Semantic Web 0 (0) 1
IOS Press

Teaching Knowledge Graph for Knowledge Graphs Education

Eleni Ilkou [a,*], Hasan Abu-Rasheed [b], David Chaves-Fraga [c], Erna Engelbrecht [d], Ernesto Jiménez-Ruiz [e,f] and Jose Emilio Labra-Gayo [f]
[a] *L3S Research Center, Leibniz University Hannover, Hannover, Germany*
E-mail: ilkou@l3s.de
[b] *Department of Electrical Engineering and Computer Science, University of Siegen, Siegen, Germany*
E-mail: hasan.abu.rasheed@uni-siegen.de
[c] *Department of Electronics and Computing, Universidade de Santiago de Compostela, Santiago de Compsotela, Spain*
E-mail: david.chaves@usc.es
[d] *Department of Spatial Planning and Strategy (SPSS), Delft University of Technology, Delft, The Netherlands*
E-mail: E.Engelbrecht@tudelft.nl
[e] *Department Computer Science, City St George's, University of London, London, United Kingdom*
E-mail: ernesto.jimenez-ruiz@city.ac.uk
[f] *Department of Computer Science, WESO, University of Oviedo, Oviedo, Spain*
E-mail: labra@uniovi.es

Abstract.
The effective education and training of knowledge graph (KG) practitioners and students are critical for sustaining the growth and utility of the KG ecosystem. This paper introduces the Teaching KG, a novel educational resources catalog designed to enhance KGs education. The Teaching KG integrates pedagogically enriched components such as skills, knowledge topics, courses, instructors, and educational materials, offering a semantically structured framework to support diverse educational needs. By leveraging a higher ontology and semantic constraints, it facilitates the integration of educational content, ensuring consistency and scalability. We present use cases that demonstrate the resource's applicability as an instructional tool for KG instructors, learners, and technology-enhanced learning applications. Preliminary evaluations highlight its potential to support KG course instructors in enhancing their courses and locating relevant resources. This work contributes to the ongoing project on the role of educational KGs [1], proposing a methodology to consolidate the collaborative initiatives within the Semantic Web community. The result is a centralized hub that interlinks resources, courses, and instructors to foster a more cohesive educational KG ecosystem.

GitHub project: https://github.com/naiayti/TeachingKG

Keywords: Educational KG, Personal KGs, Domain Model, User Model

1. Introduction

Ontologies and Knowledge Graphs (KGs) are widely used across multiple domains and are becoming more popular in educational applications [2]. By providing a structured and interconnected framework for representing

---

[*] Corresponding author. Besides the first author, the order is alphabetical. E-mail: ilkou@l3s.de.

domain knowledge, ontologies and KGs can facilitate more effective navigation, discovery, and reuse of educational resources [3]. This can help instructors create personalized learning paths, identify learners' knowledge gaps, and track student progress with greater precision [4]. Additionally, ontologies and KGs establish clear relationships between concepts, allowing for more nuanced understanding and application of complex ideas [5]. Furthermore, the machine-readable nature of these frameworks enables seamless integration with educational technologies, such as intelligent tutoring systems [6], leading to a more efficient, effective, and engaging educational experience [7].

In parallel, the demand for digital education has grown significantly with an increase in internet users taking online courses or learning via online educational material, as reported by Eurostat[1]. However, online learners often struggle to find high quality material aligned to their learning goals. This challenge arises for several reasons. Firstly, simple web searches are often insufficient, as relevant resources are not semantically interlinked with related topics or courseware. Secondly, there is a lack of open, centralized hubs offering educational resources classified by topic and interlinked with similar courses or materials. Moreover, university courses covering similar skills and topics frequently differ in their course titles, creating additional barriers to discovery. This means that education is characterized by the compartmentalization of teaching resources, where instructors often rely on isolated and disconnected materials, resulting in a fragmented and often incomplete learning experience for students. Consequently, this fragmentation hinders meaningful semantic connections between materials and concepts, ultimately limiting the effectiveness of learning and teaching.

More specifically in KG university education, we find no central hub for teaching resources for KG courses. Furthermore, identifying relevant courses for specific topics can be challenging, as KG courses often have various titles unrelated to common keywords, such as "Data Governance", "Web AI" and "Information Management". Additionally, the diverse backgrounds and research interests of instructors teaching KG courses result in inconsistent coverage of skills and knowledge topics, resulting in a fragmented knowledge landscape. This lack of coherence complicates effective discovery, retrieval, and connection of educational resources.

To address these issues, we propose the use of an educational Teaching Knowledge Graph (Teaching KG) focused on skills, knowledge topics, courses and material related to KGs. We argue that KGs can provide a structured representation of teaching resources, highlight their connections, and enable instructors to discover pedagogically classified, high quality, relevant resources for KG education. Our goal is to provide an integrated framework that connects teaching and learning resources with KG-related concepts, facilitating navigation through a vast array of educational content while providing semantic relationships between the different resources. By leveraging an ontology schema, the Teaching KG ensures effective data retrieval. This initiative aims to enhance the quality and effectiveness of university-level courses on KGs and contribute to the development of educational KGs, which can assist researchers, instructors, and practitioners in advancing their work.

Summarizing, the main contributions of this study are the following:

- a methodology for creating an educational KG for teaching resources
- pedagogical parameters that describe courses and their encoding in the ontology schema
- an ontology for supporting a teaching KG in any educational domain, that interconnects skills, topics, lectures and lecturers with educational resources
- SPARQL queries to retrieve the competency questions
- a KG with resources for KGs education
- rules and constraints that validate the semantic structure of the KG
- an RML pipeline for the KG creation and maintenance
- a user-interface and use cases description that further improve the usability of the system.

2. Related Works

Plenty of ontology-based systems for curriculum modeling have been proposed in the literature [8]. While some approaches developed an ontology for describing a specific curriculum [9], most research focused on high-level de-

---

[1] News article by Eurostat: "Increase in online education in the EU in 2023" at https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20240124-2

scription of main concepts. For instance, Chi [10] proposed an ontology-based curriculum knowledge base system, leveraging structured knowledge representation and rules to enable semantic querying for curriculum management. Many similar ontologies aimed to map functionalities present in learning management systems [11], syllabi [12, 13] or at institutional level the creation of program credentials [14]. A series of ontology-based applications that utilized Semantic Web tools were developed by mEducator [15, 16]. The proposed approach aimed to share and reuse medical educational content across institutions using Linked Data principles. Although these initiatives set the tone for Semantic Web applications in e-learning domain, the projects were not maintained and the resources were not shared as open source, limiting their reuse.

Recently, the deployment of KGs for educational applications has gained traction. Based on a recent survey by Fettach et al. [17], KG applications in education can be categorized into three clusters: assisted instruction, including institutional concept and knowledge management, assisted learning, with the personalized learning and question-answering, and educational assessment. These structured representations of knowledge allow for the formalization and interlinking of educational content, offering opportunities for improved educational outcomes and management. Towards assisted learning, recent studies have demonstrated the use of KGs in representing and managing knowledge in higher education. Li et al. [18] proposed a framework for constructing and fusing KGs tailored for Electronic Information majors, addressing the relationships between courses and their key concepts. This approach highlights interconnections among curricula while providing a mechanism for identifying critical knowledge points, thus improving the efficiency of student learning. Similarly, Hubert et al. introduced EducOnto, an ontology for modeling university curricula and student profiles, which serves as the foundation for their KG called EduKG [19]. Abu-Rasheed et al. [20] transformed hierarchical learning object (LO) models into contextual KGs using customized text mining and relation extraction pipelines. This transformation supported personalized learning by enhancing the representation and linkage of LOs' contexts, scaffolding learners' progression from recall-level learning objectives to higher-order objectives, such as application and analysis. The Education-Oriented Knowledge Graph (EOKG) [21] quantified personalized learning by defining boundaries of mastered and unmastered knowledge points. This approach enables dynamic modeling of student learning profiles, offering tailored educational experiences. Chen et al. [22] proposed the KnowEdu system, which employs heterogeneous educational data, such as pedagogical and learning assessment datasets, to automate the construction of KGs. To achieve this, neural sequence labeling is used for instructional concept extraction and probabilistic association rule mining to establish meaningful educational relationships. Similarly, EDUKG [23] was proposed as a heterogeneous K-12 educational KG. It introduced a fine-grained ontology consisting of 635 classes and 1,314 properties, and incorporated methodologies for interactively extracting factual knowledge from textbooks and dynamically maintaining the KG using a generalized entity linking system. Efforts to construct KGs for elementary education have emphasized the importance of ontology design and resource integration. By employing unsupervised methods for instance extraction and incorporating diverse learning materials, these initiatives demonstrate the feasibility and efficiency of KG-based educational tools [24].

Further, semantic summarization frameworks have been developed to support intelligent retrieval and comprehension of educational resources. For instance, Yu et al. [25] develop a system targeting entrepreneurship education, integrating NLP techniques such as text embedding and graph convolution to improve question-answering capabilities. The application of KGs extends beyond content structuring to address broader educational needs. Research has proposed models for integrating KGs into teaching resources, personalized instruction, and educational decision-making. These models aim to optimize teaching methodologies and administrative strategies, furthering the impact of AI technologies in education [26].

Despite these advancements, there are challenges that remain unresolved in developing and applying KGs in education. The extraction of high-quality, domain-specific knowledge remains an unresolved issue. Li et al. [18] noted the complexity involved in integrating data from diverse sources into a coherent KG. Additionally, while LLMs show promise in enhancing the adaptability of educational technologies, issues such as scalability, interpretability, and ethical considerations related to AI integration persist [18, 27] Therefore, our proposed ontology KG construction approach aim to address these issues by creating high-quality, pedagogically sound content for teaching KGs. This approach is designed to support higher education and self-regulated learning environments, bridging existing gaps in educational resource integration.

3. Methodology

3.1. Motivation

Although KGs have been applied in educational settings, they are usually tailored to a specific downstream application. To the best of our knowledge, no existing work outlines the requirements or methodologies for deploying an educational KG. This highlights the need for standards that bridge the gap between Semantic Web technologies and pedagogical frameworks to accommodate the adaptation of semantic solutions in technology-enhanced learning community and provide smart learning analytics. Moreover, new instructors often face significant challenges when designing modules for new courses, particularly in sourcing material for lab sessions and coursework. A course title alone is insufficient to reveal the skills and knowledge topics a resource covers, making it difficult to link related additional materials. These challenges can reduce the potential quality of content offered to the learners and are highly time consuming for instructors.

Our work aims to address these issues by proposing a methodology for generating an educational teaching KG and populated example of a teaching KG for KG education. In this paper, we focus on the development of a teaching KG specific to KG-related skills and courses in higher education, as offered by experts in the Semantic Web community. In Section 1, we highlighted the absence of a centralized repository for KG-related educational resources and the challenges posed by inconsistent course naming conventions in discovering similar resources online. Unlike programming or machine learning courses, KG courses are scarce on platforms such as Coursera[2] or edX[3]. Although there are some recent courses focus on learning on graphs, these do not cover the foundational aspects of KGs or align with the content taught in university courses by the Semantic Web community. Furthermore, based on input we received, the Data Science students at City St. George's reported difficulties in finding alternative code examples and supplementary KG materials to complement their studies. In contrast, resources for other topics, such as machine learning, computer vision, and data visualization, are far more many and online available, highlighting the relative scarcity of KG educational content online.

To address these limitations, the Interest Group on KGs at the Alan Turing Institute and the COST Action on Distributed Knowledge Graphs (DKG) have initiated independent efforts to collect and publicly share teaching materials on KGs. Previously, these materials were restricted to learning management systems such as Moodle. We provide a detailed analysis on the data collection process in Section 3.3.

3.2. Context of Research

The research was conducted to address key challenges associated with integrating and connecting complex, unstructured educational data from multiple sources within a cohesive framework. There are limitations of current open educational resources and educational KGs in capturing intricate relationships, and enabling semantic interoperability. Further, there is a need for supporting dynamic, real-time updates in educational catalogs which have highlighted the importance of semantically advanced approaches. The study identifies two primary challenges, which are addressed through the proposed research questions.

Challenge 1: One challenge is balancing the need for comprehensive representation of educational resources versus the associated workload with providing all the resource metadata. This challenge is reflected additionally in the use case where the instructor uses the KG to search for relevant content. The relevance of the search results retrieved from the catalog is affected by the level of detail in which its content is described. As the educational resources can be found in a variety of formats and description, this requires a considerable amount of time and effort to insert all the appropriate metadata values, which often is not an automated procedure. Therefore, a reasonable minimum of the values that instructors need to provide is needed to ensure algorithmic capabilities while keeping the user experience sufficiently simple and enabling educators to adopt the proposed schema.

---

[2] As per December 2024, Coursera has only one KG course, the "Knowledge Graphs for RAG".

[3] As per December 2024, edX has no KG courses.

Challenge 2: Contextual relevance versus feasible execution. Educational institutions and educators use different terminology, theories and educational design approaches. The challenge lies in ensuring that the KG effectively captures nuanced, contextual meanings behind course content, learning outcomes and instructional approaches across multiple sources, allowing instructors to contribute and find resources that are relevant and pedagogically aligned with their course design goals.

To tackle these challenges, a set of research questions (RQs) was formulated, aiming to guide the development of the ontology and a robust pipeline for populating a KG. These research questions focused on defining the ontology's structure, ensuring the scalability of the data pipeline, and assessing the KG's effectiveness in enabling efficient data retrieval and supporting downstream use case scenarios. By addressing these questions, the study seeks to create a flexible, pedagogically-founded, and semantically enriched data environment that enhances knowledge discovery and supports informed decision-making in the targeted domain. Therefore, we formulate the following research questions for our study:

RQ1: Which are the minimum pedagogical requirements a teaching KG catalog should include?

RQ2: How can the Knowledge Graph ensure contextual alignment of educational resources to enable cross-course knowledge and content sharing?

RQ3: What ontology classes and properties constitute the minimum requirements for representing an educational schema to enable effective KG-based information retrieval?

3.3. Data Collection

Courses and topics related to KGs were collected with the help of Semantic Web experts. The resources were primarily gathered from two initiatives: the Alan Turing Institute KG resources[4], and the European Union innovation network of COST Action "Distributed Knowledge Graphs" (COST DKG), CA19134[5]. Additionally, supplementary resources were obtianed the authors private lists with resources.

For the course collection, resources were resources collected via the Alan Turing Institute, the COST Action DKG and publicly available online courses. Through COST Action DKG, a data collection form was distributed to Semantic Web experts, requesting personal information and course details, such as course names, URLs, and weekly lecture schedules' titles. The topics were described by Semantic Web experts who teach KG courses within two COST Action DKG initiatives. The experts described the skills and knowledge topics based on an optional template they were provided. Experts provided descriptions of skills and knowledge topics using an optional template designed to capture core educational aspects. The template intentionally avoided restrictions or mandatory fields, allowing flexibility for brainstorming and creativity in expert contributions. In Table 1 we present a portion of the template used to guide the descriptions of skills and topics, which were populated collaboratively by groups of at least three participants. The collection of educational materials began with voluntary contributions from the course instructors. To expand the scope and depth of the repository, additional resources were incorporated, including Python libraries, visualization tools, and tutorials.

Table 1

Part of template for describing skills and topics that was provided to the experts.

Description	**Property**
Skills and objectives of the course (learning goals and objectives)	schema: assesses
Content and subtopics into 1. Theoretical level and 2. Practical level	schema: teaches
Prerequisites for the course	schema: competencyRequired
Educational level and level of difficulty of the certain topic	schema: educationalLevel

---

[4] You can access the gathered resources of the Knowledge Graphs Interest Group at the Alan Turing Institute at https://github.com/turing-knowledge-graphs

[5] The COST Action "Distributed Knowledge Graphs" ended on September 2024. More information about the Action can be found at https://cost-dkg.eu/

3.4. Research Design

The analysis of the collected data and its diverse formats revealed the need for a uniform schema to represent and interconnect courses, topics and educational material effectively. To address this, we focused on defining a schema that could standardize the representation of core educational aspects. Consequently, in the initial steps of this project, we formulated use cases as competency questions that a teaching KG could address, as detailed in our previous work [1].

To achieve high retrieval capabilities and unify the schema for structuring and organizing the collected data, we propose an ontology to support the teaching KG. The ontology serves as the foundation for standardizing relationships between courses, topics, and resources. Following the ontology design, we implemented a data pipeline to enable the integration of raw data into the KG. This pipeline populates the ontology with individuals and supports the life cycle of data loading in the KG, enabling dynamic updates and scalability as new data become available. To validate our approach, we propose semantic constraints, conduct stakeholder interviews, and describe detailed use cases for utilizing the teaching KG through the system interface.

3.5. Pedagogical Alignment

As Semantic Web experts are typically not trained in pedagogical methodologies or familiar with the educational parameters required to classify their courses effectively, it is important to identify the key components to pedagogically classify the resources provided. Also, educational institutions employ diverse educational design approaches, resulting in considerable variation in the data collected from the experts. For example, while some courses specified learning objectives, others listed only topics or included both. Additionally, courses originated from different countries, each with distinct standards for indicating notional hours and licensing agreements. Most courses also lacked detailed educational parameters. To address these challenges, we developed a comprehensive classification scheme that accommodates varying contexts without compromising data integrity.

The goal of our approach is to ensure pedagogical alignment between various elements, such as topics, skills and learning outcomes, while enabling educators to select and integrate content that meets their teaching goals and students' learning needs. The system supports dynamic and flexible updates, facilitating ongoing pedagogical alignment to maintain the relevance of both the KG and the courses informed by it. Additionally, the guidance is embedded in the user interface, assisting educators in including essential information when logging their course details and content. This aims to strike a balance between structured input and free-form text, providing contextual flexibility.

4. Ontology

4.1. Ontology Development

We chose to develop an ontology as the semantic structure to interconnect the populated data within the KG. This decision was motivated by a number of reasons. First, creating an ontology facilitates future data integration by establishing a common vocabulary and framework, ensuring the cohesiveness of the KG. A semantic structure enables efficient data retrieval and addition, while also promoting semantic interoperability across systems and enhancing data reusability. Our goal is to generate an open-source semantic resource capable of supporting multiple systems with a schema that accommodates the specification of elements present in the KG. Secondly, an ontology can improve data quality control via constraints and simplify maintenance. As we gather the resources from experts in the Semantic Web community, we can guarantee the initial quality of content of the produced resource, which further allows Artificial Intelligence and Machine Learning applications. Finally, we choose to support our KG with an ontology to enable detailed mapping and precise retrieval of the data. By developing a detailed ontology we can keep track of the data present in the KG and have the ability to query and perform question-answering with high retrieval precision.

For the ontology development we followed the workflow suggested by LOT [28], incorporating the conceptualization step described in Sabiox [29]. Initially, use cases were specified in previous work as competency questions [1]. The ontology was designed to describe resources that can be found in an educational KG with the emphasis on teaching material. To create a semantically structured catalog that accommodates education broadly—rather than being confined to a single domain—we developed a high-level ontology. This approach allows for future adaptation from other areas in computer science and educational domains. The requirements for reusing our ontology include the presence of predefined skills, knowledge topics and/or learning objectives of the educational domain, as well as a catalogue of courses, the instructors, and educational material, such as PDF slides. The end users of our ontology are researchers and practitioners interested in developing teaching KGs for any educational domain. Finally, the ontology was developed in English using both the online and desktop versions of Protégé, with the domain area of focus being higher education KGs.

4.2. Ontology Description

We reused vocabularies from Semantic Web resources, as we report in Table 2. Primarily, we leveraged classes and properties from schema.org[6] and EduCOR ontology [30]. A visualization of the ontology is provided in Figure 1. Our ontology consists of 13 classes, 16 object, and 16 datatype properties. It consists of different modules that connect the main classes in the ontology: the skill module, the course or the lecturer model, and the educational material. At this stage, the ontology does not include hierarchical structures or prerequisites at any level. The ontology is described in Turtle syntax and presented in English. Below, we analyze the main modules in more detail, derived from the primary classes in our schema.

Table 2

Vocabularies in our ontology.

Domain	**URI**
schema:	http://schema.org/
educor:	https://github.com/tibonto/educor#
xsd:	http://www.w3.org/2001/XMLSchema#
owl2:	http://www.w3.org/2006/12/owl2#
owl:	http://www.w3.org/2002/07/owl#
rdf:	http://www.w3.org/1999/02/22-rdf-syntax-ns#
rdfs:	http://www.w3.org/2000/01/rdf-schema#
skos:	http://www.w3.org/2004/02/skos/core#
courseskos:	https://w3id.org/coursesonto/kos/
base:	https://w3id.org/coursesonto#

The ontology uses higher-level pedagogical and semantic constraints to ensure consistency and interoperability across its components. Courses are structured as learning paths designed to teach specific skills or topics, and lecturers are modeled as entities that deliver these courses. Educational resources are classified as core or supplementary, which could include textbooks, videos, and tutorials. The ontology is designed to map the provided data and support various educational scenarios for educators and learners.

The Skill module of the ontology consists of the skills and topics as they were described by the experts in the Semantic Web community, as described in Section 3.3. In Table 3, we present the skills and their associated knowledge topics, categorized into into main and secondary groups. The main skills were identified through an analysis of the collected resources, highlighting three key areas frequently addressed in KG courses: the Resource Description Framework (RDF), the Knowledge Graph Standards and Tools, and SPARQL. The differentiation of the skills and knowledge topics can be complex and occasionally confusing, even for the experts. To address this,

---

[6] https://schema.org/

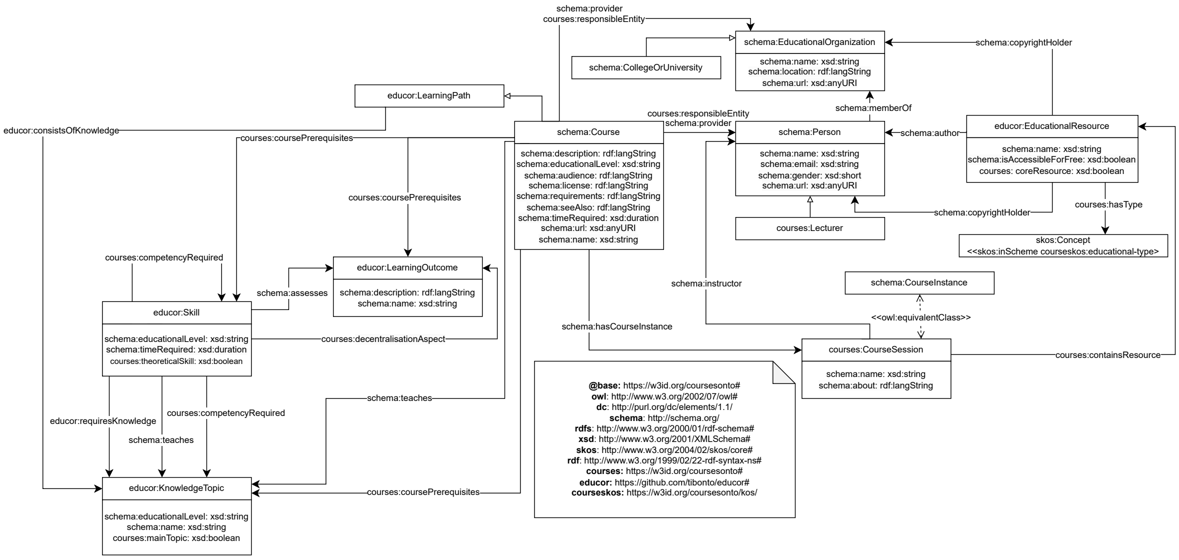

Figure 1. Visualization of the ontology classes and properties using Chowlk [31].

we adopted the following definitions:

- Skill: An abstract concept consisting of multiple elements and in different levels and scopes to be mastered, the knowledge topics.
- Knowledge topic: A concrete unit that contains theoretical and practical (educational) resources and can be part of multiple skills.

The Courses and Lecturers module could be viewed as a personal KGs around the persons who teach a course in order to achieve connection of the course content behind university paywalls with researchers' contact information. The courses are structured as comprehensive learning units that cover a subset of skills and knowledge topics. Each course is linked to its lecturer, which allows for tracking the availability of subject-matter experts and mapping their contributions to the educational ecosystem. The Educational Resources module contains the supporting materials for both learners and educators. These material include textbooks, videos, and libraries, and are annotated with metadata to foster discoverability and reuse. In order to distinguish if an educational resource is theoretical or practical we use the skos:Concept class. Our goal is to achieve a conceptual structure of the KG individuals that works as a classification scheme. For this reason, we define the educational type as "theory", "practical" or "mix" of the Educational Resources.

4.2.1. Pedagogical Classification

To inform the technical development of the KG, a pedagogical classification was created to define the key entities, attributes, and relationships that should be represented. This classification aimed to ensure that the KG would be comprehensive, contextually relevant, and adaptable to various educational practices. The classification was developed through an iterative process, leveraging both pedagogical expertise and feedback from project stakeholders.

The foundation for this classification was drawn from an analysis of common course structures and teaching practices across different institutions. Recognizing that courses often vary in the way they present information and the types of data they include, we prioritized creating a schema that balances detail and flexibility. This classification framework ensures that essential educational data can be captured while maintaining enough adaptability to accommodate courses from diverse educational systems and teaching approaches. In Table 4 we present the resulting classification, which consists of five main categories:

- Course Data: This includes basic information about the course, such as the title and a descriptive overview. Details such as notional hours and copyright information were also included to provide context and support legal compliance for the reuse and sharing of course content. Additionally, since notional hours are calculated differently

Table 3

After our analysis we found out that the majority of the KG courses teach three main skills, and plenty secondary. We present the topics taught in KG courses as their corresponding skills, and their potential objectives as knowledge topics.

Main Skills	**Knowledge Topics**
RDF	RDF (Resource Description Framework) introduction. RDF syntax (Turtle, RDF/XML, JSON-LD).
Knowledge Graph Standards and Tools	Overview of standards like RDF, RDFS, OWL, and SHACL, FAIR principles. Popular knowledge graph tools and frameworks (e.g., Apache Jena, Neo4j).
SPARQL	SPARQL query language for querying RDF data.
Secondary Skills	**Knowledge Topics**
Linked Data	Linked Data principles and best practices. Ontologies and their role in the Semantic Web.
Graph Theory	Introduction to graph theory concepts (nodes, edges, etc.) Types of graphs (property, directed, etc.) Graph algorithms for KGs (e.g., community detection, centrality measures) Graph traversal techniques.
Validation of KGs	Data accuracy, consistency, and adherence to defined schemas or constraints like SHACL and SHEx. Schema validation, data quality dimensions (e.g., correctness, completeness), and resolving inconsistencies through logical inference and debugging.
Ontology Engineering	Ontology design patterns and best practices. Ontology alignment and merging.
Knowledge Graph Construction	Data extraction and integration techniques. Entity recognition and disambiguation. Knowledge graph population and curation.
Knowledge Graph Reasoning.	Inference and reasoning in KGs. Rule-based and ontology-based reasoning.
Scalability and Performance	Knowledge graph storage and processing. Scalability challenges and solutions. Benchmarking and performance optimization.
NLP, ML and Knowledge Graphs	Named entity recognition, extraction and linking. Relation extraction from text. Graph neural networks for KG analytics. Knowledge graph completion, link prediction and entity embeddings
Case Studies	Industry-specific use cases (e.g., healthcare, e-commerce). Practical examples of KG implementations (e.g. fact checking, QnA)
Knowledge Graph Visualization	Visualization techniques for KGs. Tools and libraries for creating interactive visualizations.
AI Ethics, Bias and Legal Considerations	Data privacy and security in KGs.Ethical concerns related to KG construction and usage. Addressing bias and fairness issues in KGs Compliance with regulations (e.g., GDPR).
Blockchain and Decentralized Knowledge Graphs	Decentralized KG architectures using blockchain technology. Data ownership and security in decentralized KGs.
Domain Knowledge Graphs	In-depth exploration of KGs in specialized fields (e.g., healthcare, geospatial). Domain-specific standards and ontologies (e.g., HL7 FHIR in healthcare).
Knowledge Graph Governance and Quality	Techniques for assessing and improving data quality in KGs. Governance models and policies for maintaining KGs.

in each country's educational system, this will be indicated on the user interface as the 'time needed to complete the course'. From a legal perspective, clear copyright details help define the ownership and permissible use of educational content, protecting both the rights of the content creator and the users. Without proper copyright attribution, there is a risk of unintentional infringement, which can lead to legal complications for both individual educators and institutions. Additionally, copyright information outlines any restrictions on the redistribution or modification of course materials, which is particularly important when sharing resources across different educational contexts. In the context of this project, where educators contribute their course content to populate a shared KG, including copyright information is essential for maintaining transparency and trust. It allows contributors and users to understand the terms under which the materials can be accessed, used, and adapted. By specifying this information, educators enable a seamless integration of resources, fostering an environment of responsible and sustainable content sharing.

- Course Content: This category covers the primary educational materials used in the course, such as lecture notes, assessments, and datasets. Additional resources like external links and supplementary assessments were also considered to support a richer learning experience.
- Learning Aspects: Understanding the pedagogical objectives is vital for effective course design. Therefore, this section captures the main topics taught, the skills developed through the course, and the intended learning outcomes. This classification helps ensure that course content aligns with specific educational goals and facilitates cross-course comparability. We opted not to make 'learning outcomes' compulsory, as not all courses are required to list them (or call them as such), as mentioned in Section 3.
- Facilitators: Information about the contact person(s) is also needed, so information about the main facilitators of the course is included, as well as information about teaching assistants and tutors.
- Requirements for Following the Course: This section outlines entry requirements and other prerequisites necessary for participants, such as required software and the intended target audience. This information helps potential learners assess their readiness for the course and assists educators in aligning their teaching with learner needs.

Table 4

Course Data Categories and Descriptions

Category	**Description**	**Key Data**
Course Data	Basic information about the course, such as the title and a descriptive overview. Details such as notional hours and copyright information are included to provide context and support legal compliance for the reuse and sharing of course content. Since notional hours are calculated differently across countries, this is presented on the user interface as the "time needed to complete the course." Clear copyright details define ownership and permissible use of educational content, protecting both creators and users from unintentional infringement. In the context of this project, copyright information is essential for maintaining transparency, trust, and legal clarity when contributing or accessing course content in the shared KG.	Title AND description AND notional hours AND license
Course Content	Covers primary educational materials, such as lecture notes, assessments, and datasets. Additional resources, such as external links and supplementary assessments, enhance the learning experience.	Core resources, such as educational resources OR assessments OR datasets
Learning Aspects	Focuses on the pedagogical objectives, including the main topics taught, skills developed, and intended learning outcomes. This ensures alignment with educational goals and facilitates cross-course comparability. Learning outcomes are not mandatory, as not all courses are required to list them explicitly.	Knowledge topics OR skills or learning outcomes
Facilitators	Includes information about the contact person(s) responsible for the course, as well as details about teaching assistants and tutors.	Name AND email AND affiliation of the course provider OR tutors/lectures
Requirements for Following the Course	Outlines entry requirements and prerequisites, such as required software and the intended target audience. This information helps learners assess readiness and assists educators in aligning teaching with learner needs.	Entry requirements OR required software OR audience

This structure ensures that educators can input essential information while also allowing for optional details that enhance the depth and contextual flexibility of the data. The classification was used as the basis for constructing the KG and defining entities, attributes, and relationships. Each of the categories was designed to capture both mandatory and optional data points. A bare minimum of entities per classification was made compulsory for educators to complete when inputting their course data to ensure a pedagogically sound KG that facilitates knowledge sharing and enhances course design across different contexts.

4.3. Competency Questions into SPARQL Queries

The competency questions describe the potential use cases of the teaching KG and the ontology that structures it. In Table 5, we present the competency questions that were previously described by Ilkou and Jimenez [1], along with their coverage in this paper. We then present the SPARQL queries corresponding to these competency questions grouped by topic, course, dataset and material. We disclose the SPARQL queries for the competency questions we cover in this paper. In Table 5, the Yes* answers refer to slight change in the competency question to retrieve the theoretical educational material instead of a specific multimedia type (e.g. slides). As we report, we cover most of the competency questions and allow a great range of retrieval resources from the ontology and the system interface.

Table 5

The use cases presented as competency questions taken from Ilkou and Jimenez [1], and the coverage in the Teaching KGs. The Yes* annotation refers to a slight change of the competency question.

Competency questions from Ilkou and Jimenez [1]	**Ontology answered by SPARQL queries**	**System answered by Interface**
For (sub)topic X		
Who is teaching X?	Yes	Yes
Which are the materials for X?	Yes	Yes
Which are the prerequisites of X?	Yes	Yes
Which are the laboratories for X?	Yes	Yes
For course Y		
Who is the target audience for Y?	Yes	Yes
Which educational level does Y target?	Yes	Yes
Who is teaching Y?	Yes	Yes
Which slides are linked to Y?	Yes*	Yes*
Which labs are part of Y?	Yes	Yes
Which courses are similar to Y?	No	Yes
Which are suggested resources for Y?	Yes	Yes
For dataset D		
Which exercises exist for D?	No	No
Which courses use D?	No	No
For material M		
Which courses use a material similar to M?	No	Yes
How much similar is M to another material?	No	Yes
Which topics does M cover?	No	Yes*
Is M open access?	Yes	Yes

Below, we describe the SPARQL queries used to retrieve answers for each competency question. We group the SPARQL queries by thematic as presented in Table 5.

4.3.1. For (Sub)topic X

The competency questions which are referring to topics and subtopics are mapped as skills and knowledge topics in our ontology. For retrieving the information we formulate four SPARQL queries.

a. Who is teaching X? Retrieve all instructors teaching courses that cover the topic X.

```
PREFIX base: <https://w3id.org/coursesonto#>

SELECT ?person
WHERE {
  ?course base:teaches ?topic .
  ?course base:responsibleEntity ?person .
  FILTER(?topic = <X>) # replace X with URI of the given Knowledge Topic
}
```

b. Which are the materials for X? List all educational resources (materials) associated with the (sub)topic X.

```
PREFIX base: <https://w3id.org/coursesonto#>
PREFIX schema: <http://schema.org/>

SELECT DISTINCT ?material
WHERE {
   ?course base:teaches ?topic .
   ?course schema:hasCourseInstance ?courseSession .
   ?coursesSession base:containsResource ?material .

   FILTER(?topic = <X>) # replace X with URI of the given Knowledge Topic
}
```

c. Which are the prerequisites of X? Identify all prerequisites required for the topic X, meaning the skills or knowledge topics that the experts identified as prerequisites.

```
PREFIX base: <https://w3id.org/coursesonto#>
PREFIX educor: <https://github.com/tibonto/educor#>

SELECT DISTINCT ?prerequisite
WHERE {
   ?prerequisite educor:requiresKnowledge ?topic .

   FILTER(?topic = <X>) # replace X with URI of the given Skill
}
```

d. Which are the laboratories for X? Find all laboratories, meaning the practical educational material, associated with the knowledge topic X.

```
PREFIX base: <https://w3id.org/coursesonto#>
PREFIX schema: <http://schema.org/>
PREFIX edutype: <https://w3id.org/coursesonto/kos/educational-type/>
PREFIX skos: <http://www.w3.org/2004/02/skos/core#>

SELECT DISTINCT ?practice
WHERE {

   ?course base:teaches ?topic .
   ?course schema:hasCourseInstance ?courseSession .
   ?coursesSession base:containsResource ?material .
   ?material base:hasType ?practice .
   ?practice skos:hasTopConcept edutype:practice .

   FILTER(?topic = <X>) # replace X with URI of the given Knowledge Topic
}
```

4.3.2. For Course Y

a. Who is the target audience for Y? Retrieve the audience for course Y, such as software engineering master students.

```
PREFIX schema: <http://schema.org/>

SELECT ?audience
WHERE {
  ?course schema:audience ?audience .
  FILTER(?course = <Y>) # replace Y with URI of the given Course
}
```

b. Which educational level does Y target? Determine the educational level targeted by course Y, such as Master.

```
PREFIX schema: <http://schema.org/>

SELECT ?educationalLevel
WHERE {
  ?course schema:educationalLevel ?educationalLevel .
  FILTER(?course = <Y>) # replace Y with URI of the given Course
}
```

c. Who is teaching Y? Identify all instructors teaching course Y.

```
PREFIX schema: <http://schema.org/>
PREFIX base: <https://w3id.org/coursesonto#>

SELECT ?person
WHERE {
  ?course schema:hasCourseInstance ?courseSession .
  ?courseSession base:schema:instructor ?person
  FILTER(?course = <Y>) # replace Y with URI of the given Course
}
```

d. Which slides are linked to Y? made into "Which theory material are linked to Y?" List all theoretical educational material associated with course Y.

```
PREFIX base: <https://w3id.org/coursesonto#>
PREFIX schema: <http://schema.org/>
PREFIX edutype: <https://w3id.org/coursesonto/kos/educational-type/>
PREFIX skos: <http://www.w3.org/2004/02/skos/core#>

SELECT DISTINCT ?theory
WHERE {

   ?course schema:hasCourseInstance ?courseSession .
   ?coursesSession base:containsResource ?material .
   ?material base:hasType ?theory .
   ?theory skos:hasTopConcept edutype:theory .

   FILTER(?course = <Y>) # replace Y with URI of the given Course
}
```

e. Which labs are part of Y? Find all laboratories associated with course Y.

```
PREFIX base: <https://w3id.org/coursesonto#>
PREFIX schema: <http://schema.org/>
PREFIX edutype: <https://w3id.org/coursesonto/kos/educational-type/>
PREFIX skos: <http://www.w3.org/2004/02/skos/core#>

SELECT DISTINCT ?practice
WHERE {

   ?course schema:hasCourseInstance ?courseSession .
   ?coursesSession base:containsResource ?material .
   ?material base:hasType ?practice .
   ?practice skos:hasTopConcept edutype:practice .

   FILTER(?course = <Y>) # replace Y with URI of the given Course
}
```

g. Which are suggested resources for Y? Retrieve all educational resources offered for course Y.

```
PREFIX base: <https://w3id.org/coursesonto#>
PREFIX schema: <http://schema.org/>

SELECT DISTINCT ?material
WHERE {
   ?course schema:hasCourseInstance ?courseSession .
   ?coursesSession base:containsResource ?material .

   FILTER(?course = <Y>) # replace Y with URI of the given Course
}
```

4.3.3. For Material M

d. Is M open access? -> Which M is open access? List all the educational material M that are publicly accessible.

```
PREFIX schema: <http://schema.org/>
PREFIX educor: <https://github.com/tibonto/educor#>

SELECT ?material
WHERE {
  ?material a educor:EducationalResource ;
            schema:isAccessibleForFree true .
}
```

5. Teaching KG

5.1. Pipeline for Creating a Teaching KG

In this section, we present the pipeline used for constructing the KG. The main aim is to develop a sustainable and maintainable workflow that allow incorporating changes at any level of the KG, including ontology changes, new input data sources, removing old data, etc. For this aim, we rely on the RDF Mapping Language (RML) [32], a declarative mapping language that provides the support for transforming any (semi)structured data format into RDF. For the workflow, we use the methodology proposed in [33][7], which extends the Linked Open Terms methodology (LOT) [34] for ontology development and incorporates the creation of KGs and the life cycle of all related artifacts.

---

[7] https://lot.linkeddata.es/LOT4KG/

5.1.1. Knowledge Graph Construction

We transform (semi)structured input data into a RDF-based knowledge graph using RML mappings. To facilitate this task we use first OWL2YARRRML tool[8] to draft the mappings. This tool takes an ontology as input and generates a set of YARRRML [35] templates that can be then filled by the knowledge engineer with the data references. Once the mappings are complete, we translate them into RML using Yatter [36]. Yatter is a Python-based tool that translates YARRRML into a easy to read RML, facilitating its debugging. Finally, the input data is transformed into RDF using Morph-KGC [37]. With this approach, we ensure to be as flexible as possible, as the new specification of RML [32] allows the integration of any kind of semi-structutured format and we also ensure high performance and scalability during the transformation process thanks to Morph-KGC [38].

5.1.2. Life Cycle

As any computational resource, the created KG can suffer changes. The changes can be: i) at the schema level if the ontology evolves into new versions; ii) at the data level, if there are new input sources to be integrated or the ones already transformed into RDF are modified. For mitigate the impact of the first one, we rely on the LOT4KG methodology [33] and the tool OCP2KG [39], that automatically propagate ontology changes over RML mappings. For the second one there is still missing a novel solution that facilitates the integration of new input sources, but the procedure will be similar to the one performed for generating the first version of the KG. Mapping rules will need to be modified manually, adding the new sources and associated rules, and then the KG is regenerated again. For the regenaration, our aim is to use IncRML [40], ensuring that only the new triples are actually generated.

5.2. The System

5.2.1. System Interface Design

Although the Semantic Web experts are the primary users of the Teaching KG, we anticipate potential interest form educators and users with non-technical backgrounds. Therefore, a simple and user-friendly approach of interacting with the KG is essential to allow users to utilize the Teaching KG's content and capabilities effectively. To accomplish this goal, we developed a web-based interface, which supports predefined use cases, designed to allow users to create a complete new course on KGs, and to extend their already existing courses with content from the KG.

The interface design was conducted through an interdisciplinary approach, in which experts in pedagogy, Semantic Web, teaching, and software development collaborated to determine the content and structure of the interface to support a solution towards the main challenges discussed in section 3.2. To that end, we focus on the ease of use, the minimum input requirements from the user, and the alignment with the pedagogical requirements of the course.

Course details and metadata are organized into categories within the interface to enhance the clarity of the required fields, see Figure 2. These include the facilitators information, target audience information, course information, as well as the educational themselves in the third use case. Output data, for the use cases creating or extending a course, are presented in both textual and visual formats, such as tables and network diagrams of retrieved KG content. This dual presentation ensures clarity and usability for users of varying technical expertise.

To meet the requirement of capturing essential metadata for each course, topic, and educational resource without overwhelming educators with excessive input fields, we employ a multi-step strategy. This approach simplifies data entry while encouraging users to provide additional information. The strategy is implemented through the following interface features:

- Mandatory input-fields for the minimal metadata, such as the copyrights of the course. The system will prompt the user to complete those fields before continuing to the steps of storing course content os requesting content from the KG.
- Auto-completion feature for the user input in fields with predefined input, and in those whose test might be similar to the one in the KG. In the former case, metadata fields such as the level of the course can have a set of predefined options including bachelor, master, etc. which allows a quick input from the user. The later case can be seen when adding course topics, where the interface provides suggestions on similar topic titles that exist in the KG.

---

[8] https://github.com/oeg-upm/owl2yarrrml/

- Highlighting the missing information, where the user gets a reminder as the request information from the KG that some fields are left empty. The interface provides a list of those fields, and highlights the role they play in making the information retrieval and recommendations from the KG more accurate.
- Providing a sample of an ideal course which offers the users with a complete structure and content of an exemplary course, showing the input in each of the metadata and information fields.

While using the interface, the user receives the feedback from the system when they request a submission or retrieval of KG data. Except for the mandatory fields, the user is always given the option to submit or request data even if some input fields are empty. At each stage of the interface use, the user has the chance to check an information tab providing guidelines of the use of the interface, the sample idea structure of the course, as well as additional information on the copyright input and requirements.

5.2.2. Similarity Computation

When requesting educational content from the KG, the system retrieves relevant courses and their topics on the users input. Course relevancy includes two aspects:

- Direct relevance by identifying matching courses in the KG, whose titles match one or more of the search terms written by the user.
- Textually semantic relevance by searching for courses, whose titles are semantically similar to the user's search terms. Semantic textual similarity (STS) is calculated using the Sentence Transformers (SBERT)[9] models in Python. We calculate the embeddings of each title and use a cosine similarity between the embedding vectors, to identify similar course titles. This approach is also implemented on the (sub)topic, and the educational resource (material) levels, in a similar manner, to identify (sub)topic and material similarity and answer the competency questions in Table 5 from the interface side.

In order to ensure computational efficiency, STS is pre-calculated among all courses, (sub)topics, and materials only once, as in contract to calculating the similarities real-time at each user request. Those similarities are then stored in a latent space, enabling the search for relevant content to the user query, directly or transitively.

The retrieved data from the search is displayed in a table, including the metadata and structural information on the courses and their educational resources. Users can access these through their openly shared links. Due to copyright restrictions, the educational resources cannot be downloaded directly from the system, but are accessible via their original sources on the web, which the instructors provided when inserting new content in the KG. The interface itself was implemented using Streamlit[10] web-app development framework in Python, and can be accessed at our GitHub.

6. Evaluation

6.1. Validation

We selected OWL as the modeling language for the ontology as it comes with unambiguous semantics and brings reasoning capabilities. Enabling an ontology-driven data creation pipeline will minimise errors in the data. Reasoning capabilities of OWL will ensure that the correctness of the generated data is maintained at least with respect to obvious logical errors.

There are, however, some desired reasoning features that are not captured by OWL given its Open World Assumption nature. For example, integrity constraints used for data validation are not available in OWL. There currently exist a number of efforts to represent data constraint using languages like SHACL[11] and ShEx [41]. Alternatively, as proposed by Kharlamov et al. [42], a subset of the OWL axioms can also be interpreted as integrity constraints. These OWL axioms can be represented as SHACL or ShEx constraints or transformed into (datalog) rules with

---

[9] https://sbert.net/

[10] https://streamlit.io/

[11] https://www.w3.org/TR/shacl/

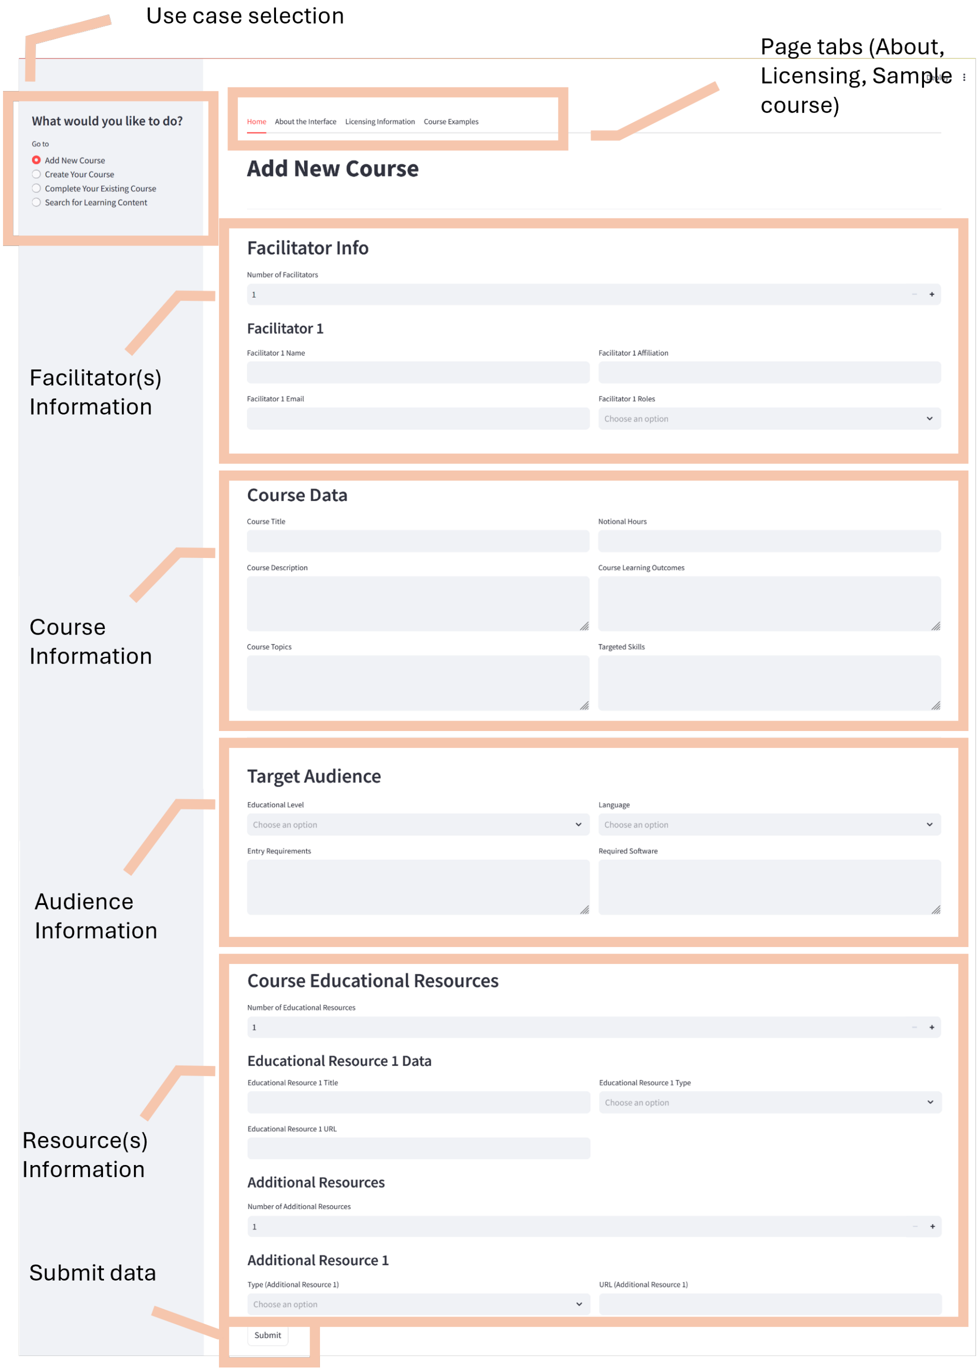

Use case selection

Page tabs (About, Licensing, Sample course)

Facilitator(s) Information

Course Information

Audience Information

Resource(s) Information

Submit data

What would you like to do?

Go to

● Add New Course
○ Create Your Course
○ Complete Your Existing Course
○ Search for Learning Content

Figure 2. Interface structure, tabs, and use cases

(stratified) negation [42].[12] The use of datalog rules to represent integrity constrains enables the use of efficient and

---

[12] Codes to interpret a subset of OWL axioms into datalog rules are available here: https://gitlab.com/ernesto.jimenez.ruiz/ontology-services-toolkit

Table 6

OWL Axioms as Integrity Constraints.

OWL Axiom	**Integrity Constraint as Datalog Rules**
Course **SubClassOf:** license **some** xsd:string	hasLicense(?x) ← license(?x, ?y) MissingLicense(?x) ← Course(?x) ∧ **not** hasLicense(?x)
EducationalResource **SubClassOf:** author **some** Person	hasAuthor(?x) ← author(?x, ?y) ∧ Person(?y) MissingAuthor(?x) ← EducationalResource(?x) ∧ **not** hasAuthor(?x)

Table 7

ShEx and SHACL shapes to identify missing properties.

ShEx

```
<Course> {
  a [ schema:Course ] ;
  schema:license rdf:langString + ;
  # . . .
}
```

SHACL

```
<Course> a sh:NodeShape ;
 sh:targetClass schema:Course ;
 sh:property [
  sh:path schema:license ;
  sh:minCount 1
  ] .
```

ShEx

```
<EducationalResource> {
 a [ educor:EducationalResource ] ;
 schema:author @<Person> ;
 # . . .
}
<Person> {
 # . . .
}
```

SHACL

```
<EducationalResource> a sh:NodeShape ;
 sh:targetClass educor:EducationalResource ;
 sh:property [
  sh:path schema:author ;
  sh:class schema:Person
 ] # . . .
.
<Person> a sh:NodeShape ;
  sh:targetClass schema:Person ;
  # . . .
```

sound datalog engines (e.g., RDFox [43])[13].

For example, Table 6 shows the interpretation of existential restrictions on the right-hand side of OWL subsumption axioms as integrity constraints via datalog rules. The rules highlight missing information in the data (e.g., the license of a course and the author of an educational resource), otherwise assumed to exist somewhere using OWL's open-world semantics. These constraints can also be expressed in ShEx or SHACL as we show in Table 7. It is also possible to define shapes that can be used to explicitly identify nodes with some missing properties which can be used for recommendations as expressed in Table 8.

The use of datalog also enables the addition of rules that can enhance the reasoning beyond OWL reasoning and the integrity constraints. For example, in our setting, in the form of recommendations or *shortcuts* to facilitate the querying. For example, Table 9 shows a series of rules to enable the recommendation of a topic for a course.

6.2. Stakeholders Feedback

To further ensure that the Teaching KG meets the needs of its intended users, we gather insights from stakeholders. This study aimed to investigate the expectations, preferences and concerns of educators and researchers within the semantic web community regarding the development of a KG resource catalog. Through semi-structured interviews with two stakeholders, who agreed to share their feedback anonymously for the purposes of our study, we sought to identify key themes, patterns, and recommendations that would inform the design and implementation of an effective KG catalog. The interview was formulated around questions related to the future use of the teaching KG catalog by stakeholders and consisted of the following questions:

---

[13] The datalog engines should support stratified negation-as-failure.

Table 8

ShEx and SHACL shapes to identify missing properties.

ShEx

```
<MissingLicense> {
  a [ schema:Course ] ;
  schema:license . { 0 }
}
```

SHACL

```
<MissingLicense> a sh:NodeShape ;
 sh:targetClass schema:Course ;
 sh:property [
  sh:path schema:license ;
  sh:maxCount 0
 ] .
```

ShEx

```
<MissingAuthor> {
 a [ educor:EducationalResource ] ;
 schema:author . { 0 }
}
```

SHACL

```
<MissingAuthor> a sh:NodeShape ;
 sh:targetClass educor:EducationalResource ;
 sh:property [
  sh:path schema:author ;
  sh:maxCount 0
 ] .
```

Table 9

Datalog rules for recommendations.

Recommendation	**Datalog Rules**
Recommendation of RDF as topic	hasTopicOWL(?x) ← consistsOfKnowledge(?x, OWL)
	hasTopicRDF(?x) ← consistsOfKnowledge(?x, RDF)
	RecommendRDF(?x) ← Course(?x) ∧ hasTopicOWL(?x) ∧ **not** hasTopicRDF(?x)

Q1 What specific topics or areas within knowledge graphs do you believe should be emphasized in the catalogue to align with your teaching objectives and expertise?
Q2 Can you share insights on how the structure and organization of the catalogue could best support students' learning process and facilitate their exploration of knowledge graph?
Q3 As a main user and stakeholder, what criteria would you use to assess the effectiveness and usefulness of the knowledge graph catalogue in enhancing student engagement and comprehension?
Q4 How do you envision leveraging the catalogue's features, such as interactive visualizations or query functionalities, to enhance students' hands-on learning experiences with knowledge graphs?
Q5 How do you plan to incorporate real-world case studies and examples from the catalogue into your teaching to illustrate the practical applications of knowledge graphs in various domains?
Q6 As lecturer in teaching knowledge graphs, what opportunities do you foresee for collaboration and knowledge sharing among educators within the catalogue's community?
Q7 What specific teaching objectives and expertise in semantic web would you expect to be present in the catalogue?
Q8 Could you share insights on how the catalogue's structure and organization would best be?
Q9 In your opinion, how could it best facilitate students' comprehension and retention of semantic web concepts, based on your teaching experience?
Q10 In your opinion, what types of supplementary resources or materials should be included in the catalogue to complement your teachings and provide comprehensive coverage of semantic web topics?
Q11 As a main user and stakeholder, what metrics or indicators would you consider essential for evaluating the success and effectiveness of the semantic web knowledge catalogue in supporting student learning and engagement?
Q12 Other examples of catalogues or other resources that you are currently using?
Q13 How can the catalogue facilitate interdisciplinary connections?

6.2.1. Stakeholders Key Findings and Insights

From the interview 1, highlighted the importance of creating a KG catalog that is structured and organized to support students' learning process. Participants suggested including labs and datasets that are easy for students to use, such as exercises in link prediction (Q2). Consistency across datasets used in labs was also highlighted as crucial, with suggestions to connect the datasets and present the same resource to students for different practical assessments. Additionally, stakeholders recommended providing pointers to extra supplementary material that is easily accessible, public, online, and well-classified (Q2). Collaboration among instructors within the community was another key theme, with suggestions that open-source learning materials and collaborative development could reduce time spent on preparing a course and finding course material. They also provided insights on how to assess the effectiveness and usefulness of the catalog in enhancing student engagement and comprehension. They suggested using publicly available resources to ensure the discoverability of resources (Q3) and providing user rights and content rights for editing presentations (Q3). Interactive visualizations and query functionalities were seen as opportunities to enhance students' hands-on learning experiences with KGs (Q4). Further, they emphasized the importance of real-world case studies and examples from the catalog, highlighting potential research directions and collaborations (Q6-8). To evaluate the success and effectiveness of the catalog, they suggested using metrics such as metadata, context descriptions, publicly available and editable resources, and connections between lectures and exercises (Q12).

From interview 2, the emphasis was on the importance of including specific topics in the KG catalog that align with their teaching objectives and expertise. They suggested focusing on SPARQL, query optimization, and solutions to problems of working with large and decentralized KGs (Q1). To support students' learning process, they recommended practical tutorials, such as Jupyter notebooks (Q2), and interactive visualizations or query functionalities (Q4) to enhance hands-on learning experiences. They also suggested incorporating real-world case studies and examples from the catalog into teaching (Q5) to illustrate practical applications of KGs in various domains. Consequentially, they provided insights on how the catalog's structure and organization could best support students' comprehension and retention of semantic web concepts. They proposed using collaborative editing tools, such as Semantic Wikis (Q7), and defining templates based on an ontology to add resources and use Semantic MediaWiki extensions for querying and visualization. To evaluate the success and effectiveness of the catalog, they suggested considering metrics such as improvement in student performance in offline tests (Q3) and assessments that test whether students have fulfilled course goals (Q12). They also mentioned the importance of facilitating interdisciplinary connections by linking to other open learning resources on the Web, such as Wikidata (Q13).

6.2.2. Stakeholders' Expectations and Alignment with the Teaching KG

In Table 10, we group the stakeholders' expectations as gathered by the interviews with the experts and we report the alignment with elements present in the Teaching KG. The four common themes in the stakeholder expectations are the: 1) structured and organized catalog, 2) practical learning material, 3) real-world case studies and examples, and 4) collaboration and knowledge sharing.

As it is shown in Table 10, the Teaching KG aligns with the stakeholder expectations in the major common themes. Firstly, the Teaching KG is based on an ontology and semantic rules, which provide a structured and organized framework for representing the concepts and relationships. This satisfies the stakeholders' requirement for a well-organized catalog. Secondly, the Teaching KG includes labs and practical educational material that can be used in assessments and hands-on learning experiences. Additionally, our project includes supplementary educational resources which include practical downstream applications and usage of KG tools in real-world examples. Finally, by adopting an open-source collaborative editing approach, we are creating a community-driven resource that allows KG experts and educators to contribute by sharing their courses, and editing the future versions of the KG. Overall, the Teaching KG fulfills the needs of KG educators, and provides a comprehensive learning framework for KG education concepts.

Table 10

The stakeholders' expectations as grouped and reported by each interview, and the alignment of these elements with the Teaching KG.

Stakeholder Expectations	**Interview 1**	**Interview 2**	**Teaching KG Alignment**
Structured and Organized Catalog	Suggesting consistency across resources so they are easy to find and reuse	Proposing linking to semantic wikis for collaborative editing, defining templates based on ontology	Ontology supporting the KG
Practical Learning Material	Recommending interactive visualizations and query functionalities	Suggesting practical tutorials (e.g., Jupyter notebooks), interactive UI for SPARQL endpoint	Labs and practical educational material
Real-world Case Studies and Examples	Highlighting the importance of real-world case studies and examples from the catalog	Proposing incorporation of real-world case studies and examples into teaching to illustrate practical applications	Supplementary educational material
Collaboration and Knowledge Sharing	Suggesting open-source learning materials and collaborative development	Mentioning potential collaboration among educators within the community and collaborative editing	Open-source collaborative edited resource

7. Use Cases

The Teaching KG is an application driven resource. In order to further validate this and demonstrate its functionality, we describe four use cases. Each use case contains a scenario description, the functionality proposed with the user actions, and the output of the user's actions.

7.1. Use Case I - New Teacher

This use case handles the scenario where a teacher wishes to create a new course on KGs. The teacher is assumed to have no previous content about the content for their course. Therefore, they are looking for an overview of a reasonable sequence of topics to be taught, along with their available educational resources.

7.1.1. Scenario Description

A teacher, Maria, is assigned a teaching responsibility for a course on KGs. She is an experience teacher, with a good understanding of semantic web, ontologies, and KGs, but has not taught this topic before. Maria wishes to start designing her course by surveying existing courses on the same domain, to get insights, inspiration, and relevant content for her course. The course she is designing targets higher education students in the Bachelor level. The course is an introductory course for that target group, which is taught a blended-learning method, in which the access to open educational resources (OERs) is important.

7.1.2. Activities for Creating a New Course

To accomplish her goal, Maria begins by setting goals and objectives of the KG course, as in Figure 3. Then, she accesses the interface of the Teaching KG, to search for relevant content. On the interface, Maria selected the "Create Your Course" option and gets a web-page with clustered fields to insert the information about her targeted course. In the first group of fields, Maria is prompted to insert the main course information, which are the title of the course, course description, and any available keywords. She is also prompted to insert the learning goals of the course within this category.

In the following group of input fields, Maria enters information about the target audience of the course, and thus selected Bachelor from the drop-down list of options. Maria also selects an additional option (Master) in from the same list, to account for a potential future extension of the course. For the target audience, she also sets

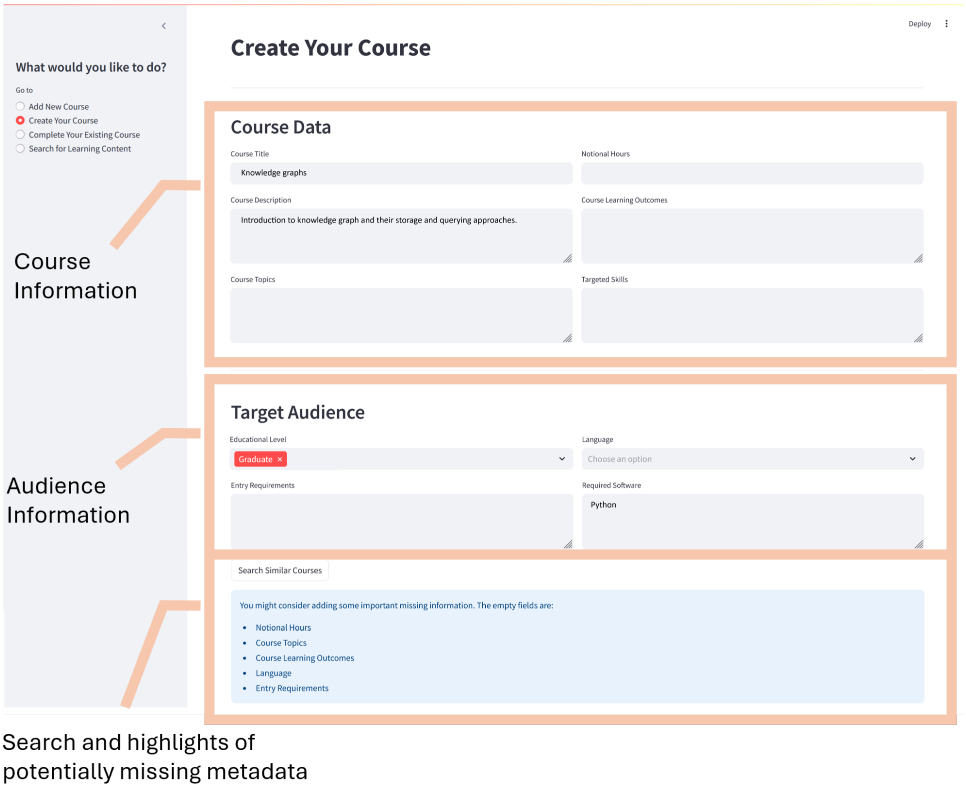

Course Title: Knowledge graphs

Course Description: Introduction to knowledge graph and their storage and querying approaches.

Educational Level: Graduate

Required Software: Python

You might consider adding some important missing information. The empty fields are:

- Notional Hours
- Course Topics
- Course Learning Outcomes
- Language
- Entry Requirements

Course Information

Audience Information

Search and highlights of potentially missing metadata

Figure 3. Interface structure for creating new course.

the prerequisites of the course she intends to teach. Here, Maria writes a list including the basic knowledge of databases, and basic knowledge of text mining. She also add Python knowledge as a software requirement expected from students for entering the course.

Maria then clicks on the button "Find relevant content" to begin the search in the KG. She receives a notification on the page stating that she left the "Keywords" field empty. In the notification, the system informed Maria that keywords play an important role in identified thematically similar content in a more specific way especially from courses that have matching keywords. Maria then adds "RDF" and "Ontology" as keywords in the respective field and clicks on the search button again.

The system receives the search request and scans the KG for relevant content, based on the title, description, learning goals and matching keywords of the course. The system also sets search constraints, which are concluded from the target audience information, and thus retrieves only courses for the Bachelor and Master levels. The system also excludes the topics on databases, text mining and Python software, as they were set as prerequisites for the course, and therefore will not be taught.

Search results are given to the teacher in two formats:

1) a table: in which the relevant course titles, metadata, topics, and a list of educational resources and their links are provided.
2) a network visualization: of the parts of the KG that were retrieved by the search function. This partial KG visualization shoes the hierarchical structure of the relevant courses, as well as the connections among courses and educational resources, when available.

The table enables Maria of clicking on OER links on topics she finds interesting and potentially suitable for her course. From the KG visualization, Maria notices that one OER is used by two of the retrieved curses, which she considers as an level of credibility of that OER.

7.2. Use Case II - Missing Content

An essential piece of information that is missing from many course descriptions is copyright information. While most course design approaches can be adapted to suit the context of implementation, copyright information has both legal implications and implications for developing and using the KG.

7.2.1. Scenario Description

The second use case focuses then on the completion of missing content and metadata from a preliminary course structure. Here, the teacher Maria wishes to extend the content she already prepared for a course on SPARQL. While the introductory parts of the course are ready, more advanced topics will be useful to her target students in the Master's degree of computer science.

The course that Maria is constructing will be hosted on the university's website as well as the local Moodle instance. While student interaction with the course content will take place on Moodle, the links to the course materials and the recorded video lectures will be available online. Therefore, the teacher aims to verify if all essential course metadata are covered for her course page.

7.2.2. Activities for Content Extension and Completion

To accomplish this goal, Maria uses the interface to access additional content from the KG, as shown in Figure 4. She selected the option "Complete your existing Course" and gets a set of fields to fill about the course itself and the existing content she already prepared.

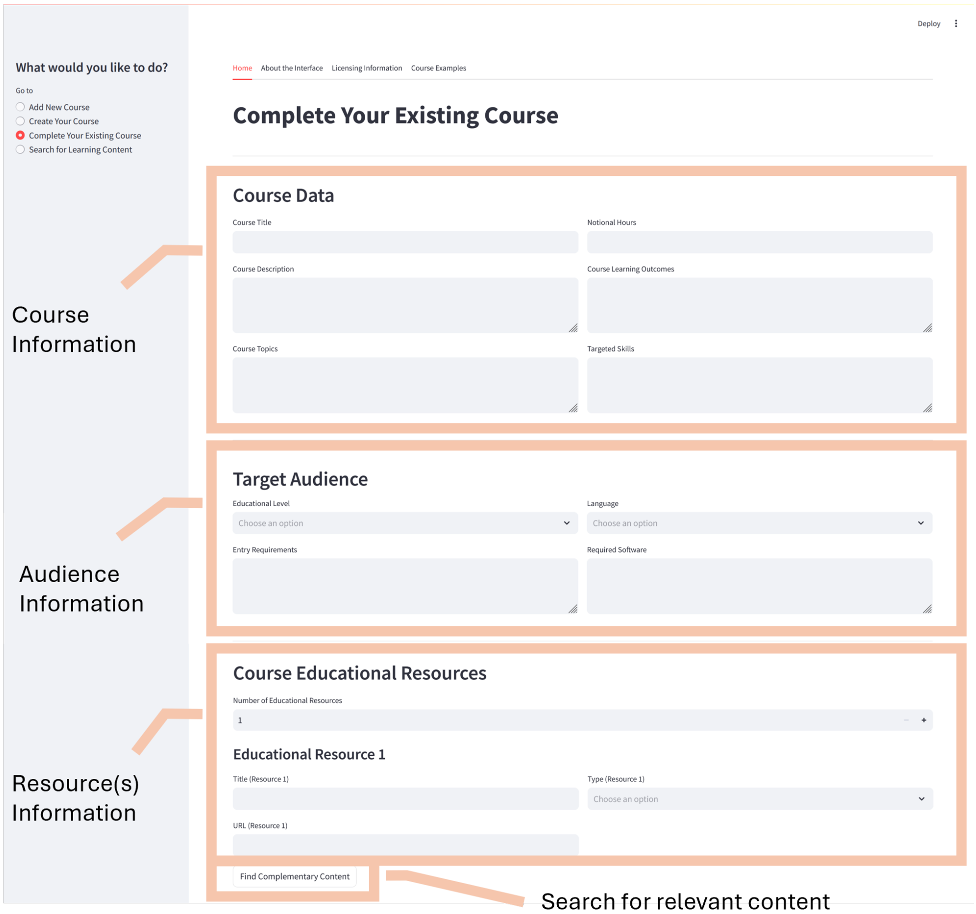

Figure 4. Interface structure for completing an existing course.

She begins by filling the information about course facilitators, where she provides the name, affiliation, roles and the contact information about herself, and any other teaching assistant. In this field, Maria notices that a contact person must be assigned to the course in the online version. After the course facilitators, the course information need to be provided. She adds the title, description, and learning outcomes of her course. Then, she enters the number of educational resource she already prepares. On the interface, a corresponding number of fields is created dynamically, in order to add the information about each educational resource. Here, the title, description, type, and URL of each educational resource are to be provided. Maria fills the information and clicks on the submission button to get additional content.

By analyzing the data that the teacher entered, the system identifies that copyrights of the course are not provided in the course metadata, and Maria receives a notification that the copyright field is empty. Furthermore, the message informs Maria about the importance of this field for courses that are hosted in the KG, or on other openly accessible repositories. The message guides Maria to the interface tab, where information about different copyright licenses is provided, along with links to more detailed resources on selecting the suitable copyright license for the course.

Based on this information, Maria reads the provided information and makes a selection of a license that suits her course and university, fills the empty fields and renews the submission of the information to receive the additional content she is looking for. The systems searches the KG, based on the course title, description, and the topics covered by the available educational resources, and retrieves relevant content, which other courses that covered SPARQL are also offering. Information on the educational resource are retrieved in a table, which includes the titles, descriptions and URLs of the educational resources, as well as the courses to which they belong in the KG.

7.3. Use Case III - Finding Educational Resources

The third use case covers the scenario where a user seeks specific learning materials for a topic of interest, such as preparing for an exam, completing an assignment, or exploring a subject in depth. The user can use the interface to find curated content in the KG, including Open Educational Resources (OERs), and related topics and courses, all contextualized within the KG.

7.3.1. Scenario Description

Alex, a computer science student at the bachelor's level, is preparing for an exam on KGs. Alex realizes that he needs to strengthen his understanding of SPARQL queries, a topic covered briefly in his course. While the lecture slides and notes provide a basic introduction, Alex is searching for detailed explanations, practice exercises, and possibly video tutorials to enhance his understanding of the topic. Since Alex prefers self-paced learning and aims for high-quality sources, he turns to the KG interface to explore relevant resources.

7.3.2. Activities for Content Extension and Completion

Alex accesses the KG interface and selects the "Search for Learning Content" option, see Figure 5. A search field and filter options are presented to guide the search process. In the search field, Alex writes the terms "SPARQL", "query examples" and "querying KGs". The feature of adding multiple terms in the same search is offered to the student to get more information on what they are looking for, and to reduce the amount of searches done for one topic. Alex also chooses to filter the results to keep the search limited to topic he can directly use. The interface offers a set of filters categorized based on the course, topic and facilitator's metadata. Here, Alex selects courses that are for the bachelor knowledge level, and those that are in English language, Moreover, Alex also tries to filter based on the existence of "additional resources" that have "external links" and are "OERs".

By clicking the "Search for content" button, the system analyzes the search terms and the selected filters, then queries the KG to find relevant sources. The search results are presented in two formats:

1) Table View: which lists the titles, descriptions, types, and direct URLs of the retrieved resources. Each entry also includes metadata such as the author, and the course to which it belongs.
2) Graph Visualization: of the retrieved content shows how resources are interconnected, highlighting their relationships to other topics and courses in the KG.

The table enables Alex to quickly review and access specific resources, while the visualization helps him understand the broader context of the topic and discover related concepts they might explore next. Through the interface, Alex not only finds the immediate resources he needs but also builds a pathway for continued exploration of the topic, utilizing the comprehensive and interconnected nature of the KG.

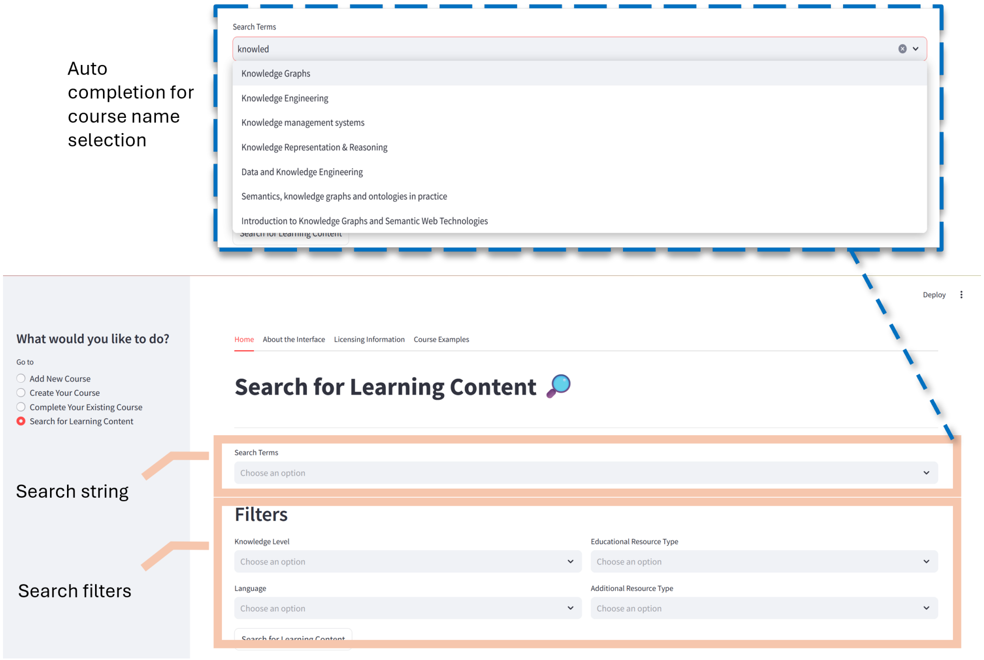

Search Terms: knowled

Go to
○ Add New Course
○ Create Your Course
○ Complete Your Existing Course
● Search for Learning Content

Figure 5. Interface components for searching for KG content.

8. Potential Impact, Limitations and Future Work

We presented the project plan and motivation on November 2024 in the 23rd International Semantic Web Conference, ISWC 2024, as part of the poster presentations [1]. The response the project received was overwhelmingly positive and highlighted the significant interest of a Teaching KG for KG courses from the Semantic Web community. More specifically, 10 new lectures stepped forward to contribute their courses to our project, and many researchers expressed eagerness in using the KG once it is published. In addition, the discussions revealed a shared vision for exploring new research directions related to the Teaching KG. The feedback we received from ISWC 2024 revealed the prospect of building on this momentum to further develop and enhance the project with new features and functionality.

Despite its potential, our study has several limitations. Firstly, not all topics and skills relevant to KG education are currently described within the KG, resulting in gaps in coverage. Our attempt to align our topics and skills with the multilingual classification of European Skills, Competences, and Occupations, the ESCO classification[14] proved challenging due to its limited coverage of KG-specific topics, which restricts our ability to capture a broader range of pedagogical content. Furthermore, the scope of educationally relevant data encoded within the skills and topics is limited to pre-encoded information, which does not cover all aspects of teaching that are interesting to pedagogical researchers. Additionally, the KG and ontology lack hierarchy and prerequisite relationships between skills, topics, and content, making it difficult to facilitate a nuanced understanding of complex concepts. Our analysis also reveals that many courses and educational resources are not open-source or accessible, limiting our ability to properly classify and interlink them. Finally, the variety of formats in which educational resources are presented, such as pdf

---

[14] The ESCO classification can be accessed at https://esco.ec.europa.eu/en

slides and Jupiter notebooks, poses challenges in integrating them into a unified multi-modal KG, which highlights the need for further standardization and research development.

This project defines the semantic structure for the deployment of KGs in education, which opens new opportunities for neurosymbolic AI applications in education [44]. In future work, we aim to expand the capabilities of the Teaching KG by implementing assessment analysis. In this regard, difficulty prediction mechanisms could be utilized to better estimate the complexity of topics and assessments, tailoring course materials to the user's level. An extension could involve creating personalized student KGs to generate individualized learning paths based on assessments and background knowledge, which are currently only minimally covered in the audience metadata. The semantic student model could be developed in collaboration with technology-enhanced learning experts to provide personalized recommendations and smart learning analytics aligned with each student's learning preferences and goals. Furthermore, based on the feedback we received, a potential future step is the integration of a question-answering system to support users in real-time. This would be particularly beneficial for new learners in the Semantic Web who are not yet familiar with SPARQL queries. Additionally, one of our goals is to implement automatic semantic similarity detection to analyze new input sources and recommend related content based on existing resources and knowledge topics in the KG. This approach could further enable dynamic content updates and uncover hidden connections across resources.

9. Conclusion

We presented the Teaching KG for KG education which contains skills, knowledge topics, courses, instructors and KG course material. Our approach is grounded in a robust methodology for constructing pedagogically enriched educational knowledge graphs, leveraging a higher ontology and semantic constraints to ensure consistency and reusability. We presented our preliminary evaluation and use cases, which demonstrated the potential impact of this project. These results underscore to be a valuable resource for Semantic Web experts, and further underline its capacity to enable smart education applications. Moving forward, we are committed to the continuous development and maintenance of the Teaching KG, aiming to expand its utility and support for diverse educational contexts.

Acknowledgments

Eleni Ilkou would like to thank Tobias Käfer for his contribution in gathering material and classification of skills. David Chaves-Fraga is supported from the Agencia Estatal de Investigación - Spain (PID2023-149549NB-I00), the Xunta de Galicia - Consellería de Cultura, Educación, Formación Profesional e Universidades (Centro de investigación de Galicia accreditation 2024-2027 ED431G-2023/04 and Reference Competitive Group accreditation 2022-2025, ED431C2022/19) and the European Union (European Regional Development Fund - ERDF). Jose Emilio Labra Gayo has been partially funded by the Project ANGLIRU: ANGLIRU: Applying kNowledge Graphs to research data interoperabiLIty and ReUsability, code: PID2020-117912RB. Erna Engelbrecht acknowledges financial support from the 4TU.Centre for Engineering Education (4TU.CEE), which funded the contribution to this research. Ernesto Jimenez-Ruiz would like to thank the organisers and contributor of the Alan Turing Institute Interest Group on Knowledge Graphs, as well as The Turing for supporting our activities.

References

[1] Eleni Ilkou and Ernesto Jimenez. Towards a knowledge graph for teaching knowledge graphs. In *ISWC (Posters/Demos/Industry)*, 2024.

[2] Kechen Qu, Kam Cheong Li, Billy TM Wong, Manfred MF Wu, and Mengjin Liu. A survey of knowledge graph approaches and applications in education. *Electronics*, 13(13):2537, 2024.

[3] Hadhemi Achour and Maroua Zouari. Multilingual learning objects indexing and retrieving based on ontologies. In *2013 World Congress on Computer and Information Technology (WCCIT)*, pages 1–6. IEEE, 2013.

[4] Yangrui Yang, Sisi Chen, Yaping Zhu, Hao Zhu, and Zhigang Chen. Knowledge graph empowerment from knowledge learning to graduation requirements achievement. *Plos one*, 18(10):e0292903, 2023.

[5] Björn Schembera, Frank Wübbeling, Thomas Koprucki, Christine Biedinger, Marco Reidelbach, Burkhard Schmidt, Dominik Göddeke, and Jochen Fiedler. Building ontologies and knowledge graphs for mathematics and its applications. In *Proceedings of the Conference on Research Data Infrastructure*, volume 1, 2023.

[6] Maiga Chang, Giuseppe D'Aniello, Matteo Gaeta, Francesco Orciuoli, Demetrios Sampson, and Carmine Simonelli. Building ontology-driven tutoring models for intelligent tutoring systems using data mining. *IEEE Access*, 8:48151–48162, 2020.

[7] Nan Li, Qiang Shen, Rui Song, Yang Chi, and Hao Xu. Medukg: a deep-learning-based approach for multi-modal educational knowledge graph construction. *Information*, 13(2):91, 2022.

[8] Kristian Stancin, Patrizia Poscic, and Danijela Jaksic. Ontologies in education–state of the art. *Education and Information Technologies*, 25(6):5301–5320, 2020.

[9] Lillian N Cassel, Gordon Davies, Richard LeBlanc, Lawrence Snyder, and Heikki Topi. Using a computing ontology as a foundation for curriculum development. In *Proceedings of the sixth international workshop on ontologies & semantic web for E-learning*, pages 21–29, 2008.

[10] Yu-Liang Chi. Ontology-based curriculum content sequencing system with semantic rules. *Expert Systems with Applications*, 36(4): 7838–7847, 2009.

[11] Olivier Palombi, Fabrice Jouanot, Nafissetou Nziengam, Behrooz Omidvar-Tehrani, Marie-Christine Rousset, and Adam Sanchez. Ontosides: Ontology-based student progress monitoring on the national evaluation system of french medical schools. *Artificial intelligence in medicine*, 96:59–67, 2019.

[12] Hyunsook Chung and Jeongmin Kim. An ontological approach for semantic modeling of curriculum and syllabus in higher education. *International Journal of Information and Education Technology*, 6(5):365, 2016.

[13] Evangelos Katis, Haridimos Kondylakis, Giannis Agathangelos, and Kostas Vassilakis. Developing an ontology for curriculum and syllabus. In *The Semantic Web: ESWC 2018 Satellite Events: ESWC 2018 Satellite Events, Heraklion, Crete, Greece, June 3-7, 2018, Revised Selected Papers 15*, pages 55–59. Springer, 2018.

[14] Muthana Zouri and Alex Ferworn. An ontology-based approach for curriculum mapping in higher education. In *2021 IEEE 11th Annual Computing and communication workshop and conference (CCWC)*, pages 0141–0147. IEEE, 2021.

[15] Stathis Th Konstantinidis, Lazaros Ioannidis, Dimitris Spachos, Charalampos Bratsas, and Panagiotis D Bamidis. meducator 3.0: combining semantic and social web approaches in sharing and retrieving medical education resources. In *2012 Seventh International Workshop on Semantic and Social Media Adaptation and Personalization*, pages 42–47. IEEE, 2012.

[16] Dmitry Mouromtsev, Fedor Kozlov, Olga Parkhimovich, and Maria Zelenina. Development of an ontology-based e-learning system. In *Knowledge Engineering and the Semantic Web: 4th International Conference, KESW 2013, St. Petersburg, Russia, October 7-9, 2013. Proceedings 4*, pages 273–280. Springer, 2013.

[17] Yousra Fettach, Mounir Ghogho, and Boualem Benatallah. Knowledge graphs in education and employability: A survey on applications and techniques. *IEEE Access*, 10:80174–80183, 2022.

[18] Y. Li et al. Multi-source education knowledge graph construction and fusion for college curricula. *arXiv preprint arXiv:2305.04567*, 2023. URL https://arxiv.org/abs/2305.04567.

[19] Nicolas Hubert, Armelle Brun, and Davy Monticolo. Educonto: An ontology and knowledge graph for higher education, October 2022. URL https://hal.archives-ouvertes.fr/hal-03768154.

[20] Hasan Abu-Rasheed, Mareike Dornhöfer, Christian Weber, Gábor Kismihók, Ulrike Buchmann, and Madjid Fathi. Building Contextual Knowledge Graphs for Personalized Learning Recommendations Using Text Mining and Semantic Graph Completion. In *2023 IEEE International Conference on Advanced Learning Technologies (ICALT)*, pages 36–40, Orem, UT, USA, July 2023. IEEE. ISBN 979-8-3503-0054-3. . URL https://ieeexplore.ieee.org/document/10260850/.

[21] Guosheng Hao, Tingting Zhang, and Xia Wang. Description of Personalized Learning Characteristics in Education Oriented Knowledge Graph. In *2021 4th International Conference on Information Systems and Computer Aided Education*, pages 2495–2498, Dalian China, September 2021. ACM. ISBN 978-1-4503-9025-5. . URL https://dl.acm.org/doi/10.1145/3482632.3487457.

[22] Penghe Chen, Yu Lu, Vincent W. Zheng, Xiyang Chen, and Boda Yang. KnowEdu: A System to Construct Knowledge Graph for Education. *IEEE Access*, 6:31553–31563, 2018. ISSN 2169-3536. . URL https://ieeexplore.ieee.org/abstract/document/8362657. Conference Name: IEEE Access.

[23] Bowen Zhao, Jiuding Sun, Bin Xu, Xingyu Lu, Yuchen Li, Jifan Yu, Minghui Liu, Tingjian Zhang, Qiuyang Chen, Hanming Li, Lei Hou, and Juanzi Li. EDUKG: a Heterogeneous Sustainable K-12 Educational Knowledge Graph, October 2022. URL http://arxiv.org/abs/2210.12228. arXiv:2210.12228 [cs].

[24] Wei Zheng, Zhichun Wang, Mingchen Sun, Yanrong Wu, and Kaiman Li. Building a Large-Scale Knowledge Graph for Elementary Education in China. In Xin Wang, Francesca A. Lisi, Guohui Xiao, and Elena Botoeva, editors, *Semantic Technology*, pages 1–12, Singapore, 2020. Springer. ISBN 978-981-15-3412-6. .

[25] Haiyang Yu, Entai Wang, Qi Lang, and Jianan Wang. Intelligent Retrieval and Comprehension of Entrepreneurship Education Resources Based on Semantic Summarization of Knowledge Graphs. *IEEE Transactions on Learning Technologies*, 17:1210–1221, 2024. ISSN 1939-1382. . URL https://ieeexplore.ieee.org/abstract/document/10430464?casa_token=FbJND_n3nccAAAAA:FdCIBfBpiwSQSiCkzfPF3SppfwAFcE4FAhpE--nicEtCfoA8MfNlusTWWojw_4ZQQ-Yrd1fO4Gk. Conference Name: IEEE Transactions on Learning Technologies.

[26] Yaning Hou, Bo Liu, Qianxi Fan, and Jianqiang Zhou. Research on the application mode of knowledge graph in education. In *Proceedings of the 2023 6th International Conference on Educational Technology Management*, pages 215–220, Guangzhou China, November 2023. ACM. ISBN 979-8-4007-1667-6. . URL https://dl.acm.org/doi/10.1145/3637907.3637976.

[27] K. Ng and A. Fung. Educational personalized learning path planning with large language models. *arXiv preprint arXiv:2407.11773*, 2024. URL https://arxiv.org/abs/2407.11773v1.

[28] María Poveda-Villalón, Alba Fernández-Izquierdo, Mariano Fernández-López, and Raúl García-Castro. LOT: an industrial oriented ontology engineering framework. *Eng. Appl. Artif. Intell.*, 111:104755, 2022. . URL https://doi.org/10.1016/j.engappai.2022.104755.

[29] Camila Zachhé Aguiar and Vítor E Silva Souza. Sabiox: the extended systematic approach for building ontologies. 2024.

[30] Eleni Ilkou, Hasan Abu-Rasheed, Mohammadreza Tavakoli, Sherzod Hakimov, Gábor Kismihók, Sören Auer, and Wolfgang Nejdl. Educor: An educational and career-oriented recommendation ontology. In *International Semantic Web Conference*, pages 546–562. Springer, 2021.

[31] Serge Chávez-Feria, Raúl García-Castro, and María Poveda-Villalón. Chowlk: from uml-based ontology conceptualizations to owl. In *European Semantic Web Conference*, pages 338–352. Springer, 2022.

[32] Ana Iglesias-Molina, Dylan Van Assche, Julián Arenas-Guerrero, Ben De Meester, Christophe Debruyne, Samaneh Jozashoori, Pano Maria, Franck Michel, David Chaves-Fraga, and Anastasia Dimou. The rml ontology: A community-driven modular redesign after a decade of experience in mapping heterogeneous data to rdf. In *International Semantic Web Conference*, pages 152–175. Springer, 2023.

[33] Romana Pernisch, Marıa Poveda-Villalón, Diego Conde-Herreros, David Chaves-Fraga, and Lise Stork. When ontologies met knowledge graphs: Tale of a methodology. *The Semantic Web: ESWC*, 2024.

[34] María Poveda-Villalón, Alba Fernández-Izquierdo, Mariano Fernández-López, and Raúl García-Castro. Lot: An industrial oriented ontology engineering framework. *Engineering Applications of Artificial Intelligence*, 111:104755, 2022.

[35] Pieter Heyvaert, Ben De Meester, Anastasia Dimou, and Ruben Verborgh. Declarative rules for linked data generation at your fingertips! In *The Semantic Web: ESWC 2018 Satellite Events: ESWC 2018 Satellite Events, Heraklion, Crete, Greece, June 3-7, 2018, Revised Selected Papers 15*, pages 213–217. Springer, 2018.

[36] Ana Iglesias-Molina, David Chaves-Fraga, Ioannis Dasoulas, and Anastasia Dimou. Human-friendly rdf graph construction: Which one do you chose? In *International Conference on Web Engineering*, pages 262–277. Springer, 2023.

[37] Julián Arenas-Guerrero, David Chaves-Fraga, Jhon Toledo, María S Pérez, and Oscar Corcho. Morph-kgc: Scalable knowledge graph materialization with mapping partitions. *Semantic Web*, (Preprint):1–20, 2024.

[38] Dylan Van Assche, David Chaves-Fraga, and Anastasia Dimou. Krown: A benchmark for rdf graph materialisation. In *International Semantic Web Conference*, pages 20–39. Springer, 2024.

[39] Diego Conde-Herreros, Lise Stork, Romana Pernisch, María Poveda-Villalón, Oscar Corcho, and David Chaves-Fraga. Propagating Ontology Changes to Declarative Mappings in Construction of Knowledge Graphs. In *Proceedings of the 5th International Workshop on Knowledge Graph Construction*, volume 3718. CEUR Workshop Proceedings, 2024.

[40] Dylan Van Assche, J Rojas, Ben De Meester, and Pieter Colpaert. Incrml: incremental knowledge graph construction from heterogeneous data sources. *Semantic Web*, (Under Review), 2024.

[41] Eric Prud'hommeaux, Jose Emilio Labra, and Harold Solbrig. Shape expressions: An RDF validation and transformation language. In *10th International Conference on Semantic Systems*, Sept. 2014.

[42] Evgeny Kharlamov, Bernardo Cuenca Grau, Ernesto Jiménez-Ruiz, Steffen Lamparter, Gulnar Mehdi, Martin Ringsquandl, Yavor Nenov, Stephan Grimm, Mikhail Roshchin, and Ian Horrocks. Capturing industrial information models with ontologies and constraints. In *The Semantic Web - ISWC 2016 - 15th International Semantic Web Conference*, volume 9982 of *Lecture Notes in Computer Science*, pages 325–343, 2016. . URL https://doi.org/10.1007/978-3-319-46547-0_30.

[43] Yavor Nenov, Robert Piro, Boris Motik, Ian Horrocks, Zhe Wu, and Jay Banerjee. Rdfox: A highly-scalable RDF store. In *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part II*, volume 9367 of *Lecture Notes in Computer Science*, pages 3–20. Springer, 2015. . URL https://doi.org/10.1007/978-3-319-25010-6_1.

[44] Chris Davis Jaldi, Eleni Ilkou, Noah Schroeder, and Cogan Shimizu. Education in the era of neurosymbolic ai. *arXiv preprint arXiv:2411.12763*, 2024.

[45] Bowen Zhao, Jiuding Sun, Bin Xu, Xingyu Lu, Yuchen Li, Jifan Yu, Minghui Liu, Tingjian Zhang, Qiuyang Chen, Hanming Li, et al. Edukg: a heterogeneous sustainable k-12 educational knowledge graph. *arXiv preprint arXiv:2210.12228*, 2022.

[46] Steffen Lohmann, Vincent Link, Eduard Marbach, and Stefan Negru. Webvowl: Web-based visualization of ontologies. In *Knowledge Engineering and Knowledge Management: EKAW 2014 Satellite Events, VISUAL, EKM1, and ARCOE-Logic, Linköping, Sweden, November 24-28, 2014. Revised Selected Papers. 19*, pages 154–158. Springer, 2015.

[47] Eleni Ilkou and Beat Signer. A technology-enhanced smart learning environment based on the combination of knowledge graphs and learning paths. In *CSEDU (2)*, pages 461–468, 2020.

[48] Eleni Ilkou, Martina Galletti, and Daniil Dobriy. Edumultikg attains 92% accuracy in k-12 user profiling! In *Proceedings of the ESWC*, volume 2043, 2023.

[49] Paolo Pareti, George Konstantinidis, Timothy J. Norman, and Murat Sensoy. SHACL constraints with inference rules. In *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference, Auckland, New Zealand, October 26-30, 2019, Proceedings, Part I*, volume 11778 of *Lecture Notes in Computer Science*, pages 539–557. Springer, 2019. . URL https://doi.org/10.1007/978-3-030-30793-6_31.

[50] Slawek Staworko, Iovka Boneva, José Emilio Labra Gayo, Samuel Hym, Eric G. Prud'hommeaux, and Harold R. Solbrig. Complexity and expressiveness of shex for RDF. In Marcelo Arenas and Martín Ugarte, editors, *18th International Conference on Database Theory, ICDT 2015, March 23-27, 2015, Brussels, Belgium*, volume 31 of *LIPIcs*, pages 195–211. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. . URL https://doi.org/10.4230/LIPIcs.ICDT.2015.195.