

Symbolic and Subsymbolic GeoAI: Geospatial Knowledge Graphs and Spatially Explicit Machine Learning

Gengchen Mai¹, Yingjie Hu², Song Gao³, Ling Cai⁴, Bruno Martins⁵, Johannes Scholz⁶, Jing Gao⁷, and Krzysztof Janowicz^{8,4}

¹Spatially Explicit Artificial Intelligence Lab, Department of Geography,
University of Georgia, Athens, Georgia, 30605, USA
gengchen.mai25@uga.edu

²GeoAI Lab, Department of Geography, University at Buffalo, Buffalo, New
York, 14260, USA

³GeoDS Lab, Department of Geography, University of Wisconsin-Madison,
Madison, Wisconsin, 53715, USA

⁴Center for Spatial Studies, University of California, Santa Barbara, Santa
Barbara, California, 93105, USA

⁵Instituto Superior Técnico and INESC-ID, University of Lisbon, Lisbon,
1649-004, Portugal

⁶Institute of Geodesy, Graz University of Technology, Graz, 8010, Austria

⁷Department of Geography and Spatial Sciences, University of Delaware,
Newark, Delaware, 19716, USA

⁸Department of Geography and Regional Research, University of Vienna,
Vienna, 1040, Austria

1 Symbolic and Subsymbolic AI

The field of Artificial Intelligence (AI) can be roughly divided into two branches: Symbolic Artificial Intelligence and Connectionist Artificial Intelligence (or so-called Subsymbolic AI). Symbolic AI focuses on research based on classical logic and higher-level symbolic (human-readable) knowledge representations. It posits the use of declarative knowledge in reasoning and learning as critical to producing intelligent behavior (Goel, 2022). Examples are logical inference, symbolic reasoning, ontology engineering (Hobbs and Pan, 2006; Allemang and Hendler, 2011; Battle and Kolas, 2011; Hu et al., 2013; Janowicz et al., 2019), and in part Knowledge Graphs (KGs) (Singhal, 2012; Hoffart et al., 2013; Noy et al., 2019; Janowicz et al., 2022; Zhu et al., 2022b). In contrast, Connectionist AI advocates the idea of explaining intelligence using artificial neural networks (van Eijck and Visser, 2012). It assumes that the learning of associations or correlations from data (with little or no prior knowledge) is crucial for understanding mental phenomena (Goel, 2022). Examples are deep learning and representation learning techniques (LeCun et al., 2015; Vaswani et al., 2017; Hamilton et al., 2017; Goodfellow et al., 2020; Mai et al., 2020c), which are currently the dominating techniques in AI. Current breakthroughs in narrow AI are largely due to these technological successes in designing artificial neural networks, while having access to large-scale datasets for training, and with high-performance computing being readily available (Goel, 2022; Janowicz et al., 2020).

More recently, an increasing number of studies have been trying to combine these two approaches to develop scalable and interpretable AI models, so called Neural-Symbolic AI. Examples include AlphaGo Zero (Dalgaard et al., 2020) and Neural Symbolic Machines (Liang et al., 2017), which made use of symbolic reasoning in conjunction with neural networks in a reinforcement learning framework for the game Go, or knowledge graph-based question answering. Knowledge graph embeddings (Nickel et al., 2012; Dong et al., 2014; Wang et al., 2014; Nickel et al., 2016; Hamilton et al., 2018; Mai et al., 2019) and relational machine learning (Schlichtkrull et al., 2018; Nickel et al., 2015; Battaglia et al., 2018) are other typical examples.

2 Symbolic and Subsymbolic GeoAI

This special issue collection is the continuation of efforts that began with the GIScience’2021 workshop on “Methods, Models, and Resources for Geospatial Knowledge Graphs and GeoAI (GeoKG & GeoAI 2021),” which were followed by an open call for full-paper submissions.

From a geospatial point-of-view, Geospatial Artificial Intelligence (GeoAI), as an interdisciplinary field combining Geography, GIScience, and AI, advocates the idea of developing and utilizing AI techniques in geography and earth science research (Janowicz et al., 2020). This, in turn, benefits downstream tasks in health (Kamel Boulos et al., 2019), urban studies (Liu and Biljecki, 2022), traffic prediction (Polson and Sokolov, 2017), earth system science (Ham et al., 2019), and so on. GeoAI can also be categorized into two branches: Symbolic GeoAI and Subsymbolic GeoAI.

2.1 Symbolic GeoAI: Geospatial Knowledge Graphs

One research direction in Symbolic GeoAI are Geospatial Knowledge Graphs (GeoKG) (Kuhn et al., 2014; Janowicz et al., 2022). GeoKGs, as symbolic representations of geospatial knowledge, are at the core of GeoAI and facilitate many intelligent applications such as geographic question answering (Chen et al., 2013; Mai et al., 2020a; Scheider et al., 2021; Nyamsuren et al., 2021; Kuhn et al., 2021), geospatial knowledge summarization (Yan et al., 2019), geospatial data integration (Trisedya et al., 2019; Sun et al., 2021; Bernard et al., 2022), or geographic knowledge discovery (Li et al., 2012; Hu et al., 2015; Jiang et al., 2018; Mai et al., 2020b; Park and Lee, 2022). In fact, geospatial data play an important role in the Linked Open Data cloud (Kuhn et al., 2014), since space and time are essential for describing the relationships among events, people, and objects. However, many knowledge graphs treat geographic entities as ordinary entities in which the spatial footprints of places and the unique spatial operations (e.g., cardinal directions, topological relations, routing operations) are neglected, which leads to suboptimal and unsatisfactory results. Compared with generic knowledge graph research, there are many unique research questions to be asked for GeoKGs, such as how to model the evolution of geographic entities (e.g., administrative regions) across space and time (Kauppinen et al., 2008; Bernard et al., 2018, 2022), how to perform spatial-aware place retrieval and place similarity assessment (Chen et al., 2013; Santos et al., 2018; Mai et al., 2020a; Park and Lee, 2022), how to take into account uncertainty and vagueness when computing (spatial) associations among geographic entities for geographic data integration and question answering (Santos et al., 2018; Regalia et al., 2019; Mai et al., 2022b), and so on. This special issue includes two papers that answer some of those questions.

In the first article, Bernard et al. (2022) present a framework called *Theseus* for automatic creation of geospatial knowledge graphs (GeoKGs) for geographic divisions as well as their evolution over time. The paper addresses three main questions: (1) how to represent geographic divisions in a GeoKG? (2) how to represent their evolution over time within a GeoKG? (3) how to detect those territorial changes and create GeoKGs automatically? The *Theseus* framework was built to answer those three questions. Given two versions of the geographic divisions at different times, represented as two Shapefiles, *Theseus* first uses a data pre-processing application named *Theseus Prep* to transform each Shapefile into a pre-formatted TSN-shapefile format, as well as a *Theseus* PostGIS database which are aligned with their defined territorial statistical nomenclature (TSN)-ontology. Next, the geographic division data are triplified into KGs with the TSN Ontology. Then, an automatic change detection algorithm is used to detect and create the KG of geographic division changes by using a TSN Change ontology. Finally, a Graphical User Interface (GUI) is developed to explore and query the geographic divisions and their evolution over time. *Theseus* represents a holistic framework for geographic division evolution over time, from data collection to automatic change detection, KG construction, and KG visualization.

In the second article, Park and Lee (2022) present a Spatial-Semantic integrated Indoor Knowledge Graph (*SSIKG*) for place-based retrieval and route guidance in indoor environments. *SSIKG* contains two layers, namely a spatial layer and a semantic layer. The spatial layer is a symbolic representation of geographic entities in the indoor environment, as well as the navigation paths among them. The semantic layer includes the semantic information about indoor places, such as place categories, opening times, user reviews, keywords, and the similarity scores between entities. By representing both the spatial and semantic information of indoor places into a KG, the authors show that *SSIKG* can answer various questions regarding indoor place-based retrieval and route

guidance, by considering both spatial proximity, connectivity, and semantic similarity.

2.2 Subsymbolic GeoAI: Spatially Explicit Machine Learning

In terms of the other branch, we have that the use of subsymbolic AI approaches such as deep neural networks, to solve geospatial problems, is also a common component of GeoAI research. Although some existing deep learning architectures for tasks such as image classification, image segmentation, question answering, modeling language and vision, or entity recognition, can be readily used for GeoAI tasks such as classification and object detection in remote sensing images (Bastani et al., 2022; Camps-Valls et al., 2021), land use classification (Camps-Valls et al., 2021), geographic question answering and question answering over earth observation products (Coelho et al., 2021; Silva et al., 2022), remote sensing image captioning (Ramos and Martins, 2022), or place name recognition and resolution (Kulkarni et al., 2021; Cardoso et al., 2021; Liu et al., 2022), some unique challenges emerge which require special model designs, training objectives, or data pre-processing techniques, for instance by incorporating spatial principles and spatial inductive biases. We call this kind of practices *spatially explicit machine learning* (Yan et al., 2017, 2019; Mai et al., 2020c; Janowicz et al., 2020; Li et al., 2021; Zhu et al., 2022a; Mai et al., 2022b; Li et al., 2021). Some of the unique challenges include how to represent different types of spatial data into the embedding/subsymbolic space (Mai et al., 2022c), how to achieve geographic generalizability for a given machine learning/deep learning model (Goodchild and Li, 2021; Li et al., 2022), how to perform transfer learning across space and tasks (Fibæk et al., 2022), how to avoid geographic biases in GeoAI models (Liu et al., 2022), and so forth. This special issue also includes two innovative papers that tackle two questions we mentioned above.

The third article in this special issue collection, by Li et al. (2022), proposes a novel few-shot transfer learning (FSTL) method to improve the geographic generalizability of deep learning models for OpenStreetMap (OSM) missing building footprint detection. After the supervised training of a building detection model on fully labeled regions, FSTL is used to transfer the learned knowledge to label-sparse regions (e.g., Cameroon and Mozambique). Previous research shows that the generalizability (including replicability) of deep learning models across geographic space is weak (Goodchild and Li, 2021). Li et al. (2022) demonstrate that, compared to the state-of-the-art approaches such as Google Open Building (Sirko et al., 2021) and Facebook’s High Resolution Settlement Layer (HRSL) (Tiecke et al., 2017), the proposed few-shot learning approach achieves better accuracy in new label-sparse areas with just a couple of few-shot samples.

The fourth article, from Fibæk et al. (2022), presents a deep learning model for population estimation in areas geographically distinct from Northern Europe. The two major contributions of this paper are: (1) the authors demonstrate how to use the same deep learning architecture and transfer learning to solve three tasks – structure area prediction, structure type classification, and population prediction for Ghana and Egypt based on Sentinel data; (2) the authors show how to use multi-sensor data to produce high resolution population estimation for both daytime and nighttime, using deep learning.

3 Conclusion and Next Steps

In this editorial, we first discussed the differences and connections between Symbolic AI and Subsymbolic AI. Next, we highlight prominent components of Symbolic GeoAI and Subsymbolic GeoAI – Geospatial Knowledge Graphs and Spatially Explicit Machine Learning. Four papers are included in this special issue, with two papers focusing on GeoKGs, while the other two address spatially explicit machine learning. We believe that both components are indispensable for GeoAI research. Special thanks are owed to the authors, and especially to all of the reviewers for their informative and insightful peer reviews.

We believe that this special issue collection provides a good overview of the current state-of-the-art in GeoAI research. Additionally, the contributions also raise many interesting GeoAI research questions such as: **(1)** How to combine deductive methods from symbolic GeoAI with the representations and induction from deep learning models used by subsymbolic GeoAI to build neuro-symbolic GeoAI models? **(2)** To improve model generalizability across space and time, instead of using few-shot transfer learning method as [Li et al. \(2022\)](#) did, can we directly learn a hypernetwork to simulate how the model’s parameters change based on the location and time with meta-learning method ([Tenzer et al., 2022](#))? **(3)** Given the increasing popularity of foundation models (FMs) in the natural language and vision communities, such as GPT-3 ([Brown et al., 2020](#)), CLIP ([Radford et al., 2021](#)), PaLM ([Wei et al., 2022](#)), and DALL·E2 ([Ramesh et al., 2022](#)), could we build a foundation model for GeoAI which, after pretrained, can be easily adapted to multiple symbolic GeoAI and subsymbolic GeoAI tasks, involving the use of different data modalities ([Mai et al., 2022a](#))? **(4)** How can we address important ethical aspects, such as better accounting for and mitigating issues of bias, fairness, and transparency ([Shin and Basiri, 2022](#); [Zheng and Sieber, 2022](#)), how to reduce the environmental footprint of model training, and how to better connect to communities studying ethics of technology ([Goodchild et al., 2022](#)).

References

- Allemang, D., Hendler, J., 2011. Semantic web for the working ontologist: effective modeling in RDFS and OWL. Elsevier.
- Bastani, F., Wolters, P., Gupta, R., Ferdinando, J., Kembhavi, A., 2022. Satlas: A large-scale, multi-task dataset for remote sensing image understanding. arXiv preprint arXiv:2211.15660 .
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al., 2018. Relational inductive biases, deep learning, and graph networks. arXiv preprint arXiv:1806.01261 .
- Battle, R., Kolas, D., 2011. GeoSPARQL: enabling a geospatial semantic web. *Semantic Web Journal* 3 (4), 355–370.
- Bernard, C., Plumejeaud-Perreau, C., Villanova-Oliver, M., Gensel, J., Dao, H., 2018. An ontology-based algorithm for managing the evolution of multi-level territorial partitions. In: *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. pp. 456–459.

- Bernard, C., Villanova-Oliver, M., Gensel, J., 2022. Theseus: A framework for managing knowledge graphs about geographical divisions and their evolution. *Transactions in GIS* .
- Brown, T., et al., 2020. Language models are few-shot learners. *NIPS 2020* 33, 1877–1901.
- Camps-Valls, G., Tuia, D., Zhu, X. X., Reichstein, M., 2021. *Deep learning for the Earth Sciences: A comprehensive approach to remote sensing, climate science and geosciences*. John Wiley & Sons.
- Cardoso, A. B., Martins, B., Estima, J., 2021. A novel deep learning approach using contextual embeddings for toponym resolution. *ISPRS International Journal of Geo-Information* 11 (1), 28.
- Chen, W., Fosler-Lussier, E., Xiao, N., Raje, S., Ramnath, R., Sui, D., 2013. A synergistic framework for geographic question answering. In: *Proceedings of the 7th IEEE International Conference on Semantic Computing*. IEEE, pp. 94–99.
- Coelho, J., Magalhães, J., Martins, B., 2021. Improving neural models for the retrieval of relevant passages to geographical queries. In: *Proceedings of the 29th International Conference on Advances in Geographic Information Systems*. pp. 268–277.
- Dalgaard, M., Motzoi, F., Sørensen, J. J., Sherson, J., 2020. Global optimization of quantum dynamics with AlphaZero deep exploration. *npj Quantum Information* 6 (1), 1–9.
- Dong, X., Gabrilovich, E., Heitz, G., Horn, W., Lao, N., Murphy, K., Strohmann, T., Sun, S., Zhang, W., 2014. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 601–610.
- Fibæk, C. S., Keßler, C., Arsanjani, J. J., Trillo, M. L., 2022. A deep learning method for creating globally applicable population estimates from sentinel data. *Transactions in GIS* .
- Goel, A., 2022. Looking back, looking ahead: Symbolic versus connectionist AI. *AI Magazine* 42 (4), 83–85.
- Goodchild, M., Appelbaum, R., Crampton, J., Herbert, W., Janowicz, K., Kwan, M.-P., Michael, K., Alvarez León, L., Bennett, M., Cole, D. G., et al., 2022. A white paper on locational information and the public interest .
- Goodchild, M. F., Li, W., 2021. Replication across space and time must be weak in the social and environmental sciences. *PNAS* 118 (35).
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2020. Generative adversarial networks. *Communications of the ACM* 63 (11), 139–144.
- Ham, Y.-G., Kim, J.-H., Luo, J.-J., 2019. Deep learning for multi-year enso forecasts. *Nature* 573 (7775), 568–572.

- Hamilton, W., Bajaj, P., Zitnik, M., Jurafsky, D., Leskovec, J., 2018. Embedding logical queries on knowledge graphs. In: *Advances in Neural Information Processing Systems*. pp. 2026–2037.
- Hamilton, W., Ying, Z., Leskovec, J., 2017. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems* 30.
- Hobbs, J. R., Pan, F., 2006. Time ontology in owl. *W3C working draft 27 (133)*, 3–36.
- Hoffart, J., Suchanek, F. M., Berberich, K., Weikum, G., 2013. YAGO2: A spatially and temporally enhanced knowledge base from wikipedia. *Artificial Intelligence* 194, 28–61.
- Hu, Y., Janowicz, K., Carral, D., Scheider, S., Kuhn, W., Berg-Cross, G., Hitzler, P., Dean, M., Kolas, D., 2013. A geo-ontology design pattern for semantic trajectories. In: *International Conference on Spatial Information Theory*. Springer, pp. 438–456.
- Hu, Y., Janowicz, K., Prasad, S., Gao, S., 2015. Metadata topic harmonization and semantic search for linked-data-driven geoportals: A case study using arcgis online. *Transactions in GIS* 19 (3), 398–416.
- Janowicz, K., Gao, S., McKenzie, G., Hu, Y., Bhaduri, B., 2020. GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. *International Journal of Geographical Information Science* 34 (4), 625–636.
- Janowicz, K., Haller, A., Cox, S. J., Le Phuoc, D., Lefrançois, M., 2019. Sosa: A lightweight ontology for sensors, observations, samples, and actuators. *Journal of Web Semantics* 56, 1–10.
- Janowicz, K., Hitzler, P., Li, W., Rehberger, D., Schildhauer, M., Zhu, R., Shimizu, C., Fisher, C. K., Cai, L., Mai, G., et al., 2022. Know, Know Where, KnowWhereGraph: A densely connected, cross-domain knowledge graph and geo-enrichment service stack for applications in environmental intelligence. *AI Magazine* 43 (1), 30–39.
- Jiang, Y., Li, Y., Yang, C., Hu, F., Armstrong, E. M., Huang, T., Moroni, D., McGibbney, L. J., Finch, C. J., 2018. Towards intelligent geospatial data discovery: a machine learning framework for search ranking. *International Journal of Digital Earth* 11 (9), 956–971.
- Kamel Boulos, M. N., Peng, G., VoPham, T., 2019. An overview of GeoAI applications in health and healthcare. *International Journal of Health Geographics* 18 (1), 1–9.
- Kauppinen, T., Väättäinen, J., Hyvönen, E., 2008. Creating and using geospatial ontology time series in a semantic cultural heritage portal. In: *European Semantic Web Conference*. Springer, pp. 110–123.
- Kuhn, W., Hamzei, E., Tomko, M., Winter, S., Li, H., 2021. The semantics of place-related questions. *Journal of Spatial Information Science* (23), 157–168.
- Kuhn, W., Kauppinen, T., Janowicz, K., 2014. Linked data—a paradigm shift for geographic information science. In: *International Conference on Geographic Information Science*. Springer, pp. 173–186.

- Kulkarni, S., Jain, S., Hosseini, M. J., Baldrige, J., Ie, E., Zhang, L., 2021. Multi-level gazetteer-free geocoding. In: Proceedings of 2nd International Combined Workshop on Spatial Language Understanding and Grounded Communication for Robotics. Association for Computational Linguistics, pp. 79–88.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *nature* 521 (7553), 436–444.
- Li, H., Herfort, B., Lautenbach, S., Chen, J., Zipf, A., 2022. Improving OpenStreetMap missing building detection using few-shot transfer learning in sub-saharan africa. *Transactions in GIS* .
- Li, N., Raskin, R., Goodchild, M., Janowicz, K., 2012. An ontology-driven framework and web portal for spatial decision support 1. *Transactions in GIS* 16 (3), 313–329.
- Li, W., Hsu, C.-Y., Hu, M., 2021. Tobler’s first law in geoi: A spatially explicit deep learning model for terrain feature detection under weak supervision. *Annals of the American Association of Geographers* 111 (7), 1887–1905.
- Liang, C., Berant, J., Le, Q., Forbus, K., Lao, N., 2017. Neural symbolic machines: Learning semantic parsers on freebase with weak supervision. In: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 23–33.
- Liu, P., Biljecki, F., 2022. A review of spatially-explicit GeoAI applications in urban geography. *International Journal of Applied Earth Observation and Geoinformation* 112, 102936.
- Liu, Z., Janowicz, K., Cai, L., Zhu, R., Mai, G., Shi, M., 2022. Geoparsing: Solved or biased? an evaluation of geographic biases in geoparsing. *AGILE: GIScience Series* 3, 1–13.
- Mai, G., Cundy, C., Choi, K., Hu, Y., Lao, N., Ermon, S., 2022a. Towards a foundation model for geospatial artificial intelligence (vision paper). In: Proceedings of the 30th International Conference on Advances in Geographic Information Systems. pp. 1–4.
- Mai, G., Janowicz, K., Cai, L., Zhu, R., Regalia, B., Yan, B., Shi, M., Lao, N., 2020a. SE-KGE: A location-aware knowledge graph embedding model for geographic question answering and spatial semantic lifting. *Transactions in GIS* 24 (3), 623–655.
- Mai, G., Janowicz, K., Prasad, S., Shi, M., Cai, L., Zhu, R., Regalia, B., Lao, N., 2020b. Semantically-enriched search engine for geoportals: A case study with arcgis online. *AGILE: GIScience Series* 1, 1–17.
- Mai, G., Janowicz, K., Yan, B., Zhu, R., Cai, L., Lao, N., 2019. Contextual graph attention for answering logical queries over incomplete knowledge graphs. In: Proceedings of the 10th International Conference on Knowledge Capture. ACM.
- Mai, G., Janowicz, K., Yan, B., Zhu, R., Cai, L., Lao, N., 2020c. Multi-scale representation learning for spatial feature distributions using grid cells. In: Proceedings of the 8th International Conference on Learning Representations.
- Mai, G., Jiang, C., Sun, W., Zhu, R., Xuan, Y., Cai, L., Janowicz, K., Ermon, S., Lao, N., 2022b. Towards general-purpose representation learning of polygonal geometries. *GeoInformatica* , 1–52.

- Mai, G., et al., 2022c. A review of location encoding for GeoAI: methods and applications. *International Journal of Geographical Information Science* 36 (4), 639–673.
- Nickel, M., Murphy, K., Tresp, V., Gabrilovich, E., 2015. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE* 104 (1), 11–33.
- Nickel, M., Rosasco, L., Poggio, T. A., et al., 2016. Holographic embeddings of knowledge graphs. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. pp. 1955–1961.
- Nickel, M., Tresp, V., Kriegel, H.-P., 2012. Factorizing YAGO: scalable machine learning for linked data. In: *Proceedings of the Web Conference*. ACM, pp. 271–280.
- Noy, N., Gao, Y., Jain, A., Narayanan, A., Patterson, A., Taylor, J., 2019. Industry-scale knowledge graphs: Lessons and challenges: Five diverse technology companies show how it’s done. *Queue* 17 (2), 48–75.
- Nyamsuren, E., Xu, H., Scheider, S., Top, E. J., Steenbergen, N., 2021. Deconstruction of geo-analytical questions in terms of measures, supports, and spatio-temporal extents. *Journal of Spatial Information Science* .
- Park, S., Lee, Y., 2022. SSIKG: A framework for spatial-semantic integrated indoor knowledge graph construction. *Transactions in GIS* .
- Polson, N. G., Sokolov, V. O., 2017. Deep learning for short-term traffic flow prediction. *Transportation Research Part C: Emerging Technologies* 79, 1–17.
- Radford, A., et al., 2021. Learning transferable visual models from natural language supervision. In: *Proceedings of the 38th International Conference on Machine Learning*. pp. 8748–8763.
- Ramesh, A., et al., 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* .
- Ramos, R., Martins, B., 2022. Using neural encoder-decoder models with continuous outputs for remote sensing image captioning. *IEEE Access* 10, 24852–24863.
- Regalia, B., Janowicz, K., McKenzie, G., 2019. Computing and querying strict, approximate, and metrically refined topological relations in linked geographic data. *Transactions in GIS* 23 (3), 601–619.
- Santos, R., Murrieta-Flores, P., Calado, P., Martins, B., 2018. Toponym matching through deep neural networks. *International Journal of Geographical Information Science* 32 (2), 324–348.
- Scheider, S., Nyamsuren, E., Kruiger, H., Xu, H., 2021. Geo-analytical question-answering with GIS. *International Journal of Digital Earth* 14 (1), 1–14.
- Schlichtkrull, M., Kipf, T. N., Bloem, P., Van Den Berg, R., Titov, I., Welling, M., 2018. Modeling relational data with graph convolutional networks. In: *European Semantic Web Conference*. Springer, pp. 593–607.

- Shin, H., Basiri, A., 2022. Geographic biases in OSM contributions: How do the geographic extent of contributions differ among demographic groups? .
- Silva, J. D., Magalhães, J., Tuia, D., Martins, B., 2022. Remote sensing visual question answering with a self-attention multi-modal encoder. In: Proceedings of the 5th ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery. pp. 40–49.
- Singhal, A., 2012. Introducing the Knowledge Graph: things, not strings.
URL <https://googleblog.blogspot.com/2012/05/introducing-knowledge-graph-things-not.html>
- Sirko, W., Kashubin, S., Ritter, M., Annkah, A., Bouchareb, Y. S. E., Dauphin, Y., Keysers, D., Neumann, M., Cisse, M., Quinn, J., 2021. Continental-scale building detection from high resolution satellite imagery. arXiv preprint arXiv:2107.12283 .
- Sun, K., Hu, Y., Song, J., Zhu, Y., 2021. Aligning geographic entities from historical maps for building knowledge graphs. *International Journal of Geographical Information Science* 35 (10), 2078–2107.
- Tenzer, M., Rasheed, Z., Shafique, K., Vasconcelos, N., 2022. Meta-learning over time for destination prediction tasks. In: Proceedings of the 30th International Conference on Advances in Geographic Information Systems. pp. 1–10.
- Tiecke, T. G., Liu, X., Zhang, A., Gros, A., Li, N., Yetman, G., Kilic, T., Murray, S., Blankespoor, B., Prydz, E. B., et al., 2017. Mapping the world population one building at a time. arXiv preprint arXiv:1712.05839 .
- Trisedya, B. D., Qi, J., Zhang, R., 2019. Entity alignment between knowledge graphs using attribute embeddings. In: Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 33. pp. 297–304.
- van Eijck, J., Visser, A., 2012. Stanford encyclopedia of philosophy .
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. *Advances in Neural Information Processing Systems* 30.
- Wang, Z., Zhang, J., Feng, J., Chen, Z., 2014. Knowledge graph embedding by translating on hyperplanes. In: Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 14. pp. 1112–1119.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Chi, E., Le, Q., Zhou, D., 2022. Chain of thought prompting elicits reasoning in large language models. arXiv preprint arXiv:2201.11903 .
- Yan, B., Janowicz, K., Mai, G., Gao, S., 2017. From ITDL to Place2Vec: Reasoning about place type similarity and relatedness by learning embeddings from augmented spatial contexts. In: Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems. pp. 1–10.

- Yan, B., Janowicz, K., Mai, G., Zhu, R., 2019. A spatially explicit reinforcement learning model for geographic knowledge graph summarization. *Transactions in GIS* 23 (3), 620–640.
- Zheng, Z., Sieber, R., 2022. Putting humans back in the loop of machine learning in Canadian smart cities.
- Zhu, R., Janowicz, K., Cai, L., Mai, G., 2022a. Reasoning over higher-order qualitative spatial relations via spatially explicit neural networks. *International Journal of Geographical Information Science* , 1–32.
- Zhu, R., Janowicz, K., Mai, G., Cai, L., Shi, M., 2022b. COVID-Forecast-Graph: An open knowledge graph for consolidating COVID-19 forecasts and economic indicators via place and time. *AGILE: GIScience Series* 3, 1–12.