Retrieval of Temporal Event Sequences from Textual Descriptions

Zefang Liu, Yinzhu Quan (Georgia Institute of Technology)
Contact: {liuzefang, yquan9}@gatech.edu

Introduction

- **Temporal event sequences** encode both event types and timestamps, essential for applications like e-commerce analysis, social media monitoring, and crime tracking.
- Retrieving event sequences from **natural language descriptions** is challenging due to the need to capture both **event semantics** and **temporal dynamics**.
- We propose **TPP-Embedding**, a model that jointly encodes temporal and textual information to enable accurate **Temporal Event Sequence Retrieval (TESR)**.

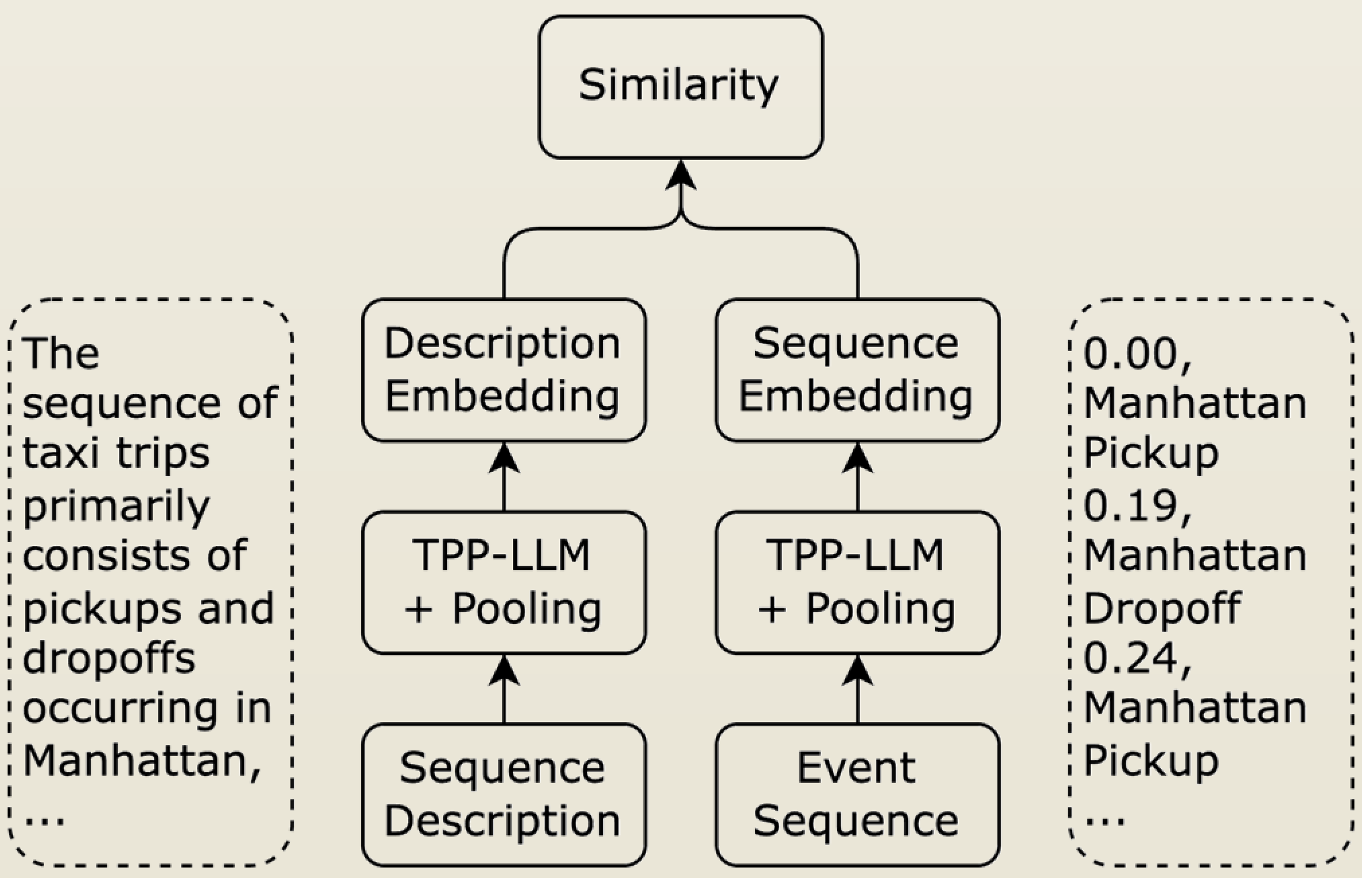

TESRBench

We introduce **TESRBench**, a benchmark for evaluating **TESR**:
- Includes **5 real-world datasets** from diverse domains.
- Provides **synthesized textual descriptions** for event sequences, generated by **GPT-4o-mini**.

Dataset	# of Types	# of Events	# of Seq.	Time Unit
Stack Overflow	25	187,836	3,336	Month
Chicago Crime	20	202,333	4,033	Month
NYC Taxi Trip	8	362,374	2,957	Hour
U.S. Earthquake	3	29,521	3,009	Day
Amazon Review	18	127,054	2,245	Week

Experiments

We evaluate **TPP-Embedding** on **TESRBench** against strong text-based baselines.

Setup:
- **Foundation models:** TinyLlama-1.1B, TinyLlama-1.1B-Chat.
- **Baselines:** MiniLM-L12, MPNet-Base, BGE-Large, MxbAI-Large, mE5-Large, Qwen2-1.5B.
- **Metrics:** Mean Reciprocal Rank (MRR) and Recall@K.

Results:
- **TPP-Embedding** consistently outperforms baselines across most datasets.
- Demonstrates strong generalization in a **multi-domain setting**.

TPP-Embedding

We propose **TPP-Embedding**, a model for **TESR** that jointly encodes temporal and textual information.

Key Features:
- Integrates **Temporal Point Process (TPP)** models with **LLMs** to capture event semantics and temporal dynamics.
- Embeds event sequences and textual descriptions in a **shared embedding space** using temporal encoding and LLM embeddings.
- Trains with **contrastive loss** to align matching sequence-description pairs.
- Fine-tuned efficiently with **QLoRA** for scalable deployment.

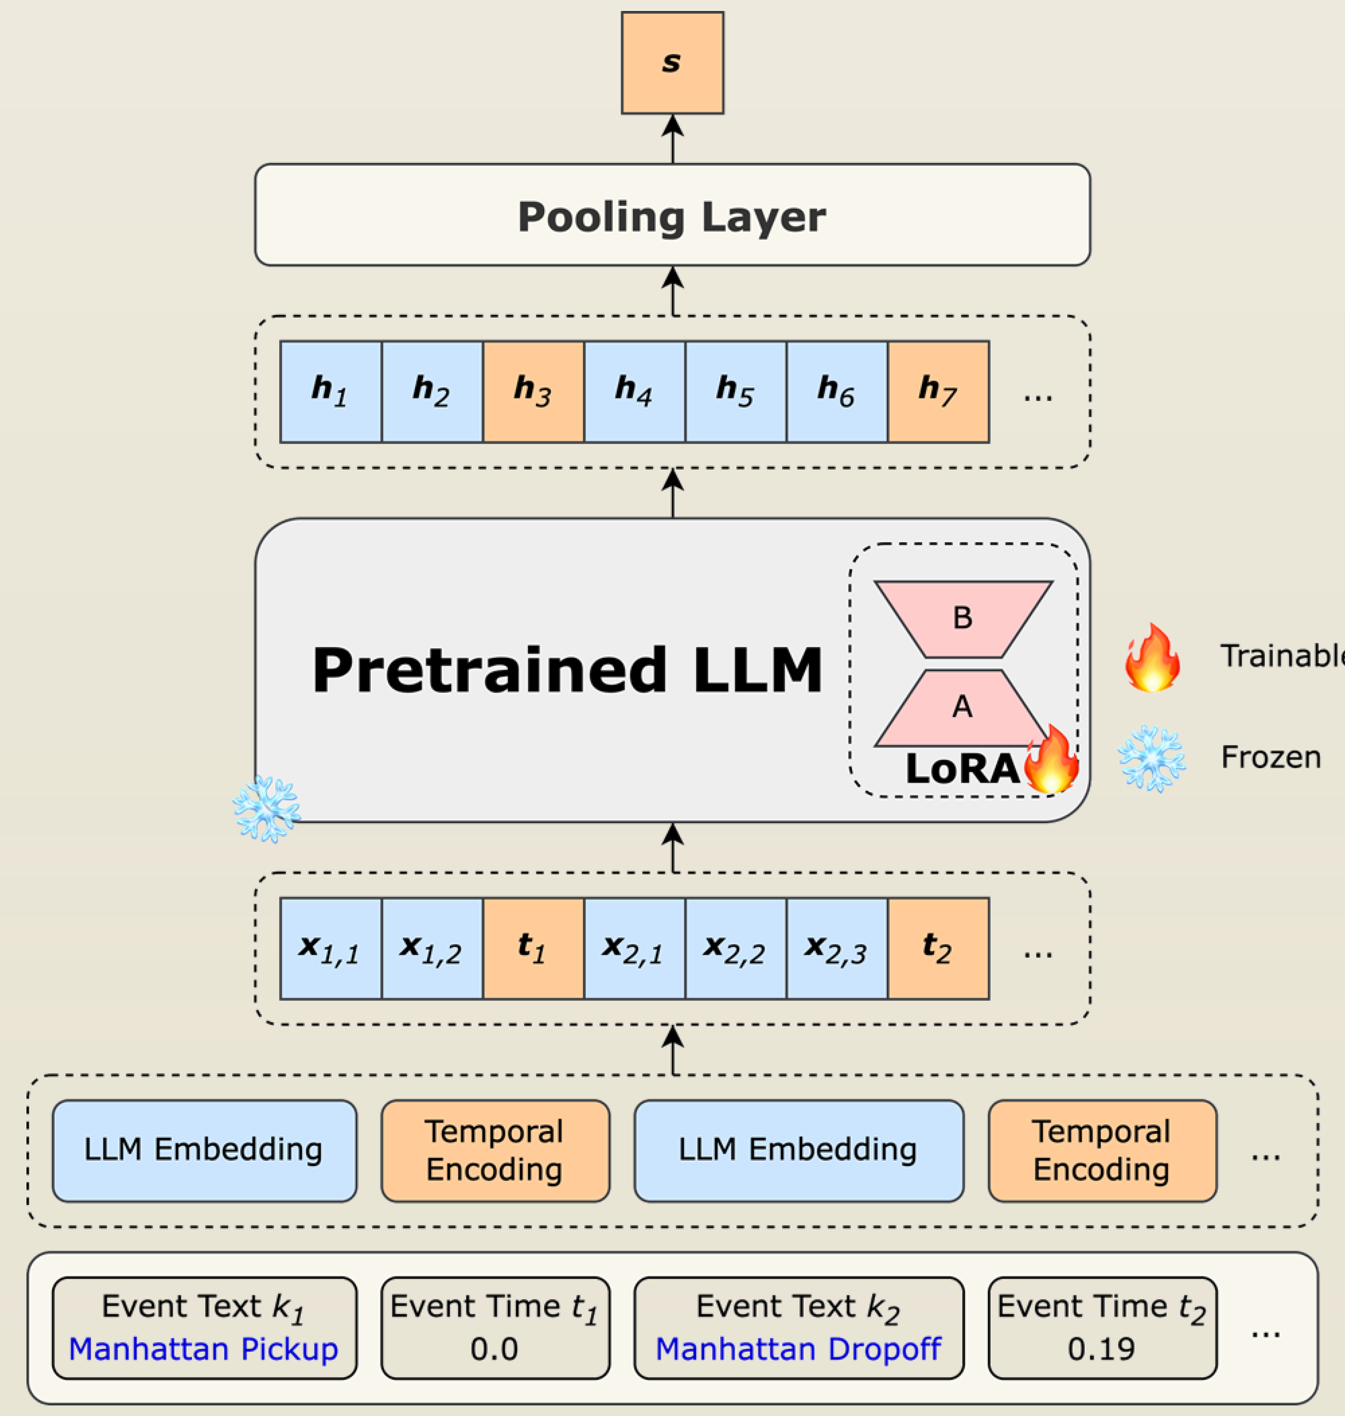

Conclusion

- We introduce **TESRBench**, a benchmark suite for evaluating **temporal event sequence retrieval (TESR)**.
- We propose **TPP-Embedding**, a retrieval model that integrates temporal information and event semantics.
- **TPP-Embedding** achieves state-of-the-art performance on **TESRBench** and generalizes well across domains.

References

- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings Using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3982–3992.
- Vinayak Gupta, Srikanta Bedathur, and Abir De. 2022. Learning Temporal Point Processes for Efficient Retrieval of Continuous Time Event Sequences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 4005–4013.
- Zefang Liu and Yinzhu Quan. 2024. TPP-LLM: Modeling Temporal Point Processes by Efficiently Fine-Tuning Large Language Models. *arXiv preprint arXiv:2410.02062*.

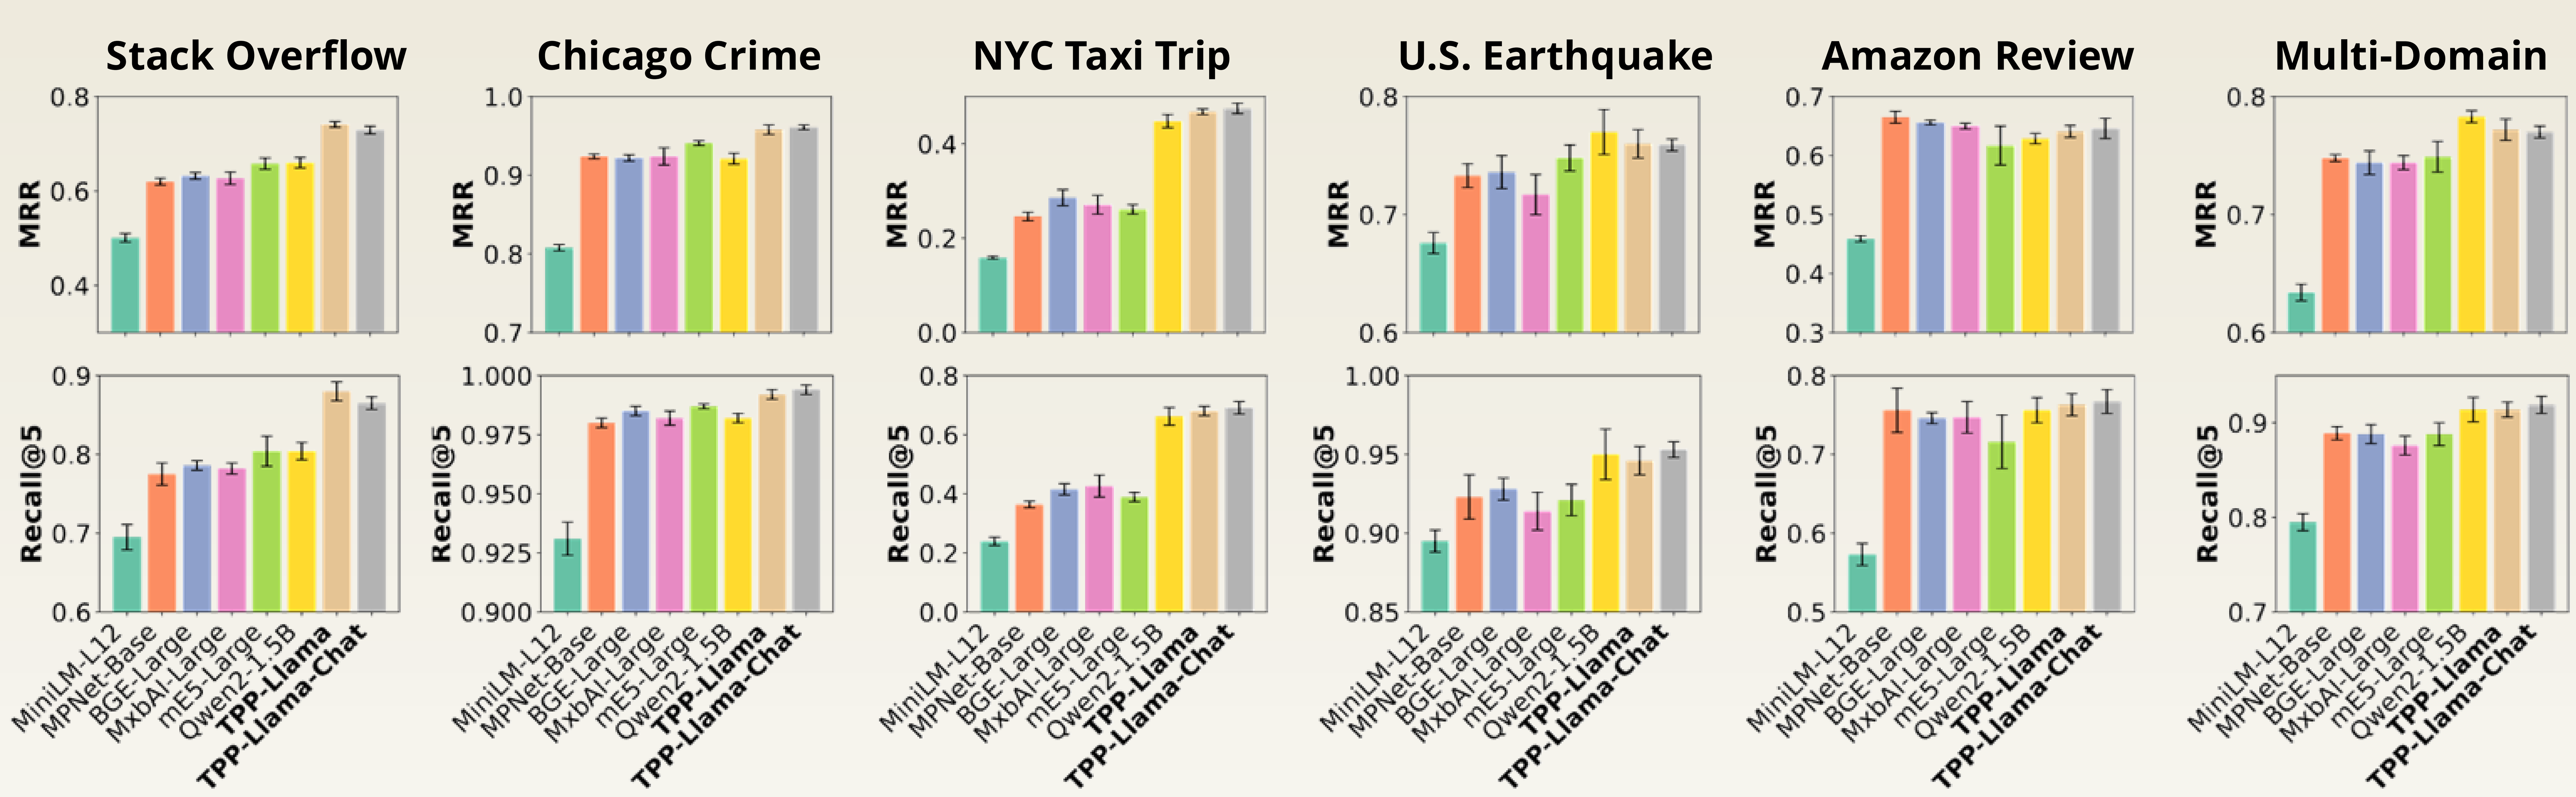

Code

Data