The Helmholtz Knowledge Graph: driving the Transition towards a FAIR Data Ecosystem in the Helmholtz Association

Jens Bröder[1][0000−0001−7939−226X], Gabriel Preuß[2][0000−0002−3968−2446], Fiona D'Mello[1][0000−0002−0465−1009], Said Fathalla[1][0000−0002−2818−5890], Volker Hofmann[1][0000−0002−5149−603X], and Stefan Sandfeld[1][0000−0001−9560−4728]

[1] Forschungszentrum Jülich GmbH, https://ror.org/02nv7yv05, Institute of Advanced Simulation - Materials Data Science and Informatics (IAS-9), Germany
{j.broeder, f.dmello, s.fathalla, v.hofmann, s.sandfeld}@fz-juelich.de
[2] Helmholtz-Zentrum Berlin für Materialien und Energie, Germany
gabriel.preuss@helmholtz-berlin.de

Abstract. The Helmholtz Knowledge Graph aggregates metadata about digital assets and research output from the various institutional and siloed digital infrastructures within the Helmholtz association. It is part of the technical backbone of the Helmholtz FAIR data space, that is established by the "Helmholtz Metadata Collaboration" (HMC). It is used to drive change towards better metadata practices, increase visibility of data and provide useful data-based services. In this paper, we present how metadata describing Helmholtz's digital assets and research outputs are harvested and uplifted. The data is made publicly accessible to both humans and machines through a user interface based text search and a SPARQL endpoint, respectively.

Keywords: Knowledge Graph · Linked Data · FAIR · Helmholtz Metadata Collaboration · Schema.org · JSON-LD · Data Mining · Metadata

1 Introduction

Research in the Helmholtz Association is carried out in inter- and multidisciplinary collaborations that span between its 18 independently operating non-university research centers across Germany. Helmholtz digital infrastructure is institutional, and thus Helmholtz's research data and other digital assets are stored and maintained in independent silos, lack visibility and accessibility with their full value remaining unavailable to scientists, managers, strategists, and policy makers. Metadata on the web is typically used to track citations not data. It often lacks completeness and semantic quality and therefore, published research data often fails to satisfy FAIR principles [7] resulting in a lack of interoperability and re-usability. The "Helmholtz Metadata Collaboration" (HMC)[3] is taking on this challenge by developing innovative technologies and tools for

[3] https://helmholtz-metadaten.de/en

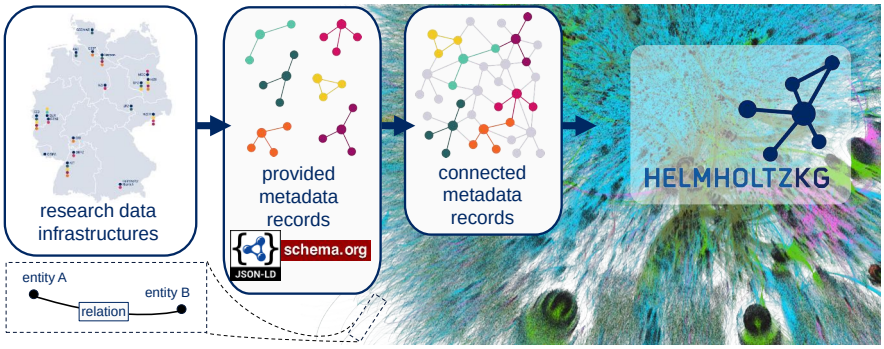

Fig. 1. Aggregation of the Helmholtz Knowledge Graph:
Data records (schema.org / JSON-LD) are continuously harvested from Helmholtz data providers and run through our data pipelines for initial uplifting, de-duplication and integration into the Helmholtz KG. The graph on the right displays the sub graph of the 2.26 mil. main entities and their connections as a semi relaxed force atlas layout using Gephi[1].

a sustainable handling of research data through high-quality metadata. HMC launched the "unified Helmholtz Information and Data Exchange (unHIDE)"[4] - an initiative to network and harmonize Helmholtz digital infrastructure, and connect Helmholtz data through a lightweight interoperability layer in form of the Helmholtz Knowledge Graph. With this, we envision to (1) provide a better cross-organizational access to Helmholtz's (meta)data and information assets on an upper semantic level, (2) harmonize and optimize the related metadata across the association, and (3) form a basis from where semantic quality and depths of metadata descriptions can be improved and extended into domain and application levels. The institutional focus and defined domain boundaries within the Helmholtz's research fields[5] differentiate the Helmholtz KG from other graphs with wider scopes, such as e.g. the OpenAIRE graph [6]. These graph partners we approach for graph-graph exchange of data and developed technologies.

2 Data Aggregation and Statistics

To aggregate the Helmholtz KG, metadata records are harvested from Helmholtz data providers and integrated (Fig. 1): we developed a library of harvesters [2] that harvests records recurrently through common web standards, such as OAI-PMH, sitemaps, feeds, or from the APIs of established data providers (DataCite, GitHub, GitLab). The data in the Helmholtz KG is aligned along https://schema.org semantics, for which exposed metadata records are preferably provided as Linked Data (e.g. serialized in JSON-LD documents). All har-

[4] https://helmholtz-metadaten.de/en/unhide_helmholtz-kg
[5] https://www.helmholtz.de/en/research/research-fields/

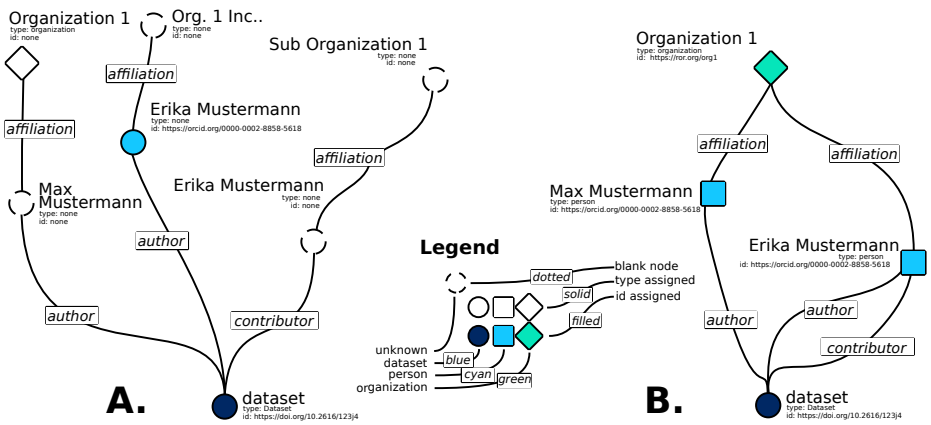

Fig. 2. Uplifting data records:
A. Often, metadata is not systematically typed or assigned with persistent identifiers (PIDs). The resultant connections show that the same entities appear as duplicated blank nodes in several instances. **B.** The same data with assigned types and PIDs allows resolution of entities leading to a higher connectivity of the data in the graph.

vested records are processed through a data pipeline utilizing the workflow manager Prefect (v2.15.0)[6], during which records are initially uplifted with inferable semantic annotations and then de-duplicated. Then, records are uploaded into an OpenLink Virtuoso triple store, and indexed into an Apache Solr database to support full federated text search. The graph is exposed as a set of triples through a SPARQL endpoint[7]. In addition, users can search the graph data through a user-friendly web front end[8] as well as an API[9]. The aggregation and deployment design was inspired by the Ocean InfoHub (OIH) project [4] whose open code base kick-started our development. The graph is deployed first on the HDF Cloud [5] and now on the JSC Cloud at the Jülich Supercomputing Centre. The first release of the Helmholtz KG contains 2.15 million metadata records, which were harvested from 32 Helmholtz data providers. At the graph level, this results in 72 million RDF triples. Of these, 16.35 million entities are associated with a semantic type (approximately seven typed entities per record). Currently, the most frequent types of entities are persons, organizations, documents, datasets, software and events. Of these, 793k entities are associated with a persistent identifier (PID) or URL. PIDs as well as correct semantic annotation (e.g. of entity types) are important to increase the connectivity in the graph by resolving entities as shown in Fig.2. Data with poor semantic quality often lacks PIDs (Fig.2A) resulting in duplicated instances of the same nodes within the same record. This impairs search queries up to a level where dupli-

[6] https://www.prefect.io/
[7] SPARQL endpoint: https://sparql.unhide.helmholtz-metadaten.de
[8] Web front end: https://search.unhide.helmholtz-metadaten.de
[9] Web API: https://api.unhide.helmholtz-metadaten.de

cated instances might not be recognized and returned as separate results for a given query. By assigning PIDs (Fig.2B) entities can be resolved leading to an increased number of connections to a single node. PID usage with data varies with entity types: ORCID and ROR PIDs are used to refer to persons respectively organizations exclusively, in contrast DOI PIDs are used heterogeneously to refer to research outputs including data and scholarly communications.

3 Outlook

In the future, we plan to continuously grow the graph by connecting more infrastructures as data providers from within Helmholtz. We further look to integrate data from Helmholtz FAIR digital objects. Through consulting and assisting data providers to expose high-quality metadata on the web, we will (1) increase their search engine optimization and (2) harmonize how top-level metadata is used in our association. Further, we will use the graph data to uplift and semantically enrich the provided data records. This will be achieved by type inference and entity resolution through logic and the application of machine learning methods. This uplifted data will be contrasted with the original data and provided back to the authoritative source of the metadata. We aim to keep the graph semantics and technology interoperable with other scientific knowledge graphs - such as the semantic pedigree ODIS [4] - to allow graph-to-graph interaction of data and federated querying.

Data and Software Availability All software related to the Helmholtz Knowledge Graph is open source and freely available [2] under MIT license. The metadata and graph data can be fully extracted via API and the SPARQL endpoints and is available under the Creative Commons Attribution 4.0 International license. Versioned dumps of the graph data are pushed to Zenodo [3].

Acknowledgments. This project was funded by the Helmholtz Metadata Collaboration (HMC), an incubator-platform of the Helmholtz Association within the framework of the Information and Data Science strategic initiative. The authors acknowledge the Helmholtz Data Federation (HDF) for providing services and compute on the HDF Cloud at the Jülich Supercomputing Centre (JSC). We thank our colleagues Pier Luigi Buttigieg, Thomas Jejkal, Oonagh Mannix, Anton Pirogov, Silke Gerlich and Mustafa Soylu for contributions to the unHIDE initiative, as well as the ODIS/ OceanInfoHub team for advice and sharing their architecture code.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bastian, M., Heymann, S., Jacomy, M.: Gephi: an open source software for exploring and manipulating networks. In: Proceedings of the international AAAI conference on web and social media. vol. 3, pp. 361–362 (2009)

2. Bröder, J., D'Mello, F., Preuß, G., Mannix, O., Buttigieg, P.L., Fathalla, S., Hofmann, V., Sandfeld, S.: Software repositories of unhide and the knowledge graph, https://codebase.helmholtz.cloud/hmc/hmc-public/unhide
3. Bröder, J., Preuß, G., D'Mello, F., Fathalla, S., Hofmann, V., Sandfeld, S.: The helmholtz knowledge graph dataset (Apr 2024), https://doi.org/10.5281/zenodo.10948205
4. Fils, D., Scott, L., Buttigieg, P.L., Spears, T., Lambert, A., Provoost, P., Pissierssens, P.: Ocean infohub: A global knowledge network for the ocean data and information system (odis). In: AGU Fall Meeting Abstracts. vol. 2021, pp. IN45H–0523 (2021)
5. Hagemeier, B.: Hdf cloud–helmholtz data federation cloud resources at the jülich supercomputing centre. Journal of large-scale research facilities JLSRF **5**, A137–A137 (2019), https://doi.org/10.17815/jlsrf-5-173
6. Manghi, P., Atzori, C., Bardi, A., Baglioni, M., Dimitropoulos, H., La Bruzzo, S., Foufoulas, I., Mannocci, A., Horst, M., Iatropoulou, K., Kokogiannaki, A., De Bonis, M., Artini, M., Lempesis, A., Ioannidis, A., Manola, N., Principe, P., Vergoulis, T., Chatzopoulos, S., Pierrakos, D.: Openaire graph dataset (Jan 2024), https://doi.org/10.5281/zenodo.10488385
7. Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.W., da Silva Santos, L.B., Bourne, P.E., et al.: The fair guiding principles for scientific data management and stewardship. Scientific data **3**(1), 1–9 (2016), https://doi.org/10.1038/sdata.2016.18