Towards the Theatre Migrants Knowledge Graph

Jorit Hopp[1], Berenika Szymanski-Düll[1], Yan Lin[1] and Daniel Hernández[2]

[1]Theatre Studies, LMU, Munich, Germany
[2]Institute for Artificial Intelligence, University of Stuttgart, Stuttgart, Germany

Abstract

We present the Theatre Migrants Knowledge Graph, a structured Linked Data resource documenting the professional migrations of European theatre practitioners in the 19th century. The knowledge graph is developed within the ERC-funded T-MIGRANTS project at LMU Munich and currently comprises 3,163 persons, 1,173 migration events, 1,128 locations, 2,901 organizations, and 3,265 interpersonal relationships. We build on well-known vocabularies to describe migrations, religious affiliations, and temporal uncertainty in historical data. Entities are systematically linked to external authority files, including Wikidata, GeoNames, GND, VIAF, and ISNI. The knowledge graph is generated from a relational database through a reproducible RDF mapping pipeline that uses agentic AI (Claude Opus 4.6) as a design assistant. We investigate three questions: whether an implementer without domain knowledge can use agentic AI to produce a functional knowledge graph, what role the AI tool plays in the process, and what risks this approach introduces. We report our findings and discuss how symbolic validation layers can mitigate the identified risks.

Keywords

knowledge graph, digital humanities, theatre history, semantic web, linked data, cultural heritage, agentic AI

1. Introduction

The nineteenth century was a period of extraordinary mobility for theatre professionals. Actors, singers, directors, impresarios, and stagehands moved between cities and across international borders driven by labor opportunities, educational pursuits, and family ties, but also as a result of forced migration caused by wars, political upheavals, and religious persecution. Despite the historical significance of this phenomenon, the careers and networks of most theater migrants remain dispersed across archival sources in multiple languages and institutions, making systematic study difficult.

The T-MIGRANTS project, funded by the European Research Council (ERC) and based in Theatre Studies at LMU Munich, addresses this gap by systematically collecting and publishing data on European theater migrants in the 19th century—a period marked by mass migration, the rise of theater as an urban mass medium, and the emergence of modern performing arts professions [1]. This project investigates the mobility of theatre practitioners and examines how migrants shaped cultural institutions and aesthetic developments, and how cross-cultural flows of theater practitioners contributed to European cultural history.

A central objective of the project is to make this research data accessible and reusable in accordance with FAIR principles [2]. In an initial work, the project's team implemented a database, T-Migrants [3], over the Percona MySQL Server relational engine, and populated it with forms that were manually filled and uploaded into the database with ad-hoc tools. The T-Migrants database contains information about persons, their migration trajectories, professional activities, social relationships, organisational affiliations, and religious backgrounds as interconnected records. In addition to the T-Migrants database, the team developed several methods to interact with the database. Figure 1 presents the T-Migrants website, including the map visualization, the person detail pages, and charts with statistics about the migrants.

SemDH 2026: International Workshop on Semantic Digital Humanities, co-located with ESWC 2026
j.hopp@campus.lmu.de (J. Hopp); berenika.szymanski@lmu.de (B. Szymanski-Düll); yan.lin@campus.lmu.de (Y. Lin); daniel.hernandez@ki.uni-stuttgart.de (D. Hernández)
https://www.ki.uni-stuttgart.de/institute/team/Hernandez/ (D. Hernández)
0000-0002-3484-9120 (J. Hopp); 0009-0007-8669-7193 (B. Szymanski-Düll); 0009-0001-9484-6351 (Y. Lin); 0000-0002-7896-0875 (D. Hernández)

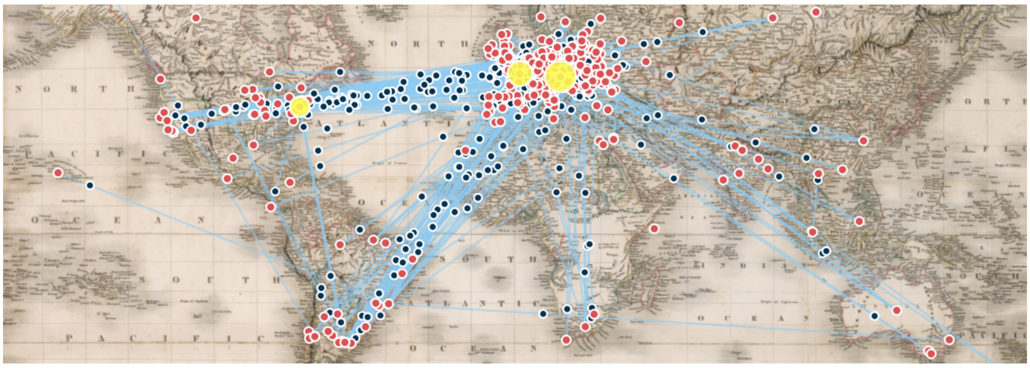

(a) Migration visualization. Red circles: locations; yellow circles: most frequent destinations (Vienna, Paris, New York); black circles: migrating people; lines: migration paths. The interface supports 11 historical cartographies from 1800 to the present.

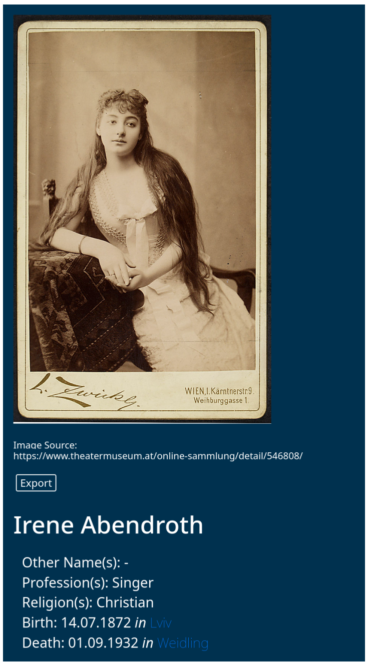

(b) Person view for opera singer Irene Abendroth.

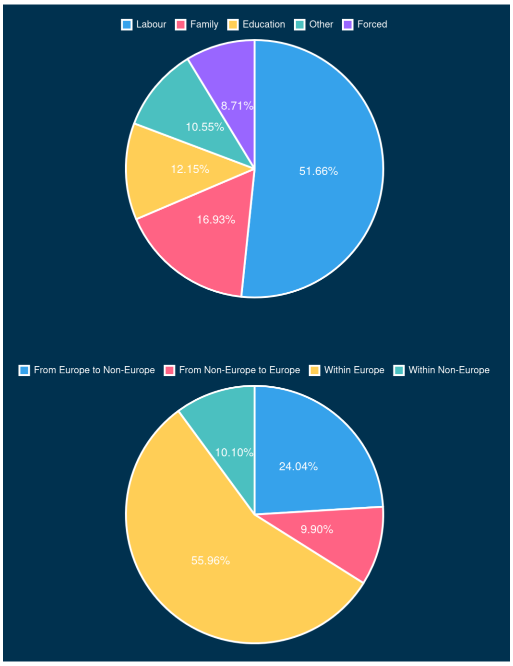

(c) Migration reasons and directions charts.

Figure 1: The T-Migrants website: (a) interactive map visualization, (b) detail page for Irene Abendroth, and (c) statistical charts of recorded migrations.

However, the relational model hinders the integration of the T-Migrants database with other data sources, in particular with data published on the Semantic Web. To be part of the Semantic Web, we developed the T-Migrants Knowledge Graph: an RDF representation of the T-Migrants database that is described in this paper.

This short paper presents the T-Migrants Knowledge Graph and investigates the use of agentic AI in its construction. We address three questions:

Q1. Can an implementer with no prior domain knowledge use an agentic AI tool to produce a functional knowledge graph from a relational database?

Q2. What role does the AI tool play in the process—code executor or design partner?

Q3. What are the risks of this approach, and how can they be mitigated?

We investigate these questions in the following setting: three of the authors developed the T-Migrants dataset as part of the T-MIGRANTS project and have deep domain knowledge. The fourth author, who implemented the mapping to RDF, had no prior knowledge of the dataset or of specific related ontologies (e.g., CIDOC-CRM, discussed in Section 2), and used Claude Opus 4.6 [13] as an agentic AI assistant throughout the process.

The main contributions of this paper are the following:

1. A preliminary version of the T-Migrants knowledge base: a curated RDF dataset integrated with Wikidata [4], GeoNames [5], DBpedia [6], and YAGO [7].
2. A preliminary version of an ontology to describe migrants, aligned with FOAF [8], Schema.org [9], OntoBio [10], the W3C Organization Ontology [11], and the W3C Time Ontology [12].
3. Empirical observations on the use of agentic AI in ontology design and data mapping, addressing Q1–Q3 above.

2. Related Work

Knowledge graphs for cultural heritage and the performing arts. The CIDOC Conceptual Reference Model (CRM) [14] is the foundational event-based ontology for cultural heritage, now standardised as ISO 21127. Large-scale knowledge graphs such as ArCo [15], which describes over 820 000 Italian cultural entities in 169 million triples, demonstrate the feasibility of national-scale cultural heritage KG engineering. In the performing arts domain, Mitsopoulou et al. [16] propose a dedicated ontology that addresses the ephemeral nature of performance events. The InTaVia knowledge graph [17] is particularly relevant to our work: it integrates 112 000 persons from four European national biographies using CIDOC-CRM and Bio CRM, enriched with Europeana and Wikidata links. Our approach differs from these CIDOC-CRM–based projects by building on ontologies like Schema.org as a lighter-weight foundation. Our decision to build on Schema.org was not motivated by simplifying adoption and alignment with the broader Linked Open Data ecosystem—which could be a benefit—but rather an initial decision made in the process of understanding the data we have at hand. As future work, we need to decide whether to follow the current path or align with CIDOC-CRM.

Historical migration and biographical data. Structured datasets for historical migration are still rare. The Mapping Manuscript Migrations project [18] tracks the provenance of medieval manuscripts across collections, aggregating over 20 million RDF triples from heterogeneous sources using CIDOC-CRM. Vesalainen et al. [19] create a dataset of over six million migration entries extracted from 19th-century Finnish church records, covering a time period similar to ours. For biographical modeling, Bio CRM [20] extends CIDOC-CRM with classes for biographical events and distinguishes between enduring roles and perduring events—a distinction that our ontology also makes, though without formal alignment to a foundational ontology. OntoBio [10] provides a biographical ontology grounded in Linked Data principles, with which our ontology is aligned. This is not a final design decision, but a preliminary one in the current development stage of the T-Migrants Knowledge Graph.

AI-assisted ontology engineering. The use of large language models (LLMs) for ontology engineering is a rapidly growing area. Caufield et al. [21] survey 34 papers and identify 46 distinct tasks where LLMs have been applied, from specification to maintenance. Shimizu and Hitzler [22] argue that LLM-based ontology engineering is an emerging research direction, with modular approaches being particularly promising. On the practical side, Meyer et al. [23] report experiments in which GPT-4 successfully substituted custom classes and properties with standard vocabularies such as FOAF and the W3C Organization Ontology—a task similar to the vocabulary alignment step in our pipeline. Gonzalez-Carvajal et al. [24] present OntoGenix, a multi-agent GPT-4 pipeline that generates OWL ontologies and RML mappings from tabular datasets. Lippolis et al. [25] propose Ontogenia, which uses metacognitive prompting to generate ontologies from LLMs, demonstrating that structured prompting strategies can improve the quality of LLM-generated ontological artifacts. Our work contributes to this line of research by reporting empirical observations on the use of Claude Opus 4.6 in a real-world ontology design and data mapping project. Unlike prior work that focuses on specific subtasks (e.g., vocabulary alignment [23] or ontology generation [24, 25]), we examine the end-to-end process: from relational data to a linked, vocabulary-aligned knowledge graph, using SPARQL UPDATE queries as the transformation language. We focus on how the human–AI interaction evolved during the process and what validation gaps this approach introduces.

It is worth noting that several established approaches exist for specific subtasks in our pipeline. Virtual Knowledge Graph systems such as Ontop [26] provide an alternative to materialized RDF mappings by exposing relational data through SPARQL without data duplication. For ontology alignment, dedicated tools like BERTMap [27] offer validated, reproducible results. For entity linking of historical figures—a particularly challenging task due to name variations and sparse records—specialized methods such as those proposed by Graciotti et al. [28] for musical heritage entities are more reliable than general-purpose LLMs. In our current work, the LLM handled these subtasks in an end-to-end fashion; integrating specialized tools into the pipeline is a natural direction for improvement (see Section 6).

3. Ontology Design

The ontology was designed to capture the core entities and relationships present in the T-Migrants database (cf. the person view and statistical charts in Figures 1b and 1c of Figure 1). The ontology distinguishes between two kinds of entities. On the one hand, *endurants* are entities that are wholly present at each moment of their existence: schema:Person, schema:Place, and schema:Organization. On the other hand, *perdurants* are entities that unfold over time and have temporal parts: Migration, Membership, Relationship, PersonName, and ReligionAffiliation. In the current version of the ontology, this distinction is implicit in the modeling but not formally represented through inheritance from a foundational ontology such as DOLCE [29] or BFO [30]. Aligning the ontology with a foundational ontology is planned as future work.

Figure 2 presents an overview of the ontology we have developed so far using the Chowlk conceptual diagramming language [31]. Colors are added to highlight the RDF vocabularies each term comes from (e.g., yellow for FOAF, green for Schema.org, and blue for our proposed ontology, denoted with prefix tm). The perdurant classes are the ones that are included in the union class (denoted with the circle with the symbol $\sqcup$). Each of these perdurant instances can be annotated with time (time:hasTime) and location (tm:placedIn) properties. We extended the W3C Time ontology with a class to introduce string descriptions of a temporal entity, as well as to denote that the available information was imprecise.

Figure 3 presents some example data describing Irene Abendroth. In particular, it shows how she migrates from Karlovy Vari to Vienna because of her work (individual :AbeIre-00-migration-6). The start and destination places are annotated with the properties tm:startPlace and tm:destinationPlace, respectively. The migration reason is classified as tm:LabourMigrationReason, which is a subclass of tm:MigrationReason.[1] To locate the migration in the time, we use the :time-1 individual with data property tm:dateStartFuzzy "1889," which represents the starting date of the migration. That is, it is known that Abendroth migrated in 1889 but the exact date is unknown. This uncertain information is also represented with the object property time:hasBeginning, whose value represents a date occurring between two certain boundary dates, the first and last day of year 1889.

Figure 4 describes how Abendroth's relationships, occupations, and memberships are described. There is a concrete relationship, the sister relationship :relationship-16839, and a relationship whose type is not identified, relationship :relationship-3. As shown, some relationships can include a location. Abendroth's occupation is singer, and she has worked in this occupation in :work-2, when working for *Hofoper Wien*. For the sake of space, we did not include class tm:EmploymentTour in Figure 2. Class tm:EmploymentTour should be included as a subclass of class org:Membership, like many other types of work described in the original dataset.

[1]For the sake of readability, Figure 2 does not include class `tm:LabourMigrationReason`, as well as other subclasses of class `tm:MigrationReason` like, e.g., `tm:EducationMigrationReason`.

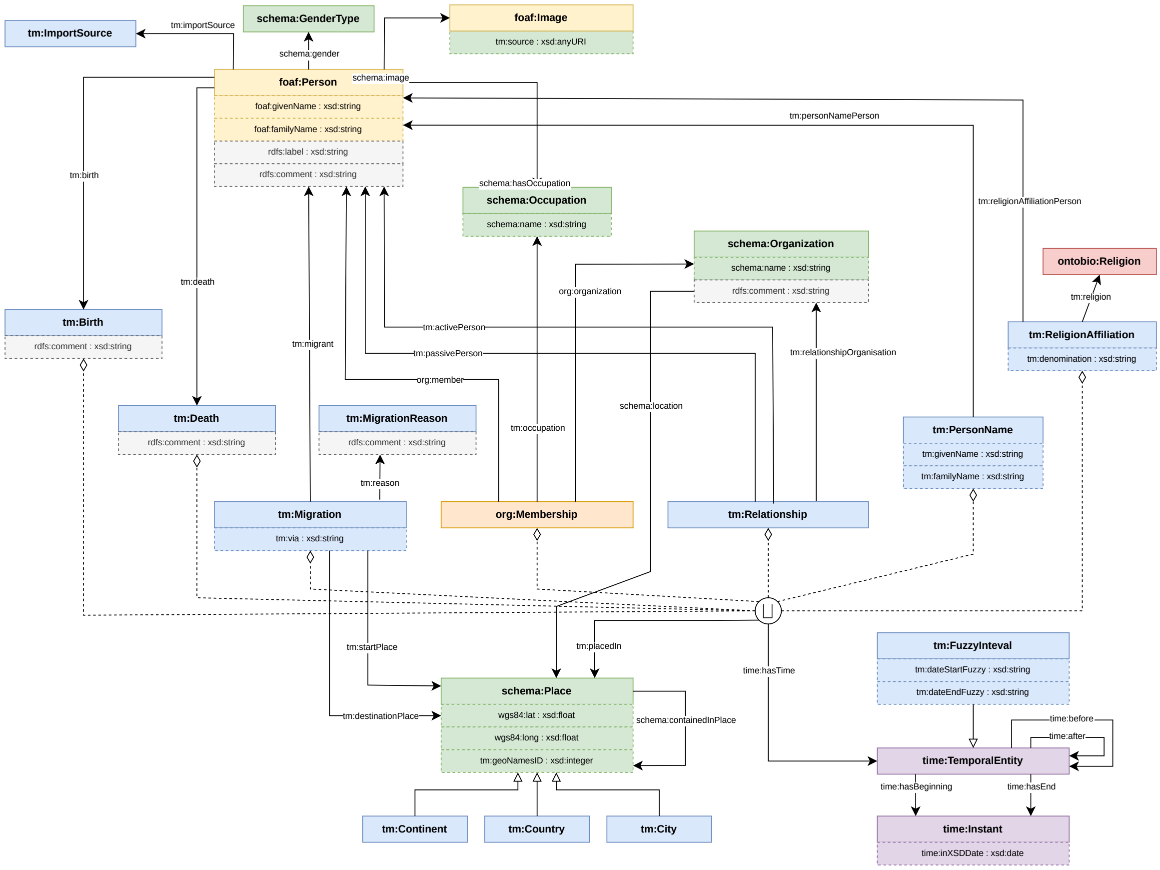

Figure 2: Ontology overview. Classes are shown as colored rectangles. Arrows denote object properties; attached boxes show selected datatype properties. The circle with symbol ⊔ represents a union class.

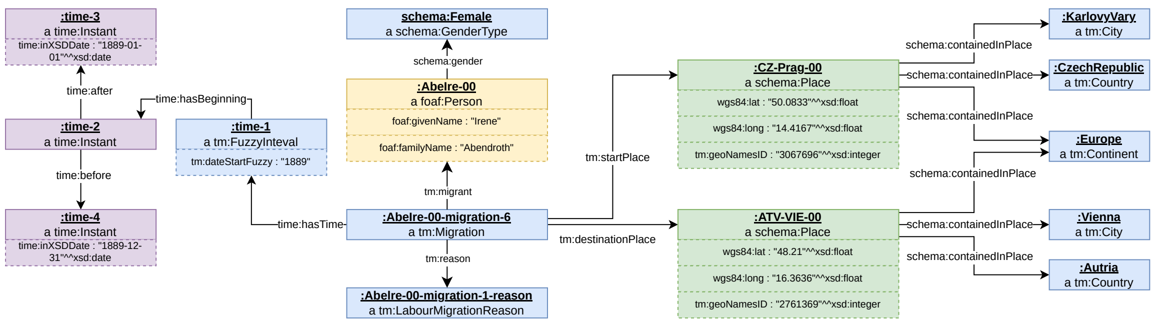

Figure 3: Example subgraph for the opera singer Irene Abendroth, showing her migration from Karlovy Vari to Vienna. We underscore individual entity names and include individual classes in the entities.

4. Mapping Pipeline

The T-Migrants knowledge graph is produced from the relational engine through a multi-step pipeline, managed by a Rakefile that tracks file dependencies and ensures reproducibility. In each step k, we provide a task description to the agentic AI model and ask it to generate: the updated RDF graph G_k, the SPARQL UPDATE [32] queries that produce G_k from G_{k-1}, a summary subgraph centered on Irene Abendroth (used as a running example), the inferred ontology, and the corresponding diagrams. We next describe the four task descriptions, S1–S4, to illustrate our workflow.

S1. *Direct mapping to RDF.* Load the MariaDB dump into a Docker container and write a Rust script

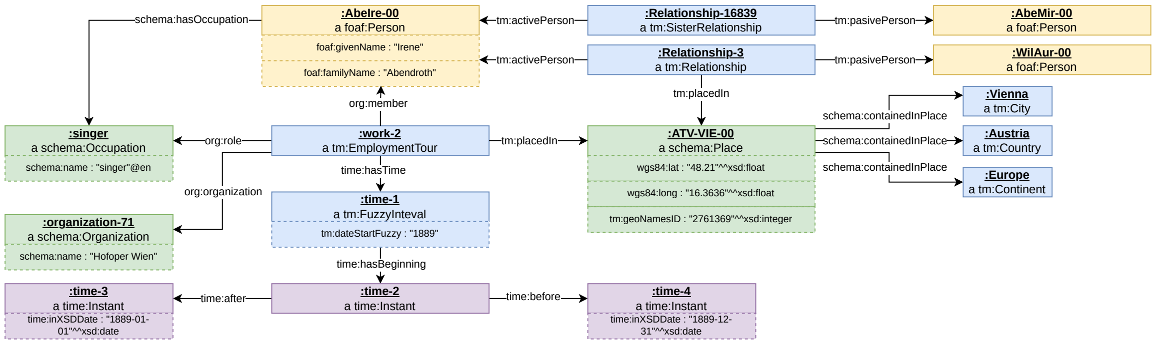

Figure 4: Example subgraph showing Irene Abendroth's occupation and relationships.

that exports the data as RDF, using prefixes tm and tmd for the ontology and data.[2]

S2. *Normalise the data* consolidating redundant nodes, rewriting string-valued foreign keys as IRI references, and adding rdfs:label values to enumeration instances.

S3. *Datatype annotation and cleaning.* Annotate plain literals with explicit XSD datatypes (e.g., xsd:date, xsd:float, xsd:anyURI) and remove empty string literals introduced by NULL columns.

S4. *Vocabulary alignment.* Replace custom properties and classes with well-known vocabulary equivalents from Schema.org, WGS84 [34], and SKOS [35]. Map authority identifiers to Wikidata normalised properties (e.g., wdtn:P227 for GND), and Wikidata entity links to schema:sameAs. Reify occupation strings into schema:Occupation instances. For enumeration classes (professions, migration reasons, relationship types) use skos:Concept typing.

We used Claude Opus 4.6 for its large context window (allowing the full database schema and ontology state to be loaded at each step), its ability to generate SPARQL and Rust code, and its support for iterative dialogue. The SPARQL UPDATE queries are executed by the in-memory Oxigraph [36] store. Some transformations caused information loss, which we recovered by modifying the corresponding step. For example, in S4 the model replaced the custom property tmd:religion with schema:memberOf, discarding temporal metadata about religion affiliation; after inspecting the Abendroth subgraph, we reverted to a reified tm:ReligionAffiliation class. At each step, the implementer inspected the generated subgraph and ontology diagrams and resolved issues with the domain expert coauthors. The decisions made in S4 are not definitive, as this is a work in progress.

5. Conclusion

We return to the questions posed in the introduction. Regarding **Q1** (feasibility): the resulting knowledge graph covers the core entities of the T-Migrants database and is aligned with well-known vocabularies. The process did not follow established ontology engineering methodologies, yet produced a functional knowledge graph in approximately 10 hours of implementer effort. This suggests that agentic AI can substantially lower the barrier for domain outsiders to produce an initial KG from relational data. Regarding **Q2** (role of the AI tool): the interaction evolved from fine-grained task specification (e.g., writing concrete SPARQL UPDATE queries) to a discussion mode in which the model proposed design alternatives and the implementer made the final decisions. The AI tool was most effective as a design assistant in dialogue with the human expert, not as an autonomous pipeline. Regarding **Q3** (risks and mitigation): the main risk is the absence of formal validation, as discussed in Section 6. If we were to repeat this process, we would integrate symbolic validation tools (reasoners, pitfall scanners, SHACL shapes) from the start, rather than relying solely on manual inspection of intermediate results.

A secondary observation is that the rigid structure of the relational data model complicates the

[2]We did not tell the model what a direct mapping is; it found the W3C Direct Mapping specification [33] on the web.

curation of heterogeneous biographical data. We are therefore exploring a more flexible, ontology-driven curation interface generated from SHACL shape descriptions.

Reproducibility Statement. The supplementary materials including code and ontologies generated in this work are archived for long-term preservation in DaRUS [37].

6. Limitations

The ontology and knowledge graph have not yet been validated with automated tools. Validation in the current pipeline relies on manual inspection of subgraphs and ontology diagrams at each step, complemented by discussion with the domain experts. While this review caught several issues (see Section 4), it does not substitute systematic validation with pitfall scanners such as OOPS! [38], OWL 2 reasoners [39], or SHACL shapes [40, 41].

Our pipeline delegates substantial design work—vocabulary selection, property mapping, data cleaning—to a commercial LLM (Claude Opus 4.6). The model may hallucinate plausible but incorrect mappings, overlook established tools (e.g., Virtual Knowledge Graph systems [26], dedicated alignment tools [27], or specialized historical entity linking methods [28]), and its proprietary nature raises reproducibility concerns. These risks can be mitigated by reinforcing the symbolic validation layers, following the neuro-symbolic principle that neural components should be complemented by symbolic verification [42]: just as software quality relies on automated checks (linters, tests, code review), AI-generated ontological artifacts can be verified by symbolic tools. In such a hybrid approach, the LLM proposes transformations while symbolic tools verify their correctness. Implementing this multi-layer validation is our immediate next step.

Acknowledgments

This work was partially funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant agreement No. 850742 (T-MIGRANTS—Crossing Borders: The Agency of Nineteenth-Century European Theatre Migrants). The work of Hernández is funded by projects where he researches the development of domain ontologies and knowledge graphs: the German Research Foundation (DFG): SFB 1574—471687386 and EXC 2120/1—390831618; and the European Research Council (ERC) under the European Union's Horizon Europe Research and Innovation Programme, grant agreement No. 101200433, project META-LEARN.

Declaration on Generative AI

During the preparation of this work, the authors used Claude Opus 4.6 as an agentic AI tool for the ontology design, data mapping pipeline, and knowledge graph construction described in Section 4. Additionally, it was used to improve the grammar and clarity of the text, fix the references in the bibliography, and explore the research area for finding related work. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] B. Szymanski-Düll, L. Skwirblies (Eds.), European Theatre Migrants in the Age of Empire: Personal Experiences, Transnational Trajectories, and Socio-Political Impacts, Palgrave Studies in Performance and Migration, Palgrave Macmillan, 2025. doi:10.1007/978-3-031-69836-1.

[2] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne, et al., The FAIR guiding principles for scientific data management and stewardship, Scientific Data 3 (2016) 160018. doi:10.1038/sdata.2016.18.

[3] B. Szymanski-Düll, et al., T-Migrants View Data, 2026. URL: https://www.t-migrants.gwi.uni-muenchen.de/, accessed: 2026-03-07.
[4] D. Vrandečić, M. Krötzsch, Wikidata: A free collaborative knowledgebase, Communications of the ACM 57 (2014) 78–85. doi:10.1145/2629489.
[5] M. Wick, GeoNames geographical database, 2006. URL: https://www.geonames.org/, accessed: 2026-03-07.
[6] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, C. Bizer, DBpedia – a large-scale, multilingual knowledge base extracted from Wikipedia, Semantic Web 6 (2015) 167–195. doi:10.3233/SW-140134.
[7] F. M. Suchanek, G. Kasneci, G. Weikum, YAGO: A core of semantic knowledge, in: Proceedings of the 16th International Conference on World Wide Web, 2007, pp. 697–706. doi:10.1145/1242572.1242667.
[8] D. Brickley, L. Miller, FOAF vocabulary specification 0.99, 2014. URL: http://xmlns.com/foaf/spec/, accessed: 2026-03-07.
[9] R. V. Guha, D. Brickley, S. Macbeth, Schema.org: Evolution of structured data on the web, Communications of the ACM 59 (2016) 44–51. doi:10.1145/2844544.
[10] B. Dutta, S. Arzoo, Towards a biographical ontology: The OntoBio framework and its applications, Knowledge Organization 53 (2026). doi:10.31083/KO45190.
[11] D. Reynolds, The organization ontology, 2014. URL: https://www.w3.org/TR/vocab-org/, W3C Recommendation, 16 January 2014.
[12] S. Cox, C. Little, Time ontology in OWL, 2017. URL: https://www.w3.org/TR/owl-time/, W3C Recommendation, 19 October 2017.
[13] Anthropic, Claude Opus 4.6 system card, 2026. URL: https://www.anthropic.com/claude/opus.
[14] M. Doerr, The CIDOC conceptual reference module: An ontological approach to semantic interoperability of metadata, AI Magazine 24 (2003) 75–92. doi:10.1609/aimag.v24i3.1720.
[15] V. A. Carriero, A. Gangemi, M. L. Mancinelli, L. Marinucci, A. G. Nuzzolese, V. Presutti, C. Veninata, ArCo: The Italian cultural heritage knowledge graph, in: Proceedings of the 18th International Semantic Web Conference (ISWC 2019), volume 11779 of *LNCS*, Springer, 2019, pp. 36–52. doi:10.1007/978-3-030-30796-7_3.
[16] E. Mitsopoulou, K. Kyprianos, P. Brattis, Documenting the ephemeral: An ontology for the performing arts, Journal of Information Science (2024). doi:10.1177/01655515241271052.
[17] M. Schlögl, K. Lejtovicz, J. Tuominen, et al., The InTaVia knowledge graph – European national biographical and cultural heritage data, Semantic Web (2025). URL: https://www.semantic-web-journal.net/system/files/swj3851.pdf, pre-print.
[18] T. Burrows, D. Emery, et al., Mapping manuscript migrations knowledge graph: Data for tracing the history and provenance of medieval and renaissance manuscripts, Journal of Open Humanities Data (2020). doi:10.5334/johd.14.
[19] A. Vesalainen, J. Kanerva, A. Nitsch, K. Korsu, I. Larkiola, L. Ruotsalainen, F. Ginter, Creating a historical migration dataset from Finnish church records, 1800–1920, Journal of Open Humanities Data (2025). doi:10.5334/johd.345.
[20] J. Tuominen, E. Hyvönen, P. Leskinen, Bio CRM: A data model for representing biographical data for prosopographical research, in: Proceedings of the Second Conference on Biographical Data in a Digital World (BD2017), volume 2119 of *CEUR Workshop Proceedings*, 2018, pp. 59–66. URL: https://ceur-ws.org/Vol-2119/paper10.pdf.
[21] J. H. Caufield, et al., Large language models for ontology engineering: A systematic literature review, Semantic Web (2025). URL: https://www.semantic-web-journal.net/content/large-language-models-ontology-engineering-systematic-literature-review, pre-print.
[22] C. Shimizu, P. Hitzler, Accelerating knowledge graph and ontology engineering with large language models, Journal of Web Semantics (2025). doi:10.1016/j.websem.2025.100862.
[23] L.-P. Meyer, C. Stadler, J. Frey, N. Radtke, K. Junghanns, R. Meissner, G. Dziwis, K. Bulert, M. Martin, LLM-assisted knowledge graph engineering: Experiments with ChatGPT, in: AIDRST 2023, Informatik aktuell, Springer, 2024. doi:10.1007/978-3-658-43705-3_8.

[24] S. Gonzalez-Carvajal, et al., OntoGenix: Leveraging large language models for enhanced ontology engineering from datasets, Information Processing & Management 62 (2024). doi:10.1016/j.ipm.2024.104042.

[25] A. S. Lippolis, M. Ceriani, S. Zuppiroli, A. G. Nuzzolese, Ontogenia: Ontology generation with metacognitive prompting in large language models, in: European Semantic Web Conference, Springer, 2024, pp. 259–265. doi:10.1007/978-3-031-60635-9_15.

[26] G. Xiao, D. Lanti, R. Kontchakov, S. Komla-Ebri, E. Güzel-Kalaycı, L. Ding, J. Corman, B. Cogrel, D. Calvanese, E. Botoeva, The virtual knowledge graph system Ontop, in: Proceedings of the 19th International Semantic Web Conference (ISWC 2020), volume 12507 of *LNCS*, Springer, 2020, pp. 259–277. doi:10.1007/978-3-030-62466-8_17.

[27] Y. He, J. Chen, D. Antonyrajah, I. Horrocks, BERTMap: A BERT-based ontology alignment system, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 36, 2022, pp. 5684–5691. doi:10.1609/aaai.v36i5.20510.

[28] A. Graciotti, N. Lazzari, V. Presutti, R. Tripodi, Musical heritage historical entity linking, Artificial Intelligence Review 58 (2025) 140. doi:10.1007/s10462-025-11132-9.

[29] S. Borgo, R. Ferrario, A. Gangemi, N. Guarino, C. Masolo, D. Porello, E. M. Sanfilippo, L. Vieu, DOLCE: A descriptive ontology for linguistic and cognitive engineering, Applied Ontology 17 (2022) 45–69. doi:10.3233/AO-210259.

[30] R. Arp, B. Smith, A. D. Spear, Building Ontologies with Basic Formal Ontology, MIT Press, 2015. doi:10.7551/mitpress/9780262527811.001.0001.

[31] S. Chávez-Feria, R. García-Castro, M. Poveda-Villalón, Chowlk: From UML-based ontology conceptualizations to OWL, in: The Semantic Web. ESWC 2022, volume 13261 of *LNCS*, Springer, 2022. doi:10.1007/978-3-031-06981-9_20.

[32] P. Gearon, A. Passant, A. Polleres, SPARQL 1.1 update, 2013. URL: https://www.w3.org/TR/sparql11-update/, W3C Recommendation, 21 March 2013.

[33] M. Arenas, A. Bertails, E. Prud'hommeaux, J. Sequeda, A direct mapping of relational data to RDF, 2012. URL: https://www.w3.org/TR/rdb-direct-mapping/, W3C Recommendation, 27 September 2012.

[34] D. Brickley, Basic geo (WGS84 lat/long) vocabulary, 2003. URL: https://www.w3.org/2003/01/geo/, W3C Semantic Web Interest Group.

[35] A. Miles, S. Bechhofer, SKOS simple knowledge organization system reference, 2009. URL: https://www.w3.org/TR/skos-reference/, W3C Recommendation, 18 August 2009.

[36] T. Pellissier Tanon, Oxigraph, 2024. URL: https://oxigraph.org/. doi:10.5281/zenodo.7408022.

[37] J. Hopp, B. Szymanski-Düll, Y. Lin, D. Hernández, Towards the Theatre Migrants Knowledge Graph – Transformation Code and Dataset, 2026. URL: https://doi.org/10.18419/DARUS-5918. doi:10.18419/DARUS-5918.

[38] M. Poveda-Villalón, A. Gómez-Pérez, M. C. Suárez-Figueroa, OOPS! (OntOlogy Pitfall Scanner!): An on-line tool for ontology evaluation, International Journal on Semantic Web and Information Systems 10 (2014) 7–34. doi:10.4018/ijswis.2014040102.

[39] R. C. Jackson, J. P. Balhoff, E. Douglass, N. L. Harris, C. J. Mungall, J. A. Overton, ROBOT: A tool for automating ontology workflows, BMC Bioinformatics 20 (2019) 407. doi:10.1186/s12859-019-3002-3.

[40] H. Knublauch, D. Kontokostas, Shapes constraint language (SHACL), 2017. URL: https://www.w3.org/TR/shacl/, W3C Recommendation, 20 July 2017.

[41] P. Seifer, D. Hernández, R. Lämmel, S. Staab, From shapes to shapes: Inferring SHACL shapes for results of SPARQL CONSTRUCT queries, in: Proceedings of the ACM Web Conference 2024, 2024. doi:10.1145/3589334.3645550.

[42] V. Belle, G. Marcus, The future is neuro-symbolic: Where has it been, and where is it going?, Proceedings of the AAAI Conference on Artificial Intelligence 40 (2026) 40954–40961. doi:10.1609/aaai.v40i48.42130.