Geographic Places to Semantic Spaces: Analysis of Geospatial Embeddings

Dilip Kumar Dalei[1], Sangeeta Shrivastava[2], Narayan Panigrahi[3]

[1,2,3]CAIR, DRDO, Bengaluru, India
[1]dilipkumar.cair@gov.in, [2]sangeeta.cair@gov.in, [3]pani.cair@gov.in

Keywords: Word Embeddings, Spatial Embeddings, Location Embeddings, Word2Vec, Word Vectorization.

ABSTRACT

Geospatial data representation has evolved significantly over the years, from basic points, lines, and polygons to more complex embeddings. Geospatial embeddings, a technique used in spatial analysis, map geographic locations to vectors of real numbers, enabling the integration of diverse data types and facilitating advanced spatial analysis tasks. By mapping geographic entities to vectors of real numbers, embeddings capture not only the spatial coordinates but also the semantic meaning and relationships embedded in the data. This transformation enables the integration of diverse spatial data types, such as satellite imagery, GIS layers, textual descriptions, and sensor data, into a unified representation that preserves the unique characteristics, underlying patterns, and relationships between data. These embeddings enable machine learning algorithms to perform tasks such as location prediction, change detection, and semantic analysis with unprecedented accuracy. These representation methods facilitate the integration of geospatial data into deep learning models and provide a mechanism for efficiently comparing, indexing, and classifying geometric entities.

This paper explores various spatial embedding techniques, their applications, challenges, and future directions. The paper also provides a comparative analysis of different approaches and discusses their effectiveness in diverse geospatial domains. Finally, we identify key insights, research gaps, and research scope in the field of geospatial embeddings.

1. Introduction

1.1 General Instructions

The advent of Geographic Information Systems (GIS) in the 1960s marked a significant step forward in geospatial data representation. However, as the volume and complexity of geospatial data continued to grow, the limitations of traditional GIS representations became increasingly apparent. The need for more efficient and effective ways to analyze and utilize spatial data led to the exploration of new approaches, culminating in the emergence of geospatial embeddings.

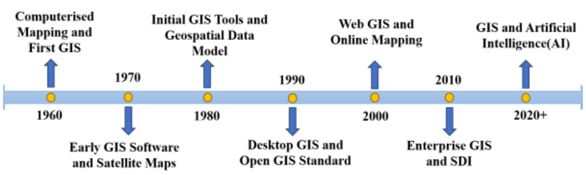

Figure 1: Evolution of GIS

Spatial embeddings are vector representations of spatial data and become the core building blocks of GeoAI applications. The crucial aspect of spatial embeddings is that they aim to capture far more than just the explicit coordinates or overt attributes of a geographic entity. A well-constructed spatial embedding encodes the semantic meaning of the entity, its inherent spatial relationships with other entities, and its dependencies on various contextual factors. These embeddings can be used for tasks such as geographic information retrieval, spatial data analysis, and semantic understanding. For example, word embeddings for geolocated text include techniques like Word2Vec and GloVe that have been applied to geotagged text data, such as social media posts, to capture semantic relationships between locations based on co-occurrence patterns.

There are three important contributions of this paper. First, it explores the basic principles and concepts of word embeddings and their applicability in the geospatial domain. Second, the paper analyses the embeddings for important spatial primitives like points, lines, and polygons. It looks into various existing research studies for such geometric embeddings. Third, the paper presents a comparative study and evaluation of these approaches. It presents the various insights and observations that may enrich the understanding geospatial embedding and their applicability in building a comprehensive and efficient GeoAI application.

The structure of the paper is as follows. Section 2 explores various geospatial data types. This is followed by a detail study of spatial embeddings in Section 3. Section 4 describes the evaluation methods and metrics for geospatial embeddings. Section 5 accounts key observations and insights from the work. At the end, the paper summarizes with a conclusion and future avenues of research in section 6.

2. Geospatial Data

2.1 Geospatial Data Type

Geographic data can take many forms such as text, images, points, lines, polygons, and graphs. Each data type requires specific embedding techniques to capture its characteristics and spatial relationships. The sheer volume and multifaceted nature of modern geospatial data makes traditional analytical approaches insufficient and inefficient, because it often rely on manual intervention and hypothesis-driven exploration. There is a necessity of advanced and autonomous methods for such scale and complex geospatial data. The focus is not not just about handling more data, but about extracting deeper understanding of spatial relations from its inherent complexity.

Geospatial data, with its mix of geometric types, attribute tables (text), and topological information (point, line, polygon), does not readily fit into the input formats expected by many standard

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-5-W2-2025-131-2025

machine learning models. These basic representations, while effective for fundamental spatial analysis, have limitations in capturing the complex relationships and contextual information embedded in geospatial data. For instance, representing a city as a point provides limited information about its size, shape, or internal structure. Similarly, representing a road as a line does not capture information about its traffic flow, speed limits, or surrounding environment. Traditionally, geospatial data structures have relied on explicit, rigid formats, coordinate pairs, topological graphs, and raster grids. The focus of this paper is to explore the embeddings for fundamental primitive vector data formats, namely, points, lines, and polygons, as shown in Figure 2.

2.2 Text

Textual descriptions of locations, such as addresses, place names, or geographical features, can be embedded using natural language processing techniques. These embeddings help to capture the semantic meaning of the text and its association with spatial locations. For example, words like "behind," "to the east of," or "in front of park" can be embedded to reflect their spatial properties.

2.3 Images

Satellite imagery, aerial photographs, and other visual data can be embedded using computer vision techniques, such as convolutional neural networks (CNNs) or vision transformers. These embeddings capture the visual features and highlight spatial patterns in the images.

2.4 Points

Points represent discrete locations in space, such as countries, cities, landmarks, or individual places. Along with locations, each point can be associated with various attributes for more context, such as population, elevation, forests, or agricultural land. These attributes can be incorporated into its embedding to enable spatial query and analysis.

2.5 Lines

Lines represent linear features, such as roads, streams, boundaries or flight paths. Embeddings of lines can capture their length, direction, and connectivity, as well as attributes such as traffic flow, speed or road types.

2.6 Polygons

Polygons represent areas with boundaries such as cities, districts, lakes, parks, and forests. Each polygon contains both spatial and attribute data, that helps to study and analyse the underlying relationships, patterns, and distributions. Their embedding captures their geometric shape, attributes and their topological relation to other polygonal structure. This helps to explore various analysis such as shape-based classification or object detection.

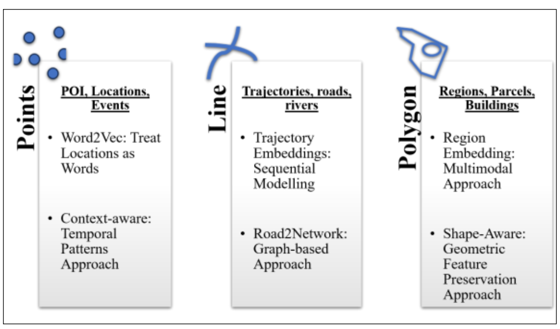

Figure 2: primitive vector data formats

3. Spatial Embedding

3.1 Overview

The transition from word embeddings to geospatial embeddings involves adapting methods to account for spatial relationships. This shift is necessitated by the growing recognition that spatial data is inherently different from textual data. Geospatial embeddings must capture not only the semantic meaning of locations but also their spatial relationships, such as proximity, accessibility, and regional characteristics. This evolution has led to the development of various geospatial embedding techniques that leverage spatial data sources, including GIS, satellite imagery, and social media data.

Spatial embeddings are effectively bridging the gap between unstructured textual descriptions of places (on the web, in social media, and in reports) and structured geospatial data (such as coordinates and GIS layers). This convergence enables a deeper semantic understanding of locations by machines. Instead of just knowing where a place is (its coordinates) or what its basic category is (e.g., "military site"), models leveraging semantic embeddings can understand qualitative aspects of the place derived from human language - for example, cultural aspects, in text. This leads to better semantic analysis capabilities. This marks a fundamental evolution in geospatial analysis: from analyzing discrete data about places to learning comprehensive representations of places themselves. Such representations encapsulate a more holistic understanding of geographic context, facilitating inferential approaches that consider multiple facets of geography.

TABLE 1: Comparison of geospatial embeddings

Method	Type	Techniques	Applications
GPS2Vec	Points	Kernel encoding	GPS encoding
Trajectory Embedding	Points/ Lines	RNN/ Transformer	Mobility prediction
Mob2Vec	Points	ST Patterns	Location Recommendation
Urban2Vec	Points	Hierarchal clustering	Urban POI
Graph-Based	Lines	GNN	Road networks
Poly2Vec	Polygon	NUFT	Shape class
Space2Vec	Polygon	Multi-scale	Region simulation
Mot2vec	Points	Word2Vec	Mobility pattern

3.2 Point Embedding

Points represent the simplest form of geospatial data. Points represent discrete locations (e.g., GPS coordinates, landmarks) and are foundational in geospatial analysis. Embedding methods for points aim to encode both spatial coordinates and contextual attributes. The goal of point embedding is to map these locations to vectors such that geometric and semantic relationships from the original spatial data are preserved as geometric relationships (typically proximity or distance) in the latent vector space. This allows for the efficient comparison, indexing, and classification of these locations. The concept is inspired by techniques like Word2Vec in NLP, where words are embedded into a vector space based on their context. In this adaptation, locations are treated as "words" and sequences of locations, such as trajectories, are treated as "sentences". Point embeddings aim to preserve essential spatial relationships and capture diverse forms of similarity. They represent the characteristics of locations, with highly related locations ideally being close in the vector space. A key challenge in embedding point data is maintaining high coordinate precision, which is critical for many applications. Techniques like discretization into grid cells, token location encoding, and incorporating additional sensor attributes are used to address this while compressing raw coordinates into compact embeddings. Five important methods are discussed below for point embedding.

GPS2Vec (McKenzie & Adams, 2018) is a kernel-based encoder that captures spatial proximity and contextual information for GPS data. By using kernel functions, GPS2Vec generates embeddings that preserve local and global spatial relationships. This method is effective for small-scale datasets but faces scalability challenges.

Trajectory Embeddings model sequential point data (e.g., GPS trajectories) using recurrent neural networks (RNNs) or transformers (Crivellari & Beinat, 2019). Temporal and spatial dependencies are encoded into vectors, enabling applications like mobility prediction and anomaly detection. When dealing with large-scale mobility data, such as GPS traces or check-in data from location-based social networks, Mot2vec can be applied. This method treats locations as "words" and trajectories as "sentences," using a Word2vec model to construct location embeddings that capture behavioural relationships between points [9]. This is particularly useful for understanding movement patterns and recommending POIs.

Mob2Vec (Yao, Zhang, Huang & Bi, 2017) leverages spatio-temporal trajectory patterns to generate embeddings that capture individual and collective mobility behaviours. Mob2Vec is particularly useful for location recommendation systems.

Urban2Vec (Liu, Liu, Li, & Li, 2020) uses hierarchical density-based clustering to stabilize POI (Point of Interest) embeddings across diverse urban regions. This method improves robustness in heterogeneous urban environments.

Loc2Vec (Sentience, 2025) is an early approach that encoded location surroundings by rasterizing the region around a point and using a deep convolutional neural network on a multi-channel tensor representing features like road networks, land cover, and amenities. It used a self-supervised method with a triplet loss based on defining positive (close) and negative (far) instances.

3.3 Line Embedding

Unlike points, line geometries are inherently sequential; the order and connectivity between the points define the overall shape of the line (e.g., a river, road or trajectory). Line data represents linear features in geographic space. Embedding techniques for lines aim to capture not only their spatial coordinates but also their shape, direction, length, and potentially associated attributes like speed limits or traffic flow. Basically, we require methods that preserve connectivity and flow within spatial networks. Hence, the spatial continuity and directional progression of vertices along a line are critical for accurate geometric representation. Traditional CNNs are less effective for line embeddings. Architectures like Recurrent Neural Networks (RNNs), Long Short-Term Memory networks (LSTMs), and Transformers are ideal because they capture dependencies along ordered sequences. Two key methods are analysed below for line embedding.

Graph-Based Approaches [5] model linear features as graphs, where nodes represent points along the line and edges encode connectivity. Graph Neural Networks (GNNs) learn embeddings by aggregating node and edge features, making them suitable for road network analysis.

Deep Learning for Trajectories (Feng, Li, Zhang, Sun, & Meng, 2018) includes RNNs and transformers that encode sequential line data (e.g., vehicle trajectories) by capturing temporal and spatial dependencies. These models enable route prediction and congestion detection. They can be used to process lines as sequences of points, capturing the temporal and spatial dependencies between them.

3.4 Polygon Embedding

Polygons have complex geometries that are more challenging to embed than points and polylines. Unlike a polyline, which is an ordered sequence of points, a polygon is represented by all the points within it. Polygons represent areal features (e.g., administrative boundaries, land parcels) and pose challenges due to their geometric complexity (e.g., holes, multipolygons) and topological relationships. These embeddings are designed to capture Topological relationships (e.g., adjacency, containment). They are invariant to geometric transformations like rotation, scaling, or vertex reordering. They retain Semantic meaning, such as land-use patterns or functional characteristics. Polygon topology can be preserved using graph-based approaches that represent a polygon as a graph, where vertices become nodes and edges capture boundary relationships. Two key methods are explained below for polygon embedding.

Poly2Vec (Mai, Janowicz & Yan, 2020) combines geometric and contextual encoding. NUFTspec method is a spectral domain encoder that leverages Non-Uniform Fourier Transformations (NUFT) to encode polygons, ensuring robustness to shape modifications while preserving topological features. This method outperforms spatial approaches in tasks like shape classification and spatial relation prediction. Conversely, spatial-domain approaches like ResNet1D use 1D convolutional neural networks (CNNs) with circular padding to achieve loop-origin invariance.

Space2Vec (Mai, Janowicz & Yan, 2020) employs neural networks and kernel functions to encode spatial relationships at local and global scales. This method is effective for tasks requiring multi-scale analysis, such as regional similarity measurement.

4. Evaluation Methods and Metrics

4.1 Evaluation Methods

Evaluating geospatial embeddings involves intrinsic and extrinsic approaches. Intrinsic evaluation measures the quality of embeddings based on their ability to preserve spatial relationships, using metrics like cosine similarity or Euclidean distance. Extrinsic evaluation assesses performance in downstream tasks, such as location prediction or POI recommendation, providing practical insights into embedding utility.

4.2 Metrics

Cosine Similarity is used metric to measure the similarity between two vectors in the embedding space. A higher cosine similarity (closer to 1) between the embedding vectors of two geospatial entities indicates that the model considers them more similar based on the learned representation. This allows quantitative measurement of relatedness based on mobility, function, or other embedded properties.

Distance Metrics in Embedding Space such as Euclidean distance and other distance measures are applied to vectors in the latent

https://doi.org/10.5194/isprs-annals-X-5-W2-2025-131-2025

space. The goal is often for the distance between vectors to be proportional to the semantic or geometric dissimilarity of the original geospatial objects.

Geometric Distance Metrics, as summarised in TABLE 1 are used to quantify the similarity between the geometric shapes of spatial objects themselves, rather than their embeddings. Key metrics include Hausdorff Distance, Euclidean Distance (between corresponding points), Fréchet Distance, Chamfer Distance, and F-Score. These can be used as benchmarks to evaluate if the embedding process preserves the original geometric relationships and distances.

TABLE 2: Metrics for Embeddings

Metric	Description
Hausdorff Distance	Measures the maximum deviation between two geometrical shapes.
Euclidean Distance	Evaluates the average distance between corresponding points in two geometries.
Fréchet Distance	Accounts for the ordering of points, quantifying similarity between trajectories.
Chamfer Distance and F-Score	Commonly applied in evaluating point cloud segmentation to compare predicted outputs against ground truth.

Each of these metrics provides different insights. Hausdorff distance is sensitive to extreme variations. Euclidean distance sheds light on the overall average behaviour of the geometric representations.

5. Observations and Insights

5.1 Observations

5.1.1 Points: Two key observations for point embeddings are *Scalability* and *Contextual Sparsity*. For Scalability case, kernel-based methods like GPS2Vec face computational bottlenecks with large-scale GPS datasets. In contextual sparsity, Trajectory embeddings often overlook semantic context (e.g., user demographics), limiting their utility in personalized recommendations.

5.1.2 Lines: Two key observations for line embeddings are *Topological Complexity* and *Dynamic Environments*. For topological complexity, current methods simplify intersections or bifurcations, losing critical network properties. In dynamic environments, Real-time updates (e.g., traffic accidents) are not yet handled effectively by static embeddings.

5.1.3 Polygons: Two key observations for polygons embeddings are *Topological Preservation* and *Multi-Scale Representation.* For topological preservation, simplifying polygons for computational efficiency often sacrifices containment or adjacency relationships. In Multi-Scale Representation, few methods dynamically adjust embeddings for tasks requiring both local detail (e.g., parcel boundaries) and global context (e.g., regional climate).

5.2 Insights

5.2.1 Interpretability vs. Scalability: Techniques like GPS2Vec and Poly2Vec prioritize either interpretability or the integration of geometric and semantic information. However, this often comes at the cost of scalability and topological accuracy. This highlights a conflict between computational efficiency and model precision.

5.2.1 Adaptability Issues in Rural Areas: Graph-based approaches like Mob2Vec struggle in low-data rural environments, demonstrating difficulties in adapting to diverse geographical conditions.

5.2.2 Limitations of Static Embeddings: Static embeddings are inadequate for dynamic scenarios, such as real-time traffic updates, which exposes their limitations in practical applications.

6. Conclusion

The future of geospatial embeddings is closely tied to advancements in machine learning techniques. As algorithms become more sophisticated, the ability to generate and analyze geospatial embeddings will improve, leading to more accurate and insightful analyses. Future embedding architectures might also incorporate heterogeneous data (e.g., imagery, text, sensor data) to create truly unified geospatial representations. This work has provided a comprehensive overview of existing methods for geospatial embeddings. By detailing their applications, challenges, and future directions, this paper underscores the potential of geospatial embeddings to revolutionize spatial analysis.

The evolution of geospatial embeddings addresses the limitations of traditional methods, such as points, lines, and polygons, which struggle to capture the complexity and interconnectedness of real-world spatial phenomena. The integration of geospatial embeddings with other data types, such as temporal data and social media data, presents exciting opportunities for comprehensive analyses. This paper seeks to serve as a systematic study of geospatial embeddings for researchers, practitioners, and policy-makers interested in GeoAI.

Acknowledgment

We extend our gratitude to members of our team for their timely support and advice. We are grateful to the director of our organisation for the continuous encouragement during the work.

References

Crivellari, A., & Beinat, E. (2019). From Motion Activity to Geo-Embeddings: Generating and Exploring Vector Representations of Locations, Traces and Visitors through Large-Scale Mobility Data. ISPRS International Journal of Geo-Information, 8(3), 134. https://doi.org/10.3390/ijgi8030134

Feng, J., Li, Y., Zhang, C., Sun, F., & Meng, F. (2018). DeepMove: Predicting human mobility with attentional recurrent networks. Proceedings of the 2018 World Wide Web Conference (WWW), 1459–1468. DOI:10.1145/3178876.3186058.

INDIGIS Team (2022), “INDIGIS, an Indigenous GIS for Defence Applications”, Technology Focus, Vol 30, Issue 2, April 2022, ISSN No. 0971 - 4413, www.drdo.gov.in.

Liu, X., Liu, Y., Li, X., & Li, Y. (2020). Urban2Vec: A hierarchical clustering approach for stable POI embeddings. Proceedings of the 28th International Conference on Advances in Geographic Information Systems (SIGSPATIAL), 1–4. DOI:10.1145/3397536.3422247.

McKenzie, G., & Adams, B. (2018). GPS2Vec: A kernel-based approach to location embedding. Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, 1–4. DOI:10.1145/3274895.3274976.

Mai, G., Janowicz, K., & Yan, B. (2020). Poly2Vec: Learning polygon embeddings for spatial similarity measurement.

Proceedings of the 28th International Conference on Advances in Geographic Information Systems (SIGSPATIAL), 1–4. DOI:10.1145/3397536.3422245.

Mai, G., Janowicz, K., & Yan, B. (2020). Space2Vec: A spatial embedding method for geographic information systems. Proceedings of the 28th International Conference on Advances in Geographic Information Systems (SIGSPATIAL), 1–4. DOI:10.1145/3397536.3422246.

Panigrahi N. (2014), Computing in Geographic Information Systems, 1st ed. Boca Raton, FL: CRC Press, 2014. ISBN 978-1-4822-2314-9.

Sentience (2025). Loc2Vec: Learning Location Embeddings with Triplet-loss Networks', https://sentiance.com/loc2vec-learning-location-embeddings-w-triplet-loss-networks

Yao, D., Zhang, C., Huang, J., & Bi, J. (2017). Mob2Vec: A trajectory embedding approach for mobility prediction. Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, 1–10. DOI:10.1145/3139958.3139965.

Zhou, J., Cui, G., Zhang, Z., Yang, C., Liu, Z., & Sun, M. (2020). Graph neural networks: A review of methods and applications. AI Open, 1, 57–81. DOI: 10.1016/j.aiopen.2021.01.001.