

EmpiRE-Compass: A Neuro-Symbolic Dashboard for Sustainable and Dynamic Knowledge Exploration, Synthesis, and Reuse

Navigating the Knowledge Landscape of Empirical Research in Requirements Engineering

Oliver Karras^{1,*}, Amirreza Alasti², Lena John¹, Sushant Aggarwal² and Yücel Celik²

¹TIB - Leibniz Information Centre for Science and Technology, Germany

²Leibniz University Hannover, Germany

Abstract

[Background.] Software engineering (SE) and requirements engineering (RE) face a significant increase in secondary studies, particularly literature reviews (LRs), due to the ever-growing number of scientific publications. Generative artificial intelligence (GenAI) exacerbates this trend by producing LRs rapidly but often at the expense of quality, rigor, and transparency. At the same time, secondary studies often fail to share underlying data and artifacts, limiting replication and reuse. **[Objective.]** This paper introduces *EmpiRE-Compass*, a neuro-symbolic dashboard designed to lower barriers for accessing, replicating, and reusing LR data. Its overarching goal is to demonstrate how LRs can become more sustainable by semantically structuring their underlying data in research knowledge graphs (RKGs) and by leveraging large language models (LLMs) for easy and dynamic access, replication, and reuse. **[Method.]** Building on two RE use cases, we developed *EmpiRE-Compass* with a modular system design and workflows for curated and custom competency questions. The dashboard is freely available online, accompanied by a *demonstration video*. To manage operational costs, a limit of 25 requests per IP address per day applies to the default LLM (GPT-4o mini). All source code and documentation are released as an *open-source project* to foster reuse, adoption, and extension. **[Results.]** *EmpiRE-Compass* provides three core capabilities: (1) Exploratory visual analytics with visualizations, interpretations, and explanations for curated competency questions; (2) Neuro-symbolic synthesis for custom competency questions with dynamic visualizations, interpretations, explanations, and direct access to underlying RKG data; and (3) Reusable knowledge with all queries, analyses, and results openly available. **[Conclusions.]** By unifying RKGs and LLMs in a neuro-symbolic dashboard, *EmpiRE-Compass* advances sustainable LRs in RE, SE, and beyond. It lowers technical barriers, fosters transparency and reproducibility, and enables collaborative, continuously updated, and reusable LRs.

Keywords

Neuro-symbolic dashboard, research knowledge graph, large language model, sustainable literature review

1. Introduction

Software engineering (SE), including requirements engineering (RE), faces a significant increase in secondary studies, particularly literature reviews (LRs), due to the ever-growing number of scientific publications [1, 2, 3]. The rise of generative artificial intelligence (GenAI) has exacerbated this problem, as it is increasingly used to generate LRs rapidly and automatically, but often at the expense of quality, rigor, and transparency – a trend so concerning that major preprint servers have changed their policies¹ [4].

The challenge of LRs, however, goes beyond the recent GenAI trend: The core problem is that extracted and analyzed data, along with the corresponding research artifacts, are often unavailable [5]. This lack of underlying data and artifacts poses a substantial obstacle to collaboration and updating of LRs, as the analyses and results cannot be replicated and reused [3]. A recent systematic mapping study found that secondary studies in SE, published between 2013 and 2023, are increasingly adopting

Joint Proceedings of REFSQ-2026 Workshops, Doctoral Symposium, Posters & Tools Track, and Education and Training Track. Co-located with REFSQ 2026. Poznan, Poland, March 23-26, 2026

*Corresponding author.

✉ oliver.karras@tib.eu (O. Karras); amirreza.alasti@stud.uni-hannover.de (A. Alasti); lena.john@tib.eu (L. John); sushant.aggarwal@stud.uni-hannover.de (S. Aggarwal); yucel.celik@stud.uni-hannover.de (Y. Celik)

ORCID 0000-0001-5336-6899 (O. Karras); 0009-0002-1165-773X (A. Alasti); 0009-0007-2097-9761 (L. John)



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

¹arXiv banned non-peer-reviewed LRs from computer science to counteract the flood of low-quality, AI-generated content [4].

open science practices. However, only 31.5% of the 537 analyzed studies shared their data and artifacts, and just 12.1% deposited them in a persistent repository [6]. Although openly available data and artifacts in persistent repositories represent progress toward the broader vision of open science in SE [6], current practices still fail to leverage the full technological potential to support replication and reuse effectively [3, 5, 7, 8]. It is necessary to move beyond static, file-based representations in repositories by using semantically structured representations integrated into open science infrastructures aligned with the FAIR data principles [1, 7, 9]. In this way, data and artifacts become interpretable and actionable for humans and machines, enabling *sustainable* exploration, synthesis, and reuse [3, 8].

Building on this vision, we used the *Open Research Knowledge Graph (ORKG)* [2, 10] to demonstrate how data from an LR with 776 papers on the state and evolution of empirical research in RE can be semantically structured in an open science infrastructure to improve its availability for replication and reuse [5]. The resulting research knowledge graph (RKG), called *KG-EmpIRE*, and its *analysis script* constitute long-term available and reusable research artifacts² [11, 12]. Based on these results, we advocated for the broader use of RKGs for sustainable LRs in SE and beyond [3]. However, while RKGs ensure that LR data is openly available and semantically structured, their effective use demands advanced technical skills, such as writing SPARQL queries or navigating complex graph schemas, to represent, query, analyze, and interpret the data. Many researchers lack these skills and thus face barriers that continue to impede access, replication, and reuse.

To address these barriers, we build on the established use of dashboards to enable intuitive, human-centered access. While dashboards support exploration, synthesis, and reuse, they typically provide predefined visualizations with only minor customization that users must interpret themselves [13]. We go beyond these limitations by integrating a neuro-symbolic approach that leverages the synergy between large language models (LLMs) and RKGs. This approach unifies the neural, generative capabilities of LLMs with symbolic representations of semantically structured data in RKGs [14]. In particular, LLMs can *dynamically* generate custom visualizations, interpretations, and explanations directly from RKG data as answers to competency questions³. This synergy enables researchers to interact with RKGs easily and flexibly, lowering the barriers to access, replication, and reuse.

We applied this idea to our KG-EmpIRE use case [5], resulting in *EmpIRE-Compass*: A neuro-symbolic dashboard that allows researchers to explore, synthesize, and reuse knowledge about empirical research in RE *sustainably* and *dynamically*. *EmpIRE-Compass* offers three core capabilities: (1) *Exploratory visual analytics*, investigating curated competency questions using predefined and dynamic visualizations, interpretations, and explanations; (2) *Neuro-symbolic synthesis*, combining RKGs and LLMs to answer custom competency questions with dynamically generated visualizations, interpretations, explanations, and direct access to the underlying data; and (3) *Reusable knowledge*, making all structured data, queries, analyses, and results openly available for replication, reuse, and sharing. Beyond KG-EmpIRE, we already extended the dashboard with a second use case on *empirical research in NLP4RE* [15] based on 50 *NLP4RE ID Cards* [16], which are organized in the *ORKG* [17]. We make the following contributions:

1. Presenting *EmpIRE-Compass*, along with its system design and workflows, incorporating both use cases (KG-EmpIRE and NLP4RE ID Card) and demonstrating the integration of a neuro-symbolic approach into a dashboard for LR data.
2. Releasing the complete source code and documentation as *open-source project* [18] to foster reuse, adoption, and extension within the research community.

2. Related Work

Dashboards are established tools for making complex data available and accessible across domains [19]. We identified 8 dashboards⁴ developed to support exploration, synthesis, and reuse of LR data in various disciplines, such as environmental science [20], *software architecture* [21], *biodiversity* [22], *atmospheric science* [23], *social science* [24], educational science [25], *invasion biology* [26], and *virology* [27].

²Recognized at *RE'24 Artifact Track* with the badges “Available” and “Reusable” as well as the *Best Artifact Award*.

³A competency question is a natural language query expressing an information need that an RKG must be able to answer.

⁴Two dashboards are no longer accessible. Links to the remaining six dashboards are included in the listed disciplines.

Among these eight dashboards, two approaches [20, 21] rely on spreadsheets, three [22, 23, 24] use different RKGs, and three [25, 27, 26] build directly on the ORKG. Kaplan et al. [21], while currently using spreadsheets, plan to migrate their data into the ORKG. The prominent use of RKGs reflects the envisioned broader transition toward semantically structured representations of LR data integrated into open science infrastructures aligned with the FAIR data principles.

Although these dashboards demonstrate clear benefits for improving data availability, accessibility, and transparency, they remain limited to predefined visualizations with only minor customization and provide no interpretations or explanations to users. Our work builds on these efforts and benefits, but goes further by enhancing exploration, synthesis, and reuse through unification of LLMs and RKGs – combining semantically structured data with flexible and intuitive user interaction using natural language. To the best of our knowledge, dashboards for LR data have not previously employed a neuro-symbolic approach. *EmpiRE-Compass* demonstrates how this integration enables sustainable and dynamic knowledge exploration, synthesis, and reuse.

3. The Neuro-Symbolic Dashboard: *EmpiRE-Compass*

EmpiRE-Compass is a progressive web application designed to support researchers in exploring, synthesizing, and reusing knowledge about empirical research in RE through curated and custom competency questions. Its overarching goal is to demonstrate how LRs can become more sustainable by semantically structuring their underlying data in RKGs and by leveraging LLMs for easy and dynamic access. This neuro-symbolic approach enhances availability, accessibility, and transparency for replication and reuse. Figure 1 illustrates the system design and the two main workflows, explained below.

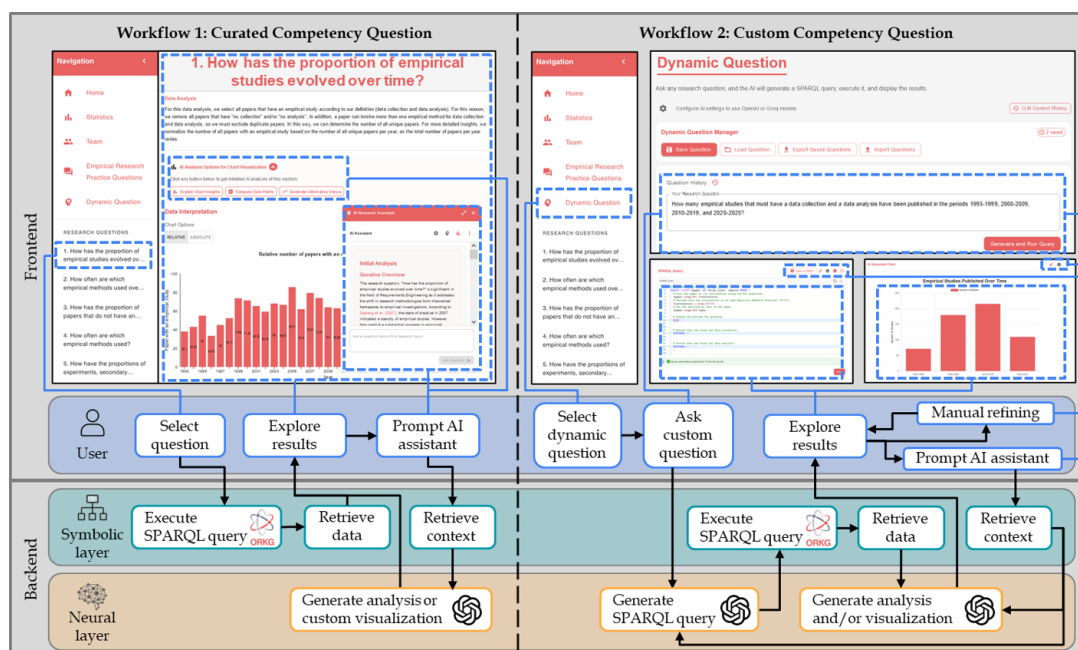


Figure 1: System design of *EmpiRE-Compass* with its two main workflows for knowledge exploration, synthesis, and reuse based on curated (*Workflow 1*) and custom (*Workflow 2*) competency questions⁵.

3.1. System Design: Requirements, Architecture, and Implementation

EmpiRE-Compass is designed to meet a set of functional and non-functional requirements derived from its intended use. Functionally, the dashboard must (FR1) support curated competency questions, (FR2) allow custom natural language competency questions, (FR3) execute SPARQL queries against the ORKG, (FR4) process and visualize retrieved data, (FR5) provide contextual interpretations and explanations, (FR6) enable manual and prompt-based refinement, (FR7) support export and import of

⁵The custom competency question used is provided online as example question “*Number of empirical studies per decade*”.

results for replication, reuse, and sharing, (FR8) persist user interactions and history across sessions, and (FR9) provide documentation and developer support. Non-functionally, the dashboard emphasizes (NFR1) usability through a responsive interface, (NFR2) transparency and reproducibility of queries and interactions, (NFR3) extensibility for new LR datasets, competency questions, and LLMs, (NFR4) interoperability with external services, (NFR5) reliability in retrieving structured LR data, and (NFR6) maintainability through modular component libraries and comprehensive API specifications. Below, we describe how the frontend and backend components implement the requirements in their architecture.

Frontend. Through a responsive web interface, researchers can explore, synthesize, and reuse LR data by selecting curated competency questions (FR1), posing custom ones (FR2), interacting with the LLM and manually refining the dynamically generated visualizations, interpretations, and explanations (FR5, FR6) (cf. Section 3.2). The frontend integrates a *workflow engine* that orchestrates the competency question-based pipeline in five steps: (1) Selecting curated competency questions and their corresponding SPARQL queries (FR1), or analyzing custom natural language competency questions and iteratively generating SPARQL queries (FR2); (2) Executing SPARQL queries against the ORKG (FR3); (3) Processing the retrieved data (FR4); (4) Generating interactive visualizations enriched with interpretations and explanations (FR4, FR5); and (5) Supporting manual and prompt-based refinement of the results (FR6). In addition, the frontend provides a *history manager* for comprehensive control over LLM memory and context, offering detailed history for selective restoration, contextual memory management to retain or discard information, and cross-session persistence to maintain progress across browser sessions (FR8). The system also supports exporting the underlying data retrieved for curated competency questions (FR7), as well as exporting and importing the complete content of custom competency questions (FR7), facilitating replication, reuse, and sharing. Together, these components ensure that users can iteratively refine queries, visualizations, interpretations, and explanations in a transparent and reproducible manner (NFR2). Built with modern web technologies, *React*, *Vite*, *TypeScript*, and *Material-UI*, and a *Storybook*-based *component library* for consistent user interface and user experience design (FR9), the frontend prioritizes usability (NFR1) and maintainability (NFR6). These qualities ensure effective interaction, reliable presentation of LR data, and adaptability for future developments.

Backend. The backend consists of the symbolic layer and the neural layer. The symbolic layer connects the dashboard to the ORKG via SPARQL queries (FR3), ensuring transparency, reproducibility, and reliability in retrieving structured LR data and contextual information (NFR2, NFR5). In particular, it implements the *competency question framework*, which comprises pairs of curated natural language competency questions investigated by an LR and the corresponding SPARQL queries to retrieve the required data (FR1). For the KG-Empire use case, the framework contains 16 pairs⁶, while the NLP4RE ID Card use case provides 10 pairs⁷. Each use case relies on its own graph schema⁸ to ensure semantic consistency (NFR2). The neural layer integrates LLMs from leading providers (*OpenAI* for ChatGPT models, *Groq* for LLaMA and DeepSeek models, *Mistral AI* for Mistral models, and *Google Generative AI* for Gemini and Gemma models). These LLMs generate SPARQL queries from custom natural language competency questions based on the respective graph schema (FR2), interpret retrieved data (FR3, FR4), and provide visualizations enriched with interpretations and explanations tailored to the users' needs (FR4, FR5, FR6). After login, users can use the system's default LLM (GPT-4o mini) or select their preferred LLM (NFR3). Together, both layers balance semantic accuracy with adaptive flexibility (NFR1, NFR2). In addition, the backend also manages the real-time ORKG access (FR3), the persistence of user interactions in a *Firebase Firestore* (FR8), automated *statistics* via Python scripts (FR9), and a comprehensive *Swagger*-based *API documentation* for an interactive, machine-readable specification of all backend endpoints (FR9, NFR6). The Swagger interface enables developers to explore request/response structures, test queries directly, and seamlessly connect to external services (NFR4). Overall, the backend prioritizes extensibility (NFR3), interoperability (NFR4), and maintainability (NFR6), enabling integration of new LR datasets, competency questions, LLMs, and external services.

⁶We systematically derived and validated these 16 competency questions from literature in our previous work [5].

⁷We collected these 10 competency questions directly from domain experts, namely the authors of the NLP4RE ID Card [16].

⁸Graph schema of the *KG-Empire* use case and the *NLP4RE ID Card* use case.

3.2. Workflows and User Interaction

As shown in Figure 1, the dashboard supports two main workflows: Answering *curated competency questions* ([Workflow 1](#)) and answering *custom competency questions* ([Workflow 2](#)).

Workflow 1: Curated Competency Question. Users select a curated competency question, which executes the corresponding SPARQL query against the ORKG. The retrieved data is processed and an interactive visualization is rendered. The visualization is enriched with pre-formulated interpretations and explanations⁹. Users can explore results further and prompt the LLM in the context of the selected competency question, i.e., to obtain additional interpretations, explanations, and visualizations. The underlying data can be exported for replication and reuse.

Workflow 2: Custom Competency Question. Users ask a natural language competency question tailored to their analysis needs. The selected LLM generates a SPARQL query consistent with the corresponding graph schema, which can be iteratively refined manually and prompt-based. The query is executed against the ORKG, and the retrieved data is processed into an interactive visualization. Interpretations and explanations are generated by the LLM to enrich the visualization¹⁰, while users can further edit outputs manually or by prompting the LLM. The complete content of the custom competency questions, including prompts, queries, visualizations, interpretations, explanations, and the underlying data, can be exported for replication, reuse, and sharing, and imported to resume analyses.

Both workflows follow a common pipeline of querying the ORKG, processing retrieved data, rendering interactive visualizations, providing interpretations and explanations, and enabling manual and prompt-based iterative refinement. Export and import functionality ensures replication, reuse, sharing, and continuity across sessions, with detailed history tracking for transparency and reproducibility.

4. Evaluation Plan

We organized the evaluation plan for *EmpiRE-Compass* into short-, mid-, and long-term actions to cover all aspects from immediate validation to large-scale evaluation. These actions are stepwise aligned with the tool’s central objectives of knowledge *exploration*, *synthesis*, and *reuse*, clarifying which activities can be conducted directly and which require further development and community involvement.

The **short-term actions** focus on content-wise *exploration* of the two use cases (KG-EmpiRE and NLP4RE ID Card). We will examine whether the curated competency questions, visualizations, interpretations, and explanations provided by the dashboard are relevant, correct, and useful for their respective fields. These evaluations will involve expert reviews with researchers in empirical RE and NLP4RE, using surveys similar to our prior work in engineering sciences [28]. For these actions, we plan to benefit from the interactive setting of the posters and tools session at REFSQ’26 and the co-located NLP4RE workshop to attract domain experts as participants.

The **mid-term actions** will conduct broader pilot studies for both use cases. Participants will be asked to perform *synthesis* tasks by exploring LR data and combining evidence for their custom competency questions. We will measure task completion time, coverage of relevant information, and user experience. For these actions, we plan to again benefit from the interactive setting at REFSQ’26 to attract participants, gather custom competency questions, and collect initial feedback from the conference attendees. These inputs will guide the design of the long-term actions and help identify participants for follow-up studies.

The **long-term actions** will address the large-scale evaluation of knowledge *reuse* for both use cases. We will systematically examine the performance of LLMs generating visualizations, interpretations, and explanations by reusing the knowledge to answer the collected custom competency questions. This evaluation requires numerous custom competency questions and answers, and thus significant community support to ensure that the assessment reflects genuine research interests rather than

⁹For curated questions, interpretations and explanations are currently written manually to guarantee correctness. In the long term, LLM-generated ones will be adopted once their accuracy and trustworthiness have been validated (cf. Section 4).

¹⁰For custom questions, interpretations and explanations are LLM-generated and remain subject to validation and refinement.

author-defined perspectives. We will combine quantitative metrics, such as completeness, accuracy, precision, and recall collected with automated approaches, with qualitative insights from domain experts gathered through focus groups and surveys.

This staged evaluation plan enables a systematic assessment of *EmpiRE-Compass*, by addressing immediate content exploration, intermediate piloting of the synthesis capabilities, and large-scale evaluation of knowledge reuse. By leveraging diverse studies and community involvement, we aim to demonstrate the technical robustness, user-friendliness, and scientific value of the *EmpiRE-Compass*.

5. Conclusion and Future Work

EmpiRE-Compass demonstrates how the integration of a neuro-symbolic approach into a dashboard can make LR data more transparent, reusable, and sustainable. By unifying the symbolic representation of semantically structured data in RKGs with the neural, generative capabilities of LLMs, researchers can explore, synthesize, and reuse knowledge without requiring advanced technical skills. The dashboard provides a modular system design with clearly defined workflows for curated and custom competency questions, ensuring that users can interact with LR data easily and dynamically while maintaining transparency and reproducibility. In addition, releasing *EmpiRE-Compass* as an *open-source project* fosters reuse, adoption, and extension by the broader community. This openness ensures that the dashboard can evolve beyond the presented use cases and be adapted to diverse research domains.

Beyond the evaluation plan, we plan two main actions for future work. First, we aim to generalize the dashboard into a technical framework that can be applied to arbitrary RKGs, enabling other communities to set up their own instances for knowledge exploration, synthesis, and reuse. Second, we will extend the synergy between LLMs and RKGs by incorporating a complementary perspective in our upcoming research project *SciD-QuEst*. So far, *EmpiRE-Compass* focuses on the use of LLMs to improve access to and reuse of RKG data. In *SciD-QuEst*, we will extend the dashboard so that users can contribute new, semantically structured data to the ORKG. Based on our two use cases, we will develop a questionnaire-based approach that empowers researchers to extract scientific knowledge from publications, with LLMs suggesting extraction candidates and humans validating them before import into the ORKG. This human-in-the-loop workflow lowers the barrier to contributing FAIR (meta-)data while ensuring quality and trustworthiness. Together, these actions will allow *EmpiRE-Compass* to fully exploit the synergistic potential of LLMs and RKGs, laying the foundation for sustainable LRs that can be collaborative, continuously updated, and reused, thereby ensuring the quality, reliability, and timeliness of research results from LRs in RE, SE, and beyond.

Declaration on Generative AI

During the preparation of this work, the author(s) used Microsoft Copilot and Grammarly in order to: Grammar and spelling check. After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] B. M. Napoleão, et al., Towards Continuous Systematic Literature Review in Software Engineering, in: 48th Euromicro Conference on Software Engineering and Advanced Applications, 2022.
- [2] S. Auer, et al., Organizing Scholarly Knowledge in the Open Research Knowledge Graph: An Open-Science Platform for FAIR Scholarly Knowledge, in: Proceedings of the Conference on Research Data Infrastructure, volume 1, 2023.
- [3] O. Karras, Research Knowledge Graphs for Sustainable Literature Reviews in Software Engineering and Beyond, in: Software Engineering 2025 – Proceedings, Gesellschaft für Informatik, 2025.
- [4] D. Castelvechi, Preprint Site arXiv is Banning Computer-Science Reviews, Nature (2025).

- [5] O. Karras, et al., Divide and Conquer the EmpiRE: A Community-Maintainable Knowledge Graph of Empirical Research in Requirements Engineering, in: 17th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement, IEEE, 2023.
- [6] A. Huotala, et al., Research Artifacts in Secondary Studies: A Systematic Mapping in Software Engineering, *Information and Software Technology* 187 (2025).
- [7] V. dos Santos, et al., Sustainable Systematic Literature Reviews, *Information and Software Technology* 176 (2024).
- [8] L. John, et al., SciMantify - A Hybrid Approach for the Evolving Semantification of Scientific Knowledge, in: *International Conference on Web Engineering*, Springer, 2025.
- [9] M. D. Wilkinson, et al., The FAIR Guiding Principles for Scientific Data Management and Stewardship, *Scientific Data* 3 (2016).
- [10] S. Auer, et al., Open Research Knowledge Graph: A Large-Scale Neuro-Symbolic Knowledge Organization System, in: *Handbook on Neurosymbolic AI and Knowledge Graphs*, 2025.
- [11] O. Karras, KG-EmpiRE: A Community-Maintainable Knowledge Graph for a Sustainable Literature Review on the State and Evolution of Empirical Research in Requirements Engineering, in: *IEEE 32nd International Requirements Engineering Conference*, 2024.
- [12] O. Karras, Divide and Conquer the EmpiRE: A Community-Maintainable Knowledge Graph of Empirical Research in Requirements Engineering, 2026. doi:10.5281/zenodo.18154586.
- [13] M. Alhamadi, et al., Data Quality, Mismatched Expectations, and Moving Requirements: The Challenges of User-Centred Dashboard Design, in: *NordiCHI Conference*, ACM, 2022.
- [14] S. Pan, et al., Unifying Large Language Models and Knowledge Graphs: A Roadmap, *IEEE Transactions on Knowledge and Data Engineering* 36 (2024).
- [15] A. Alasti, et al., From PDFs to Structured and FAIR Knowledge: Organizing NLP4RE ID Cards in the Open Research Knowledge Graph, in: *Joint Proceedings of REFSQ-2026 Workshops, Doctoral Symposium, Posters & Tools Track, and Education and Training Track*, 2026.
- [16] S. Abualhaija, et al., Replication in Requirements Engineering: The NLP for RE Case, *ACM Transactions on Software Engineering and Methodology* 33 (2024).
- [17] O. Karras, et al., Comparison of NLP4RE ID Cards, 2025. doi:10.48366/R1584378.
- [18] O. Karras, et al., EmpiRE-Compass, 2026. doi:10.5281/zenodo.18170203.
- [19] L. John, et al., SciKGDash: The Scientific Knowledge Graph Dashboard for Supporting Knowledge Curation, in: *2025 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, IEEE, 2025.
- [20] S. Scheuer, et al., Creating Accessible Evidence Bases: Opportunities through the Integration of Interactive Tools into Literature Review Synthesis, *MethodsX* 8 (2021).
- [21] A. Kaplan, et al., Structuring Scientific Knowledge in Software Engineering Using the Open Research Knowledge Graph: A Use Case and Experience Report, in: *Software Engineering 2025 – Proceedings*, Gesellschaft für Informatik, 2025.
- [22] L. Penev, et al., OpenBiodiv: A Knowledge Graph for Literature-Extracted Linked Open Data in Biodiversity Science, *Publications* 7 (2019).
- [23] G. I. Gkatzelis, et al., The Global Impacts of COVID-19 Lockdowns on Urban Air Pollution: A Critical Review and Recommendations, *Elementa: Science of the Anthropocene* 9 (2021).
- [24] G. Spadaro, et al., The Cooperation Databank: Machine-Readable Science Accelerates Research Synthesis, *Perspectives on Psychological Science* (2022).
- [25] O. Lezhnina, et al., A Scholarly Knowledge Graph-Powered Dashboard: Implementation and User Evaluation, *Frontiers in Research Metrics and Analytics* 7 (2022).
- [26] M. Bernard-Verdier, et al., Knowledge Synthesis in Invasion Biology: From a Prototype to Community-Designed Templates, *Cuvillier Verlag*, 2024.
- [27] M. Shamsabadi, et al., Large Language Models for Scientific Information Extraction: An Empirical Study for Virology, in: *Findings of the Association for Computational Linguistics: EACL*, 2024.
- [28] O. Karras, et al., Organizing Scientific Knowledge from Engineering Sciences Using the Open Research Knowledge Graph: The Tailored Forming Process Chain Use Case, *Data Science Journal* (2024).