The Future of Graph-based Spatial Pattern Matching (Vision Paper)

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

Kent O'Sullivan
University of Maryland
College Park, Maryland
Email: osullik@umd.edu

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

***Abstract*—Spatial Pattern Matching is an important problem in information retrieval that involves reasoning about the relative position, distance, and orientation of objects with respect to each other. Most spatial pattern matching approaches use large, complex graphs or multigraphs to explicitly encode rich spatial information. The downside of this complexity is that search over spatial patterns remains badly constrained by computationally intensive classes of algorithms, like subgraph matching and constraint satisfaction. This paper highlights the recent approaches to graph-based spatial pattern matching, and presents a vision of the way forward, using graph-based Artificial Intelligence as a flexible, approximate approach to the otherwise intractable problem.**

I. Introduction

Many information retrieval problems in the spatial domain involve *Spatial Pattern Matching (SPM)*, the task of selecting entities that match a given set of spatial constraints. Performing spatial pattern matching requires automatically and efficiently reasoning over complex spatial data, including the relative position, distance, and orientation of objects with respect to one another. For future systems to answer complex queries with spatial components that require spatial pattern matching, the underlying pattern matching algorithms will need to be efficient enough to support response times within a few seconds. Recent work successfully applies neural approaches to related spatial problems, like geospatial factoid question answering, by leveraging spatial semantic graph representations [1]. However, no existing neural methods can handle the complexity of spatial pattern matching.

Most of the existing spatial pattern matching approaches use graphs to encode the data, since they are a natural structure to capture the complex pairwise spatial constraints that apply to a set of entities in space. The downside of this complexity is that search over spatial patterns remains badly constrained by computationally intensive algorithm classes like subgraph matching (SGM) and constraint satisfaction problems (CSP), which are broadly intractable for any reasonable number of database objects and query constraints. The tension between the natural expressiveness of graphs for encoding spatial relationships and the intractability of resolving queries against them for large collections of spatial data prevents modern search engines from supporting complex spatial queries at scale.

Recent graph-based spatial pattern matching approaches tend to leverage heuristic methods, custom data structures, pruning, and early stopping to achieve better than exponential complexity in practice [2]–[5]. However, even these improvements are not sufficient to support real-time spatial pattern matching. Machine learning has sped up many on-the-fly search tasks involving geospatial data [6] but has yet to reach spatial pattern matching. To bridge this gap, we propose leveraging graph-based machine learning algorithms, which align well with the naturally graph-based spatial pattern matching task. In particular, Graph Neural Networks (GNNs) offer a promising avenue to improve the efficiency of spatial pattern matching. Recent work applies GNNs to general subgraph matching problems as an efficient, noise-resistant approximate method [7], [8], and we envision extending this to spatial pattern matching.

We must address a few challenges to make GNN-based spatial pattern matching possible. GNN subgraph matching methods show promise on synthetically generated undirected graphs, but spatial pattern matching requires matching on complex spatial relations represented by (possibly multiple) directed edges between nodes. More work remains to extend the flexibility of GNN-based methods, especially to include matching on more complicated edge relations. Likewise, classical graph theoretic methods offer correctness guarantees that neural methods lack. There are many avenues of active research aiming to address the lack of guarantees, including by devising hybrid approaches, improving model explainability, and adding confidence scores and alignment witnesses to model output [9].

This paper presents a vision of the way forward for spatial pattern matching, highlighting many of the key challenges and opportunities associated with applying graph-based machine learning to the problem. The rest of this paper describes the relevant background on spatial search and pattern matching in sections II and III. Then section IV presents an overview of recent state of the art spatial pattern matching approaches, focusing on subgraph matching and constraint satisfaction approaches. Finally, in section V we provide a vision for the future of graph-based spatial pattern matching that draws on recent GNN subgraph matching methods, before concluding in section VI.

II. Spatial Entities and Relations

We begin by defining the core concepts of spatial data and spatial pattern matching.

A. Graph Encoding

Many classical and recent spatial pattern matching approaches represent spatial data using graphs. A graph $\mathcal{G}$ consists of a set of nodes or vertices V and edges E such that $\mathcal{G} = \{V, E\}$. Edges connect Nodes and possess one or many labels, representing spatial relations (defined below).

B. Spatial Entities (Nodes)

Spatial entities are the basic elements of spatial data, and are typically encoded as nodes in the graph. They fall into three main classes: **(i)** *Points*, which consist of an (x, y) coordinate pair in Cartesian space, **(ii)** *Lines*, which represent the shortest path between two points, and **(iii)** *Regions*, which represent the area inside a polyline joining several points. Typically, points represent objects in the world, like trees, fire hydrants, and street signs, lines represent ways, like roadways, waterways, and railways, and regions represent the extent of lakes, stadiums, large buildings, and counties. Most graph-based spatial pattern matching methods operate on point entities, with a few also supporting line and region data.

C. Spatial Relations (Edges)

Points, lines and regions relate to each other spatially through the following types of relationships, (typically encoded as edges in the graph [10]–[12]): **(i)** *Metric relations* that describe the distances between spatial entities (ten miles, near, far, etc.), **(ii)** *Topological relations* that describe how regions, lines, and points interact (intersect, contain, touch, etc.), and **(iii)** *Directional relations* that describe how entities are positioned relative to each other in space (north, left, behind, etc.).

Some methods also address *Order relations* [5], which describe the cyclic order in which objects appear with respect to their centroid. Hybrid relations that combine two or more of the primitive spatial relations are also possible [12]. To encode spatial relations, most methods [3]–[5], [13], [14] create a *Graph* or *Qualitative Constraint Network* to explicitly encode pairwise spatial relations between objects as edges. More complex spatial relations, like directional relations, require a fully connected graph or multigraph to capture all the relevant relations.

III. Spatial Pattern Matching

Spatial search is the task of retrieving spatial entities from a database that meet the constraints specified by a spatial query. *Spatial pattern matching* is a type of spatial search that is often defined as a graph matching problem, where nodes have keyword labels representing objects in the world or points of interest (POIs) and edges represent desired spatial relationships between the nodes. Most spatial pattern matching methods formulate the problem as a *constraint satisfaction* problem or a *subgraph-matching* problem, both of which scale poorly with the number of relations in the database and constraints in the query.

A. Constraint Satisfaction Problems (CSP)

Constraint satisfaction problems for spatial pattern matching seek an assignment of valid variables (database objects) given the query constraints, typically using backtracking or forward-checking approaches. Most CSPs are exponential in the number of objects, the number of constraints, or the number of query terms. In the worst case, where constraints are numerous, specificity is low, and constraints are poorly ordered, CSP approaches prove intractable.

B. Subgraph Matching (SGM)

Subgraph matching approaches encode the database as a graph of objects and relations, and queries as a subgraph of objects and constraints. These methods then seek to identify where the query pattern exists in the database graph, often using tree search, look-ahead functions, and aggressive pruning to avoid exponential worst-case complexity [2]–[4]. However, even recent methods are prohibitively slow on large database and query graphs, or do not support complex spatial relation types, like topological and directional relations [2], [3].

C. Other Methods

A few spatial pattern matching methods that have been proposed avoid using graphs or qualitative constraint networks entirely. These methods typically encode the data by segmenting space and indexing each object along with the set of objects that it relates to spatially [15]–[18]. Set intersection [15], [16], [19], star calculus [17], or recursive matrix search [18] then identify candidates matching the query pattern. None of these methods overcome the complexity barrier of the graph-based methods without reducing the problem to a "find any matching pattern" rather than the standard "find all matching patterns" [18].

IV. Current Spatial Pattern Matching Approaches

Researchers have been applying CSP and SGM approaches to graph-based spatial pattern matching for several decades. In this section we give an overview of the prominent methods which captures the trajectory of the two main threads of work.

A. Constraint Satisfaction Problem Approaches

Multi-relation Spatial Joins (MSJ) [13] and their window reduction variants are early methods that enable spatial pattern matching over point, line, and region data. They support the eight major topological relations and handle both qualitative metric relations and directional relations for richer constraints. Their spatial join constraint satisfaction methods are exponential in complexity, and even when run in parallel using distributed in-memory computation, response times are on the order of minutes for large datasets [20].

SkectchMapia [5], [21] takes a simplified sketch-map approach, performing spatial pattern matching over point and limited line and region data consisting of street segments and disjoint city blocks. Their CSP method supports some topological constraints and directional constraints for very localized objects, which limits the size of their qualitative constraint network. They further use a heuristic method to find candidates with high overlap to the query qualitative constraint network, which is sub-exponential in time complexity, but it remains intractable for large sets of objects.

Qualitative Spatial Reasoning with Uncertain Evidence [22] addresses probabilistic directional spatial reasoning by formulating the problem as a Markov Logic Network. Despite the noise-tolerance of this approach, it remains exponential in complexity, and does not handle topological or metric relations.

B. Subgraph Matching Approaches

Pictorial Queries Using Isomorphic Subgraphs [14] uses a bottom-up subgraph matching approach that supports metric and directional spatial pattern matching over point data. It simplifies the problem of directional pattern matching by treating directional constraints as angles between vertices, which requires accurate relative distances between query objects. The exponential complexity in the number of database objects makes this approach prohibitively slow, and the accurate distance assumption on the query objects proves unrealistic given humans are notoriously poor at estimating distance [5].

Spacekey [4] provides the SGM-based *Multi-Pair Join* and *Multi-Star Join* algorithms which enable spatial pattern matching over point data and metric relations defined by a distance interval and a sign. Their Multi-Pair-Join and Multi-Star-Join algorithms use multi-way distance joins to find matching edges and join them to find matching subgraphs. Their evaluation shows that both of their approaches are faster to execute than generic subgraph matching algorithms, but further testing on large datasets show response times greater than 11 minutes [3], which is too slow for real-time spatial reasoning.

ESPM: Efficient Spatial Pattern Matching [3] extends the work of Fang et al. [4] by adding a step that uses a set of Inverted Linear Quadtrees, one per object keyword, to prune unpromising edges before running Multi-Star-Join. By eliminating unpromising nodes early and carefully constructing the join order, ESPM scales much better than the original Multi-Star-Join in practice, despite having a similar theoretical worst-case complexity. Although ESPM achieves five second response times on databases of 10,000,000 objects, their tests use queries of size six nodes or smaller, and their approach only supports metric relations.

QQESPM [2] extends the work of Chen et al. [3] to account for topological relations in addition to metric relations, checking both the quantitative (metric) and qualitative (topological) constraints at each step. The theoretical complexity is the same as ESPM with an extra constant factor for the additional qualitative checks, but the response time in practice is not known for databases larger than 40,000 objects.

V. Opportunities in Spatial Pattern Matching

In this section we describe the common threads in recent spatial pattern matching approaches and describe opportunities to incorporate spatial pattern matching into future AI pipelines by solving the problem approximately using graph AI.

A. Common Threads and Observations

Spatial pattern matching is becoming more efficient as newer approaches leverage heuristic methods, custom data structures, pruning, and early stopping to improve response times in practice. The recent SGM-based spatial pattern matching works build off the work of Fang et al. [4], using Inverted Linear Quadtrees [3], and new methods follow a similar approach but support additional spatial relations (i.e. topological in addition to metric) [2]. The recent CSP-based work can account for uncertain information using probabalistic methods, but remains bounded by exponential complexity [22]. Even the most recent spatial pattern matching approaches do not support all four spatial relation types (metric, topological, directional, and order relations), and have not been shown to be efficient on large databases and large query patterns.

To push spatial pattern matching forward substantially enough to support applications that require real-time spatial reasoning, approximate solutions offer a promising direction to explore. In addition to efficiency, approximate solutions are more robust to noise in the database and query graphs, returning valuable partial matches when an exact match cannot be found, a characteristic that is critical when operating on noisy real-world geodata, where tags are sometimes missing or inaccurate. The obvious downside to approximate methods is the lack of correctness guarantees. For spatial pattern matching, the consequence of an incorrect or incomplete result is minimal, given that the alternative exact methods remain prohibitively slow, making them unusable for live systems. Further, hybrid methods can be devised, wherein we find an initial approximate solution optimized for recall, and then use an exact method to verify the correctness of the results, since the candidate pool is greatly reduced by the initial search.

B. Graph AI for Spatial Pattern Matching

Since we already typically formulate spatial pattern matching as a graph problem and solve it either using subgraph matching or constraint satisfaction approaches, a natural extension to the existing line of work is to apply graph-based AI as an approximate solution to one of those formulations of the problem. The obvious choice, Graph Neural Networks (GNNs), are traditionally good at similarity measurement tasks, like finding the similarity between two graphs [23], [24]. However, they do not perform as well at substructure extraction problems like shortest path extraction or subgraph matching (which is how spatial pattern matching is typically formulated). Some work uses GNNs to approximately solve subgraph matching tasks on synthetically generated graphs, starting with a simple node classification framing of the problem, where each node is classified as a match or non-match with the query graph [7], [24]. Early methods require access to the query graph during training, but more recent approaches overcome this limitation [8], [25]. Foundational node embedding methods sample the graph using random walks and aggregate neighborhood information, enabling fast inference (classifying 80,000 nodes in about one second) [25]. However, even the most promising approaches cannot support large directed graphs with multiple edge labels [8], [9], which is required for spatial pattern matching.

We envision the path to using GNNs for spatial subgraph matching involves further improvement in the flexibility of the current methods, especially to include matching on multiple directed edges. To address these challenges, GNN-based SPM approaches will likely need community and edge-level embeddings in addition to the common node and graph-level embeddings. Further, additional strategies for efficiently generating embeddings of unseen nodes on the fly without aggressive sampling should be explored.

Beyond the benefits of speed and flexibility, we envision that GNN and other graph AI-based spatial pattern matching approaches will integrate into existing learned information retrieval pipelines and systems as a step towards achieving general spatial reasoning. Such systems could leverage pretrained models that already contain geographic world knowledge, and enable powerful spatial reasoning over that information. Further work in multi-modal learning could enable reasoning over multiple modalities of geospatial data, like remote sensing imagery (RSI), structured geodatabase tables, eyewitness video recordings of world events, and so on. Models could combine and synthesize diverse sources of geospatial data that are traditionally treated separately to perform spatial pattern matching and other spatial reasoning tasks over them, unlocking a wealth of spatial knowledge.

VI. Conclusion

This paper highlights recent trends in graph-based spatial pattern matching, describing how constraint satisfaction and subgraph matching approaches currently solve the problem. We present our vision of future graph AI-based spatial pattern matching, and detail how the proposed way forward could enable spatial pattern matching in systems already excelling at standard reasoning tasks. Realizing this vision would enable systems and search engines to support queries involving complex spatial patterns at scale, which is a key step toward achieving broader spatial reasoning over the wealth of geodata presently available.

Acknowledgment

This work was sponsored in part by NSF Grants IIS-18-16889, IIS-20-41415, and IIS-21-14451 and the Australian-American Fulbright Commission. Its contents do not necessarily represent the official views of the Fulbright Program.

References

[1] H. Li, E. Hamzei, I. Majic, H. Hua, J. Renz, M. Tomko, M. Vasardani, S. Winter, and T. Baldwin, "Neural factoid geospatial question answering," *Journal of Spatial Information Science*, no. 23, pp. 65–90, 2021.

[2] C. Minervino, C. Campelo, M. Oliveira, and S. Silva, "Qqespm: A quantitative and qualitative spatial pattern matching algorithm," *arXiv preprint arXiv:2312.08992*, 2023.

[3] H. Chen, Y. Fang, Y. Zhang, W. Zhang, and L. Wang, "Espm: Efficient spatial pattern matching," *IEEE Transactions on Knowledge and Data Engineering*, vol. 32, no. 6, pp. 1227–1233, 2019.

[4] Y. Fang, Y. Li, R. Cheng, N. Mamoulis, and G. Cong, "Evaluating pattern matching queries for spatial databases," *The VLDB Journal*, vol. 28, pp. 649–673, 2019.

[5] A. Schwering, J. Wang, M. Chipofya, S. Jan, R. Li, and K. Broelemann, "Sketchmapia: Qualitative representations for the alignment of sketch and metric maps," *Spatial cognition & computation*, vol. 14, no. 3, pp. 220–254, 2014.

[6] C. Lülf, D. M. L. Martins, M. A. V. Salles, Y. Zhou, and F. Gieseke, "Rapidearth: A search-by-classification engine for large-scale geospatial imagery," in *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems*, ser. SIGSPATIAL '23. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: https://doi.org/10.1145/3589132.3625601

[7] X. Liu, H. Pan, M. He, Y. Song, X. Jiang, and L. Shang, "Neural subgraph isomorphism counting," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 1959–1969.

[8] Z. Lan, L. Yu, L. Yuan, Z. Wu, Q. Niu, and F. Ma, "Sub-gmn: The subgraph matching network model," in *2023 16th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, 2023, pp. 1–7.

[9] I. Roy, V. S. B. R. Velugoti, S. Chakrabarti, and A. De, "Interpretable neural subgraph matching for graph retrieval," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 7, 2022, pp. 8115–8123.

[10] A. C. Carniel, "Spatial information retrieval in digital ecosystems: A comprehensive survey," in *Proceedings of the 12th International Conference on Management of Digital EcoSystems*, 2020, pp. 10–17.

[11] P. G. K. Bertella, Y. K. Lopes, R. A. P. de Oliveira, and A. C. Carniel, "A systematic review of spatial approximations in spatial database systems," *Journal of Information and Data Management*, vol. 13, no. 2, 2022.

[12] A. Chaves Carniel, "Defining and designing spatial queries: the role of spatial relationships," *Geo-spatial Information Science*, pp. 1–25, 2023.

[13] D. Papadias, N. Mamoulis, and B. Delis, "Algorithms for querying by spatial structure," in *Proceedings of Very Large Data Bases Conference (VLDB), New York*, 1998.

[14] A. Folkers, H. Samet, and A. Soffer, "Processing pictorial queries with multiple instances using isomorphic subgraphs," in *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, vol. 4. IEEE, 2000, pp. 51–54.

[15] A. Soffer and H. Samet, "Pictorial query specification for browsing through spatially referenced image databases," *Journal of Visual Languages & Computing*, vol. 9, no. 6, pp. 567–596, 1998.

[16] K. O'Sullivan, N. R. Schneider, and H. Samet, "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*, ser. GeoSearch'23. Association for Computing Machinery, 2023.

[17] X. Liu, S. Shekhar, and S. Chawla, "Object-based directional query processing in spatial databases," *IEEE transactions on knowledge and data engineering*, vol. 15, no. 2, pp. 295–304, 2003.

[18] K. O'Sullivan, N. R. Schneider, and H. Samet, "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*, ser. GeoSearch'23. Association for Computing Machinery, 2023.

[19] A. Soffer and H. Samet, "Pictorial query specification for browsing through image databasess," in *Proceedings of the Second International Conference on Visual Information Systems*, 1997, pp. 117–124.

[20] Z. Du, X. Zhao, X. Ye, J. Zhou, F. Zhang, and R. Liu, "An effective high-performance multiway spatial join algorithm with spark," *ISPRS International Journal of Geo-Information*, vol. 6, no. 4, p. 96, 2017.

[21] S. Jan and A. Schwering, "Sketchmapia: a framework for qualitative alignment of sketch maps and metric maps," in *Proceedings of the AGILE conference on Geographic Information Science*, vol. 18. Association of Geographic Information Laboratories in Europe, 2015.

[22] M. Duckham, J. Gabela, A. Kealy, R. Kyprianou, J. Legg, B. Moran, S. K. Rumi, F. D. Salim, Y. Tao, and M. Vasardani, "Qualitative spatial reasoning with uncertain evidence using markov logic networks," *International Journal of Geographical Information Science*, vol. 37, no. 9, pp. 2067–2100, 2023.

[23] Y. Bai, H. Ding, S. Bian, T. Chen, Y. Sun, and W. Wang, "Simgnn: A neural network approach to fast graph similarity computation," in

Proceedings of the twelfth ACM international conference on web search and data mining, 2019, pp. 384–392.
[24] D. Krleža and K. Fertalj, "Graph matching using hierarchical fuzzy graph neural networks," *IEEE Transactions on Fuzzy Systems*, vol. 25, no. 4, pp. 892–904, 2016.
[25] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," *Advances in neural information processing systems*, vol. 30, 2017.