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] and Jose Emilio Labra-Gayo [f]

[a] *TIB - Leibniz Information Center for Science and Technology, Hannover, Germany*
E-mail: eleni.ilkou@tib.eu
[b] *Studiumdigitale, Goethe University Frankfurt, Frankfurt am Main, Germany*
E-mail: rasheed@sd.uni-frankfurt.de
[c] *CiTIUS, Universidade de Santiago de Compostela, Santiago de Compostela, 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@citystgeorges.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 is critical for sustaining the growth and utility of the KG ecosystem. This paper introduces the Teaching KG, a novel educational resource catalogue 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 that supports diverse educational needs. By leveraging a high-level 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 finding similar courses, 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/Teaching-Knowledge-Graph/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 increasingly adopted in educational applications [2]. By providing a structured and interconnected framework for representing domain

*Corresponding author. Besides the first author, the order is alphabetical. E-mail: eleni.ilkou@tib.eu.

1570-0844/$35.00 © 0 – IOS Press. All rights reserved.

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 materials, 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. Furthermore, there is a notable scarcity of high-quality open educational resources (OER) in the field of KG education. Open Educational Resources (OER) are educational materials that reside in the public domain or are released under an open license, permitting free access and use with minimal or no restrictions [8]. Teaching materials within this context may include lecture slides, tutorials, datasets, exercises, assignments, and case studies that support the understanding of KG concepts, tools, and applications. Such resources play a vital role in enhancing the quality of educational provision.

In university-level Knowledge Graph (KG) education, there is currently no centralized repository that aggregates teaching resources for KG courses. Identifying relevant materials is often difficult because such courses are offered under diverse titles, such as "Data Governance", "Web AI", or "Information Management", that do not share common keywords. The varied backgrounds and research interests of instructors further contribute to inconsistencies in the coverage of key knowledge areas and skills, resulting in a fragmented educational landscape. Establishing a central hub for KG teaching resources could address these challenges by improving the discovery, retrieval, and integration of materials. Such a platform would enable instructors to access high-quality resources and metadata to enrich their courses, while providing students with targeted learning materials aligned with their educational goals, whether within formal curricula or through independent study.

To address these issues, we propose the use of an educational Teaching Knowledge Graph (Teaching KG) focused on knowledge topics (concrete units that can contain theoretical and practical educational resources), skills (abstract concepts consisting of multiple knowledge topics in different levels and scopes), courses, and educational material, such as slides and tests, 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. While the primary target audience of the Teaching KG is teachers and course instructors, it also benefits students directly by interlinking educational resources. Indirectly, it provides a pedagogically informed framework that helps students better understand whether a course aligns with their learning needs and interests.

From a pedagogical perspective, the Teaching KG supports instructors in aligning learning objectives with teaching materials and assessment-related resources. This alignment enables more coherent course structures and im-

[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

proved discoverability of supplementary materials that support target competencies. As a result, learners benefit indirectly through more consistent curricula and better-aligned instructional resources.

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 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 modelling have been proposed in the literature [9]. While some approaches developed an ontology for describing a specific curriculum [10], most research focused on a high-level description of main concepts. For instance, Chi [11] 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 [12], syllabi [13, 14] or, at institutional level the creation of program credentials [15]. A series of ontology-based applications that utilized Semantic Web tools were developed by mEducator [16, 17]. The proposed approach aimed to share and reuse medical educational content across institutions using Linked Data principles. Their proposed framework connects lectures with testing material; however, it does not reference the corresponding knowledge topics, skills, or lecturers. Although these initiatives set the tone for Semantic Web applications in the e-learning domain, the projects were not maintained and the resources were not shared as open source, limiting their reuse. In contrast, the ontology that serves as the basis of our KG is not domain or use case specific and is published under a persistent URI.

Recently, the deployment of KGs for educational applications has gained traction. Based on a recent survey by Fettach et al. [18], KG applications in education can be categorized into three clusters: assisted instruction, including institutional concept [19] and knowledge management [20], 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. Chen et al. [19] 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 labelling is used for instructional concept extraction and probabilistic association rule mining to establish meaningful educational relationships. It is worth acknowledging that their work represents a pioneering approach in educational knowledge graphs. Consequently, subsequent research has expanded toward richer relations and content, with a stronger focus on practical use cases, enabling educational knowledge graphs to function as useful systems rather than merely collections of triples. Similarly to KnowEdu, EDUKG [21] 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. The EDUKG ontology is not published, and the work primarily focuses on Chinese texts, some drawn from educational resources but largely sourced from external materials, such as The New York Times. In contrast, our approach relies exclusively on data collected from domain experts, ensuring high data quality, and employs a domain-independent ontology that is publicly available under a persistent URI. In primary schooling, 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 Concept Graph-based educational tools [22]. Their approach focuses on the hierarchical relations of knowledge units

and concepts, and models textbooks, MOOCs and exams. Inspired by their approach, we include extra curricular material and further includes the skills, courses, and lecturers.

Furthermore, towards assisted learning, recent studies have demonstrated the use of KGs in representing and managing knowledge in higher education. Li et al. [23] 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. Our approach also models the courses, knowledge topics and content and further connects topics, skills and educators. Similarly, Hubert et al. introduced EducOnto, an ontology for modelling university curricula and student profiles, which serves as the foundation for their KG called EduKG [24]. Although their approach is useful for modelling connections between university curriculum and high school major, it lacks support for teaching as it does not connect to any courses, topics, or lecturers. Abu-Rasheed et al. [25] 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. Inspired by their approach, we introduce pedagogical elements to achieve learning and pedagogical alignment and further extend on interconnecting skills, topics, courses, lecturers, and material. The Education-Oriented Knowledge Graph (EOKG) [26] quantified personalized learning by defining boundaries of mastered and unmastered knowledge points. This approach enables dynamic modelling of student learning profiles, offering tailored educational experiences. However, their work is limited to mapping knowledge points in the EOKG, while our approach models a broader and more complete range of elements essential to both learning and teaching. Overall, none of these educational KGs bridges the gap of containing and interconnecting essential resources for teaching, namely the skills and topics, courses, lecturers and educational material in English from official resources. Our approach addresses this limitation by integrating all of these elements, and further enriches them with expert-provided descriptions of skills, topics, and courses from the represented domain.

Further, semantic summarization frameworks have been developed to support intelligent retrieval and comprehension of educational resources. For instance, Yu et al. [27] 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 [28]. Inspired by those works, we incorporate a semantic search functionality in our platform to accommodate the retrieval of similar resources.

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. [23] 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 [23, 29]. Therefore, our proposed ontology KG construction approach aims 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 the 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 creating an educational teaching KG and a populated example of a teaching KG for KGs 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, such as the "Knowledge Graphs for RAG" on Coursera, these do not cover the foundational aspects of KGs or align with the content taught in university courses by the Semantic Web community as they emphasize on the application of KGs and not on the fundamentals of KGs education, such as OWL or SPARQL. Furthermore, based on informal feedback we received from the Data Science students at City St. George's, the students face difficulties in finding alternative code examples and supplementary KG materials to complement their studies. In contrast, resources for other topics, such as machine learning and data visualization[4], are far more numerous and readily available online, highlighting the relative scarcity of KG educational content online.

To address these limitations, the Alan Turing Institute KG resources[5], and the European Union innovation network of COST Action "Distributed Knowledge Graphs" (COST DKG), CA19134[6] 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 of 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 in current open educational resources and educational KGs in capturing intricate relationships, such as connections among resources covering similar topics, and enabling semantic interoperability, such as automatically connecting material at the metadata level. Further, there is a need for supporting dynamic, real-time updates in educational catalogues, highlighting the importance of semantically advanced approaches. The study identifies two primary challenges, which are addressed through the proposed research questions.

Challenge 1: The first 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 catalogue 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 descriptions, this requires a considerable amount of time and effort to insert all the appropriate metadata values, which is often not an automated procedure. Therefore, a reasonable minimum set of 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 September 2025, Coursera has only one KG course, the "Knowledge Graphs for RAG".

[3] As per September 2025, edX has no KG courses.

[4] Machine learning has 357 courses, and data visualisation has 130 associated courses in edX, as per September 2025.

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

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

Challenge 2: Contextual relevance versus feasible execution. Educational institutions and educators use different terminologies, 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: What are the minimum pedagogical requirements a teaching KG catalogue 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?

RQ1 focuses on creating a pedagogically-grounded representation of educational content, which facilitates KG schema to ensure that the content is described meaningfully from a pedagogical perspective. RQ2 follows up with the focus on the interoperability and connectedness of educational content, to ensure better alignment between the course within different teaching and learning contexts. RQ3 then addresses the ontological structure needed to embed the meta data and relations between the entities that represent the educational content; it investigates the ontological schema that the Teaching KG can offer to better represent, connect, and retrieve the educational content for the teachers and learners. The requirements considered in the research questions are determined by the characteristics and availability of the collection data. The following subsections analyse the pedagogical and contextual alignment pursued in the design of the Teaching KG. In the subsequent sections, we detail the specific measures implemented to ensure this alignment.

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, and the European Union innovation network of COST Action "Distributed Knowledge Graphs" (COST DKG), CA19134. Additionally, supplementary resources were obtained from the authors' private lists with resources for open-source tutorials and book recommendations related to KGs.

For course collection, resources were gathered via the Alan Turing Institute, COST Action DKG, and publicly available online courses. Data collection for the Turing Interest Group on Knowledge Graphs was directly performed via its GitHub repository, where members could create pull requests with their module materials.[7] The (teaching) GitHub repository was promoted via the group's regular meet-ups, annual symposiums, and mailing list.

Through COST Action DKG, an online data collection form was distributed to Semantic Web experts, requesting personal information, such as their full name, affiliation, and email for the purposes of connecting them sufficiently with their courses, and course details, such as course names, URLs, and weekly lecture schedules' titles. We relied on the mature network that the COST Action DKG had cultivated of experts in distributed KGs across Europe for the identification, and contact of semantic web experts via its newsletter. All data collected through this form were provided voluntarily, and participants gave their informed consent for the use of their course-related information for research purposes. The topics were described by Semantic Web experts who teach KG courses within two COST Action DKG initiatives, during an in-person interactive Workshop in Malaga, Spain on September 2023, and a follow-up online session on May 2024. During those sessions, the experts were provided with a structured template designed

[7] https://github.com/turing-knowledge-graphs/teaching/

to guide their input on relevant topics and skills. For each topic, the template included specific fields to be completed, such as educational level and topic difficulty level (aligned with the https://schema.org/educationalLevel). This step was important to provide semantic information about what each field aims to represent, and according to which definition. From the topics collection, we have a detailed description of theoretical and practical skills, learning objectives, educational level alignment, and knowledge topics.

The template intentionally avoided restrictions or mandatory fields, allowing flexibility for brainstorming and creativity in experts' 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 template was presented in an online collaboratively writing document, with the description and the property it aligns with in a parenthesis. The collection of educational materials began with contributions from the course instructors. To expand the scope and depth of the repository, additional resources in diverse formats 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

3.3.1. Ethical and Legal Concerns

The development and deployment of the Teaching KG inevitably raise important ethical and legal considerations, particularly in relation to copyright, licensing, privacy, and data ownership [30]. Since the catalogue incorporates content from institutional courses and open educational materials, we have adopted a cautious and transparent approach to data handling and reuse.

To address copyright and licensing issues, the Teaching KG does not host or redistribute course materials locally. Instead, it interlinks to externally available resources through persistent identifiers or URLs, ensuring that original content remains under the control of its rightful owners. All interlinked materials are subject to the licensing conditions defined by their respective creators or institutions. Whenever possible, we prioritize publicly available and open resources, thereby aligning the catalogue with established open science and open education principles [8]. We make no claim of ownership over any externally sourced content and fully acknowledge the intellectual property rights of content creators. Therefore, the Teaching KG does not currently aim to integrate directly with existing learning platforms, such as Coursera.

In relation to privacy and personal data, the catalogue strictly limits the inclusion of personally identifiable information. Only the minimal set of public data, namely the instructor's name, institutional affiliation, and professional contact email, is included, and solely when such information is already publicly accessible through institutional or course websites. No sensitive or non-public data are stored or shared.

Our approach is further guided by ethical principles of transparency, accountability, and respect for academic authorship. The design of the Teaching KG complies with relevant data protection regulations and institutional policies, ensuring that all information is handled responsibly and lawfully. By prioritizing interlinking over replication and attribution over appropriation, we aim to balance openness and accessibility with the ethical stewardship of educational data.

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 materials 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 performance 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 as pedagogy experts are, 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.

The pedagogical alignment supported by the Teaching KG stems from its ability to semantically and contextually interconnect learning content, learning outcomes, and, where available, learning activities such as assessment or practice. This form of connectedness is consistent with established pedagogical practices, including, but not limited to, principles discussed in the constructive alignment literature, which emphasize coherence between learning objectives, instructional activities, and learning resources [31]. While the Teaching KG was not explicitly designed as a constructive-alignment framework, its explicit representation of relationships between skills, knowledge topics, and educational resources (theoretical and practical) can support instructors in reflecting on the alignment of their course components. In particular, the KG can assist instructors in examining whether intended learning outcomes are sufficiently covered by available resources, identifying potential gaps, and selecting materials that are appropriate for a given educational level. By making such relationships explicit and reusable, the Teaching KG may contribute to more coherent and informed course design, reducing reliance on ad hoc resource selection and supporting pedagogically grounded instructional decisions.

Through the interaction model of the proposed web application, the Teaching KG adopts a tiered metadata policy. Only a minimal set of pedagogically essential fields is mandatory to ensure KG consistency and retrieval functionality, while all other pedagogical descriptors remain optional and context-dependent. This design avoids forced over-specification and supports gradual enrichment of the KG over time. While the user interface enforces completion of a minimal set of mandatory fields, it does not require exhaustive pedagogical descriptions. Users may submit partial but valid entries, which can be incrementally enriched later without violating KG constraints.

4. Ontology

4.1. Ontology Development

We chose to develop an ontology as the semantic structure to interconnect the populated data in 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 of data, 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 through 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 enables 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 [32], namely the process of requirements specification, and implementation to publication and maintenance, incorporating the conceptualization step described in Sabiox [33]. 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 materials. To create a semantically structured catalogue that accommodates education broadly—rather than being confined to a single domain—we developed a high-level ontology. This approach allows for future adaptation to 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. 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[8] and EduCOR ontology [34]. 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 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. In Table 2, we analyse the main modules in more detail, derived from the primary classes in our schema.

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 modelled 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 explicit links between learning outcomes, skills, knowledge topics, and educational resources operationalize pedagogical alignment and enable reasoning over coverage, prerequisites, and instructional coherence. The input form for adding new courses to the system ensures that newly added courses fulfil the required characteristics.

As the ontology's scalability depends on how flexible its schema and constraints are, we enforce only a minimum of required fields, and ask for the additional (meta)data as optional values. To remain scalable, our ontology supports

[8] https://schema.org/

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/def/courses/kos/
base:	https://w3id.org/def/courses#

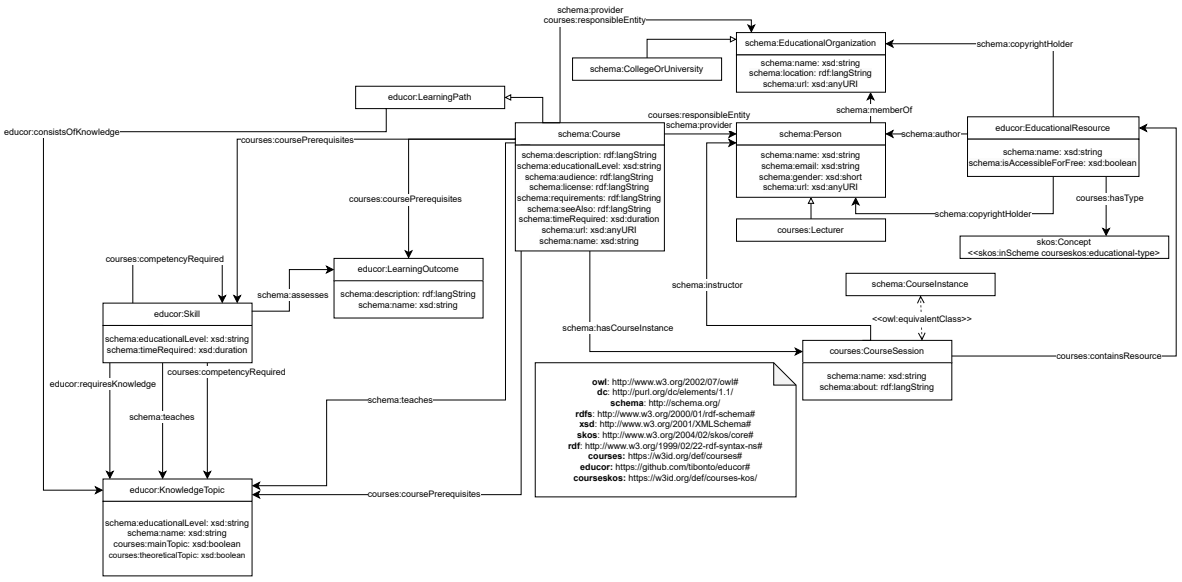


Figure 1. Visualization of the ontology classes and properties using Chowlk [35].

modularity and mechanisms for gradual schema evolution. In practice, new courses could be incorporated through partial alignment (capturing only shared core attributes) while allowing domain-specific extensions.

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 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, we adopted the following definitions:

• Skill: An abstract concept consisting of multiple elements across different levels and scopes to be mastered, namely 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 KG centred on the individuals who teach courses 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 materials include textbooks, videos, and libraries, and are manually annotated with metadata to foster discoverability and reuse. In order to distinguish whether 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 with respect to the project's requirements and data availability.

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 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 refer to them using this terminology, 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.

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

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.

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

the KG and defining entities, attributes, and relationships. Each of the categories was designed to capture both mandatory and optional data points. Table 4 explicitly distinguishes between mandatory and optional elements; the logical "AND" denotes compulsory fields, while "OR" denotes alternative optional descriptors that may be provided depending on course context. A balanced minimum set of entities per classification was made compulsory for educators to complete when inputting their course data, to ensure that sufficient entities, attributes, and relationships can be defined to construct the KG and facilitate course connections and thus knowledge sharing.

The pedagogical classification presented in Table 4 structures the core educational entities and relationships that can be represented in the Teaching KG, while explicitly distinguishing between mandatory and optional elements. This classification is designed to ensure that a minimal, pedagogically meaningful description of courses and resources can be captured, while still allowing richer contextual information to be provided when available. By organizing learning outcomes, skills, knowledge topics, and educational resources within a shared semantic structure, the Teaching KG enables these elements to be interrelated in a consistent and reusable manner, supporting subsequent reflection on the coherence of course design.

For example, a learning outcome such as *"students can write SPARQL queries to retrieve data from an RDF graph"* can be represented in the KG through its association with relevant skills (e.g., SPARQL querying), knowledge topics (e.g., SPARQL syntax and querying RDF data), and corresponding educational resources. These resources may include theoretical materials introducing the SPARQL language, as well as practical materials such as guided query exercises or hands-on laboratory sessions using RDF datasets. Where available, assessment-related resources (e.g., practical assignments or exams requiring SPARQL query formulation) can be linked to the same learning outcome. While the Teaching KG does not enforce a specific instructional design methodology, this explicit representation of relationships can support instructors in examining whether learning outcomes are adequately supported by appropriate teaching materials and learning activities. In this way, the pedagogical classification facilitates reflective course design, helps identify missing or weakly supported outcomes, and promotes more coherent alignment between objectives, resources, and educational level.

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 provide the SPARQL queries corresponding to these competency questions grouped by topic, course, dataset, and educational material. We disclose the SPARQL queries for the competency questions we cover in this paper. In Table 5, the Yes* answers refer to a slight change in the competency question to retrieve the theoretical educational material instead of a specific multimedia type (e.g. slides). As reported, we cover most of the competency questions and allow a wide 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?	Yes	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/def/courses#>

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/def/courses#>
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/def/courses#>
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/def/courses#>
PREFIX schema: <http://schema.org/>
PREFIX edutype: <https://w3id.org/def/courses/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/def/courses#>

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

d. The query "Which slides are linked to Y?" has been reformulated as "Which theoretical materials are linked to Y?". In the original formulation [1], educational resources were categorized into slides and labs. This terminology has been generalized to theoretical and practical materials, since the Knowledge Graph now incorporates a wider variety of content formats beyond slides and labs, including textbooks and other resources.

List all theoretical educational material associated with course Y.

```
PREFIX base: <https://w3id.org/def/courses#>
PREFIX schema: <http://schema.org/>
PREFIX edutype: <https://w3id.org/def/courses/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/def/courses#>
PREFIX schema: <http://schema.org/>
PREFIX edutype: <https://w3id.org/def/courses/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/def/courses#>
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 dataset D

a. Which courses use D? Identify all courses that utilize dataset D.

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

SELECT DISTINCT ?course
WHERE {
  ?course schema:hasCourseInstance ?courseSession .
  ?courseSession base:containsResource ?dataset .
  FILTER(?dataset = <D>)  # replace D with URI of the given Dataset
}
```

4.3.4. For Material M

a. 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 allows incorporating changes at any level of the KG, including ontology changes, new input data sources, and removing old data. For this purpose, we rely on the RDF Mapping Language (RML) [36], a declarative mapping language that provides support for transforming any (semi)structured data format into RDF. For the workflow, we use the methodology proposed in [37][9], which extends the Linked Open Terms methodology (LOT) [38] for ontology development and incorporates the creation of KGs and the life cycle of all related artifacts. An overview of the proposed approach is presented in Figure 2.

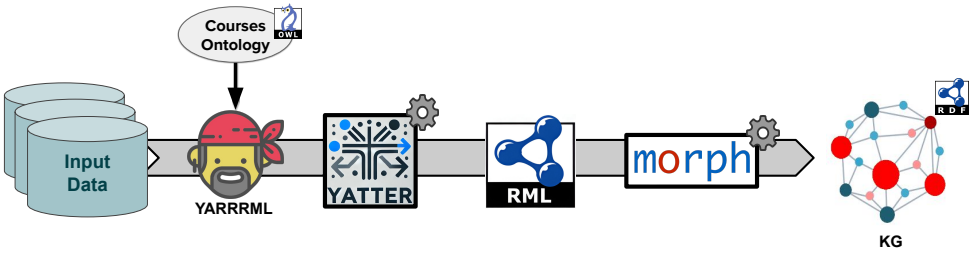


Figure 2. The followed pipeline to transform educations resources and integrate them into the Courses Ontology.

5.1.1. Knowledge Graph Construction

We transform (semi)structured input data into an RDF-based knowledge graph using RML mappings. For this first version, the data are provided in CSV files, but the pipeline can be adapted to any kind of common data format (CSV, JSON, XML, RDB). To facilitate this task, we use first OWL2YARRRML tool[10] to draft the mappings. This tool takes an ontology as input and generates a set of YARRRML [39] templates (see Listing 1 as an example for the `educor:Skill` class) that can be then filled by the knowledge engineer with the data references (Listing 2). Once the mappings are complete, we translate them into RML using Yatter [40] (see Listing 3). 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 [41]. With this approach, we ensure maximum flexibility, as the new specification of RML [36] allows the integration of any kind of semi-structured format. To ensure high performance and scalability during the transformation process we decided to select Morph-KGC based on the last benchmark results [42, 43].

```
map_skill:
  sources:
    - [ ]
  s: http://example.org/resource/skill/$()
  po:
   - [ rdf:type, educor:Skill ]
   - [ schema:educationalLevel, $() ]
   - p: schema:assesses
     o:
      - function: join(map_learning_outcome, equal($(), $()))
   - p: schema:teaches
     o:
```

[9] https://lot.linkeddata.es/LOT4KG/
[10] https://github.com/oeg-upm/owl2yarrrml/

```
      - function: join(map_knowledge_topic_t, equal($(), $()))
  - p: schema:teaches
    o:
     - function: join(map_knowledge_topic_f, equal($(), $()))
  - p: educor:requiresKnowledge
    o:
     - function: join(map_knowledge_topic, equal($(), $()))
  - p: courses:competencyRequired
    o:
     - function: join(map_competency_req, equal($(), $()))
```

Listing 1: YARRRML mapping template (empty placeholders)

```
map_skill:
  sources:
   - [kg_topics.csv~csv]
  s: http://example.org/resource/skill/$(skill)
  po:
   - [rdf:type, educor:Skill]
   - [schema:educationalLevel, $(educationalLevel)]
   - p: schema:assesses
     o:
     - function: join(map_learning_outcome, equal($(assesses),$(assesses)))
   - p: schema:teaches
     o:
     - function: join(map_knowledge_topic_t, equal($(teaches_t),$(teaches_t)))
   - p: schema:teaches
     o:
     - function: join(map_knowledge_topic_f, equal($(teaches_f),$(teaches_f)))
   - p: educor:requiresKnowledge
     o:
     - function: join(map_knowledge_topic, equal($(requiresKnow),$(requiresKnow)))
   - p: courses:competencyRequired
     o:
     - function: join(map_competency_req, equal($(competencyReq),$(competencyReq)))
```

Listing 2: YARRRML mapping instantiated for `kg_topics.csv`

```
<map_skill> a rml:TriplesMap;

  rml:logicalSource [
    rml:source "kg_topics.csv";
    rml:referenceFormulation ql:CSV
  ];
  rml:subjectMap [
    rml:template "http://example.org/resource/skill/{skill}";
    rml:class educor:Skill;
  ];
  rml:predicateObjectMap [
    rml:predicate schema:educationalLevel;
    rml:objectMap [
      rml:reference "educationalLevel";
```

```
    ];
  ];
  rml:predicateObjectMap [
    rml:predicate schema:assesses;
    rml:objectMap [
      rml:parentTriplesMap <map_learning_outcome>;
      rml:joinCondition [ rml:child "assesses";rml:parent "assesses"; ];
    ];
  ];
  rml:predicateObjectMap [
    rml:predicate schema:teaches;
    rml:objectMap [
      rml:parentTriplesMap <map_knowledge_topic_t>;
      rml:joinCondition [ rml:child "teaches_true"; rml:parent "teaches_true"; ];
    ];
  ];
  rml:predicateObjectMap [
    rml:predicate schema:teaches;
    rml:objectMap [
      rml:parentTriplesMap <map_knowledge_topic_f>;
      rml:joinCondition [ rml:child "teaches_false"; rml:parent "teaches_false"; ];
    ];
  ];
  rml:predicateObjectMap [
    rml:predicate educor:requiresKnowledge;
    rml:objectMap [
      rml:parentTriplesMap <map_knowledge_topic>;
      rml:joinCondition [ rml:child "requiresKnow"; rml:parent "requiresKnow"; ];
    ];
  ];
  rml:predicateObjectMap [
    rml:predicate courses:competencyRequired;
    rml:objectMap [
      rml:parentTriplesMap <map_competency_req>;
      rml:joinCondition [ rml:child "competencyReq"; rml:parent "competencyReq"; ];
    ];
  ].
```

Listing 3: RML mapping for `map_skill` in Turtle

5.1.2. *Life Cycle and Sustainability*

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 those already transformed into RDF are modified. To mitigate the impact of the first one, we rely on the LOT4KG methodology [37] and the tool OntoRipple[11] [44], that helps to propagate ontology changes over RML mappings in a semi-automatic way. For the second case, there is still no 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 updated manually, adding the new sources and associated rules, and then the KG is regenerated. For the re-generation, our aim is to use IncRML [45], ensuring that only the new triples are actually generated. IncRML builds on the use of KG snapshots and Linked Data Event Streams (LDES) to capture changes over time and to incrementally adapt RML mappings accordingly, enabling the selective generation of RDF triples without

[11] https://github.com/oeg-upm/OntoRipple

full re-materialization. In this work, IncRML is considered as a promising direction for reducing regeneration costs rather than as a fully integrated component.

From a long-term sustainability perspective, the project explicitly adopts an open and community-driven approach. All core artefacts produced within the project, including ontologies, RML mappings, transformation workflows, and supporting documentation, are made publicly available under open licenses, ensuring transparency, reproducibility, and long-term accessibility. This openness lowers the barrier to entry for third-party contributions, enabling external researchers and practitioners to actively participate in the maintenance and evolution of the KG by proposing new data sources, mappings, or refinements.

In addition, ontology evolution is explicitly addressed through the adoption of the OWL Change Ontology (OCH) vocabulary[12] to formally represent and publish ontology versions and changes over time. By making ontology updates explicit, machine-readable, and publicly accessible, OCH facilitates systematic change tracking, supports automated impact analysis, and enables downstream users to adapt their mappings and applications accordingly. Combined with open tooling and incremental regeneration mechanisms, this approach ensures that the KG can evolve in a controlled, transparent, and community-supported manner well beyond the lifetime of the project.

5.2. Web Application

Although the Semantic Web experts are the primary users of the Teaching KG, we anticipate potential interest from educators and users with non-technical backgrounds. Therefore, a simple and user-friendly approach to interact 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 application, which supports predefined use cases, designed to allow users to create a new course on KGs, extend their already existing courses with content from the KG, or browse existing content of other courses in the KG.

The web application (see Figure 3) is designed around the features of reading from, and writing to the KG, supporting adding new course facilitators, exploring KG content, or creating new content, among other features (Figure 3A).

5.2.1. Interface Design

The design of the Teaching KG interface followed an iterative and interdisciplinary approach, involving experts in pedagogy, Semantic Web technologies, and software development. The primary goal of this process was to develop an interaction model that enables instructors to insert and manage course data in a manner that is both pedagogically meaningful and technically consistent with the underlying ontology.

Pedagogy experts contributed by defining the pedagogical classification presented in Section 4.2.1 and by prioritizing the educational entities that reflect learning design principles, such as learning outcomes, skills, and knowledge topics. Their input guided decisions on which elements should be mandatory to ensure a minimal pedagogical description and which could remain optional to accommodate institutional and contextual diversity. Semantic Web experts ensured that the interface workflows aligned with the ontology structure and semantic constraints, preventing inconsistent or incomplete representations during data entry. Software developers translated these requirements into concrete UX/UI design decisions, focusing on reducing cognitive load, supporting incremental data entry, and adhering to established usability principles for web applications. The interface design was refined through multiple design sessions, in which interaction flows, data entry forms, and validation mechanisms were jointly reviewed. These sessions informed decisions such as optional fields, pre-defined lists of options, contextual guidance embedded in the interface, and validation rules reflecting the pedagogical classification. All experts involved in this process were recruited from the authoring team of the paper, ensuring close collaboration and shared domain understanding throughout the design and implementation phases.

To operationalize the pedagogical classification while maintaining a low entry barrier for educators, the interface adopts a progressive data entry strategy that balances minimal mandatory input with optional enrichment. Only a set of essential metadata fields is required to ensure KG consistency and improved course connectedness, while additional information can be provided incrementally, e.g., through completing and extending the course content in

[12] http://w3id.org/def/och

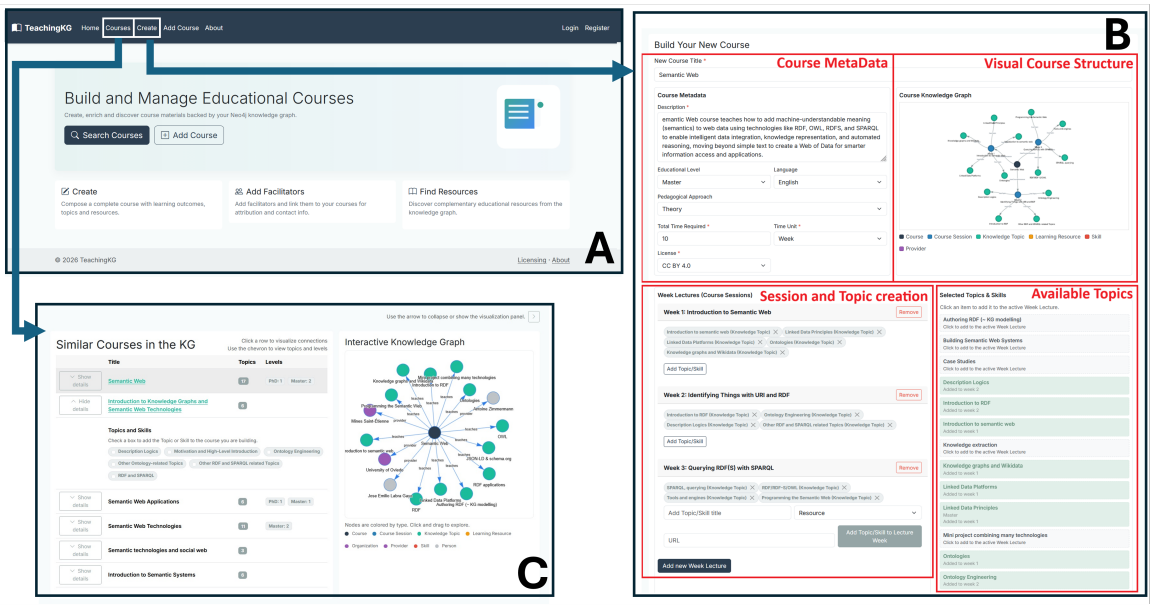


Figure 3. Web application interface structure, tabs, and use cases. A: Hhome page, B: Creating a new course, C: Searching resources within the KG

separate sessions. The interface supports this process through usability-oriented features such as auto-completion for predefined or previously used values, contextual highlighting of missing information relevant to retrieval quality, and the visualization of exemplary similar courses from the KG course as a reference structure. Throughout the interaction, users receive feedback during data submission and retrieval, and are provided with contextual guidance for the input of course data and content. The web application is implemented utilizing a lightweight Flask[13] backend, combining read-only access to an RDF knowledge graph with authenticated data entry workflows backed by a Neo4j[14] graph database. RDF KG data are used to support catalogue browsing and exploration of the online KG, while Neo4j is employed as a space to persist courses authored through the interface, store associated course embeddings, and support an offline, computationally efficient calculation of course similarities. The web application is mainly written in Python and can be accessed at our GitHub repository..

5.2.2. User Experience

To illustrate how users interact with the Teaching KG and to demonstrate the system's core functionality, we describe two representative use cases. These use cases are application-driven; each emphasizes on the user's goal, the sequence of interface interactions, and the system responses generated through KG-based retrieval, validation, and visualization mechanisms.

Use Case I: Creating a New Course and Validating Mandatory Fields Suppose a teacher wants to design a new course on KGs without relying on previously prepared teaching material. The primary interaction goal is to explore existing courses, topics, and educational resources in order to construct a coherent course outline. To initiate the course design process, a user selects the "Create Your Course" option in the interface (Figure 3B). They are guided through a structured form that groups input fields into logical categories, including general course information (title, description, keywords, and learning goals) and target audience specifications.

After specifying the intended audience level and defining prerequisites such as basic database knowledge and familiarity with Python, the user initiates a content search. The system detects missing keywords and prompts the

[13] https://flask.palletsprojects.com
[14] https://neo4j.com

user to provide them, highlighting their importance for identifying thematically similar courses. Once the search is submitted, the system queries the Teaching KG by combining course metadata, learning goals, keywords, and audience constraints, while excluding prerequisite topics from the retrieved results.

Upon submission, the system validates the completeness of the provided metadata. If mandatory fields, such as copyright or licensing information, are missing, the interface presents targeted feedback explaining the relevance of the missing metadata for course sharing and reuse. The user is directed to contextual help elements that summarize licensing options and provide links to external guidance.

The interface presents the results in two complementary views, first in a tabular overview listing relevant courses, topics, and linked educational resources, and secondly on a graph-based visualization that shows how courses and resources are interconnected within the KG, as shown in Figure 3B "Visual Course Structure". Through these views, the user can explore reused resources, identify commonly referenced materials, and assess their relevance and credibility for inclusion in her own course.

Use Case II: Resource Discovery When a learner or a teacher seeks targeted educational resources for a specific topic, they can use the Teaching KG to locate high-quality, relevant learning materials while maintaining contextual awareness of related topics and courses. Initially, the learner begins the interaction by selecting the Search for Learning Content option (Figure 3C). The interface provides a search field that supports multiple query terms, along with filter options based on course level, language, resource type, and openness. These controls are designed to balance expressive search capabilities with usability and to reduce the need for repeated queries.

When the search is executed, the system analyses the query terms and selected filters to retrieve relevant content from the Teaching KG. The results are displayed in both a table view and a graph visualization. The table view enables quick inspection of resource metadata and direct access to materials, while the visualization reveals how the retrieved resources are connected to other topics and courses within the KG.

This dual presentation supports both goal-oriented retrieval and exploratory learning. Learners can efficiently access specific resources while also discovering adjacent concepts and related learning paths, leveraging the interconnected structure of the KG.

5.2.3. Similarity Computation

When requesting educational content from the KG, the system retrieves relevant courses and their topics on the users input. Course relevance comprises 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.

• Textual 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)[15] 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. Embeddings are generated and stored separately for those entity types. Similarities are calculated on the same entity type and translated into KG relations among similar nodes. This approach allows flexible and customizable retrieval of, e.g. similar courses, through controlling the KG queries that utilize the transitive relations between courses, e.g. through their skills or topics.

In order to ensure computational efficiency, STS is pre-calculated among all courses, (sub)topics, and materials only once, in contrast to calculating the similarities in 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. To avoid a system fail when the content of the KG does not include materials relevant to the user query, or when SPARQL queries fail

[15] https://sbert.net/

for any reason, error handling in the system's back-end falls back to empty response structure and prioritizes a clear communication with the user through the interface, to provide clearer or different search terms, or try again due to a query-error.

SBERT is used as a pre-trained model, accessed via the HuggingFace Transformers Python module, without additional domain-specific fine-tuning. This choice was motivated by prior work demonstrating that SBERT produces high-quality sentence embeddings, with an average accuracy of 87.41 for the base model and 87.69 for the large model [46]. This accuracy is suitable for semantic similarity tasks while maintaining computational efficiency. To ensure scalability and efficient query performance, the latent embedding space is structured in accordance with the ontological organization of the KG and implemented using a Neo4j graph database. Textual embeddings are computed offline for course-related textual content, specifically titles and descriptions, and stored directly as properties of the corresponding graph entities. This design minimizes the computational overhead associated with real-time embedding generation and cosine-similarity calculation, and enables fast similarity-based retrieval over large numbers of courses.

6. Evaluation

6.1. Validation

We selected OWL as the modelling language for the ontology as it provides unambiguous semantics and brings reasoning capabilities. Enabling an ontology-driven data creation pipeline will minimise errors in the data [47] (e.g., out-of-range data values, incorrect domains and ranges for a property in a triple). The reasoning capabilities of OWL will ensure that the correctness of the generated data is maintained, at least with respect to obvious logical errors [48–50] (i.e., detection and repair of unsatisfied classes and data inconsistencies).

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 approaches to represent data constraints using languages like SHACL[16] and ShEx [51]. Alternatively, as proposed by Kharlamov et al. [52], 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 (stratified) negation [52][17]. The use of datalog rules to represent integrity constrains enables the use of efficient and sound datalog engines (e.g., RDFox [53])[18].

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), which would otherwise be 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 explore whether the Teaching KG meets the needs of its intended users, we gathered 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 catalogue. Through

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

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

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

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 + ;
  # . . .
}
```

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

SHACL

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

```
<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 ;
  # . . .
```

Table 8
ShEx and SHACL shapes to identify missing properties.

ShEx

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

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

SHACL

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

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

semi-structured interviews with two stakeholders, who teach semantic web courses, and 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 catalogue. The interview was formulated around questions related to the future use of the Teaching KG catalogue by stakeholders and consisted of twelve questions (Q1 - Q12) with no responses omitted, as shown in Table 10.

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)

Table 10
Knowledge Graph Catalogue Interview Questions by Theme and Classification

Code	Question	Theme	Classification
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?	Content	Topics and Materials
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?	Layers	Semantic Layer Components
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?	Purpose	Usefulness Components
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?	Content	Use-Cases and Examples
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?	Purpose	Usage of Teaching KG Content
Q6	As lecturer in teaching knowledge graphs, what opportunities do you foresee for collaboration and knowledge sharing among educators within the catalogue's community?	Purpose	Potential of Teaching KG
Q7	What specific teaching objectives and expertise in semantic web would you expect to be present in the catalogue?	Layers	Learning Layer Components
Q8	Could you share insights on how the catalogue's structure and organization would best be?	Layers	Re-use of Resources for Semantic Layer
Q9	In your opinion, how could it best facilitate students' comprehension and retention of semantic web concepts, based on your teaching experience?	Purpose	Learning Outcomes
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?	Content	Linkage to Additional Resources
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?	Purpose	Measuring Pedagogical Efficiency
Q12	Other examples of catalogues or other resources that you are currently using or find effective in your teaching practice?	Layers	Literature Paradigms and Reusage of Resources

The interviews revealed that the semantic layer of the catalogue is fundamental to supporting student learning, and allows connecting different types of materials, such as open learning resources on the Web with tests. Regarding content-related aspects, both experts underscored the importance of incorporating interactive visualizations and query functionalities to promote hands-on engagement. They also highlighted the value of real-world case studies and domain-specific applications as essential for bridging theoretical knowledge with practical understanding. In terms of purpose, both participants agreed that ensuring the usefulness and effectiveness of the catalogue is crucial, framing it as a dynamic pedagogical tool that enhances student engagement and comprehension rather than serving merely as a static repository. Finally, both experts identified collaboration among educators as a core objective of the catalogue and as a central driver for community-oriented and open educational resource development within this

project.

Furthermore, we group the answers we received based on the classification of the questions into four common themes: (1) structured semantic layer, (2) practical learning material and usages, (3) real-world use-cases and examples, and (4) collaboration and knowledge sharing. While the first themes correspond to the classification of the interview questions, the latter, (4) collaboration and knowledge sharing) was introduced by the experts' responses. In Table 11, we report the alignment of each stakeholders' expectations with the elements present in the Teaching KG. 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 catalogue. 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 that 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 addresses the needs of KG educators, and provides a comprehensive learning framework for KG education concepts.

Table 11
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**
Layers	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
Content: Practical Learning Material, and Usage	Recommending interactive visualizations and query functionalities	Suggesting practical tutorials (e.g., Jupyter notebooks), interactive UI for SPARQL endpoint	Labs and practical educational material
Content: Use-Cases and Examples	Highlighting the importance of real-world case studies and examples from the catalogue	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

6.3. User Study

Further, we conducted a second round of user feedback collection to evaluate the system interface's usability. It is important to note that this evaluation provided insights into the front-end presentation of the system, rather than the KG, ontology or data integration pipeline. The participants were asked to complete a predefined user scenario in which they created a new course using the Teaching KG, as seen in Figure 3. The participants were provided with a demo of the interface[19] and the link to the testing platform[20]. They could only access the survey to provide feedback via the testing interface. We distributed the recruiting details via mailing lists and social media. During this task, users could search the system to discover and reuse relevant courses, topics and skills already available in the Knowledge Graph. The evaluation focused on assessing how effectively and efficiently users could complete the course creation process, as well as their perceived usability of the system. After completing the task, participants provided feedback through the system usability scale questionnaire (SUS) [54]. All responses were collected anonymously, and no personally identifiable information was recorded or stored. Participation in the study was entirely voluntary.

[19] You can access the demo at https://www.veed.io/view/224d8176-95cf-4512-82b4-159261d809b3?panel=share
[20] You can access the testing interface at https://teaching-kg-onto.vercel.app/test

We collected 19 valid responses from December 15th, 2025, until January 5th, 2026. The participants were in majority aged between 29 and 44 years old, either on postgraduate or PhD educational level, and they had predominantly advanced usage and familiarity with KGs on a daily or weekly basis, with advanced (3+ years) or intermediate (1-3 years) expertise in Semantic Web and KGs, and were teachers in KGs (i.e. Teaching Assistants).

The SUS score 58.82 (SD = 22.95, median = 65, max = 95) indicates moderate perceived usability with high variability which indicates heterogeneous user experiences as some liked the interface well (max=95.00), while others were confused on its functionalities. The approximately 95% confidence interval for the mean is [48.50, 69.13] which indicates an above marginal and below average usability for the interface.

Qualitative feedback from optional comments field indicated that participants found the system conceptually strong, but with functionality constrains: "*I would like some options not to be excluding each other. For example a course can be for both masters and phd. Also a course can both theoretical and practical so with hands-on experience. This happens especially in the semantic web field. There were also not so much examples to draw from (eg the topics) so the system might be much more complex when there are more examples. Overall however it was very easy to use and very interesting!*", and *"I am not sure it is realistic that a topic (from a similar course) can only be added to one week by default (of course one could add it to multiple weeks more manually, but that would be some unnecessary effort). I can imagine many scenarios where a broader or larger topic is discussed over multiple weeks and one might want to add that to multiple weeks. Similarly, I am not sure the activation of weeks for adding previous topics is very intuitive. Maybe, instead a pop-up asking to which week(s) a topic should be added would be a more intuitive way."*. These findings may indicate that the system provides a solid foundation for assisting users, particularly by offering access to relevant educational resources and information. However, future iterations should further explore alternative user experience designs. For example, future versions could simplify educational level classifications to undergraduate and postgraduate categories only, while also incorporating more detailed subtopic visualisations, such as distinguishing RDF serialisation formats into specific representations like JSON-LD and N3.

Additionally, a participant indicated that the use-case tested, namely creating a course, was not the ideal for our platform: *"This tool is a great help if I want to add information about courses to a knowledge graph. Nevertheless, I don't see the big advantage if i would do it. If I wanted to prepare a course about KGs, I would first look for other courses of the same subject. I could see an advantage if these courses were listed and added to a KG, so that I can easily find them and see, what is taught. I would only use this tool if I wanted to add my course to a KG of other courses, but not for building a course."*. Therefore, the most suitable use cases may involve the platform functioning primarily as a lookup or dictionary-style educational resource for quality, community-approved, institutional courses, rather than supporting the dynamic creation of entirely new courses directly within the platform.

Moreover, there are indications of participant's confusion regarding the interface's functionality: *"The key confusion is why we need such an interface, which seems to build a graph with nodes of parts in a lectures."* and *"The division of materials into ressources, assessment, or database is unclear to me. How about uploading a pdf as a ressource?"*. Some degree of user confusion regarding the presented interface elements is not unexpected, given the novelty of the application and the absence of directly comparable systems with which users may already be familiar. Furthermore, the inclusion of graph-based visualisations and the categorisation of educational materials by type may not align equally well with all user preferences or interaction styles; an observation we aim to further investigate into alternative interface designs and adaptive presentation strategies in future iterations of the platform.

One of the comments indicated that the goal and coverage of the interface was unclear and they expected to find a greater availability of courses: *"The system is not too difficult to use but I don't find it particular usable, useful or necessary. It would be useful if you could provide some background on the system: why you are building it, what purposes does it serve, who is the target user, what unique features do you offer and why are useful. For example, I don't understand why I need a course knowledge graph. Also, why I need a specialized course management system for KGs? Why can I not use this system for other courses?"*. It is important to note that participants were provided with a video demonstration explaining both the system functionality and the purpose of the evaluation scenario. Nevertheless, we acknowledge that the proposed approach may also be applicable to courses and educational resources beyond the specific domain explored in this study, a step that will require storing more personal information and material that we do not own, such as slides to a course. Indeed, this is a current limitation of our system which we aim to address in future work by focusing on inclusion of more courses, and non-personal data. However, currently

Table 12
Future development feedback from a 5-point Likert scale (1 - strongly disagree, 5 - strongly agree).

Item	Mean	SD	Interpretation
Reviewing similar courses helped identify overlooked topics	**3.00**	1.38	Moderately positive, but not strong
Graph visualization on the side is helpful	**3.11**	1.44	Most positively perceived feature
Task prompted reflection on required metadata	**2.89**	1.52	Mixed but meaningful reflection
Course formulation had too many fields	**2.16**	1.07	Leaning toward disagreement

as we comply with the General Data Protection Regulation (GDPR) we do not process, reuse or crawl resources and personal lectures information that have not been voluntarily added on our system [30].

Moreover, we collected further feedback for future developments in a structured way as seen in Table 12. The goal was to evaluate current elements and design decisions that affect the effectiveness of the interface and could secondarily affect its usability. Reviewing existing courses before authoring a new one was perceived as moderately helpful (M = 3.00, SD = 1.38), suggesting that reuse of prior course structures supported topic discovery for many participants. The graph visualization component received the highest mean rating (M = 3.11, SD = 1.44), indicating its perceived usefulness, albeit with substantial variability across users. Participants also reported that the task prompted reflection on the metadata required for course development (M = 2.89, SD = 1.52). Concerns regarding the number of required fields were less pronounced overall (M = 2.15, SD = 1.07), though some users did perceive the formulation process as demanding.

Concluding, there are a few important things to consider. First, the objective of the user study was formative rather than summative. Given that the Teaching KG interface is currently an early-stage (alpha) research prototype, the primary goal of the evaluation was to identify usability challenges and design limitations, rather than to demonstrate high or benchmark-level usability. Also, it is critical to note that the usability evaluation took place for one use case and not for the whole functionality of our platform and KG, due to resource constrains. Secondly, at this stage of development, the system has not yet undergone iterative interface refinement or user-centred optimization. Moreover, there is no directly comparable or established reference system for course authoring based on educational KGs that could serve as a usability benchmark or provide transferable design patterns. As a result, design decisions were necessarily exploratory and driven by conceptual and pedagogical considerations rather than by prior usability evidence. Finally, within this context, the reported usability results should be interpreted as diagnostic indicators highlighting aspects of the interface that require improvement. The findings provide actionable insights that directly inform future redesign efforts and iterative development of the system. Consequently, the usability evaluation is intended to support system maturation rather than to position the current prototype as a finalized or production-ready interface.

7. Potential Impact, Limitations and Future Work

We presented the project plan and motivation in 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 in a Teaching KG for KG courses from the Semantic Web community. More specifically, 10 new lecturers stepped forward to contribute their courses to our project, and many researchers expressed eagerness to use 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 underlined the prospect of building on this momentum to further develop and enhance the project with new features and functionality. Further, the integration of Teaching KG with the Open Research Knowledge Graph (ORKG)[21] and other structured data sources has the potential to significantly enhance smart learning applications by tightly coupling research knowledge with pedagogical design [55]. By linking learning objectives, courses, and assessments

[21] https://orkg.org/

in Teaching KGs to semantically structured research contributions in ORKG, educators can more easily identify relevant, up-to-date scientific content and align it with instructional goals. This connection will support evidence-informed curriculum design, enable semi-automated course and task generation, and facilitates reflective teaching practices by making implicit pedagogical decisions explicit and machine-readable.

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[22] 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 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 [56]. 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 represented 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 analyse 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. An alternative scalability strategy would be to expand beyond semantic web courses and include educational materials and topic classifications across broader AI domains. As AI literacy becomes an increasingly important concern, the Teaching KG could be extended to support newcomers in AI learning. Furthermore, by utilising Language Models, we could offer a grounded generation of educational material inspired by the educators content to assist assessment and accessibility [57–60].

8. 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 high-level 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 preliminary results indicate 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.

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

Acknowledgments

Eleni Ilkou would like to thank Tobias Käfer for his contribution in gathering material and classification of skills, and COST Action CA23147 GOBLIN – Global Network on Large-Scale, Cross-domain and Multilingual Open Knowledge Graphs, supported by COST (European Cooperation in Science and Technology, https://www.cost.eu). David Chaves-Fraga acknowledges funding from the Spanish State Research Agency (Agencia Estatal de Investigación) through the REOPEN (PID2023-149549NB-I00). He also acknowledges the support of the Galician Ministry of Education, Universities and Vocational Training and the European Regional Development Fund (ERDF, "A way of making Europe") through the "Centro de Investigación de Galicia accreditation (ED431G-2023/04, 2024-2027)" and the "Reference Competitive Group accreditation (ED431C 2026/52, 2026-2029)". 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). Ernesto Jimenez-Ruiz was supported by Turing Innovations Limited and The Alan Turing Institute's Defence and Security Programme via the project GUARD. Ernesto would also like to thank the organisers and contributors 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] In UNESCO. Recommendation on open educational resources (oer). *Legal Instruments*, 2019.

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

[10] 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.

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

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] 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.

[17] 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.

[18] 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.

[19] 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.

[20] Ismail Aliyu, A Kana, Salisu Aliyu, et al. Development of knowledge graph for university courses management. *International Journal of Education and Management Engineering*, 10(2):1, 2020.

[21] 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].

[22] 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. .

[23] Zeju Li, Linya Cheng, Chunhong Zhang, Xinning Zhu, and Hui Zhao. Multi-source education knowledge graph construction and fusion for college curricula. pages 359–363, 2023.

[24] 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.

[25] 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/.

[26] 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.

[27] 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.

[28] 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.

[29] 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.

[30] Paul Voigt and Axel Von dem Bussche. The eu general data protection regulation (gdpr). *A practical guide, 1st ed., Cham: Springer International Publishing*, 10(3152676):10–5555, 2017.

[31] John Biggs. Enhancing teaching through constructive alignment. *Higher Education*, 32(3):347–364, 1996. URL https://doi.org/10.1007/BF00138871.

[32] 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.

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

[34] 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.

[35] 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.

[36] 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.

[37] Romana Pernisch, Maria 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.

[38] 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.

[39] 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.

[40] 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.

[41] 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.

[42] 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.

[43] Enrique Iglesias, David Chaves-Fraga, Sören Auer, and Maria-Esther Vidal. Normalization-driven optimization of knowledge graph creation for semantically grounded information fusion. *Available at SSRN 5280029*.

[44] Diego Conde-Herreros, Oscar Corcho, and David Chaves-Fraga. Ontoripple: Making waves in the knowledge graph lifecycle. *SoftwareX*, 33:102542, 2026.

[45] Dylan Van Assche, Julián Andrés Rojas, Ben De Meester, and Pieter Colpaert. Incremental knowledge graph construction from heterogeneous data sources. *Semantic Web*, 17(2):22104968251412270, 2026.

[46] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China, 2019. Association for Computational Linguistics. URL https://doi.org/10.48550/arXiv.1908.10084.

[47] Michael Uschold. Finding and avoiding bugs in enterprise ontologies. In Heiko Paulheim, Jens Lehmann, Vojtech Svátek, Craig A. Knoblock, Matthew Horridge, Patrick Lambrix, and Bijan Parsia, editors, *Joint Proceedings of the 5th Workshop on Data Mining and Knowledge Discovery meets Linked Open Data and the 1st International Workshop on Completing and Debugging the Semantic Web (Know@LOD-2016, CoDeS-2016) co-located with 13th ESWC 2016, Heraklion, Greece, May 30th, 2016*, volume 1586 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2016. URL https://ceur-ws.org/Vol-1586/codes1.pdf.

[48] Aditya Kalyanpur, Bijan Parsia, Evren Sirin, and Bernardo Cuenca Grau. Repairing unsatisfiable concepts in OWL ontologies. In York Sure and John Domingue, editors, *The Semantic Web: Research and Applications, 3rd European Semantic Web Conference, ESWC 2006, Budva, Montenegro, June 11-14, 2006, Proceedings*, volume 4011 of *Lecture Notes in Computer Science*, pages 170–184. Springer, 2006. . URL https://doi.org/10.1007/11762256_15.

[49] Matthew Horridge, Bijan Parsia, and Ulrike Sattler. Explaining inconsistencies in OWL ontologies. In Lluís Godo and Andrea Pugliese, editors, *Scalable Uncertainty Management, Third International Conference, SUM 2009, Washington, DC, USA, September 28-30, 2009. Proceedings*, volume 5785 of *Lecture Notes in Computer Science*, pages 124–137. Springer, 2009. . URL https://doi.org/10.1007/978-3-642-04388-8_11.

[50] Matthew Horridge. *Justification based explanation in ontologies*. PhD thesis, University of Manchester, UK, 2011. URL http://www.manchester.ac.uk/escholar/uk-ac-man-scw:131699.

[51] 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.

[52] 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.

[53] 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.

[54] John Brooke. Sus: a retrospective. *Journal of usability studies*, 8(2), 2013.

[55] Eleni Ilkou and Sören Auer. Smart learning applications with the combination of the open research knowledge graph and teaching knowledge graphs. In *Proceedings of the 1st GOBLIN Workshop on Knowledge Graph Technologies (GOBLIN25)*, Leipzig, Germany, June 2025. . URL https://doi.org/10.5281/zenodo.16912917.

[56] Chris Davis Jaldi, Eleni Ilkou, Noah Schroeder, and Cogan Shimizu. Education in the era of neurosymbolic ai. *Journal of Web Semantics*, 85:100857, 2025.

[57] CD Jaldi, A Saini, S Zhang, N Schroeder, C Shimizu, and E Ilkou. Small, private language models as teammates for educational assessment design. In *Proceedings of the International Conference on Artificial Intelligence in Education (AIED 2026)*, 2026.

[58] Eleni Ilkou, Thomai Alexiou, Grigoris Antoniou, and Olga Viberg. Dyslexia and ai: Do language models align with dyslexic style guide criteria? In *International Conference on Artificial Intelligence in Education*, pages 32–47. Springer, 2025.

[59] Eleni Ilkou, Maria Angela Pellegrino, and Olga Viberg. Feedback dynamics in agentic llm-based text accessibility. In *Proceedings of the International Conference on Artificial Intelligence in Education (AIED 2026)*, 2026.

[60] Eleni Ilkou, Stephan Linzbach, and Jonas Wallat. Hybrid evaluation of socratic questioning for teaching. In *Proceedings of ISWC*, volume 24, 2024.

[61] 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.

[62] 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.

[63] 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.

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

[65] 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.

[66] 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.