Project title	Augmented Intelligence for Pedagogically Sustained Training and Education
Project acronym	augMENTOR
Grant agreement No.	101061509
Start date of project	01/01/2023
Duration of project	36 months
Project website	https://augmentor-project.eu/

D5.1 Data Mapping, Fusion and Orchestration Toolbox

Related work package	WP5
Document reference	D5.1
Related deliverables	D5.2, D2.2
Status	Final
Version	1.0
Due date	31/08/2024
Submission date	21/08/2024
Lead partner	NVCR
Contributing partners	UNIGR, MSX
Reviewers	UNIGR, MSX
Keywords	Knowledge Graphs, Ontologies, Data Fusion

Dissemination level	
✔	PU: Public
	Sen: Sensitive
	R-UE/EU-R: EU Classified
	S-UE/EU-S: EU Classified
	C-UE/EU-C: EU Classified

Funded by the European Union

The augMENTOR project (GA No. 101061509) is funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. The European Union cannot be held responsible for them.

Document information

List of contributors	
Name	**Partner (short name)**
Dimitris Charalampakis	NVCR
George Domalis	NVCR
Nikos Alimpertis	NVCR
Dimitris Tsakalidis	NVCR
George Koukoulas	UNIGR
Gerasimos Spyratos	MSX

Document History			
Version	**Date**	**Change editors**	**Changes**
0.1	03/04/2024	Dimitris Charalampakis (NVCR)	TOC initiation
0.2	10/05/2024	Nikos Alimpertis (NVCR), Koukoulas Georgios (UNIGR)	First draft
0.3	15/06/2024	Gerasimos Spyratos (MSX), George Domalis (NVCR)	Second draft
0.4	25/07/2024	Nikos Alimpertis (NVCR)	Draft ready for internal review
0.5	13/08/2024	Eleftheria Tsourlidaki (UNIGR) Michael Spyratos (MSX)	Comments added by internal reviewers
0.6	17/08/2024	Dimitris Charalampakis (NVCR)	Reviewed version
0.7	18/08/2024	Dimitris Tsakalidis (NVCR)	Pre-final version
0.8	19/08/2024	Eleftheria Tsourlidaki (UNIGR)	Pre-final check
0.9	21/08/2024	George Garofalakis (UNI)	Final check
1.0	21/08/2024	George Garofalakis (UNI)	Release and submission

Quality control		
Role	**Partner (short name)**	**Approval date**
Internal reviewers	UNIGR, MSX	13/08/2024
Deliverable leader	NVCR	18/08/2024
Quality manager	UNIGR	19/08/2024
Project coordinator	UNI	21/08/2024

Table of contents

List of figures

List of tables

List of acronyms

Acronym	Description
AI	Artificial Intelligence
augMENTOR	Augmented Intelligence for Pedagogically Sustained Training and Education
CSCL	Computer-Supported Collaborative Learning
ETL	Extract, Transform, Load
FP Growth Algorithm	Frequent Pattern Growth Algorithm
GNNs	Graph Neural Networks
JSON	JavaScript Object Notation
KG	Knowledge Graph
ML	Machine Learning
LMS	Learning Management System
MOOC	Massive Open Online Course
NER	Named Entity Recognition
NLP	Natural Language Processing
NLQ	Natural Language Query
OWL	Web Ontology Language
POS	Part-Of-Speech
RDF	Resource Description Framework
RDFS	Resource Description Framework Schema
QA	Question - Answering
SCORM	Sharable Content Object Reference Model

Executive summary

This deliverable D5.1 reports on the activities undertaken in task T5.1 (Data and Educational Resources Mapping, Fusion and Orchestration). T5.1 focuses on designing and implementing a set of interoperable software services that enable and orchestrate the required data and educational resources mapping and fusion for the needs of the augMENTOR solution. Data and educational resources were derived from the learning management systems and educational resources deployed by the pilots. By using mapping and fusion techniques we accomplished their efficient representation through semantically-rich knowledge graphs that exploit the advantages of complementarity while removing redundancies created by different modalities. D5.1 reports on the design of two knowledge graphs, one for each of the learning management systems used by the pilots. To design the augMENTOR Knowledge Graphs our team performed an extensive literature review on knowledge graphs and followed a comprehensive step-by-step methodology.

1 Introduction

1.1 Scope and objective of the deliverable

The scope of this deliverable is to present the work done in T5.1 on

a) the data mapping, fusion and orchestration software services developed
b) the content, semantics and structure of the augMENTOR Knowledge Graph (KG).

Chapter 2 presents the methodology followed for developing the augMENTOR Knowledge Graphs. As our pilots deploy two different Learning Management Systems (LMS)(Moodle for UPATRAS, IASIS, EASD and Try-Hack-Me for KTU) our team developed two separate KGs, one for each LMS. Chapter 3 presents the KG for the Moodle LMS while Chapter 4 presents the KG for the Try-Hack-Me LMS.

1.2 Relation to other tasks and deliverables

D5.1 draws input from deliverable D2.2, particularly chapter 5 which presents the main components of the augMENTOR software architecture in detail. The work presented is going to serve as input for T5.2, T5.3 and T5.4 as well as for D5.2 -"The AI-boosted augMENTOR platform".

1.3 Knowledge graphs in a nutshell and their deployment in augMENTOR

KGs have become instrumental in the big data era, helping to organise complex knowledge in structured and insightful ways. These graphs utilise nodes to represent entities and edges to illustrate the relationships between them, capturing the semantics of data through schemas, ontologies, and logic. KGs are crucial for various applications, such as enhancing search engines, enabling knowledge discovery, developing explanation services of the Artificial Intelligence (AI) outcomes and facilitating data analysis. The integration power of KGs is evident as they unify data from diverse sources, leading to the discovery of previously unseen relationships and insights. KGs in education play a pivotal role in adaptive and personalised learning, curriculum design, concept mapping, semantic search, and broader applications. They customise learning experiences, support curriculum planning, visualise complex concepts, and enhance search and recommendation systems. Annex 1 discusses the "state-of-the-art" in KGs, presenting the major findings of an extensive literature review of KGs within the education domain.

In the context of the augMENTOR project, data modelling is crucial for constructing KGs. This process entails defining the entities, attributes, and relationships that will populate the graph, ensuring interoperability through standards like Resource Description Framework (RDF) Schema and Web Ontology Language (OWL). This structured approach allows for comprehensive integration of diverse data sources, revealing hidden dependencies and patterns. Robust schema alignment and semantic enrichment further enhance the accuracy and utility of these graphs. The data framework of KGs is presented in Annex 2.

Storing and managing educational knowledge graphs involve using specialised databases and platforms designed to handle complex, interconnected datasets. Graph databases and RDF triple stores efficiently manage the graph structures, while Extract, Transform, Load (ETL) processes support integration from multiple sources. Through these meticulously maintained operations, educational knowledge graphs serve as dynamic resources, aiding educators, policymakers, and researchers in data-driven decision-making and comprehensive educational analysis. Annex 3 presents the major operations of KGs.

The process of deploying KGs and ontologies for educational platforms such as MOODLE and Try-Hack-Me involves several meticulous and interconnected steps. Initially, relevant data is collected and preprocessed to ensure quality and consistency. This is followed by the critical phase of ontology design, where entities and their relationships are defined. Once the data is cleaned and the ontology is established, the next step is to map and integrate this data into the KG, ensuring alignment with the defined schema. Chapter 2 highlights the methods conducted for the deployment of augMENTOR KGs.

1.4 Terminology

The foundation of KGs is associated with some main terms and concepts that are essential for understanding, building, and utilising KGs to represent and leverage complex knowledge in a structured and interconnected manner. The key terms and concepts commonly related to KGs mentioned in this deliverable, are presented in the table below:

Table 1. Terminology and concepts in KGs

Name	Description
Entity	An entity represents a specific object, concept, or instance in the KG. Entities can be tangible objects (e.g., a person, a city) or abstract concepts (e.g., a concept in a domain, a class). Entities are essentially the fundamental building blocks of a graph, and in this project each dataset is treated as a distinct entity.
Attribute	Attributes are properties or characteristics associated with entities in the KG. They provide additional information about the entities. For example, a person entity may have attributes like name, age, and occupation. The metadata describing each dataset are treated as attributes.
Relationship	Relationships define the connections or associations between entities in the KG. They represent the semantic links and dependencies between entities. Datasets can be associated with each other based on different similarity metrics. For example, two datasets can be considered to be related if their titles are semantically close ("Associated by title") or because their descriptions are close ("Associated by description").
Classes and types	Classes and types categorise entities into groups based on their shared characteristics or attributes. In ontologies, classes represent concepts or categories within the educational domain.
Taxonomies	Taxonomies are hierarchical structures that organise entities into parent-child relationships based on their similarities and differences. Taxonomies provide a way to classify entities into broader and more specific categories.
Triplet	A triplet is the fundamental building block of a KG and represents a single piece of information. It consists of three components: subject, predicate, and object. The subject and object are entities, while the predicate represents the relationship between them. Each relationship is represented as a triplet between two entities along with a relationship type. Inherently, the metrics used to establish these triplets are not always symmetrical, meaning that fun(a,b)≠fun(b,a) for some functions. To go around this problem, since associations have to be symmetrical by definition (if dataset A is related to dataset B, then it follows that dataset

	B is related to dataset A), Max(fun(A,B), fun(B,A)) is used to decide whether datasets A and B are associated.
Graph	A graph is a collection of entities, attributes, relationships, and triples organised in a network structure. It represents the overall structure and connections in the KG. Nodes in the graph represent entities, and edges represent relationships between the entities.
Ontology	An ontology defines the schema or vocabulary for the KG. It specifies the types of entities, relationships, and attributes that can exist in the graph. The ontology provides a formal representation of the domain knowledge and helps ensure consistency and interoperability within the KG.
Ontology hierarchies	Ontologies often include hierarchies that define relationships between classes and types. These hierarchies help establish a structured and logical organisation of knowledge within the graph.
Domain specific vocabulary	KGs often incorporate domain-specific vocabularies and terminologies relevant to Learning Management Systems. These vocabularies ensure consistency and clarity in data representation.
Data sources	KGs include information about the sources of data, such as data providers and data acquisition methods. This metadata helps users trace the origin of information.
Semantic annotations	Annotations within the ontology may provide additional context or metadata about entities, relationships, or classes. Semantic annotations enhance the understanding and interoperability of data.
Semantic reasoning	Semantic reasoning involves drawing logical inferences and making deductions based on the semantic relationships and rules defined in the KG. It enables the discovery of implicit knowledge, filling in missing information, and performing advanced queries and analysis.
Linked data	Linked Data is a set of best practices and standards for publishing and interconnecting structured data on the web. It emphasises the use of Uniform Resource Identifiers (URIs) to identify and link data elements, enabling the creation of a web of interconnected data sources.

2 Design of the augMENTOR Knowledge Graphs

2.1 Methodology

Our team followed a set of clear steps to develop the project's KGs. As it was mentioned above, we deliver a separate KG and ontology for each of the LMSs deployed in the framework of the project. The development of a common ontology for both LMSs was not pursued due to several factors. Firstly, the nature of the courses differed significantly, requiring tailored approaches to data representation and knowledge modelling. Secondly, while common entities were indeed utilised to assist in the pilot courses, the specific needs and contexts of each course made it impractical to develop a one-size-fits-all ontology. Lastly, the differing structures of learning activities further complicated the creation of a unified ontology, as each LMS featured unique methodologies and frameworks for instructional design. Consequently, the augMENTOR project has opted for a more flexible approach that accommodates the unique characteristics of each LMS while still leveraging shared entities to facilitate cross-pollination of ideas and practices. This careful consideration ensures that the KGs developed are not only effective but also aligned with the specific educational goals and contexts of each course.

Deploying the KGs for MOODLE and Try-Hack-Me involved several interconnected steps. More analytically the following methodological steps were followed.

a. Ontology design

The design of the ontology was a critical step in defining the structure of the KGs. For MOODLE, essential entities were identified, including Teacher, Learner, Course, Module, Activity (such as Forum, Sharable Content Object Reference Model (SCORM), Assignment, and Quiz), Resource, and Profile. Each entity was meticulously defined with specific attributes; for example, a Teacher entity included attributes like username, country, institution, and user ID, while a Course entity included description, title, and ID. Furthermore, the relationships between these entities were established to reflect the pedagogical and organisational structure of the educational content. Key relationships identified including CREATE (linking Teachers to Courses), STUDY (linking Learners to Resources), REGISTERED (linking Learners to Courses), and PARTICIPATE (linking Learners to Activities). Additional relationships such as HAS_MODULE, BELONGS, HAS_ACTIVITY, and HAS_RESOURCE were defined to map the hierarchical nature of courses and activities.

Designing the ontology of Try-Hack-Me involved identifying key entities specific to the learning platform, such as Room, Tutorial, Video, Task, Question, Module, Path, and Profile. Each entity was defined with attributes that described its properties; for instance, a Room entity included attributes like difficulty, timeToComplete, title, ID, description, type, and code. Relationships between these entities were also defined, such as HAS_TASK (connecting Rooms to Tasks), HAS_QUESTION (connecting Tasks to Questions), HAS_ROOM (connecting Modules to Rooms), BELONGS (connecting Learners to Profiles), and HAS_MODULE (connecting Paths to Modules). These relationships captured the hierarchical and interconnected nature of the platform's educational content.

b. Data acquisition and preprocessing

The initial step involved the collection of relevant data from MOODLE, primarily focusing on course materials, user interactions, activity logs as well as related metadata. This phase ensured that the datasets included diverse information such as course descriptions, user profiles (teachers and learners), module specifications, and activity logs encompassing forums, quizzes, assignments, and SCORM packages. Same data acquisition process was followed for "Try-Hack-Me", focusing on various elements such as room details, tasks, questions, tutorials, and user progress. After data extraction, the data underwent a thorough cleaning process to remove inconsistencies, duplicates, and missing values. Standardisation of data formats, such as ensuring uniform date representations and consistent naming conventions, was implemented to facilitate seamless integration.

c. Data mapping and integration

The next phase involved mapping and integrating the preprocessed data into the ontology for the development of the KG. Schema alignment was crucial to ensure that data attributes matched the established ontology. This involved aligning attribute names, data types, and semantics to the predefined graph schema. Semantic enrichment was also performed by linking entities to external ontologies, enhancing the contextual accuracy and overall quality of the KG. Techniques such as RDF Schema and OWL were employed to formalise the data model within the graph structure.

d. Storage and management

Selecting an appropriate graph database was essential for efficiently storing and managing the KG. Neo4j was chosen due to its capability to handle complex relationships and interconnected data structures effectively. Data was then transformed into RDF triples and loaded into the graph database, ensuring that the structure was compatible with the standardised schema and ontology. Version control mechanisms were established to track changes and ensure data lineage, which aids in maintaining the reproducibility and integrity of the KG over time.

e. KGs construction

With the MOODLE ontology and schema in place, the KG was constructed by creating nodes representing each entity (e.g., TEACHER, LEARNER) and populating them with the corresponding attributes. Relationships between entities were established to reflect the educational structure and user interactions. For instance, edges were created to denote relationships such as a Teacher creating a Course, a Learner participating in an Activity, and a Course containing Modules. Each relationship was assigned properties to capture detailed data, such as grades received by learners, time spent on activities, the number of attempts on quizzes, and other relevant metrics. Nodes representing each entity (e.g., ROOM, TASK) were created and populated for "Try-Hack-Me" with the relevant attributes. Relationships between nodes were established to reflect the structure and flow of the educational content. For example, edges were created to denote that a Room contained specific Tasks and a Task included specific Questions. Properties were assigned to these relationships to capture details such as the order of rooms within modules, enhancing the granularity and usability of the knowledge graph.

f. Querying and analysis

To extract and analyse valuable insights from the KGs, querying techniques using SPARQL and Cypher were used. These queries enabled the extraction of meaningful information such as learners' performance metrics, engagement levels in various activities, and the effectiveness of different educational resources. Visualisation and exploration tools like Neo4j's built-in visualisation features were utilised to create interactive visualisations, facilitating deeper insights and better comprehension of the data structure and relationships within the KG.

g. Continuous updates and maintenance

The KG requires continuous updates to remain relevant and accurate. Automated ETL (Extract, Transform, Load) processes were set up using tools such as Apache NiFi to ensure real-time data updates and consistent integration of new information. Regular validation checks were performed to maintain data integrity and quality, addressing any inconsistencies or anomalies promptly. Continuous monitoring and maintenance ensured that the KG evolved with new data, pedagogical trends, and changes in the educational landscape.

2.2 Deployment and publication

Following the steps mentioned above our team designed the project's two (2) KGs. In the following chapters we present the two dedicated ontologies produced. Both these ontologies have been turned to Neo4j databases and have been deployed to construct the KGs. The related code is published in a dedicated repository: https://github.com/novelcore/EU_augMENTOR.

Funded by
the European Union

3 Ontology based on Moodle

Figure 1 presents the ontology (nodes, relationships and properties) within the Moodle-based educational environments, which will be operated during the project's pilot phases from UPATRAS, EASD and IASIS and will be populated with educational resources, content and learning outcomes of the pilot users, namely educators and learners.

It includes entities like teachers, learners, courses, modules, activities (forum, SCORM, assign, quiz), and resources, each with specific properties, that have been derived based on terms used by Moodle. Relationships define how these entities interact, such as an educator creating a course, a learner participating in activities, and the structure of courses into modules and activities. The properties of relationships, particularly for participation in activities, provide detailed data on learner engagement and performance.

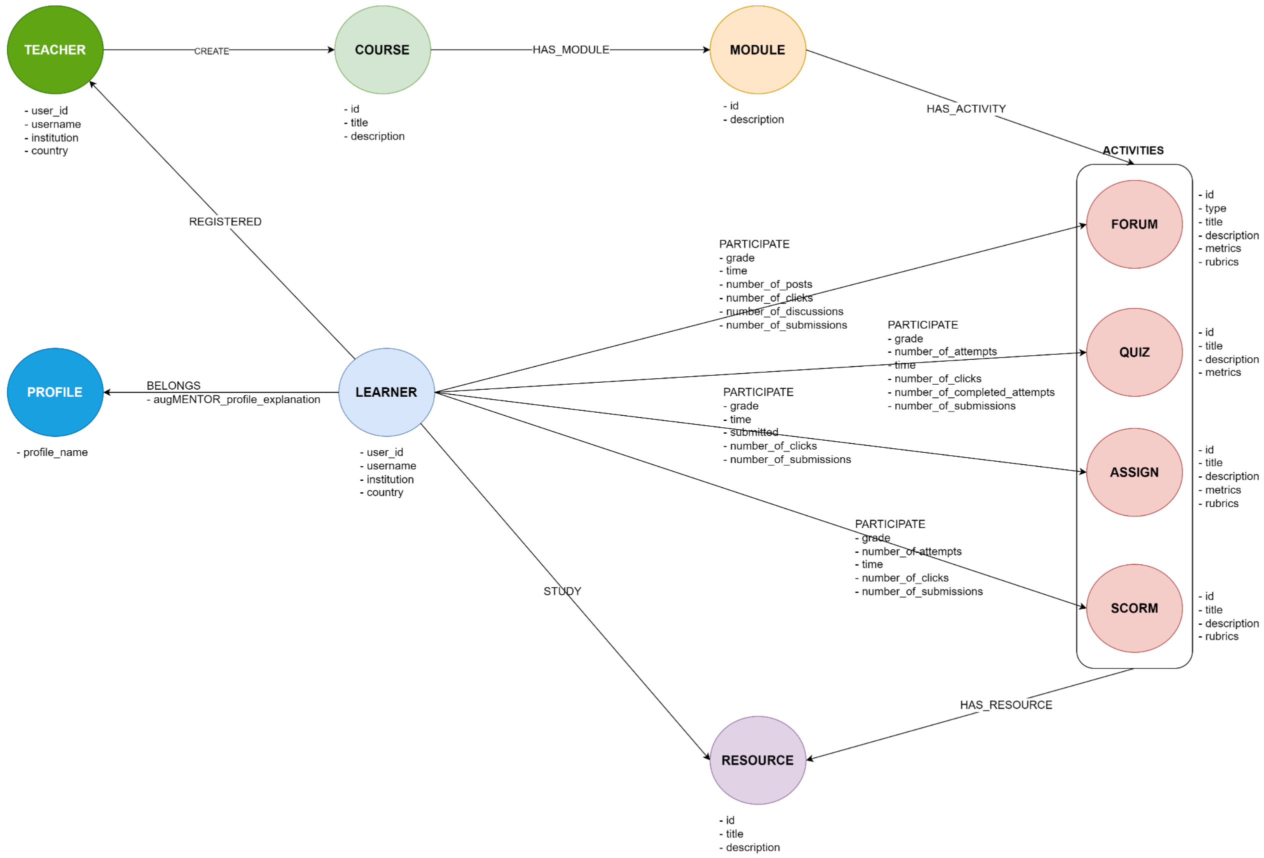

Figure 1. Moodle-based Knowledge Graph ontology

3.1 Nodes and their properties

Next, we provide a brief description of the nodes of the ontology as well as their properties.

- TEACHER
  - username (STRING): The username of the teacher.
  - country (STRING): The country where the teacher is located.
  - institution (STRING): The institution the teacher is affiliated with.
  - user_id (INTEGER): A unique identifier for the teacher.
- LEARNER
  - username (STRING): The username of the learner.
  - country (STRING): The country where the learner is located.
  - institution (STRING): The institution the learner is affiliated with.
  - user_id (INTEGER): A unique identifier for the learner.
- COURSE
  - description (STRING): A description of the course.
  - title (STRING): The title of the course.
  - id (INTEGER): A unique identifier for the course.
- MODULE
  - title (STRING): The title of the module.
  - id (INTEGER): A unique identifier for the module.
- ASSIGN
  - description (STRING): A description of the activity.
  - id (STRING): A unique identifier for the activity.
  - title (STRING): The title of the activity.
  - rubrics (STRING): Rubrics of the activity.
  - metric (STRING): One of the 4Cs ['Creativity', 'Collaboration', 'Critical thinking', 'Communication'] or 'Basic skills'.
- SCORM
  - description (STRING): A description of the activity.
  - id (STRING): A unique identifier for the activity.
  - title (STRING): The title of the activity.
  - rubrics (STRING): Rubrics of the activity

  - FORUM
  - description (STRING): A description of the activity.
  - id (STRING): A unique identifier for the activity.
  - title (STRING): The title of the activity.
  - forum_type (STRING): The type of forum (e.g., general, Q&A).
  - rubrics (STRING): Rubrics of the activity.
  - metric (STRING): One of the 4Cs ['Creativity', 'Collaboration', 'Critical thinking', 'Communication'] or 'Basic skills'.
- QUIZ
  - description (STRING): A description of the activity.
  - id (STRING): A unique identifier for the activity.
  - title (STRING): The title of the activity.
  - quiz_max_grade (INTEGER): The maximum grade achievable in a quiz activity.
  - rubrics (STRING): Rubrics of the activity.
  - metric (STRING): One of the 4Cs ['Creativity', 'Collaboration', 'Critical thinking', 'Communication'] or 'Basic skills' and 'Quiz'
- FORUM
  - description (STRING): A description of the activity.
  - id (STRING): A unique identifier for the activity.
  - title (STRING): The title of the activity.
  - forum_type (STRING): The type of forum (e.g., general, Q&A).
  - metric (STRING): One of the 4Cs ['Creativity', 'Collaboration', 'Critical thinking', 'Communication'] or 'Basic skills'
  - rubrics (STRING): Rubrics of the activity
- RESOURCE
  - description (STRING): A description of the resource.
  - id (STRING): A unique identifier for the resource.
- PROFILES
  - profile_name (STRING): profile name

3.2 Relationships and their properties

Below we present a brief description of the relationships of the ontology and their properties.

- Relation: CREATE
  - Description: Connects a TEACHER node to a COURSE node. Represents the relationship that a teacher has created a course.
  - Properties: None
- Relation: STUDY
  - Description: Connects a LEARNER node to a RESOURCE node.
  - Represents the relationship that a learner is studying or using a resource.
  - Properties: None
- Relation: REGISTERED
  - Description: Connects a LEARNER node to a COURSE node. Represents the relationship that a learner has registered for a course.
  - Properties: None
- Relation: PARTICIPATE
  - Description: Connects a LEARNER node to an ACTIVITY (FORUM, QUIZ, ASSIGN, SCORM) node. Represents the relationship that a learner is participating in an activity.
  - Properties:
    - grade (FLOAT): The grade received by the learner in the activity.
    - time (INTEGER): The time spent by the learner on the activity.
    - number_of_attempts (INTEGER): The number of attempts made by the learner in a quiz activity.
      - number_of_completed_attempts (INTEGER): The number of completed attempts made by the learner in a quiz/scorm activity.
    - number_of_posts (INTEGER): The number of posts made by the learner in a forum activity.
    - number_of_clicks (INTEGER): Indicates the number of clicks each learner conducted for an activity

        - number_of_discussions (INTEGER): Indicates the number of discussions each learner created in a forum activity.
        - number_of_submissions (INTEGER): Indicates the number of submissions each learner conducted for an activity
- Relation: HAS_MODULE
    - Description: Connects a COURSE node to a MODULE node. Represents the relationship that a course contains one or more modules.
    - Properties: None
- Relation: BELONGS
    - Description: Connect the learner with the profile he/she belongs
    - Properties:
        - augMENTOR_profile_explanation (STRING): provide an explanation why a learner has been assigned to an augMENTOR profile.
- Relation: HAS_ACTIVITY
    - Description: Connects a MODULE node to an ACTIVITY (SCORM, FORUM, ASSIGN, QUIZ) node. Represents the relationship that a module contains one or more activities.
    - Properties: None
- Relation: HAS_RESOURCE
    - Description: Connects an ACTIVITY node to a RESOURCE node.
    - Represents the relationship that an activity includes one or more resources.
    - Properties: None

4 Ontology based on Try-Hack-Me

Figure 2 depicts the developed ontology for organising the educational content and resources from "Try-Hack-Me" MOOC, which KTU will be operating for during the pilot phase of the augMENTOR project. It defines different types of nodes (ROOM, TUTORIAL, VIDEO, TASK, QUESTION, MODULE, PATH) each with specific properties that describe their attributes. Relationships between these nodes indicate how they are connected and organised, with properties on some relationships specifying the sequence or order of the connections. This ontology provides a detailed blueprint for how educational content is structured and navigated, ensuring a coherent and logical flow of information for users.

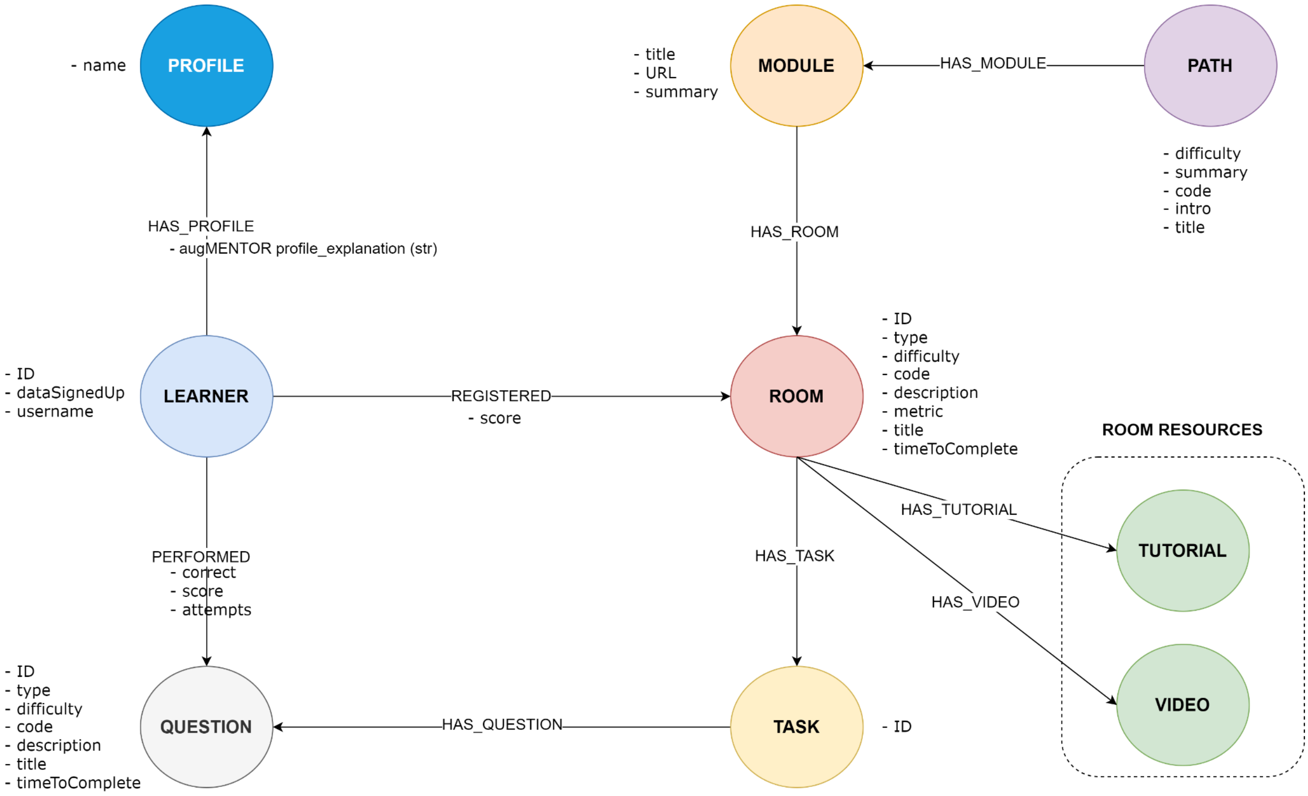

Figure 2. Try-Hack-Me Knowledge Graph ontology

4.1 Nodes and their properties

Follows a brief description of the nodes of the ontology and their properties.

- ROOM
  - difficulty (STRING): Level of difficulty for the content in the room (e.g., easy, medium, hard).
  - timeToComplete (STRING): Estimated time required to complete the room's content.
  - title (STRING): The title of the room.
  - ID (STRING): Unique identifier for the room.
  - description (STRING): Detailed description of the room.
  - type (STRING): The type or category of the room.
  - code (STRING): Any specific code associated with the room, possibly for access or reference.
  - metric: One of "Critical thinking" and "Communication"
- TUTORIAL
  - tutorials (STRING): URLs or identifiers for tutorials available for rooms.
- VIDEO
  - videos (STRING): URLs or identifiers for video content available for rooms.
- TASK
  - ID (STRING): Unique identifier for the task.
- QUESTION
  - questionNo (STRING): The number or identifier of the question within a task.
  - answer (STRING): The correct answer to the question.
  - extraPoints (STRING): Any additional points that can be earned by answering the question.
  - hint (STRING): Hints or tips to help solve the question.
  - question (STRING): The text of the question itself.
- MODULE
  - title (STRING): The title of the module.
  - moduleURL (STRING): URL linking to the module's content.
  - summary (STRING): A brief summary of the module.
- PATH
  - title (STRING): The title of the path.
  - difficulty (STRING): Overall difficulty level of the path.
  - summary (STRING): A brief summary of the path.

  - intro (STRING): An introduction to the path.
  - code (STRING): Specific code associated with the path, possibly for access or reference.
- PROFILES
  - profile_name (STRING): profile name

4.2 Relationships and their properties

Follows a brief description of the relationships of the ontology and their properties.

- Relation: HAS_TASK
  - Description: Connects ROOM nodes to TASK nodes.
  - Meaning: Indicates that a specific room contains a certain task.
  - Properties: None specified for this relationship.
- Relation: HAS_QUESTION
  - Description: Connects TASK nodes to QUESTION nodes.
  - Meaning: Indicates that a specific task contains a certain question.
  - Properties: None specified for this relationship.
- Relation: REGISTERED
  - Description: Connect the learner with the rooms he/she registered
  - Properties: score (INTEGER): The score achieved by the learner
- Relation: HAS_ROOM
  - Description: Connects MODULE nodes to ROOM nodes.
  - Meaning: Indicates that a module includes or is associated with a particular room.
  - Properties: order (INTEGER): The order in which the rooms appear within the module.
- Relation: BELONGS
  - Description: Connect the learner with the profile he/she belongs
  - Properties:
    - augMENTOR_profile_explanation (STRING): provide an explanation why a learner has been assigned to an augMENTOR profile.
- Relation: HAS_MODULE
  - Description: Connects PATH nodes to MODULE nodes.
  - Meaning: Indicates that a specific path includes a certain module.
  - Properties: order (INTEGER): The order in which the modules appear within the path.

5 Conclusions

The objective of this deliverable was to present the work done in T5.1 focusing particularly on the data mapping, fusion and orchestration software services developed as well as the the content, semantics and structure of the augMENTOR KGs.

The augMENTOR project has adopted and is currently applying the core principles and benefits of KGs in the educational domain through the establishment of robust frameworks and methodologies for data modelling and integration. These KGs enable the exploration and discovery of hidden relationships within vast datasets, thereby facilitating personalised learning, learning analytics and intelligent recommendation.

The project underscores the significance of KGs in organising educational data, improving data interoperability, and enhancing query performance through semantic web technologies and machine learning.

Our team developed two (2) separate KGs deploying two (2) different ontologies, which were designed based on the two different LMSs to be used in the pilots, namely Moodle for UPATRAS, IASIS, EASD pilots and Try-Hack-Me for the KTU pilot.

Following the delivery of the KGs, they will feed the machine learning pipelines to facilitate the deployment and use of the augMENTOR AI-boosted solution. Thus, the work presented in this deliverable is going to serve as input for T5.2, T5.3 and T5.4 as well as for D5.2 -"The AI-boosted augMENTOR platform".

Annex 1 - Overview and "state-of-the-art" of knowledge graphs

In the era of big data and interconnected information, knowledge graphs have emerged as a powerful tool for representing and organising complex knowledge in a structured and meaningful way. KGs are a way of representing and integrating data using a graph-structured data model or topology, where nodes are entities and edges are relationships between them. They also capture the semantics or meaning of the data, using schemas, ontologies, and logic.

KGs can be used for various purposes, such as enhancing search engines, powering question-answering systems, enabling knowledge discovery, and facilitating data analysis. KGs can be built from different sources, such as Wikipedia, web pages, public datasets, or domain-specific corpora.

The foundation of a KG lies in the concept of linked data, which emphasises the principles of interconnectedness and machine-readable semantics. KGs employ a schema, or ontology, that defines the types of entities, attributes, and relationships that can exist within the graph. The graph structure enables efficient navigation and exploration of the knowledge, facilitating both breadth-first and depth-first traversals. This graph-centric approach allows for rich semantic modelling, where entities can be connected through various types of relationships, forming a web of interconnected knowledge. Such relationships can represent simple associations, hierarchical taxonomies, temporal dependencies, or even complex semantic associations that capture nuanced connections between entities.

Benefits and limitations of knowledge graphs

KGs offer a multitude of benefits across various domains by integrating data from structured, semi-structured, and unstructured sources into a unified framework. This integration enables the discovery of hidden relationships and insights that might not be apparent from individual datasets. KGs support the modelling of contextual information, allowing for precise and personalised knowledge representation. They enhance data interoperability and information sharing by adhering to standard semantic web technologies and open standards [1]. Thus, they facilitate data integration, collaborative knowledge construction, and advanced analytics.

One of the major benefits of KGs is their ability to link and harmonise data from diverse sources, fostering data sharing and collaboration among organisations. They provide a comprehensive view of information, enabling insights from hierarchical data, representing sequences of decisions, and portraying characteristics of entities and their relationships [2]. This capability is crucial for understanding macro relationships and dynamics in business data. The flow of information can be visualised to show how data is transformed, processed, and consumed by different actors or systems. Additionally, KGs facilitate direct and expandable knowledge representation [3]. They perform better in data-intensive environments by allowing easier schema expansion, improved knowledge extraction and representation, masking of data complexity, integration from external sources, and exploitation of graph algorithms [4]. Furthermore, KGs effectively represent both relationships between entities and information about individual entities [5].

By combining structured and unstructured data, KGs support richer data services and enable the extraction and discovery of deeper patterns. This allows end-users to make informed decisions by providing relevant facts and contextualised answers [6]. KGs have the advantage of summarising relationships between entities concisely, illustrating how everything is related at a macro level. They also improve the time to knowledge with the application of machine learning algorithms and natural language processing techniques. Graph Neural Networks (GNNs) [6], graph embedding methods [7], graph-based neural network models [8], random walk-based techniques [9], and semantic web technologies [10] enhance the exploration and querying of connections within the network. Additionally, KGs reduce the complexity and cost of data preparation by using a common schema and ontology to map and align data from different sources.

While KGs offer powerful capabilities for representing and organising knowledge, they are not without limitations. Integrating data from multiple sources can be complex and time-consuming due to diverse data formats, varying data quality, and conflicting representations [11]. Data cleansing, resolving semantic heterogeneity, and aligning ontologies require significant effort and expertise. As the size and complexity of KGs increase, scalability and query performance become crucial concerns. Efficient queries and traversals on large-scale KGs require specialised indexing techniques, distributed processing frameworks, or graph databases. Balancing scalability and performance remains a challenge, especially for dynamic and evolving knowledge.

Another limitation is incomplete knowledge representation. Capturing the entirety of knowledge within a domain is challenging, resulting in gaps and limitations during complex reasoning or inference tasks. Additionally, uncertainty and ambiguity in knowledge representation can arise due to conflicting or incomplete information sources. KGs are dynamic and need continuous updates to reflect evolving knowledge. This requires significant effort to ensure data consistency, schema evolution, and data governance. Maintaining the integrity and accuracy of a KG over time is particularly challenging in dynamic or rapidly changing environments. Achieving accurate and consistent semantic representations of entities and relationships is also complex. Resolving the semantic gap between natural language and structured knowledge, disambiguating entities, understanding context, and capturing nuanced semantics are ongoing challenges. These require advances in natural language processing, entity disambiguation, and semantic understanding.

Despite these limitations, KGs remain a valuable tool for knowledge representation and organisation. Ongoing research and advancements continue to address these challenges, improving the effectiveness of KGs in various domains.

Importance of knowledge graphs in education

This section explores substantial academic efforts leveraging KGs to enhance education. It reviews significant achievements and challenges across five core domains: Adaptive and Personalized Learning, Curriculum Design and Planning, Concept Mapping and Visualization, Semantic Search and Question Answering, and other miscellaneous applications. Each domain underscores the potential of KGs to deliver tailored, data-driven educational experiences. This review meticulously outlines methodologies, functionalities, knowledge extraction and construction techniques, resource requirements, evaluation criteria, and limitations.

a. Adaptive and personalised learning

Personalised learning customises instruction based on each student's strengths, needs, skills, and interests. KGs facilitate this by structuring and organising information about learners and their progress, thereby enhancing the adaptation of instruction. For instance, [12-33] discusses a personalised learning resource recommendation system that utilises KGs. Authors outlined the problem of personalised learning resource recommendation in online education, underscored the significance of KGs, presented an algorithm for personalised recommendations, and elaborated on the construction of a personalised learning resource recommendation system.

Lu et al. introduced "RadarMath," an intelligent tutoring system providing automatic grading and personalised learning support [13]. This system uses two grading models to cater to individual learner preferences and progress. Similarly, Bai et al. demonstrated an adaptive learning model using KGs and an improved FP-Growth algorithm for data mining in online learning systems [14]. This model addresses challenges such as cognitive overload and confusion in online learning environments.

Further applications of KGs in adaptive learning involve learner profiling, Machine Learning (ML) integration for performance prediction, and the development of early warning systems for at-risk students. For example, Albreiki et al. applied KGs and machine learning to predict students' academic performance, identify students at risk of failing a course, and deliver personalised interventions [15]. Another study demonstrated the efficacy of integrating deep learning with KGs to provide interpretive early warnings in interactive learning environments [16]. Additionally, KGs have been utilised for clustering students, recommending courses, and suggesting learning paths based on diverse learner features [17-19]. For example, Xue et al. explored clustering students and recommending courses based on semantic and statistical dimensions, which confirmed the strategy's practicality and efficiency in different educational contexts [17].

Incorporating KG technology for adaptive education, personalised learning, and education recommendations is discussed and conveyed across various references [20-26], demonstrating consistent improvements in educational outcomes. The combination of KGs and machine learning has shown improvements in predicting students' performance and identifying those at risk of course failure.

b. Curriculum design and planning

Curriculum design encompasses the systematic development of educational programs, aligning learning objectives, content, instructional methods, assessment strategies, and resources to foster meaningful learning among students. Incorporating KGs into curriculum design and planning offers a data-driven, learner-centric approach that enhances the quality, relevance, and effectiveness of educational offerings. For example, Yang et al. constructed a KG to aid water conservancy students in understanding relationships between various educational components [27]. This KG provided a structured representation of educational data, enhancing efficient learning and knowledge retrieval.

Funded by
the European Union

EducOnto and EduKG were developed to assist in university curriculum recommendation, modelling the transition from high school to university by incorporating concepts like student, curriculum, major, and specialty [28]. Zablith designed a KG aiming to establish semantic links between social media content and formal course entities, enabling students to integrate and access interdisciplinary social media content within formal courses [29]. This integration created a comprehensive and interconnected learning environment.

Further applications involve multimodal KGs, such as ModelsKG, which enhance intelligent education through integrated multi-modal knowledge [30]. ModelsKG integrates PaddleOCR and DeepKE to create a structured and intelligent framework for linking and reorganising multi-modal knowledge, facilitating intelligent search, link construction, quantitative analysis, and intelligent recommendation. Specific KGs like Course KG and Teacher KG are constructed to facilitate course-teacher matching, optimise course offerings, and ensure curriculum coherence [31].

Additional studies have focused on developing visual KGs to aid in teaching programming languages, providing feedback in collaborative learning, and integrating theoretical and practical knowledge [32-39]. For instance, Gao et al. created a visual KG to display relationships between curriculum components such as chapters, sections, and knowledge points [40]. These visualisations enhance understanding and retention, making educational content more accessible and engaging.

c. Concept mapping and visualisation

The synergistic integration of concept mapping and visualisation techniques with KGs has garnered significant attention for its potential to revolutionise the learning experience. Concept mapping allows students to visually organise complex information, fostering deeper understanding and retention. The combination with visualisation techniques, which graphically represent intricate relationships, enhances cognitive engagement by breaking down complex concepts and extracting meaningful insights from interconnected information.

Various studies have explored the interplay between concept mapping, visualisation techniques, and KGs. For instance, Li et al. developed a multi-source education KG for college curricula in the major of Electronic Information, utilising visualisation techniques to enhance learning efficiency [41]. Similarly, Su et al. aimed to represent interconnected knowledge points within educational subjects, enabling non-linear teaching approaches, microteaching, learning assessment, learning navigation, and personalised resource recommendations [42].

In programming education, a KG was constructed to assist students with Python programming, helping them understand logical linkages between knowledge points based on their questions [43]. Additionally, KGs have been utilised to integrate educational content, such as the interdisciplinary framework connecting historical events, literature, geography, and other related concepts [88]. Integrating heterogeneous data sources into a unified KG has also been reported, with EDUKG representing knowledge topics, educational resources, and external heterogeneous data sources related to K-12 education [44].

KG visualisation in education enhances comprehension, exploration, and engagement with complex educational content. For instance, Tang et al. constructed a KG to visually represent the curriculum content system, integrating online learning resources to help students locate and understand the relationships between knowledge points [45]. Another study created a KG from unstructured course materials to facilitate learning in cybersecurity education, utilising a chatbot to answer learners queries [46].

d. Semantic search, QA, and recommender systems

KGs have significantly transformed semantic search and Question-Answering (QA) in education by enhancing search results and answers' precision, contextuality, and depth. Leveraging their interconnected structure and semantic richness, KGs offer contextual and relevant information for teachers and students. For instance, Fang et al. proposed a model that improves multi-hop reasoning in educational KGs using a bilinear graph neural network and a two-teacher knowledge distillation approach [47]. Nguyen et al. integrated an ontology-based model to organise and query educational content for a database systems course, laying the groundwork for an intelligent querying system [48].

Other studies have developed intelligent QA systems and course content recommendation systems based on KGs. For example, a KG was constructed to support an intelligent QA system for science and technology intermediary services, providing relevant answers based on structured data [49]. Another system integrated KG, smart Q&A, and big data technologies to enable knowledge modelling, knowledge acquisition, and teaching feedback [50]. In a prototype for remote schools, authors demonstrated the need for a KG to develop a QA system to support remote education [51].

Moreover, a Massive Open Online Course (MOOC)-KG was presented as a solution to enhance online learning resource utilisation by collecting and organising information from various MOOCs, platforms, universities, teachers, and courses [52]. A four-step approach using semantic web technologies was demonstrated to identify and evaluate prerequisite relationships between concepts accurately [53]. Additionally, a multi-view KG was developed to facilitate real-time learning and instructional support, incorporating information from student learning activities, search patterns, and course-specific content [54].

e. Miscellaneous applications

Beyond defined categories, KGs are applied innovatively across diverse fields, demonstrating their adaptability and transformative potential. For instance, an "entity-event KG" was used to manage education personnel data in public administrations [55]. Another study enhanced piano teaching by integrating curriculum resources and constructing a multimodal knowledge atlas combining deep neural networks [56]. Similarly, a domain KG for primary school mathematical operation literacy was proposed to assist in assessing and improving students' mathematical skills [57].

KG technology has also been instrumental in lifelong learning ecosystems. For example, Yu created an ecological chain of supply for lifelong learning resource bases using KGs, encompassing resource utilisation, visual presentation, intelligent recommendation, question answering, and cross-resource association [58]. Zheng et al. developed an Automatic Activated and Unactivated KG for collaborative knowledge building within Computer-Supported Collaborative Learning (CSCL) contexts [59].

Additionally, the development of KGs for power grid education has been reported, facilitating access to educational resources and enabling intelligent teaching platforms [60]. The automated construction of course ontologies using a combination of automated data acquisition and manual annotation methods ensures comprehensive and rich course representations [61]. Moreover, integrating theoretical and practical knowledge in computer networking courses for secondary vocational schools was discussed, involving data acquisition and transformation techniques [62]. Comprehensive educational KGs have been constructed to support data-driven decision-making, resource aggregation, problem-solving, and adaptive learning systems [63].

Examples include CKGG, a Chinese KG for high school geography curriculum designed to enhance students' computer-aided education by integrating diverse geographic data types [64]. These varied applications demonstrate KGs' versatility and potential to reshape education across various contexts.

f. Literature review conclusion and implications

This literature review underscores KGs' critical role in enhancing personalised learning, curriculum design, concept mapping, and educational content recommendation systems. Recent advancements involve integrating semantic web technologies, natural language processing, and machine learning algorithms to create dynamic, contextually relevant KGs. These KGs are now capable of capturing evolving knowledge within the educational landscape. Despite these advancements, several limitations and challenges persist. These include:

Knowledge extraction techniques: Limited discussion and documentation restrict reproducibility and progress, as many studies lack transparency in their methods for converting unstructured information into structured knowledge.

Lack of standardisation: The absence of universally accepted formats or ontologies hampers sharing, integrating, and comparing KGs across different educational platforms and institutions. The diversity of educational contexts further complicates standardisation efforts.

Limited interoperability: Most KGs operate in isolation, limiting their potential for creating broader educational networks that harness KGs for various applications. The absence of common ontologies also impedes seamless integration of data across institutions.

Sparse and incomplete data: Constructing comprehensive KGs requires extensive data collection and curation, often resulting in sparse and incomplete representations that hinder effectiveness.

Scalability challenges: As educational content grows, maintaining updated and comprehensive KGs becomes increasingly challenging, often resulting in outdated or incomplete knowledge.

Educational semantic heterogeneity: Integrating diverse data sources while maintaining coherence and meaning is a technical challenge, as educational data is semantically heterogeneous.

Real-time updates: Ensuring that KGs are continuously updated to reflect new knowledge and pedagogical trends is a significant technical demand that existing approaches often neglect.
Poor evaluation techniques: Assessing KGs' educational effectiveness is challenging due to reliance on subjective factors like learner engagement and knowledge retention. Standardised evaluation metrics are lacking.
Privacy and security concerns: KGs often include sensitive student data, necessitating stringent measures to ensure data protection and compliance with regulations. This aspect is frequently underexplored.

Addressing these challenges requires collaborative efforts towards transparency, standardisation, interoperability, comprehensive data integration, scalable solutions, context-aware knowledge extraction, real-time updates, rigorous evaluation metrics, and stringent privacy and security measures. By overcoming these limitations, KGs can significantly transform educational landscapes, fostering personalised learning, informed decision-making, and enhanced educational outcomes.

Annex 2 - Data mapping and modelling for educational knowledge graph

In the era of abundant data and information, organising and extracting meaningful insights from complex knowledge sources is a significant challenge. KGs have emerged as a powerful approach to address this challenge by providing a structured and interconnected representation of knowledge. Data modelling is a crucial step in building a KG for educational environments. It involves defining the structure, entities, attributes, and relationships that will be represented in the graph.

Building a KG based on LMS data and pedagogical guidelines requires collaboration among educational stakeholders, domain experts, data scientists and policymakers. Data modelling in an educational KG is an iterative process that may evolve as new data sources, research findings, or requirements emerge. The benefits of building such a KG are manifold. It enables the integration of disparate data sources, including structured, semi-structured, and unstructured data, into a unified framework. By capturing the relationships and connections between entities, a KG reveals hidden patterns, insights, and dependencies that might not be apparent in isolated data silos.

Data mapping, acquisition and integration

Building KGs for educational purposes involves the crucial steps of data acquisition and integration. This process begins with collecting relevant data from diverse sources and follows with transforming and integrating this data into a unified KG. Initially, data is extracted through various methods such as querying databases or employing APIs, and then preprocessed to ensure high quality and consistency. Preprocessing may include cleaning, normalising, and transforming the data to fit the desired schema and ontology of the KG. The next step involves mapping the preprocessed data to the schema, aligning it with the designated entities, attributes, and relationships. Following schema mapping, the data is integrated into the graph by creating or updating its structure to include the new information. Semantic enrichment techniques, such as entity disambiguation and relationship extraction, further enhance the graph's quality and contextual accuracy. Given the dynamic nature of KGs, establishing mechanisms for continuous updates is essential to maintain their relevance and accuracy over time.

Graph data schema and ontology

Graph data modelling and schema design are essential for structuring educational data, representing entities, attributes, and relationships in a graph format, ensuring data consistency and quality. This involves analysing data sources, entities, and relationships pertinent to the domain and graph goals, utilising standards like RDF Schema and Web Ontology Language (OWL) to formalise the data model within a graph structure. Semantic web technologies, such as URIs and linked data principles, enhance interoperability and reusability. The design may incorporate ontologies to standardise the concepts, properties, and relationships within the domain.

In analysing data sources for KGs in education, it is crucial to assess the characteristics, quality, and relevance of data from various origins like databases, APIs, and web sources. Understanding the structure and semantics of these data sources is essential to align them with the domain-specific ontology or schema of the graph. Techniques for data profiling help evaluate completeness, accuracy, and consistency, addressing any quality issues like missing values or anomalies. It is vital to map the data sources to the KG's scheme, identifying commonalities and resolving conflicts to ensure accurate integration. Analysing the semantic content and aligning it with external knowledge bases aids in enriching data with semantic information. Ultimately, a tailored data integration strategy is needed to effectively merge and link data from different sources within the KG.

Conceptualising entities, attributes, and relationships in online education

Conceptualising entities, attributes, and relationships in online education is fundamental to designing the ontology or data model for the augMENTOR KGs. This process involves identifying and defining key elements, their properties, and their connections to reflect real-world educational data accurately. The primary entities relevant to online education are first identified, followed by determining the attributes or descriptive characteristics associated with each entity. Relationships between entities, including their cardinality and multiplicity, are then defined to illustrate how they interact and connect. Hierarchies or taxonomies may be created for entities to represent levels of abstraction or specialisation, providing a clearer understanding of their relationships.

Ensuring clarity and consistency involves defining the semantics of all entities, attributes, and relationships using standardised vocabularies or ontologies specific to the online education domain. Collaboration with domain experts, educators, and stakeholders is crucial to ensure the conceptualization accurately reflects the real-world domain. This process is iterative, meaning new entities, attributes, or relationships may be added or refined as the project evolves. The ultimate goal is to create a KG that organises educational data in an accessible and meaningful manner for analysis and querying.

Data transformation and cleaning

Data transformation and cleaning are pivotal when preparing data for integration into KGs. These processes encompass activities such as conversion, standardisation, and the assurance of data quality and consistency. In more detail, the steps towards integrating data in a KG are presented in Table 2.

Table. Data transformation and cleaning steps

Integration steps	Description
Data format conversion	Analysis of the data sources, identification of their formats (e.g., CSV, JSON, XML) and conversion into a standardised format suitable for the KG. This step ensures consistency and facilitates data integration.
Data profiling	Examination of the data to gain insights into its structure, quality, and content. It includes identifying data types, missing values, outliers, and potential data quality issues.
Data standardisation	Examination of the data for inconsistencies, variations, or discrepancies in attribute values, units, or representations and standardisation of the data to adhere to a common format or set of rules. This may involve normalising dates, transforming measurements to a common unit, or resolving inconsistencies in naming conventions.
Data cleaning	Identification and handling of data quality issues such as missing values, outliers, and erroneous entries and apply data cleaning techniques such as removing duplicates, correcting errors, and filling in missing values using appropriate strategies (e.g., mean imputation, data inference, or external data sources).
Schema alignment	Analysis of the data attributes and relationships and alignment with the schema or ontology of the KG, while ensuring that attribute names, types, and semantics are consistent with the intended graph structure. Afterwards,

	the data attributes are mapped to the corresponding entities and attributes defined in the KG schema.
Entity disambiguation	Address of entity disambiguation challenges by resolving ambiguities and identifying unique entities. This may involve applying entity resolution techniques to match and merge similar entities based on their attributes or leveraging external knowledge bases or ontologies for disambiguation.
Semantic enrichment	Enhancement of the data with additional semantic information to enrich the KG. This can involve annotating the data with additional attributes or linking the entities to external ontologies, controlled vocabularies, or knowledge bases. Semantic enrichment improves the contextual understanding and interoperability of the data in the KG.
Data integration	Combination of the data from multiple sources into a unified format. This may involve resolving data conflicts, mapping different schemas, or aligning timestamps.
Data validation and quality assurance	Performance of validation checks to ensure the integrity and quality of the transformed data and validate the data against predefined rules, constraints, or validation criteria. Data quality issues or inconsistencies identified during the validation process must be addressed.
Data versioning	Maintaining a record of data versions and changes made during the transformation and cleaning process. This helps track data lineage and facilitates reproducibility.
Automation	Automation of data transformation and cleaning processes where possible to ensure consistency and efficiency, especially when dealing with large volumes of data.
Iterative process	Data transformation and cleaning are often iterative processes. As new data becomes available or as requirements change, revisit and update these steps accordingly.

Entity recognition, relationship extraction and linking

Entity recognition and extraction are essential for building KGs, involving the identification and extraction of entities from unstructured or semi-structured data. The first step is defining the entity types (e.g., people, organisations, locations, or domain-specific entities like products or events). Utilising Natural Language Processing (NLP) techniques such as Named Entity Recognition (NER), Part-Of-Speech (POS) tagging, and chunking helps in identifying and extracting these entities and their attributes. Pretrained models like spaCy or Stanford NER, or custom-trained models, can be used for this purpose.

Once entities are identified, they are linked to corresponding entities in the KG, enriching it with information from various data sources. Continuous iteration and refinement based on feedback and performance evaluation are crucial to enhance accuracy and coverage.

Relationship extraction involves identifying and capturing connections between entities. Techniques like dependency parsing, pattern matching, and information extraction are used to analyse data and reveal meaningful relationships. Dependency parsing examines sentence structures to find syntactic connections, while pattern matching captures specific linguistic patterns indicating relationships.

Relationship linking matches these extracted relationships to those in the KG, enabling applications like question answering and text summarization. This involves creating links between entities in the text and their counterparts in the graph, leveraging disambiguation techniques and external knowledge bases. Proper labelling and creating edges to represent relationships are essential steps to enhance the graph's semantic richness and interconnectivity.

In summary, entity recognition and relationship extraction are crucial for integrating and connecting diverse information within a KG, enabling advanced analysis, reasoning, and discovery.

Annex 3 - Operations of the KGs

The operations of educational KGs are fundamental to leveraging and optimising the value of extensive educational data. This section delves into two critical components: the storage and management of educational KGs, and the processes involved in constructing and populating these graphs with relevant data. By focusing on effective data integration, organisation, and analysis, educational KGs serve as structured repositories that facilitate coherent access to interconnected information. The systematic acquisition, preprocessing, and integration of data from diverse sources are essential for building a comprehensive knowledge framework that supports educators, policy makers, and researchers in their efforts to make informed decisions and drive educational initiatives. The following sections will provide insights into the methodologies and technologies that underpin these operations, paving the way for more effective data management in the educational landscape.

Storing and managing educational KGs

Storing and managing educational KGs is vital for leveraging educational data effectively. These graphs organise, integrate, and analyse vast and intricate educational data, providing a structured repository that captures complex relationships between entities, attributes, and phenomena.

KGs integrate diverse data sources, ensuring consistency and interoperability across different datasets by mapping data to a common schema. This unified structure allows educators, policy makers, and stakeholders to access and analyse integrated data from multiple sources coherently.

A well-defined data model is essential, defining entities, their properties, and relationships using standards like RDF or custom schemas. Graph databases, such as Neo4j, Amazon Neptune, and JanusGraph, efficiently store and manage these structures. RDF triple stores like Apache Jena and Stardog manage semantic data using triples (subject-predicate-object).

For large-scale KGs, distributed graph processing frameworks like Apache Spark GraphX and Apache Giraph enable parallel processing and computations. Ontology languages like RDF Schema (RDFS) and Web Ontology Language (OWL) define the structure, semantics, and reasoning capabilities, providing a formal way to define domain-specific concepts.

Specialised platforms like AllegroGraph and Stardog combine storage, querying, and visualisation with features like entity resolution, reasoning, and integration with external data sources. ETL (Extract, Transform, Load) processes, utilising tools like Apache Kafka and Apache NiFi, are commonly used to integrate data from various sources into the KG.

Constructing and populating educational KGs

The process of constructing and populating educational KGs involves a systematic approach to integrating diverse educational data sources into a cohesive framework. This critical phase transforms raw data into a structured format, allowing for better organisation and analysis. The following sections detail the processes of data acquisition, preprocessing, and integration into the KG, providing essential insights for establishing and maintaining an educational knowledge repository.

Data acquisition begins with the collection of relevant information from various sources, such as educational institutions or online platforms. The subsequent step involves extracting pertinent datasets and variables necessary for building the KG. Once collected, the data undergoes cleaning to eliminate inconsistencies, missing values, duplicates, and errors. Techniques such as filling in gaps and correcting inaccuracies ensure that the data quality remains high. Data transformation follows, standardising formats and units to ensure compatibility across diverse datasets, enhancing their utility when integrated into the KG.

Enrichment of the data through linkage with external ontologies or knowledge bases adds semantic context, making the data more valuable for analysis. Additionally, semantic annotation provides essential metadata that describes the data's meaning and relationships, facilitating alignment with the graph's ontology. Entity recognition techniques are employed to accurately identify and disambiguate entities within the data, while relationship extraction helps to define the connections between them. Once cleaned, transformed, and annotated, the data can be mapped to the graph schema, establishing a clear structure for integration. Maintaining data versioning throughout this process allows for tracking changes and ensuring reproducibility.

Preparing the data for analysis presents its own challenges, as educational data is often not readily available in formats suitable for immediate use. Persistent identifiers can lead to landing pages requiring manual navigation, which complicates automated systems' efforts to access and utilise the data effectively. To address these challenges, the growth of research infrastructures has facilitated ongoing curation and publication of educational data. These infrastructures support both research initiatives and policymaking, ensuring continuous access to relevant information.

Extracting information from data sources is a multifaceted endeavour. It begins with selecting the most relevant datasets, informed by domain expertise to align with the project's objectives. After careful selection, the data must be parsed and filtered to isolate specific content of interest. Quality assessments identify and rectify any anomalies or errors, ensuring reliable information is integrated into the KG. The integration process involves structuring the data in accordance with the graph's ontology and semantics, adding semantic annotations to enhance context and meaning.

Transforming data into a suitable format for the KG involves defining relationships among various attributes, variables, and entities. This structuring phase typically utilises standardised formats like JSON or RDF, which clearly delineate entities and their connections. The process also includes normalising data to ensure uniformity, linking it to external ontologies for added context and interoperability. Throughout this transformation, quality checks verify the data's integrity, ensuring accuracy and compliance with predefined standards.

Loading and updating the KG is vital for its ongoing relevance and utility. The integration of transformed data into the graph structure ensures that all elements fit correctly within the established schema. As the graph is populated with new data, it expands and evolves, enhancing the breadth of information available. Continuous updates can occur through automated processes or manual contributions from domain experts, ensuring that the KG remains current. To facilitate reproducibility, version control is established, documenting changes and maintaining the lineage of data.

In conclusion, the construction and population of educational KGs is a systematic process that emphasises data acquisition, preparation, transformation, and integration. Each phase plays a crucial role in establishing a dynamic resource, allowing for comprehensive analysis and informed decision-making within educational contexts. The ongoing maintenance and updating of the KG are essential to its success, ensuring the integrity, accuracy, and usability of the educational data it encompasses.

Querying and analysing information from educational KGs

In the field of online education, extracting valuable insights from extensive data repositories is crucial. Educational KGs facilitate this through various querying and analysis techniques. These techniques enable users to construct precise queries to retrieve relevant information efficiently.

Graph querying methods, such as keyword search, Natural Language Queries (NLQs), semantic search, and pattern matching, offer different ways to explore the KG. Keyword search is straightforward and user-friendly, ideal for quick data retrieval, while NLQs leverage natural language processing to translate user queries into structured database queries, making the data accessible to non-technical users. Semantic search improves accuracy by understanding the context and relationships within user queries, offering more relevant results. Pattern matching allows for the identification of specific subgraphs or patterns, using defined queries to extract structured information and compare graph structures.

Graph algorithms and analytics are pivotal for uncovering insights and patterns within educational data. By applying these algorithms, users can explore data relationships, enhance understanding, and support data-driven decision-making. For instance, in the augMENTOR project, graph algorithms are used to derive meaningful analytics, improving data analysis and pattern recognition, valuable for educators and policymakers.

Semantic reasoning and inferencing are essential for deriving new knowledge from existing data within the graph. These processes use formal logic, rules, and ontologies to make automated deductions, generating new insights that might not be explicitly represented in the original data. Techniques like deductive, inductive, and abductive inference enhance the system's reasoning capabilities and support complex problem-solving.

Visualisation and exploration tools encompass a wide range of software and applications used to visually represent and interact with data in a meaningful and informative way. These tools are essential for gaining insights, identifying patterns, and communicating findings from complex datasets. Querying and visualisation tools allow for exploring and analysing the KG. SPARQL is a widely used query language for RDF-based KGs, while Cypher is commonly used for querying graph databases. Visualisation tools like Cytoscape[1], Gephi[2], Bloom[3] and Graphistry[4] allow the creation of interactive network visualisations,

[1] https://cytoscape.org/
[2] https://gephi.org/
[3] https://neo4j.com/product/bloom/
[4] https://www.graphistry.com/

analysis of node and edge properties, and discovery of patterns within the graph. Libraries like D3.js[5], Matplotlib[6], and Plotly[7] provide developers with the tools to create custom data visualisations and interactive charts. These libraries are often used to visualise trends, correlations, and patterns within datasets. For exploring linked data and ontologies, OntoWiki[8] and Marbles[9] are useful tools that assist in navigating complex semantic structures and understanding relationships between entities. For dashboards, platforms like Tableau[10], PowerBI[11] and Grafana[12] allow users to combine data from diverse sources to provide real and near-real time analytics.

[5] https://d3js.org/
[6] https://matplotlib.org/
[7] https://plotly.com/
[8] https://aksw.org/Projects/OntoWiki.html
[9] https://mes.github.io/marbles/
[10] https://www.tableau.com/
[11] https://powerbi.microsoft.com/en-us/
[12] https://grafana.com/

References

[1] Malinverni, E.S., Naticchia, B., Lerma Garcia, J.L. et al. A semantic graph database for the interoperability of 3D GIS data. Appl Geomat 14 (Suppl 1), 53–66 (2022), https://doi.org/10.1007/s12518-020-00334-3.

[2] Dilmegani C. (2023). In-depth Guide to Knowledge Graph: Benefits, Use Cases & examples.

[3] Wang, Z., Zhang, J., Feng, J., & Chen, Z. (2014). Knowledge Graph Embedding by Translating on Hyperplanes. Proceedings of the AAAI Conference on Artificial Intelligence, 28(1). https://doi.org/10.1609/aaai.v28i1.8870.

[4] Horst H. and Cimiano P. (2020). Incorporating Semantic Dependencies Extracted from Knowledge Graphs into Joint Inference Template-based Information Extraction. In Proc. of the 24th European Conference on Artificial Intelligence - ECAI 2020.

[5] Hailun Lin, Yong Liu, Weiping Wang, Yinliang Yue, Zheng Lin, Learning Entity and Relation Embeddings for Knowledge Resolution, Procedia Computer Science, Volume 108, 2017, Pages 345-354, ISSN 1877-0509, https://doi.org/10.1016/j.procs.2017.05.045.

[6] Phillip Schneider, Tim Schopf, Juraj Vladika, Mikhail Galkin, Elena Simperl, and Florian Matthes. 2022. A Decade of Knowledge Graphs in Natural Language Processing: A Survey. In Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 601–614, Online only. Association for Computational Linguistics.

[7] Amit Kumar Gupta and Rajesh Prasad. (2023). Graph Learning and Network Science for Natural Language Processing, Edited By Muskan Garg, Copyright 2023, 1st Edition.

[8] Lingfei Wu, Yu Chen, Heng Ji, and Bang Liu. 2021. Deep Learning on Graphs for Natural Language Processing. In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD '21). Association for Computing Machinery, New York, NY, USA, 4084–4085. https://doi.org/10.1145/3447548.3470820.

[9] Rajendra Akerkar and Springer Cham. (2020). Big Data in Emergency Management: Exploitation Techniques for Social and Mobile Data, Springer Nature Switzerland AG 2020, ISBN 978-3-030-48098-1, 15 September 2020.

[10] Collarana, Diego & Galkin, Mikhail & Traverso, Ignacio & Lange, Christoph & Vidal, Maria-Esther & Auer, Sören. (2017). Semantic Data Integration for Knowledge Graph Construction at Query Time. 109-116. 10.1109/ICSC.2017.85.

[11] Peng, C., Xia, F., Naseriparsa, M. et al. Knowledge Graphs: Opportunities and Challenges. Artif Intell Rev 56, 13071–13102 (2023). https://doi.org/10.1007/s10462-023-10465-9.

[12] Q. Wei, X. Yao, Personalised recommendation of learning resources based on knowledge graph, in: 2022 11th International Conference on Educational and Information Technology (ICEIT), IEEE, 2022.

[13] Lu, Y., Pian, Y., Chen, P., Meng, Q. and Cao, Y. 2021. RadarMath: An Intelligent Tutoring System for Math Education. Proceedings of the AAAI Conference on Artificial Intelligence. 35, 18 (May 2021), 16087-16090. https://doi.org/10.1609/aaai.v35i18.18020.

[14] Z. Bai, Variable incremental adaptive learning model based on knowledge graph and its application in online learning system, Int. J. Comput. Appl. 44 (7) (2022) 650–658.

[15] Albreiki, Balqis & Habuza, Tetiana & Palakkal, Nishi & Zaki, Nazar. (2023). Clustering-based knowledge graphs and entity-relation representation improves the detection of at risk students. Education and Information Technologies. 29. 1-30. 10.1007/s10639-023-11938-8.

[16] X. Xia, W. Qi, Interpretable early warning recommendations in interactive learning environments: a deep-neural network approach based on learning behaviour knowledge graph, Humanities and Social Sciences Communications 10 (1) (2023) 1–13.

[17] Y. Xue, Knowledge graph based recommendation by adversarial learning algorithm in application of lifelong education data classification, ACM Transactions on Asian and Low-Resource Language Information Processing (2023).

[18] L. Min, Research on personalized recommendation technology for learning resources combining knowledge graph and interest diffusion mechanism, in: 2022 3rd International Conference on Information Science and Education (ICISE-IE), IEEE, 2022.

[19] M. Mu, M. Yuan, Research on a personalized learning path recommendation system based on cognitive graph with a cognitive graph, Interact. Learn. Environ. (2023) 1–19.

[20]] S. Bhaskaran, N. Bharathiraja, K. Pradeepa, M. Vinoth Kumar, N. V. Ravindhar and R. Marappan, "New Recommender System for Online Courses Using Knowledge Graph Modeling," 2023 International Conference on Computer Communication and Informatics (ICCCI), Coimbatore, India, 2023, pp. 1-6, 10.1109/ICCCI56745.2023.10128262.

[21] Z. Ni, F. Ni, Research on knowledge graph model of diversified online resources and personalized recommendation, in: Journal of Physics: Conference Series, IOP Publishing, 2020.

Funded by
the European Union

[22] H. Wu, F. Meng, Research on the application of personalized course recommendation of learn to rank based on knowledge graph, in: Mobile Wireless, Middleware, Operating Systems and Applications: 9th EAI International Conference, MOBILWARE 2020, Hohhot, China, July 11, 2020, Proceedings, vol. 9, Springer, 2020.

[23] Feng Wang, Lingling Zhang, Xingchen Chen, Ziming Wang, Xin Xu, Research on personalized learning path discovery based on differential evolution algorithm and knowledge graph, in: Data Science: 6th International Conference, ICDS 2019, Ningbo, China, May 15–20, 2019, Revised Selected Papers 6, Springer, 2020.

[24] Weng, J., Gao, Y., Qiu, J., Ding, G., Zheng, H. (2019). Construction and Application of Teaching System Based on Crowdsourcing Knowledge Graph. In: Zhu, X., Qin, B., Zhu, X., Liu, M., Qian, L. (eds) Knowledge Graph and Semantic Computing: Knowledge Computing and Language Understanding. CCKS 2019. Communications in Computer and Information Science, vol 1134. Springer, Singapore. https://doi.org/10.1007/978-981-15-1956-7_3.

[25] Yuan Sun, Jiaya Liang, and Pengchao Niu. 2021. Personalized Recommendation of English Learning Based on Knowledge Graph and Graph Convolutional Network. In Artificial Intelligence and Security: 7th International Conference, ICAIS 2021, Dublin, Ireland, July 19–23, 2021, Proceedings, Part II. Springer-Verlag, Berlin, Heidelberg, 157–166. https://doi.org/10.1007/978-3-030-78612-0_13.

[26] Zayet TMA, Ismail MA, Almadi SHS, Zawia JMH, Mohamad Nor A. What is needed to build a personalized recommender system for K-12 students' E-Learning? Recommendations for future systems and a conceptual framework. Educ Inf Technol (Dordr). 2023;28(6):7487-7508, 10.1007/s10639-022-11489-4.

[27] Y. Yang, Y. Zhu, P. Jian, Application of knowledge graph in water conservancy education resource organisation under the background of big data, Electronics 11 (23) (2022) 3913.

[28]] N. Hubert, A. Brun, D. Monticolo, New ontology and knowledge graph for university curriculum recommendation, in: ISWC 2022-The 21st International Semantic Web Conference, 2022.

[29] F. Zablith, Constructing social media links to formal learning: a knowledge Graph Approach, Educ. Technol. Res. Dev. 70 (2) (2022) 559–584.

[30] L. Feng, Z. Ke, N. Wu ModelsKG, A design and research on knowledge graph of multimodal curriculum based on PaddleOCR and DeepKE, in: 2022 14th, International Conference on Advanced Computational Intelligence (ICACI), IEEE, 2022.

[31] S. Shanshan, L. Wenfei, L. Lijuan, Applying lean six sigma incorporated with big data analysis to curriculum system improvement in higher education institutions, International Journal of System Assurance Engineering and Management (2022) 1–16.

[32]] X. Li, M. Chen, Management course knowledge graph construction based on ontology, in: 2022 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT), IEEE, 2022.

[33] X. Yao, Research on the construction and application of multi-modal curriculum knowledge graph for blended teaching, in: Proceedings of the 2021 5th International Conference on E-Education, E-Business and E-Technology, 2021.

[34] J. Gao, L. Wang, F. Xu, Research on the construction of course knowledge graph of high school information technology, in: 2020 International Conference on Artificial Intelligence and Education (ICAIE), IEEE, 2020.

[35] J. Bai, L. Che, Construction and application of database micro-course knowledge graph based on Neo4j, in: The 2nd International Conference on Computing and Data Science, 2021.

[36] L. Hu, Y. Jiang, Y. Li, Optimization design of internet fraud case based on knowledge graph and case teaching, in: Proceedings of the 2019 7th International Conference on Information and Education Technology, 2019.

[37] G. Ding, X. Shi, S. Li, Integrating Programming Errors into Knowledge Graphs for Automated Assignment of Programming Tasks, Education and Information Technologies, 2023, pp. 1–34.

[38] Y. Li, J. Qiu, S. Gui, Y. Song and T. Liu, "Analytics 2.0 for Precision Education Driven by Knowledge Map," 2022 IEEE Frontiers in Education Conference (FIE), Uppsala, Sweden, 2022, pp. 1-5, doi: 10.1109/FIE56618.2022.9962461.

[39] Zheng, L., Long, M., Chen, B. et al. Promoting knowledge elaboration, socially shared regulation, and group performance in collaborative learning: an automated assessment and feedback approach based on knowledge graphs. Int J Educ Technol High Educ 20, 46 (2023). https://doi.org/10.1186/s41239-023-00415-4.

[40] Z. Li, L. Cheng, C. Zhang, X. Zhu and H. Zhao, "Multi-Source Education Knowledge Graph Construction and Fusion for College Curricula," 2023 IEEE International

Conference on Advanced Learning Technologies (ICALT), Orem, UT, USA, 2023, pp. 359-363, doi: 10.1109/ICALT58122.2023.00111.

[41] Ying Su and Yong Zhang. 2020. Automatic Construction of Subject Knowledge Graph based on Educational Big Data. In Proceedings of the 2020 3rd International Conference on Big Data and Education (ICBDE '20). Association for Computing Machinery, New York, NY, USA, 30–36. https://doi.org/10.1145/3396452.3396458.

[42]]Gong, Z., Yu, X., Fu, W., Che, X., Mao, Q., & Zheng, X. (2021). The construction of knowledge graph for personalized online teaching. In Data Mining and Big Data: 6th International Conference, DMBD 2021, Guangzhou, China, October 20–22, 2021, Proceedings, Part II 6 (pp. 98-107). Springer Singapore.

[43] Yu, H., Chen, Y., Wang, X., Hu, S., & Chang, L. (2021, May). A capability-oriented method for integration of liberal arts knowledge graph in primary education. In 2021 International Conference on Computers, Information Processing and Advanced Education (CIPAE) (pp. 342-346). IEEE.

[44] Bowen Zhao, Jiuding Sun, Bin Xu, Xingyu Lu, Yuchen Li, Jifan Yu, Minghui Liu, Tingjian Zhang, Qiuyang Chen, Hanming Li, Lei Hou, Juanzi Li. (2022). EDUKG: a heterogeneous sustainable k-12 educational knowledge graph. arXiv preprint arXiv:2210.12228.

[45] Tang, Y., Yuan, E., Chen, W., Zhang, S., Liu, L., & Wu, Y. (2023, June). Analysis of learning effectiveness based on knowledge graph. In 2023 4th International Conference on Education, Knowledge and Information Management (ICEKIM 2023) (pp. 1742-1749). Atlantis Press.

[46] Agrawal, G., Deng, Y., Park, J., Liu, H., & Chen, Y. C. (2022). Building Knowledge Graphs from Unstructured Texts: Applications and Impact Analyses in Cybersecurity Education. Information 2022, 13, 526.

[47]] Fang, Y., Deng, J., Zhang, F., & Wang, H. (2023). An intelligent question-answering model over educational knowledge graph for sustainable urban living. Sustainability, 15(2), 1139.

[48]] Nguyen, H. D., Truong, D., Vu, S., Nguyen, D., Nguyen, H., & Tran, N. T. (2023). Knowledge management for information querying system in education via the combination of rela-ops model and knowledge graph. Journal of Cases on Information Technology (JCIT), 25(1), 1-17.

[49] Feng, S., Tu, Z., Huang, M., & Wu, Y. (2021, August). Knowledge graph construction and intelligent question answering on science and technology intermediary service. In 2021 4th International Conference on Pattern Recognition and Artificial Intelligence (PRAI) (pp. 192-196). IEEE.

[50] Yang, Z., Wang, Y., Gan, J., Li, H., & Lei, N. (2021). Design and research of intelligent question-answering (Q&A) system based on high school course knowledge graph. Mobile Networks and Applications, 1-7.

[51] L.S. Nair, M. Shivani, Knowledge graph based question answering system for remote school education, in: 2022 International Conference on Connected Systems & Intelligence (CSI), IEEE, 2022.

[52] F. Dang, J. Tang, S. Li, MOOC-KG: a MOOC knowledge graph for cross-platform online learning resources, in: 2019 IEEE 9th International Conference on Electronics Information and Emergency Communication (ICEIEC), IEEE, 2019.

[53] R. Manrique, B. Pereira, O. Marino, ˜ Exploring knowledge graphs for the identification of concept prerequisites, Smart Learning Environments 6 (1) (2019) 1–18.

[54] Bhattacharjee, S. D., Gokaraju, J. S. A. V., Yuan, J., & Kalwa, A. (2022, August). Multi-view knowledge graph for explainable course content recommendation in course discussion posts. In 2022 26th International Conference on Pattern Recognition (ICPR) (pp. 2785-2791). IEEE.

[55]] Bourmpoulias, S., Zeginis, D., & Tarabanis, K. (2023, June). An entity event knowledge graph for human resources management in public administration: the case of education personnel. In 2023 IEEE 25th Conference on Business Informatics (CBI) (pp. 1-8). IEEE.

[56] Li, H. (2022). Piano Teaching Knowledge Graph Construction Based on Cross-Media Data Analysis and Semantic Network. Computational Intelligence and Neuroscience, 2022(1), 5499593.

[57] Geng, J., Zhu, J., Yang, H., Hu, S., & Huang, T. (2022, May). Research on diagnosis and improvement of mathematical operation literacy based on the domain knowledge graph. In Proceedings of the 7th International Conference on Distance Education and Learning (pp. 67-73).

[58] Yu, Y. (2022). Research on the application of knowledge graph in constructing ecological chain of supply of lifelong learning resource base. Open Access Library Journal, 9(9), 1-11.

[59] Zheng, L., Niu, J., Long, M., & Fan, Y. (2023). An automatic knowledge graph construction approach to promoting collaborative knowledge building, group performance, social interaction and socially shared regulation in CSCL. British Journal of Educational Technology, 54(3), 686-711.

[60] Ma, Y., Hong, D., Dan, F., Yang, X., & Li, X. (2021, June). Research on the construction method of knowledge graph for power grid education resources. In 2021 IEEE 3rd International Conference on Computer Science and Educational Informatization (CSEI) (pp. 99-103). IEEE.

[61] Y. Chen, Z. Chen, Research on the key technology of course ontology construction, in: Journal of Physics: Conference Series, IOP Publishing, 2022.

[62] Li, G., Wang, H., & Liu, H. (2022, November). Knowledge graph construction for computer networking course group in secondary vocational school based on multi-source heterogeneous data. In 2022 12th International Conference on Information Technology in Medicine and Education (ITME) (pp. 99-103). IEEE.

[63] Wu, L., Tu, F., Liu, Q., Ma, X., Jia, X., & Gao, Y. (2022, July). A review of intelligent application of educational knowledge graph from the perspective of artificial intelligence. In 2022 International Symposium on Educational Technology (ISET) (pp. 195-199). IEEE.

[64] Shen, Y., Chen, Z., Cheng, G., & Qu, Y. (2021). CKGG: A Chinese knowledge graph for high-school geography education and beyond. In The Semantic Web–ISWC 2021: 20th International Semantic Web Conference, ISWC 2021, Virtual Event, October 24–28, 2021, Proceedings 20 (pp. 429-445). Springer International Publishing.

Funded by
the European Union