Towards a Joint Schema for Organizing Knowledge about Empirical Research in Computer Science

Moayad Almohaishi [1], Angelika Kaplan [2], Melanie Scholz[1], Maximilian Britton [1], Agnes Koschmider [1], Anne Koziolek [2], and Oliver Karras [3]

Abstract: Empirical research in computer science (CS) often employs heterogeneous, domain-specific schemas, which limits interoperability and reuse of research knowledge. To align with the FAIR principles, namely Findable, Accessible, Interoperable, and Reusable, there is a need for a unified representation of research concepts and relationships. We propose a joint schema that captures core concepts of empirical research practice and their relationships. This schema serves as a foundation for knowledge graph population based on a bottom-up unification of diverse schemas from the research domains Software Architecture and Design, Requirements Engineering, and Business Informatics. The proposed schema was evaluated through a survey conducted within the NFDIxCS consortium, collecting feedback from 18 respondents. The results indicate that the schema is perceived as relevant and useful for harmonizing research metadata across domains. Our findings suggest that the joint schema provides a promising step toward FAIR-compliant infrastructures and cross-domain knowledge integration in CS. In future work, we aim extending the schema for additional CS disciplines and further validate it.

Keywords: Schema integration, research data, FAIR, Empirical study in CS

1 Introduction

Disciplines in Computer Science (CS) have an increasing complexity of research data. A significant challenge lies in the prevalence of heterogeneous, domain-specific schemas used to represent research data. Each CS discipline has developed its own specialized research data structures leading to schemas that are often incompatible. This fragmentation makes it difficult to integrate findings across CS disciplines and complicates research data management by limiting the structured reuse of research data knowledge. Also the heterogeneity of the domain-specific schemas for describing research data poses an obstacle to achieving the FAIR principles (Findable, Accessible, Interoperable, Reusable). To align the research data of CS disciplines with the FAIR principles requires to develop a unified

[1] Universität Bayreuth, Bayreuth, Germany,
moayad.almohaishi@uni-bayreuth.de, https://orcid.org/0009-0002-1758-3153;
melanie.scholz@uni-bayreuth.de;
maximilian.britton@uni-bayreuth.de, https://orcid.org/0009-0007-8926-7543;
agnes.koschmider@uni-bayreuth.de, https://orcid.org/0000-0001-8206-7636

[2] Karlsruher Institut für Technologie (KIT), Karlsruhe, Germany,
angelika.kaplan@kit.edu, https://orcid.org/0009-0009-9101-5833;
koziolek@kit.edu, https://orcid.org/0000-0002-1593-3394

[3] TIB - Leibniz Information Centre for Science and Technology, Hannover, Germany,
oliver.karras@tib.eu, https://orcid.org/0000-0001-5336-6899

representation of research data and its relationships capturing the essence of research data across CS disciplines.

This paper presents a joint schema for research data of the three CS disciplines Software Architecture and Design (SWA), Requirements Engineering (RE), and Business Informatics (BI), supporting interoperable descriptions of empirical studies. We identified common research data concepts and relationships between these disciplines and developed a unified minimal representation of these concepts and relationships. To develop the joint schema structure, we followed a bottom-up approach. Finally, the schema consists of elements common in all three sub-domains, elements occurring only in one of the three CS sub-domains or elements overlapping in two sub-domains (BI and RE or RE and SWA). Additionally, the schema includes a core layer, a domain-specific layer, and a metadata layer. We evaluated the proposed schema through a survey conducted within the NFDIxCS consortium, collecting feedback from 18 respondents. The results indicate that the schema is perceived as relevant and useful for harmonizing research metadata across domains. The proposed schema structure can be used as a foundation for knowledge graph (KG) population and can facilitate the reuse and integration of research results across domains.

The paper is structured as follows. Section 2 provides an overview of the proposed schema. Section 3 presents the empirical study conducted to validate the schema. We discuss related work in Section 4, and finally, in Section 5 we offer a discussion and outlook.

2 Joint Structure Schema

2.1 Previous Works

Classification and data schema in SWA research. This taxonomy offers a structured schema for classifying research artifacts in empirical software engineering (SE) and has been instantiated for SWA research [KKK+22a] across three core dimensions: (i) *Findings*, which describe the nature of the reported results and are represented through the Research Object, Evaluation Method (empirical and non-empirical), and investigated Property (categorized as Quality in Use and Product Quality according to ISO/IEC 25010:2011 [IS11], as well as Quality of Analytical Method according to Taverniers et al. [Ta04]); (ii) *Validity* of Findings, which reflects the extent to which the results are credible and generalizable, considering common validity threat categories such as internal, external, and construct validity; and (iii) *Evidence* of Findings, which characterizes the empirical foundation of the claims, including tool support, input data, and the availability of replication packages. A detailed description of the classification concepts outlined above is provided in the corresponding replication package [KKK+22b].

Classification and data schema in RE research. The schema developed supports the semantic description of publications with respect to their empirical research practices applied in RE by focusing on the *research paradigm*, *research design*, *research method*, *data collection*, *data analysis*, and *bibliographic metadata* [Ka23]. The schema is implemented in the Open Research Knowledge Graph[4] (ORKG) [Au25] as a so-called ORKG template, which defines and enforces the data structures, formats, and constraints as graph patterns. In this way, we determined which data is to be extracted and ensured that all semantic descriptions follow a uniform structure. We provide an interactive overview of the schema online[5].

Classification and data schema in BI research, inspired by [Re16], is an ongoing effort aimed at classifying literature within the BI sub-field with respect to research artifacts and the semantic aspects of the literature, such as research questions and outcomes. This domain-specific schema also incorporates empirical evidence related to research artifacts. The BI schema is designed to promote the FAIR principles within the BI sub-domain. Particularly, by leveraging the advantages of KGs. This approach enhances research transparency and facilitates the reuse of research artifacts by a broader community.

2.2 Process of Agreement

For the development of the joint schema, we followed an iterative approach with five iterations over three months. In the first iteration, the three main authors of the individual schemas develop an initial draft of the joint schema by merging all three schemas into one schema. This merge was based solely on combining concepts and relationships with completely matching names and definitions. The following two iterations served to first merge concepts and then relationships without exact matches, which required intensive discussions between the three main authors of the individual schemas, resolving conflicts by consensus to find a common denominator. In the fourth iteration, all merged concepts and relationships have been integrated into the initial draft, and the three main authors of the individual schemas discussed the intermediate schema again to revise and improve it. The resulting joint schema was presented to the remaining authors of this paper in the fifth iteration. After an intensive discussion, all authors of this paper agreed on the final version of the joint schema presented in Fig. 1.

2.3 Joint Schema to Describe Empirical Research Practice in Computer Science

To characterize empirical research practices in CS, we derive eight high-level/first-level concepts based on an analysis of the individual concepts identified across the three CS sub-disciplines, as introduced in subsection 2.1 and discussed in subsection 2.2. The resulting joint schema is shown in Figure 1. Each main concept in this schema can be further refined through facets, providing a more nuanced and detailed description. The following are the core concepts:

[4] https://orkg.org/ [Last accessed on 2026-03-03]
[5] https://empire-compass.tib.eu/R186491/schema [Last accessed on 2026-03-03]

Research Problem: represents a statement that motivates the research work in a publication (e.g. research gaps). *Research Question*: is a clearly formulated question that guides the study and defines what the researchers aim to answer. *Hypothesis*: is a precise, testable statement of what the researcher predicts will be the outcome. It may represent an alternative or complementary approach to a research question. *Data Collection*: is a research process of collecting data that can be used to gain insights and information regarding the research topic. *Data Analysis*: refers to the process of systematically applying statistical, computational, or logical techniques to describe, summarize, and interpret data. *Result*: refers to the research outcome or findings derived from the data analysis process or the whole research. *Threats to Validity*: is the set of reported limitations affecting the validity, soundness, results, and outcome of a research. In addition to these concepts, we introduce one central concept that is essential for the alignment with empirical research practices: the *Research Object*.

The research object of a paper refers to the central object of investigation by the researchers and is typically explicitly stated (e.g.,"In our paper, we present a design method ..."). This concept represents the primary deviation from the core empirical method design and serves as a key knowledge-based descriptive metadata element, enabling domain-specific statements to be discoverable. Across the three considered sub-domains of CS, we found no established de facto standard for this concept, except for the community-driven initiative in SWA research [KKK+22a; KKK+22b].

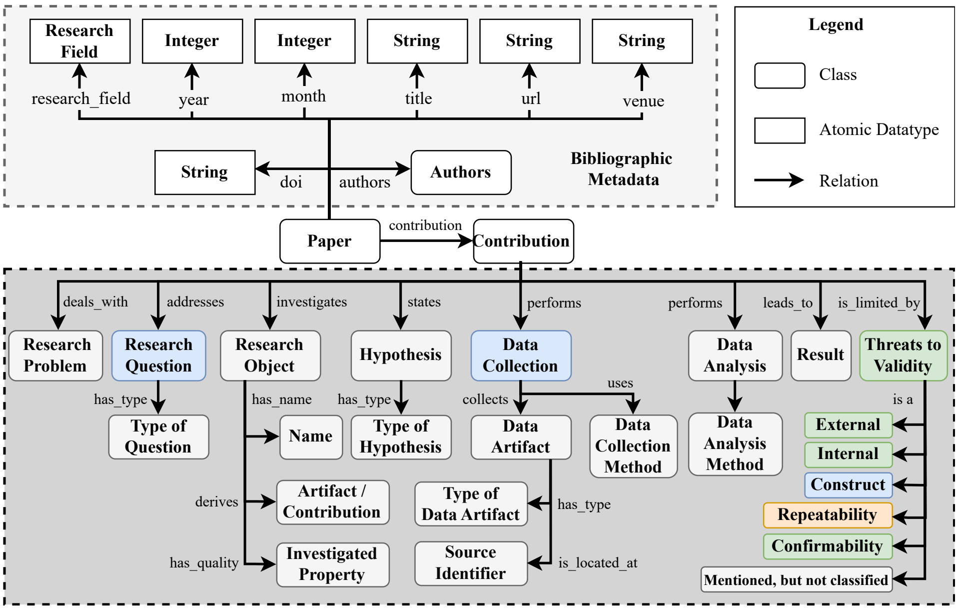

Fig. 1: Overview of the joint schema describing empirical research practices in CS. Gray boxes show elements unique to one CS sub-domain. Green boxes denote shared elements across all three. Blue boxes indicate overlap between BI and RE, and orange boxes between RE and SWA.

3 Empirical Study

To validate the proposed joint structure schema, we sought feedback from community experts through a carefully designed, web-based questionnaire. The survey was distributed to the NFDIxCS consortium[6] between August and September, 2025. Participation in the survey was voluntary and anonymous, and participants had the freedom to abandon the survey at any time without having their response recorded.

3.1 Descriptive Information

A challenge in designing the questionnaire is the diversity in the sub-domains and its research data artifacts or concepts respectively. Artifacts may use different terminology across sub-domains; for example, in SWA, the output of research is referred to as "research contribution"whereas in BI, it is called "research artifact". To address these heterogeneity, we established the following guidelines: i) inclusion of a clear explanation of our motivation, ii) inclusion of each concept, iii) clear definitions for each concept, iv) a visualization of the schema diagram, and v) a completion time of no more than 15 minutes. The resulting questionnaire is designed as follows.

Structure: The survey consists of two main sections: demographic questions and evaluation of the schema's concepts. The demographic section collects participants' background information, such as sub-domain according to the DFG subject area classification [7], career level, years of experience, and the number of empirical studies conducted in the CS field. The evaluation section begins with a presentation of the schema and its purpose, followed by a brief explanation of each concept and a Likert-scale feedback questionnaire. To reduce the time, the main statement *'The concept [...] is useful for describing empirical research practices in computer science'* was presented once in the task description. Then each concept is presented with a concept, a short definition, an example when necessary, and a multiple choice input as Likert scale feedback. Before submitting the feedback, the participants are invited to indicate their willingness to participate in an interview. If they agree, their name and email are requested for follow-up contact.

Evaluation of feasibility: a pilot version of the questionnaire was distributed to five PhD students to obtain informal feedback regarding clarity and the average completion time.

3.2 Results

In response, 18 experts from various CS sub-domains provided feedback through the questionnaire, 12 of whom hold a PhD or are professional researchers. The majority (15 out

6 https://nfdixcs.org [Last accessed on 2026-03-03]
7 https://www.dfg.de/resource/blob/331950/85717c3edb9ea8bd453d5110849865d3/fachsystematik-2024-2028-en-data.pdf [Last accessed on 2026-03-03]

of 18) have conducted empirical studies in CS. The feedback, summarized in Fig. 2, reveals strong agreement among respondents on almost all concepts of the proposed minimal joint structure schema. Notably, the concept *'Type of Hypothesis'* fell below the 50% acceptance threshold, despite a disagreement from only one participant, due to 61% of participants being *'Neutral'*. In contrast, the related concept *'Hypothesis'* received 55.6% agreement, suggesting that specifying the type of hypothesis may be relevant for fewer sub-domains. Generally, with the exception of one participant who strongly disagreed with the *'Source Identifier'* concept, there was considerable consensus on all other aspects of the proposed structure.

Interpretation of results Overall, the survey results indicate support for most concepts in our joint schema. Notably, the respondents come from several CS sub-disciplines beyond the three domains used to derive the schema (SWA, RE, BI), which suggests that the identified concepts capture a common understanding of empirical research practice in CS and are, to some extent, generalizable across subcommunities; however, recruiting participants via NFDIxCS could introduce selection bias. One exception is *Type of Hypothesis*, which received comparatively low acceptance and many neutral responses. We interpret this as an indication that this level of granularity is not equally relevant across CS subdisciplines. Consequently, we will consider moving *Type of Hypothesis* from the joint core to subdiscipline-specific refinements.

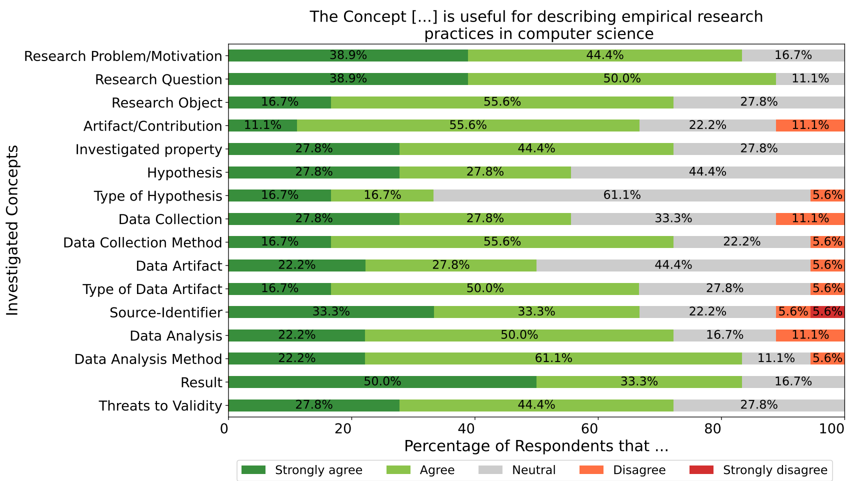

Fig. 2: Feedback from the minimal joint structure schema survey

4 Related Work

In this section, we discuss how our work relates to prior research. Ramesh et al. provide an overview of CS research by analyzing 628 papers from 13 leading journals and characterizing topics, research approaches, methods and levels of analysis [Ra04]. Their study offers a broad picture of how research in CS is conducted, but it does not focus specifically on empirical research or on a reusable schema for the core concepts. Galster and Weyns map empirical research in SWA by analyzing which empirical methods are used, how human participants are involved and how validity is reported [Ga16]. In a follow-up, Galster and Weyns position their survey as a complement to such state of practice studies by examining how SWA researchers perceive current empirical research [Ga23]. Taken together, these works provide an overview of empirical research in one of our research domains. Vassiliou et al. introduce a hybrid Dublin Core–DDI metadata framework for documenting empirical research in student theses, with the aim of improving their discoverability, classification and reuse in institutional repositories [Va25]. Their work demonstrates how information about empirical research can be represented in a dedicated metadata framework, however it is tailored to theses in the social sciences. Another line of research proposes infrastructures and frameworks for publishing and using structured metadata about research outputs. Kelley and Garijo introduce a framework that extracts scientific software metadata from documentation such as README files and organizes it into a KG [Ke21]. Da Silva Santos et al. present the FAIR Data Point as a SWA for publishing semantically rich, machine-actionable metadata in accordance with the FAIR principles and discuss how metadata schemas can be defined within this infrastructure, which is relevant for FAIR oriented handling of research data [Si23]. While existing literature analyzes research practices within specific sub-domains or proposes general metadata frameworks, there remains a lack of a unified, cross domain schema specifically to describe empirical research data across CS disciplines.

5 Discussion and Outlook

We presented a joint schema to model empirical CS research contributions with the purpose of annotating research papers in a KG (specifically the ORKG). This allows more systematic comparison of contributions and facilitates finding relevant literature beyond keyword- or citation-based search. In the age of Large Language Models (LLMs), literature research is increasingly supported by generative AI assistants. We argue that schema-based KG annotations can provide a more reliable basis for comparisons and related work discovery than pure LLM-based search. In particular, AI-supported literature search that leverages structured annotations can be seen as an instance of intersymbolic (or neuro-symbolic) AI [Pl24], combining LLM capabilities with explicit symbolic representations.

The scope of our schema is intentionally limited to concepts shared across multiple CS subdisciplines. For more detailed comparisons within a single subdiscipline (or for specific research problems), specialized schemata can be derived from the joint schema. As future work, we will explore how to best establish and evolve such a hierarchy of schemata from the joint schema down to subdiscipline-specific refinements, extend alignment to additional CS subdisciplines, and populate the schema with more papers to study its application in practice (e.g., systematic literature studies created by humans and/or LLM-supported workflows).

Acknowledgments

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under the National Research Data Infrastructure – NFDIxCS 52/1 – 501930651 and NFDI4Ing - 442146713. All authors contributed equally to this work.

References

[Au25] Auer, S. et al.: Open Research Knowledge Graph: A Large-Scale Neuro-Symbolic Knowledge Organization System. In: Handbook on Neurosymbolic AI and Knowledge Graphs. IOS Press, 2025.

[Ga16] Galster, M. et al.: Empirical Research in Software Architecture: How Far have We Come? In: 13th Working IEEE/IFIP Conference on Software Architecture (WICSA). 2016.

[Ga23] Galster, M. et al.: Empirical Research in Software Architecture — Perceptions of the Community. Journal of Systems and Software 202, 2023.

[IS11] ISO/IEC: Systems and software engineering - Systems and software Quality Requirements and Evaluation (SQuaRE) - System and software quality models, tech. rep., International Organization for Standardization, 2011.

[Ka23] Karras, O. 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 (ESEM). IEEE, 2023.

[Ke21] Kelley, A. et al.: A Framework for Creating Knowledge Graphs of Scientific Software Metadata. Quantitative Science Studies 2 (4), 2021.

[KKK+22a] Konersmann, M.; Kaplan, A.; Kühn, T., et al.: Evaluation Methods and Replicability of Software Architecture Research Objects. In: 19th IEEE International Conference on Software Architecture, ICSA 2022. IEEE, 2022.

[KKK+22b] Konersmann, M.; Kaplan, A.; Kühn, T., et al.: Replication Package of "Evaluation Methods and Replicability of Software Architecture Research Objects". In: 19th IEEE International Conference on Software Architecture. IEEE, 2022.

[Pl24] Platzer, A.: Intersymbolic AI: Interlinking Symbolic AI and Subsymbolic AI. In: International Symposium on Leveraging Applications of Formal Methods. Springer, 2024.

[Ra04] Ramesh, V. et al.: Research in Computer Science: An Empirical Study. Journal of Systems and Software 70 (1), 2004.

[Re16] Recker, J. et al.: The State of the Art of Business Process Management Research as Published in the BPM Conference. Bus. Inf. Syst. Eng. 58, 2016.

[Si23] da Silva Santos, L. O. B. et al.: FAIR Data Point: A FAIR-Oriented Approach for Metadata Publication. Data Intelligence 5 (1), 2023.

[Ta04] Taverniers, I. et al.: Trends in Quality in the Analytical Laboratory. II. Analytical Method Validation and Quality Assurance. TrAC Trends in Analytical Chemistry 23 (8), 2004.

[Va25] Vassiliou, G. et al.: Enhancing Discoverability: A Metadata Framework for Empirical Research in Theses. Algorithms 18 (8), 2025.