 2     US STOCKS-Wall St starts the year with a dip; ... (For a Reuters live blog on U.S.
# 3     Buy 4 January Dogs Of The Dow, Watch 4 More     Home\nDividends\nDividend Quick Picks\nf
# 4     Apple's stock market value falls below $2 trill... Jan 3 (Reuters) - Apple Inc's \r
# 5     CORRECTED-UPDATE 1-Apple's stock market value ... Jan 3 (Reuters) - Apple Inc's \r
# 6     Apple Stock Falls Amid Report Of Product Order... Apple stock got off to a slow st
# 7     US STOCKS-Wall St starts the year with a dip; ... Summary\nCompanies\nTesla shares
# 8     More than $1 trillion wiped off value of Apple... apple store\nMore than $1 trill:
# 9     McLean's Iridium inks agreement to put its sat... The company hasn't named its pai
```

FinGPT-Benchmark: Instruction Tuning Paradigm for Financial Data

Current limitation: Focused on Single-Task Instruction Tuning

Our Solution: Expanding Instruction Tuning Paradigm on **Task-Specific, Multi-Task, and Zero-Shot Tuning**

Contributions: An Instruction Tuning paradigm, specifically **tailored for open-source Large Language Models (LLMs) in the financial sector**. Promotion of openness and reproducibility.



FinGPT-Benchmark: Single-Task Tuning

Data: <https://huggingface.co/datasets/FinGPT/fingpt-sentiment-train>

Model: https://huggingface.co/FinGPT/fingpt-sentiment_internlm-20b_lora

Code: https://github.com/AI4Finance-Foundation/FinGPT/blob/master/fingpt/FinGPT_Benchmark/train_lora.py

Model Name	Base-Model	FPB	FiQA	TFNS	NWGI
InternLM-20b-1gpu_8epochs_lr2e4_bs8_fp16	internlm-20b	0.878	0.892	0.904	0.646
FinGPT v3.3	llama2-13b	0.882	0.874	0.903	0.643
FinGPT v3.2	llama2-7b	0.850	0.860	0.894	0.636
FinGPT v3.1	chatglm2-6b	0.855	0.850	0.875	0.642

Run summary:

```
eval/loss 0.00395
eval/runtime 428.8234
eval/samples_per_second 35.807
eval/steps_per_second 4.477
train/epoch 7.99
train/global_step 7672
train/learning_rate 0.0
train/loss 0.0021
train/total_flos 5.522384844073468e+18
train/train_loss 0.05217
train/train runtime 49268.0316
```

We used a newly release model **InternLM-20B** to **fine-tune the sentiment analysis task** and achieved SOTA

Run history:

```
eval/loss
eval/runtime
eval/samples_per_second
eval/steps_per_second
train/epoch
train/global_step
train/learning_rate
train/loss
```

FinGPT-Benchmark: Multi-Task Tuning

==== Financial Sentiment Analysis ====

Instruction: What is the sentiment of this news? Please choose an answer from {negative/neutral/positive}.

Input: Glaxo's ViiV Healthcare Signs China Manufacturing Deal With Desano

Answer: positive

==== Financial Relation Extraction ====

Instruction: Given phrases that describe the relationship between two words/phrases as options, extract the word/phrase pair and the corresponding lexical relationship between them from the input text. The output format should be "relation1: word1, word2; relation2: word3, word4". Options: product/material produced, manufacturer, distributed by, industry, position held, original broadcaster, owned by, founded by, distribution format, headquarters location, stock exchange, currency, parent organization, chief executive officer, director/manager, owner of, operator, member of, employer, chairperson, platform, subsidiary, legal form, publisher, developer, brand, business division, location of formation, creator.

Input: Wednesday, July 8, 2015 10:30AM IST (5:00AM GMT) Rimini Street Comment on Oracle Litigation Las Vegas, United States Rimini Street, Inc., the leading independent provider of enterprise software support for SAP AG's (NYSE:SAP) Business Suite and BusinessObjects software and Oracle Corporation's (NYSE:ORCL) Siebel, PeopleSoft, JD Edwards, E-Business Suite, Oracle Database, Hyperion and Oracle Retail software, today issued a statement on the Oracle litigation.

Answer: product_or_material_produced: PeopleSoft, software; parent_organization: Siebel, Oracle Corporation; industry: Oracle Corporation, software; product_or_material_produced: Oracle Corporation, software; product_or_material_produced: Oracle Corporation, software

==== Financial Headline Classification ====

Instruction: Does the news headline talk about price in the past? Please choose an answer from {Yes/No}.

Input: april gold down 20 cents to settle at \$1,116.10/oz

Answer: Yes

==== Financial Named Entity Recognition ====

Instruction: Please extract entities and their types from the input sentence, entity types should be chosen from {person/organization/location}.

Input: Subject to the terms and conditions of this Agreement, Bank agrees to lend to Borrower, from time to time prior to the Commitment Termination Date, equipment advances (each an "Equipment Advance" and collectively the "Equipment Advances").

Answer: Bank is an organization, Borrower is a person.

Challenges: Task Interference & Hallucination

Task reformulation: We implement a strategy of task reformulation, we reform the instructions of all tasks into classification format

Instruction: [prompt] **Input:** [input] **Answer:** [output]



Instruction: [prompt] **Options:** [options]

Input: [input] **Answer:** [output]

FinGPT-Benchmark: Multi-Task Tuning

Phase	Dataset	Llama2	Falcon	MPT	BLOOM	ChatGLM2	Qwen
Task-Specific	FPB	0.863	0.846	0.872	0.810	0.850	0.854
	FiQA	0.871	0.840	0.863	0.771	0.864	0.867
	TFNS	0.896	0.893	0.907	0.840	0.859	0.883
	NWGI	0.649	0.636	0.640	0.573	0.619	0.638
	Avg	0.820	0.804	0.821	0.748	0.798	0.811
Multi-Task	FPB	0.861↓	0.845↓	0.870↓	0.766↓	0.836↓	0.873 ↑
	FiQA	0.825↓	0.881 ↑	0.863-	0.737↓	0.822↓	0.870↑
	TFNS	0.890↓	0.880↓	0.892 ↓	0.789↓	0.858↓	0.890↑
	NWGI	0.652↑	0.647↑	0.651↑	0.530↓	0.618↓	0.653 ↑
	Avg	0.807	0.813	0.819	0.701	0.784	0.822
Performance Gain		-1.3%	+0.7%	-0.2%	-4.7%	-1.4%	+1.1%

Table 3: Sentiment Analysis Instruction Tuning Results: The table reports detailed F1-scores for base models tuned during task-specific and multi-task phases on each sentiment analysis dataset. Arrows (↑↓) denote the influence of multi-task settings on Instruction Tuning results, with performance gains calculated between phases based on average F1 scores across all datasets.

For Single-Task job:

- No single model dominates across all tasks.
- The effectiveness of models varies depending on the specific task.

Task	Phase	Llama2	Falcon	MPT	BLOOM	ChatGLM2	Qwen
NER	Task-Specific	0.637	0.619	0.615	0.729	0.645	0.679
	Multi-Task	0.678↑	0.600↓	0.682↑	0.709 ↓	0.629↓	0.666↓
	Performance Gain	+4.1%	-1.9%	+6.7%	-2.0%	-1.6%	-1.3%
HC	Task-Specific	0.942	0.940	0.938	0.930	0.942	0.936
	Multi-Task	0.938 ↓	0.932↓	0.928↓	0.898↓	0.932↓	0.922↓
	Performance Gain	-0.4%	-0.8%	-1.0%	-3.2%	-1.0%	-1.4%
RE	Task-Specific	0.395	0.428	0.309	0.425	0.340	0.371
	Multi-Task	0.674↑	0.576↑	0.667↑	0.697 ↑	0.557↑	0.640↑
	Performance Gain	+27.2%	+14.8%	+35.8%	+27.2%	+21.7%	26.9%

Table 4: Multi-Task Instruction Tuning Summary: The table reports entity-level F1 scores for NER, relation-only F1 for RE, and standard classification F1 for HC. It includes both task-specific and multi-task models for comparison. Arrows (↑↓) signify performance gains from multi-task settings, calculated in each task's last row.

For Multi-Task job:

- Some models that perform exceptionally well in single-task jobs may not excel in multi-task job scenarios.
- The performance of models can be different in multi-task settings compared to single-task settings.

FinGPT-Benchmark: Zero-shot Results

- FinGPT showcases **an advanced level of fine-tuning** in financial sentiment analysis, **leveraging llama2-7b as base models**
- The benchmarks reveal **a striking performance advantage**, both in terms of **prediction scores and efficiency in training time and cost**

Metrics	Accuracy	F1
BloombergGPT	-	0.51
ChatGPT 4.0	0.64	0.51
ChatGLM2-6B	0.47	0.40
Llama-7B	0.60	0.40
Ours (Instruct-FinGPT)	0.76	0.74

Training Cost Comparison Between LLMs				
LLMs	BloombergGPT	ChatGLM2	Llama2	Instruct-FinGPT
Devices	512 × A100	64 × A100	2048 × A100	8 × A100
Time	53 days	2.5 days	21 days	2 hours
Cost	\$2.67 million	\$ 14,976	\$4.23 million	\$65.6

\$4.1 per GPU per hour

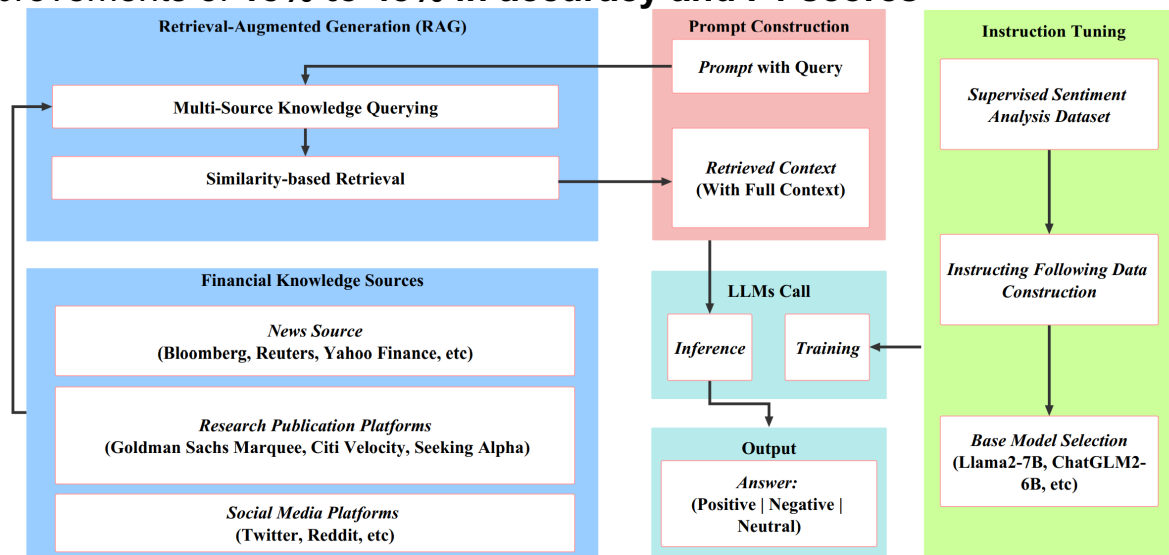
- Zero-shot evaluation between BloombergGPT, general-purpose LLMs ChatGPT, ChatGLM2-6B, Llama-7B, and our model on the **dataset of financial phaseBank (FPB)**

FinGPT-RAG: Retrieval Augmented Generation Framework

Current limitation: most financial news lack of adequate context information

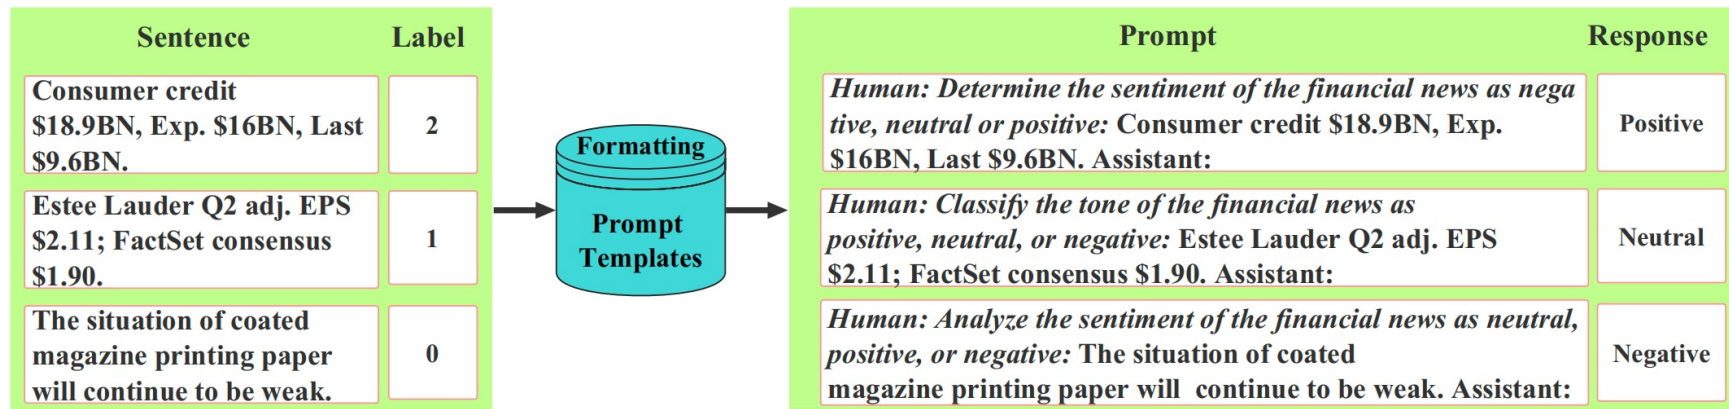
Our Solution: use **instruction tuning + retrieval augmented generation (RAG)** to fill up contexts

Contributions: Integrate external knowledge retrieval to **enhances information depth and context**. Utilizing specific instruction tuning processes, **the LLMs respond more accurately** to financial sentiment analysis tasks, achieving performance improvements of **15% to 48% in accuracy and F1 scores**



FinGPT-RAG: Retrieval Augmented Generation Framework

Format Instruction-following Dataset

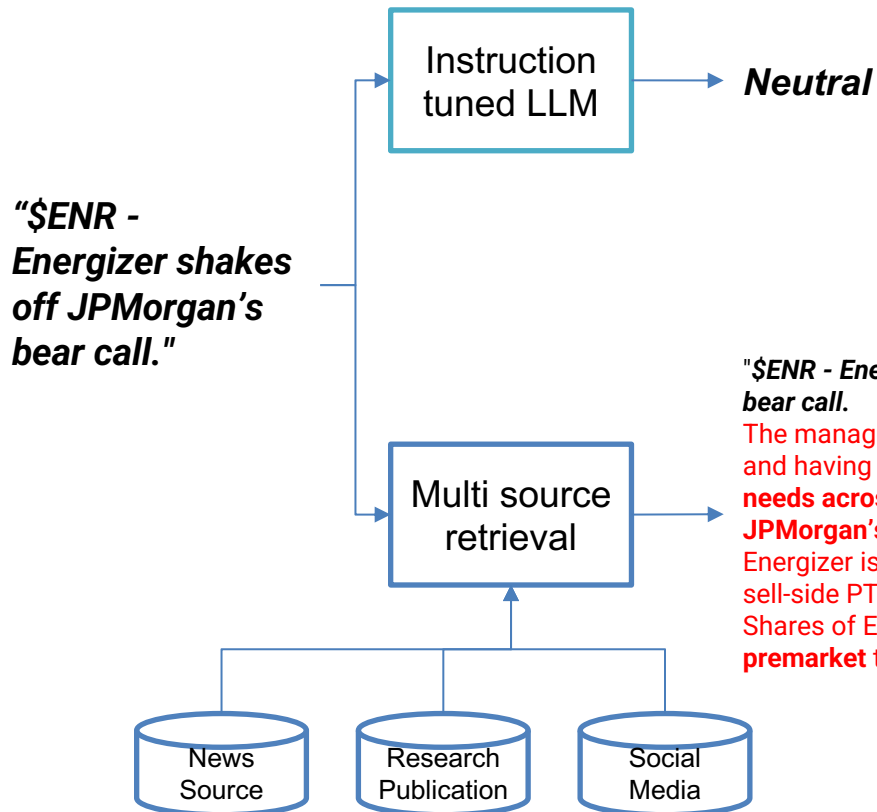


Training Objective

$$\mathcal{L}_{\text{CausalLM}} = - \sum_{t=1}^T \log P(w_t | w_1, w_2, \dots, w_{t-1}; \theta)$$

$$\nabla_{\theta} \mathcal{L}_{\text{CausalLM}} = - \sum_{t=1}^T \frac{\partial \log P(w_t | w_1, w_2, \dots, w_{t-1}; \theta)}{\partial \theta}$$

Examples and Results using RAG Framework for FinLLMs



"\$ENR - Energizer shakes off JPMorgan's bear call.
 The management discussed their, and having a portfolio **that fits consumer needs across the value spectrum.** JPMorgan's price target of \$51 on Energizer is based on, The average sell-side PT on Energizer is **\$56.17.** Shares of Energizer **are up 0.46% premarket** to \$50.44."

Metrics	Acc	F1
ChatGPT 4.0 w/o RAG	0.788	0.652
ChatGPT 4.0 w/ RAG	0.813	0.708
Ours w/o RAG	0.863	0.811
Ours w/ RAG	0.881	0.842

Table 2. Experimental results on the Twitter Val dataset.

FinGPT-Forecaster: The Future of Robo-Advisory Services

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
Companies from Dow-30 are recommended

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

n_weeks
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

Information

Response

FinGPT-Forecaster: The Future of Robo-Advisory Services

Kweichow Moutai 600519

Forecaster for Kweichow Moutai on 2024-01-28

[积极发展]:

1. 贵州茅台的购买率持续上涨，表明市场对茅台的需求持续增长。(News)
2. 公司的新闻发布，如“美茅台美世界”2024全球品牌文化活动，表明公司继续追求国际化发展，并且将继续以“美”的国际表达，坚持追求美，实现看好未来的发展。(News)
3. 茅台的营业总收入同比增长率达到了19.80%，表明公司的经营状况良好，具有稳定的盈利能力。(Basic Financials)
4. 茅台的流动比率和速动比率都相对较高，表明公司的资金流动性较好，可以更好地应对市场的压力。(Basic Financials)

[潜在担忧]:

1. 茅台的资产负债率达到了17.98%，显示公司的负债压力较大，可能会对公司的盈利能力产生影响。(Basic Financials)
2. 公司未来可能面临的市场竞争压力可能会增加，如果公司无法继续提供高质量的产品和服务，可能会受到市场反应。(News)
3. 茅台的购买率上涨，可能会导致市场价格的偏高，可能会影响到公司的股价。(Stock Price)

[预测和分析]:

预测涨跌幅: 上涨0-1%

总结分析: 虽然贵州茅台的资产负债率较高，但公司的经营状况良好，营业总收入同比增长率达到了19.80%，显示公司在市场的需求和竞争中具有稳定的盈利能力。考虑到茅台的购买率持续上涨，表明市场对茅台的需求持续增长，可能会对公司的股价产生积极影响。同时，公司的新闻发布，如“美茅台美世界”2024全球品牌文化活动，表明公司继续追求国际化发展，可能会带来未来的发展机会。因此，预计茅台的股价在下一周将上涨0-1%。

NVDA

Forecaster for NVDA on 2024-01-29

[Positive Developments]:

1. NVDA's stock price has been steadily increasing over the past weeks, suggesting a strong investor sentiment. (Stock Price)
2. The company's CEO, Jensen Huang, has expressed optimism about the future of AI and Nvidia's role in it, which could boost investor confidence. (News)
3. The company has received attention from Zacks.com users, indicating a strong performance in the market. (News)
4. The company has a strong balance sheet with a healthy gross margin and a high EV. (Basic Financials)

[Potential Concerns]:

1. The company's long-term debt-to-total asset and debt-to-total capital ratios are relatively high, indicating a potential risk of financial instability. (Basic Financials)
2. The company's high PE ratio and high net debt-to-total equity ratio suggest high valuation and potential financial risk. (Basic Financials)
3. The company's quick ratio is relatively low, indicating a potential liquidity risk. (Basic Financials)

[Prediction & Analysis]:

Prediction: Up by 0-1%

Analysis: Despite some potential concerns, the positive developments for NVDA are significant. The company's strong balance sheet and increasing stock price suggest a strong investor sentiment. The optimistic outlook from the CEO and the attention from Zacks.com users also suggest a positive future outlook. However, the high long-term debt-to-total asset and debt-to-total capital ratios are a cause for concern. These high ratios could potentially increase the company's financial risk and affect its liquidity.

FinGPT-Forecaster: The Future of Robo-Advisory Services

FinGPT-Forecaster is an LLMs model that synthesizes **recent market news** and **relevant financial ratios** of a given company to provide a dual output – a rundown of the **company's latest positive strides and potential concerns**, as well as a **forecast of the stock price movements** for the upcoming week, complete with an analytic summary.

Developed with a **fine-tuned Llama-2-7b-chat-hf with LoRA**, leveraging the latest year's market data from the **DOW 30**, **FinGPT-Forecaster** not only brings **precision** to predictions for **these blue chips** but also demonstrates **remarkable generalization capabilities** across various stock symbols.

FinGPT-Forecaster stands as a testament to the promise and potential of AI in finance, a **junior robo-advisor** that combines ease of deployment with strategic foresight, **marking a significant milestone on our path to the future of automated financial advisory.**

FinGPT-Forecaster: How to Use?

FinGPT-Forecaster is hosted on Hugging Face Spaces, allowing anyone with internet access to use it without any cost. It's the embodiment of open-source philosophy – shared, improved, and used by a community of developers and financial analysts alike.

Inputting Data

To start your forecasting journey, you simply need to:

- 1. Select Your Ticker Symbol:** Enter the ticker symbol for the company you are interested in, such as '**AAPL**' for Apple Inc. or '**MSFT**' for Microsoft.
- 2. Set Your Date:** Choose the specific day (formatted as **yyyy-mm-dd**) from which you want the prediction to commence.
- 3. Determine the News Timeframe:** Decide on the number of **past weeks** for which you'd like to gather market news. This helps the model to understand recent trends and sentiments.
- 4. Incorporate Financials:** Opt-in to add the **latest basic financials** for a more informed prediction, if desired.

<https://huggingface.co/spaces/FinGPT/FinGPT-Forecaster>

https://huggingface.co/FinGPT/fingpt-forecaster_dow30_llama2-7b_lora

FinGPT Project 2023 Academic Achievements

October 2023

- **FinGPT: Instruction Tuning Benchmark for Open-Source Large Language Models in Financial Datasets;** Instruction Workshop @ NeurIPS 2023; <https://arxiv.org/abs/2310.04793>
- **FinGPT: Democratizing Internet-scale Data for Financial Large Language Models;** Instruction Workshop @ NeurIPS 2023; <https://arxiv.org/abs/2307.10485>


September 2023

- **Enhancing Financial Sentiment Analysis via Retrieval Augmented Large Language Models;** ICAIF 2023; <https://arxiv.org/abs/2310.04027>

July 2023

- **Instruct-FinGPT: Financial Sentiment Analysis by Instruction Tuning of General-Purpose Large Language Models;** FinLLM 2023@IJCAI 2023; <https://arxiv.org/abs/2306.12659>
- **FinGPT: Open-Source Financial Large Language Models;** FinLLM 2023@IJCAI 2023; <https://arxiv.org/abs/2306.06031>

FinRobot: An Open-Source AI Agent Platform for Financial Applications using LLM



FinRobot

Conclusion

The models and data pipeline are open-sourced on huggingface :

Model: <https://huggingface.co/FinGPT>

Data Pipeline: <https://github.com/AI4Finance-Foundation/FinNLP>

FinGPT Github Repo: <https://github.com/AI4Finance-Foundation/FinGPT>, 10.1k stars

Tutorials for Beginners: [\[Training\] Beginner's Guide to FinGPT: Training with LoRA and ChatGLM2-6B One Notebook, \\$10 GPU](#)

Education Channel: <https://byfintech.medium.com/>

- <https://medium.datadriveninvestor.com/introducing-fingpt-forecaster-the-future-of-robot-advisory-services-50add34e3d3c>
- <https://medium.datadriveninvestor.com/fingpt-powering-the-future-of-finance-with-20-cutting-edge-applications-7c4d082ad3d8>
- <https://medium.datadriveninvestor.com/fingpt-ii-cracking-the-financial-sentiment-analysis-task-using-instruction-tuning-of-3333bce428c4>

Discord Group: <https://discord.gg/trsr8SXpW5>

LinkedIn Group: <https://www.linkedin.com/groups/14297568/>



FINGPT