

# Graph-Enhanced Large Language Models for Spatial Search

Nicole R. Schneider  
nsch@umd.edu  
University of Maryland  
College Park, USA

Kent O’Sullivan  
kosu0918@uni.sydney.edu.au  
University of Sydney  
Sydney, Australia

Hanan Samet  
hjs@cs.umd.edu  
University of Maryland  
College Park, USA

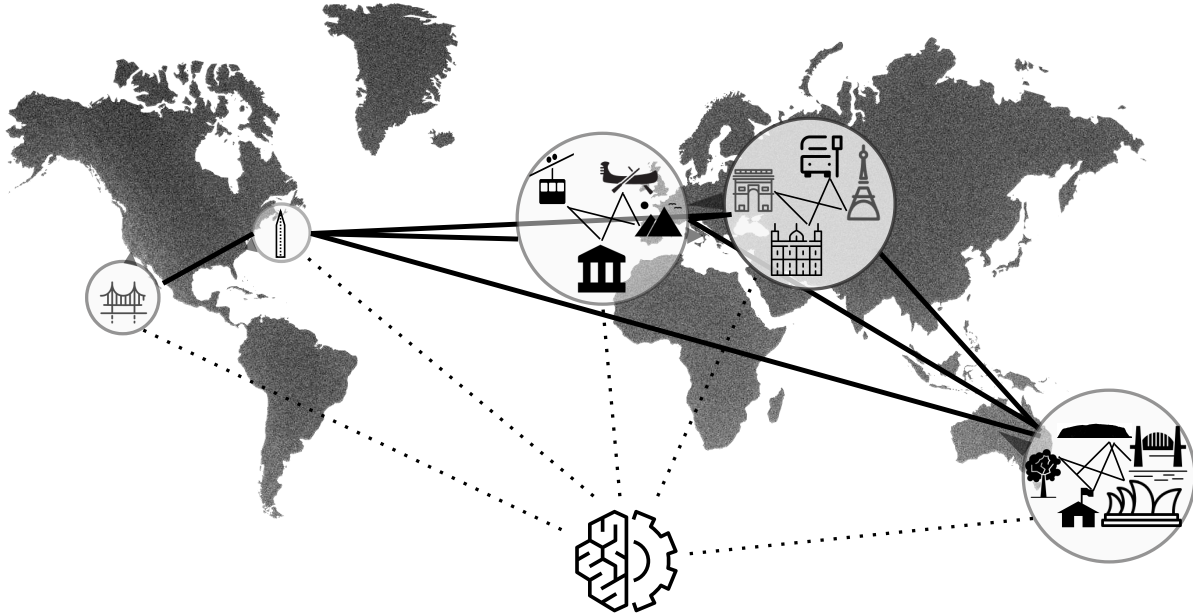


Figure 1: A world model for LLMs.

## ABSTRACT

There have been many recent improvements in the ability of Large Language Models (LLMs) to perform complex tasks and answer domain-specific questions through techniques like Retrieval Augmented Generation (RAG). However, reasoning abilities of LLMs, including spatial reasoning abilities, are still lacking. Spatial reasoning is a key component required to answer questions in a variety of domains that are grounded in the physical world, including urban planning, civil engineering, travel, and many others. To advance the development of LLMs and facilitate an impact in these domains, new research techniques must be developed to enable LLMs to reason over spatial data, which is commonly stored in the form of a graph. In this paper we outline the challenges associated with spatial reasoning through LLMs and envision a future in which search engines integrate with LLMs to answer complex spatial questions through graph-enhanced reasoning.

## CCS CONCEPTS

• **Computing methodologies** → **Spatial and physical reasoning**; *Natural language generation*; • **Information systems** → **Language models**; *Geographic information systems*; Query representation; Searching with auxiliary databases.

## KEYWORDS

Graph reasoning, spatial pattern matching, LLMs, spatial search

### VLDB Workshop Reference Format:

Nicole R. Schneider, Kent O’Sullivan, and Hanan Samet. Graph-Enhanced Large Language Models for Spatial Search. VLDB 2025 Workshop: LLM+Graph.

## 1 INTRODUCTION

There have been many recent improvements in the ability of Large Language Models (LLMs) to perform complex tasks and answer domain-specific questions through Retrieval Augmented Generation (RAG) [21] and other techniques. LLMs have been used with great success in areas that involve human abstractions, like writing code [22], summarizing and translating text [20, 30], and performing math calculations [9]. However, LLMs struggle to perform well when a task requires spatial or temporal reasoning because they

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.  
Proceedings of the VLDB Endowment. ISSN 2150-8097.

Relation	Domain	Task	Potential Use Case or Application	Examples
Metric	Realty	Housing Search	“Which homes are within 10km of a school and more than 5km from a highway?”	[6]
	Travel	Itinerary Planning	“Find a (library, museum, restaurant) or a (museum, library, restaurant) where the distance from one place to the next is less than 400m.”	[12]
Topological	Civil Engineering	Road Network Alignment	“Given a sketch drawing of street intersections and highways, which cities are a match?”	[35]
	Agriculture	Multi-Scale Data Fusion	“Which land cover datasets at different scales are a spatial match to this one?”	[17]
Directional	Event Planning	Landmark Search	“Which restaurant has an ocean to the East, a parking lot to the North, and a golf course to the south?”	[27]

**Table 1: Application domains of spatial-pattern search where graph-enhanced LLMs can provide potential benefit.**

lack a coherent, consistent world model, like the ones designed in reinforcement learning settings [1, 15]. Further, in many domains, search queries are naturally expressed through an abstract representation of spatial constraints as opposed to the keyword search terms that most search engines and LLMs are accustomed to handling. In urban planning, road network alignment is used to identify existing places with similar road layouts to a query location or proposed road plan [35]. In the travel domain, a user may recall a specific Point of Interest (POI) by the spatial configuration of the landmarks around it, without remembering the exact address or name of the place [26, 27]. In event planning, a user may not have a specific place in mind, but instead have some spatial criteria they seek to satisfy, such as finding the set of restaurants in Washington D.C. that are located next to a park and within half a kilometer of a metro stop for easy accessibility. In these examples, as well as other use cases outlined in Table 1, an abstract spatial pattern is the most natural expression of the user’s search intention.

To advance the development of LLMs to aid in answering spatial questions like the ones we describe, new research techniques must be developed to enable LLMs to reason over spatial data. Such a development would require LLMs to internalize an accurate spatial model of the world that would allow them to compare places based on their physical locations and infer spatial relationships through transitivity (Figure 1). Although recent work has shown that LLMs possess a rough model of the world intrinsically through training [2, 29], they lack knowledge of less common places and cannot generalize or infer spatial relationships, such as the distances between places [28].

The need for spatially-aware search is not presently filled by mainstream search and mapping platforms, leaving a gap that LLM-aided spatial search may be able to address. Current search engines and mapping platforms like Google Maps<sup>1</sup> are designed for keyword search, not spatial search. That is, they can support top-k

keyword queries that seek co-located objects relevant to query keywords, but they cannot support queries involving spatial relationships between keyword entities [8]. Taking the previous example of searching for restaurants located next to a park and within half a kilometer of a metro stop, top-k search would return any restaurants, parks, and metro stops in the area of interest, ignoring the desired distance and topological constraints. Even spatial group keyword search, which is more complex than top-k search, would find groups of those entities in a minimum enclosing circle, but still would not satisfy the topological constraint (restaurant adjacent to park) [4, 23, 24]. Instead, to fulfill search intentions expressed by abstract spatial patterns, an NP-hard form of search called spatial pattern matching, or pattern-based spatial search, is needed [12].

Pattern-based spatial search seeks to match query keywords and spatial constraints to the data entities that satisfy them. Query constraints can capture metric, topological, and directional relationships between entities, based on how they are positioned in space. These constraints can describe explicit relationships between pairs of entities (i.e. X north of Y), in the form of edges in a graph, or implicit relationships, such as through a pictorial representation or sketch map query [13, 35]. Spatial pattern matching supports queries with arbitrary spatial constraints, making it useful for the spatial search scenarios we describe, but it has several challenges that inhibit its use in search engines and LLMs.

In this paper, we first outline the challenges of reasoning about spatial relationships, then describe research directions that can be undertaken to address these challenges and enable pattern-based spatial search using LLMs. Specifically, we highlight:

- (1) The computational challenges associated with resolving spatial queries,
- (2) The heterogeneous nature of spatial entities and relations, which makes reasoning over them difficult, and
- (3) The graph or pictorial query formats, which are not naturally compatible with text-centric search engines and LLMs.

After describing these challenges, we present a vision for how graph-enhanced LLMs can play a role in enabling spatial search,

<sup>1</sup><https://developers.google.com/maps/documentation/places/web-service/search-text>

thus equipping users with a form of search not currently available with mainstream search engines.

### 1.1 Vision

We envision a future where search engines integrate with LLMs to answer complex spatial questions using graph-enhanced reasoning. Realizing this vision will make pattern-based spatial search accessible for widespread use in the applications we describe. To accomplish accessible pattern-based spatial search, each of the challenges outlined above must be addressed; computational complexity, heterogeneity of entity/relation types, and custom query formats. In this vision paper, we describe several research directions that can be undertaken to address these challenges by leveraging graph-enhanced LLMs. We propose that LLMs can be used to interpret spatial questions and formulate queries that are answerable by tailored spatial reasoning models, which can use spatial databases and graphs to find and rank results that answer the question. By tailoring specific learning objectives or data stores to different categories of spatial information, the key spatial properties, like directional transitivity, can be learned and then leveraged in an ensemble approach to spatial reasoning. We further envision the use of Natural Language Processing (NLP) techniques to interpret spatial patterns described in natural language and convert them into graph or pictorial queries, which can be reasoned about, enabling seamless integration with the pre-existing text interfaces used by LLMs and by most search interfaces.

### 1.2 Contribution

Our primary contributions are an analysis of the present challenges preventing the broad use of spatial pattern search and a road map of a future where LLMs provide a means to solve those challenges. We identify real-time search scenarios that would benefit from the adoption of pattern-based spatial search and describe several research directions that may lead graph-enhanced LLMs to enable fast, robust approximate pattern-based spatial search. We further highlight the need for a benchmark to measure the accuracy and scalability of pattern-based spatial search approaches that rely on LLMs compared to traditional spatial pattern matching methods that have correctness guarantees. We see this work as a natural extension to the current research efforts to achieve reasoning through LLMs using techniques like RAG. We hope that by identifying interesting and worthwhile open problems in graph-enhanced LLM spatial reasoning, this work will spark further efforts to equip LLMs with new abilities that make them better suited to serve everyday users in a variety of applications.

## 2 CHALLENGES OF SPATIAL SEARCH

Spatial search has a variety of challenges that make it difficult to support by current search engines and mapping platforms.

### 2.1 Computational Cost

Geospatial data is naturally stored as a graph, with edge attributes explicitly capturing the relationships between places (nodes). Spatial pattern matching over geospatial graphs is generally regarded as intractable for large search spaces [31, 32]. The computational complexity stems from the quadratic number of relations that can

be defined on a set of entities, which makes the process of matching query and data patterns NP-hard [12]. Unlike less expressive forms of spatial search, such as spatial group keyword search, which is also NP-hard, there are no approximate methods that reduce the complexity of spatial pattern matching broadly. Since search engines and mapping applications typically support real-time search, current spatial pattern matching methods are too computationally burdensome to incorporate into their pipelines.

### 2.2 Heterogeneous Spatial Relationships


Spatial patterns can be defined by metric, topological, and distance relationships between entities. These relations can be quantitative or qualitative, and they can be defined over point, line, and/or region data. With many possible combinations to account for, it is difficult to reason about all of them with one generic approach. Instead there are many different methods that work on different subsets of data and relation types. For example, tailored spatial pattern matching systems like Spacekey [12] and SketchMapia [35] have been developed to handle use cases like housing search (e.g. “Find all the houses in an area that are within 10km of a school and greater than 5km from any highway.”) or road network alignment (e.g. “Given a sketch drawing of street intersections and highways, which cities are a match?”). While these systems can handle some pattern-based spatial search queries, they are scoped to single use cases that involve only a subset of possible entity and relation types. Hence the methods they use are not generalizable.

### 2.3 Graph Query Input

Finally, spatial pattern matching typically requires a query pattern, given in the form of a sketch map or pictorial query that is drawn or constructed on a digital canvas, by dragging and dropping entities and arranging them in the spatial layout of interest [11, 26]. This query format offers exceptional expressivity to capture the user’s search intention, and it aligns with the intuitive visual nature of spatial pattern queries, mimicking the process of hand-drawing a sketch map on paper or imagining oneself at a location, capturing spatial relationships between objects in view [35]. However, accepting queries in a pictorial or sketch map format necessitates having a custom interface with a mechanism to construct the graph or pictorial representation, including by selecting, dragging, and dropping entities, and possibly labeling edges with relationship constraints, such as “less than 3 km.” Such an interface requires capabilities beyond the standard text input used by most LLMs and search engines, making it structurally difficult to accomplish spatial search through existing platforms.

## 3 VISION FOR SPATIALLY-AWARE LLMs

We envision future research using graph-enhanced LLMs to address the challenges causing the current gap in pattern-based spatial search (Figure 2). Specifically, we envision the computational challenges being addressed through approximation, the heterogeneous nature of spatial entities and relations being addressed through an ensemble of methods, learning objectives, and data sources, and the graph query input challenge being addressed through natural language to graph conversion using pre-trained language models. Below we describe each of these ideas in greater depth.



**Figure 2: Vision for graph-enhanced search using an LLM. A user interacts with the LLM, which interprets the query and leverages appropriate spatial models and graphs to answer the search question.**

### 3.1 Computational Efficiency Through Approximation

To alleviate the computational challenges associated with resolving spatial queries that require searching for patterns in a large quantity of spatial entities, we envision achieving efficient, real time search through approximation. Traditional spatial pattern matching methods typically perform exact search, often using pruning and other techniques to reduce some of the computational burden. A few approaches perform approximate search, but these are typically limited in expressivity, supporting only point entities and metric relations [7], or localizing the search to the entities that are relevant for a particular problem space, like roads and city blocks for the road network alignment task [35]. We envision leveraging graph-enhanced LLMs that can access reference spatial information to perform flexible, approximate spatial pattern search across a variety of domains. Using LLMs would avoid the NP-hard complexity of exact spatial pattern search, while maintaining flexibility, which is critical to support multiple domains and deal with noisy data tags, which are commonplace in spatial data.

To ensure approximate LLM-based methods achieve reasonable accuracy, a spatial pattern matching benchmark should be developed. State of the art spatial pattern matching approaches typically use pruning techniques, binning, localization, and approximate methods to search through candidate entities efficiently. However, the accuracy and efficiency of these approaches can depend on the size of the localized region or pruned entity set, which may vary significantly based on the query, dataset, and method being used. Presently, empirical testing is done using ad-hoc datasets which vary from paper to paper. Without a common benchmark or leader-board, spatial pattern matching methods cannot be rigorously evaluated with respect to one another and the trade-off between accuracy and efficiency of different techniques cannot be fully understood [28]. To elicit worst-case behavior in traditional methods for a fair comparison, we envision a spatial pattern matching benchmark that includes complex datasets and large query patterns covering the major types of spatial entities and relations. Further, such a benchmark should vary query and dataset characteristics that may influence the size of the pruned candidate set for some methods, which can significantly impact runtime.

Of particular importance, based on initial exploration, is the size of the query pattern, which has a large impact on performance. We further observed that most of the current spatial pattern matching approaches test their method on ad-hoc datasets containing query

patterns of no more than six entities [6, 25]. Based on the theoretical complexities derived in Schneider et al., spatial pattern search over larger query patterns or complex questions with many spatial constraints remains an open problem, with most existing methods having a complexity that is exponential in the size of the query [31]. As a result, including large query patterns in a benchmark will be critical to evaluating existing approaches and weighing the benefits of approximate methods, like graph-based LLM methods, that likely scale better, at the cost of search precision and recall.

### 3.2 Reasoning Over Heterogeneous Spatial Relationships with Graph-Enhanced LLMs

To address the heterogeneity in spatial entities and relations that can appear in spatial reasoning questions, we suggest using an ensemble of models or reference graph databases, each of which is tailored to a certain type of spatial relation (Figure 3). For instance, following Schneider et al., one graph may include distances on the edges, making it useful for metric spatial reasoning questions [33, 34], and another may include nodes representing points, lines, and regions with qualitative relationships between them, making it useful for topological spatial reasoning. By doing retrieval from each of these separate data stores, or by developing expert models that are trained to do different types of spatial reasoning, the results can be combined to achieve rigorous reasoning over the heterogeneous set of possible spatial relationships.

A spatial reasoning mechanism can be developed in a variety of ways, including by using an Heterogeneous Graph Neural Network (HGNN) [36] or Heterogeneous Graph Transformer [16] architecture, where the learning objective complements the goal of spatial reasoning. Instead of the masked language model or next token prediction objectives that enable LLMs to learn grammatical structure of language, a spatial learning objective, like a distance reconstruction learning objective, could be developed. For example, to accomplish metric spatial reasoning, random walks could be taken through a fully connected graph of places and distances between them, where some edges and nodes are masked. Then, the model could reconstruct the full walk, using the masked version, allowing the model to learn to predict how far apart different places are to each other, by identifying them in a reference graph and aligning the masked walk to its corresponding reference subgraph.

A topological spatial reasoning model would need to encode region and line data to enable topological reasoning. These types of geographic data are usually stored in trees, so a learning objective could seek to reconstruct the topological relationships by traversing a masked region quadtree. Since there are 9 distinct topological operators [5], one approach may be to develop one topological sub-model per operator, and have a topological relationship identifier to determine the appropriate model to use given clues in the query, such as the word “contains” or “intersects.” Tailored approaches can also be developed for directional spatial pattern matching, using graphs to encode the critical spatial properties needed to answer the questions, and developing models capable of reasoning over that information.

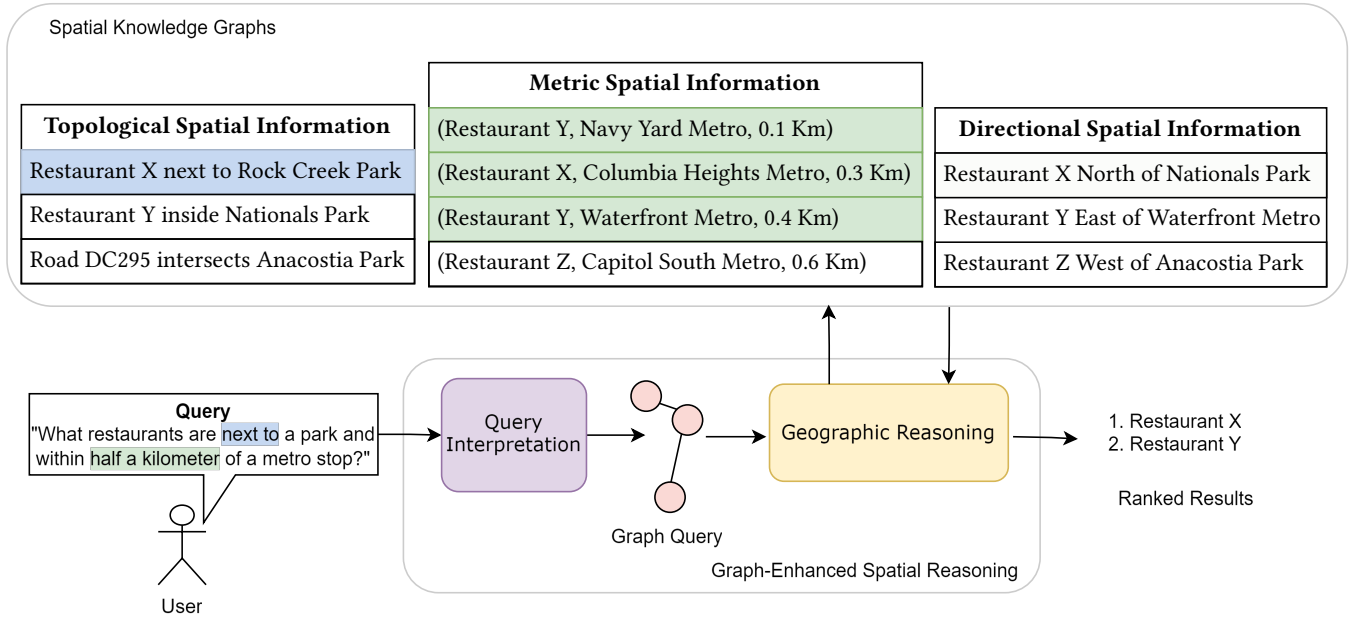


Figure 3: Vision for spatial search using graph-enhanced LLMs. User inputs a spatial query, which is interpreted and converted into a graph query. Geographic reasoning agent or module uses query to select geospatial knowledge stored in graph(s) and perform graph reasoning to resolve the query. Ranked results are outputted.

### 3.3 Converting Natural Language into Spatial Graph Queries Using Language Models

To bridge the gap between pattern-based spatial search and traditional search engines and platforms that require textual input, we draw inspiration from the Natural Language to Structured Query Language (NL2SQL) literature which uses LLMs to generate structured queries based on natural language input [3, 14, 19]. We envision an LLM-based conversion process that turns natural language text into a pictorial spatial pattern representation, enabling pattern-based spatial search via natural language input. While previous work has investigated methods for semantic parsing of spatial information from natural language for robotics and other applications [10, 18], no studies have tackled the problem of generating pictorial or graph queries for spatial pattern matching given text input. Moreover, many semantic parsing methods for extracting spatial relations are brittle, making them unsuitable for arbitrary text input from users describing spatial patterns for search. Instead, to enable spatial pattern matching over natural language input, an LLM could be used to interpret a user’s text input and construct a query or series of queries over graph data stores containing spatial entities and relation constraints.

The challenge of applying LLMs to tasks involving geospatial reasoning is further complicated by the lack of representation of lesser-known places in training data. Many obscure places have long, complex names that are out-of-vocabulary for LLMs, making it even harder for LLMs to associate them with their physical location in the world and reason about them. Given the vast array of possible places that can be included in a question requiring spatial reasoning, a basic ‘world model’ is needed to complement

the existing linguistic knowledge held in LLMs, with the ability to pull more obscure places from a data store, rather than relying on trained knowledge of every place on earth.

## 4 CONCLUSION

Spatial reasoning is a key component required to answer many questions. New research techniques must be developed to enable LLMs to reason over spatial data, which is commonly stored in the form of a graph. We outline several challenges associated with spatial reasoning through LLMs, including computational complexity, heterogeneity of spatial entities and relations, and non-textual query input format. To address these issues we envision a future in which search engines integrate with LLMs that use graph-enhanced reasoning to answer complex spatial questions. We hope that by identifying interesting and worthwhile open problems in graph-enhanced LLM spatial reasoning, this work will spark further efforts to equip LLMs with new abilities that make them better suited to serve everyday users in a variety of applications.

## ACKNOWLEDGMENTS

This work was sponsored in part by the NSF under Grants IIS-18-16889, IIS-20-41415, and IIS-21-14451.

## REFERENCES

- [1] Axel Backlund and Lukas Petersson. 2025. Vending-Bench: A Benchmark for Long-Term Coherence of Autonomous Agents. arXiv:2502.15840 [cs.AI] <https://arxiv.org/abs/2502.15840>
- [2] Prabin Bhandari, Antonios Anastasopoulos, and Dieter Pfoser. 2023. Are Large Language Models Geospatially Knowledgeable?. In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems* (Hamburg, Germany) (SIGSPATIAL '23). Association for Computing Machinery, New

- York, NY, USA, Article 75, 4 pages. <https://doi.org/10.1145/3589132.3625625>
- [3] Ursin Brunner and Kurt Stockinger. 2021. ValueNet: A Natural Language-to-SQL System that Learns from Database Information. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 2177–2182. <https://doi.org/10.1109/ICDE51399.2021.00220>
  - [4] Xin Cao, Gao Cong, Tao Guo, Christian S. Jensen, and Beng Chin Ooi. 2015. Efficient Processing of Spatial Group Keyword Queries. *ACM Trans. Database Syst.* 40, 2, Article 13 (jun 2015), 48 pages. <https://doi.org/10.1145/2772600>
  - [5] Anderson Chaves Carniel. 2024. Defining and designing spatial queries: the role of spatial relationships. *Geo-spatial Information Science* 27, 6 (2024), 1868–1892. <https://doi.org/10.1080/10095020.2022.2163924> arXiv:<https://doi.org/10.1080/10095020.2022.2163924>
  - [6] Hongmei Chen, Yixiang Fang, Ying Zhang, Wenjie Zhang, and Lizhen Wang. 2019. ESPM: Efficient spatial pattern matching. *IEEE Transactions on Knowledge and Data Engineering* 32, 6 (2019), 1227–1233.
  - [7] Lisi Chen, Gao Cong, Christian S Jensen, and Dingming Wu. 2013. Spatial keyword query processing: An experimental evaluation. *Proceedings of the VLDB Endowment* 6, 3 (2013), 217–228.
  - [8] Yue Chen, Kaiyu Feng, Gao Cong, and Han Mao Kiah. 2022. Example-based spatial pattern matching. *Proceedings of the VLDB Endowment* 15, 11 (2022), 2572–2584.
  - [9] Aniket Didolkar, Anirudh Goyal, Nan Rosemary Ke, Siyuan Guo, Michal Valko, Timothy Lillicrap, Danilo Jimenez Rezende, Yoshua Bengio, Michael C Mozer, and Sanjeev Arora. 2024. Metacognitive capabilities of llms: An exploration in mathematical problem solving. *Advances in Neural Information Processing Systems* 37 (2024), 19783–19812.
  - [10] Kais Dukes. 2014. SemEval-2014 Task 6: Supervised Semantic Parsing of Robotic Spatial Commands. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, Preslav Nakov and Torsten Zesch (Eds.). Association for Computational Linguistics, Dublin, Ireland, 45–53. <https://doi.org/10.3115/v1/S14-2006>
  - [11] Yixiang Fang, Reynold Cheng, Jikun Wang, Lukito Budiman, Gao Cong, and Nikos Mamoulis. 2018. SpaceKey: exploring patterns in spatial databases. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*. IEEE, IEEE, 1577–1580.
  - [12] Yixiang Fang, Yun Li, Reynold Cheng, Nikos Mamoulis, and Gao Cong. 2019. Evaluating pattern matching queries for spatial databases. *The VLDB Journal* 28 (2019), 649–673.
  - [13] Andre Folkers, Hanan Samet, and Aya Soffer. 2000. Processing pictorial queries with multiple instances using isomorphic subgraphs. In *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, Vol. 4. IEEE, IEEE, 51–54.
  - [14] Han Fu, Chang Liu, Bin Wu, Feifei Li, Jian Tan, and Jianling Sun. 2023. CatSQL: Towards Real World Natural Language to SQL Applications. *Proc. VLDB Endow.* 16, 6 (Feb. 2023), 1534–1547. <https://doi.org/10.14778/3583140.3583165>
  - [15] David Ha and Jürgen Schmidhuber. 2018. World models. <https://doi.org/10.5281/ZENODO.1207631>
  - [16] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous Graph Transformer. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) (WWW '20). Association for Computing Machinery, New York, NY, USA, 2704–2710. <https://doi.org/10.1145/3366423.3380027>
  - [17] Zhekun Huang, Haizhong Qian, Xiao Wang, Defu Lin, Junwei Wang, and Limin Xie. 2023. Graph neural network-based identification of ditch matching patterns across multi-scale geospatial data. *Geocarto International* 38, 1 (2023), 2294900. <https://doi.org/10.1080/10106049.2023.2294900> arXiv:<https://doi.org/10.1080/10106049.2023.2294900>
  - [18] Dohyun Kim, Nayoung Oh, Deokmin Hwang, and Daehyung Park. 2024. LINGO-Space: Language-Conditioned Incremental Grounding for Space. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 9 (Mar. 2024), 10314–10322. <https://doi.org/10.1609/aaai.v38i9.28898>
  - [19] Hyeonji Kim, Byeong-Hoon So, Wook-Shin Han, and Hongrae Lee. 2020. Natural language to SQL: where are we today? *Proc. VLDB Endow.* 13, 10 (June 2020), 1737–1750. <https://doi.org/10.14778/3401960.3401970>
  - [20] Philippe Laban, Wojciech Kryscinski, Divyansh Agarwal, Alexander Fabbri, Caiming Xiong, Shafiq Joty, and Chien-Sheng Wu. 2023. SummEdits: Measuring LLM Ability at Factual Reasoning Through The Lens of Summarization. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 9662–9676. <https://doi.org/10.18653/v1/2023.emnlp-main.600>
  - [21] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33 (2020), 9459–9474.
  - [22] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. *Advances in Neural Information Processing Systems* 36 (2023), 21558–21572.
  - [23] Ahmed Mahmood and Walid G. Aref. 2017. Query Processing Techniques for Big Spatial-Keyword Data. In *Proceedings of the 2017 ACM International Conference on Management of Data (Chicago, Illinois, USA) (SIGMOD '17)*. Association for Computing Machinery, New York, NY, USA, 1777–1782. <https://doi.org/10.1145/3035918.3054773>
  - [24] Ahmed R. Mahmood and Walid G. Aref. 2019. *Querying Spatial-Keyword Data*. Springer International Publishing, Cham, 7–23. [https://doi.org/10.1007/978-3-031-01867-1\\_2](https://doi.org/10.1007/978-3-031-01867-1_2)
  - [25] Carlos Minervino, Claudio Campelo, Maxwell Oliveira, and Salatiel Silva. 2024. QQESPM: A Quantitative and Qualitative Spatial Pattern Matching Algorithm. arXiv:2312.08992 [cs.DB] <https://arxiv.org/abs/2312.08992>
  - [26] Kent O'Sullivan, Nicole R. Schneider, Aleza Rasheed, and Hanan Samet. 2023. GESTALT: Geospatially Enhanced Search with Terrain Augmented Location Targeting. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on Searching and Mining Large Collections of Geospatial Data (Hamburg, Germany) (GeoSearch '23)*. Association for Computing Machinery, New York, NY, USA, Article 1, 8 pages. <https://doi.org/10.1145/3615890.3628535>
  - [27] Kent O'Sullivan, Nicole R. Schneider, and Hanan Samet. 2023. COMPASS: Cardinal Orientation Manipulation and Pattern-Aware Spatial Search. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on Searching and Mining Large Collections of Geospatial Data (Hamburg, Germany) (GeoSearch '23)*. Association for Computing Machinery, New York, NY, USA, Article 5, 8 pages. <https://doi.org/10.1145/3615890.3628537>
  - [28] Kent O'Sullivan, Nicole R. Schneider, and Hanan Samet. 2024. Metric Reasoning in Large Language Models. In *Proceedings of the 32nd ACM International Conference on Advances in Geographic Information Systems (Atlanta, GA, USA) (SIGSPATIAL '24)*. Association for Computing Machinery, New York, NY, USA, 501–504. <https://doi.org/10.1145/3678717.3691226>
  - [29] Jianzhong Qi, Zuqing Li, and Egemen Tanin. 2023. MaaSDB: Spatial Databases in the Era of Large Language Models (Vision Paper). In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems (Hamburg, Germany) (SIGSPATIAL '23)*. Association for Computing Machinery, New York, NY, USA, Article 54, 4 pages. <https://doi.org/10.1145/3589132.3625597>
  - [30] Nicole R. Schneider, Avik Das, Kent O'Sullivan, and Hanan Samet. 2024. Cross-Lingual Clustering Using Large Language Models. In *Proceedings of the 7th ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery (Atlanta, GA, USA) (GeoAI '24)*. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3687123.3698280>
  - [31] Nicole R. Schneider, Kent O'Sullivan, and Hanan Samet. 2024. Graph-based Spatial Pattern Matching: A Theoretical Comparison. In *Proceedings of the 32nd ACM International Conference on Advances in Geographic Information Systems (Atlanta, GA, USA) (SIGSPATIAL '24)*. Association for Computing Machinery, New York, NY, USA, 505–508. <https://doi.org/10.1145/3678717.3691227>
  - [32] Nicole R. Schneider, Kent O'Sullivan, and Hanan Samet. 2024. The Future of Graph-based Spatial Pattern Matching (Vision Paper). In *40th IEEE International Conference on Data Engineering, ICDE 2024 – SEAGraph Workshop*. IEEE, IEEE, Utrecht, Netherlands, 360–364.
  - [33] Nicole R. Schneider, Nandini Ramachandran, Kent O'Sullivan, and Hanan Samet. 2025. DistRAG: Towards Distance-Based Spatial Reasoning in LLMs. arXiv:2506.03424 [cs.CL] <https://arxiv.org/abs/2506.03424>
  - [34] Nicole R. Schneider, Nandini Ramachandran, Kent O'Sullivan, and Hanan Samet. 2025. Engineering Prompts for Spatial Questions. In *Companion Proceedings of the ACM on Web Conference 2025 (Sydney NSW, Australia) (WWW '25)*. Association for Computing Machinery, New York, NY, USA, 1633–1634. <https://doi.org/10.1145/3701716.3717807>
  - [35] Angela Schwing, Jia Wang, Malumbo Chipofya, Sahib Jan, Rui Li, and Klaus Broelemann. 2014. SketchMapia: Qualitative representations for the alignment of sketch and metric maps. *Spatial cognition & computation* 14, 3 (2014), 220–254.
  - [36] Chuxu Zhang, Dongjin Song, Chao Huang, Ananthram Swami, and Nitesh V. Chawla. 2019. Heterogeneous Graph Neural Network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (Anchorage, AK, USA) (KDD '19)*. Association for Computing Machinery, New York, NY, USA, 793–803. <https://doi.org/10.1145/3292500.3330961>