GeoKG for GeoFM: Neuro-Symbolic Integration of Semantically-Enriched Geospatial-Temporal Data

Kai Moltzen
AIX, Leuphana University Lüneburg, Germany

Debayan Banerjee
AIX, Leuphana University Lüneburg, Germany
HITeC e.V., Hamburg, Germany

Julian Burmester
AIX, Leuphana University Lüneburg, Germany

Anna Ehrenberg
AIX, Leuphana University Lüneburg, Germany

Martin Kohler
AIX, Leuphana University Lüneburg, Germany

Cedric Möller
University of Hamburg, Germany

Jann Pfeifer
AIX, Leuphana University Lüneburg, Germany

Tilahun Abedissa Taffa
AIX, Leuphana University Lüneburg, Germany

Hanna Marlene Schulz
AIX, Leuphana University Lüneburg, Germany

Aida Usmanova
AIX, Leuphana University Lüneburg, Germany

Patrick Westphal
AIX, Leuphana University Lüneburg, Germany
HITeC e.V., Hamburg, Germany

Ricardo Usbeck
AIX, Leuphana University Lüneburg, Germany
HITeC e.V., Hamburg, Germany

Abstract

Solving complex real-world problems requires integrating heterogeneous data. Although recent Geospatial Foundation Models (GeoFMs) excel at pixel-level analysis, they make only limited use of the geographic information contained in natural language texts and therefore often lack the semantic context needed for complex geospatial reasoning. This vision paper argues for reconceiving unstructured texts as a core component of a neuro-symbolic framework that populates a geospatial-temporal knowledge graph (GeoKG). We outline an architecture in which modality-specific extractors produce multi-granular footprints that are fused into a GeoKG, synchronizing symbolic representations with embeddings, geometries, and grid cells. This GeoKG supports geospatial reasoning and provides enriched training data for truly multimodal GeoFMs, for which we pose a set of research challenges.

2012 ACM Subject Classification Information systems → Ontologies; Computing methodologies → Spatial and physical reasoning; Computing methodologies → Information extraction; Computing methodologies → Learning latent representations

Keywords and phrases GeoKG, Foundation Models, GeoFM, neuro-symbolic Data Integration

Digital Object Identifier 10.4230/LIPIcs.COSIT.2026.20

Category Short Paper

Funding The authors acknowledge financial support from the Federal Ministry for Economic Affairs and Energy (BMWe) of Germany for the project CoyPu (project number 01MK21007G), and from the Ministry of Research and Education (BMBF) for the project "RESCUE-MATE: Dynamische Lageerstellung und Unterstützung für Rettungskräfte in komplexen Krisensituationen mittels Datenfusion und intelligenten Drohnenschwärmen" (FKZ 13N16844). This work is also supported by the German Research Foundation (DFG) project NFDI4DS under Grant No.: 460234259, and further funded by the German Federal Ministry for Research, Technology, and Space (BMFTR) under Grant No. 13N17698 (PROVIDER). This publication was produced as part of the project "Leuphana AI Campus (LAICA)", funded by the Foundation for Innovation in Higher Education.

Declaration on GenAI/LLM use During the preparation of this work, the authors used generative models hosted by the ChatAI service of the Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen (GWDG) and Grammarly to check grammar and spelling, paraphrase and reword, and improve writing style. After using these tools and services, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

1 Introduction

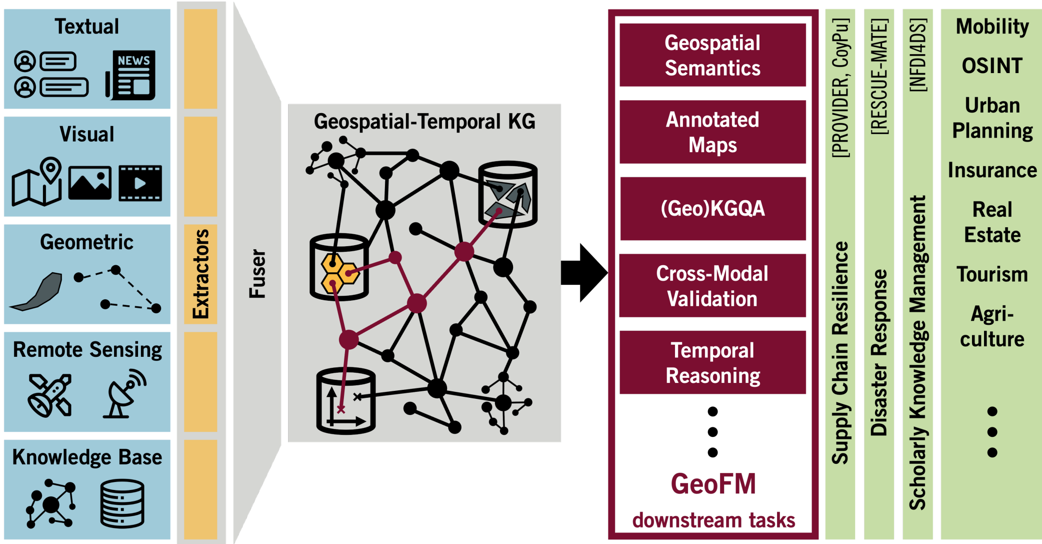

Figure 1 Vision architecture **GeoKG for GeoFM**. We integrate geospatial-temporal information extracted from textual, visual, RS, and geometric data sources via extractor modules and fuse the results into an ontology-based multimodal GeoKG, synchronized with cell, geometric, raster, and vector stores. The red links illustrate cross-modal alignment by geospatial relation, connecting grid cells, geometries, and embedding vectors within the GeoKG. The resulting enriched, spatially and temporally aligned data backbone is intended as training data for GeoFMs and diverse applications.

Addressing complex real-world challenges across domains, from disaster response and supply chain resilience to scholarly knowledge management, requires the integration of heterogeneous data. As we were able to demonstrate in interdisciplinary research consortia such as CoyPu, RESCUE-MATE, PROVIDER, and NFDI4DS[1], decision support in these domains relies on aligning diverse modalities: unstructured text, remote sensing (RS) imagery, structured geometrical data, and entire knowledge bases. However, a fundamental barrier remains: while these modalities can describe the same physical *reality*, they differ significantly in spatial granularity, temporal resolution, and semantic structure.

To bridge this gap, raw data must be geographically, temporally, and semantically aligned and enriched. Throughout, we use the term alignment in this data-centric sense: co-locating heterogeneous representations of the same place or event in a shared spatial, temporal, and semantic reference, and distinguish it from the separate notion of aligning model behavior with human values, which we discuss in Section 3.2. Semantic data enrichment is an established process for enhancing FAIRness [36] and providing meaning and context for data by means of an ontology [9, 35]. In this work, we envision a novel framework that extends enrichment beyond entity and relation extraction. Source- and modality-specific extractor modules derive the geospatial-temporal dimension of the input data at different granularities. By fusing enriched information into a multimodal geospatial-temporal Knowledge Graph (GeoKG), various input data can be semantically aggregated, filtered, and further processed using reliable spatiotemporal footprints.

[1] `https://coypu.org/en/`; `https://www.rescue-mate.de/`; `https://projekt-provider.de/`; `https://www.nfdi4datascience.de/`

Throughout this paper, we return to one running example to illustrate this framework: a flood that disrupts a supply chain. News reports describe inundated areas and explain how rescue services and organizations react to it based on past events; RS imagery delineates the flooded extent; in-situ sensors report rising water levels at regular timestamps; and geometric data locate the affected roads, warehouses, and supplier sites. An operational question such as "which product deliveries and transport routes are at risk, and why?" can only be answered by bringing all of these sources into a common spatial, temporal, and semantic frame.

In this vision paper, we argue that the current trajectory of Geospatial Artificial Intelligence (GeoAI) [9] is constrained by the limited role that geographic information from texts plays in the emerging generation of Geospatial Foundation Models (GeoFMs). While recent surveys indicate a surge in multimodal GeoFMs [37, 18], these models are predominantly trained on RS data. When text is included, it is often limited to vision-language tasks like image captioning or visual grounding [37]. Consequently, current systems struggle to perform complex semantic geospatial reasoning tasks such as question answering (GeoQA) [26, 27] that require understanding the "why" and "how" of an event – for instance, for whom a particular road is impassable and how that propagates to downstream deliveries in our flood scenario. Such information is rarely visible in pixels alone but abundant in textual narratives. Natural language (NL) texts can therefore enhance GeoFMs in reasoning, in the retrieval of geographically similar documents, and as a valuable source of information about events and place semantics [23].

To enable genuinely multimodal GeoFMs that reason across the full spectrum of spatial, temporal, and textual evidence, this vision paper argues for reconceiving texts containing geographical information as an integral component of geospatial-temporal, multimodal Knowledge Graphs (KGs), as shown in Figure 1. Such GeoKGs can in turn serve as a rich source of training data for the next generation of GeoFMs. We present concrete research questions for the community to discuss, and look forward to receiving feedback, as our framework will serve as the foundation for future research.
Our contributions are as follows:

1. **Unified geospatial-temporal extraction and fusion as a fundamental step:** Our extractors first derive information about the geospatial extent (e.g., polygons, Discrete Global Grid System (DGGS) cells) and temporal (meta-)data (e.g., validity intervals, event timestamps) of the diverse inputs, then fuse and map it to a central GeoKG.
2. **Neuro-symbolic data integration of all core concepts:** By coupling structured, symbolic KG construction with dense-vector embeddings and raster maps, our framework integrates all core concepts of spatial information [13] and supports semantic similarity search [28, 8] across modalities while preserving the logical rigor of ontological reasoning.
3. **Data backbone for GeoFMs:** The GeoKG provides a structured training ground of richly annotated, geospatially and temporally aligned textual descriptions linked to multi-granular spatial footprints. We expect this grounding to enhance truly multimodal GeoFMs and improve tasks such as GeoQA, reasoning, and cross-modal retrieval.

2 Neuro-Symbolic Geospatial-Temporal Architecture

The architecture in Figure 1 ranges from diverse input sources to domain-specific use cases. At its core, the GeoKG is designed for geospatial multimodality, interlinking KG nodes with textual descriptions to standardized grid cells, geometric shapes, or dense vector embeddings.

2.1 Input Sources

Geospatial-temporal data can be captured in various forms [17]. NL textual descriptions take the form of news articles, social media posts, or more formalized modes such as encyclopedic descriptions; in our running example, these include the news reports describing the flood and its consequences for the supply chain. These sources also provide visual modalities, such as images, videos, and illustrated maps. Street-view and drone imagery, as well as RS, are essential sources within the geospatial domain, with the latter typically further categorized into, e.g., optical RGB, spectral, or SAR [37]. Geometric data types like points, polylines, and (multi-)polygons represent, e.g., geographic regions, trajectories, or building footprints, such as the routes and warehouse outlines in the flood scenario.

Text is a prominent modality in KGs in general, but geographic and place information in texts is rarely captured in GeoKGs beyond entity linking. We suppose this is because such spatial language must first be expensively extracted and interpreted. Consequently, we rely on an explicit ingestion layer as a fundamental step for our GeoKG.

2.2 Ingestion Layer

Modality-specific extractors are not only necessary for aligning the various data types, but also for enriching and interpreting the information for representation in the GeoKG. For example, NL descriptions of real-world spatiotemporal phenomena are inherently vague. Therefore, complex location descriptions must be interpreted reliably, and the extents of, e.g., the impact area of natural disasters, estimated based on the context [19, 14]. In the flood scenario, an extractor must turn a phrase such as “the lower river valley is under water” into a concrete, multi-granular footprint. Other examples include the spatial and temporal aggregation of social media posts or narrative extraction [6]. After extraction, all information is aligned and enriched semantically within the fuser module, thereby mapping it to the GeoKG ontology.

2.3 Data Layer: GeoKG

We populate a GeoKG by reusing existing ontologies [4] (e.g., KWG-Ont [29], PROV-O [1], or CoyPu [31]). We expect to further develop novel ontologies to represent heterogeneous spatial data types [13, 27] with different degrees of granularity, temporality, and vagueness. Among those data types, existing GeoKGs focus on objects, events, and networks, but rarely consider fields. To enable fully four-dimensional, spatiotemporal ’hyperobjects’ [5], however, fields indicating presence in space-time constitute a promising means. Fields enable us to consider the temporal dimension of events more adequately than with mere object-timestamps, as common in spatiotemporal KGs [22]. Besides geometric shapes like points, lines, or polygons stored in well-known text (WKT) strings, we will thus implement synchronization with standardized cells within DGGSs [21], and dense vector embeddings for downstream usage.

A DGGS partitions the Earth into a hierarchy of nested, approximately equal-area cells, each with a stable identifier [25, 21], and this plays a central role in our design for three reasons. First, a cell identifier acts as a modality-independent join key: a flooded pixel in an RS scene, a polygon describing an inundated district, and a sentence georelated to a disaster area can all be attached to the same cell and thereby related without committing to a single geometry. Second, the cell hierarchy provides multi-granular footprints, so that coarse, vague textual descriptions and fine-grained sensor readings can coexist and be reconciled at a shared resolution. Third, aggregation over cells – for example, collecting all NL labels and embeddings that refer to a given area within a time window – yields compact, queryable summaries that function as supervision for downstream models. In our running example, the affected region is thus represented as a set of cells that links the news- and RS-derived event, extent, and the geometric footprints of the threatened warehouses and routes, while the corresponding text embeddings remain attached for similarity search. Fundamentally, this synchronization allows us to align arbitrary spatiotemporal representations, extending the foundational work on combining KGs and DGGS in the KnowWhereGraph project [40]. Therefore, our GeoKG constitutes a valuable source of training data for truly multimodal GeoFMs (see Section 3.1). Technically, this capability is based on the synchronization between a triple store for the KG, geospatial relational databases (DB), and DB extensions for vector embeddings and DGGS cells.

2.4 Interface and Service Layer

The interface layer provides a stable, unified interface to the GeoKG and external services. It offers direct graph access (e.g., SPARQL over named graphs), modality-specific endpoints for spatial queries and data access (e.g., GeoSPARQL), vector similarity search, and access to external resources and interfaces (e.g., routing, weather APIs). On top of the interface layer, the service layer provides reusable capabilities that can be shared across downstream domain applications. Services orchestrate multi-step operations that cannot be expressed as single queries or isolated store operations. Examples include services for geospatial semantics, temporal reasoning, cross-modal validation, geographic visualization, and semantic interfaces like GeoKG Question Answering (GeoKGQA), which incorporates spatial inference and enables discovery and reasoning about geospatial relationships between entities – for instance, answering which suppliers fall within the flooded cells and which routes connect them.

2.5 Domains

Each application implements domain-specific functions while reusing common services for geospatial-temporal grounding. The research project PROVIDER develops a decision-support system for strengthening supply chain resilience, the setting of our running example. Real-time news streams must be processed to extract events and explicitly link them to spatiotemporal footprints, such as facilities, routes, and assets, integrating existing knowledge on dependencies and constraints. This domain, therefore, relies on aligning heterogeneous sources and provenance information, supported by the GeoKG.

The research project RESCUE-MATE provides an exemplary disaster response application, the perspective from which the same flood is monitored as it unfolds. It maintains a digital twin updated from real-time sources (rescue radio traffic, environmental sensors), complemented by social media data to serve as the foundation for decision-making in crises. The shared geospatial-temporal services support aligning incoming observations with the GeoKG and enable aggregating local reports into an operational picture.

NFDI4DS constitutes a research consortium for scholarly knowledge management, for which our framework targets the extraction and linking of geographic aspects of research artifacts and their alignment with a DGGS. NFDI4DS will provide interfaces for geospatial-temporal searching and filtering of scholarly data, e.g., all papers on a certain regularly-flooded region across various research fields in a given time span.

Our framework will support further domains that depend on geospatial-temporal grounding and multimodal data integration, such as urban planning, insurance, real estate, tourism, agriculture, and mobility, all of which require integrating data with authoritative geospatial-temporal reference data.

3 Overarching Perspectives

3.1 Towards (Truly) Multimodal GeoFMs

Recent years have seen a surge in GeoFM publications and performance [18, 37, 16]. Advancements are driven by models such as Prithvi-EO-x [32] or AlphaEarthFoundations [2], which excel at pixel-level tasks, including segmentation and scene classification. However, the former model only considers RS modalities and no text at all; the latter relies mostly on Wikipedia articles with document-level point coordinates. Vision-language GeoFMs largely come down to RS visual grounding and image QA tasks [39, 15]. We observe that geographic information from texts so far plays only a minor role in the current generation of GeoFMs, despite transformer models for NL processing having sparked the recent surge of interest in foundation models in the first place.

We expect that the wealth of unstructured geographic insights captured in NL texts will contribute greatly to the next generation of truly multimodal GeoFMs. Individuals and societies attribute meaning to places, largely expressed through natural language and captured in texts. As Purves and colleagues argue [23], places shall thus be jointly represented by textual attributes and locations, for which the GeoKG depicts an ideal technology to store, align, and semantically enrich place information. Using the NL semantic labels aggregated by DGGS cells within the GeoKG will serve as a much-needed standardization. It constitutes an important step toward adopting GeoFMs, or earth embeddings, to NL processing tasks such as geographic information retrieval [11, 8], for which GeoFMs can then serve as dense prior information.

Returning to our running example, a query for documents about places similar to the flooded district can be answered by combining the textual descriptions and the earth embeddings aggregated per cell. Consequently, mitigation and disaster relief policies from these places might be transferred to the current impact area.

Our proposed architecture makes two contributions: (1) we include extractors to interpret and align geospatial-temporal information before ingestion into a GeoKG, and (2) we design the GeoKG to represent and synchronize different data types, accounting for the heterogeneity of spatial data as an important data source for truly multimodal GeoFMs capable of complex geospatial reasoning and retrieval.

3.2 Representation of Heterogeneous Geographic Data

Geographical data is the translation of the complexity of spatial *realities* into computable representations. Because of this dependence on users' cultures [33], pluralistic values, or legal regulations [7], we acknowledge that committing our framework to responsible AI practices alone is insufficient to address systemic power imbalances [24, 3]. Drawing on feminist AI literature [34], we incorporate principles such as critical assessment of power structures, consideration of environmental impacts, context, and consent [12]. Yoo [38] translates GeoAI-specific principles into a practical framework consisting of guidelines and checklists. The framework draws on contributions from industry, academia, and global governance institutions. It addresses risks that are particular to GeoAI, such as linkage attacks, surveillance, and geographic representational imbalances, resulting in twelve ethical GeoAI axes. In our future work, we aim to test this framework in practice by applying the proposed checklists.

Further, the synchronization and representation across spatial data formats help us to draw on geographical fuzziness as an important source of knowledge rather than as mere noise to be removed. Vague extents such as the referent of a place name or the cognitive boundary of a region carry meaning of their own [20], and our shared representation preserves this graded structure instead of forcing exact boundaries. Indigenous perspectives on tribal land, e.g., create not only different data, but can also show alternative concepts of spatial representations. The dynamic nature of land-shaping forces, the fractal nature of boundaries, or the uneven distribution of data collection result in inherent uncertainty. Place names are rarely neutral but bound to hierarchies and power through their implicit political nature. Distinguishing, e.g., between coarse cells or clearly defined boundaries, in combination with the semantical richness of textual descriptions in the GeoKG, enables the representation of ambiguous [10] or contested regions [30] to account for this heterogeneity.

4 Conclusions Towards Future Research Agendas

Realizing the proposed framework leads to a set of foundational research questions. On the representation side, we need principled models of time and change, identity across evolving geometries, and a clear semantics of events with potentially fragmented or recurrent time spans. On the learning side, we require embedding spaces that integrate heterogeneous geospatial-temporal modalities, including GeoKG embeddings, and causal or neuro-symbolic representations that capture more than surface-level similarity. On the integration side, a key challenge is bidirectional translation between symbolic and neural representations, namely, when and how to map learned representations back into structured representations and formats used in GIS workflows. Finally, for GeoFMs, we lack consensus on training tasks, supervision signals, and evaluation protocols that reward complex semantic geospatial reasoning, provenance, and auditability, while meeting FAIRness requirements towards heterogeneity, representativeness, privacy, and accountable use in high-stakes settings.

References

1 Khalid Belhajjame, James Cheneyand David Corsar, Daniel Garijo, Stian Soiland-Reyes, Stephan Zednik, and Jun Zhao. PROV-O: The PROV ontology. W3C recommendation, World Wide Web Consortium (W3C), April 2013. URL: `https://www.w3.org/TR/2013/REC-prov-o-20130430/`.

2 Christopher F. Brown, Michal Kazmierski, Valerie Pasquarella, and colleagues. AlphaEarth Foundations: An embedding field model for accurate and efficient global mapping from sparse label data. *CoRR*, abs/2507.22291, 2025. `doi:10.48550/arXiv.2507.22291`.

3 Eleanor Drage, Kerry McInerney, and Jude Browne. Engineers on responsibility: feminist approaches to who's responsible for ethical AI. *Ethics and Information Technology*, 26(1):4, 2024. `doi:10.1007/s10676-023-09739-1`.

4 Mariano Fernández-López, Asunción Gómez-Pérez, and Mari Carmen Suárez-Figueroa. Methodological guidelines for reusing general ontologies. *Data Knowl. Eng.*, 86:242–275, 2013. `doi:10.1016/J.DATAK.2013.03.006`.

5 Antony Galton. Fields and objects in space, time, and space-time. *Spatial Cogn. Comput.*, 4(1):39–68, 2004. `doi:10.1207/s15427633scc0401_4`.

6 Junbo Huang and Ricardo Usbeck. Narration as functions: from events to narratives. In *Proceedings of the 6th Workshop on Narrative Understanding*, pages 1–7, Miami, Florida, USA, November 2024. Association for Computational Linguistics. `doi:10.18653/v1/2024.wnu-1.1`.

7 Krzysztof Janowicz, Zilong Liu, Gengchen Mai, Zhangyu Wang, Ivan Majic, Alexandra Fortacz, Grant McKenzie, and Song Gao. Whose Truth? Pluralistic Geo-Alignment for (Agentic) AI. In *SIGSPATIAL 2025*, pages 799–803. ACM, 2025. `doi:10.1145/3748636.3760465`.

8 Krzysztof Janowicz, Martin Raubal, and Werner Kuhn. The semantics of similarity in geographic information retrieval. *Journal of Spatial Information Science*, 2(1):29–57, May 2011. `doi:10.5311/JOSIS.2011.2.3`.

9 Krzysztof Janowicz, Simon Scheider, Todd Pehle, and Glen Hart. Geospatial semantics and linked spatiotemporal data - past, present, and future. *Semantic Web*, 3(4):321–332, 2012. doi:10.3233/SW-2012-0077.

10 Zhe Jiang, Yan Li, Shashi Shekhar, Lian P. Rampi, and Joseph F. Knight. Spatial ensemble learning for heterogeneous geographic data with class ambiguity: A summary of results. In *SIGSPATIAL 2017*, pages 23:1–23:10. ACM, 2017. doi:10.1145/3139958.3140044.

11 Christopher B. Jones and Ross S. Purves. Geographical information retrieval. *International Journal of Geographical Information Science*, 22(3):219–228, 2008. doi:10.1080/13658810701626343.

12 Lauren F. Klein and Catherine D'Ignazio. Data feminism for AI. In *FAccT 2024*, pages 100–112. ACM, 2024. doi:10.1145/3630106.3658543.

13 Werner Kuhn. Core concepts of spatial information for transdisciplinary research. *Int. J. Geogr. Inf. Sci.*, 26(12):2267–2276, 2012. doi:10.1080/13658816.2012.722637.

14 Egoitz Laparra and Steven Bethard. A dataset and evaluation framework for complex geographical description parsing. In *COLING 2020*, pages 936–948. ICCL, 2020. doi:10.18653/V1/2020.COLING-MAIN.81.

15 Sylvain Lobry, Diego Marcos, Jesse Murray, and Devis Tuia. RSVQA: visual question answering for remote sensing data. *IEEE Trans. Geosci. Remote. Sens.*, 58(12):8555–8566, 2020. doi:10.1109/TGRS.2020.2988782.

16 Gengchen Mai, Weiming Huang, Jin Sun, and colleagues. On the opportunities and challenges of foundation models for geoai (vision paper). *ACM Trans. Spatial Algorithms Syst.*, 10(2):11, 2024. doi:10.1145/3653070.

17 Gengchen Mai, Xiaobai Angela Yao, Yiqun Xie, Jinmeng Rao, Hao Li, Qing Zhu, Ziyuan Li, and Ni Lao. SRL: towards a general-purpose framework for spatial representation learning. In *SIGSPATIAL 2024*, pages 465–468. ACM, 2024. doi:10.1145/3678717.3691246.

18 Valerio Marsocci, Yuru Jia, Georges Le Bellier, and colleagues. Pangaea: Assessing geospatial foundation models capabilities through a global and inclusive benchmark. *IEEE Geoscience and Remote Sensing Magazine*, pages 2–43, 2025. doi:10.1109/MGRS.2025.3628194.

19 Kai Moltzen, Junbo Huang, and Ricardo Usbeck. LLM Agents for Georelating - A New Task for Locating Events. In *SIGSPATIAL 2025*, pages 277–280. ACM, 2025. doi:10.1145/3748636.3762733.

20 Daniel R. Montello, Michael F. Goodchild, Jonathon Gottsegen, and Peter Fohl. Where's Downtown?: Behavioral Methods for Determining Referents of Vague Spatial Queries. *Spatial Cognition & Computation*, 3(2-3):185–204, September 2003. doi:10.1080/13875868.2003.9683761.

21 OGC. Discrete global grid systems abstract specification, 01.08.2017. URL: https://docs.ogc.org/as/15-104r5/15-104r5.html.

22 Philipp Plamper, Hanna Köpcke, and Anika Groß. A survey on spatio-temporal knowledge graph models. *CoRR*, abs/2512.16487, 2025. doi:10.48550/arXiv.2512.16487.

23 Ross S. Purves, Stephan Winter, and Werner Kuhn. Places in information science. *J. Assoc. Inf. Sci. Technol.*, 70(11):1173–1182, 2019. doi:10.1002/asi.24194.

24 Paola Ricaurte and Mariel Zasso. AI, ethics, and coloniality: A feminist critique. In *What AI Can Do: Strengths and Limitations of Artificial Intelligence*, pages 39–58. Chapman and Hall/CRC, New York, NY, 2023. doi:10.1201/b23345.

25 Kevin Sahr, Denis White, and A. Jon Kimerling. Geodesic discrete global grid systems. *Cartography and Geographic Information Science*, 30(2):121–134, 2003. doi:10.1559/152304003100011090.

26 S. Scheider, A. Ballatore, and R. Lemmens. Finding and sharing GIS methods based on the questions they answer. *International Journal of Digital Earth*, 12(5):594–613, May 2019. doi:10.1080/17538947.2018.1470688.

27 Simon Scheider, Rogier Meerlo, Vedran Kasalica, and Anna-Lena Lamprecht. Ontology of core concept data types for answering geo-analytical questions. *Journal of Spatial Information Science*, 20(1):167–201, June 2020. `doi:10.5311/JOSIS.2020.20.555`.

28 Angela Schwering. Approaches to Semantic Similarity Measurement for Geo-Spatial Data: A Survey. *Transactions in GIS*, 12(1):5–29, February 2008. `doi:10.1111/j.1467-9671.2008.01084.x`.

29 Cogan Shimizu, Shirly Stephen, Rui Zhu, and colleagues. The KnowWhereGraph Ontology: A showcase. In *FOIS 2023*, volume 3637 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2023. URL: `https://ceur-ws.org/Vol-3637/paper46.pdf`.

30 Alistair Sisson. Territory and territorial stigmatisation: On the production, consequences and contestation of spatial disrepute. *Progress in Human Geography*, 45(4):659–681, 2021. `doi:10.1177/0309132520936760`.

31 Nils Steinert. CoyPu Ontology. Revision: 2.3, 2023. URL: `https://schema.coypu.org/global/2.3`.

32 Daniela Szwarcman, Sujit Roy, Paolo Fraccaro, and colleagues. Prithvi-EO-2.0: A Versatile Multitemporal Foundation Model for Earth Observation Applications. *IEEE Transactions on Geoscience and Remote Sensing*, 64:1–20, 2026. `doi:10.1109/TGRS.2025.3642610`.

33 Yan Tao, Olga Viberg, Ryan S. Baker, and René F. Kizilcec. Cultural bias and cultural alignment of large language models. *PNAS Nexus*, 3(9):346, September 2024. `doi:10.1093/pnasnexus/pgae346`.

34 Sophie Toupin. Shaping feminist artificial intelligence. *New Media Soc.*, 26(1):580–595, 2024. `doi:10.1177/14614448221150776`.

35 Mike Uschold and Michael Gruninger. Ontologies: principles, methods and applications. *Knowl. Eng. Rev.*, 11(2):93–136, 1996. `doi:10.1017/S0269888900007797`.

36 Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, and colleagues. The FAIR guiding principles for scientific data management and stewardship. *Scientific Data*, 3(1):1–9, 2016. `doi:10.1038/sdata.2016.18`.

37 Liling Yang, Ning Chen, Jun Yue, Yidan Liu, Jiayi Ma, Pedram Ghamisi, Antonio Plaza, and Leyuan Fang. Survey of multimodal geospatial foundation models: Techniques, applications, and challenges. arXiv preprint, 2025. `doi:10.48550/arXiv.2510.22964`.

38 Suhong Yoo. An operational ethical framework for geoai: A prisma-based systematic review of international policy and scholarly literature. *ISPRS International Journal of Geo-Information*, 15:51, January 2026. `doi:10.3390/ijgi15010051`.

39 Yang Zhan, Zhitong Xiong, and Yuan Yuan. RSVG: exploring data and models for visual grounding on remote sensing data. *IEEE Trans. Geosci. Remote. Sens.*, 61:1–13, 2023. `doi:10.1109/TGRS.2023.3250471`.

40 Rui Zhu, Cogan Shimizu, Shirly Stephen, and colleagues. The KnowWhereGraph: A Large-Scale Geo-Knowledge Graph for Interdisciplinary Knowledge Discovery and Geo-Enrichment. *Transactions in GIS*, 30(1):e70184, 2026. `doi:10.1111/tgis.70184`.