How to Create a Portal for Digital Humanities Research Using a Linked Open Data Cloud of Cultural Heritage Knowledge Graphs: Case SampoSampo

Eero Hyvönen[1,2,*], Petri Leskinen[1], Annastiina Ahola[1], Heikki Rantala[1] and Jouni Tuominen[3,2,1]

[1]*Aalto University, Department of Computer Science, Semantic Computing Research Group (SeCo), https://seco.cs.aalto.fi*
[2]*University of Helsinki, Helsinki Centre for Digital Humanities (HELDIG)*
[3]*University of Helsinki, Helsinki Institute for Humanities and Social Sciences (HSSH)*

Abstract

This paper presents a novel approach and first results of creating a global data service and portal, SampoSampo, based on a data alignment service on top of a cloud of interlinked Cultural Heritage knowledge graphs and data services of different application domains. In this way, a more comprehensive global view for searching, exploring, and analyzing entities with enriched linked data and their semantic connections can be provided than by using local KGs separately. The portal can be used for searching and exploring a cloud of linked KGs with a single user interface (UI) and for finding semantic "interesting" connections (relations) between their entities with natural language explanations.

Keywords

linked data, digital humanities, entity alignment, semantic portal, data analysis, knowledge discovery

1. Connecting Everything to Everything Else

Leonardo da Vinci (1452–1519) has said: "*Learn how to see. Realize that everything connects to everything else*" [1]. This wisdom is very true regarding Cultural Heritage (CH) data due to its rich linkedness—and regarding Linked Open Data (LOD) where enriching data by linking is the great underlying promise, as promoted by the 5th star in Tim Berners-Lee 5-star model[1]. However, exposing and learning the connections (links) between resources in CH LOD for researchers and the public to study and learn is a challenge [2]: 1) The data are typically available in distributed data silos. 2) The data is heterogeneous based on different data models. 3) Different identifiers for entities are used for the same entities in different data silos, which cuts off links and connections.

This paper addresses the problem of finding connections between entities in distributed heterogeneous cloud of knowledge graphs (KG) to enrich data for Digital Humanities (DH) research. More specifically, the research problems addressed are as follows:

1. How to search and link data about entities from different KGs in a LOD cloud?
   - For enriched descriptions about entities, such as persons, organizations, and places.
   - To find possibly conflicting information about the entities in the data sources.
2. How to search for relations between entities within and across KGs?
   - For finding out, e.g., how people are related to places or each other.
   - For the discovery of 'interesting' or even 'serendipitous' connections [3].

SemDH 2025: Second International Workshop of Semantic Digital Humanities. Co-located with ESWC 2025, June 02, 2025, Portoroz, Slovenia.
*Corresponding author.
eero.hyvonen@aalto.fi (E. Hyvönen)
https://seco.cs.aalto.fi/u/eahyvone (E. Hyvönen)
0000-0003-1695-84 (E. Hyvönen)

[1]5-star model: https://5-star.info

3. How to create a LOD service needed in applications addressing the questions 1–2 above?
4. How to create an application on top of the data service 3 above?

As a case study, the Sampo series[2] of mutually related CH LOD services and semantic portals based on the so-called "Sampo model" [4] are considered. Our research hypothesis is to re-use the Sampo model itself to establish an entity alignment LOD service on top of which a meta-level Sampo, "SampoSampo", is created, based on different Sampo systems and related datasets and KGs. For searching entities globally, faceted semantic search on aligned entities is used, and for searching and discovering new connections between entities, faceted search is applied to knowledge-based relational search on explainable connections between entities, a step toward the "Web of Wisdom" [3]. First results of this work are presented and discussed based on a new demonstrator SampoSampo on the Web. The data, LOD service, and portal code will be published using CC BY 4.0 and MIT licenses by the end of 2025.

The paper starts with presenting related works (Section 2) that have inspired our research. After this, use cases of SampoSampo are motivated and explained. Our solution approach is then presented (Section 4) followed by Section 5 on its demonstrational implementation as a portal on top of a SPARQL endpoint. Using the portal is illustrated with examples (Section 6) before conclusions in Section 7.

2. Related Work

Aligning heterogeneus KGs involves 1) mapping (aligning) data models used onto each other and 2) aligning entities and concepts used for populating the datamodel, i.e., resources with their local identifiers (IRIs) across datasets. For mapping datamodels one can use, e.g., Dublin Core[3] and its dumb down principle, or alternatively transform KGs into a harmonizing underlying ontologial model, such as CIDOC CRM [5] based on modeling the real world using events. This paper focuses on the problem of aligning resources often referred to as ontology mapping [6]. This problem has been addressed, e.g., in libraries where different identifiers are often used for, e.g., authors and places and literary works. The Virtual International Authority File service VIAF[4] has been created to mitigate the problem [7] and is a source of inspiration in our own work. The Linked Open Data Cloud[5] is a prominent example of aligning resources in different local KGs on a global level using `owl:sameAs` and other properties. Various ontology services, such as BioPortal [8] and the KOKO ontology cloud served by ONKI/Finto ontology services [9] provide alignments between different Knowledge Organizations (KOS) systems that are used for indexing data. Linked Open Vocabularies (LOV) is an example of a catalog of reusable vocabularies, their alignment, and version histories [10]. The Europeana Data Model [11] presents an approach for aligning entities of local datasets using a proxy model.

3. Use Cases of SampoSampo

The focus of our work is on entity resources for historical persons, organizations, and places that are widely used in virtually all Sampo systems and related CH datasets. For example, Figure 1 illustrates the linkedness of some biographical Sampo systems and other systems, including the Kanto authority file system as a KG[6] by the National Library of Finland, Wikidata[7], and the German Integrated Authority File system of the Deutsche National Bibliothek (GND)[8]. The numbers on the connection arcs tell the number of shared person resources between the connected datasets.

The main use cases for creating SampoSampo are (cf. research questions 1 and 2 in Section 1):

[2]Sampo series of over 20 LOD services and CH portals: https://seco.cs.aalto.fi/applications/sampo/
[3]Dublin Core Metadata Initiative: https://dublincore.org
[4]Virtual International Authority File system: https://viaf.org
[5]Linked Open Data Cloud: https://lod-cloud.net/
[6]Kanto authorities: https://finto.fi/finaf/en/
[7]Wikidata: https://wikidata.org
[8]Gemainsame Normdatei: https://www.dnb.de/EN/Professionell/Standardisierung/GND/gnd_node.html

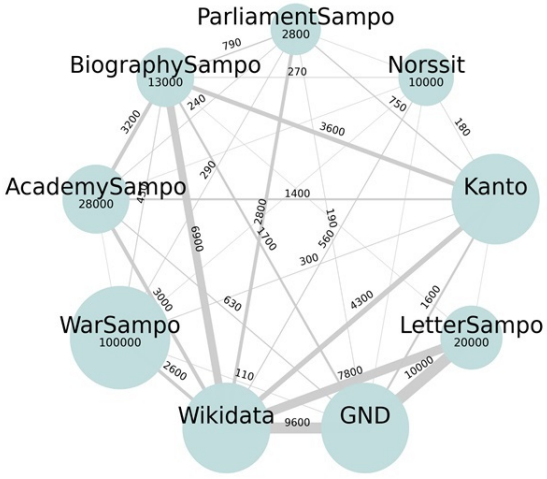

Figure 1: Interlinked actor resources in some biographical Sampos and beyond [12]

1. Firstly, it is often useful to search for data about entities from different data sources, as different datasets may contain complementary enriching or conflicting information about the entities. For example, BiographySampo [13] publishes general biographical data about 13 000 prominent Finns but there is more data available on them in other Sampos as shown in Figure 1: AcademySampo contains more details about the academic careers for 3200 persons in BiographySampo, ParliamentSampo about the speeches and activities in Parliament of Finland for 790 politicians, WarSampo supplements data for 290 prominent Finns as WW2 soldiers that have an entry in BiographySampo, and Wikidata contains entities for 6900 people in BiographySampo.
2. Secondly, searching for explainable explicit or implicit connections (relations) using relational search can be interesting within a single data alone, as demonstrated by the relational search application perspective for BiographySampo [14]. However, it is also possible to combine connection search results from different datasets, based on different data models and criteria, and when searching for new link paths across dataset boundaries [15].

SampoSampo aims to make data enrichment and semantic connection search available for the end users as a portal based on an entity alignment LOD service aggregated from a set of underlying KGs. The goal is to provide the user with enriched global descriptions about entities, such as persons, organizations, and places. In addition to that, comparing data imported from multiple databases also facilitates detecting contradictions and errors in the data sources [16], which is valuable, e.g., in biography research and prosopography [17]. Furthermore, search for explained "interesting" relations between entities within and across datasets can be provided. For example: how are people related to places, say Finnish artist to France or Jean Sibelius, the composer, to Berlin?

The entity resources used in Sampo systems are often based on the same infrastructure resources available in the Finnish ontology services ONKI[9] and Finto[10]. However, due to various reasons also application-specific KOS have been used for populating the metadata models, and the data is linked to international datasets, too. The LOD service of SampoSampo creates a kind of universal reference service aa a SPARQL endpoint for using the resources of Sampos and related datasets. This kind of LOD service is useful when creating and aligning new datasets with existing ones by FAIR principles[11]. For example, the person data of SampoSampo has been used as a biographical dictionary in the new LetterSampo Finland (1809-1917) system[12] containing data about 1 200 000 letters sent in the Grand Duchy of Finland.

[9]ONKI service: https://onki.fi
[10]Finto service: https://finto.fi
[11]FAIR principles: https://www.go-fair.org/
[12]LetterSampo Finland (1809–1917): https://seco.cs.aalto.fi/projects/coco/

4. Solution Approach

4.1. Using Sampo Model and Sampo-UI Framework

Our approach to address the research questions above is to use the Sampo model [4] and the Sampo-UI framework [18, 19]. In this model, data are aggregated from multiple data sources, aligned, and published in an open SPARQL endpoint. The endpoint can be used directly for data analyses using tools, such as the Yasgui editor [20] or Jupyter notebooks[13], or by applications created on top of the data service, such as the Sampo portals. The data service is completely separated from the applications; only the external SPARQL endpoint is used to access the data in the portal.

The Sampo-UI framework implements the principles of the Sampo model regarding end-user interface design. Sampo-UI aims at "standardizing" UI development by providing a framework where faceted search and browsing are seamlessly integrated with data visualization and analysis tools. This makes it easy for end users to use the UI and for application developers to implement it.

We followed the steps below for creating the SampoSampo system:

1. Create a KG for aligning entities in different Sampos and related data sources by using a proxy model adapted from the Europeana Linked Data model [11].
2. Create KGs for inferred explained relations between entities.
3. Publish KGs as a LOD service (SPARQL endpoint) at the Linked Data Finland platform[14] [21].
4. Create a Sampo portal using the Sampo-UI framework on top of the LOD service with faceted search perspectives for Persons, Places, Organizations, and (historical) Events.
5. Create instances pages for entities that contain aggregated information from different KGs with provenance information about the sources, highlighting possible issues of conflicting information.
6. Create relational search perspectives to discover new implicit semantic connections between entities with natural language explanations, using knowledge-based relational search [22, 3].

4.2. Data Model Based on Proxies

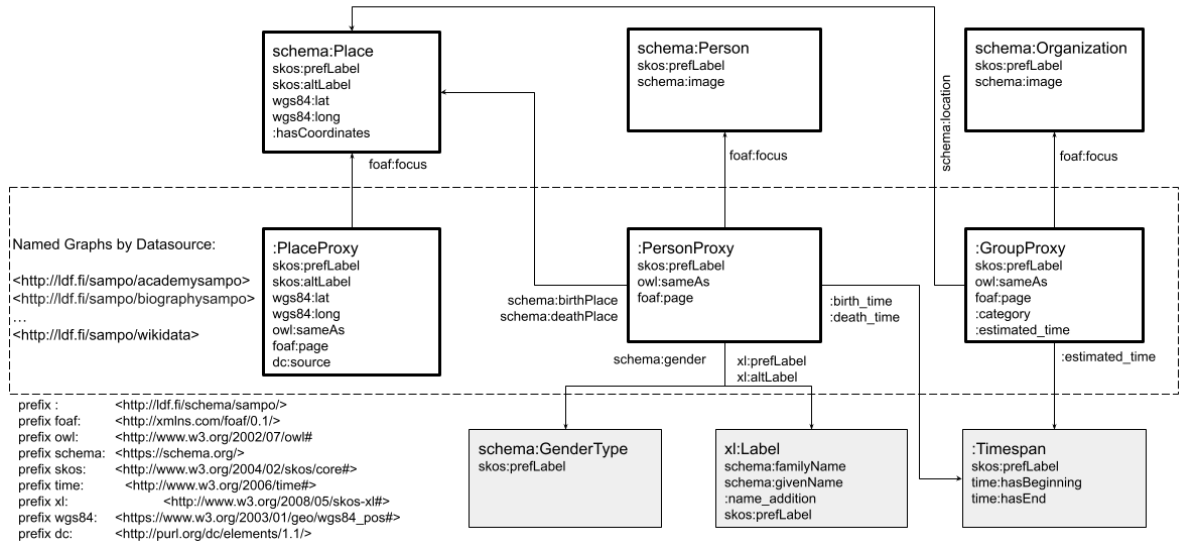

Figure 2: SampoSampo data model based on proxies

The data model underlying SampoSampo is depicted in Figure 2. The focus of the model is on shared core classes that are used widely in the underlying KGs of the cloud: Persons, Organizations, and Places.

[13] Jypyter notebooks: https://jupyter.org/
[14] Linked Data Finland platform: https://ldf.fi

The properties harvested are the key properties by which the entities can be identified and searched for on a global level in the portal using faceted search of the Sampo model.

4.3. Aligning Entities Using the Proxy Model

The person data collected from the data sources was first aggregated using SPARQL queries, API services, and web scraping, when data was not available in structured form. For the people data, we chose only those deceased individuals of whom the full name and the times of birth and death were known. Results from Wikidata were limited to those having some connection to Finland, e.g., by nationality, place of living, or having an entry in Finnish databases or Finnish relatives. In addition to this basic information, also the name variations and their types (preferable name, birth name, married name, pseudonym), gender, and the places of birth and death were included in the data.

The person data was disambiguated and linked first based on existing links between the datasets, and secondly by finding the matching candidates based on the string distances between their names as well as on the matching e.g. nearby years of living. The original total count of person entries from the data sources was 387 000 which was first reduced to 323 000 people by focusing on those with existing database linkage and finally, after a deduplication process, to 257 000 people. For the data publication of the SampoSampo portal, the person data was further limited by filtering out those who do not have a biographical description text. For example, relatives of which only the name and years of lifetime were know were excluded. This process revealed many errors in data, duplicate entries, and inconsistencies between the source datasets, e.g., contradicting gender or years of living, and misspelled name variations.

A similar process was applied in the case of organizations. However, in this case, different language variations of organization names were used. Some data sources did not provide any information about the type of organization (company, school, etc.) or its time of activity, so when possible, this information was approximated using the known memberships of related people. The deduplication process of organizations required more processing than in the case of people, like extracting the place names from the full names and adding a list of founders or members of the organization. For the deduplication the Python module Dedupe[15] was used by providing a small amount of human training data. The original total count of organization entries from the data sources was 62 000 which was reduced first by using the existing database linkage to 54 000 entries, and finally after deduplication to 53 000.

For places, the data was limited to those places that had at least one reference to the place in the original data (e.g., had an event taking place there). In the case of one particular dataset, BookSampo, an additional restriction was added to get only the places with coordinates to ensure that fictitious places present among the data were filtered out. For places with links to Wikidata and the places of the Genenral Finnish Ontology YSO[16], information was also collected from those sources, but was limited only to those particular place entities. The place disambiguation process followed a process similar to that of the person-disambiguation process above by first disambiguating based on links between datasets and then finding potential matches based on string distances of the place names from different sources. Potential matches were then either accepted or rejected based on whether the difference in coordinates fell within a set threshold. During this process 44 000 place proxies were grouped under 23 000 place provided instances.

5. Implementation of SampoSampo Data Service and Portal

The SampoSampo data service hosts the LOD for the SampoSampo portal. The data for different entity types was aggregated from various Sampo systems as well as external sources. At the moment, SampoSampo aggregates data about 80 566 people entities from the following Sampo systems and other data sources (with the number of instances aligned given in brackets): Wikidata [36027], KANTO

[15] https://docs.dedupe.io/en/latest/
[16] YSO Places: https://finto.fi/yso-paikat/en/

[31406], ISNI [30788], Wikipedia [27985], Geneanet [23269], AcademySampo [23094], BiographySampo [21540], VIAF [17824], BookSampo [9097] [], Geni [5985] genoalogy web service[17], Wikitree [5720], WarSampo [4357] [], Edelfelt [3319] letters in LetterSampo Finland, ParliamentSampo [2101], ULAN [1718] KG provided by Getty Research, Snellman [1342] letters in LetterSampo Finland, ArtSampo [957], Åbo Academy [801] letters in LetterSampo Finland, Norssit Sampo [578], OperaSampo [575], and HISTO [521] Finnish history ontology. The same datasets were used for extracting and aligning 52 161 organizations, 20 676 place entities, and 1198 major historical events of the HISTO ontology. Based on classes of people, groups/organizations, places, and events, application perspective were created for searching the class instances based on their shared core properties. For example, the people perspective has facets for names, birth and death dates, birth and death places, and the data sources.

As for searching connections, data from our earlier relational search systems [23, 23, 24] can be aggregated, based on different methods and approaches for creating connections and explanations for them. Here faceted search is applied to filter out instances of "interesting" or even "serendipitous" connections between people, organizations, places, and events, that have been pre-compiled from the data, The properties of the connection instances include the endpoints of the connection and an explanation of the connection that has been created when searching for the connections. In the portal, these connections can then be searched and analyzed using faceted search as customary in the Sampo model. The demonstrator includes separate perspectives for searching connections: one based on the BiographySampo relational search for finding relations between people and places, and one based on links between entities in the Finnish Wikipedia, where explanations are based on the textual context in which the links occur.

The first demonstrator of the portal was created using the Sampo-UI framework [18, 19]. It offers software developers a starting base to build a JavaScript web application, which can be customized with minimal effort to create LD applications. The framework consists of two main components: (1) a client-side interface built using the well-established React[18] and Redux[19] libraries and (2) a Node.js[20] back-end developed with the Express[21] framework. The portal will be published by the end of 2025.

6. Using SampoSampo Portal

The landing page of SampoSampo provides access to application perspectives where the instances of KG classes can be searched using semantic faceted search where the facets correspond to the properties of the class. After filtering results by making selections on the facets, the result set can displayed as a table or using a variety of data-analytic tools and visualizations, such as charts, maps, and timelines. By selecting an instance from the result set, aggregated linked data related to it can be displayed and data-analyses and visualizations pertaining to the entity instance be shown.

The initial demo implementation contains perspectives for searching People, Organizations, Places, and Historical events, Connections between people and places, and Connections between entities based on links available in the Finnish Wikipedia and Wikidata. Figure 3 depicts the People perspective with 7 facets on the left and search results on four alternative tabs on the right. In this case OperaSampo was selected from that data source facet and a result set of 575 people there that are linked with other data sources are shown on the right. The user has selected from the tabs the MIGRATIONS tab that shows the life lines of the people on a map. One can learn from the visualization, for example, that quite a few opera people have moved to Paris and died there.

Figure 4 illustrates how a perspective for discovering implicit semantic connections is used. The perspective is based on the class of connections whose instances represent connections between people and places with properties for them, as well as an explanation about the connection. The facets on the left are: Person, Occupation or title of the person, Place, and Type of the connection, based on the

[17] https://geni.com
[18] https://reactjs.org
[19] https://redux.js.org
[20] https://nodejs.org/en
[21] https://expressjs.com

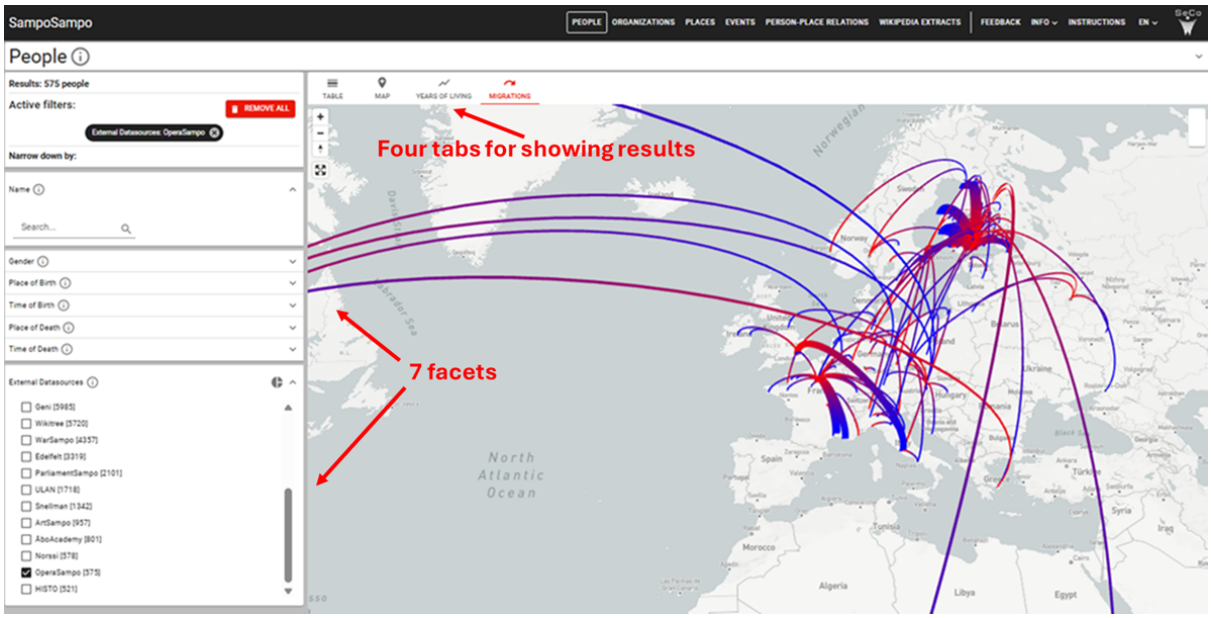

Figure 3: Visualizing the life lines of people in OperaSampo from their place of birth (blue end of the line) to the place of death (red end).

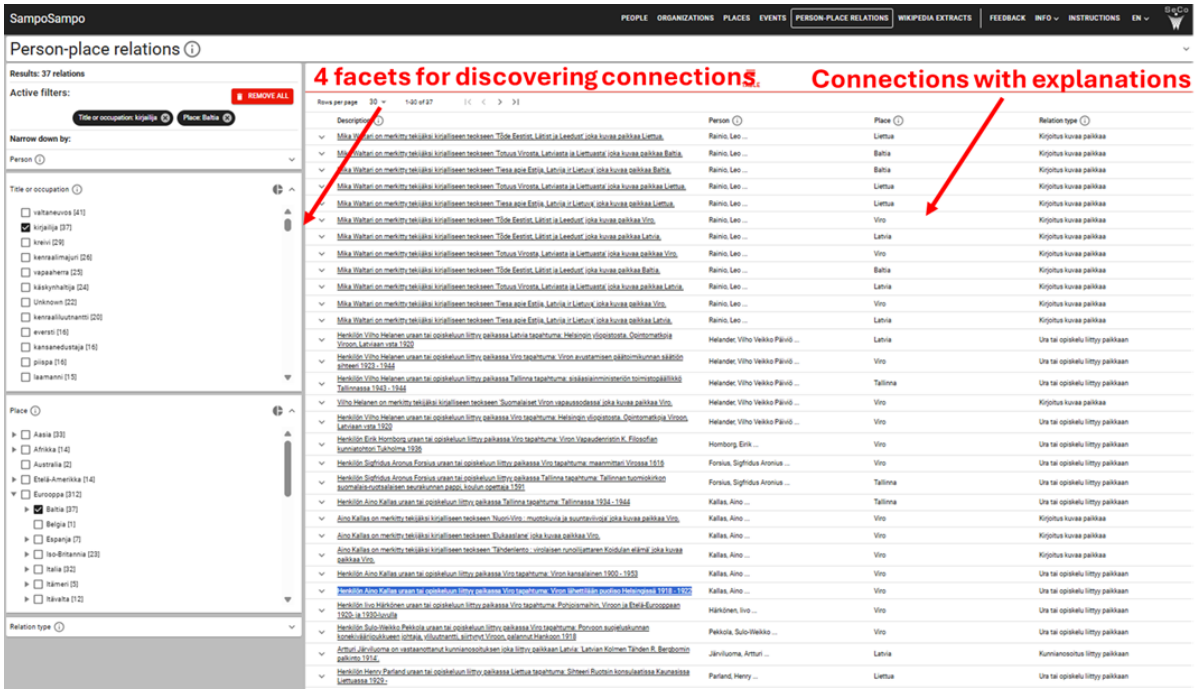

Figure 4: Findig connections with explanations between people and places in SampoSampo.

events through which the connections have been established. Ten event types have been used, such as getting an honorary price in a place and creating a painting depicting a place or a book about it. The user has selected from the Occupation facet the category “Author” and “Baltic countries” from the Place facet. On the right, 37 connections between people and places with their explanations are listed, such as “Relating to the studies or career of the person Aino Kallas, she was the spouse of the ambassador of Estonia in Finland in 1918 - 1922” (In Finnish: “Henkilön Aino Kallas uraan tai opiskeluun liittyy paikassa Viro tapahtuma: Viron lähettilään puoliso Helsingissä 1918 - 1922”.)‘

7. Conclusions

Entity alignment services, as discussed in Section 2 on Related Work, have been used for enriching data in new applications. In the same way, the SampoSampo LOD service presented in this paper can and has has been used for enriching data in the Finnish infrastructure context in new Sampo applications, such as the LetterSampo Finland (1809–1917) system. In addition to that, a major contribution of this paper is to present novel use cases and possibilities of entity alignment services: they can be used as a basis for application development in Digital Humanities research as demonstrated by the SampoSampo portal. This portal provides the end user with a single semantic portal and UI to access entities and their implicit "interesting" relations based on shared resources in a cloud on interlinked KGs. Using the portal, it is possible to search, browse, and analyze several local KGs on a global level with enriched or conflifcting descriptions of entities, including provenance information. Furthermore, the portal is capable for cross-cultural knowledge discovery of relatrions between entities [25, 26] and explaining them in natural language, using the knowledge-based approach to relational search [23, 15, 24].

Directions for further research include, e.g., extending the event perspective further, integrating new datasets in the relational search perspectives, and more thorough evaluations of the quality of the entity alignment processes presented and of the user experience of using the portal. SampoSampo will be published open source in 2025 as part of the Finnish DH research infrastructure FIN-CLARIAH.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

Acknowledgments

This research was supported by the Research Council of Finland FIN-CLARIAH funding from the European Union NextGenerationEU instrument. Computing resources of the CSC – IT Center for Science were used in our work.

References

[1] L. Askins, Learn how to see. Realize that everything connects to everything else. - Leorardo da Vinci, Independent publication, 2021. URL: https://www.amazon.com/Realize-everything-connects-Leonardo-Notebook/dp/B09L9Y5CP3.

[2] E. Hyvönen, Publishing and using cultural heritage linked data on the Semantic Web, Morgan & Claypool, Palo Alto, California, 2012. doi:10.2200/S00452ED1V01Y201210WBE003.

[3] E. Hyvönen, Serendipitous knowledge discovery on the Web of Wisdom based on searching and explaining interesting relations in knowledge graphs, Journal of Web Semantics (2025). doi:DOI:10.1016/j.websem.2024.100852.

[4] E. Hyvönen, Digital humanities on the Semantic Web: Sampo model and portal series, Semantic Web 14 (2022) 729–744. doi:10.3233/SW-223034.

[5] M. Doerr, The CIDOC CRM—an ontological approach to semantic interoperability of metadata, AI Magazine 24 (2003) 75–92.

[6] N. F. Noy, Ontology mapping, in: S. Staab, R. Studer (Eds.), Handbook on ontologies (2nd Edition), Springer–Verlag, 2009, pp. 573–590.

[7] T. B. Hickey, J. A. Toves, Managing ambiguity in VIAF, DLib Magazine 20 (2014). doi:doi:10.1045/july2014-hickey.

[8] W. Xia, E. Jiménez-Ruiz, V. V. Cross, Using BioPortal as a repository for mediating ontologies in ontology alignment, in: Workshop on Semantic Web Applications and Tools for Life Sciences, 2015. URL: https://api.semanticscholar.org/CorpusID:37359417.

[9] M. Frosterus, J. Tuominen, S. Pessala, E. Hyvönen, Linked open ontology cloud: managing a system of interlinked cross-domain light-weight ontologies, International Journal of Metadata, Semantics and Ontologies 10 (2015) 189–201. doi:10.1504/IJMSO.2015.073879.

[10] M. Dumontier, P.-Y. Vandenbussche, G. A. Atemezing, M. Poveda-Villalón, B. Vatant, Linked open vocabularies (lov): A gateway to reusable semantic vocabularies on the web, Semant. Web 8 (2017) 437–452. URL: https://doi.org/10.3233/SW-160213. doi:10.3233/SW-160213.

[11] A. Isaac, Europeana Data Model Primer, Technical Report, Europeana, 2023. URL: https://pro.europeana.eu/files/Europeana_Professional/Share_your_data/Technical_requirements/EDM_Documentation/EDM_Primer_130714.pdf.

[12] P. Leskinen, Modeling and Using Biographical Linked Data for Prosopographical Data Analysis, Ph.D. thesis, Aalto University, School of Science, Department of Computer Science, Espoo, Finland, 2024. URL: https://seco.cs.aalto.fi/publications/2024/leskinen-phd-2024.pdf.

[13] E. Hyvönen, P. Leskinen, M. Tamper, H. Rantala, E. Ikkala, J. Tuominen, K. Keravuori, BiographySampo – Publishing and enriching biographies on the Semantic Web for digital humanities research, in: The Semantic Web. 16th International Conference, ESWC 2019, Springer–Verlag, 2019, pp. 574–589. doi:10.1007/978-3-030-21348-0_37.

[14] E. Hyvönen, H. Rantala, Knowledge-based relation discovery in cultural heritage knowledge graphs, in: DHN 2019 Digital Humanities in Nordic Countries. Proceedings of the Digital Humanities in the Nordic Countries 4th Conference, CEUR Workshop Proceedings, Vol-2364, 2019, pp. 230–239. URL: http://www.ceur-ws.org/Vol-2364/.

[15] H. Rantala, E. Hyvönen, P. Leskinen, Finding and explaining relations in a biographical knowledge graph based on life events: Case biographysampo, in: Joint Proceedings of the ESWC 2023 Workshops and Tutorials co-located with 20th European Semantic Web Conference (ESWC 2023), volume 3443, CEUR Workshop Proceedings, 2023. URL: https://ceur-ws.org/Vol-3443/ESWC_2023_SEMMES_relations.pdf.

[16] M. Laouenan, P. Bhargava, J.-B. Eyméoud, O. Gergaud, G. Plique, E. Wasmer, A cross-verified database of notable people, 3500bc-2018ad, Scientific Data 9 (2022) 290. URL: https://www.nature.com/articles/s41597-022-01369-4.pdf.

[17] K. Verboven, M. Carlier, J. Dumolyn, A short manual to the art of prosopography, in: Prosopography approaches and applications. A handbook, Unit for Prosopographical Research (Linacre College), 2007, pp. 35–70.

[18] E. Ikkala, E. Hyvönen, H. Rantala, M. Koho, Sampo-UI: A full stack JavaScript framework for developing semantic portal user interfaces, Semantic Web 13 (2022) 69–84.

[19] H. Rantala, A. Ahola, E. Ikkala, E. Hyvönen, How to create easily a data analytic semantic portal on top of a SPARQL endpoint: introducing the configurable Sampo-UI framework, in: VOILA! 2023 Visualization and Interaction for Ontologies, Linked Data and Knowledge Graphs 2023, volume 3508, CEUR Workshop Proceedinbgs, 2023. URL: https://ceur-ws.org/Vol-3508/paper3.pdf.

[20] L. Rietveld, R. Hoekstra, The YASGUI family of SPARQL clients, Semantic Web 8 (2017) 373–383. doi:10.3233/SW-150197.

[21] E. Hyvönen, J. Tuominen, 8-star linked open data model: Extending the 5-star model for better reuse, quality, and trust of data, in: Posters, Demos, Workshops, and Tutorials of the 20th International Conference on Semantic Systems (SEMANTiCS 2024), volume 3759, CEUR Workshop Proceedings, 2024. URL: https://ceur-ws.org/Vol-3759/paper4.pdf.

[22] E. Hyvönen, H. Rantala, Knowledge-based relational search in cultural heritage linked data, Digital Scholarship in the Humanities (DSH) 36 (2021) 155–164. doi:https://doi.org/10.1093/llc/fqab042.

[23] E. Hyvönen, H. Rantala, Knowledge-based relational search in cultural heritage linked data, Digital Scholarship in the Humanities (DSH) 16 (2021) ii155–ii164. URL: https://doi.org/10.1093/llc/fqab042. doi:10.1093/llc/fqab042.

[24] H. Rantala, P. Leskinen, L. Peura, E. Hyvönen, Representing and searching associations in cultural heritage knowledge graphs using faceted search, in: SEMANTiCS 2024, 20th International Conference on Semantic Systems, proceedings, IOS Press, 2024. URL: https://seco.cs.aalto.fi/

publications/2024/rantala-et-al-searching-interesting-relations-2024.pdf, in press.
[25] J. Lehmann, J. Schüppel, S. Auer, Discovering unknown connections—the DBpedia relationship finder, in: Proc. of the 1st Conference on Social Semantic Web (CSSW 2007), volume 113 of *LNI*, GI, 2007, pp. 99–110. URL: http://subs.emis.de/LNI/Proceedings/Proceedings113/gi-proc-113-010.pdf.
[26] G. Tartari, A. Hogan, WiSP: Weighted shortest paths for RDF graphs, in: Proceedings of VOILA 2018, CEUR Workshop Proceedings, vol. 2187, 2018, pp. 37–52. URL: https://ceur-ws.org/Vol-2187/paper4.pdf.