# Explainable Scientific Resource Recommendation through Aspect-based Relatedness over a Spatio-temporal Knowledge Graph

Anna Brauer[1]

[1]Chair of Geoinformatics, TUD Dresden University of Technology, 01069 Dresden, Germany

**Abstract**

Relatedness is inherently multi-faceted for datasets and other resources in the Earth sciences, spanning concepts such as thematic similarity and spatio-temporal overlap, which motivates explainable approaches to academic resource recommendation. Using a schema-rich knowledge graph of resource metadata as a case study, this work in progress explores the decomposition of relatedness into domain-specific aspects (topic, methodology, spatial coverage, temporal extent, provenance) to enable explainable recommendation across heterogeneous similarity measures of varying transparency. Rather than resolving the tension between explainability and functionality globally, aspect decomposition allows for a hybrid approach tailored to each aspect's characteristics and enables users to steer discovery by selecting aspects relevant to them. We propose a scoring strategy for each aspect and for their combination, and present early results suggesting that keyword annotations based on controlled vocabularies can shift content-based similarity rankings to emphasize a chosen aspect.

**Keywords**

explainable recommendation, knowledge graphs, resource discovery, spatio-temporal metadata

## 1. Introduction

Empowering researchers to discover relevant scientific resources, such as datasets, publications, or software, is essential to progress in a domain as multi-disciplinary as the Earth sciences. Generic data repositories and portals rarely meet domain-specific requirements, e.g., spatio-temporally enabled search functionalities, which motivated the establishment of domain-specific alternatives [1, 2]. Among these are the services developed by the consortium for Earth System Sciences (NFDI4Earth) of the German National Research Data Infrastructure: the community portal OneStop4All [3] and its knowledge base, a knowledge graph containing linked metadata of scientific resources.

Finding relevant resources within such portals is not always linear: researchers often possess only a high-level understanding of their problem and cannot describe the precise type of resource they need [4]. For this kind of explorative search, related resource recommendations could help researchers find what they did not know they were looking for. Relatedness is heavily context-dependent [5]: two scientific resources may be related through similar metadata properties (e.g., a common author) or close coupling (e.g., a dataset and a publication describing it) [6]. In the Earth sciences, spatial (e.g., data recorded in the same country) and temporal (e.g., reanalyses over overlapping periods) relatedness are also relevant.

The existence of these different types of relatedness suggests that designing a resource-to-resource recommender system upon the NFDI4Earth-developed services should include a strategy for explaining recommendations. This may not only enhance the user experience but may also help to build trust: recommendations perceived as random, biased, or unreliable may harm a data portal's acceptance in the scientific community and among data providers.

*XAI-KG@ISWC2026 - 2nd International Workshop on eXplainable AI and Knowledge Graphs, co-located with ISWC 2026, October 25-29, 2026, Bari, Italy*

anna.brauer@iki.fi (A. Brauer)

0000-0002-7092-1492 (A. Brauer)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

published 2026-11-07

CEUR-WS.org/Vol-4318/paper9.pdf

### 1.1. Problem statement

Finding an explainable approach to recommendation across multiple types of relatedness is difficult when relatedness is an opaque concept for some types, yet comparatively well-defined for others. For example, while there are several approaches to computing spatial or temporal similarity, existing measures such as the area of overlap are transparent and require no further explanation. At the other end of this spectrum are similarity measures over learned vector spaces, which are typically explained only post-hoc but are state of the art for identifying semantically related resources based on textual properties.

The knowledge-graph structure of the NFDI4Earth knowledge base also enables explanations supported by reasoning paths. However, such explanations are applicable only to types of relatedness that are (a) manifested as links in the graph and (b) sufficiently populated. The latter condition is not always met, given the heterogeneous quality of resource metadata.

These challenges suggest that a unified approach intended to capture any type of relatedness would force compromises on either explainability or recommendation quality. Instead, this work in progress explores an approach that aims to provide explainable recommendations by decomposing relatedness into domain-specific aspects. Under this framework, each aspect's relatedness scoring strategy may fall anywhere between symbolic and sub-symbolic, depending on how explicitly the aspect is modeled in the schema and how completely the corresponding metadata is populated. Per-aspect relatedness rankings would then be presented to the user alongside a combined ranking over the aspects they select.

Two guiding questions shape this work in progress: (1) What are relevant aspects of relatedness for scientific resources in the Earth sciences? (2) To what degree can inherently explainable approaches be utilized to model these aspects over a schema-rich yet unevenly populated spatio-temporal knowledge graph?

### 1.2. Contributions

This work-in-progress paper develops aspect decomposition as a means of hybrid resource-to-resource recommendation over unstructured and structured properties that establishes a baseline level of explainability, using the NFDI4Earth knowledge graph of Earth science resources as a case study.

For this, we introduce five aspects of relatedness for Earth-science resources: topic, methodology, spatial coverage, temporal extent, and provenance, derived from metadata commonly provided in the domain. We discuss strategies for computing resource relatedness, considering the opportunities and challenges each aspect poses, and outline how per-aspect scores could produce explainable recommendations. We also describe preliminary evaluation results and planned next steps to assess this approach's feasibility and potential to improve resource discovery.

## 2. Related work

Enabling scholars to discover relevant resources is a core concern of research data infrastructures [1, 2]. Much effort has been directed at improving query-based search, e.g., through LLM-based semantic search [7] and transformer-generated explanations of why specific resources are retrieved [8, 9]. In comparison, recommendations for related resources are an underexplored route to supporting resource discovery, although item-to-item recommendation is well established in recommender-systems research [10]. Academic recommender systems usually rely on content-based filtering and often provide global recommendations that are not tailored to user profiles [6]. As user profiles are an integral part of most recommender systems, research on explainable recommendation transfers only partly to profile-free academic settings, although it offers guidance. In particular, Zhang and Chen [11] give an overview of post-hoc and model-intrinsic approaches; the latter are structurally transparent and self-explaining, relying on, e.g., factorization, rule mining, or knowledge graphs.

Over knowledge graphs, paths between entities can explain their relatedness post hoc [12, 13], but they may also identify related objects in the first place. Relatedness measures over meta paths are

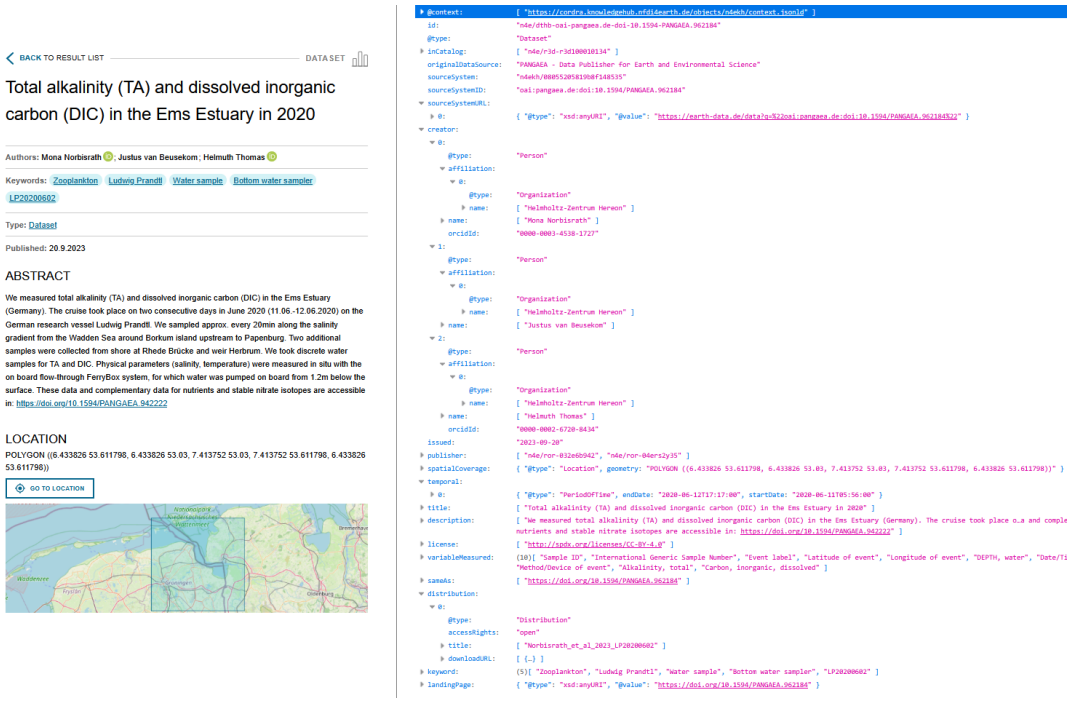



**Figure 1:** A sample dataset's detail page in the OneStop4All (left) and its metadata (right).

intrinsically explainable [14, 15], as meta paths are predefined sequences of relations between entity types, which may be learned or defined based on domain knowledge. Their rigid structure, however, makes the discovery of not explicitly stated patterns impossible [16]. Alternative approaches transform knowledge graphs into high-dimensional feature spaces, such as graph kernel-based approaches [17] and knowledge graph embedding (KGE) models [18, 19, 20]. These approaches recover implicit patterns, but generally do not distinguish the nature of the relatedness they capture.

A similar observation was made by Ostendorff et al. [21] for text transformers, leading them to propose aspect-based similarity scores for scientific literature. Later work fine-tuned pre-trained language models with respect to one or multiple domain-specific aspects, arguing that vector spaces governed by these aspects are more interpretable [22, 23]. Still, these approaches fold the aspects into a single trained model, while we explore keeping aspect scores separate. Furthermore, the textual nature of the aspects considered in these works sets them apart from our approach, as we additionally include aspects defined over structured properties, for which transparent, widely established similarity measures exist.

## 3. Data

The NFDI4Earth Knowledge Hub (KH) [24] is a linked data service that organizes metadata of scientific resources relevant for the Earth sciences in an RDF graph. It is searchable through a SPARQL endpoint or indirectly through the OneStop4All portal [3]. The property-rich ontology of the KH [25] aligns with metadata schemas of national data portals and aggregators that are relevant for the domain. Furthermore, it is built on widely adopted metadata standards, such as DCAT, Dublin Core, and ISO 19115. At the time of writing, the KH contains metadata on 968,967 data sets. Other resource types are currently available in smaller numbers, e.g., software (259), repositories (155), learning resources (532), and publications (190), but are expected to grow in the future. Metadata quality and level of detail vary, which poses a challenge for the identification of related resources. For example, authors often lack persistent identifiers (e.g., ORCIDs), so that relationships between resources created by the same author cannot always be traced within the KH.

# 4. Approach

We explore an inherently explainable approach to global resource-to-resource recommendation through a multi-aspect relatedness score whose components reflect aspects of relatedness relevant for geospatial resource discovery. The aspects we identified are anchored in the ontology of the KH, which was modeled on metadata commonly provided in the Earth sciences and thus reflects what the community considers essential for data discovery. We describe how each aspect is represented in the KH and how we plan to obtain its relatedness score, then discuss combining these scores into ranked recommendations.

## 4.1. Aspects of relatedness

**Content-related aspects: topic and methodology** Content-related metadata is rarely explicit and is instead carried in free-text fields such as *title*, *description*, and *keywords*. In preliminary tests, embeddings of these concatenated fields, produced with Sentence Transformer models [26], identified semantically related resources well. However, this method cannot characterize the *type* of relatedness; for example, it collapses the distinction between topically related resources (sharing the subject studied) and methodologically related ones (sharing the means of study).

We hypothesized that the type of relatedness could be recovered by adding aspect-specific keyword annotations to shift the embeddings toward a particular aspect. To verify this experimentally, we used the domain-specific controlled vocabulary GCMD [27], which includes keyword schemes related to *topic* (*Science Keywords*, *Measurement Name*) and keyword schemes related to *methodology* (*Platforms*, *Instruments*). We assigned semantically close keywords from each aspect to a random sample of dataset descriptions ($N = 5000$) using transformer-based classification [28]. Then, we used an architecturally different Sentence Transformer model to embed the dataset descriptions, as well as the annotated keywords and their GCMD definitions. Through a linear combination of the dataset and keyword embeddings, we created two embedding variants, each emphasizing one of the aspects.

Preliminary experiments comparing the top-$k$ ($k \in \{5, 10, 20\}$) most similar resources under both aspects showed that, for resources that received keyword annotations in both aspects, the rankings were meaningfully distinct. This result was observed consistently across all $k$ and combination weights. However, the vocabulary's coverage of a resource's content was decisive for the method's success. As no suitable methodology-related keyword was found for 72% of datasets, further work will expand the vocabularies used to create annotations and experiment with lower thresholds for annotation assignment. Moreover, the experiments highlighted that presenting the annotated keywords alongside the ranked results aids in explaining rankings that are otherwise difficult to attribute to either aspect due to the strong correlation between methodological and topical keywords (e.g., some instruments are predominantly used within one topical context).

**Spatio-temporal aspects: spatial and temporal extent** Spatial and temporal extent are usually provided as point geometry or minimum bounding rectangle and time interval, respectively. Traditional measures for spatial similarity include the Haversine distance, Hausdorff distance, area of overlap, and combinations thereof [29]; analogous measures exist for temporal overlap and distance. Recent work aims to embed geometries and time intervals into distance-, shape-, and direction-preserving vector spaces [30, 31, 29]. Still, for simple rectangles and intervals, the established traditional measures are exact and transparent. We therefore aim to utilize a combination of traditional measures.

**Provenance-related aspect** This aspect concerns a resource's relation to agents, such as publishers, authors, and their affiliations. Ideally, it should also capture lineage, e.g., if one resource has been used in producing another, but relevant properties are rarely explicit in the metadata and currently unpopulated in the KH. Provenance is the aspect where the knowledge graph's linked structure offers the clearest advantage, and its well-defined relationships make meta paths (e.g., Sun et al. [14]) feasible. However, as agent-affiliation links suffer from duplicates in the KH, we propose using KGEs of entities over a provenance-restricted subgraph of the KH to capture underlying patterns of relatedness and

partially recover missing links. Meta paths, combined with an intermediate node-unification layer, serve as an explainable fallback. KGE model selection and hyperparameter tuning will be carried out empirically using a training library such as PyKEEN [32], where TransE [33] is a promising candidate given its success in similar use cases [34, 35]. We will assess this through qualitative evaluation and a link-recovery test, removing known links before training.

### 4.2. Providing explainable recommendations

Decomposing relatedness into aspects shapes how we present recommendations to the user, which in turn imposes requirements on aspect-combination strategies. On a resource's detail page, the user will receive recommendations ranked by the aspects they select; thus, the user can switch individual aspects on or off. Each recommendation will also include a per-aspect breakdown that explains the ranking, e.g., through a radar chart.

This design requires that aspect scores are normalized to be comparable and globally meaningful rather than calibrated to the data. Adding an outlier, e.g., a Pleistocene dataset to a collection of contemporary ones, should not alter the relatedness scores among the other resources. Cosine similarity already fulfills these requirements; for the remaining aspects, domain-informed transformations can be used for normalization. The combined score is then the mean over all normalized aspect scores; here, a mean with disjunctive characteristics, such as the quadratic mean, may be preferable to the arithmetic mean, so that resources scoring exceptionally high on one aspect are not harshly penalized for low scores on others.

## 5. Conclusion and future work

In this work in progress, we have identified five aspects of relatedness for Earth science resources: topic, methodology, spatial coverage, temporal extent, and provenance, and presented an approach to providing explainable recommendations through aspect decomposition, which allows for the hybrid selection of an optimal scoring method per aspect. We have briefly outlined per-aspect evaluation strategies, yet selecting an appropriate methodology for evaluating the approach as a whole remains future work. Aspect decomposition may improve resource discovery by bringing resources to the forefront that are relevant under only one or a few aspects and may be missed entirely by general-purpose entity-similarity approaches; this assumption could be evaluated through a result-diversity analysis and a user study. The latter could also reveal whether users perceive steering discovery through aspect selection as an asset or as a burden due to its complexity. Moreover, emphasizing domain-relevant aspects may allow users to find true links between resources that are absent from the knowledge graph through recommendation (e.g., a publication and the dataset it uses), which could be evaluated using links not yet materialized in the knowledge graph.

## Acknowledgments

The work was funded by the German Research Foundation through the project NFDI4Earth (DFG project no. 460036893, https://www.nfdi4earth.de/) within the German National Research Data Infrastructure (NFDI, https://www.nfdi.de/).

## Declaration on Generative AI

During the preparation of this work, the author used Claude Opus 4.8 in order to: Improve writing style, grammar and spelling check. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the publication's content.

## References


[1] C. Henzen, A. Rümmler, H. Figgemeier, M. Wagner, L. Bernard, R. Müller-Pfefferkorn, Research data infrastructures in environmental sciences — challenges and implementation approaches focusing on the integration of software components, Transactions in GIS 28 (2024) 176–199. doi:10.1111/tgis.13131.

[2] S. Ankama, A. Degbelo, E. Naoseb, C. Henzen, L. Bernard, How well do current geoportals support geodata discovery? An empirical study, ISPRS International Journal of Geo-Information 15 (2026) 82. doi:10.3390/ijgi15020082.

[3] C. Henzen, A. Brauer, J. Grieb, R. Klammer, M. Konkol, R. Koppe, K. Liu, J. Munke, T. Niers, D. Nüst, T. Schäfer, A. Wellmann, Developing central support services for the German National Research Data Infrastructure in Earth System Sciences through a community-driven effort, in: EGU General Assembly Conference Abstracts, EGU25–5136, 2025. doi:10.5194/egusphere-egu25-5136.

[4] M. Hulsebos, W. Lin, S. Shankar, A. Parameswaran, It took longer than i was expecting: Why is dataset search still so hard?, in: Proceedings of the 2024 Workshop on Human-In-the-Loop Data Analytics, 2024, pp. 1–4.

[5] D. Bär, T. Zesch, I. Gurevych, A reflective view on text similarity, in: Proceedings of the International Conference Recent Advances in Natural Language Processing 2011, 2011, pp. 515–520.

[6] J. Beel, B. Gipp, S. Langer, C. Breitinger, Research-paper recommender systems: a literature survey, International Journal on Digital Libraries 17 (2016) 305–338. doi:10.1007/s00799-015-0156-0.

[7] N. W. Paton, J. Chen, Z. Wu, Dataset discovery and exploration: A survey, ACM Computing Surveys 56 (2023) 1–37. doi:10.1145/3626521.

[8] P. Yu, R. Rahimi, J. Allan, Towards explainable search results: A listwise explanation generator, in: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Association for Computing Machinery, New York, NY, USA, 2022, pp. 669–680. doi:10.1145/3477495.3532067.

[9] X. Shao, T. K. Wijaya, DS4RS: Community-driven and explainable dataset search engine for recommender system research, 2025. doi:10.48550/arXiv.2508.10238. arXiv:2508.10238.

[10] Z. Zhang, B. G. Patra, A. Yaseen, J. Zhu, R. Sabharwal, K. Roberts, T. Cao, H. Wu, Scholarly recommendation systems: A literature survey, Knowledge and Information Systems 65 (2023) 4433–4478. doi:10.1007/s10115-023-01901-x.

[11] Y. Zhang, X. Chen, Explainable recommendation: A survey and new perspectives, Foundations and Trends in Information Retrieval 14 (2020) 1–101. doi:10.1561/1500000066.

[12] Y. Xian, Z. Fu, S. Muthukrishnan, G. de Melo, Y. Zhang, Reinforcement knowledge graph reasoning for explainable recommendation, in: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, Association for Computing Machinery, New York, NY, USA, 2019, pp. 285–294. doi:10.1145/3331184.3331203.

[13] G. Balloccu, L. Boratto, G. Fenu, M. Marras, Reinforcement recommendation reasoning through knowledge graphs for explanation path quality, Knowledge-Based Systems 260 (2023) 110098. doi:10.1016/j.knosys.2022.110098.

[14] Y. Sun, J. Han, X. Yan, P. S. Yu, T. Wu, PathSim: Meta path-based top-k similarity search in heterogeneous information networks, in: Proceedings of the VLDB Endowment, volume 4, VLDB Endowment, 2011, pp. 992–1003. doi:10.14778/3402707.3402736.

[15] C. Shi, X. Kong, Y. Huang, P. S. Yu, B. Wu, HeteSim: A general framework for relevance measure in heterogeneous networks, IEEE Transactions on Knowledge and Data Engineering 26 (2014) 2479–2492. doi:10.1109/TKDE.2013.2297920.

[16] L. Luo, Y. Fang, M. Lu, X. Cao, X. Zhang, W. Zhang, GSim: A graph neural network based relevance measure for heterogeneous graphs, IEEE Transactions on Knowledge and Data Engineering 35 (2023) 12693–12707. doi:10.1109/TKDE.2023.3271425.

[17] N. M. Kriege, F. D. Johansson, C. Morris, A survey on graph kernels, Applied Network Science 5 (2020) 6. doi:10.1007/s41109-019-0195-3.

[18] J. Cao, J. Fang, Z. Meng, S. Liang, Knowledge graph embedding: A survey from the perspective of representation spaces, ACM Computing Surveys 56 (2024) 1–42.
[19] P. Ristoski, H. Paulheim, Rdf2vec: Rdf graph embeddings for data mining, in: International semantic web conference, Springer, 2016, pp. 498–514.
[20] N. Hubert, H. Paulheim, A. Brun, D. Monticolo, Do similar entities have similar embeddings?, in: The Semantic Web: 21st International Conference, ESWC 2024, Hersonissos, Crete, Greece, May 26–30, 2024, Proceedings, Part I, Springer-Verlag, Berlin, Heidelberg, 2024, pp. 3–21. doi:10.1007/978-3-031-60626-7_1.
[21] M. Ostendorff, T. Ruas, T. Blume, B. Gipp, G. Rehm, Aspect-based document similarity for research papers, in: Proceedings of the 28th International Conference on Computational Linguistics, 2020, pp. 6194–6206. doi:10.18653/v1/2020.coling-main.545.
[22] T. Schopf, E. Gerber, M. Ostendorff, F. Matthes, AspectCSE: Sentence embeddings for aspect-based semantic textual similarity using contrastive learning and structured knowledge, in: Proceedings of the 14th International Conference on Recent Advances in Natural Language Processing, INCOMA Ltd., Shoumen, Bulgaria, 2023, pp. 1054–1065. doi:10.26615/978-954-452-092-2_113.
[23] M. F. Brinner, S. Zarrieß, SemCSE-Multi: Multifaceted and decodable embeddings for aspect-specific and interpretable scientific domain mapping, in: Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics, volume 1, Association for Computational Linguistics, San Diego, USA, 2026, pp. 40565–40586. doi:10.18653/v1/2026.acl-long.1884.
[24] A. Degbelo, C. Henzen, J. Grieb, R. Klammer, C. Weiland, L. Bernard, R. Müller-Pfefferkorn, NFDI4Earth Knowledge Hub - Concept, NFDI4Earth Deliverable D4.3.2 (2023). doi:10.5281/zenodo.7950860.
[25] NFDI4Earth, NFDI4Earth Knowledge Hub schema (NFDI4Earth Ontology), https://nfdi4earth.pages.rwth-aachen.de/knowledgehub/nfdi4earth-kh-schema/, 2026. Accessed: 2026-07-22.
[26] N. Reimers, I. Gurevych, Sentence-BERT: Sentence embeddings using siamese BERT-networks, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, 2019. URL: https://arxiv.org/abs/1908.10084.
[27] Global Change Master Directory (GCMD), GCMD Keywords, Version 24.8, Greenbelt, MD: Earth Science Data and Information System, Earth Science Projects Division, Goddard Space Flight Center, NASA, 2026. URL: https://forum.earthdata.nasa.gov/app.php/tag/GCMD+Keywords, accessed via GCMD Keyword Forum Page.
[28] A. Hamm, Strategy comparison for semantic zero-shot taxonomy filters, in: 4th International Open Search Symposium (OSSYM 2022), 2022. URL: https://elib.dlr.de/190966/.
[29] Y. Yuan, F. Liu, G. Yang, M. Wang, A review of representation and similarity measurement methods for geospatial scenes, International Journal of Geographical Information Science 40 (2026) 1519–1546.
[30] G. Mai, K. Janowicz, Y. Hu, S. Gao, B. Yan, R. Zhu, L. Cai, N. Lao, A review of location encoding for GeoAI: Methods and applications, International Journal of Geographical Information Science 36 (2022) 639–673. doi:10.1080/13658816.2021.2004602.
[31] G. Mai, C. Jiang, W. Sun, R. Zhu, Y. Xuan, L. Cai, K. Janowicz, S. Ermon, N. Lao, Towards general-purpose representation learning of polygonal geometries, GeoInformatica 27 (2023) 289–340. doi:10.1007/s10707-022-00481-2.
[32] M. Ali, M. Berrendorf, C. T. Hoyt, L. Vermue, S. Sharifzadeh, V. Tresp, J. Lehmann, PyKEEN 1.0: A Python Library for Training and Evaluating Knowledge Graph Embeddings, Journal of Machine Learning Research 22 (2021) 1–6. URL: http://jmlr.org/papers/v22/20-825.html.
[33] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, O. Yakhnenko, Translating embeddings for modeling multi-relational data, Advances in neural information processing systems 26 (2013).
[34] G. Mai, K. Janowicz, B. Yan, Combining text embedding and knowledge graph embedding techniques for academic search engines., in: Semdeep/NLIWoD@ ISWC, 2018, pp. 77–88.
[35] M. Nayyeri, S. Vahdati, X. Zhou, H. Shariat Yazdi, J. Lehmann, Embedding-based recommendations on scholarly knowledge graphs, in: European semantic web conference, Springer, 2020, pp. 255–270.