



Requirements Analysis for an Open Research Knowledge Graph

Arthur Brack¹(✉) , Anett Hoppe¹(✉) , Markus Stocker¹(✉) ,
Sören Auer^{1,2}(✉) , and Ralph Ewerth^{1,2}(✉) 

¹ TIB – Leibniz Information Centre for Science and Technology, Hannover, Germany
{arthur.brack,anett.hoppe,markus.stocker,auer,ralph.ewerth}@tib.eu

² L3S Research Center, Leibniz University, Hannover, Germany

Abstract. Current science communication has a number of drawbacks and bottlenecks which have been subject of discussion lately: Among others, the rising number of published articles makes it nearly impossible to get a full overview of the state of the art in a certain field, or reproducibility is hampered by fixed-length, document-based publications which normally cannot cover all details of a research work. Recently, several initiatives have proposed knowledge graphs (KGs) for organising scientific information as a solution to many of the current issues. The focus of these proposals is, however, usually restricted to very specific use cases. In this paper, we aim to transcend this limited perspective by presenting a comprehensive analysis of requirements for an Open Research Knowledge Graph (ORKG) by (a) collecting daily core tasks of a scientist, (b) establishing their consequential requirements for a KG-based system, (c) identifying overlaps and specificities, and their coverage in current solutions. As a result, we map necessary and desirable requirements for successful KG-based science communication, derive implications and outline possible solutions.

Keywords: Scholarly communication · Research Knowledge Graph · Design science research · Requirements analysis

1 Introduction

Today's scholarly communication is a document-centred process and as such, rather inefficient. Scientists spend considerable time in finding, reading and reproducing research results from PDF files consisting of static text, tables, and figures. The explosion in the number of published articles [12] aggravates this situation further: It gets harder and harder to stay on top of current research, that is to find relevant works, compare and reproduce them and, later on, to make one's own contribution known for its quality.

Some of the available infrastructures in the research ecosystem already use *knowledge graphs* (KG)¹ to enhance their services. Academic search engines, for instance, such as *Microsoft Academic Knowledge Graph* [24] or *Literature Graph* [3] employ metadata-based graph structures which link research articles based on citations, shared authors, venues and keywords.

Recently, initiatives have promoted the usage of KGs in science communication, but on a deeper, semantic level [4, 32, 37, 48, 51, 54]. They envision the transformation of the dominant document-centred knowledge exchange to knowledge-based information flows by representing and expressing knowledge through semantically rich, interlinked KGs. Indeed, they argue that a shared structured representation of scientific knowledge has the potential to alleviate some of the science communication’s current issues: Relevant research could be easier to find, comparison tables automatically compiled, own insights rapidly placed in the current ecosystem. Such a powerful data structure could, more than the current document-based system, also encourage the interconnection of research artefacts such as datasets and source code much more than current approaches (like DOI references etc.); allowing for easier reproducibility and comparison. To come closer to the vision of knowledge-based information flows, research articles should be enriched and interconnected through machine-interpretable semantic content. Jaradeh et al.’s study [37] indicates that authors are also willing to contribute structured descriptions of their research articles.

The work of a researcher is manifold, but current proposals usually focus on a specific use case (e.g. the above-named examples focus on enhancing academic search). In this paper, we provide a detailed analysis of common work tasks in a scientist’s daily life and analyse (a) how they could be supported by an ORKG, (b) what requirements result for the design of (b1) the KG and (b2) the surrounding system, (c) how different use cases overlap in their requirements and can benefit from each other. Our analysis is led by the following research questions:

1. What functionalities should be provided by ORKG interfaces?
 - (a) Which user interfaces are necessary?
 - (b) Which machine interfaces are necessary?
2. What requirements can be defined for the underlying ontologies?
 - (a) Which granularity of information representation is needed?
 - (b) To what degree is domain specialisation needed?
3. What requirements can be defined for the instance data?
 - (a) Which approaches (human vs. machine) are suitable to populate the KG?
 - (b) Which coverage of research artefacts is necessary for the instance data?
 - (c) Which quality is necessary for the instance data?

¹ Acknowledging that knowledge graph is vaguely defined, we adopt the following definition: A *knowledge graph* (KG) consists of (1) an *ontology* describing a conceptual model, and (2) the corresponding *instance data* following the constraints posed by the ontology. The construction of a KG involves *ontology design* and *population* with instances.

We follow the design science research (DSR) methodology [33]. In this study, we focus on the first phase of DSR conducting a requirements analysis. The objective is to chart necessary (and desirable) requirements for successful KG-based science communication, and, consequently, provide a map for future research.

The remainder of the paper is organised as follows. Section 2 summarises related work on research knowledge graphs, scientific ontologies and methods for KG construction. The requirements analysis is presented in Sect. 3, while Sect. 4 discusses implications and possible approaches for ORKG construction. Finally, Sect. 5 concludes the requirements analysis and outlines areas of future work.

2 Related Work

This section provides a brief overview of (a) existing research KGs, (b) ontologies representing scholarly knowledge, and (c) approaches for KG construction.

2.1 Research Knowledge Graphs

Academic search engines (e.g. Google Scholar, Microsoft Academic, Semantic Scholar) exploit graph structures such as the Microsoft Academic Knowledge Graph [24], SciGraph [68], or the Literature Graph [3]. These graphs interlink research articles through metadata, e.g. citations, authors, affiliations, grants, journals, or keywords.

To help reproducing research results, initiatives such as Research Graph [2], Research Objects [7] and OpenAIRE [48] interlink research articles with research artefacts such as datasets, source code, software, and presentation videos. Scholarly Link Exchange (Scholix) [16] aims to create a standardised ecosystem to collect and exchange links between research artefacts and literature.

Some approaches were proposed to interlink articles at a more semantic level: Paperswithcode.com is a community-driven effort to link machine learning articles with tasks, source code and evaluation results to construct leaderboards. Ammar et al. [3] interlink entity mentions in abstracts with DBpedia [43] and Unified Medical Language System (UMLS) [10], and Cohan et al. [17] extend the citation graph with semantic citation intents (e.g. cites as background or as used method).

Various scholarly applications benefit from semantic content representation, e.g. academic search engines by exploiting general-purpose KGs [67], and graph-based research paper recommendation systems [8] by utilising citation graphs and mentioned genes. However, the coverage of science-specific concepts in general-purpose KGs is rather low [3], e.g. the task “geolocation estimation of photos” from Computer Vision is neither present in Wikipedia nor in CSO (Computer Science Ontology) [59].

2.2 Scientific Ontologies

Various ontologies have been proposed to model metadata such as bibliographic resources and citations [53]. Iniesta and Corcho [58] reviewed ontologies to

describe scholarly articles. In the following, we describe some ontologies that conceptualise the semantic content in research articles.

Several ontologies focus on rhetorical [19, 30, 66] (e.g. Background, Methods, Results, Conclusion), argumentative [45, 63] (e.g. claims, contrastive and comparative statements about other work) or activity-based [54] (e.g. sequence of research activities) aspects and elements of research articles. Others describe scholarly knowledge with interlinked entities such as problem, method, theory, statement [15, 32], or focus on the main research findings and characteristics of research articles described in surveys with concepts such as problems, approaches, implementations, and evaluations [25, 64].

There are various domain-specific ontologies, for instance, mathematics [42] (e.g. definitions, assertions, proofs) and machine learning [40, 49] (e.g. dataset, metric, model, experiment). The EXPeriments Ontology (EXPO) is a core ontology for scientific experiments conceptualising experimental design, methodology, and results [61].

Taxonomies for domain-specific research areas support the characterisation and exploration of a research field. Salatino et al. [59] provide an overview, e.g. Medical Subject Heading (MeSH), Physics Subject Headings (PhySH), Computer Science Ontology (CSO). Gene Ontology [1] and Chemical Entities of Biological Interest (CheBi) [21] are KGs for genes and molecular entities.

2.3 Construction of Knowledge Graphs

Automatic Construction from Text: Petasis et al. [55] provide a review on *ontology learning*, that is ontology creation from text, while Lubani et al. [47] review *ontology population systems*. Pajura and Singh [56] provide an overview of the involved tasks for *KG population*: (a) *knowledge extraction* to extract a graph from text with *entity extraction* and *relation extraction*, and (b) *graph construction* to clean and complete the extracted graph, as it is usually ambiguous, incomplete and inconsistent. *Coreference resolution* [46] clusters different mentions of the same entity and *entity linking* [41] maps them to entities in the KG. For *taxonomy population* Salatino et al. [59] provide an overview of methods based on rule-based natural language processing (NLP), clustering and statistical methods. In particular, the Computer Science Ontology (CSO) has been populated automatically from research articles [59].

Information Extraction from Scientific Text: Nasar et al. [50] provide a survey about scientific information extraction. Beltagy et al. [9] present benchmarks for several datasets.

There are datasets which are annotated at *sentence level* for several domains, e.g. biomedical [22, 38], computer graphics [28], computer science [18], chemistry and computational linguistics [63]. They focus either on the rhetorical structure in abstracts [18, 22, 38] or full articles [28, 45], or on the argumentative structure of full articles [63]. The datasets differentiate between five and twelve concept classes (e.g. Background, Objective, Results). On abstracts and

full articles machine learning approaches achieve an F1 score of 83–92% [18] or 51–80% [28, 44], respectively.

More recent corpora, annotated at *phrasal level*, aim at constructing a fine-grained KG from scholarly abstracts with the tasks of concept extraction [5, 13, 31, 46], relation extraction [5, 29, 46], and coreference resolution [46]. They cover several domains, e.g. computational linguistics [29, 31]; computer science, material sciences, and physics [5]; machine learning [46]; or a set of ten scientific, technical and medical domains [13]. The datasets differentiate between four to seven concept classes (like Task, Method, Tool) and between two to seven relation types (like used-for, part-of, evaluate-for). Concept extraction, coreference resolution and relation extraction achieve an F1 score of 45–89% [5, 9, 13], 48% [46] and 28–50% [5, 29, 46], respectively, and the inter-coder agreement is 60–76% [5, 13, 46], 68% [46] and 60%–90% [5, 29, 31, 46], respectively. *This indicates, that these tasks are not only difficult for machines but also for humans.*

Manual Curation: WikiData [65] is one of the most popular KGs with semantically structured, encyclopaedic knowledge curated manually by a community. As of March 2020, WikiData comprises 80M entities curated by almost 25.000 active contributors. The community also maintains a taxonomy of categories and “infoboxes” which define common properties of certain entity types. Paperswithcode.com is a further community-driven effort to interlink machine learning articles with tasks, source code and evaluation results. KGs such as Gene Ontology [1] or Wordnet [26] are curated by domain experts. Research article submission portals such as easychair.org enforce the submitter to provide machine-readable metadata. Librarians and publishers tag new articles with keywords and subjects [68]. Virtual research environments enable the execution of data analysis on interoperable infrastructure and store the data and results in KGs [62].

3 Requirements Analysis

As the discussion of related work reveals, existing research KGs focus on specific use cases (e.g. improve search engines, help to reproduce research results) and mainly manage metadata and research artefacts about articles. We envision a KG in which research articles are interlinked through a deep semantic representation of their content to enable further use cases. In the following, we formulate the problem statement and describe our research method. This motivates our use case analysis in Sect. 3.1, from which we derive requirements for an ORKG.

Problem Statement: Scholarly knowledge is very heterogeneous and diverse. Therefore, an ontology that conceptualises scholarly knowledge comprehensively does not (and unlikely will) exist. Besides, due to the complexity of the task, the population of comprehensive ontologies requires domain and ontology experts. Current automatic approaches can only populate rather simple ontologies and achieve moderate accuracy (see Sect. 2.3). *On the one hand, we desire an ontology that can comprehensively capture scholarly knowledge and instance data with high quality and coverage. On the other hand, we are faced with a “knowledge acquisition bottleneck”.*

Research Method: To illuminate the above problem statement we perform a *requirements analysis*. We follow the *design science research (DSR)* methodology [14,35]. The requirements analysis is a central phase in DSR, as it is the basis for design decisions and selection of methods to construct effective solutions systematically [14]. DSR’s objective in general is the innovative, rigorous and relevant design of information systems for solving important business problems or the improvement of existing solutions [14,33]. To elicit requirements, we studied guidelines for systematic literature reviews [27,39,52] and interviewed members of the ORKG team at TIB (<https://projects.tib.eu/orkg/project/team/>), who are software engineers and researchers in the field of computer science and environmental sciences. Based on the requirements, we elaborate possible approaches to construct an ORKG, which were identified through a literature review (see Sect. 2.3). To verify our assumptions on the presented requirements and approaches, ORKG team members reviewed them.

3.1 Overview of the Use Cases

We define functional requirements with use cases [11]. A use case describes the interaction between a user and the system from the user’s perspective to achieve a certain goal. As a motivating scenario it also guides the design of a supporting ontology [20].

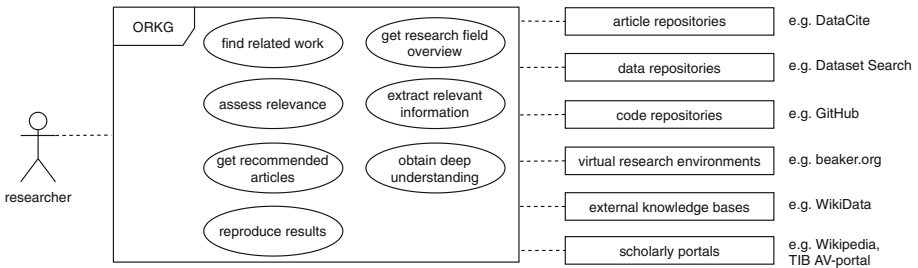


Fig. 1. UML use case diagram for the main use cases between the actor researcher, an Open Research Knowledge Graph (ORKG), and external systems.

There are many use cases (e.g. literature reviews, plagiarism detection, peer reviewer suggestion) and several stakeholders (e.g. researchers, librarians, peer reviewer, practitioners) that may benefit from an ORKG. In this study, we focus on use cases that support *researchers* (a) conducting literature reviews, (b) obtaining a deep understanding of a research article and (c) reproducing research results. A full discussion of all possible use cases of graph-based knowledge management systems in the research environment is far beyond the scope of this article. With the chosen focus, we hope to cover the most frequent, literature-oriented tasks of scientists. Figure 1 depicts the main identified use cases, which are described briefly in the following. Please note that we focus on how *semantic content* can improve these use cases and not further metadata.

Get Research Field Overview: Survey articles provide an overview of a particular research field, e.g. a certain research problem or a family of approaches. The results in such surveys are sometimes summarised in structured and comparative tables (an approach usually followed in domains such as computer science, but not as systematically practised in other fields). However, once survey articles are published they are no longer updated. Moreover, they usually represent only the perspective of the authors, i.e. very few researchers in the field. To support researchers to obtain an up-to-date overview of a research field, the system should maintain such surveys in a structured way, and allow for dynamics and evolution. A researcher interested in such an overview should be able to search or to browse the desired research field. Then, the system should provide related articles and available overviews, e.g. in a table or a leaderboard chart. While the user interface shows tabular, leaderboards, or other visual representations the backend should semantically represent information to allow for exploiting overlaps in conceptualisations between research problems or fields.

Find Related Work: Finding relevant research articles is a daily core activity of researchers. It should be possible to pose queries for related work, which can be fine-grained or broad search intents. Systems should preferably support natural language queries as approached by semantic search and question answering engines [6]. The system has to return a set of relevant articles.

Assess Relevance: Given a set of relevant articles the researcher has to assess whether the articles match the criteria of interest. Usually researchers skim through the title and abstract. Sometimes, the introduction and conclusions have to be considered. However, this is usually cumbersome and time-consuming. Presenting the researcher only the most important zones in the article in a structured way can boost this process. This includes, for instance, text passages that describe the problem tackled in the research work, the employed methods or materials, or the yielded results. Also, faceted drill-down methods based on the properties of semantic descriptions of research approaches will empower researchers to quickly filter and zoom into the most relevant literature.

Extract Relevant Information: To tackle a particular research question, the researcher has to extract relevant information from research articles. Such information is usually compiled in written text or comparison tables in a related work section or survey articles. For instance, for the question *Which datasets exist for scientific sentence classification?* a researcher who focuses on a new annotation study could be interested in (a) domains covered by the dataset and (b) the inter-coder agreement. Another researcher might follow the same question but with a focus on machine learning could be interested in (c) evaluation results and (d) feature types used. The system should support the researcher with tailored information extraction from a set of research articles: (1) the researcher defines a data extraction form as proposed in systematic literature reviews [39] (e.g. the above fields (a)–(d)) and (2) the system presents the extracted information for the corresponding data extraction form and articles in a table.

Get Recommended Articles: When the researcher focuses on a particular article, further related articles should be recommended by the system, for instance, articles that address the same research problem or apply similar methods.

Obtain Deep Understanding: The system should help the researcher to obtain a deep understanding of a research article (e.g. equations, algorithms, diagrams, datasets). For this purpose, the system should interlink the article with artefacts such as conference videos, presentations, source code, datasets, etc., and visualise the artefacts appropriately. Also text passages can be interlinked, e.g. method explanations in Wikipedia, source code snippets implementing algorithms or equations described in the article.

Reproduce Results: The system should provide the researcher links to all necessary artefacts to reproduce research results, e.g. datasets, source code, virtual research environments, materials describing the study, etc. Further, the system should maintain semantic descriptions of domain-specific and standardised evaluation protocols and guidelines.

3.2 Knowledge Graph Requirements

The non-functional requirements for the respective use cases are discussed in the light of the following dimensions.

1. *Domain specialisation of the ontology:* How domain-specific should the concepts be in the ontology? Various ontologies (e.g. [13, 54]) propose domain independent concepts (e.g. Process, Method, Material). In contrast, Klampanos et al. [40] present a very domain-specific ontology for artificial neural networks.
2. *Granularity of the ontology:* Which granularity is required to conceptualise scholarly knowledge? For instance, the annotation schemes for scientific corpora (see Sect. 2.3) have a rather low granularity, as they do not have more than 10 classes and 10 relation types. In contrast, various ontologies (e.g. [32, 54]) with more than 20–35 classes and over 20–70 relations and properties are fine-grained and have a relatively high granularity.
3. *Coverage of the instance data:* Given an ontology, to which extent do *all* possible instances in *all* research articles have to be represented in the KG? For instance, given an ontology with a class “Task”, the instance data for that ontology would have a high coverage if all tasks mentioned in all research articles are present.
4. *Quality of the instance data:* Given an ontology, which quality is necessary for the corresponding instances? In a KG with high quality all present instances must conform to the ontology and reflect the content of the research articles properly, e.g. an article is correctly assigned to the task addressed in the article, the F1 score in the evaluation results is correctly extracted, etc.

Next, we discuss the seven main use cases with regard to the required level of ontology domain specialisation and granularity, as well as coverage and quality

Table 1. Requirements and approaches for the main use cases. The upper part describes the minimum requirements for the ontology (domain specialisation and granularity) and the instance data (coverage and quality). The bottom part provides possible approaches for manual, automatic and semi-automatic curation of the KG for the respective use cases. “X” indicates that the approach is suitable for the use case while “(x)” means that the approach is only appropriate with human supervision. The left part (delimited by the vertical triple line) groups use cases suitable for manual, and the right side for automatic approaches. Vertical double lines group use cases with similar requirements.

		Extract relevant info	Research field overview	Deep understanding	Reproduce results	Find related work	Recommend articles	Assess relevance
Ontology	Domain specialisation	high	high	med	med	low	low	med
	Granularity	high	high	med	med	low	low	low
Instance data	Coverage	low	low	low	med	high	high	med
	Quality	high	high	high	high	low	low	med
Manual curation	Maintain terminologies	-	X	-	-	X	X	-
	Define templates	X	X	-	-	-	-	-
	Fill in templates	X	X	X	X	-	-	-
	Maintain overviews	X	X	-	-	-	-	-
Automatic curation	Entity/relation extraction	(x)	(x)	(x)	(x)	X	X	X
	Entity linking	(x)	(x)	(x)	(x)	X	X	X
	Sentence classification	(x)	-	(x)	-	-	-	X
	Template-based extraction	(x)	(x)	(x)	(x)	-	-	-
	Cross-modal linking	-	-	(x)	(x)	-	-	-

of instance data. Table 1 summarises the requirements for the use cases along the four dimensions at ordinal scale. The use cases are grouped together, when they have (1) similar justifications for the requirements, and (2) a high overlap in ontology concepts and instances.

Extract Relevant Information and Get Research Field Overview: The information to be extracted from relevant research articles for a data extraction form is very heterogeneous and depends highly on the intent of the researcher and the research questions. Thus, the ontology has to be domain-specific and fine-grained to offer all possible kinds of desirable information. In addition, the provided information has to be of high quality, e.g. a provided F1 score of an evaluation result must not be wrong. However, missing information for certain questions in the KG may be tolerable for a researcher.

Obtain Deep Understanding and Reproduce Results: The provided information for these use cases has to be of high quality (e.g. accurate links to dataset, source code, videos, articles, research infrastructures). The ontology for representing default artefacts can be rather domain-independent (e.g. Scholix [16]). However, semantic representation of evaluation protocols require domain-dependent ontologies (e.g. EXPO [61]). Missing information is tolerable for these use cases.

Find Related Work and Get Recommended Articles: When searching for related work, it is essential not to miss relevant articles. Previous studies revealed that more than half of search queries in academic search engines refer to scientific entities [67] and the coverage of scientific entities in KGs is rather low [3]. Despite the low coverage, Xiong et al. [67] could improve the ranking of search results by exploiting KGs. Hence, the instance data for the “find related work” use case should have high coverage with fine-grained scientific entities. However, semantic search engines employ latent representations of KGs and text (e.g. graph and word embeddings) [6]. Since a non-perfect ranking of the search results is tolerable for a researcher, lower quality of the instance data is acceptable. Furthermore, due to latent feature representations, the ontology can be kept rather simple and domain-independent. For instance, the STM corpus [13] proposes four domain-independent concepts. Graph- and content-based research paper recommendation systems [8] have similar requirements since they also leverage latent feature representations, require fine-grained scientific entities, and non-perfect recommendations are tolerable.

Assess Relevance: To help the researcher to assess the relevance of an article according to her needs, the system should highlight the most essential zones in the article to get a quick overview. The coverage and quality of the presented information must not be too low, as otherwise the user acceptance may suffer. However, it can be suboptimal, since it is acceptable for a researcher when some of the highlighted information is not essential or when some important information is missing. The ontology to represent essential information should be rather domain-specific and quite simple (cf. ontologies for scientific sentence classification in Sect. 2.3).

4 Implications for ORKG Construction

In this section, we discuss the implications for the design and construction of an ORKG and outline possible approaches, which are mapped to the use cases in Table 1. Based on the discussion in the previous section, we can subdivide the use cases into two groups: (1) requiring high quality and high domain specialisation with only low requirements on the coverage (left side in Table 1), and (2) requiring high coverage with rather low requirements on the quality and domain specialisation (right side in Table 1). The first group requires manual approaches while the second group could be accomplished with fully automatic approaches. However, manually curated data can also support use cases with

automatic approaches, and vice versa. Besides, automatic approaches can complement manual approaches by providing suggestions in user interfaces.




Fig. 2. Conceptual meta-model in UML for templates and interface design for an external template-based information extractor.

4.1 Manual Approaches

Ontology Design: The first group of use cases requires rather domain-specific and fine-grained ontologies. We suggest to develop novel or reuse ontologies that fit the respective use case and the specific domain (e.g. EXPO [61] for experiments). Moreover, appropriate and simple user interfaces are necessary for efficient and easy population.

However, such ontologies can evolve with the help of the community, as demonstrated by WikiData and Wikipedia with “infoboxes” (see Sect. 2.3). Therefore, the system should enable the maintenance of *templates*, which are pre-defined and very specific forms consisting of fields with certain types (see Fig. 2). For instance, to automatically generate leaderboards for machine learning tasks a template would have the fields Task, Model, Dataset and Score, which can then be filled in by a curator for articles providing such kind of results in a user interface generated from the template. Such an approach is also called *meta-modelling* [11], as the meta-model for templates enables the definition of concrete templates, which are then instantiated for articles.

Knowledge Graph Population: Several user interfaces are required to enable manual population: (1) populate semantic content for a research article by (1a) choosing relevant templates or ontologies and (1b) fill in the values; (2) terminology management (e.g. domain-specific research fields); (3) maintain research field overviews by (3a) assigning relevant research articles to the research field, (3b) define corresponding templates and (3c) fill in the templates for the relevant research articles.

Further, the system should also provide *APIs* to enable population by third-party applications, e.g. (i) submission portals such as easychair.org during submission of an article; (ii) authoring tools such as overleaf.com during writing; (iii) virtual research environments [62] to store evaluation results and links to datasets and source code during experimenting and data analysis.

To *encourage crowd-sourced content*, we see the following options: (a) *top-down enforcement* via submission portals and publishers; (b) *incentive models*: Researchers want their articles to be cited; semantic content helps other

researchers to find, explore and understand an article; (c) provide *public acknowledgements* for curators.

4.2 (Semi-)automatic Approaches

The second group of use cases require a high coverage while a rather low quality and domain specialisation are acceptable. For these use cases, rather simple and domain-independent ontologies should be developed or reused.

Various approaches can be used to populate an ORKG (semi-)automatically. Methods for *entity and relation extraction* (see Sect. 2.3) can help to populate fine-grained KGs with high coverage and *entity linking* approaches can link mentions in text with entities. For cross-modal linking, Singh et al. [60] propose an approach to detect URLs to datasets in research articles automatically, while the Scientific Software Explorer [34] interlinks text passages in research articles with code fragments. To extract relevant information at sentence level, approaches for *sentence classification* in scientific text can be applied (see Sect. 2.3). To support the curator fill in templates semi-automatically, *template-based extraction* can (1) suggest relevant templates for a research article and (2) pre-fill fields of templates with appropriate values. For pre-filling, approaches such as for natural language inference used in leaderboard construction [36] or end-to-end question answering [23, 57] can be employed.

Further, the system should enable to plugin *external information extractors*, developed for certain scientific domains to extract specific types of information. For instance, as depicted in Fig. 2, an external template information extractor has to implement an interface with three methods. This enables the system (1) to filter relevant template extractors for an article and (2) extract field values from an article.

5 Conclusions

In this paper, we have presented a requirements analysis for an Open Research Knowledge Graph (ORKG). An ORKG should represent the content of research articles in a semantic way to enhance or enable a wide range of use cases. We identified literature-related core tasks of a researcher that can be supported by an ORKG and formulated them as use cases. For each use case, we discussed specificities and requirements for the underlying ontology and the instance data. In particular, we identified two groups of use cases: (1) the first group requires high-quality instance data and rather fine-grained, domain-specific ontologies, but with moderate coverage; (2) the second group requires a high coverage, but the ontologies can be kept rather simple and domain-independent, and a moderate quality of the instance data is sufficient. Based on the requirements, we have described possible manual and semi-automatic approaches (necessary for the first group), and automatic approaches (appropriate for the second group) for KG construction. In particular, we propose a framework with lightweight ontologies that can evolve by community curation. Further, we have described

the interdependence with external systems, user interfaces, and APIs for third-party applications to populate an ORKG.

The results of our work aim to provide a holistic view of the requirements for an ORKG and be a guideline for further research. The suggested approaches have to be refined, implemented and evaluated in an iterative and incremental process (see www.orkg.org for the current progress). Additionally, our paper can serve as a foundation for a discussion on ORKG requirements with other researchers and practitioners.

References

1. Harris, M.A.M.A., et al.: Gene ontology consortium: The gene ontology (GO) database and informatics resource. *Nucleic Acids Res.* **32**, D258–D261 (2004)
2. Amir, A., Jing-bo, W.: Research graph: building a distributed graph of scholarly works using research data switchboard. In: Open Repositories CONFERENCE (2017)
3. Ammar, W., et al.: Construction of the literature graph in semantic scholar. In: NAACL-HLT (2018)
4. Auer, S.: Towards an open research knowledge graph (2018). <https://doi.org/10.5281/zenodo.1157185>
5. Augenstein, I., Das, M., Riedel, S., Vikraman, L., McCallum, A.: Semeval 2017 task 10: scienceie - extracting keyphrases and relations from scientific publications. In: SemEval@ACL (2017)
6. Balog, K.: Entity-Oriented Search. The Information Retrieval Series. Springer, Heidelberg (2018). <https://doi.org/10.1007/978-3-319-93935-3>
7. Bechhofer, S., et al.: Why linked data is not enough for scientists. In: 2010 IEEE 6th International Conference on e-Science (2010)
8. Beel, J., Gipp, B., Langer, S., Breitingner, C.: Research-paper recommender systems: a literature survey. *Int. J. Digit. Libr.* **17**(4), 305–338 (2015). <https://doi.org/10.1007/s00799-015-0156-0>
9. Beltagy, I., Lo, K., Cohan, A.: Scibert: pretrained language model for scientific text. In: EMNLP (2019)
10. Bodenreider, O.: The unified medical language system (UMLS): integrating biomedical terminology. *Nucleic Acids Res.* **32**, D267–D270 (2004)
11. Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. *Commun. ACM* **57**(10), 78–85 (2014)
12. Bornmann, L., Mutz, R.: Growth rates of modern science: a bibliometric analysis based on the number of publications and cited references. *J. Assoc. Inf. Sci. Technol.* **66**(11), 2215–2222 (2015)
13. Brack, A., D’Souza, J., Hoppe, A., Auer, S., Ewerth, R.: Domain-Independent extraction of scientific concepts from research articles. In: Jose, J.M., et al. (eds.) ECIR 2020. LNCS, vol. 12035, pp. 251–266. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45439-5_17
14. Braun, R., Benedict, M., Wendler, H., Esswein, W.: Proposal for requirements driven design science research. In: Donnellan, B., Helfert, M., Kenneally, J., VanderMeer, D., Rothenberger, M., Winter, R. (eds.) DESRIST 2015. LNCS, vol. 9073, pp. 135–151. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-18714-3_9
15. Brodaric, B., Reitsma, F., Qiang, Y.: Skiing with DOLCE: toward an e-science knowledge infrastructure. In: FOIS (2008)

16. Burton, A., et al.: The scholix framework for interoperability in data-literature information exchange. *D-Lib Mag.* **23**(1/2) (2017)
17. Cohan, A., Ammar, W., van Zuylen, M., Cady, F.: Structural scaffolds for citation intent classification in scientific publications. In: *NAACL-HLT* (2019)
18. Cohan, A., Beltagy, I., King, D., Dalvi, B., Weld, D.S.: Pretrained language models for sequential sentence classification. In: *EMNLP* (2019)
19. Constantin, A., Peroni, S., Pettifer, S., Shotton, D.M., Vitali, F.: The document components ontology (DoCO). *Seman. Web* **7**(2), 167–181 (2016)
20. Degbelo, A.: A snapshot of ontology evaluation criteria and strategies. In: *SEMANTICS*, pp. 1–8. *ACM* (2017)
21. Degtyarenko, K., et al.: Chebi: a database and ontology for chemical entities of biological interest. *Nucleic Acids Res.* **36**, 344–350 (2008)
22. Deroncourt, F., Lee, J.Y.: PubMed 200k RCT: a dataset for sequential sentence classification in medical abstracts. In: *IJCNLP* (2017)
23. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: *NAACL-HLT* (2019)
24. Färber, M.: The microsoft academic knowledge graph: a linked data source with 8 billion triples of scholarly data. In: Ghidini, C. (ed.) *ISWC 2019*. LNCS, vol. 11779, pp. 113–129. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30796-7_8
25. Fathalla, S., Vahdati, S., Auer, S., Lange, C.: Towards a knowledge graph representing research findings by semantifying survey articles. In: Kamps, J., Tsakonias, G., Manolopoulos, Y., Iliadis, L., Karydis, I. (eds.) *TPDL 2017*. LNCS, vol. 10450, pp. 315–327. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67008-9_25
26. Fellbaum, C. (ed.): *WordNet: An Electronic Lexical Database*. Language, Speech, and Communication. MIT Press, Cambridge (1998)
27. Fink, A.: *Conducting Research Literature Reviews*. SAGE Publications, Thousand Oaks (2014)
28. Fisas, B., Saggion, H., Ronzano, F.: On the discursive structure of computer graphics research papers. In: *LAW@NAACL-HLT* (2015)
29. Gábor, K., Buscaldi, D., Schumann, A.K., QasemiZadeh, B., Zargayouna, H., Charnois, T.: Semeval-2018 task 7: semantic relation extraction and classification in scientific papers. In: *Proceedings of The 12th International Workshop on Semantic Evaluation* (2018)
30. Groza, T., Kim, H., Handschuh, S.: Salt: semantically annotated latex. In: *SAAW@ISWC* (2006)
31. Handschuh, S., QasemiZadeh, B.: The ACL RD-TEC: a dataset for benchmarking terminology extraction and classification in computational linguistics. In: *COLING 2014: 4th international workshop on computational terminology* (2014)
32. Hars, A.: *Structure of Scientific Knowledge*. Springer, Heidelberg (2003)
33. Hevner, A.R., March, S.T., Park, J., Ram, S.: Design science in information systems research. *MIS Q.* **28**(1), 75–105 (2004)
34. Hoppe, A., Hagen, J., Holzmann, H., Kniesel, G., Ewerth, R.: An analytics tool for exploring scientific software and related publications. In: Méndez, E., Crestani, F., Ribeiro, C., David, G., Lopes, J.C. (eds.) *TPDL 2018*. LNCS, vol. 11057, pp. 299–303. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00066-0_27
35. Horváth, I.: Comparison of three methodological approaches of design research. In: *ICED* (2007)
36. Hou, Y., Jochim, C., Gleize, M., Bonin, F., Ganguly, D.: Identification of tasks, datasets, evaluation metrics, and numeric scores for scientific leaderboards construction. In: *ACL* (2019)

37. Jaradeh, M.Y., Oelen, A., Prinz, M., Stocker, M., Auer, S.: Open research knowledge graph: a system walkthrough. In: Doucet, A., Isaac, A., Golub, K., Aalberg, T., Jatowt, A. (eds.) TPDL 2019. LNCS, vol. 11799, pp. 348–351. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30760-8_31
38. Kim, S., Martínez, D., Cavedon, L., Yencken, L.: Automatic classification of sentences to support evidence based medicine. In: BMC Bioinformatics (2011)
39. Kitchenham, B., Charters, S.: Guidelines for performing systematic literature reviews in software engineering. Keele University and Durham University Joint Report, Technical report (2007)
40. Klampanos, I.A., Davvetas, A., Koukourikos, A., Karkaletsis, V.: Annett-o: an ontology for describing artificial neural network evaluation, topology and training. *IJMSE* **13**, 24–49 (2018)
41. Kolitsas, N., Ganea, O.E., Hofmann, T.: End-to-end neural entity linking. In: CoNLL (2018)
42. Lange, C.: Ontologies and languages for representing mathematical knowledge on the semantic web. *Semant. Web* **4**, 119–158 (2013)
43. Lehmann, J., et al.: Dbpedia - a large-scale, multilingual knowledge base extracted from wikipedia. *Semant. Web* **6**, 167–195 (2015)
44. Liakata, M., Saha, S., Dobnik, S., Batchelor, C., Rebholz-Schuhmann, D.: Automatic recognition of conceptualization zones in scientific articles and two life science applications. *Bioinformatics* **28**(7), 991–1000 (2012)
45. Liakata, M., Teufel, S., Siddharthan, A., Batchelor, C.R.: Corpora for the conceptualisation and zoning of scientific papers. In: LREC (2010)
46. Luan, Y., He, L., Ostendorf, M., Hajishirzi, H.: Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction. In: EMNLP (2018)
47. Lubani, M., Noah, S.A.M., Mahmud, R.: Ontology population: Approaches and design aspects. *J. Inf. Sci.* **45**(4), 502–515 (2019)
48. Manghi, P., et al.: The OpenAIRE research graph data model (2019). <https://doi.org/10.5281/zenodo.2643199>
49. Mesbah, S., Fragkeskos, K., Lofi, C., Bozzon, A., Houben, G.J.: Semantic annotation of data processing pipelines in scientific publications. In: Blomqvist, E., Maynard, D., Gangemi, A., Hoekstra, R., Hitzler, P., Hartig, O. (eds.) ESWC 2017. LNCS, vol. 10249, pp. 321–336. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-58068-5_20
50. Nasar, Z., Jaffry, S.W., Malik, M.K.: Information extraction from scientific articles: a survey. *Scientometrics* **117**(3), 1931–1990 (2018). <https://doi.org/10.1007/s11192-018-2921-5>
51. Oelen, A., Jaradeh, M.Y., Farfar, K.E., Stocker, M., Auer, S.: Comparing research contributions in a scholarly knowledge graph. In: SciKnow@K-CAP (2019)
52. Okoli, C.: A guide to conducting a standalone systematic literature review. *CAIS* **37**, 43 (2015)
53. Peroni, S., Shotton, D.M.: Fabio and cito: ontologies for describing bibliographic resources and citations. *J. Web Semant.* **17**, 33–43 (2012)
54. Pertsas, V., Constantopoulos, P.: Scholarly ontology: modelling scholarly practices. *Int. J. Digit. Libr.* **18**(3), 173–190 (2016). <https://doi.org/10.1007/s00799-016-0169-3>
55. Petasis, G., Karkaletsis, V., Paliouras, G., Krithara, A., Zavitsanos, E.: Ontology population and enrichment: state of the art. In: Paliouras, G., Spyropoulos, C.D., Tsatsaronis, G. (eds.) Knowledge-Driven Multimedia Information Extraction and

- Ontology Evolution. LNCS (LNAI), vol. 6050, pp. 134–166. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20795-2_6
56. Pujara, J., Singh, S.: Mining knowledge graphs from text. In: WSDM 2018 (2018)
 57. Rajpurkar, P., Zhang, J., Lopyrev, K., Liang, P.: Squad: 100, 000+ questions for machine comprehension of text. In: EMNLP (2016)
 58. Ruiz Iniesta, A., Corcho, O.: A review of ontologies for describing scholarly and scientific documents. In: 4th Workshop on Semantic Publishing (SePublica) (2014)
 59. Salatino, A.A., Thanapalasingam, T., Mannocci, A., Birukou, A., Osborne, F., Motta, E.: The computer science ontology: a comprehensive automatically-generated taxonomy of research areas. In: Data Intelligent (2019)
 60. Singh, M., et al.: Ocr++: a robust framework for information extraction from scholarly articles. In: COLING (2016)
 61. Soldatova, L.N., King, R.D.: An ontology of scientific experiments. *J. R. Soc. Interface* **3**, 795–803 (2006)
 62. Stocker, M., Prinz, M., Rostami, F., Kempf, T.: Towards research infrastructures that curate scientific information: a use case in life sciences. In: Auer, S., Vidal, M.E. (eds.) DILS 2018. LNCS, vol. 11371, pp. 61–74. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-06016-9_6
 63. Teufel, S., Siddharthan, A., Batchelor, C.: Towards discipline-independent argumentative zoning: evidence from chemistry and computational linguistics. In: EMNLP (2009)
 64. Vahdati, S., Fathalla, S., Auer, S., Lange, C., Vidal, M.E.: Semantic representation of scientific publications. In: Doucet, A., Isaac, A., Golub, K., Aalberg, T., Jatowt, A. (eds.) TPD 2019. LNCS, vol. 11799, pp. 375–379. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30760-8_37
 65. Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. *Commun. ACM* **57**(10) (2014)
 66. de Waard, A., Tel, G.: The ABCDE format enabling semantic conference proceedings. In: SemWiki (2006)
 67. Xiong, C., Power, R., Callan, J.P.: Explicit semantic ranking for academic search via knowledge graph embedding. In: WWW (2017)
 68. Yaman, B., Pasin, M., Freudenberg, M.: Interlinking scigraph and dbpedia datasets using link discovery and named entity recognition techniques. In: LDK (2019)