Large Language Models for Spatial Analysis Queries

Youssef Hussein
University of Minnesota, USA
husse408@umn.edu

Mohamed Hemdan
University of Minnesota, USA
hemda001@umn.edu

Mohamed F. Mokbel
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.

PVLDB Reference Format:
Youssef Hussein, Mohamed Hemdan, and Mohamed F. Mokbel. Large Language Models for Spatial Analysis Queries. PVLDB, 18(12): 5451 - 5454, 2025.
doi:10.14778/3750601.3750693

1 INTRODUCTION

Spatial analysis queries aiming to provide insights and efficient retrieval of large spatial data have been an active research area in the database community over the last few decades [3, 8, 10, 59, 74]. Such queries have been a cornerstone for enabling a wide range of widely used 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 [16], GPT [58], and Llama [63], as large-scale deep learning architectures trained on vast amounts of data, have triggered the database community to exploit the use of LLMs in various data management problems [13, 18, 20, 54, 56, 57, 65].

This tutorial provides a comprehensive overview of the research landscape of the rapidly evolving area of employing Large Language Models (LLMs) to support spatial analysis queries. Figure 1 gives an overview of the approaches that will be covered in this tutorial. The vertical axis of the figure represents spatial analysis tasks that have exploited the usage of LLMs. These tasks are grouped into four subcategories based on the type of spatial query they address: (1) *traffic-related queries*, (2) *multi-modal queries*, (3) *trajectory-related queries*, and (4) *PoI-related queries*. The horizontal axis of the figure categorizes existing work into four categories based on their use of LLMs. 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 retrain it from scratch with spatial data. The fourth category does the same as the third category, yet in addition, the techniques in this category change the model internals to fit various spatial analysis tasks. Each approach in Figure 1 is labeled by the LLM model it is based on.

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

Proceedings of the VLDB Endowment, Vol. 18, No. 12 ISSN 2150-8097.
doi:10.14778/3750601.3750693

2 TUTORIAL OUTLINE

Figure 2 gives the detailed outline and timing of this 90-minutes Tutorial. The tutorial is mainly composed of six parts. The first part serves as a background and introduction. The next four parts correspond to the four categories in the horizontal axis of Figure 1. The last part concludes the tutorial by covering various evaluation and benchmarking papers, while pointing out to future research directions. This section briefly discusses the contents of each part.

2.1 Part 1: LLMs and Spatial Queries: Why?

This part of the tutorial starts with a necessary background about Large Language Models (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 [5, 14, 43, 44, 52, 53]. This part will also explain Figure 1 in details pointing out to 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 the techniques in the third column of Figure 1. Such techniques use pre-trained Large Language Models (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 [67, 73], traffic signal control [15, 42, 75], and traffic and trip prediction [38] and for trajectory-related queries as trajectory generation [66].

2.3 Part 3: Fine-Tuning an LLM

This part of the tutorial covers the techniques that appear in the fourth column of Figure 1. 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

		Using an LLM As Is	Fine-Tuning An LLM	Using Architecture-Only LLMs	New Loss Function
Traffic Related Queries	Traffic Analysis	♣TransGPT [67] ♣Trafficgpt [73]			
	Traffic Signal Control	♣Open-TI [15] ♣GeoGPT[75] ♣PromptGAT [42]	♣LLMLight[30]		
	Traffic & Trip Prediction	♣LLM-MPE [38]	★♣AuxMobLCast [68] □LLM-COD [70] ♣UrbanGPT [36]	◇PromptST [77]	♣ST-LLM [40] ◇ UniST [71] ★BERT-Trip [2] ★MultiCast [11]
Multi-Modal Queries	Street Navigation		♣VELMA [60]		
	Urban Region Profiling				□UrbanCLIP [69]
Trajectory Related Queries	Trajectory Generation	♣LLMob [66]	□Geo-Llama [34] ♣UMA-LM [7]		◇TrajGPT [25]
	Trajectory Imputation			★KAMEL [50, 51] ◇RNTrajRec [12] ★TrajBERT [62]	
	Anomaly Detection		□CLARA [29]	♣LM-TAD [47]	◇TranAD [61] ◇GADFormer [41]
	Generic Trajectory Models				◇STPT [26] ◇CTLE [39]
POI Related Queries	Next PoI Recommendation		□LLM4POI [33]		
	Query-POI Matching				♣MGeo [17] ◇FIR-PT [48]
	Generic PoI Models		□ UrbanKGent [55]	★SpaBERT [35] ★GeoBERT [72]	◇ERNIE-GeoL [27] ★ CityFM [5] ★GeoLM [37]

Base Model: ◇ Transformer ♣ GPT ★ BERT □ Llama

Figure 1: Overview of Generic and Task-Specific Models in Spatial Analysis Queries

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 [30], and traffic and trip prediction [36, 68, 70], in multi-modal queries as street navigation [60], in trajectory-related queries as trajectory generation [7, 34] and anomaly detection [29], or in PoI-related queries as next PoI recommendation [33] and generic PoI models [55].

2.4 Part 4: Using Architecture-Only LLMs

This part covers the techniques that appear in the fifth column of Figure 1. 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 [77], trajectory-related queries as trajectory imputation [12, 51, 62] and anomaly detection [47], or in PoI-related queries as generic PoI models [35, 72].

2.5 Part 5: New Loss Function

This part of the tutorial covers the techniques that appear in the last column of Figure 1. 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 loss function gives the techniques in this category an edge in adapting the behaviour 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 on 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 [2, 11, 40, 71],, multi-modal queries as urban profiling [69], trajectory-related queries as anomaly detection [41, 61] and generic trajectory models [26, 39], or PoI-related queries as query PoI matching [17, 48] and generic PoI models [5, 27, 37].

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 [1], time-series data generation [4], and geospatial knowledge [6], (b) evaluating LLMs ability for generating code in geospatial applications [22], representing spatial relations and geometries as text [28], and supporting spatial queries in dynamic graphs [76], and (c) evaluating LLMs' bias in spatial domain and its impact to various tasks [45, 46] and LLMs representation of space and time in various spatial datasets [24]. 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.

Part 1: LLMs and Spatial Data: why? (20 minutes)
- Introduction to LLMs and FMs (6 minutes)
- The Vision of LLM and FM for Spatial Analysis Queries (7 minutes)
- Explaining Figure 1 (7 minutes)

Part 2: Using an LLMs As Is (15 minutes)
- What does this category mean (5 minutes)
- Employing LLMs as is for traffic queries (5 minutes)
- Employing LLMs as is for trajectory queries (5 minutes)

Part 3: Fine-Tuning an LLM (15 minutes)
- What does this category mean (5 minutes)
- Fine tuning LLMs for traffic queries and multi-modal queries (5 minutes)
- Fine tuning LLMs for trajectory and PoI queries(5 minutes)

Part 4: Using Architecture-Only LLMs (15 minutes)
- What does this category mean (5 minutes)
- Architecture-Only LLMs for traffic queries (5 minutes)
- Architecture-Only LLMs for trajectory and PoI queries (5 minutes)

Part 5: New Loss Function (15 minutes)
- What does this category mean (5 minutes)
- New Loss Function for traffic and multi-modal queries (5 minutes)
- New Loss Function for trajectory and PoI queries (5 minutes)

Part 6: Benchmarking, Evaluation, Open Problems (10 minutes)
- Benchmarking and Evaluating LLMs for Spatial Analysis Queries (5 minutes)
- Research Gaps and Open Problems (5 minutes)

Figure 2: Outline & Timing of the 90-minute Tutorial.

3 TARGET AUDIENCE AND BACKGROUND

This tutorial targets researchers, developers, and practitioners, who are interested in understanding and assessing the impact of Large Language Models on spatial analysis queries. The tutorial will have enough introductory material to ensure that the audience get the necessary background needed to understand its contents, hence no prior knowledge is required to understand the techniques and approaches presented in this tutorial. The tutorial will also be very beneficial for graduate students as it will help them in identifying various challenges for PhD topics. Practitioners will get to know the state-of-the-art techniques to exploit the use of LLMs in spatial analysis queries. Finally, the tutorial will act as an invitation to the data management community to exploit the rapidly evolving field of Large Language Models as a means of boosting the accuracy and efficiency of spatial analysis queries.

4 RELATED RECENT TUTORIALS

Over the last three years, there were five tutorials presented in VLDB and SIGMOD about LLMs [21, 23, 31, 49, 64]. These tutorials have either discussed the role that LLMs can play in boosting the performance of various database components or the role that database systems can do in scaling up the computations of LLMs. None of these five tutorials have discussed the topic of spatial analysis queries. Meanwhile, over the last three years, there were four tutorials in VLDB and SIGMOD addressing various challenges and innovations in spatial data analysis [9, 19, 32, 74], however, none of them discussed LLM topics. Our tutorial will be the first to bridge the two areas of LLMs and spatial data analysis.

5 BIOGRAPHICAL SKETCHES

Youssef Hussein is a PhD student at University of Minnesota. He has a bachelors 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.
Mohamed F. Mokbel is a Distinguished McKnight University Professor at 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.

REFERENCES

[1] Mohamed Aghzal, Erion Plaku, and Ziyu Yao. 2024. Can Large Language Models Be Good Path Planners? A Benchmark and Investigation on Spatial-Temporal Reasoning. In *ICLR*.
[2] Kuo Ai-Te, Chen Haiquan, and Ku Wei-Shinn. 2023. BERT-Trip: Effective and Scalable Trip Representation using Attentive Contrast Learning. In *ICDE*.
[3] A. B. Siddique andAhmed Eldawy and Vagelis Hristidis. 2019. Comparing Synopsis Techniques for Approximate Spatial Data Analysis. *PVLDB* (2019).
[4] Yihao Ang, Yifan Bao, Qiang Huang, Anthony K. H. Tung, and Zhiyong Huang. 2024. TSGAssist: An Interactive Assistant Harnessing LLMs and RAG for Time Series Generation Recommendations and Benchmarking. *VLDB J.* (2024).
[5] Pasquale Balsebre, Weiming Huang, Gao Cong, and Yi Li. 2024. City Foundation Models for Learning General Purpose Representations from OpenStreetMap. In *CIKM*.
[6] Prabin Bhandari, Antonios Anastasopoulos, and Dieter Pfoser. 2023. Are Large Language Models Geospatially Knowledgeable?. In *SIGSPATIAL*.
[7] Prabin Bhandari, Antonios Anastasopoulos, and Dieter Pfoser. 2024. Urban Mobility Assessment Using LLMs. In *SIGSPATIAL*.
[8] Tsz Nam Chan, Pak Lon Ip, Leong Hou U, Weng Hou Tong, Shivansh Mittal, Ye Li, and Reynold Cheng. 2021. KDV-Explorer: A Near Real-Time Kernel Density Visualization System for Spatial Analysis. *PVLDB* (2021).
[9] Tsz Nam Chan, Leong Hou U, Byron Choi, Jianliang Xu, and Reynold Cheng. 2023. Large-scale Geospatial Analytics: Problems, Challenges, and Opportunities (Tutorial). In *SIGMOD*.
[10] Tsz Nam Chan, Leong Hou U, Yun Peng, Byron Choi, and Jianliang Xu. 2022. Fast Network K-function-based Spatial Analysis. *PVLDB* (2022).
[11] Georgios Chatzigeorgakidis, Konstantinos Lentzos, and Dimitrios Skoutas. 2024. MultiCast: Zero-Shot Multivariate Time Series Forecasting Using LLMs. In *ICDE*.
[12] Yuqi Chen, Hanyuan Zhang, Weiwei Sun, and Baihua Zheng. 2023. RNTrajRec: Road Network Enhanced Trajectory Recovery with Spatial-Temporal Transformer. In *ICDE*.
[13] Zui Chen, Lei Cao, and Sam Madden. 2023. Lingua Manga: A Generic Large Language Model Centric System for Data Curation. *PVLDB* (2023).
[14] Shushman Choudhury, Abdul Rahman Kreidieh, Ivan Kuznetsov, and Neha Arora. 2024. Towards a Trajectory-powered Foundation Model of Mobility. In *SIGSPATIAL*.
[15] Longchao Da, Kuanru Liou, Tiejin Chen, Xuesong Zhou, Xiangyong Luo, Yezhou Yang, and Hua Wei. 2024. Open-TI: Open Traffic Intelligence with Augmented Language Model. *ICMLC* (2024).
[16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*.
[17] Ruixue Ding, Boli Chen, Pengjun Xie, Fei Huang, Xin Li, Qiang Zhang, and Yao Xu. 2023. MGeo: Multi-Modal Geographic Language Model Pre-Training. In *SIGIR*.
[18] Raul Castro Fernandez, Aaron J. Elmore, Michael J. Franklin, Sanjay Krishnan, and Chenhao Tan. 2023. How Large Language Models Will Disrupt Data Management. *PVLDB* (2023).
[19] Raul Castro Fernandez and Arnab Nandi. 2024. Responsible Sharing of Spatiotemporal Data (Tutorial). In *SIGMOD*.

[20] Benjamin Feuer, Yurong Liu, Chinmay Hegde, and Juliana Freire. 2024. ArcheType: A Novel Framework for Open-Source Column Type Annotation Using Large Language Models. *PVLDB* (2024).

[21] Badaro Gilbert and Papotti Paolo. 2022. Transformers for Tabular Data Representation: A Tutorial on Models and Applications (Tutorial). *PVLDB* (2022).

[22] Piotr Gramacki, Bruno Martins, and Piotr Szymanski. 2024. Evaluation of Code LLMs on Geospatial Code Generation. In *SIGSPATIAL*.

[23] Li Guoliang, Zhou Xuanhe, and Zhao Xinyang. 2024. LLM for Data Management (Tutorial). *PVLDB* (2024).

[24] Wes Gurnee and Max Tegmark. 2024. Language Models Represent Space and Time. In *ICLR*.

[25] Shang-Ling Hsu, Emmanuel Tung, John Krumm, Cyrus Shahabi, and Khurram Shafique. 2024. TrajGPT: Controlled Synthetic Trajectory Generation Using a Multitask Transformer-Based Spatiotemporal Model. In *SIGSPATIAL*.

[26] Mingzhi Hu, Zhuoyun Zhong, Xin Zhang, Yanhua Li, Yiqun Xie, Xiaowei Jia, Xun Zhou, and Jun Luo. 2023. Self-supervised Pre-training for Robust and Generic Spatial-Temporal Representations. In *ICDM*.

[27] Jizhou Huang, Haifeng Wang, Yibo Sun, Yunsheng Shi, Zhengjie Huang, An Zhuo, and Shikun Feng. 2022. ERNIE-GeoL: A Geography-and-Language Pre-trained Model and its Applications in Baidu Maps. In *SIGKDD*.

[28] Yuhan Ji and Song Gao. 2023. Evaluating the Effectiveness of Large Language Models in Representing Textual Descriptions of Geometry and Spatial Relations. In *GIScience*.

[29] Chan Young Koh, Kyle DeMedeiros, and Abdeltawab Hendawi. 2025. CLARA: Context-aware RAG-LLM Framework for Anomaly Detection in Mobile Device Sensors. In *MDM*.

[30] Siqi Lai, Zhao Xu, Weijia Zhang, Hao Liu, and Hui Xiong. 2025. LLMLight: Large Language Models as Traffic Signal Control Agents. In *SIGKDD*.

[31] Guoliang Li, Jiayi Wang, Chenyang Zhang, and Jiannan Wang (Tutorial). 2025. Data+AI: LLM4Data and Data4LLM. In *SIGMOD*.

[32] Huan Li, Bo Tang, Hua Lu, Muhammad Aamir Cheema, and Christian S. Jensen. 2022. Spatial Data Quality in the IoT Era: Management and Exploitation (Tutorial). In *SIGMOD*.

[33] Peibo Li, Maarten de Rijke, Hao Xue, Shuang Ao, Yang Song, and Flora D. Salim. 2024. Large Language Models for Next Point-of-Interest Recommendation. In *SIGIR*.

[34] Siyu Li, Toan Tran, Haowen Lin, John Krumm, Cyrus Shahabi, and Li Xiong. 2025. Geo-Llama: Leveraging LLMs for Human Mobility Trajectory Generation with Spatiotemporal Constraints. In *MDM*.

[35] Zekun Li, Jina Kim, Yao-Yi Chiang, and Muhao Chen. 2022. SpaBERT: A Pretrained Language Model from Geographic Data for Geo-Entity Representation. In *EMNLP*.

[36] Zhonghang Li, Lianghao Xia, Jiabin Tang, Yong Xu, Lei Shi, Long Xia, Dawei Yin, and Chao Huang. 2024. UrbanGPT: Spatio-Temporal Large Language Models. In *SIGKDD*.

[37] Zekun Li, Wenxuan Zhou, Yao-Yi Chiang, and Muhao Chen. 2023. GeoLM: Empowering Language Models for Geospatially Grounded Language Understanding. In *EMNLP*.

[38] Yuebing Liang, Yichao Liu, Xiaohan Wang, and Zhan Zhao. 2024. Exploring Large Language Models for Human Mobility Prediction under Public Events. *Elsevier Computers, Environment and Urban Systems* (2024).

[39] Yan Lin, Huaiyu Wan, Shengnan Guo, and Youfang Lin. 2021. Pre-training Context and Time Aware Location Embeddings from Spatial-Temporal Trajectories for User Next Location Prediction. In *AAAI*.

[40] Chenxi Liu, Sun Yang, Qianxiong Xu, Zhishuai Li, Cheng Long, Ziyue Li, and Rui Zhao. 2024. Spatial-Temporal Large Language Model for Traffic Prediction. In *MDM*.

[41] Andreas Lohrer, Darpan Malik, Claudius Zelenka, and Peer Kröger. 2024. GADformer: A Transparent Transformer Model for Group Anomaly Detection on Trajectories. In *IJCNN*.

[42] Minchiuan Gao Longchao Da, Hao Mei, and Hua Wei. 2024. Prompt to Transfer: Sim-to-Real Transfer for Traffic Signal Control with Prompt Learning.. In *AAAI*.

[43] Gengchen Mai, Chris Cundy, Kristy Choi, Yingjie Hu, Ni Lao, and Stefano Ermon. 2022. Towards a Foundation Model for Geospatial Artificial Intelligence (Vision Paper). In *SIGSPATIAL*.

[44] Gengchen Mai, Weiming Huang, Jin Sun, Suhang Song, Deepak Mishra, Ninghao Liu, Song Gao, Tianming Liu, Gao Cong, Yingjie Hu, Chris Cundy, Ziyuan Li, Rui Zhu, and Ni Lao. 2024. On the Opportunities and Challenges of Foundation Models for GeoAI (Vision Paper). *TSAS* (2024).

[45] Rohin Manvi, Samar Khanna, Marshall Burke, David B. Lobell, and Stefano Ermon. 2024. Large Language Models are Geographically Biased. In *ICML*.

[46] Rohin Manvi, Samar Khanna, Gengchen Mai, Marshall Burke, David Lobell, and Stefano Ermon. 2024. Geollm: Extracting Geospatial Knowledge from Large Language Models. In *ICLR*.

[47] Jonathan Mbuya, Dieter Pfoser, and Antonios Anastasopoulos. 2024. Trajectory Anomaly Detection with Language Models. In *SIGSPATIAL*.

[48] Lang Mei, Jiaxin Mao, Juan Hu, Naiqiang Tan, Hua Chai, and Ji-Rong Wen. 2023. Improving First-stage Retrieval of Point-of-interest Search by Pre-training Models. *ACM Transactions on Information Systems* (2023).

[49] Xupeng Miao, Zhihao Jia, and Bin Cui. 2024. Demystifying Data Management for Large Language Models (Tutorial). In *SIGMOD*.

[50] Mashaal Musleh and Mohamed F. Mokbel. 2023. A Demonstration of KAMEL: A Scalable BERT-based System for Trajectory Imputation. In *SIGMOD*.

[51] Mashaal Musleh and Mohamed F. Mokbel. 2023. Kamel: A Scalable BERT-Based System for Trajectory Imputation. In *PVLDB*.

[52] Mashaal Musleh and Mohamed F. Mokbel. 2024. Let's Speak Trajectories: A Vision to Use NLP Models for Trajectory Analysis Tasks. In *TSAS*.

[53] Mashaal Musleh, Mohamed F. Mokbel, and Sofiane Abbar. 2022. Let's Speak Trajectories. In *SIGSPATIAL*.

[54] Avanika Narayan, Ines Chami, Laurel J. Orr, and Christopher Ré. 2022. Can Foundation Models Wrangle Your Data? *Proc. VLDB Endow.* (2022).

[55] Yansong Ning and Hao Liu. 2024. UrbanKGent: A Unified Large Language Model Agent Framework for Urban Knowledge Graph Construction. In *NeurIPS*.

[56] Matteo Paganelli, Francesco Del Buono, Andrea Baraldi, and Francesco Guerra. 2022. Analyzing How BERT Performs Entity Matching. *VLDB J.* (2022).

[57] Matteo Paganelli, Donato Tiano, and Francesco Guerra. 2023. A multi-facet analysis of BERT-based entity matching models. *PVLDB* (2023).

[58] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. In *OpenAI*.

[59] Ibrahim Sabek, Mashaal Musleh, and Mohamed F. Mokbel. 2019. Flash in Action: Scalable Spatial Data Analysis Using Markov Logic Networks. *PVLDB* (2019).

[60] Raphael Schumann, Wanrong Zhu, Weixi Feng, Tsu-Jui Fu, Stefan Riezler, and William Yang Wang. 2024. VELMA: Verbalization Embodiment of LLM Agents for Vision and Language Navigation in Street View. In *AAAI*.

[61] Tuli Shreshth, Casale Giuliano, and Jennings Nicholas R. 2022. TranAD: deep transformer networks for anomaly detection in multivariate time series data. *PVLDB* (2022).

[62] Junjun Si, Jin Yang, Yang Xiang, Hanqiu Wang, Li Li, Rongqing Zhang, Bo Tu, and Xiangqun Chen. 2024. TrajBERT: BERT-Based Trajectory Recovery With Spatial-Temporal Refinement for Implicit Sparse Trajectories. *IEEE Transactions on Mobile Computing* (2024).

[63] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and Efficient Foundation Language Models. *ArXiv* (2023).

[64] Immanuel Trummer. 2022. From BERT to GPT-3 Codex: Harnessing the Potential of Very Large Language Models for Data Management (Tutorial). *PVLDB* (2022).

[65] Alexander van Renen, Mihail Stoian, and Andreas Kipf. 2024. DataLoom: Simplifying Data Loading with LLMs. *PVLDB* (2024).

[66] Jiawei Wang, Renhe Jiang, Chuang Yang, Zengqing Wu, Makoto Onizuka, Ryosuke Shibasaki, Noboru Koshizuka, and Chuan Xiao. 2024. Large Language Models as Urban Residents: An LLM Agent Framework for Personal Mobility Generation. In *NeurIPS*.

[67] Peng Wang, Xiang Wei, Fangxu Hu, and Wenjuan Han. 2024. TransGPT: Multimodal Generative Pre-trained Transformer for Transportation. In *CLNLP*.

[68] Hao Xue, Bhanu Prakash Voutharoja, and Flora D. Salim. 2022. Leveraging Language Foundation Models for Human Mobility Forecasting. In *SIGSPATIAL*.

[69] Yibo Yan, Haomin Wen, Siru Zhong, Wei Chen, Haodong Chen, Qingsong Wen, Roger Zimmermann, and Yuxuan Liang. 2024. UrbanCLIP: Learning Text-enhanced Urban Region Profiling with Contrastive Language-Image Pretraining from the Web. In *WWW*.

[70] Chenyang Yu, Xinpeng Xie, Yan Huang, and Chenxi Qiu. 2024. Harnessing LLMs for Cross-City OD Flow Prediction. In *SIGSPATIAL*.

[71] Yuan Yuan, Jingtao Ding, Jie Feng, Depeng Jin, and Yong Li. 2024. UniST: A Prompt-Empowered Universal Model for Urban Spatio-Temporal Prediction. In *SIGKDD*.

[72] Siqi Wang Yunfan Gao, Yun Xiong and Haofen Wang. 2022. GeoBERT: Pre-Training Geospatial Representation Learning on Point-of-Interest. In *Applied. Sciences*.

[73] Siyao Zhang, Daocheng Fu, Wenzhe Liang, Zhao Zhang, Bin Yu, Pinlong Cai, and Baozhen Yao. 2024. TrafficGPT: Viewing, Processing and Interacting with Traffic Foundation Models. *ElSevier Transport Policy* (2024).

[74] Xin Zhang and Ahmed Eldawy. 2024. Spatial Query Optimization With Learning (Tutorial). *PVLDB* (2024).

[75] Yifan Zhang, Cheng Wei, Shangyou Wu, Zhengting He, and Wenhao Yu. 2024. GeoGPT: Understanding and Processing Geospatial Tasks through An Autonomous GPT. *ElSevier International Journal of Applied Earth Observation and Geoinformation* (2024).

[76] Zeyang Zhang, Xin Wang, Ziwei Zhang, Haoyang Li, Yijian Qin, and Wenwu Zhu. 2024. LLM4DyG: Can Large Language Models Solve Spatial-Temporal Problems on Dynamic Graphs?. In *SIGKDD*.

[77] Zijian Zhang, Xiangyu Zhao, Qidong Liu, Chunxu Zhang, Qian Ma, Wanyu Wang, Hongwei Zhao, Yiqi Wang, and Zitao Liu. 2023. PromptST: Prompt-Enhanced Spatio-Temporal Multi-Attribute Prediction. In *CIKM*.