Large Language Models for Spatial Analysis Queries

Mohamed F. Mokbel[1], **Youssef Hussein**[1], **Mohamed Hemdan**[1]

[1]University of Minnesota - USA

mokbel@umn.edu

Abstract. *This tutorial provides a comprehensive overview of the research landscape of employing Large Language Models (LLMs) to spatial analysis queries. The tutorial categorizes the research in this area based on how LLMs are employed to serve such queries. This goes from employing LLMs as is, to fine-tuning LLMs, to completely retrain LLM architectures, to modifying the LLM internals to fit spatial queries. The tutorial concludes by a set of benchmarks and pointing out to research gaps and future research directions.*
Key-words: *Large Language Models (LLMs), Spatial Analysis Queries, Spatial Query Processing, Spatial Information Retrieval*

Resumo. *Este tutorial oferece uma visão abrangente do panorama de pesquisas sobre o uso de Grandes Modelos de Linguagem (LLMs) em consultas de análise espacial. O tutorial classifica as pesquisas nessa área com base na forma como os LLMs são empregados para atender a tais consultas, abrangendo desde o uso de LLMs em seu estado original e o ajuste fino ("fine-tuning") até o retreinamento completo de arquiteturas de LLM e a modificação de seus componentes internos para adequá-los a consultas espaciais. O tutorial é encerrado com a apresentação de benchmarks e a indicação de lacunas na pesquisa e de direções para estudos futuros.*
Palavras-chave: *Modelos de Linguagem (LLMs), Consultas de análise espacial, Processamento de consultas espaciais, Recuperação de informação espacial*

1. Introduction

Spatial analysis queries that aim to provide insights and enable efficient retrieval of large spatial data have been an active research area in the database community over the last few decades. Such queries have been a cornerstone of enabling a wide range of applications, including location-based services, trajectory data analysis, transportation, and urban planning. Meanwhile, the recent tremendous success and wide deployment of Large Language Models (LLMs), e.g., BERT, GPT, and Llama, as large-scale deep learning architectures trained on vast amounts of data, have prompted the database community to explore the use of LLMs in various data management problems.

This tutorial provides a comprehensive overview of the research landscape of the rapidly evolving area of employing LLMs to support spatial analysis queries. We analyze existing work according to the analysis tasks and how LLMs are leveraged to support them. We discuss four spatial analysis tasks that have exploited the usage of LLMs: (1) traffic-related queries, (2) multi-modal queries, (3) trajectory-related queries, and (4) PoI-related queries. We classify existing work according to how LLMs are leveraged to support for these tasks. In the first category, LLMs are used as is to support spatial analysis tasks. The second category goes one step further and fine-tunes LLMs to fit spatial analysis. The third category goes one step further by using an LLM solely as an architecture, and completely retraining it from scratch with spatial data. The fourth category does the same as the third, but in addition, the techniques in this category modify the model internals to fit various spatial analysis tasks. We also cover various evaluation and benchmarking papers, while pointing out future research directions.

A previous version of this tutorial was presented at VLDB 2025. Further details on the extensive list of related work examined can be found in [Hussein et al. 2025].

2. Tutorial Outline

2.1. Part 1: LLMs and Spatial Queries: Why?

This part of the tutorial starts with a necessary background about LLMs and their architectures. It also presents how LLMs fit under the wider umbrella of Foundation Models (FMs) along with various visions that laid out the foundation and needs for exploiting LLMs and FMs for various spatial analysis queries. This part also explains the rationale of categorizing existing work based on their underlying LLMs and the tasks they aim to support. This would require going a bit deeper into LLMs architecture as it will help in explaining the differences among the various categories that will be discussed later.

2.2. Part 2: Using an LLM As Is

This part of the tutorial covers techniques that use pre-trained LLMs as is without altering any of its trained data, parameters, or loss function. Yet, the techniques in this category still involve crafting several iterations of input prompts to guide the LLMs towards specific output types that fit spatial analysis tasks; a process known as prompt-engineering. Though such category of techniques is somehow straightforward in terms of using LLMs as is, it was shown that it is still powerful for many spatial analysis tasks. Examples of work in this category that will be covered in this tutorial is designed for traffic-related queries such as traffic analysis, traffic signal control, and traffic and trip prediction and for trajectory-related queries as trajectory generation.

2.3. Part 3: Fine-Tuning an LLM

As with the previous category, techniques in this category start from an already trained LLM. Yet, techniques in this category go one step further by modifying various parameters and weights in the LLMs internals to make them better fit for various spatial analysis tasks; a process known as fine-tuning. Such supervised fine-tuning process allows the employed LLM to learn task-specific patterns, and hence significantly enhances its performance. Apparently, this is a bit more complicated than just using LLMs as is, and hence it gives more power in boosting the performance of various spatial analysis tasks. Examples of work in this category that will be covered in this tutorial are either designed for traffic-related queries as traffic signal control, and traffic and trip prediction, in multimodal queries as street navigation, in trajectory-related queries as trajectory generation and anomaly detection, or in PoI-related queries as next PoI recommendation and generic PoI models.

2.4. Part 4: Using Architecture-Only LLMs

Unlike the previous two categories, these techniques do not employ any trained models. Instead, they use the vanilla architecture of LLMs without any training. Then, they completely train the employed architecture using some form of spatial data. This makes the employed models fully trained on spatial data, and hence they become more suitable to spatial analysis tasks. As this category is more complicated than the previous category, where retraining an LLM form scratch would also require fine-tuning, techniques in this category are able to natively support spatial analysis tasks as they are trained on spatial data, and hence results in a much higher accuracy. Examples of existing work in this category that will be covered in this tutorial are either designed for traffic related queries as traffic and trip prediction, trajectory-related queries as trajectory imputation and anomaly detection, or in PoI-related queries as generic PoI models.

2.5. Part 5: New Loss Function

In addition to only using the LLM architecture as in the previous category, techniques in this category go one step further and modify the loss function of the LLM architecture to fit spatial analysis tasks. LLMs employ a loss function to assess the difference between their actual output and either the target output or the ground truth. Modifying such a loss function gives the techniques in this category an edge in adapting the behavior of LLMs to fit spatial applications. This is in addition to adapting the prompt, parameters, and data, which took place in our first, second, and third categories, respectively. Apparently, this provides more native support for spatial applications, and hence better control over the efficiency and accuracy of such applications. Examples of existing work in this category covered in this tutorial is either designed for traffic-related queries as traffic and trip prediction, multi-modal queries as urban profiling, trajectory related queries as anomaly detection and generic trajectory models, or PoI-related queries as query PoI matching.

2.6. Part 6: Benchmarking and Open Problems

This part of the tutorial discusses the work that benchmarks the use of LLMs in various spatial analysis tasks, which include: (a) benchmarking LLMs for path planning, time-series data generation, and geospatial knowledge, (b) evaluating LLMs ability for generating code in geospatial applications, representing spatial relations and geometries as text, and supporting spatial queries in dynamic graphs, and (c) evaluating LLMs' bias in spatial domain and its impact to various tasks and LLMs representation of space and time in various spatial datasets. The tutorial concludes by pointing out to several research gaps, open problems, and future directions in the research landscape of LLMs for spatial analysis queries, which will help in guiding the future research agenda for the VLDB community.

3. Biography Sketches

Mohamed F. Mokbel is a Distinguished McKnight University Professor at the University of Minnesota. He is a recipient of the ACM SIGSPATIAL 10-Year Impact Award, the IEEE ICDE 10-Year Influential Paper Award, and the VLDB 10-Year Best Paper Award. He is the Editor-in-Chief for ACM Transactions on Spatial Algorithms and Systems (TSAS) and a past elected Chair of ACM SIGPATIAL. Mohamed is an IEEE Fellow and ACM Distinguished Scientist.

Youssef Hussein is a PhD student at University of Minnesota. He has a bachelor's degree from the American University in Cairo with Highest Honors, with a year-long internship at Dell Technologies. His research interests include large language models and their use for spatiotemporal applications.

Mohamed Hemdan is a PhD student at University of Minnesota.He has a bachelors degree in Computer Engineering from The American University in Cairo with Highest Honors, with internships at Microsoft and CERN. His research interests include spatial data management and spatial data science.

Acknowlegments: This work is supported by NSF grants IIS-2203553 and OAC-2118285.

Referências

Hussein, Y., Hemdan, M., and Mokbel, M. F. (2025). Large language models for spatial analysis queries. *Proc. VLDB Endow.*, 18(12):5451–5454.