TimeCAP: Learning to Contextualize, Augment, and Predict Time Series Events with Large Language Model Agents

KAIST

Geon Lee, Wenchao Yu, Kijung Shin, Wei Cheng, and Haifeng Chen

{geonlee0325, kijungs} @ kaist.ac.kr {wyu, weicheng, haifeng} @ nec-labs.com

https://github.com/geon0325/TimeCAP

NEC Laboratories America

Summary

- **Summary** We utilize **large language models (LLMs)** for **time series event prediction** focusing on three key purposes: (1) Contextualization, (2) Augmentation, and (3) Prediction.
- **Method** We introduce **TimeCAP**, an effective framework for **time series event prediction** using **LLMs agents** through three key steps:
  - **Contextualize** time series data into a textual summary;
  - **Augment** raw time series data and prompts;
  - **Predict** the outcome of future events.
- **Experiments** **TimeCAP** demonstrates outstanding performance with:
  - **Accurate**: Achieves up to 28.75%↑ F1 score over SOTA methods;
  - **Effective**: Employs LLMs beyond their typical roles as predictors;
  - **Interpretable**: Provides clear rationales behind its predictions.

Time Series Event Prediction

- Real-world **time series data** often involves **contextual information**.
  - **e.g. 1,** Hourly temperature is associated with geographical factors.
  - **e.g. 2,** Daily stock prices are affected by market trends.
- → **Contextual insights beyond raw time series data are crucial.**

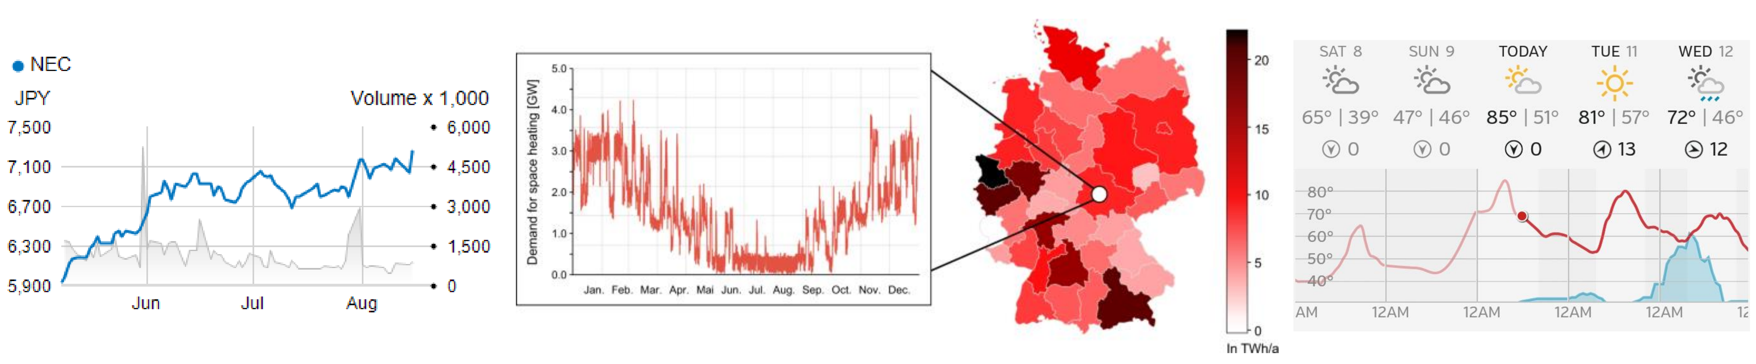

Stock Price Energy Consumption Weather

- **Time series event prediction** is crucial in various applications.
  - **Input:** Time series data;
  - **Output:** Predicted outcome of the future events.
  - **Goal:** To provide ***accurate*** and ***interpretable*** predictions.

Input Time Series

Will it rain tomorrow?

Will the # of infected people increase in the next season?

Will the stock price increase tomorrow?

LLMs as Time Series Predictors

- **Large language models (LLMs)** exhibit following strengths:
  - Sophisticated **reasoning and pattern recognition** capabilities;
  - Remarkable **few-shot and zero-shot learning** capabilities.

GPT-4 Gemini Llama 3.2

- As a result, **LLMs** have been used for **time series analysis**.
  - Approach 1: **Prompt LLMs** with time series data.
  - Approach 2: **Fine-tune LLMs** using time series data.

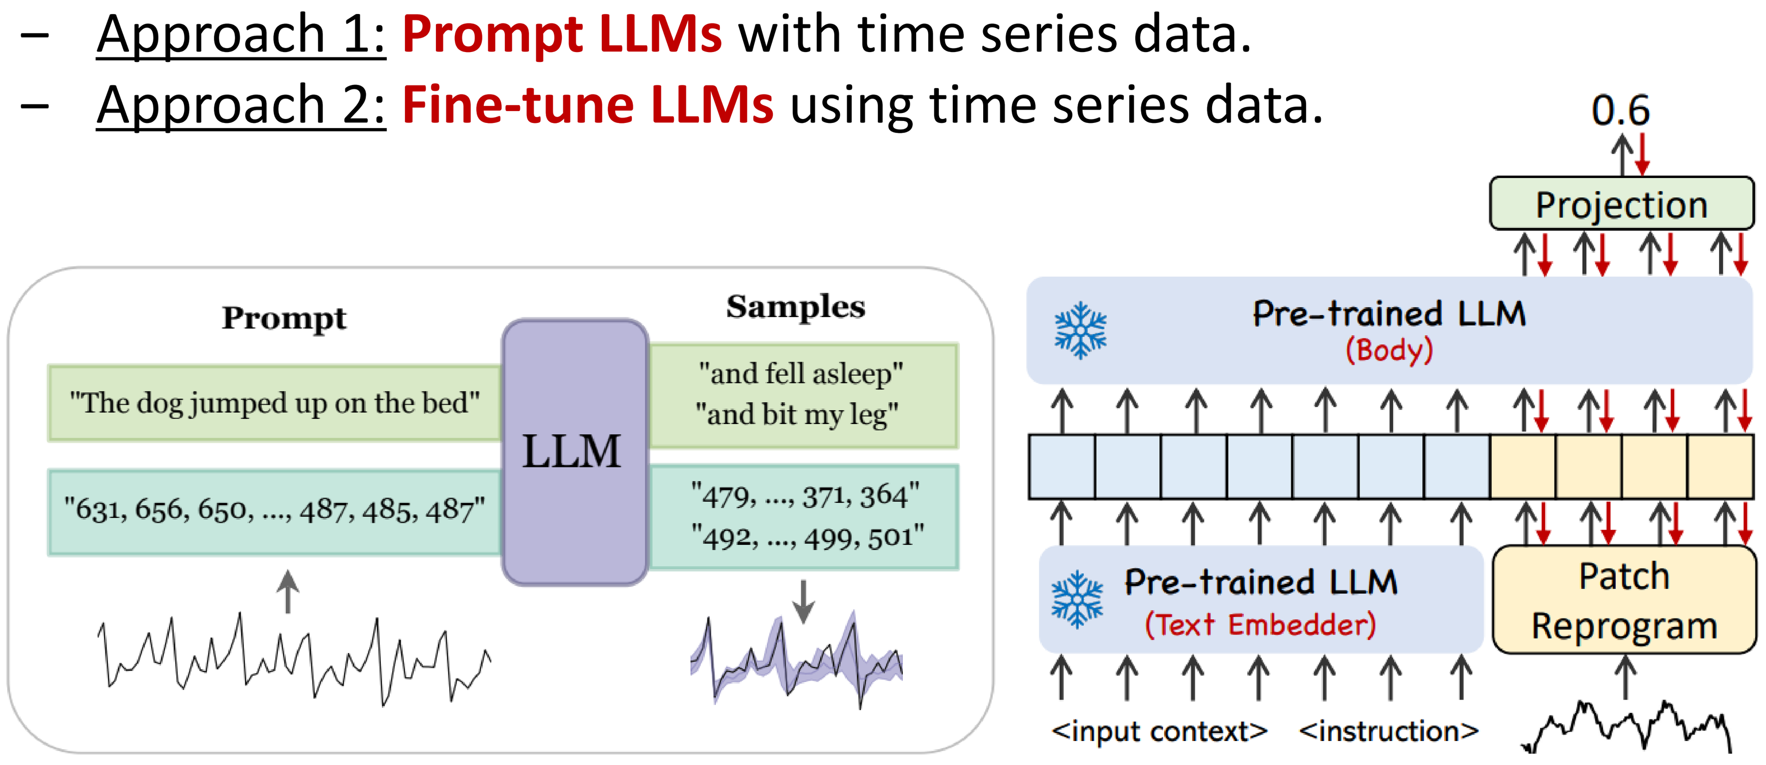

Prompt LLMs with Time Series (e.g., LLMTime; NeurIPS'23)

Fine-tune LLMs using Time Series (e.g., TimeLLM; ICLR'24)

- Existing approaches have focused on using **LLMs as "predictors."**
  - LLMs' **contextual understanding capabilities** are not fully utilized.

TimeCP: Contextualize & Predict

- We present **TimeCP**, our preliminary method for time series event prediction by introducing **two LLM agents**:
  - (1) **A contextualizer** generates a textual summary of input time series;
  - (2) **A predictor** predicts future events based on the summary.
- → **Contextual insights beyond raw time series data are incorporated.**

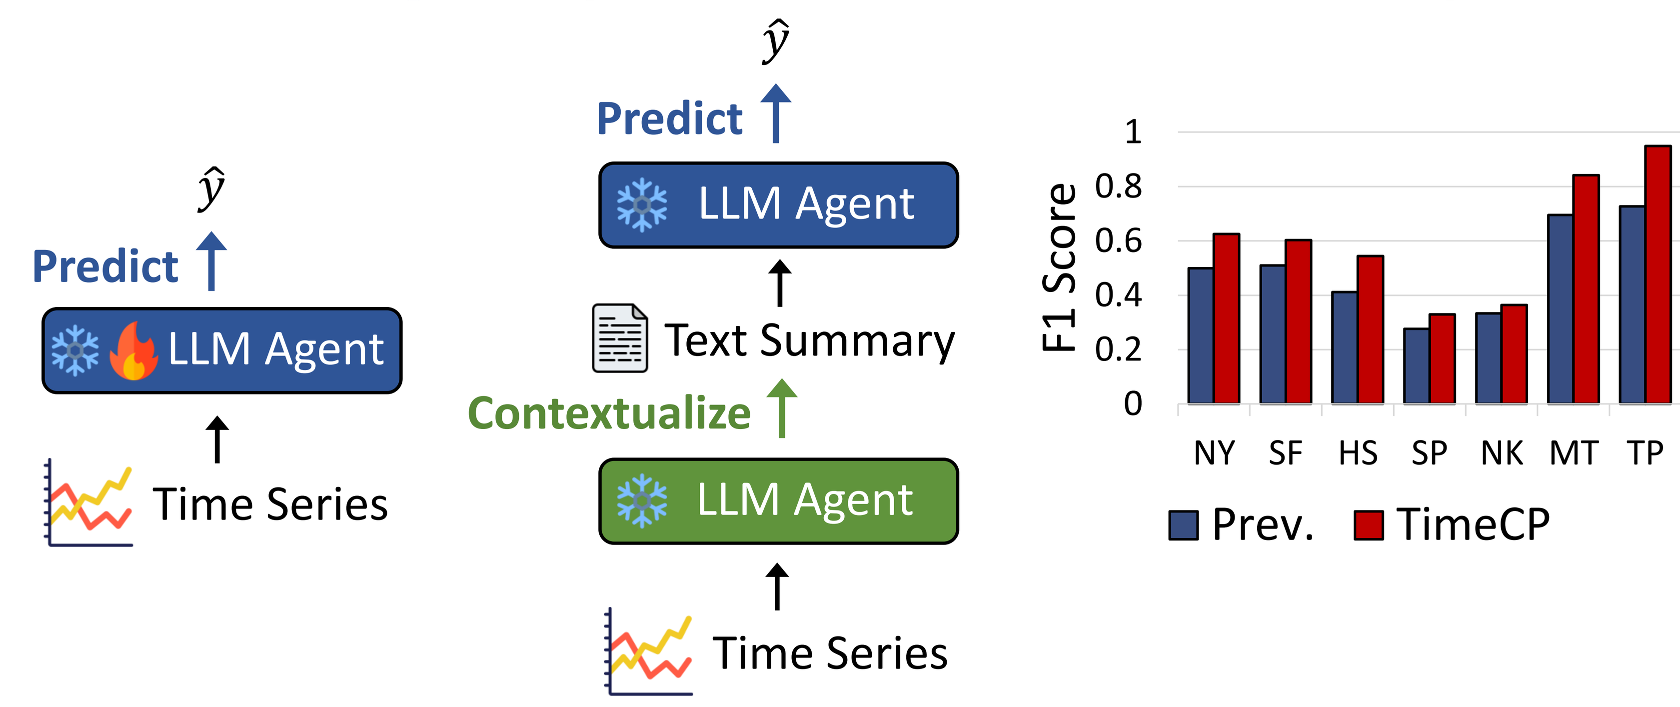

TimeCAP: Contextualize, Augment, and Predict

- We present **TimeCAP**, our advanced version of our framework.
  - It employs **a multi-modal encoder** that ***synergizes*** **with LLM agents**.

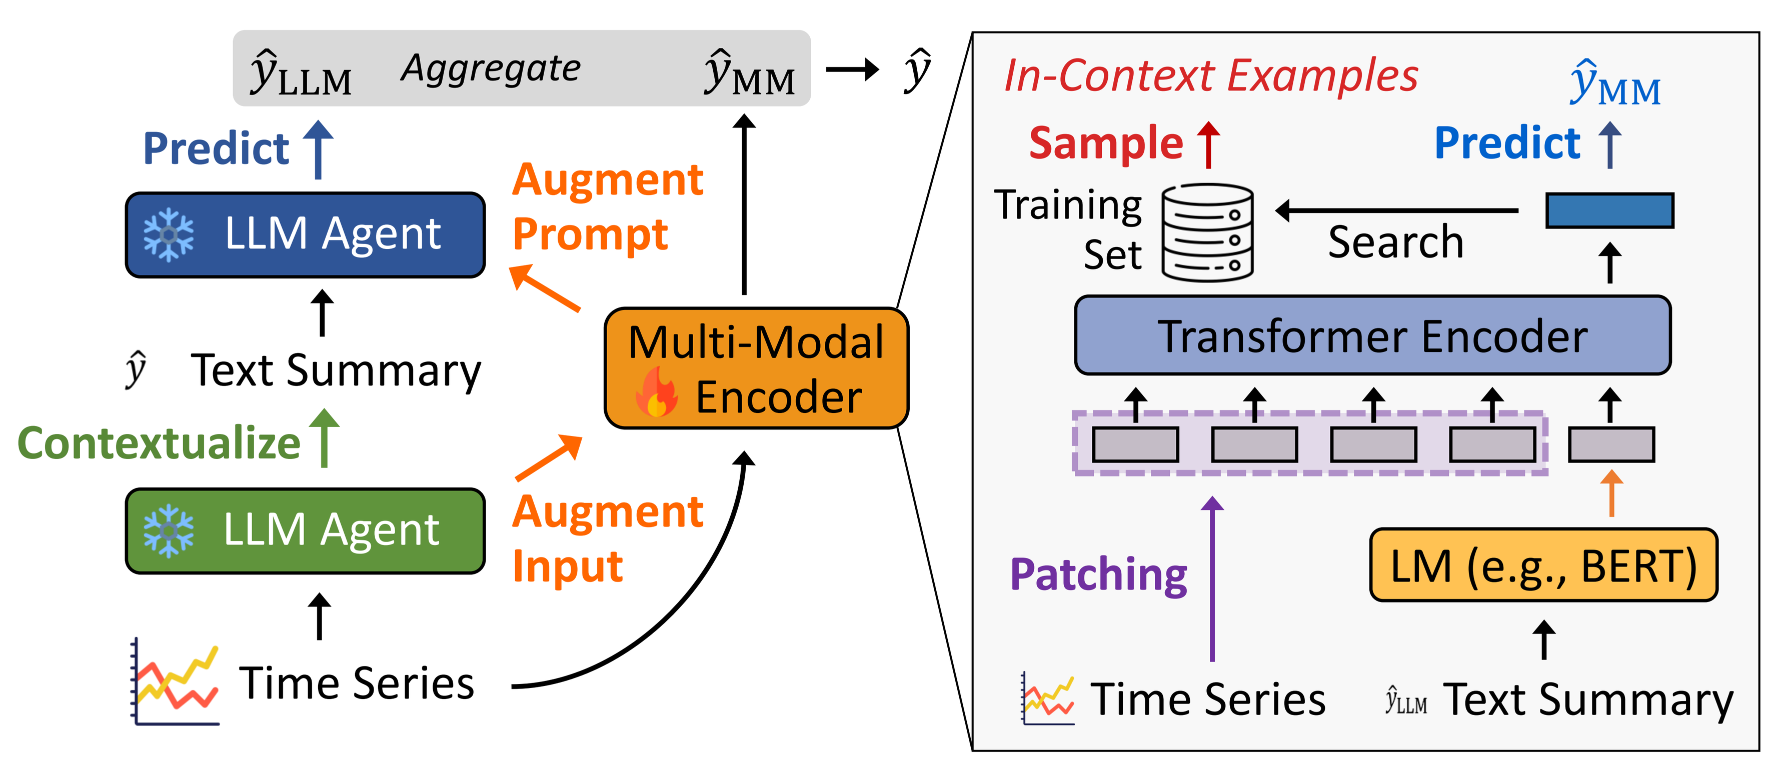

- **Input augmentation:** The textual summaries generated by the LLM agent provide contextual insights to the multi-modal encoder.
- **Prompt augmentation:** The multi-modal encoder learns enhanced input representations to retrieve highly relevant in-context examples.

Experimental Results

- **Accuracy** **TimeCAP** outperforms its competitors (e.g., PatchTST, GPT4TS) under various training ratios.

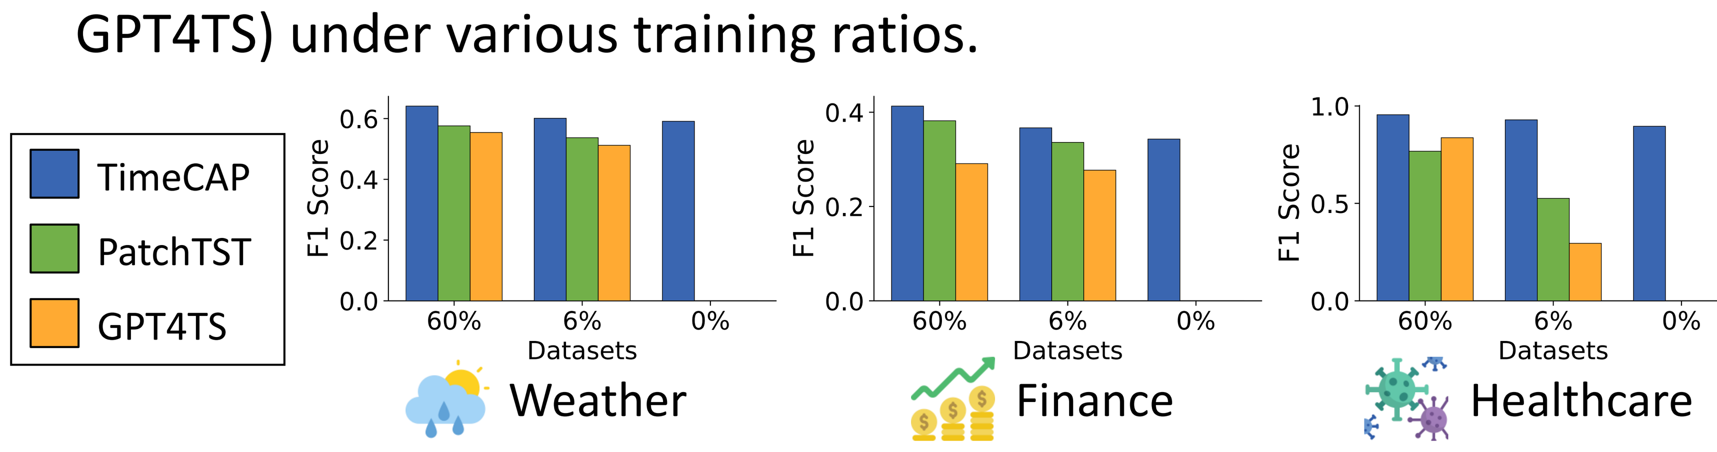

- **Effectiveness** Three key steps of **TimeCAP** are effective.

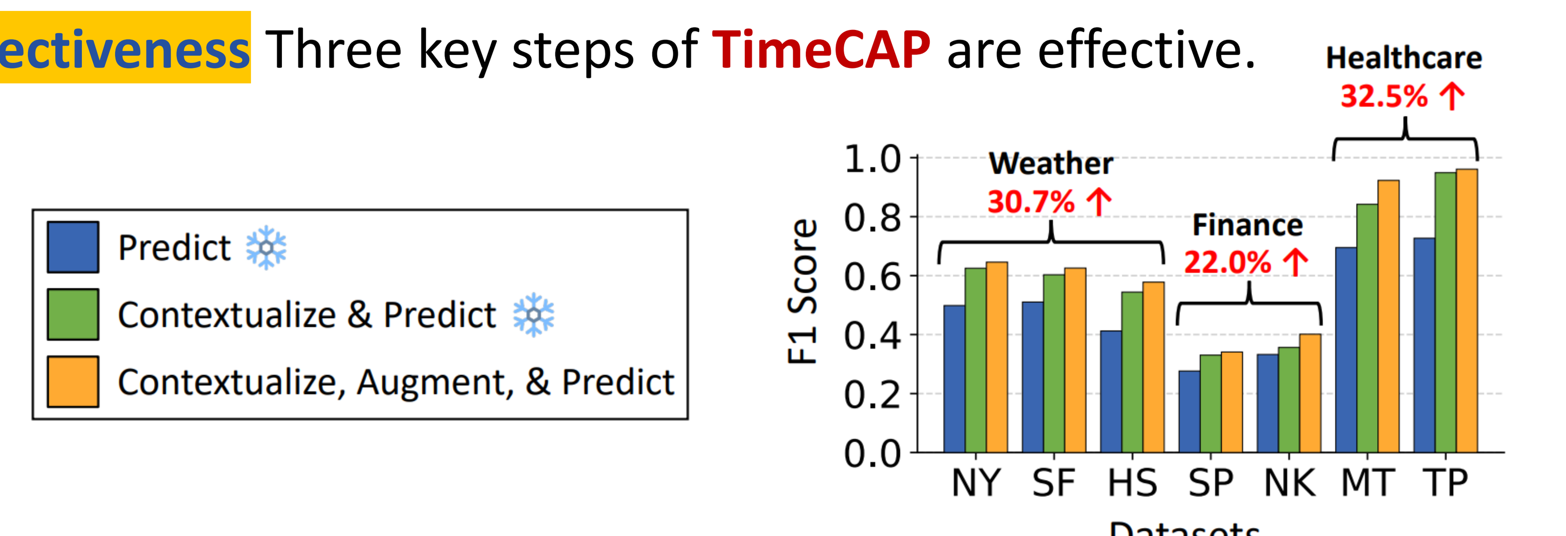

- **Interpretability** **TimeCAP** provides implicit and explicit interpretations.

Text Summary (June 26th, 2017) **Answer: Rain**

Over the past 24 hours, New York City has experienced a gradual decrease in temperature followed by a moderate increase, indicating a typical diurnal variation with cooler conditions overnight and warmer conditions during the day. Humidity levels showed a general upward trend throughout the night, reaching a peak before starting to decrease, which could be indicative of a passing weather front. Air pressure remained relatively stable with only slight fluctuations, suggesting a steady state of the atmosphere without any significant weather systems moving in or out of the area. Wind speeds varied, with some calmer periods and occasional gusts, but no sustained high winds were observed, which would have been noteworthy for weather impacts. Wind direction shifts were not consistent, indicating the presence of variable winds rather than a dominant weather pattern influencing the direction over the period.

Explicit Interpretation (September 4th, 2014) **Prediction: Rain**

Over the past 24 hours in New York City, temperatures have experienced a notable increase, indicative of a warming trend that peaked in the late afternoon before gradually declining toward the evening. Humidity levels have shown fluctuations, starting relatively low, increasing significantly during the early hours, and then decreasing during the day, which could be associated with a front passing through the area. Air pressure remained relatively stable throughout the period, with only slight variations, suggesting a period of settled weather. Wind speeds varied modestly, with calmer conditions prevailing for most of the day before a slight increase later on, while wind direction shifted from primarily westerly to more variable, including southerly and easterly directions, which may influence the transport of air masses and possibly lead to changes in weather patterns.