UNIVERSITÀ
DEGLI STUDI
DI CAGLIARI

Ph.D. DEGREE IN
Mathematics and Computer Science

Cycle XXXVIII

TITLE OF THE Ph.D. THESIS

Knowledge-aware Methods for Explainable Decision Support in Lifelong Learning

Scientific Disciplinary Sector(s)

INFO-01/A

Ph.D. Student:	Neda Afreen
Supervisor:	Prof. Gianni Fenu
Co-Supervisors:	Prof. Ludovico Boratto, Dr. Mirko Marras

Final exam. Academic Year 2024/2025
Thesis defence session: May 2026

Statement of Authorship

I declare that this thesis entitled "Knowledge-aware Methods for Explainable Decision Support in Lifelong Learning" and the work presented in it are my own. I confirm that:

- this work was done while in candidature for this PhD degree;
- when I consulted the work published by others, this is always clearly attributed;
- when I quoted the work of others, the source is always given;
- I have acknowledged all main sources of help;
- with the exception of the above references, this thesis is entirely my own work;
- appropriate ethics guidelines were followed to conduct this research;
- for work done jointly with others, my contribution is clearly specified.

Abstract

Scaling up digital education presents several critical challenges, including the management of large learner populations, the abundance of learning resources, and the difficulty of supporting informed learning decisions at scale. While digital platforms provide unprecedented access to education, learners often face overwhelming choices and limited guidance while navigation. The increasing availability of learning data and technological advancement creates significant opportunities for artificial intelligence (AI) to support lifelong learning. To make such intelligent support effective in real-world educational settings, careful planning, representation, reasoning are essential.

This thesis addresses the design, implementation, and evaluation of explainable artificial intelligence methods for lifelong learning environments. The focus is on how AI systems can represent educational knowledge, reason over learning-related data, and communicate recommendations and explanations in a transparent and user-centered manner. The proposed contributions leverage learner-centered ontologies, educational knowledge graphs, path-based reasoning methods, conversational interfaces, and large language models to support explainable educational recommendation and knowledge exploration. Rather than prioritizing predictive accuracy alone, the thesis emphasizes explainability and learner trust as key requirements for AI-driven educational systems.

Through a series of studies, this work shows how structured knowledge representations can enable explainability by construction, how conversational interfaces can facilitate intuitive interaction with educational knowledge graphs, and how path-based reasoning can support transparent recommendations. Furthermore, the thesis explores how large language models can transform structured, path-based explanations into user-friendly natural language narratives while preserving faithfulness to the underlying reasoning process. The findings provide methodological insights and design guidelines for combining structured reasoning and language-based explanation in educational AI systems. Overall, this thesis contributes to advancing explainable AI for lifelong learning by encouraging transparency, improving user acceptance, and strengthening the role of structured educational data as an asset for scalable and trustworthy learning support.

Biography

Neda Afreen, was born on December 29, 1996 in Chiknauta (India). She is pursuing her PhD in Artificial Intelligence in Lifelong Learning at the *University of Cagliari* (Italy), under a scholarship funded by the Ministerial Decree as part of the National Recovery and Resilience Plan (NRRP), advised by *Prof. Gianni Fenu* and co-supervised by *Prof. Ludovico Boratto* and *Dr. Mirko Marras*. She completed her Master of Technology degree in Computer Engineering from *Jamia Millia Islamia* (India), where she graduated with First Division with Honors, and a Bachelor's degree in Information Technology from the *Women's Institute of Technology*, an institution with a strong academic legacy in India.

Before starting her doctoral journey, she collaborated with Indian and international researchers on projects in the healthcare domain and contributed technical writing for top *IEEE* conferences and journals. As part of her PhD, she has undertaken two research visits, one at *KU Leuven* (Belgium), where she contributed to the Flemish government-supported i-Learn project on Human-Centered Explainable AI for personalized digital learning, and another at the *Autonomous Region of Sardinia* (Italy). She has co-authored research in international venues, including *RecSys*, *UMAP*, *CIKM*, and *IIR*, addressing topics on path-based explainable recommendation, ontology development for education, and conversational interfaces for exploring educational knowledge graphs. Her research contributions reflect a strong commitment to bridging the gap between AI methods and their real-world educational applications. She also contributed to the academic community as a volunteer at *RecSys* 2024 and as part of the organizing team for *UMAP* 2024.

Beyond academia, she has been involved in science outreach and community initiatives since her undergraduate days. She volunteered at the *Agastya Mobile Science Lab* and the *Vikshit Bharat Foundation*, conducting hands-on science demonstrations and supporting community development activities. She also coordinated environmental awareness events under the *Aagosh Welfare Trust* in India.

Dissemination

The research that contributed to the skills mastered for and the content part of this Ph.D. thesis has resulted from papers fully published in national and international journals and conference proceedings. I would sincerely thank my co-authors for their precious contribution, and such a gratitude would be demonstrated with the adoption of the scientific ‘We’ throughout the thesis.

I firstly make it clear my contribution. I envisioned the research presented in this thesis and completed the majority of the work. I designed the approaches and chose the research directions. I collected the datasets and was responsible for data analysis. Moreover, I was in charge of the implementation of the related scripts. Finally, I wrote the papers for submission, managed the peer-review pipeline, and subsequently revised them. I collaborated closely with the listed co-authors throughout all stages. Co-authors provided feedback on the approaches, offered technical support, discussed techniques, and contributed to the preparation of the submitted work. Furthermore, I was in charge of the presentation of *3* papers at conferences and workshops. Detailed references to the produced papers are provided below, with presented papers indicated by an asterisk (*).

Peer-reviewed Publications in Journals

i. **Afreen, N.**, Boratto, L., Fenu, G., Marras, M. (2025). *Reimagining Personalization through Explainable Decision Support for Learners.* In: Journal of Inclusive Methodology and Technology in Learning and Teaching, V.5 N.3, ISSN 2785-5104. https://www.inclusiveteaching.it/index.php/inclusiveteaching/article/view/457

ii. **Afreen, N.**, Boratto, L., Fenu, G., Malloci, F. M., Marras, M., Soccol, A. (2025). *Explainable Course Recommendation with Knowledge Graphs: A Comparative Audit of Diverse Modeling Paradigms.* In: Data Mining and Knowledge Discovery, SI Knowledge Discovery from Graphs (under review).

Peer-reviewed Publications in International Conference Proceedings

iii. **Afreen, N.**, Balloccu, G., Boratto, L., Fenu, G., Malloci, F.M., Marras, M., Martis, A.G. (2024). *Learner-centered Ontology for Explainable Educational Recommendation.*. In: In Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization (UMAP Adjunct ’24). Association for Computing Machinery, NewYork, NY, USA, 567–575. https://doi.org/10.1145/3631700.3665226

iv. **Afreen, N.** (2024). *Explainable and Faithful Educational Recommendations through Causal Language Modelling via Knowledge Graphs.* In: Proceedings of the 18th ACM Conference on Recommender Systems (RecSys '24). Association for Computing Machinery, New York, NY, USA, 1358–1360. https://doi.org/10.1145/3640457.3688022 *

v. **Afreen, N.**, Balloccu, G., Boratto, L., Fenu, G., Malloci, F.M., Marras, M., Martis, A.G. (2024). *EDGE: A Conversational Interface driven by Large Language Models for Educational Knowledge Graphs Exploration.* In: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24). Association for Computing Machinery, New York, NY, USA, 5159-5163. https://dl.acm.org/doi/abs/10.1145/3627673.3679231

vi. **Afreen, N.**, Balloccu, G., Boratto, L., Fenu, G., Malloci, F.M., Marras, M., Martis, A.G. (2025). *Can Path-Based Explainable Recommendation Methods based on Knowledge Graphs Generalize for Personalized Education?* In: Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization (UMAP '25). Association for Computing Machinery, New York, NY, USA, 273-278. https://doi.org/10.1145/3699682.3728323

Peer-reviewed Publications in National Conference Proceedings

vii. **Afreen, N.**, Balloccu, G., Boratto, L., Fenu, G., Marras, M. (2023). *Towards Explainable Educational Recommendation through Path Reasoning Methods..* In: IIR 2023. Italian Information RetrievalWorkshop 2023 Proceedings of the 13th Italian Information RetrievalWorkshop (IIR 2023). Pisa, Italy, June 8-9, 2023. https://ceurws.org/Vol-3448/paper-29.pdf *

viii. **Afreen, N.**, Boratto, L., Fenu, G., Marras, M., Soccol A. (2025). *Effective and Transparent Course Recommendation through Causal Reasoning with Language Models..* In: IIR 2025. Italian Information RetrievalWorkshop 2025 Proceedings of the 15th Italian Information RetrievalWorkshop (IIR 2025). Cagliari, Italy, September 3-5, 2025. https://ceur-ws.org/Vol-4026/paper26.pdf *

Acknowledgements

This thesis has been enriched by the professional and emotional support of many people. I have been fortunate to have incredible supervisors who supported me at every step of this journey.

First of all, I would like to express my gratitude to my advisor, Prof. *Gianni Fenu*, for believing in me and for his continuous support throughout this project. My co-supervisor and mentor, Dr. *Mirko Marras*, who guided me from the early stages of enrolment to completion, thank you for your mentorship, technical expertise, and the art of organization and writing. To my other co-supervisor, Prof. *Ludovico Boratto*, thank you for your invaluable guidance, insightful feedback, and encouragement that have been fundamental to the progress of this work. Both of you have been incredible throughout this journey, and I am deeply grateful for all the work and experiences we have shared. Your dedication, patience, and inspiration have been instrumental in my growth.

Furthermore, I gratefully acknowledge that this thesis was produced while attending the PhD programme in Computer Science at the University of Cagliari, Cycle XXXVIII, with the support of a scholarship financed by the Ministerial Decree no. 351 of 9th April 2022, based on the NRRP - funded by the European Union - NextGenerationEU - Mission 4 "Education and Research", Component 1 "Enhancement of the offer of educational services: from nurseries to universities" - Investment 4.1 "Extension of the number of research doctorates and innovative doctorates for public administration and cultural heritage". I am also thankful to the *SIREM* and *RecSys* communities for providing me with travel grants to conduct research abroad and attend international conferences. These opportunities enriched me both professionally and personally.

A heartfelt thanks to my dear colleague and sister, *Sondos Mohamed*, for her constant support. I am also grateful to my wonderful labmates- *Giacomo Medda*, *Giacomo Balloccu*, *Roberta Galici*, *Andrea Atzori*, *Andrea Martis*, *Alessandro Soccol*, who were always there to guide me whenever I needed help. I further extend my gratitude to the administrative staff at the University of Cagliari for their guidance, support, and assistance.

Thank you to my parents, *Shagufta Perween* and *Ghayas Ahmad* for your unwavering support throughout the years. To my husband, *Sarfaraz Ahmad*, who shared the ups and downs of the PhD journey with me and constantly supported me. I would also like to thank my extended family, who have provided me with unwavering support. Of course, I must also thank my friends who made this journey easier. Most importantly, I thank God for giving me the strength to take on this challenge and for surrounding me with such kind people on this journey.

Thank you all. Grazie a tutti.

Use of Generative AI Tools

During the preparation of this thesis, generative AI tools (including large language models) were used to support the writing and editing process. These tools were employed primarily for language refinement, structuring of text, and formatting assistance. All scientific content, experimental design, analysis, and conclusions presented in this thesis are the original work of the author. The author critically reviewed and validated all outputs generated by these tools to ensure accuracy, consistency, and academic integrity.

Nomenclature

Abbreviations

AI	Artificial Intelligence
API	Application Programming Interface
CSV	Comma Separated Value
DL	Deep Learning
GPU	Graphic Processing Unit
GPT	Generative Pretrained Transformer
IR	Information Retrieval
KG	Knowledge Graph
LIR	Linked Interaction Recency
LID	Linked Interaction Diversity
LLaMA	Large Language Model Meta AI
LLM	Large Language Model
L4C	Learning for Credits
L4P	Learning for Pleasure
MAE	Mean Absolute Error
MIUR	Italian Ministry of Education, University and Research
ML	Machine Learning
MOOC	Massive Open Online Course
MOOPER	Massive Open Online Practice-Oriented dataset
MRR	Mean Reciprocal Rank
MSE	Mean Squared Error
nDCG	normalized Discounted Cumulative Gain
NLP	Natural Language Processing
PEEK	Personalised Educational Engagement with Knowledge Topics
PFR	Path Faithfulness Rate
PTD	Path Type Diversity
PTC	Path Type Concentration
RS	Recommender Systems
SEP	Shared Entity Popularity
SED	Shared Entity Diversity

Latin Expressions

i.e.	id est, ‘that is’
e.g.	exempli gratia, ’for the sake of example’
ergo	’therefore’

Numerical Expressions

{n}K	{n} thousands
{n}M	{n} millions

Contents

List of Figures

List of Tables

Chapter 1
Introduction

1.1 Motivation

In an increasingly digital and interconnected economy, professional success depends on the continuous acquisition of knowledge and skills. Rapid technological innovation, evolving organizational structures, and emerging professional roles have shortened skill lifecycles, making lifelong learning a necessity rather than a choice [1]. *Digital education* has become a central ecosystem for lifelong learning, enabling broad access to educational opportunities beyond traditional constraints. It brings together diverse actors (e.g., learners, instructors), educational resources, infrastructures, and strategic elements such as learning objectives and credential models [2]. Driven by its flexibility, online learning has rapidly expanded worldwide [3]. For instance, by 2025, platforms such as *Google Classroom*, *Coursera*, and *Udemy* supported tens of millions of learners and educators globally [4, 5, 6]. While this growth highlights the transformative potential of online education, it also introduces challenges related to resource overload and academic integrity [7]. As a result, learners often struggle to identify opportunities aligned with their goals, and educators faces difficulty in providing personalized support.

The rapid advancement of *Artificial Intelligence* (AI) presents significant opportunities to address challenges in online education. AI techniques support automation, personalization, and decision-making, for example, *recommender systems* guide learners toward relevant resources based on prior experiences [8, 9]. AI-driven analytics have also been explored to support assessment and academic integrity. These developments have attracted substantial interest and investment from researchers, policymakers, and industry, reflecting both the economic and societal impact [10, 11].

Despite this promise, many AI-based educational systems offer limited transparency, raising concerns about trust and learning ecosystem—particularly in lifelong learning contexts where informed decision-making is essential [12, 13]. Recent advances in *Knowledge Graphs* (KG) and *Large Language Models* (LLMs) provide promising directions to address these concerns. KGs enable explicit and traceable modeling of learners, content, and their semantic relationships, while path-based reasoning supports grounded explanations [14]. LLMs further enhance these systems by enabling natural, flexible explanation and interaction [15]. While, path-based reasoning ensures faithfulness and transparency, explanations derived from reasoning paths may remain difficult for learners to interpret. Conversely, unconstrained natural language explanations generated by LLMs risk producing outputs that are unfaithful to the underlying reasoning process. This highlights the need for approaches that combine the structural reliability of path-based explanations with the expressive flexibility of LLMs [16].

This thesis is motivated by the need to have knowledge-aware methods that support transparent and accountable decision-making in lifelong learning contexts. It examines how learner-centered knowledge modeling can enable explainability by construction, how conversational interfaces can facilitate intuitive exploration of educational knowledge graphs, how path-based reasoning methods can support explainable recommendation, and how large language models can transform structured explanations into accessible, human-centered narratives while preserving faithfulness to the reasoning processes.

1.2 Challenges

Recent years have seen a rapid growth in machine learning research and applications, accompanied by increasing attention to the challenges of deploying such technologies in real-world contexts [17]. While these advances have significantly expanded the technical capabilities of AI systems, their application to large-scale online education remains affected by several challenges [18]. In particular, the integration of AI into educational contexts raises issues related to knowledge representation, interaction, reasoning, explainability, and trust that are not fully addressed by existing background work on ontologies, recommender systems, and explainable AI. This thesis focuses on the following four key challenges, each of which motivates the subsequent chapters of the thesis.

- **Modeling knowledge in a learner-centered and explainable manner.** Educational data is often weakly structured, and centered on resource consumption rather than learner reasoning. Designing knowledge representations that explicitly capture learners, educational content, and their semantic relationships—while remaining suitable for explainable reasoning, remains a fundamental challenge.
- **Enabling intuitive interaction with structured educational knowledge.** Even when educational knowledge is represented through structured models such as knowledge graphs, interacting with these structures typically requires technical expertise. Providing learners with accessible, natural, and flexible interaction mechanisms to explore educational knowledge for informed learning decisions is an open problem.
- **Supporting explainable recommendation through knowledge aware reasoning.** Most existing educational recommender systems rely on weak interpretable models, limiting their suitability for transparent decision support. To have recommendation methods grounded in explicit reasoning paths — capable of transparently con-

necting learners to educational resources, and learning trajectories — poses significant challenges, particularly in lifelong learning contexts.

- **Transforming structured explanations into user-friendly narratives.** While path-based explanations offer transparency and interpretability at the system level, they are not always accessible to end users. Bridging the gap between structured reasoning outputs and human-centered explanation delivery—without sacrificing faithfulness—remains a key challenge for informed decision-making.

Addressing these challenges is essential to ensure that AI-driven educational systems remain transparent, interactive, and trustworthy at scale. This thesis addresses these challenges by combining learner-centered knowledge modeling, conversational interfaces, path-based recommendation methods, and language model–based explanation generation within a unified framework.

1.3 Contributions

In this thesis, the design, development, and evaluation of explainable AI methods for lifelong learning are systematically investigated with the goal of empowering educational environments through transparent, interactive, and trustworthy AI methods. The focus is on how machines can represent educational knowledge, reason over learning-related data, and communicate recommendations and explanations in a form that is meaningful to learners. To this end, the thesis presents data structures, methods, and models to support explainable educational recommendation and exploration. The proposed contributions address different stages of the instructional and decision-making pipeline, ranging from knowledge representation and reasoning to user interaction and explanation generation. Figure 1.1 provides an overview of how these contributions form a unified framework, progressing from ontology-driven knowledge representation to explainable reasoning and natural language explanation generation. These contributions are grounded in learner-centered design principles and informed by continuous feedback from learners and educators across multiple evaluation settings. Going into detail, this thesis makes the following contributions:

- **Ontologies and knowledge graphs**. We introduce learner-centered ontology design principles and related knowledge graphs that enable explainable educational recommendation by construction. By explicitly modeling learners, educational resources, and their semantic relationships, the proposed knowledge representations support reasoning paths for transparent decision-making (*Chapter 3*).

- **Interfaces for knowledge graph exploration**. We propose conversational interfaces driven by large language models that allow learners to interact with and explore educational knowledge graphs using natural language, lowering the barrier to access structured data and supporting decision-making (*Chapter 4*).

Fig. 1.1: Integrated pipeline for explainable lifelong learning. Learner data and educational resources are structured through ontology-driven knowledge graphs, accessed through an LLM-based interface (EDGE), processed through path reasoning for explainable recommendations, and finally transformed into natural language explanations to support user understanding and decision-making

.

- **Explainable path-reasoning methods**. We investigate path-based recommendation methods over educational knowledge graphs, enabling recommendations and explanations that are grounded in transparent semantic reasoning (*Chapter 5*).
- **Natural language explainers**. Building on the above structured explanations, we explore language model–based explanation generation techniques that transform path-based reasoning outputs into user-friendly natural language narratives, while preserving faithfulness to the underlying reasoning process (*Chapter 6*).

Overall, this thesis provides design guidelines, methodological insights, and empirical evidence for researchers and practitioners working on explainable AI for education. By integrating structured reasoning with interactive and language-based explanation mechanisms, the proposed approaches aim to improve transparency, trust, and user acceptance in AI-driven educational systems. In doing so, the thesis contributes to advancing the understanding of how explainability can be systematically supported.

1.4 Outline

The remainder of this thesis is organized as follows:

Chapter 2 provides the background necessary to contextualize this work. It reviews foundational concepts and related research on knowledge graphs, educational recommender systems, and large language models. This chapter positions the thesis within the existing literature and highlights the gaps that motivate the proposed approaches.

Chapter 3 focuses on learner-centered knowledge modeling for explainable educational recommendation. It introduces the design principles and structure of an educational ontology that enables explainability by construction. The chapter discusses how learners, educational resources, and learning objectives are modeled, and how ontology-derived reasoning paths support transparent explanations.

Chapter 4 presents a conversational interface driven by large language models for educational KG exploration. The chapter investigates how natural language interaction can mediate between learners and structured educational knowledge, enabling intuitive access, and exploration of knowledge graphs.

Chapter 5 introduces methods for path-based educational recommendation over knowledge graphs. It examines how structured reasoning paths can be generated and exploited to produce explainable recommendations grounded in explicit semantic relations, providing the algorithmic foundation for faithful recommendation and explanation.

Chapter 6 explores large language model–based explanation generation and empirical evaluation. The chapter investigates how path-based explanations can be transformed into user-friendly natural language narratives while preserving faithfulness to underlying reasoning processes, and presents an empirical assessment of explanation quality.

Finally, *Chapter 7* concludes the thesis by summarizing the main contributions, discussing their implications for explainable AI in lifelong learning, and outlining directions for future research.

Chapter 2

Background

This chapter provides essential context around Knowledge graph, recommendation models, and explainability concepts leveraged by this thesis.

2.1 Ontologies and Knowledge Graphs in Education

Introduction Ontologies and knowledge graphs have become central technologies for representing, integrating, and reasoning over complex domains characterized by heterogeneous entities and rich semantic relationships [19]. In the context of lifelong learning, where educational data spans learners, competencies, learning resources, objectives, and institutional constraints, structured knowledge representations provide a principled foundation for organizing educational ecosystems in a transparent and machine-interpretable manner [20]. Unlike unstructured or purely statistical representations, ontologies and knowledge graphs explicitly encode domain semantics, enabling reasoning, interoperability, and explainability properties that are increasingly critical for AI-driven educational systems [21].

Recent advances in educational technology have highlighted the limitations of traditional data models that focus primarily on content delivery or user–item interactions [22]. Lifelong learning scenarios demand representations that reflect not only what resources are consumed, but also why they are relevant in relation to learner goals, skill gaps, and professional trajectories [23]. Ontologies and knowledge graphs offer mechanisms to capture such relationships explicitly, supporting more informed and accountable decision support in educational environments.

Related Work Ontologies provide formal, shared conceptualizations of a domain by defining entities, attributes, and relations using logic-based representations. Knowledge graphs extend these ideas by instantiating ontological concepts with large-scale data, forming graph-structured representations that can be queried and reasoned over. In recent years, knowledge graphs have gained substantial traction across AI applications, including education, due to their flexibility, interpretability, and compatibility with reasoning methods [19]. In educational contexts, knowledge graphs have been used to represent curricula, learning resources, skills, competencies, and learner profiles [24], enabling tasks such as curriculum mapping, learning path generation, and resource recommendation [25]. Recent surveys highlight their growing role in supporting intelligent tutoring systems, recommender systems, and learning analytics by providing structured and semantically rich representations of educational data [21].

Post-2020 research has increasingly emphasized learner-centered knowledge modeling, where learners are not treated merely as consumers of content but as entities with evolving competencies, objectives, and contextual constraints [26]. For example, competency-based knowledge graphs have been proposed to model skill acquisition and progression across formal and informal learning contexts, supporting lifelong learning and workforce upskilling scenarios [24]. Similarly, educational ontologies have been extended to incorporate learning objectives, assessment outcomes, and contextual metadata, enabling richer reasoning about learning relevance and outcomes. Beyond data integration, a key advantage of ontologies and knowledge graphs lies in their capacity to support explainability by construction. Because relationships between entities are explicitly modeled, reasoning processes, such as traversing paths between learners and learning resources can be inspected, traced, and communicated [27]. This property has positioned knowledge graphs as a promising foundation for explainable AI, particularly in domains where transparency and accountability are required [28].

Recent work in explainable AI has highlighted that symbolic and graph-based representations offer inherent advantages over black-box models when explanations must be faithful to underlying reasoning processes [29]. In educational systems, explainability grounded in knowledge graphs enables explanations that reference pedagogical concepts such as prerequisite relationships, competency alignment, or learning objectives, making explanations more meaningful to learners and educators alike [19]. However, existing educational knowledge graphs often remain resource-centric, emphasizing content classification or metadata annotation, while providing limited support for learner-centered reasoning and explanation [30]. Moreover, many implementations prioritize scalability and data coverage over semantic richness, resulting in shallow representations that limit the depth and interpretability of reasoning paths [31]. Additionally, explanation interfaces (even when provided) can increase perceived trust in ways that are not always calibrated to correctness, which further motivates careful faithfulness-focused design [32].

Identified Challenge The literature reviewed above directly relates to the first two challenges identified in Chapter 1 (i.e., modeling and exploring educational knowledge in a learner-centered and explainable manner). While ontologies and knowledge graphs provide powerful foundational tools, their application to lifelong learning remains affected by fragmentation, limited learner modeling, and insufficient alignment with explainable reasoning requirements. Specifically, current approaches often lack:

- Explicit modeling of learner goals and decision contexts,
- Structured representations suitable for reasoning,
- Systematic principles that align knowledge with explainable decision support.

These gaps motivate the need for learner-centered ontology design approaches that go beyond background representations and support traceable reasoning.

Contribution Building on the foundational concepts discussed in this section, the subsequent chapters of this thesis (Chapters 3 and 4) advance the state of the art by introducing learner-centered ontology design principles tailored to explainable decision-making. Rather than treating ontologies solely as data schemas, the proposed approach leverages knowledge modeling as a core mechanism for enabling explainability by construction. The educational ontology developed in this thesis explicitly models learners, competencies, learning objectives, and their semantic relationships, providing the structural basis for the path-based reasoning and explanation methods introduced in later chapters. In this way, ontologies and knowledge graphs serve not only as background technologies, but as the conceptual and methodological foundation upon which the thesis builds its contributions to knowledge-aware and explainable decision support.

2.2 Recommendation Techniques

Introduction Recommender systems play a central role in online learning environments by supporting learners in identifying relevant educational resources, learning paths, and opportunities aligned with their goals [33]. As the scale and diversity of educational content continue to grow, recommendation technologies have become essential components of learning platforms, mitigating information overload and enabling personalized learning experiences [34]. In lifelong learning, contexts characterized by heterogeneous learner profiles, evolving objectives, and long-term skill development, recommendation systems increasingly function as decision-support tools rather than simple content filters [34]. Traditional recommendation approaches, however, face significant limitations when applied to educational settings. Learning decisions often have long-term consequences, requiring recommendations that are not only accurate but also transparent, pedagogically meaningful, and trustworthy [35]. These requirements have motivated growing interest in recommendation models that incorporate domain knowledge, learner context, and explainability as priority design considerations [36].

Existing Educational Recommendation Methods Educational recommender systems have drawn from general-purpose recommendation techniques, including collaborative filtering, content-based filtering, and hybrid approaches [37, 38]. Collaborative filtering methods leverage patterns in learner resource interactions to infer preferences, while content-based methods rely on resource features and learner profiles to generate recommendations; hybrid models combine these strategies to improve robustness

[37]. While these approaches have demonstrated effectiveness in various educational scenarios [38], recent studies highlight their limitations in lifelong learning contexts [38]. Collaborative filtering models often suffer from sparsity and cold-start problems, particularly for new learners or emerging competencies [39]. Content-based approaches, although more interpretable, depend heavily on the quality and granularity of available metadata and may struggle to capture complex pedagogical relationships [38].

To address these issues, more recent work has explored neural and deep learning based recommender systems, including sequence-aware and representation learning models that capture temporal learning patterns and learner trajectories [39, 40]. While these models improve predictive performance [40], they often operate as black boxes, offering limited insight into why a particular recommendation is made an issue that is especially problematic in educational decision support scenarios [41].

Existing Knowledge-Aware and Graph-Based Recommendation Approaches In response to the limitations of purely data-driven models, knowledge-aware recommender systems have gained increasing attention [42, 43]. These approaches incorporate structured domain knowledge, often in the form of knowledge graphs, to enhance recommendation quality, robustness, and interpretability [42]. Knowledge graphs enable the explicit modeling of relationships between learners, resources, competencies, and learning objectives, providing a richer context for recommendation [44, 25]. Graph-based recommendation techniques exploit the relational structure of knowledge graphs to generate recommendations through graph traversal, embedding, or reasoning mechanisms [42, 45, 46]. Recent surveys highlight the growing adoption of knowledge graph-based recommender systems across domains (including education) due to their ability to integrate heterogeneous data sources and support explainable reasoning [42, 44].

In educational settings, graph-based models have been applied to recommend courses, learning paths, and skill development opportunities by leveraging relationships such as prerequisite structures, competency alignment, and curricular dependencies [25, 47]. Importantly, these approaches allow recommendations to be grounded in explicit semantic relations, making them more interpretable and pedagogically meaningful than purely statistical methods [42, 44].

Existing Path-Based Reasoning for Explainable Recommendation Among knowledge-aware approaches, path-based reasoning methods have emerged as a promising direction for explainable recommendation [48]. These methods generate recommendations by identifying meaningful paths in a knowledge graph that connect learners to educational resources through intermediate entities such as objectives, or prior learning experiences [49]. The resulting paths provide a natural basis for explanations, as they encode the reasoning process underlying a recommendation [50]. Recent work has shown that path-based recommendation methods can support both personalization and explainability by aligning recommendations with learner-specific contexts and educational semantics. In educational domains, path-based reasoning enables explanations such as competency alignment or prerequisite satisfaction, which are more intuitive and actionable for learners and educators [34][51] .

However, several challenges remains as generating relevant and scalable reasoning paths in large knowledge graphs is computationally demanding, and selecting paths that are both informative and educationally meaningful requires careful design. Moreover, questions remain regarding the generalizability of path-based recommendation methods across different educational contexts, learner populations, and knowledge graph structures particularly in lifelong learning.

Identified Challenge The literature reviewed in this section directly relates to the third challenge identified in Chapter 1 (i.e., supporting explainable recommendation through structured reasoning). While knowledge-aware and path-based recommendation methods offer promising solutions, existing approaches often prioritize technical performance over educational relevance, transparency, or long-term decision support. Specifically, current research reveals gaps in:

- Systematic integration of learner-centered knowledge models,
- Limited explainability and faithfulness in educational recommendations,
- Understanding how path-based reasoning methods perform for lifelong learning.

These gaps motivate the need for recommendation approaches that are explicitly designed to support explainable and accountable decision-making in education.

Contribution Building on the background presented in this section, the subsequent algorithmic chapter of this thesis (Chapter 5) investigates path-based recommendation methods over educational knowledge graphs in depth. The proposed methods leverage learner-centered knowledge representations to generate recommendations grounded in explicit semantic reasoning paths. In contrast to opaque predictive models, the approaches explored in this thesis emphasize traceability, faithfulness, and pedagogical relevance, providing a robust foundation for explainable recommendation and explanation generation in later chapters. In this way, recommendation models are treated not merely as predictive mechanisms, but as core components of a knowledge-aware decision-support framework for lifelong learning.

2.3 Explainability in Knowledge-based Systems

Introduction Explainability has emerged as a central requirement for systems deployed in high-stakes and human-centered domains, including education [52]. In lifelong learning environments, AI systems increasingly influence decisions related to learning pathways, skill development, and career progression. As a result, learners and educators must be able to understand, and trust the recommendations and guidance provided by these systems. Explainability is therefore not merely a technical add-on, but a foundational property that enables informed decision-making, accountability, and user acceptance [53].

Within educational recommender systems, explainability is particularly critical due to the pedagogical implications of recommendations [38]. Learners need to understand why a specific resource or learning track is suggested, how it relates to their goals or prior knowledge, and whether alternative options exist. Knowledge-based systems such as those built on ontologies and knowledge graphs offer a natural foundation for explainability by explicitly representing domain concepts and reasoning processes [42, 44]. However, translating this structured reasoning into explanations that are accessible and meaningful to end users remains an open challenge [54].

Existing Methods for Explainable Machine Learning Explainability has long been associated with symbolic and knowledge-based AI systems, where reasoning processes are explicitly encoded using rules, logical relations, or graph structures. Unlike black-box models, knowledge-based systems enable intrinsic explainability, as their outputs can be traced back to interpretable entities and relationships [55]. In educational contexts, this traceability supports explanations grounded in concepts such as competencies, prerequisites, learning objectives, and learner profiles [56]. Recent work has emphasized the value of knowledge graphs and rule-based reasoning as foundations for explainable recommendation and decision support [57]. By exposing reasoning chains or paths, these systems allow users to inspect how conclusions are derived, thereby supporting transparency and trust. In recommender systems, such explanations often take the form of path-based or feature-based justifications, linking learners to recommended resources through meaningful semantic relations [58].

Despite these advantages, prior research also highlights usability challenges. Explanations derived directly from symbolic reasoning structures may be technically correct but cognitively demanding for non-expert users. In educational settings, learners vary widely in background knowledge and explanatory needs, making it difficult to present raw reasoning outputs in a universally accessible manner.

Existing Post-hoc and User-Centered Explainability Approaches To address usability concerns, research in explainable AI has increasingly focused on user-centered explanation design [59]. Rather than exposing full reasoning traces, explanations are adapted to user goals, context, and cognitive load [60]. In educational systems, this includes tailoring explanations to learner expertise, learning objectives, or motivational needs [61]. Recent studies emphasize that explainability should be evaluated not only in terms of technical faithfulness, but also perceived usefulness, clarity, and actionability [62, 63]. This shift has motivated hybrid approaches that combine structured reasoning with natural language explanations, visualizations, or interactive interfaces [64].

However, a persistent challenge is balancing faithfulness and accessibility. Simplified or post-hoc explanations risk obscuring or distorting the underlying reasoning process, potentially misleading users [65]. This issue is particularly pronounced in educational recommender systems, where explanations may influence long-term learning decisions.

Existing Large Language Models and Explanation Generation LLMs have recently introduced new possibilities for explanation generation due to their strong natural language generation capabilities [27]. In educational systems, LLMs can produce fluent, context-aware explanations that adapt to learner preferences and interaction styles. As a result, LLMs are increasingly explored as interfaces between complex AI systems and human users [66]. Recent research investigates the use of LLMs for explanation generation, question answering, and conversational interaction in knowledge-based systems [67]. When combined with structured knowledge sources, LLMs can translate symbolic reasoning outputs into coherent natural language narratives, potentially lowering the barrier to understanding AI-driven decisions [68].

At the same time, the use of LLMs raises critical concerns for explainability. LLM-generated explanations may appear plausible without being faithful to the actual reasoning process, leading to risks of hallucination or overgeneralization [69]. Several studies emphasize that LLMs should not be treated as reasoning engines themselves, but rather as explanation mediators grounded in verifiable symbolic inputs.

Existing Knowledge-Grounded Explainability with LLMs To address these concerns, recent work focuses on knowledge-grounded explanation generation, where LLM outputs are constrained or guided by structured reasoning artifacts such as knowledge graph paths, rules, or provenance metadat [70]. This approach aims to combine the strengths of knowledge aware methods (faithfulness and traceability) with those of LLMs (linguistic flexibility and user adaptation) [71]. In educational contexts, this paradigm enables explanations that are both pedagogically meaningful and technically faithful. For example, a path-based recommendation can be transformed into a narrative explanation that explicitly references competencies, learning objectives, or prior activities, while remaining anchored to the underlying reasoning path [72].

Despite growing interest, open challenges remain regarding: how to systematically structure LLM prompts using reasoning artifacts, how to control explanation properties such as tone, granularity, and pedagogical intent, and how to evaluate faithfulness and learner perception of LLM-generated explanations in real educational settings. These challenges are amplified in lifelong learning environments, where learners' goals and contexts evolve over time.

Identified Challenge This section directly addresses the fourth challenge outlined in Chapter 1 (i.e., bridging structured reasoning and user-centered explanation delivery). While knowledge-based systems provide a strong foundation for explainability, and LLMs offer powerful language generation capabilities, existing approaches often fail to integrate these components in a principled and systematic manner.

Specifically, the literature reveals gaps in:

- ensuring faithfulness of explanations to underlying reasoning processes,

- adapting explanations to diverse learner contexts without sacrificing transparency,
- evaluating explanation effectiveness beyond surface-level fluency.

These gaps motivate the need for explanation frameworks that treat explainability as an end-to-end design problem, spanning representation, reasoning, and communication.

Contribution Building on this background, the final pillar of this thesis (Chapter 6) explores LLM-based explanation generation grounded in structured reasoning outputs from educational knowledge graphs. Rather than relying on unconstrained language generation, the proposed approaches use path-based explanations as explicit inputs to LLMs, enabling the generation of user-friendly narratives that preserve faithfulness to the underlying recommendation logic. Through empirical evaluation, this thesis investigates how different explanation properties—such as tone, context, and reasoning affect learner perception, trust, and understanding. In doing so, the work contributes methodological insights into how large language models can be responsibly integrated into explainable, knowledge-based educational systems.

Chapter 3

Methods for Learner-Centered Knowledge Graphs Creation

Research Highlights
- Proposed *LOXER*, a learner-centered ontology.
- Support explainability through ontology-driven reasoning paths.
- Usage of LOXER under a course recommendation task.

3.1 Introduction

Lifelong learning has become an essential requirement of modern education systems, as continuous technological change, evolving professional requirements, and the increasing need for ongoing skill development. Despite traditional academic learning, lifelong learning environments are characterized by heterogeneous learner profiles, diverse learning objectives, and non-linear learning trajectories [73]. These characteristics pose significant challenges for educational systems tasked with supporting personalized and effective learning experiences [74]. AI-driven educational recommender systems have demonstrated considerable potential in supporting personalized learning paths by aligning learners with appropriate educational content [75]. Knowledge graphs provide structured, semantically rich representations of domains by explicitly modeling entities, their attributes, and the relationships between them. In educational settings, knowledge graphs can represent learners, learning resources, competencies, subjects, and educational objectives, enabling reasoning over structured educational knowledge.

Ontologies form the core of knowledge graphs, which act as faithful, semantic-rich sources for training models in delivering explainable recommendations. These models learn to extract logical paths between learners and resources to be recommended within the knowledge graph, according to behavior- and content-based patterns. Extracted paths are then used not only to provide recommendations, but also to generate accompanying textual explanations. Including explanations in the provided recommendations is essential, not only to promote informed decision-making aligned with individual goals and interests [76] but also to comply with the legal mandate for the right to explanation [77]. Explanations can, in turn, strengthen trust in the system, lead to higher engagement and retention, ignite curiosity, and streamline decision processes [78]. Despite the potential of this approach, current ontologies derived from the traditional learner-resource interaction data fall short in terms of richness from an educational perspective. Conversely, general-purpose ontologies, while comprehensive in educational aspects, are overly complex for recommendation tasks. Unfortunately, a suboptimal ontology might prevent to articulate reasoning paths, and thus explanations, relevant for learners within the knowledge graph. Educational recommender systems can shape learners' academic and career trajectories by directing them toward pertinent university courses and facilitating lifelong competency development [73]. To offer personalized suggestions, these systems typically utilize learners' historical interactions with educational resources and content-based relationships among the latter. Despite their promising accuracy, the inner workings of existing systems is often opaque to the learner, prompting for more transparent processes. Including explanations in the provided recommendations is essential, not only to promote informed decision-making aligned with individual goals and interests [76] but also to comply with the legal mandate for the right to explanation [77]. Explanations can, in turn, strengthen trust in the system, lead to higher engagement and retention, ignite curiosity, and streamline decision processes [78].

Contemporary recommendation methods strive to enhance transparency by offering textual explanations alongside recommended items [75]. Explanations are created by applying pre-defined templates or generative methods to reasoning paths within a knowledge graph[1], identified by the recommendation model. To this end, models employ reinforcement learning [48, 49, 80, 81] or language modelling [82, 83], whose effectiveness heavily depends on the semantic structure of the knowledge graph and, hence, the quality (e.g., completeness, expressiveness, and relevance) of the adopted ontology[2] [85]. While tailored ontologies have been defined to build knowledge graphs that enable explainable recommendation in domains like music [81] or cinema [48, 49], the need for ontologies to explain educational recommendations is relatively recent. Initial applications of the above explainable recommendation methods in education have focused on models that merely rely on the implicit ontology inherently embedded into the interaction data used for traditional non-explainable recommendation [86]. Ontologies ingrained within the learner-resource interaction data, such as tracking clicks on educational content [87, 86], might lack crucial explanatory aspects like the semantic connections between learning resources and their relevance to subjects, competencies, and occupations. General-purpose ontologies, while covering a broad spectrum of educational details [88, 89], can be excessively fine-grained for explainable recommendation methods, which typically manage paths of limited length. They might also include less

[1] A knowledge graph is a structured representation of knowledge that captures entities, their attributes, and relationships in a graph-based format [79].

[2] Ontologies are used to formally model the structure of a system by defining entities, their properties, and relationships; knowledge graphs often incorporate them to establish a standardized framework for organizing and interlinking information [84].

relevant details for the recommendation purpose like administrative structures or taxes information. Moreover, existing studies tend not to engage with the end users, i.e., the learners, since the beginning of the design, resulting in ontologies (and so knowledge graphs and explainable recommendations) not aligned with the explanatory aspects that learners find relevant [90, 91]. Hence, there is a need for learner-centered ontologies tailored to explainable recommendation in education.

To address this limitation, this chapter introduces LOXER, a Learner-centered Ontology for eXplainable Educational Recommendation (LOXER). This ontology is created by standardizing the representation of the educational recommendation domain under a multidisciplinary approach to explainable artificial intelligence in education [92]. In the formulation of our ontology, we integrated insights from various sources, incorporating feedback from a local co-creation group of learners for a preliminary validation. We leveraged observations from existing datasets tailored to traditional educational recommendation [93, 94, 95, 24, 96] to avoid overlooking crucial information for model learning. We also established connections with other (educational) ontologies [88, 97], prioritizing the integration of well-known vocabularies, if feasible, to align with FAIR[3].

Subsequently, we conducted a comprehensive evaluation of the explanation types enabled by LOXER under a course recommendation task. Relevant, example meta-paths were defined from our ontology and corresponding explanation templates were created in another stage of local co-design. We then developed a questionnaire encompassing two tasks. Firstly, learners ranked explanations (one for each meta-path) related to a given recommended course, considering their perceived importance in the presented scenario. In the second task, learners individually assessed the same explanations based on seven properties (effectiveness, decision speed, motivation, satisfaction, correction ease, transparency, and confidence boost), coming from prior work on explainable artificial intelligence evaluation [76]. In this stage, we extended the validation to a gender-balanced population of 100 university and lifelong learners from around the world. Our results show LOXER's ability to adhere to learners' decision-making processes and establish a semantic structure for knowledge graphs tailored to explainable recommendations. Our findings also represent a blueprint on the extent to which learners value explanatory factors for different meta-paths, serving as actionable insights for novel, learner-centered models in this research area. The main contributions are three-fold:

- The introduction of LOXER, a learner-centered ontology for explainable educational recommendation.
- Integration of multidisciplinary insights into the ontology design to balance expressiveness, learner relevance, and structural simplicity.
- Empirical evaluation of explanation types supported by LOXER in a course recommendation scenario.

[3] https://joinup.ec.europa.eu/collection/oeg-upm/news/fair-ontologies

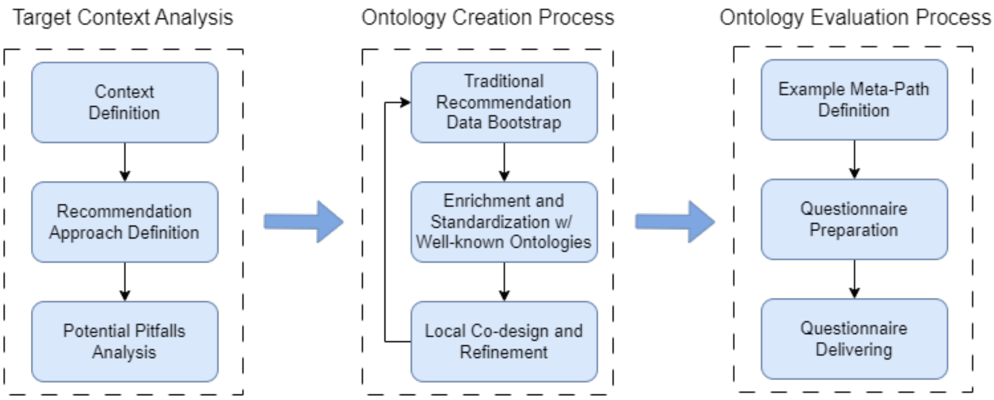

Fig. 3.1: Overview of our methodology. The study begins with an in-depth analysis of the targeted educational context. The ontology creation process involves bootstrapping from traditional recommendation data, enrichment based on existing (educational) ontologies, and local co-design and refinement. The evaluation phase includes defining meta-paths, generating textual explanations through templates, designing questionnaires, and delivering them through Prolific to assess the ontology's effectiveness.

The rest of this chapter is structured as follows: Section 3.2 outlines the methodology employed for the development and evaluation of the ontology, with a focus on a learner-centered perspective. Section 3.3 presents the outcomes of the proposed comprehensive assessment, highlighting various explanation types facilitated by the developed ontology. Lastly, Section 3.4 concludes the chapter with a discussion and implications.

3.2 Methodology

The methodology is structured around three main phases: (i) analysis of the target educational context, (ii) ontology creation through an iterative, learner-informed process, and (iii) design-oriented assessment of the explanation types enabled by the ontology. An overview of the methodology is illustrated in Figure 3.1.

3.2.1 Target Educational Context Analysis

Following the framework for explainable artificial intelligence design in education proposed by [92], the first phase of the methodology focuses on characterizing the target educational context and identifying the design constraints relevant to explainable recommendation systems based on knowledge graphs. The objective of this phase is not to restrict the applicability of the ontology, but to ground its design in realistic and representative educational scenarios. Without loss of generalizability, the analysis considers a course recommendation task as a reference scenario.[4] Course recommendation

[4]The ontology is designed to support recommendation tasks involving other educational entities, such as learning resources or occupations, as well.

is widely adopted in educational platforms and provides a suitable abstraction for analyzing learner–resource interactions, prerequisite structures, and learning objectives.

Context Definition. Explanations for educational recommendations are integrated into platforms for learners. Therefore, the provided suggestions should primarily resonate with the learners. Without loss of generality, we concentrate on learners within two educational contexts. We consider formal learning scenarios conducted in universities or training organizations, leading to diplomas and certificates. University learners typically fall within the age range of 19 to 27. Adult learners, on the other hand, are generally older and primarily include workers undergoing training as part of their workplace plans. We have also envisioned non-formal learning scenarios for the same learners, outside universities and workplaces, driven more by their intrinsic desire to learn. By encompassing different learning scopes (formal and non-formal) and two types of learners (university and adult), we aim to ensure the generality of our ontology for adoption in various scenarios. Additionally, we plan to analyze how these two factors influence the perception of the explanations facilitated by our ontology.

Recommendation Approach Definition. Path reasoning[5] in explainable recommendation for the target context involves extracting meaningful logical paths between learners and resources within a constructed knowledge graph, following an underlying ontology. Once paths are extracted, they form the foundation for generating textual explanations accompanying recommended resources. For instance, assume to consider an ontology that captures relationships such as course completion and course prerequisites. The recommendation model identifies that the learner completed the course "Introduction to Data Science," which includes the subject "Numpy," that is a prerequisite for the course "Machine Learning for Data Analysis." Based on this path, a corresponding explained recommendation for that learner might be "Consider enrolling in the Machine Learning for Data Analysis course. The Introduction to Data Science course you attended covered the subject of Numpy which serves as a prerequisite for the suggested course." The explanations provide transparent insights into why a particular course is being suggested, offering a clearer rationale grounded in the learner's historical interactions and the ontology's structure.

Potential Pitfalls Analysis. Several potential pitfalls must be considered while developing ontologies for the above explainable recommendation models based on knowledge graphs within the target educational contexts. Overly complex ontologies may hinder interpretability, emphasizing the need for simplicity and clarity in design. Failing to align ontologies with learner expectations can result in misrepresentations, highlighting the importance of learner-centric development. Lack of relevance in recommendation paths may occur if ontologies are not tailored to the educational context, potentially leading to irrelevant explanations. Ontologies must dynamically adapt to evolving learning contexts to avoid becoming outdated, ensuring accuracy and relevance. Additionally,

[5] Our choice was motivated by the fact that path-reasoning methods have been proven to achieve the best trade-off between utility and explainability [81].

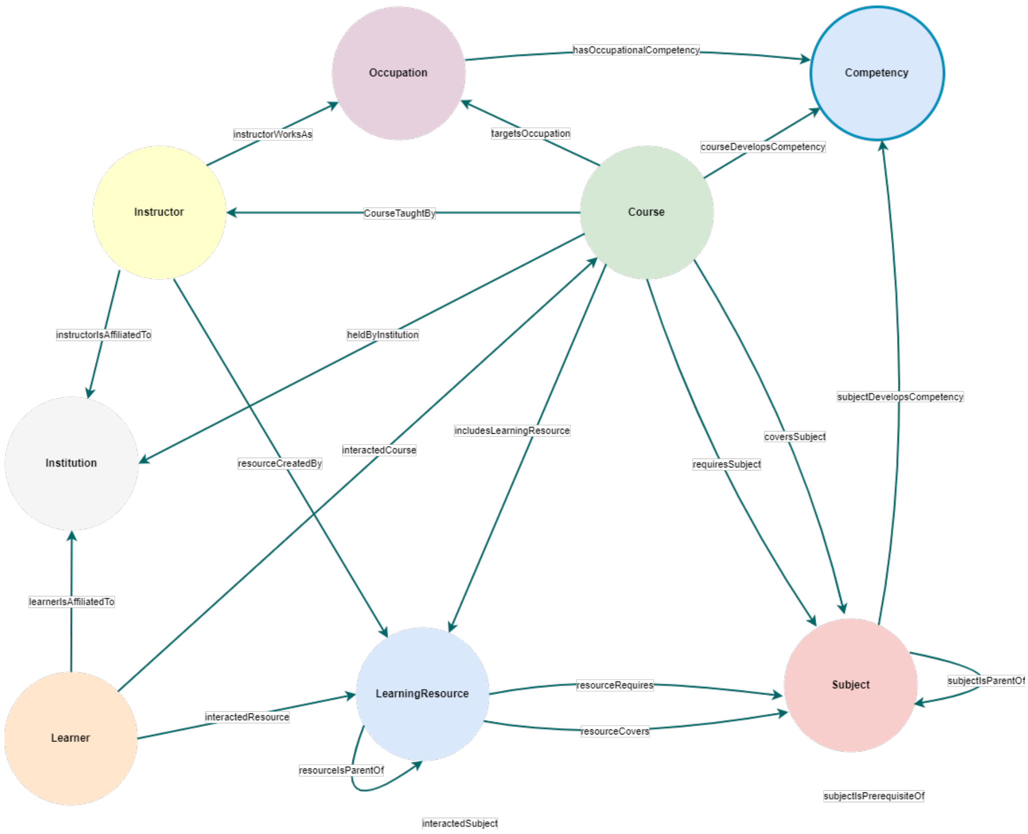

Fig. 3.2: Schematic ontology representation. We summarize the ontology proposed in this study, obtained by formally defining entities, their properties, and relationships. For conciseness, we removed the list of entity properties from this schematic view.

inaccurate or outdated information within the knowledge graph can compromise the integrity of recommendations. Limited coverage of educational aspects may lead to biased suggestions, necessitating a holistic representation of diverse domains. Handling multidisciplinary contexts is crucial, requiring ontologies to integrate seamlessly. Complexity in ontology mapping and ineffective explanations pose challenges that need consideration. Balancing these aspects is vital for the success of ontology-driven, explainable recommendation systems in education.

3.2.2 Ontology Creation Process

Recognizing potential pitfalls identified in our contextual analysis, we adopted a co-design approach involving collaboration with learners and instructors throughout the conceptualization, development, and validation phases. Our local group, consisting of two university students, two doctoral students, and three experienced researchers in teaching, played a crucial role in this process. To bootstrap the initial version of the ontology, we leveraged information from state-of-the-art, public, large-scale datasets specifically designed for recommendation tasks, including COCO [93], MOOCube [94], MOOPer [95], EduKG [98], and PEEK [96]. Our goal was to identify relevant yet common entities, relationships, and properties, resolving vocabulary conflicts, and ensuring that less frequent but significant elements of the domain were duly incorporated.

With this ontology version, we progressed to the integration phase by establishing connections with other educational ontologies. We prioritized the inclusion of well-known vocabularies, with a specific focus on concepts from ontologies such as [88], ensuring alignment with FAIR principles. This strategic integration resulted in a robust and interconnected ontology, significantly enhancing its potential for widespread use and alignment with established standards in the educational domain. Throughout the integration phase with other educational ontologies, we maintained active involvement with our local co-design group. This iterative, participatory approach contributed to the ontology's refinement, making it an effective tool for explainable recommendation.

In our resulting educational ontology[6], the `Learner` takes center stage, embodying the protagonist on a quest for specialized knowledge or skills. This personalized approach caters to the unique needs of each individual learner. Academic or training paths are represented by `Course`, emphasizing recommendations tailored to individual learners. Diverse educational content, labeled as `Learning Resource`, encompasses videos, fragments, chapters, and exercises. Interactions between learners and these resources provide valuable insights into preferences and engagement patterns. Another building block of our educational ontology are `Subject`, spanning topics. Beyond individual learning, our ontology considers the broader context, acknowledging `Institution` as educational organizations, `Instructor` as guiding figures, and `Occupations` as real-world applications of acquired knowledge. The connections within the ontology breathe life into entities. The `interacted` relation captures dynamic engagement, shaping a personalized learning profile. The `is_parent_of` relation introduces hierarchy, aiding in navigating knowledge levels. Prerequisites, highlighted by the `pre_requisite` relation, form a structured learning path. Our ontology also included the `covers` relation, offering a nuanced understanding of cognitive levels. Figure 3.2 provides a schematic overview of the ontology.

In essence, our ontology is a dynamic framework designed to understand the educational journey, recommend personalized specializations, and provide transparent explanations rooted in learner interactions, preferences, and the semantics of educational content. By embracing the learner as the central character, it aims to foster a tailored learning path that aligns with individual needs.

3.2.3 Ontology Evaluation Process

To assess whether the ontology supports meaningful explanation structures from a learner perspective, a design-oriented evaluation was conducted focusing on the types of explanations enabled by the ontology, rather than on algorithmic performance.

[6] https://u.garr.it/sWjhQ

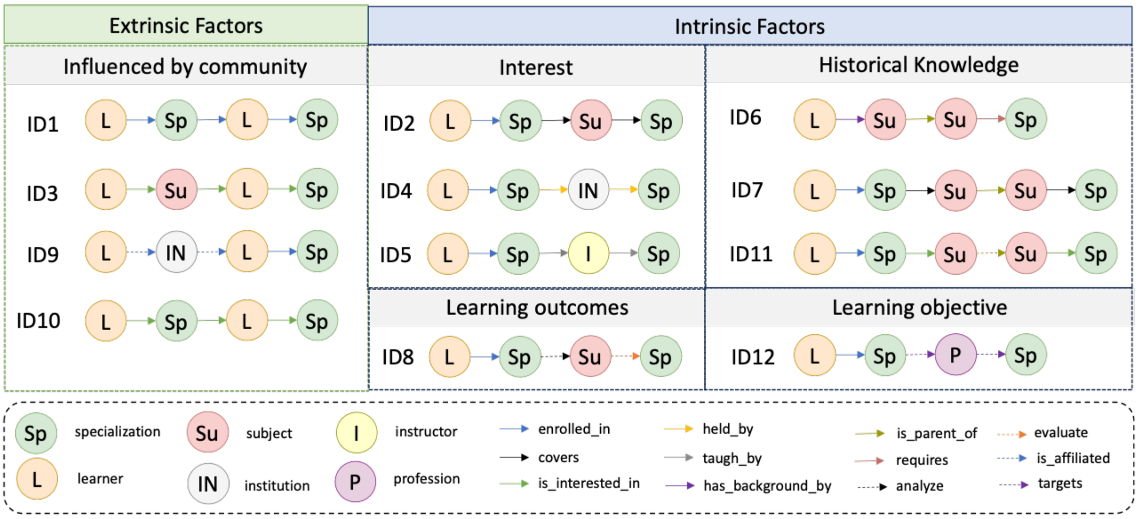

Fig. 3.3: Illustration of example meta-paths in our ontology. Relevant meta-paths, with learners as heads and courses as tails, extracted by involving the same local co-design participants. Categorization into *extrinsic* (e.g., community influence) and *intrinsic* factors (e.g., personal interests) was done.

Example Meta-Path Definition. In our investigation, we delved into meta-paths in our ontology, which represent specific patterns of relationships between entities. Specifically, we explored the meta-path "learner - enrolled in - course - held by - institution," signifying the connection between a learner and an institution through their course enrollments. Focusing on meta-paths with a learner as the head and a specialization as the tail, we selected 12 significant paths based on their relevance and number of hops involved. This meta-path extraction was conducted with the same local group of learners who actively participated in the co-design phases, ensuring a consistent and contextually relevant representation. These paths were categorized into *extrinsic* and *intrinsic* motivational factors, where *extrinsic* factors involve community influence, and *intrinsic* factors encompass personal interest, historical knowledge, learning outcomes, and learning goals. We associated each meta-path with an explanation motivating why a course was recommended. These explanations, designed for both *extrinsic* and *intrinsic* motivations, aim to provide transparent insights into the rationale behind each recommendation. Figure 3.3 describes the example meta-paths.

Questionnaire Preparation. We created a fifteen-minute questionnaire[7] with an introductory section explaining its purpose and motivation. Participants provided anonymous responses, and data confidentiality was maintained through coded identification. The study primarily explored learner preferences for receiving course explanations based on either extrinsic factors, influenced by the community, or intrinsic factors, such as interest, historical knowledge, learning outcomes, and learning objectives. In

[7] https://u.garr.it/ZJHfM

Table 3.1: Example explanation from meta-paths in our ontology. For each meta-path described in Figure 3.3, we provide an example explanation that can be derived. These textual explanations are adopted during the ontology evaluation process.

Meta-Path ID	Derived Example Textual Explanation
ID1	Learners similar to you, who attended the Programming course, have found it to be a valuable complement.
ID2	It covers the topic of object-oriented programming, aligning with content from the Programming course you attended.
ID3	It has been successfully completed by another learner who shares your interest in object-oriented programming.
ID4	It is offered by the same institution, the University of London, from which you previously attended the Programming course.
ID5	It is taught by the same instructor, Mark Bowle, with whom you previously attended the Programming course.
ID6	Your background in object-oriented programming, relevant to Java language, serves as a suitable prerequisite for this course.
ID7	It can amplify your understanding about computational complexity, going beyond its subtopic of program organization and design covered in the Programming course you attended.
ID8	It can advance your skills in basic notions of computational complexity from the analytical level (e.g., identifying algorithmic time complexity) gained in the Programming course to the evaluation level (e.g., comparing and selecting optimal algorithms).
ID9	Another learner from your institution, the University of London, has successfully completed the course.
ID10	Learners who share your interest in the Programming course have also expressed interest in this course.
ID11	Learners who share your interest in the Programming course have also expressed interest in this course.
ID12	To enhance expertise for the software engineer role. This complements the content of the Programming course on the same role.

a first part, participants ranked these explanations in order of effectiveness, depending on their motivation, whether studying for pleasure or earning credits. The second part of the questionnaire focused on evaluating explanations derived from each meta-path individually, considering seven criteria proposed by [99]: decision speed, motivation, satisfaction, correction ease, transparency, and confidence boost. Specifically, to facilitate this process, six scenarios were presented, each involving a pairwise comparison of two explanations. Participants used a 5-point Likert scale to express their agreement with each explanation. The scenarios covered various factors like community influence, interest, historical knowledge, instructor influence, institution influence, learning outcomes, learning objectives, and prerequisite knowledge.

Questionnaire Delivering. We recruited 100 participants through Prolific [8], ensuring a balanced gender distribution. Participants were diverse, representing various countries. The cohort was devised to include an equal number of university learners (50) and lifelong learners (50). Further segmentation was carried out within each group, resulting in two subgroups of 25 participants each. This division was instilled into the questionnaire by randomly assigned learners into two different scenarios: learning-for-pleasure and learning-for-credits. In the learning-for-pleasure scenario, participants were motivated by the joy of learning a new topic of interest. In the learning-for-credits scenario, participants aimed to earn credits by enrolling in a course. Except for this, the questionnaire content was the same across all participants.

[8] www.prolific.com

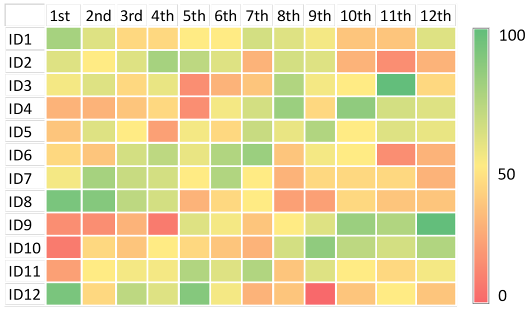

Fig. 3.4: Heatmap illustrating ranked explanations. The y-axis displays unique explanations identified by their respective IDs, while the x-axis indicates ranking positions from first to twelfth. Each cell's color corresponds to the main variable's value within its range, from red (0%) to green (100%).

3.3 Experimental Results

For a learner-centered ontology design, we target two research questions:

- **RQ1**: To what extent are meta-paths from our ontology and their corresponding textual explanations effective in aiding course selection?
- **RQ2**: How do learners perceive different types of meta-paths and the explanations derived from our ontology across explanation criteria?

3.3.1 RQ1: Perceived Meta-Paths Importance

In our initial analysis, we explored the effectiveness of the example meta-paths and their associated textual explanations in assisting learners with course selection. To this end, we considered the answers given by learners to the ranking task included in the questionnaire, considering their perceived importance of the explanations in the presented scenario. The heatmap presented in Figure 3.4 illustrates the outcomes of this analysis across the entire study population. The y-axis represents different explanations, each assigned a unique ID (Table 3.1). On the x-axis, positions from first to twelfth are marked, indicating the ranking of these explanations. Each column in the heatmap is color-coded to represent the percentage of learners who ranked the explanations in that particular position. The color spectrum ranges from red (indicating 0%) to green (indicating 100%).

In the first position, the explanation *"to enhance your expertise for the … role. This complements the content … the course related to the same role"* with 19% highlights the importance learners attribute to clear and defined objectives. This may indicate a desire for clarity and direction in the learning path, suggesting learners appreciate explanations aligned with potential occupations.

The second position, occupied by the explanation *"It can advance ... notions of computational complexity from the analytical level (e.g., ...) gained in the Programming course to the evaluation level (e.g., ...)"* with 18%, underscores learners' notable interest in attaining tangible educational outcomes. This inclination may be interpreted as a quest for explicit feedback and assessments of their progress, signifying an emphasis on measurable aspects of learning.

The third position, claimed by the explanation *"It can amplify your understanding about..., going beyond ... covered in the Programming course you attended;"* with 15%, suggests that learners prioritize expanding their knowledge into new areas while building upon a strong foundation. From a technical standpoint, this preference shows the significance of incorporating relationships such as *(subject, is_parent_of, subject)* and *(subject, requires, subject)* into the ontology, along with relationships related to the Bloom's Taxonomy. This favors *content-based* over *community-based* explanations.

The analysis of answers for formal and informal learning reveals distinct priorities among university learners in different contexts. In formal settings, where learning for pleasure is emphasized, preferences lean towards learning objectives. Conversely, when the aim is to attain educational credits, interests take precedence (16%). This adaptability implies that learners shift their focus based on the learning scenario. However, the importance of tailoring explanations to learners' professional profiles, particularly in early university experiences, highlights the potential role of an explainable recommendation as a guidance throughout various stages of their academic journey. In a professional context, when focused on formal learning, the second position is claimed by an extrinsic factor, with 20% favoring the explanation *"It has been successfully completed by another learner..."*. This suggests that, driven by external motivation, workers prioritize swift certification to meet professional requirements, making *community-based* explanations more effective for course selection in this scenario.

Overall, the ontology facilitates effective course selection for learners by offering impactful meta-paths and textual explanations, driven by preferences for clear objectives, measurable outcomes, and a focus on expanding knowledge. This adaptability extends to various learning contexts, while a notable inclination towards community-based explanations emerges, especially in professional formal learning scenarios.

3.3.2 RQ2: Multi-Criteria Meta-Path Assessment

We conducted an investigation to identify the preferred explanations and their perceived benefits through a pair-wise comparison focused on effectiveness, decision speed, motivation, satisfaction, correction ease, transparency, and confidence boost. In our statistical analysis of Likert scale responses, we computed weighted mean scores ($\bar{X}$), standard deviations (σ), and frequencies (f). The Likert scales were weighted with averages calculated from *strongly disagree* (1) to *strongly agree* (5), using a length of 0.80 (4/5) to represent the four distances between numbers on the five-point Likert scale. Table 3.2

Table 3.2: Pair-wise comparison of explanations and their perceived benefits. We conducted a statistical analysis of Likert scale responses, presenting results with weighted mean scores ($\bar{X}$), standard deviations (σ), and frequencies (f).

	UC1						**UC2**					
	$\bar{X}$		σ		Attitude		$\bar{X}$		σ		Attitude	
	ID01	ID02	ID01	ID02	ID01	ID02	ID03	ID06	ID03	ID06	ID03	ID06
Effectiveness	3,16	3,20	0,94	0,88	A	A	3,02	3,23	0,97	0,92	A	A
Decision Speed	2,93	3,23	0,97	0,85	A	A	2,89	3,27	1,02	0,85	A	A
Motivation	3,06	3,25	0,97	0,98	A	A	2,87	3,36	0,97	0,84	A	A
Satisfaction	3,09	3,21	1,00	1,03	A	A	2,88	3,24	1,00	0,78	A	A
Correction Ease	3,00	3,05	1,04	0,96	A	A	2,88	3,22	1,09	1,00	A	A
Transparency	3,21	3,45	1,00	0,98	A	**SA**	3,13	3,51	0,97	0,81	A	**SA**
Confidence Boost	3,14	3,24	1,00	1,00	A	A	3,01	3,36	0,95	0,84	A	A
	UC3						**UC4**					
	$\bar{X}$		σ		Attitude		$\bar{X}$		σ		Attitude	
	ID04	ID05	ID04	ID05	ID04	ID05	ID07	ID08	ID07	ID08	ID07	ID08
Effectiveness	3,19	2,91	0,95	1,02	A	A	3,11	3,45	0,91	0,81	A	**SA**
Decision Speed	3,22	2,84	0,91	1,01	A	A	3,14	3,33	0,91	0,90	A	A
Motivation	3,11	2,90	0,94	1,07	A	A	3,05	3,40	0,87	0,84	A	A
Satisfaction	3,14	2,89	1,02	1,06	A	A	3,20	3,42	0,92	0,88	A	**SA**
Correction Ease	3,17	2,87	1,11	1,08	A	A	3,21	3,45	0,88	0,91	A	**SA**
Transparency	3,17	3,16	1,03	1,07	A	A	3,32	3,51	0,87	0,88	A	**SA**
Confidence Boost	3,27	2,96	0,96	0,98	A	A	3,24	3,39	0,87	0,89	A	A
	UC5						**UC6**					
	$\bar{X}$		σ		Attitude		$\bar{X}$		σ		Attitude	
	ID12	ID11	ID12	ID11	ID12	ID11	ID09	ID10	ID09	ID10	ID09	ID10
Effectiveness	3,36	3,02	0,85	0,94	A	A	2,69	2,84	1,14	0,97	A	A
Decision Speed	3,35	3,11	0,82	0,96	A	A	2,61	2,86	1,18	1,03	A	A
Motivation	3,40	3,11	0,85	0,86	**SA**	A	2,74	2,91	1,18	1,06	A	A
Satisfaction	3,24	3,12	0,75	0,92	A	A	2,72	2,79	1,15	1,00	A	A
Correction Ease	3,27	3,12	0,85	0,94	A	A	2,70	2,87	1,16	1,05	A	A
Transparency	3,47	3,30	0,73	0,77	**SA**	A	2,91	3,02	1,12	1,08	A	A
Confidence Boost	3,35	3,07	0,77	0,81	A	A	2,77	2,87	1,18	1,05	A	A

summarizes the results.

UC1 Learner connection (ID01) - Content connection (ID02). In this scenario, the first recommendation (ID01) suggests a course based on the enrollment of a user similar to the learner, while the second advises the course because it delves into a topic of object-oriented programming previously encountered in another course (ID02). The $\bar{X}$ and σ highlight that, for each criterion, the predominant response is *agree*, except for transparency, where the $\bar{X}$ (3.45) for the second recommendation categorizes it as *strongly agree*. Additionally, the latter obtained a higher score in terms of motivation (45 occurrences). From the results, the intrinsic factor motivates more strongly compared to the extrinsic factor. The learners' preference for recommendations centered around intrinsic aspects emphasizes the importance of personalization and individual relevance.

UC2 Common interest (ID03) - Knowledge prerequisite (ID06). In this scenario, the *subject* entity, represented by *object-oriented programming*, is compared in two distinct scenarios: one based on shared interest (ID03) and the other on historical knowledge (`ID06`). The second explanation is perceived as more beneficial across all evaluated criteria, especially in terms of transparency ($\bar{X} = 3.51$) and motivation (47 occurrences). These findings highlight the significance for learners to follow a learning path aimed at reinforcing previously acquired knowledge rather than enrolling in a course merely because someone with similar interests has done so. This emphasizes that the decision-making process for choosing a learning path should be motivated by the desire to enhance existing expertise, rather than influenced by behaviors/choices of others.

UC3 Common instructor (`ID04`) - Common Institution (`ID05`). In this scenario, we are assessing two distinct types of explanations, one based on the similarity of the instructor (`ID04`) and the other on the institution (`ID05`). The first explanation is favored, surpassing the second in terms of effectiveness (46 occurrences) and correction ease (45 occurrences). However, it is perceived moderately in relation to motivation. In contrast, the second explanation stands out for confidence boost, with *strongly agree* occurring 39 times. While both instructor and institution are deemed important, the former was more impactful in understanding their appreciation for the process.

UC4 Subject knowledge (`ID07`) - Cognitive knowledge (`ID08`). In this comparison, we assess the impact of consolidating prior knowledge (`ID07`) against the anticipated outcomes of the learning process expressed through Bloom's Taxonomy (`ID08`). The first explanation is perceived as moderately useful across all evaluated criteria. In contrast, the second explanation receives a rating of *strongly agree* in terms of effectiveness ($\sigma = 3.45$), satisfaction ($\sigma = 3.42$), correction ease ($\sigma = 3.45$), and transparency ($\sigma = 3.51$). These findings suggest that learners may value explanations that offer a tangible understanding of learning outcomes compared to more general objectives.

UC5 Occupation objective (`ID12`) - Knowledge prerequisite (`ID11`). In this context, we compare the overall instructional intent, which focuses on specific topics for the preparation of a software engineer (`ID12`), with the depth of prior knowledge regarded as prerequisites for the recommended course (`ID11`). The first explanation receives, on average, a rating of *strongly agree* in terms of motivation ($\bar{X} = 3.4$) and transparency ($\sigma = 3.47$), while in other instances, it is perceived as *agree*, particularly for satisfaction, with 47 occurrences. In the second explanation, it is classified as *agree*, and noteworthy scores are observed for satisfaction (42 occurrences) and confidence boost (47 occurrences). Both explanations are deemed significant based on the considered criteria.

UC6 Common institution (`ID09`) - Common learner (`ID10`). In this context, we are evaluating two different types of explanations based on either the similarity of the institution (`ID09`) or the learner (`ID10`). Both explanations received an overall rating of *agree*, implying a general consensus among respondents. However, despite this agreement, the evaluations show a balanced distribution between negative and positive responses for both explanations. This suggests that neither type of explanation has a pronounced impact on the explanation properties we assessed. Moreover, results indicate that these explanations have limited effectiveness in motivating users to enroll in recommended courses.

Overall, learners perceive different meta-paths and explanations from the ontology in a positive manner for all the explanation criteria. While certain types, such as those based on instructors and tangible learning outcomes, show higher impact, others, like institution and common learner interest, exhibit more balanced effectiveness with limited impact on motivation.

3.4 Discussion and Implications

This chapter addressed the growing importance of ontologies for explainable educational recommendation. Despite their growing importance, existing ontologies for this purpose face challenges related to sufficiency and coherence, which hinder their effective use in explaining recommendations. To overcome these limitations, we proposed LOXER, a learner-centered ontology designed for both formal and non-formal education. Our approach saw the construction of a comprehensive, interdisciplinary, and flexible pipeline to transform and expand traditional educational recommendation datasets. Through a qualitative evaluation, we showed the practicality of LOXER in enabling meaningful explanations for recommendations across educational scopes and learner populations. In particular, learners may face cognitive and technical barriers when interacting directly with structured knowledge graphs or static explanation outputs. To address this limitation, the next chapter shifts the focus from knowledge representation to interaction design. Chapter 4 investigates how conversational interfaces driven by LLMs can mediate between learners and educational knowledge graphs, enabling interactive exploration, sensemaking, and explainable dialogue over these structures.

Chapter 4

Methods for Conversational Knowledge Graphs Exploration

Research Highlights

- Designed *EDGE*, a conversational interface educational KGs exploration.
- Insight extraction from KGs through natural language queries.
- Demonstration of EDGE capabilities across educational scenarios.

4.1 Introduction

As education adopts digital platforms, the vast amount of information from various sources, such as learning management systems and learning object repositories, presents challenges in navigation and elaboration. Traditional interfaces involve a steep learning curve, limited user accessibility, and lack of flexibility. Language models alone cannot address these issues as they do not have access to structured information specific to the educational organization. In this paper, we propose EDGE (EDucational knowledge Graph Explorer), a natural language interface that uses knowledge graphs to organize educational information. EDGE translates natural language requests into queries and converts the results back into natural language responses.

With the adoption of digital tools and platforms, educational activities are introducing an abundance of information within organizations [100]. Navigating this wealth of information about both learning content and behavior presents a significant challenge for learners, educators, and managers [101, 102, 103, 104]. For instance, when the former visit the learning management system with a specific goal in mind [105], the interface often necessitates navigating through multiple menus and options. This process can be overwhelming and time-consuming, as it is constrained by the predefined functional paths [106]. Similarly, educators encounter difficulties in assessing student progress or accessing summarized course data, resorting to analyzing log tables or predefined dashboard plots [107, 108, 109]. Managers face similar challenges when extracting insights

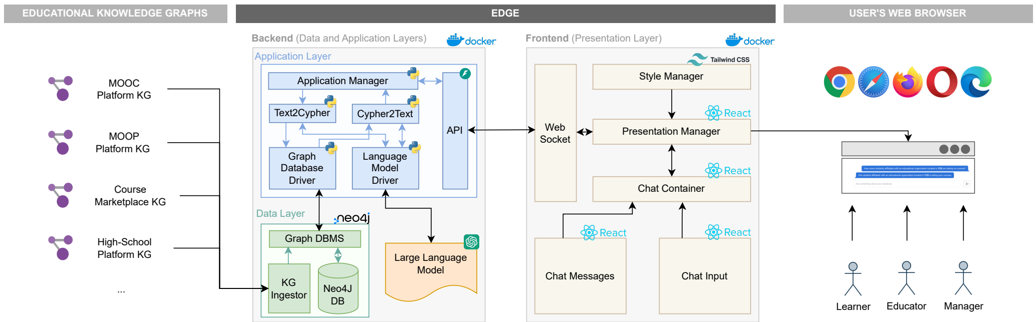

Fig. 4.1: Architecture. EDGE leverages knowledge graphs of a given educational organization. It is based on a backend-frontend setup, with a three-tier layering including the presentation layer in the frontend, and application and data layers in the backend.

across courses, constrained by the limitations of graphic interfaces [110, 111]. As such, there is a need for more natural ways of interacting with the knowledge base of the organization.

Conversational interfaces have emerged as a promising solution to address the challenges faced in navigating educational information, offering a natural way of implementing personalization [112], tutoring [113], delivery, assessment [114], and support [115]. However, existing conversational interfaces often have a narrow focus and lack the capability to provide contextually-relevant and flexible access to the knowledge base of the organization [116]. While conversational interfaces driven by language models (LMs) have gained prominence in education [117], LMs still encounter difficulties in accessing organization-specific information and modeling it. Concurrently, knowledge graphs (KGs) have garnered attention as a method for modeling educational information [118, 81]. Their integration with conversational interfaces powered by LMs has been investigated in several fields [119, 120, 121] but remains unexplored in education, being a missed opportunity for exploring learning data in a more natural, accurate, and transparent manner. In contrast to conventional tools such as relational databases, which often necessitate manual and intricate querying based on string matching or basic operations, and vector search, which demands prior computation of the embeddings, the integration of LMs with KGs not only offers a more intuitive and sophisticated approach of retrieving information but also capitalizes on the inherent, nuanced semantic structure for modeling data that KGs provide. KGs represent knowledge through a semantic framework, which relates to natural language, an area in which contemporary LMs excel.

To bridge the gap between KGs and LMs and enable experimentation in the educational field, our contributions in this chapter are three-fold:

- A *KG ontology mapping framework* able to turn educational datasets into a coherent semantic structure according to prior work, for natural language exploration.
- A *LM-powered conversational interface*, named EDGE, which enables stakeholders

to extract insights from educational KGs through natural language queries.

- A demonstration of EDGE's capabilities using knowledge graphs extracted through our framework across four different educational scenarios.

The rest of this chapter is structured as follows: Section 4.2 outlines the methodology that offers a user-friendly natural language interface, allowing to seamlessly ask for information about the KG of the educational organization ecosystem. Section 4.3 assess EDGE's capabilities via distinct scenarios, extracting and digesting knowledge graphs from four public educational datasets, and provide preliminary example interactions. Lastly, Section 4.4 concludes the chapter with a discussion and implications.

4.2 Methodology

EDGE follows a backend-frontend setup, with a three-tier layering including the *Presentation* layer in the frontend, and the *Application* and *Data* layers in the backend (Figure 4.1). The first layer provides a chat-based interface. The second layer handles natural language requests from the *Presentation* layer, converts them into queries for the *Data* layer's database, and generates natural language responses for the *Presentation* layer from query results. Lastly, the third layer stores educational KGs to be accessed by the *Application* layer.

Presentation Layer. EDGE implements a web-based conversational interface[1] using React[2], a framework renowned for responsive and interactive user interfaces. The *Presentation Manager* module is responsible for initializing the *Chat Container* component, which encompasses both the *Chat Messages* component to preserve the conversation history and the *Chat Input* component to allow sending natural language queries. Example natural language queries are provided to users to expedite their comprehension. Moreover, the integration of the React useWebSocket[3] library ensures a persistent connection with the *Application* layer. Using Tailwind CSS[4] within the *Style Manager* module enriches cross-device aesthetics.

Application Layer. We implemented the system's core logic using Python, structuring it into a series of steps as depicted in Figure 4.2. When a user submits a natural language query through the API from the *Presentation* layer, the query is forwarded to the LM through the *Text2Cypher* component[5]. The LM, in turn, generates a query in Cypher, the graph database language, tailored to the educational knowledge graph ontology. We experimented with both zero-shot and few-shot educational-specific techniques in order to tailor the approach to our target domain. Subsequently, this Cypher query is executed

[1] We leave solutions like **OpenWebUI** (`https://docs.openwebui.com/`) for future work.

[2] **React**: `https://react.dev/`

[3] **React useWebSocket**: `https://www.npmjs.com/package/react-use-websocket`

[4] **TailwindCSS**: `https://tailwindcss.com/`

[5] **NaLLM**: `https://github.com/neo4j/NaLLM`

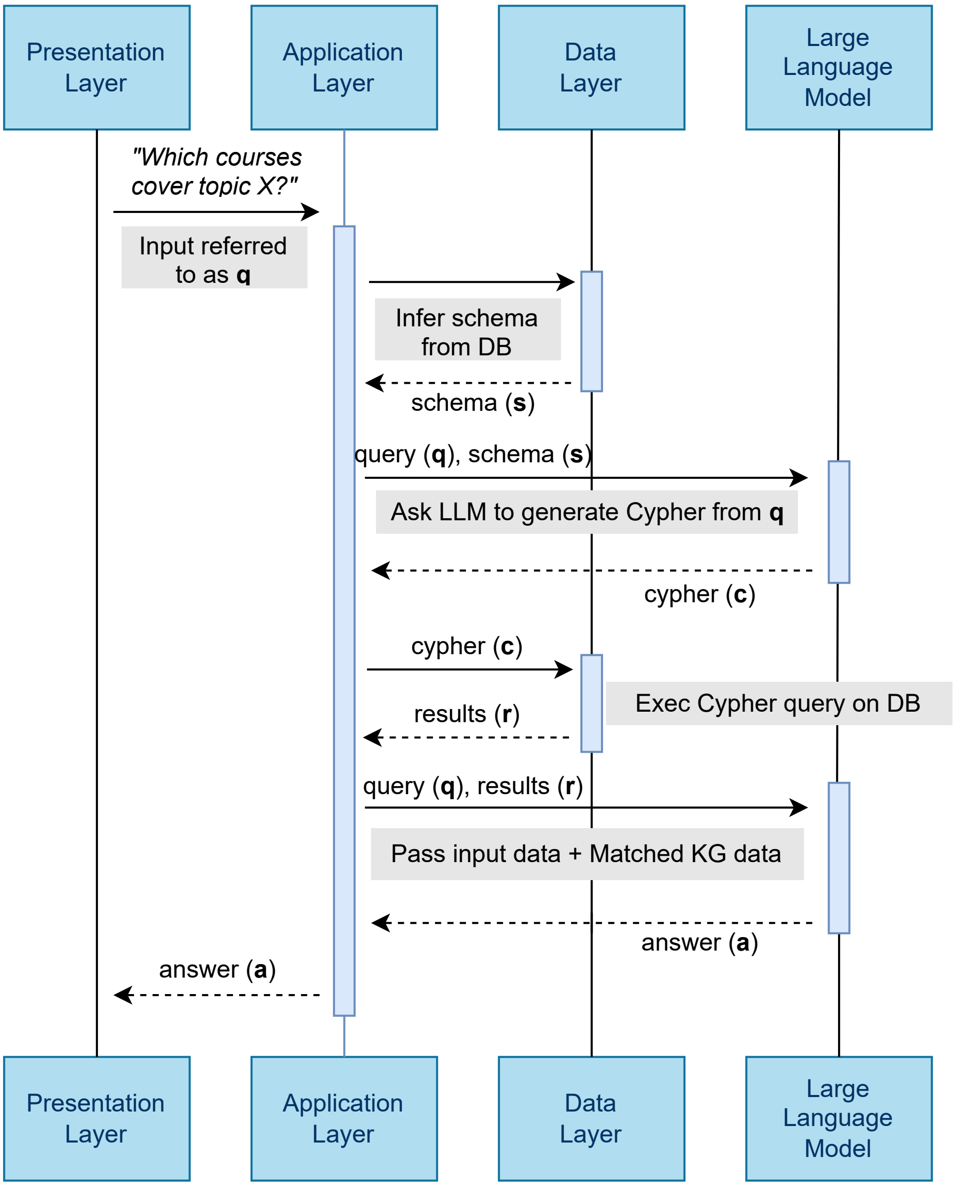

Fig. 4.2: Workflow. The user submits natural language. The LM converts it into a Cypher query based on the educational ontology. This query is run on the graph database. The LM generates a natural language response from the query result.

on the database stored in the *Data* layer, and the results are fetched using the *Graph Database Driver* component. Finally, the *Application Manager* component formulates a response to the user's query in natural language, arranging insights from the query results by means of the LM via the *Cypher2Text* component. In terms of implementation, we relied on the FastAPI[6] library to establish API endpoints, interfaced with OpenAI's GPT 3.5 Turbo LM (with 1000 as maximum tokens and 0.0 as the temperature) through the OpenAI[7] library within the *Language Model Driver*, and utilized the Neo4j[8] library for the *Graph Database Driver* to facilitate seamless communication with the graph database. Other open-sourced LMs and graph databases can be easily integrated.

Data Layer. The *Data* layer serves as the repository for pertinent organizational data, structured in the form of a knowledge graph. The latter is managed by the Neo4j[9] *Graph Database Management System* (DBMS) and stored within a Neo4j database. Neo4j was chosen for its robust scalability and optimized performance proven on graph-based data. Within this layer, the *KG Ingestor* facilitates the loading of a knowledge graph from a Turtle[10] file into a Neo4j database. This process involves several steps: initially, establishing a connection to Neo4j via API, then executing data import and translation queries. This ensures a seamless transition of data from the RDF[11] model to the property-graph model utilized by the Neo4j DBMS. By leveraging Neo4j's capabilities and the functionality provided by the *KG Ingestor* component, the *Data* layer ensures efficient management and utilization of the knowledge graphs.

4.3 Experimental Demonstration

The demonstration covers the core functionalities of EDGE through four example educational scenarios. To validate them, we conducted a qualitative analysis with local students and educators on response accuracy, completeness, and understandability.

In a first step, we introduce attendees with the process of preparing knowledge graphs from conventional educational data and uploading them into EDGE through its *Data* layer components. To this end, we start by describing the educational ontology at the basis of the use cases considered during the demonstration. Coming from extensive validation conducted by [122], the ontology[12] was created through a collaborative approach of conceptualization, development, and validation, by engaging learners, educators, and learning analytics researchers. Subsequently, we delineate the four representative educational environments under consideration, along with their corresponding knowledge graphs, namely a Massive Open Online Course (MOOC) platform [94], a Massive Open Online Practice (MOOP) platform [95], a course marketplace [93], and a high-school platform [123]. We emphasize the distinctive entities and relationships each of the environments enables in the considered ontology. We finally showcase how the encoded

[6] **FastAPI**: `https://fastapi.tiangolo.com/`
[7] **OpenAI Python**: `https://github.com/openai/openai-python`
[8] **Neo4j Python**: `https://neo4j.com/docs/api/python-driver/current/`
[9] **Neo4j**: `https://neo4j.com/`
[10] **Turtle**: `https://www.w3.org/TR/turtle/`
[11] **RDF**: `https://www.w3.org/RDF/`
[12] **KG Ontology**: `https://u.garr.it/ATBgW`

files can be transformed into the format supported by Neo4j, and the steps to perform into the Neo4J DBMS to upload the converted knowledge graph.

4.3.1 UC1: MOOC Platform (MOOCCube)

In a subsequent step, we delve into XuetangX[13], an online learning platform spearheaded by Tsinghua University, with a mission to provide learners with the high quality courses from elite Chinese universities. In this case, we demonstrate EDGE's capabilities by leveraging a knowledge graph created from the Xuetang-collected MOOCCube dataset [94]. As an icebreaker, we suggest a series of example queries to the attendees. For instance, a learner seeking to discover a course might submit the query "Which courses cover the subject forecasting model?". Traditionally, this would involve performing textual searches, applying filtering parameters, browsing through paginated results, and possibly navigating between different pages. With EDGE, the learner can send the request through the interface and receive the result shown in Figure 4.3a. As another example, an educator curious about the diversity of educational backgrounds might inquiry "Which subjects are preferred the most among learners enrolled in my courses?". Without EDGE, the educator would need to manually analyze enrollment curricula. With EDGE, they input the query and receive a breakdown detailing the most frequent subjects. Similarly, managers might need to identify the courses with the least amount of learners enrolled. They might submit the query "Are there any courses with no learners enrolled? If so, please provide a list of course IDs". Without EDGE, the manager might need to go through all the course enrollments lists or ask to do this to the platform admin. With EDGE, the response might be "Yes, there are X courses with no learners enrolled, including ... ".

4.3.2 UC2: MOOP Platform (MOOPer)

In a third step, we introduce the attendees with Educoder[14], a platform for computer science subjects in China that provides students with a practice-oriented approach to subjects. A knowledge graph created from the Educoder-collected MOOPer dataset [95] is adopted to demonstrate EDGE's capabilities. As an example, a learner seeking information about the subjects of a course can query EDGE with a request like "What subjects does course X cover?", bypassing the need to navigate via multiple pages (Figure 4.3b). Similarly, an educator's inquiry about the number of learners enrolled in their course from a given geographic area can be addressed via EDGE, with no need to manipulate enrollment files (Figure 4.3c). Lastly, the manager's need to identify courses requiring improvements can be met by querying EDGE with "Which course has received the lowest review score on average?", as an example.

[13] **XuetangX**: `https://www.xuetangx.com/global/`

[14] **Educoder**: `https://www.educoder.net/`

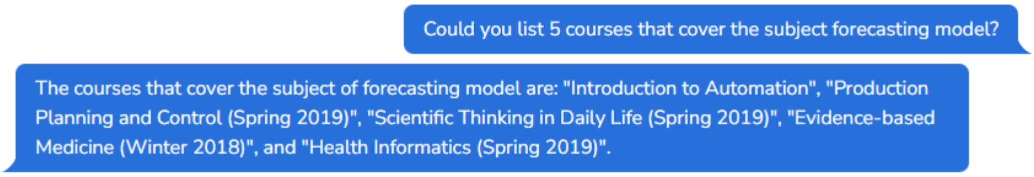

(a) Courses by topic retrieval on top of MOOCCube's KG.

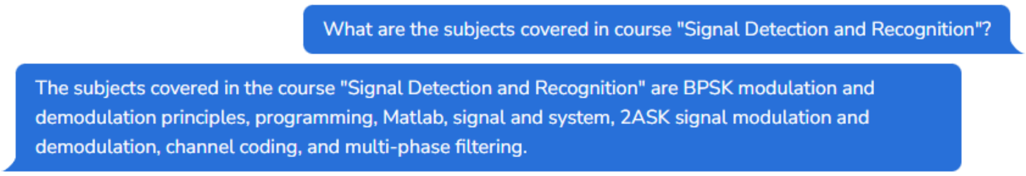

(b) Topics per course retrieval on top of MOOPer's KG.

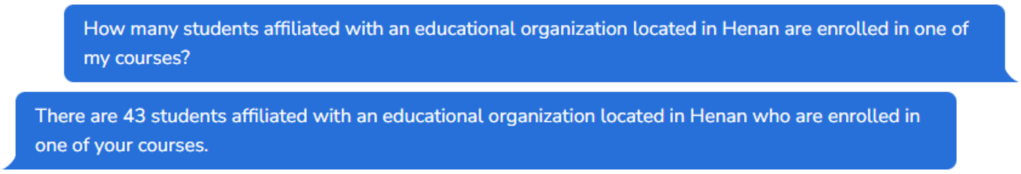

(c) Student statistics on courses on top of MOOPer's KG.

Fig. 4.3: Example Responses. Users initiate queries in natural language, and the system responds based on the retrieved data. Each response includes a toggle (available in the demo video) with the generated Cypher query, for interpretability.

4.3.3 UC3: Course Marketplace (COCO)

We then move to a course marketplace scenario. Udemy[15] is a prominent online learning platform that hosts thousands of courses. We demonstrated EDGE's capabilities under this environment by leveraging a knowledge graph extracted from the COCO dataset [93]. Learners seeking courses covering specific topics and languages can leverage EDGE to streamline their search process. Traditionally, they would need to manually check course descriptions or consult academic advisor. Search functions might be sometimes available but do not allow to perform semantically-complex searches. With EDGE, a query like "Are there any courses covering math and science and delivered in Japanese?" yields a list of relevant courses. Educators also benefit from EDGE's capabilities, e.g., in computing course statistics. By inputting queries like "On average, how many resources are included in a course? Separate them by type", educators get breakdowns for course preparation reasoning. Managers can find EDGE valuable, such as for identifying learners with heavy course workloads. Queries like "Which learners are enrolled in more than X courses?" are supported to streamline monitoring.

[15] **Udemy**: `https://www.udemy.com/`

4.3.4 UC4: High-School Platform (EduKG)

Finally, we cover the high-school scenario, specifically focusing on the transition to the university. To this end, we leverage the EduKG [98] knowledge graph created from survey answers given by learners from French high schools and universities. As an example, high-school learners exploring university opportunities can utilize EDGE to inquire about Engineering Technological Bachelor's courses, asking, "Could you give me a list of Engineering Technological Bachelor's courses?" In response, EDGE can provide a list of relevant courses such as Energy, Industrial Logistics, Physics, Electrical Engineering, Mechanical Engineering, Civil Engineering, Industrial Engineering, Materials Engineering, Packaging, and so on. Meanwhile, educators in a graduate program seeking insights into learner engagement with specific topics can query EDGE with, "How many learners enrolled in 'Bachelor's Degree - Psychology' have dealt with the subject 'digital skills'?" In this scenario, EDGE can reports that 37 learners enrolled in 'Bachelor's Degree - Psychology' have engaged with the subject 'digital skills'. Additionally, managers interested in tracking enrollment in dual bachelor's degree programs can input the query "How many learners are enrolled in Dual Bachelor's programs?" to receive the response indicating that there are 47 learners who interacted with this type of course.

4.4 Discussion and Implications

This chapter presented EDGE, a conversational interface designed to support intuitive and explainable exploration of educational knowledge graphs. By combining the structural properties of knowledge graphs with the natural language understanding capabilities of large language models, the chapter demonstrated how conversational interaction can lower the barrier for learners to access, interpret, and navigate complex educational knowledge structures. The proposed system illustrates how learners can engage in exploratory dialogue to retrieve information, refine queries, and obtain explanations grounded in structured educational data.

Beyond the interface itself, the chapter introduced a pipeline for mapping heterogeneous datasets into an educational knowledge graph ontology, enabling their integration within the conversational interface. Through illustrative demonstrations, the chapter showed how conversational mediation can streamline information-seeking tasks and support learner-driven sensemaking, complementing the ontology-based explainability mechanisms introduced in Chapter 3.

At the same time, the chapter highlights an important limitation of purely interaction-driven explainability. While conversational interfaces can effectively mediate between users and structured knowledge, the quality, reliability, and faithfulness of conversational explanations ultimately depend on the underlying reasoning mechanisms used to traverse and interpret the knowledge graph. In particular, free-form natural language generation alone does not guarantee that explanations remain grounded in valid semantic paths or that recommendations are supported by coherent reasoning structures. While the EDGE system is evaluated through multiple illustrative use cases, the current assessment remains primarily qualitative. A formal quantitative evaluation, including metrics such as response accuracy, query success rate, and user satisfaction, is left as future work. Conducting a controlled user study would provide deeper insights into the system's effectiveness and usability.

This observation motivates the need for explicit, controllable reasoning methods that can systematically connect learners to educational resources through interpretable paths in the knowledge graph. Such methods are essential not only to support recommendation generation, but also to ensure that explanations surfaced through conversational interfaces remain faithful to the underlying knowledge representation.

Accordingly, the next chapter shifts the focus from interaction design to algorithmic reasoning over educational knowledge graphs. Chapter 5 introduces methods for path-based recommendation, investigating how structured reasoning paths can be generated, constrained, and exploited to produce explainable recommendations. These methods provide the computational foundation that complements the conversational interface presented in this chapter, enabling personalization that is interactive and grounded.

Chapter 5

Methods for Explainable Path-based Recommendation

Research Highlights

- *KG* provide a shared structure for recommendation and explanation.
- Different path-based models exhibit complementary strengths.
- Recommendation performance is influenced by *data sparsity*.

5.1 Introduction

Digital advances have increased the need for intelligent systems to guide learners through large and complex educational ecosystems. Learners are often required to select learning resources from an overwhelming set of alternatives, while ensuring coherence with prior knowledge and learning objectives. In such settings, educational recommender systems play a critical role by supporting decision-making and personalization. However, the large-scale deployment of recommendation systems in education raises important concerns related to transparency and trust.

Most existing educational recommender systems uses collaborative filtering or content-based techniques, which model learner preferences through latent representations or similarity measures. While these approaches have demonstrated strong predictive performance, they typically provide limited insight into the reasoning process. As recommender systems grow in complexity, transparency becomes as crucial as effectiveness, making explanations a core part of modern systems. This dual requirement has positioned knowledge graphs (KGs) as the preferred representation of entities, along with their relationships, to learn preference patterns. Recommendation models trained over KGs can generate paths having user and item entities at their endpoints. While the item at the endpoint is the one being recommended, intermediate elements serve to construct a textual explanation, through template- or generation-based techniques.To generate informative paths over KGs, three main reasoning approaches have emerged so far. Reinforcement learning-based reasoning, such as PGPR [48], optimizes agents to discover paths within KGs that connect users to relevant, previously unseen items. Neuro-symbolic reasoning, such as CAFE [49], combines symbolic logic with neural networks to capture relationships in the KG, uncovering paths aligned with high-level semantic patterns. Meanwhile, generative-based reasoning, such as PEARLM [83], leverages the structural similarity between KG paths and natural language to generate sequences that alternate between entities and relationships within the KG, beginning with a user entity.

While these methods show promise, their evaluation has largely focused on domains outside of education, such as entertainment (e.g., `Movielens` and `LastFM` in [124, 125]), e-commerce (e.g., `Amazon` in [48, 49]), and news media (e.g., `MIND` in [126]), even in recent reproducibility research concerning KG reasoning [125]. Unlike these domains, education has unique characteristics, such as long intervals between course enrollments, and educational decision-making involve distinctive rationales, such as course prerequisites, typically absent in other domains [127]. These factors introduce uncertainty as to whether such methods generalize to education. Research on the above reasoning methods under course recommendation remains in its infancy. Reinforcement learning variants have included only minor changes, such as alternative training path sampling policies or hierarchical learning [128, 129, 130, 131, 132]. Evidence for neuro-symbolic reasoning is limited to studies focused solely on utility [133]. Generative modeling remains fully unexplored. Moreover, evaluations in the above studies are based on limited datasets (often just one), heterogeneous baselines (primarily traditional collaborative filtering methods or the original method to which the adaptations were applied), and mere utility-based metrics. This fragmentation leaves it unclear which methods, and in which conditions, are best suited for course recommendation.

In this chapter, we investigate methods for path-based explainable recommendation in education, with the goal of understanding their suitability and limitations. Building on the learner-centered ontology introduced in Chapter 3, we construct educational knowledge graphs that capture learners, educational resources, and learning relations under a unified semantic schema. This controlled setting allows us to systematically analyze different path-based recommendation paradigms and compare their behavior under consistent experimental conditions.

The approaches studied in this chapter generate recommendations by reasoning over knowledge graph paths, thereby linking recommendation utility and explainability through a reasoning process. The contribution of this chapter is three-fold:

- Systematically study state-of-the-art recommendation methods in the educational domain, analyzing their ability to generalize from other application domains to personalized education scenarios.
- Translation of educational datasets into KG structures under a unified educational knowledge graph and ontology.
- Evaluation of representative methods from major families of KG reasoning and

path language modeling.

The remainder of this chapter is structured as follows. Section 5.2 outlines the Research Methodology covering datasets, KG structures conversion and processing for recommendation. Section 5.3 discusses the empirical results, analyzing the trade-offs between recommendation utility and beyond-utility goals across different methods. Finally, Section 5.4 concludes with the discussion and implications.

5.2 Methodology

Our methodology for auditing KG reasoning methods involves five main steps. First, we selected educational datasets (1), convert them to KG structures (2), and pre-process the latter for recommendation (3). Next, we optimize reasoning models on these paths (4) and evaluate them using utility, beyond-utility, and explainability metrics, against selected baselines (5).

5.2.1 Educational Datasets Selection

Course recommendation has been widely studied, but many datasets are private, and public ones often lack KGs or suitable structures for recommendation tasks [134, 35, 135]. To address this gap, we decided to select four public educational datasets for conversion into their corresponding KG structure tailored to the course recommendation task: COCO [93], `MOOPer` [95], `MOOCube` [136], PEEK [137]. These datasets were chosen for their broad coverage of users and courses to provide enough data for learning, their richness in features to support KG creation, and the diversity of context.

The COCO dataset (data: obtained from original authors by e-mail) captures data from a marketplace where non-accredited experts offer diverse courses. Unlike academic platforms, the catalog is larger and varied, attracting a broad user base but resulting in interaction sparsity. Users typically enroll in a few courses for on-demand learning, with many courses having limited enrollments due to niche targeting or competition. Hence, the dataset includes 529, 857 ratings, 43, 021 learners, and 24, 321 courses, with course metadata (title, descriptions, categories, audiences, prerequisites, goals, lectures) and instructor affiliations.

The `MOOCCube` dataset (data: `http://moocdata.cn/data/MOOCCube`) provides data from a platform offering a structured, academic selection of courses across multiple disciplines. It attracts a learner base more focused on formal education. Its catalog diversity lies between that of a general marketplace and a specialized practice platform, though it is comparable in scale to the latter. With the largest user base among four, the dataset includes 706 courses, 199, 199 learners, and 682, 753 ratings, along with metadata on course titles, video titles, concepts, prerequisites, and affiliations.

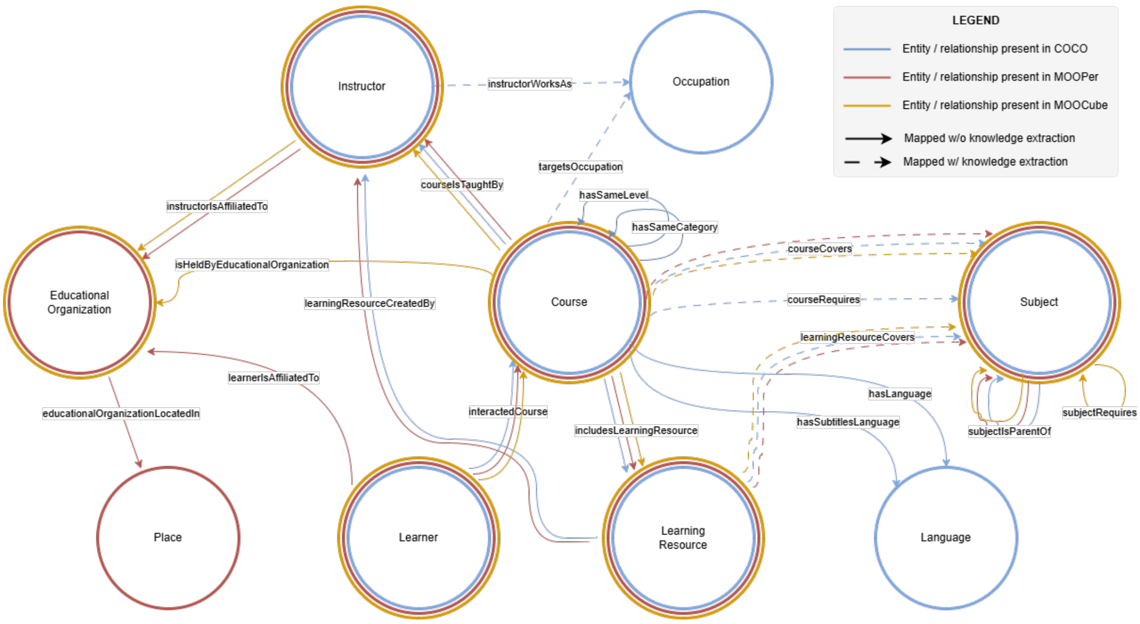

Fig. 5.1: KG Ontology. Entities and relationships represented in each constructed knowledge graph (KG). Learner, course, subject, and learning resource entities are included in all KGs.

The PEEK dataset (data: `github.com/sahanbull/PEEKC-Dataset`) comprises video metadata and learner activity data extracted from VideoLectures.Net, a repository of scientific and educational video lectures. Unlike the other contexts, it focuses on learner engagement with video-based educational content, offering insights into watch-time interactions and concept coverage. With a large-scale interaction dataset, PEEK captures 290, 535 learner interaction events from 20, 019 distinct users, each with at least five engagement events. The dataset spans 10, 233 unique lecture videos, segmented into 39, 113 fragments, averaging 3.82 fragments per video. Unlike other datasets of its kind, it provides a structured view of learner engagement with open educational resources, facilitating videolecture recommendations.

Finally, the `MOOPer` dataset (data: `https://www.educoder.net/ch/rest`) contains data from a practice-oriented platform for computer science focused on exercises and challenges. Unlike traditional academic platforms and course marketplaces, it focuses on interactive engagement, attracting a smaller but highly active user base that aims to reinforce concepts through applied practice. Hence, the dataset includes 205, 315 implicit interactions (enrollments) from 27, 992 learners across 327 courses, along with detailed course metadata on chapters, topics, disciplines, exercise titles, and institutional affiliations.

5.2.2 Knowledge Graphs Construction

This step involves constructing the KGs that serve as input for the reasoning methods. For each selected dataset, we created a recommendation-tailored KG using the original data provided, following a common ontology.

The first step involved selecting an ontology to structure the KGs, providing a formal, common representation of domain knowledge by defining entities, their properties, and relationships. For this work, we chose LOXER [122], a modern ontology specifically designed for the educational recommendation, developed collaboratively with input from learners, educators, and researchers. In this ontology, the central entity is the *Learner*, representing an individual pursuing specialized knowledge or skills. Supporting entities include *Course*, which defines structured learning paths, and *Learning Resource*, encompassing diverse content types, such as videos, chapters, and exercises. The ontology also includes *Subject* to represent academic topics, *Institution* and *Instructor* to capture the educational context, and *Occupation* to connect knowledge to job positions. Key relationships include *interacted*, which captures learner engagement; *is_parent_of*, which introduces hierarchical structures; *pre_requisite*, which supports guided learning pathways; and *covers*, contextualizing resources within subjects.

Next, all text data in each dataset were translated into English to standardize the language[1]. For instance, MOOCCube included Chinese course titles, while COCO contained multilingual course data in languages such as Spanish, German, and Arabic. After translation, each dataset's content was mapped to the ontology through a series of mechanical mappings, aligning it with the ontology's structure and relationships. These mappings consisted solely of simple format transformations without natural language processing, as represented by continuous lines in Figure 5.1. Specifically, we associated the IDs of learners, courses, instructors (where available), educational organizations (e.g., universities in MOOCCube), and their locations with corresponding entities in the ontology: *Learner*, *Course*, *Instructor*, *Institution*, and *Place*. Additionally, the IDs of learning materials, such as lessons, videos, and activities, were linked to the *Learning Resource* entity, while those of languages were assigned to the *Language* entity. Relationships between these entities were also established mechanically, covering *instructor-to-course* affiliations, *learner-to-course* and *learner-to-learning_resource* interactions, *course-to-language* affiliations, and *course-to-resource* connections.

In the final step, we enriched the initial KGs by adopting knowledge extraction techniques similar to those used in prior research [48, 49]. To this end, we merged course metadata for each course and extracted `Subject` entities[2], connecting them to `courses` via the `covers` relationship. We then leveraged the hierarchical structure of `Wikipedia` to instantiate the `is_parent_of` relationship for the `Subject` entity. This approach was applied to all datasets, using at least course titles (e.g., in

[1] We used the Google Translate tool (https://pypi.org/project/googletrans/).

[2] We relied on the Wikifier tool (https://www.wikifier.org/info.html).

Collaborative Components					Knowledge Components				
(lr)1-5(lr)6-10	COCO	MOOCCube	PEEK	MOOPer		COCO	MOOCCube	PEEK	MOOPer
Users (Core)	24,036 (10)	6,486 (10)	1,025 (5)	13,885 (6)	Entities (Types)	11,246 (7)	25,110 (6)	4,875 (5)	6,072 (7)
Courses (Core)	8,196 (10)	549 (10)	355 (5)	266 (6)	Relations (Types)	104,983 (6)	142,416 (3)	14,391 (2)	21,020 (3)
Interactions	378,469	96,579	8,111	145,850	Avg. Degree (All)	18.67	11.34	5.90	6.92
Density	0.002	0.014	0.020	0.040	Avg. Degree (Courses)	12.80	259.50	40.53	79.02

Table 5.1: KGs Statistics. COCO has the highest number of users, courses, and interactions but also the lowest density, whereas PEEK and MOOPer exhibit higher densities. MOOCCube stands out with the largest number of entities and relations.

MOOCube) and, where available, additional fields (e.g., title, long description, and short description in COCO). Additionally, in COCO, which provides further metadata on target audiences and objectives, we instantiated the Occupation entity and linked it to courses. We employed a zero-shot classifier [138] to extract occupational groups based on ESCO[3] profession labels (e.g., *Science and engineering associate professionals* and *Information and communications technology professionals*) from audiences and goals texts. Only in COCO, we extracted subjects from prerequisites metadata using the same above method to establish pre_requisite relationships between Course and Subject entities.

Our approach adheres to established practices from past KG construction efforts for recommendation. While creating a more comprehensive KG is beyond the scope of this study, our work lays the foundation for this direction. To support further research, the resulting KGs are available in our repository.

5.2.3 Knowledge Graphs Preprocessing

In this step, we prepared all the KGs for use in training KG reasoning models. This process involved cleaning, splitting, and coding meta-paths.

Similarly to previous studies [124, 125], we refined the KG by filtering out low-engagement learners and courses, focusing on those with higher activity. We defined a dataset-specific core (5, 6 or 10) subset to capture sufficient personalization while addressing data sparsity, balancing personalization with near-cold-start conditions typical in education. After filtering, we retained only entities with at least one connection to ensure meaningful KG structure. The descriptive statistics of the pre-processed KGs are reported in Table 5.1.

Next, we divided the learner-course interactions within the KG into training, validation, and test sets. Leveraging the timestamps available for each learner-course interaction in the KG, we implemented a per-user timeline split, allocating 80% of each learner's earliest interactions to the training set, 10% to the validation set, and the remaining 10%

[3] The ESCO database is available at https://esco.ec.europa.eu/en/.

to the test set. This time-based approach aligns with best practices in KG reasoning, preserving the sequential nature of interactions and maintaining consistency across learner timelines. As many KG reasoning methods (e.g., PGPR and CAFE) require knowledge graph embeddings, we then extracted embeddings for the train split of our pre-processed KGs. Given that original works used TransE [139] embeddings, we trained a distinct TransE model for each KG after having categorized entities and relationships by type.

Finally, we coded a set of meta-paths, i.e., predefined sequences of entities and relationships within the KG relevant for explaining recommendations in the educational domain. Each meta-path starts with a learner entity and progresses through related entities, ending with a course entity for recommendation. For this study, we specifically used twelve three-hop meta-paths validated in [122], which reflect educational decision-making factors. These paths guide agent behavior in reinforcement learning, predictions in neuro-symbolic models, and sequence inference in generative models.

5.2.4 Models Selection and Optimization

After constructing and preprocessing the base KGs, we evaluated several models, each representing a distinct approach, to assess their effectiveness in education. Specifically, we selected representative models from three major families: general recommendation, KG-based reasoning, and path language modeling. All models were implemented and evaluated under a unified protocol in our repository. We detail General, KG-based, and language model approaches here; additional information is available in our repository. We further divide the recommendation into four subgroups: traditional baselines, matrix factorization, graph-based, and autoencoder-based models. Similarly, KG-based methods are grouped into graph embedding, propagation-based models, graph neural networks, and reinforcement learning-based models.

The traditional baseline family includes simple yet fundamental models- POP and Random [140], which recommend the most popular and random items, respectively, serving as benchmark baselines. In the matrix factorization family, ENMF [141] enhances training efficiency using a neural matrix factorization framework, while BPR [142] learns personalized rankings through pairwise comparisons. Among graph-based models, NGCF [143] organizes user-item interactions into a graph and extracts useful information from multi-hop neighbors to refine latent representations; DGCF [144] further advances graph-based collaborative filtering with a disentangled mechanism to capture diverse user intents; and LightGCN [145] uses simple graph-based model that includes the essential component of graph convolutional networks, and uses neighborhood aggregation to improve scalability. The final model in the general recommendation category is the autoencoder-based MacridVAE [146], which adopts an item-based collaborative filtering approach, learning disentangled user behavior representations and ranking items by modeling latent preference factors.

Inside the KG-based models, first sub-family focuses on graph embeddings to represent user-item interactions. CFKG [147] merges the user-item interaction graph with the KG to jointly learn embeddings; CKE [148] enriches item representations by integrating KG information with auxiliary data; and MKR [149] adopts a multi-task learning framework to jointly optimize recommendation and KG embedding objectives. We also include propagation-based model such as KTUP [150] and Ripplenet [151], which actively combines features from neighboring nodes and relations through learned transformations. Within the GNN-based family, KGCN [152] learns entity representations by mining KG-associated attributes, while KGNNLS [153] transforms the KG into a user-specific weighted graph and applies a graph neural network to learn personalized item embeddings. For reinforcement learning-based reasoning, we selected PGPR [48] and for neural-symbolic reasoning, we included CAFE [49], and PLM [82] for generative reasoning. To address PLM's susceptibility to hallucination, we included PEARLM [83] that addresses hallucination and improves path faithfulness. All models were implemented and evaluated under a unified protocol (see our repository).

We optimized each model's hyperparameters via grid search, focusing on those highlighted in the original papers. When data characteristics affected performance, we expanded the search space. Model selection maximized utility on the validation set. Full hyper-parameter details are in our repository.

5.2.5 Path Reasoning Models Evaluation

We evaluated methods' utility, beyond-utility, and explanation metrics on top-10 courses and their reasoning paths in the KG's test set using the preprocessed data and metric computation code.

We included four utility metrics to provide an evaluation of performance from different angles. Specifically, Normalized Discounted Cumulative Gain (NDCG) [154] with binary relevance and logarithmic decay, measures the ranking quality by considering the positions of relevant items in the list. Mean Reciprocal Rank (MRR) [155] focuses on the rank of the first relevant item. Precision (PREC) and Recall (REC) [156] measure the proportion of recommended items that are relevant and of all relevant items recommended, respectively.

We also incorporated four beyond-utility metrics to evaluate additional aspects that contribute to user satisfaction not captured by utility metrics [157]. Specifically, catalog coverage (COV) was monitored as the extent to which the recommended items span the available catalog. Serendipity (SER) captures the element of surprise in recommendations by evaluating how different the recommendations are from those of a non-personalized most popular model. Novelty (NOV) measures the system's ability to recommend less popular items, while (POP) measures the ability to recommend widely consumed or frequently interacted items.

Finally, we assessed explanation-related metrics from various perspectives of path reasoning, including recency, popularity, and path type diversity [124]. The recency of linked interactions (LIR) reflects the age of past interactions in the generated paths, helping users connect recommendations to recent interactions. Linking interaction diversity (LID) captures the variety of prior interactions within the paths. Shared entity popularity (SEP), measured by node degree in the KG, assumes that more popular entities are likely more familiar to users. Shared entity diversity (SED) indicates the range of distinct entities featured. We also evaluated the path quality metric, where path type diversity (PTD) measures how many different relation types are used as the final step in the reasoning paths, indicating diversity in explanation types. Path type concentration (PTC) complements (PTD) by assessing how evenly explanation types are distributed across the recommended list. We also evaluated the path pattern type (PPT), which captures the structural template of relation sequences in the paths, helping to analyze the complexity of reasoning templates used in the explanations. Lastly, we assess the diversity of interaction types in the reasoning paths using (LITD), and measure the diversity of the shared entities in the explanation paths (SETD).

5.3 Experimental Results

This section examines the generalizability of KG reasoning methods to education. We first assess their utility in educational recommendations (**RQ1**), then analyze support for beyond-utility goals (**RQ2**), and evaluate their effectiveness in providing informative reasoning paths for explanations (**RQ3**).

Family	sub-family	Method	COCO				MOOCube				PEEK				MOOPer			
			NDCG ↑	MRR ↑	REC ↑	PREC ↑	NDCG ↑	MRR ↑	REC ↑	PREC ↑	NDCG ↑	MRR ↑	REC ↑	PREC ↑	NDCG ↑	MRR ↑	REC ↑	PREC ↑
General Recommendation	**Traditional Baseline Model**	Random	0.0006	0.0005	0.001	0.0002	0.008	0.006	0.01	0.002	0.008	0.005	0.02	0.002	0.02	0.01	0.04	0.004
		POP	0.03	0.02	0.06	0.006	0.08	0.06	0.15	0.02	0.04	0.03	0.09	0.009	0.08	0.05	0.17	0.01
	Matrix Factorization Model	ENMF	0.12	0.11	0.20	0.02	0.22	0.19	0.36	0.04	0.35	0.28	0.59	0.06	**0.87**	**0.84**	0.96	0.10
		BPR	0.12	0.10	0.20	0.02	0.23	0.20	0.37	0.04	0.29	0.22	0.52	0.05	0.80	0.76	0.95	0.10
	Graph Model	DGCF	0.10	0.08	0.17	0.02	0.22	0.19	0.36	0.04	0.35	0.28	0.58	0.06	0.81	0.77	0.96	0.10
		NGCF	0.08	0.06	0.14	0.01	0.23	0.19	0.38	0.04	0.31	0.23	0.56	0.06	0.78	0.73	0.95	0.10
		LightGCN	0.12	0.10	0.20	0.02	0.24	**0.21**	0.39	0.04	0.34	0.27	0.56	0.06	0.77	0.72	0.94	0.10
	Autoencoder Model	MacridVAE	**0.15**	**0.13**	**0.24**	**0.02**	0.24	**0.21**	0.38	**0.04**	**0.36**	**0.29**	0.60	**0.06**	0.85	0.81	**0.97**	**0.10**
Knowledge Based Recommendation	**Graph Embedding Model**	CFKG	0.10	0.08	0.18	0.02	**0.25**	**0.21**	**0.40**	0.04	0.34	0.27	**0.61**	0.06	0.81	0.76	0.95	0.09
		CKE	0.12	0.10	0.20	0.02	0.24	**0.21**	**0.40**	0.04	0.34	0.26	0.58	0.06	0.80	0.75	0.96	0.09
		MKR	0.10	0.08	0.18	0.02	0.18	0.15	0.30	0.03	0.16	0.12	0.29	0.03	0.58	0.50	0.83	0.08
	Propagation Model	KTUP	0.08	0.07	0.14	0.02	0.22	0.18	0.36	0.04	0.29	0.22	0.52	0.05	0.79	0.73	0.95	0.09
		RippleNet	0.06	0.04	0.11	0.01	0.18	0.14	0.30	0.03	0.29	0.22	0.52	0.05	0.58	0.52	0.79	0.08
	GNN model	KGCN	0.07	0.06	0.14	0.01	0.21	0.17	0.35	0.04	0.31	0.24	0.55	0.05	0.77	0.71	0.94	0.09
		KGNNLS	0.09	0.08	0.16	0.02	0.23	0.19	0.38	0.04	0.25	0.19	0.46	0.04	0.76	0.70	0.94	0.09
	Reinforcement Learning Model	PGPR	0.03	0.03	0.05	0.006	0.11	0.08	0.23	0.027	0.17	0.12	0.34	0.035	0.32	0.28	0.48	0.052
		CAFE	0.02	0.01	0.04	0.004	0.05	0.04	0.11	0.014	0.09	0.06	0.19	0.020	0.45	0.37	0.76	0.079
Path Language Based Recommendation	**Generative model**	PEARLM	0.03	0.02	0.05	0.006	0.09	0.08	0.12	0.014	0.08	0.06	0.15	0.016	0.31	0.29	0.38	0.038
		PLM	0.04	0.04	0.06	0.007	0.10	0.09	0.14	0.016	0.02	0.02	0.03	0.003	0.28	0.26	0.35	0.035

Families: General Recommendation; Knowledge Based Recommendation; Path Language Based Recommendation.
Metrics: Norm. Disc. Cum. Gain (NDCG); Mean Recip. Rank (MRR); Recall (REC); Precision (PREC).
For each dataset: best result in **bold**, second-best result underlined.

Table 5.2: RQ1: Utility. PEARLM achieves the highest utility across all metrics and datasets. In COCO and MOOCube, it outperforms traditional collaborative filtering, followed by non-explainable knowledge-aware models, with other reasoning methods showing the lowest performance. In MOOPer, both generative models lead, with the performance ranking for other methods remaining consistent to COCO and MOOCube.

5.3.1 RQ1: Generalizability of Utility Outcomes

In a first analysis, we evaluated the recommendation utility of the selected methods within the context of education. To do so, we assessed the performance of each method on four distinct datasets using standard evaluation metrics: NDCG, MRR, Recall, and Precision (Table. 5.2). This analysis aimed to provide insight into which families of methods are most effective in supporting relevant learning resource recommendations across varying data characteristics and sparsity levels.

The results shows several important patterns. In COCO, an autoencoder-based model (MacridVAE) achieves the highest performance across all metrics (e.g., NDCG (0.15), MRR (0.13), Recall (0.24)), outperforming both traditional and graph methods. This suggests that modeling latent representations with reconstruction loss is well-suited to this sparse dataset. In MOOCube, a knowledge-aware embedding model (CFKG) performs best in terms of NDCG (0.25) and Recall (0.40), indicating that explicit modeling of entity relations in a knowledge graph provides a substantial utility gain over general models like LightGCN and MacridVAE. Among general models, both MacridVAE and LightGCN remain competitive, suggesting generalizability to mid-density settings. In PEEK, traditional matrix factorization models (especially ENMF and BPR) dominate the leaderboard. ENMF achieves the highest NDCG (0.35), MRR (0.28), and Recall (0.59), while MacridVAE slightly edges out in Precision (0.06). Interestingly, path-based methods such as PGPR and CAFE show weaker performance, indicating their limited utility in this setting. In MOOPer, ENMF again achieves the strongest overall performance with scores of 0.87 (NDCG), 0.84 (MRR), and 0.96 (Recall), confirming its robustness in dense datasets. Among the generative path language-based models, PLM performs better than PEARLM, with improvements of 6.5% in MRR and 16% in Precision. However, both still lag behind traditional models, suggesting that their interpretability comes at a cost in ranking performance under high-density conditions.

Overall, collaborative filtering and autoencoder-based models demonstrate strong generalizability across all datasets. Generative models like PEARLM and PLM offer potential for explainability but tend to underperform in utility, especially on dense datasets such as MOOPer. Knowledge-aware embedding models (e.g., CFKG) excel in mid-sparsity scenarios like MOOCube, while path reasoning methods such as PGPR and CAFE show limited adaptability across domains. These findings highlight the trade-offs between utility and explainability.

5.3.2 RQ2: Generalizability of Beyond-Utility Outcomes

In our second analysis, we evaluated the effectiveness of different recommendation methods in addressing beyond-utility aspects of user experience. We specifically measured serendipity, novelty, average popularity of recommended items, and catalog coverage across four educational datasets (Tab. 5.3). This analysis aimed to understand how well each method balances between surprising the user, introducing novel items,

Family	sub-family	Method	COCO				MOOCube				PEEK				MOOPer			
			SER ↑	NOV ↑	Avg POP ↑	COV↑	SER ↑	NOV ↑	Avg POP ↑	COV ↑	SER ↑	NOV ↑	Avg POP ↑	COV ↑	SER ↑	NOV ↑	Avg POP ↑	COV ↑
General Recommendation	**Traditional Baseline Model**	Random	0.99	0.93	38.23	**0.99**	**0.98**	0.93	135.9	**0.99**	**0.96**	**0.85**	15.63	0.95	0.95	0.86	405.13	0.99
		POP	0.02	0.25	**1370.82**	0.002	0.66	0.63	801.8	0.05	0.96	0.55	47.32	0.05	0.66	0.58	**1362.9**	0.08
	Matrix Factorization Model	ENMF	0.88	0.75	447.9	0.29	0.66	0.65	759.05	0.52	0.81	0.59	42.8	0.93	0.85	0.68	1058.2	0.98
		BPR	0.86	0.76	439.7	0.55	0.65	0.65	763.6	0.78	0.85	0.62	39.9	0.92	0.87	0.68	1049.2	0.96
	Graph Model	DGCF	0.79	0.69	567.2	0.22	0.61	0.61	839.8	0.59	0.82	0.58	45.2	0.70	0.88	0.70	985.7	0.97
		NGCF	0.95	0.86	245.1	0.69	0.68	0.66	735.9	0.75	0.89	0.68	34.2	0.98	0.86	0.67	1074.0	0.98
		LightGCN	0.84	0.74	477.7	0.37	0.65	0.64	785.7	0.69	0.80	0.56	46.9	0.50	0.87	0.68	1028.7	0.92
	Autoencoder Model	MacridVAE	0.89	0.80	361.1	0.72	0.58	0.60	873.0	0.65	0.84	0.62	39.9	0.90	0.84	0.65	1138.1	0.91
Knowledge Based Recommendation	**Graph Embedding Model**	CFKG	0.78	0.68	573.7	0.49	0.53	0.56	934.6	0.81	0.52	0.49	49.6	0.96	0.86	0.67	1046.3	0.99
		CKE	0.87	0.77	410.38	0.61	0.64	0.63	784.9	0.82	0.81	0.70	29.3	**0.99**	0.82	0.69	872.1	0.99
		MKR	0.84	0.74	479.70	0.51	0.44	0.51	**1039.7**	0.18	0.20	0.33	65.5	0.18	0.78	0.65	995.1	0.75
	Propagation Model	KTUP	0.82	0.72	500.7	0.65	0.59	0.59	868.3	0.72	0.77	0.68	31.6	0.99	0.71	0.61	1125.8	0.98
		RippleNet	0.87	0.77	428.4	0.48	0.57	0.58	889.6	0.57	0.76	0.66	33.5	0.87	0.78	0.66	967.2	0.99
	GNN model	KGCN	0.84	0.73	480.3	0.53	0.64	0.62	812.27	0.67	0.77	0.66	33.2	0.98	0.76	0.65	1023.6	0.99
		KGNNLS	0.90	0.80	350.5	0.84	0.70	0.67	706.05	0.86	0.79	0.67	32.4	0.83	0.76	0.65	1010.0	**0.99**
	Reinforcement Learning Model	PGPR	0.99	0.93	113.9	0.63	0.69	0.65	760.3	0.44	0.82	0.58	43.9	0.89	**0.95**	0.82	577.4	0.98
		CAFE	**1.0**	**0.97**	61.8	0.35	0.93	0.82	399.4	0.36	0.46	0.32	**71.9**	0.31	0.93	0.75	794.4	0.91
Path Language Based Recommendation		PEARLM	0.83	0.71	517.8	0.02	0.61	0.69	671.2	0.04	0.92	0.78	22.6	0.57	0.88	0.83	556.7	0.32
		PLM	0.59	0.67	722.06	0.06	0.65	0.70	647.8	0.06	0.91	0.86	13.93	0.05	**0.95**	**0.86**	429.7	0.48

Families: General Recommendation; Knowledge Based Recommendation; Path Language Based Recommendation.
Metrics: Serendipity (SER); Novelty (NOV); Average Popularity (Avg POP); Coverage (COV).
For each dataset: best result in **bold**, second-best result underlined.

Table 5.3: RQ2: Beyond Utility. Traditional collaborative filtering methods excel in serendipity and novelty. PGPR shows stronger performance on average in diversity. PEARLM stands out with high coverage, with moderate scores across other metrics.

reducing popularity bias, and covering a broader catalog.

The results reveal that traditional baseline models like Random consistently achieve the highest serendipity and coverage scores across all datasets (e.g., SER- 0.99 and COV- 0.99 in both COCO and MOOCube). However, their utility remains extremely low (see Table. 5.3), making them impractical for real-world deployment. Similarly, POP yields high average popularity scores (e.g., 1370.82 in COCO).

Among general recommendation models, matrix factorization methods such as ENMF and BPR achieve strong performance in novelty and moderately good coverage. For example, in COCO, ENMF attains a novelty score of 0.75 and a coverage of 0.66. Likewise, graph-based models like NGCF and LightGCN show competitive scores in novelty and serendipity, reflecting strength in learning patterns with no external knowledge.

In the knowledge-based family, models such as CFKG, MKR, and RippleNet exhibit a good balance across all metrics. Notably, MKR achieves the highest novelty score (0.71) in MOOCube and maintains competitive diversity and coverage across datasets, suggesting its ability to surface underexposed but relevant content. Among propagation and GNN-based models, RippleNet and KGCN also maintain relatively high novelty and moderate serendipity, confirming the potential of knowledge graph signal propagation in promoting novel learning materials.

Focusing on path-based KG reasoning models, we observe that PGPR excels in serendipity, achieving the top score in COCO (1.0) and high values in other datasets (e.g., 0.96 in MOOCube, 0.95 in MOOPer). It also shows competitive diversity, with lower average popularity scores, suggesting it avoids recommending only well-known content. However, its coverage remains relatively limited compared to other methods.

PEARLM demonstrates the highest coverage scores across multiple datasets, achieving 0.80 in COCO, 0.99 in MOOCube, and 0.99 in MOOPer. This indicates a strong ability to recommend from a wide selection of items in the catalog. Still, it falls behind in novelty and serendipity compared to both general and graph-based models. PLM, while competitive in coverage and novelty (e.g., 0.67 in MOOCube, 0.86 in MOOPer), also shows limitations in consistently balancing all beyond-utility aspects.

Overall, KG reasoning methods show complementary strengths in beyond-utility metrics. PGPR leads in serendipity, while PEARLM achieves the highest catalog coverage. Other methods like MKR and RippleNet show a balance across novelty and popularity reduction. However, trade-offs remain, and no single model dominates.

5.3.3 RQ3: Generalizability of Explainability Outcomes

In our final analysis, we evaluated the explainability of the KG reasoning methods through six core metrics related to explanation path quality and diversity across four educational datasets (Table 5.4). The metrics linking interaction recency (LIR), linking interaction diversity (LID), shared entity popularity (SEP) and diversity (SED), path type diversity (PTD), and path pattern type (PPT) are important to understand how well the models support transparency, interpretability, and educationally relevant recommendations. PLM is retained here for completeness, although it previously showed high rates of unfaithful path generation due to invalid KG links.

Method	COCO						MOOCube						Peek						MOOper					
	LIR ↑	LID ↑	SEP ↑	SED ↑	PTD ↑	PPT ↑	LIR ↑	LID ↑	SEP ↑	SED ↑	PTD ↑	PPT ↑	LIR ↑	LID ↑	SEP ↑	SED ↑	PTD ↑	PPT ↑	LIR ↑	LID ↑	SEP ↑	SED ↑	PTD ↑	PPT ↑
PGPR	0.46	**0.68**	0.64	**0.81**	0.60	0.69	**0.41**	**0.67**	0.23	**0.99**	0.40	0.28	**0.45**	0.45	0.30	0.98	0.48	0.27	0.43	0.66	0.42	**0.99**	0.46	0.36
CAFE	**0.48**	0.04	0.64	0.21	0.42	**0.85**	0.37	0.04	0.26	0.22	0.36	0.41	0.45	0.08	**0.47**	0.26	0.17	0.15	**0.44**	0.06	0.11	0.29	0.31	0.36
PEARLM	0.40	0.25	**0.98**	0.40	0.38	0.40	0.32	0.44	0.68	0.69	0.50	**0.49**	0.44	0.33	0.35	0.65	0.49	0.31	0.36	**0.60**	**0.44**	0.78	0.68	**0.68**
PLM	0.09	0.39	0.95	0.45	**0.60**	0.83	0.07	0.31	0.41	0.28	**0.73**	0.29	0.02	**0.99**	0.14	**0.99**	**0.99**	**0.99**	0.26	0.59	0.10	0.55	**0.79**	0.49

Metrics (1): Linking interaction recency (LIR) and diversity (LID).

Metrics (2): Shared entity popularity (SEP) and diversity (SED); Path type diversity (PTD); Path pattern type (PPT).

Table 5.4: RQ3: Explainability. PGPR excels in linking interaction recency, linking interaction diversity, and shared entity diversity, especially in sparser datasets like COCO and MOOCube. CAFE performs strongly in shared entity popularity, while PEARLM leads in explanation type diversity across datasets.

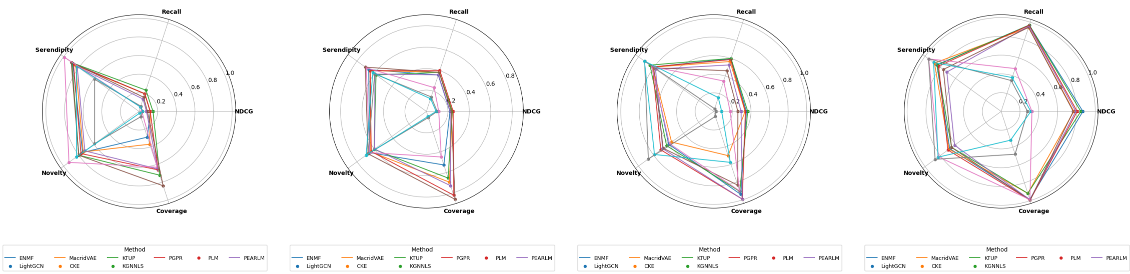

Fig. 5.2: Methods Comparison. Comparison of recommendation utility and beyond-utility goals across datasets: (a) COCO, (b) MOOCCube, (c) PEEK, and (d) MOOPer.

PGPR demonstrates strong performance in LIR and LID, indicating its ability to generate explanations rooted in recent and varied user interactions. These qualities are particularly pronounced in COCO and MOOCube, where PGPR achieves the highest LIR and LID scores. For instance, in COCO, PGPR's LID is 68%, substantially higher than the next best method (CAFE at 64%), and its SED is also the highest (0.81), indicating diverse entity usage in paths. CAFE stands out in SEP, often scoring among the top across all datasets, including 0.64 in COCO and 0.47 in Peek. This implies a reliance on popular, frequently occurring entities when constructing explanations and potentially increasing user trust. PEARLM, on the other hand, leads significantly in explanation diversity, especially in SED, PTD, and PPT. It consistently scores the highest SEP and SED across datasets such as COCO (SEP: 0.95, SED: 0.98), and maintains strong performance in PTD and PPT in MOOCube and MOOper. For example, in MOOCube, it achieves a PTD of 0.60 and PPT of 0.49, indicating a greater variety of path types and patterns that enrich the semantic structure of the explanations. This breadth can help learners connect recommendations to a wider range of pedagogical contexts (e.g., course objectives, instructors, or institutional offerings). In contrast, PLM lags behind on nearly all metrics, especially LIR and LID, reflecting its tendency to generate less relevant and temporally distant paths, and limited diversity in entity and path types.

Findings RQ3Different methods showcase unique strengths in generating explainable recommendations. PGPR excels in recency and diversity of learner interactions, grounding its recommendations in memorable user history. CAFE prioritizes popular entities, enhancing user familiarity with explanations. PEARLM offers the most diverse and pedagogically rich explanations through varied entity types and paths. These findings underline the importance of selecting explanation strategies tailored to educational goals and learner preferences.

5.4 Discussion and Implications

In this chapter, we assessed th generalizability of path reasoning recommendation within the educational domain. By extracting explicit reasoning paths from educational knowledge graphs, these methods provide transparent and faithful reasoning that link learners to recommended actions. The effectiveness of such reasoning paths is strongly influenced by the underlying learner-centered ontology introduced in Chapter 3, confirming that explainability at the reasoning level is inherently constrained by representation-level design choices. When educational entities and relations are modeled explicitly and coherently, path-based methods can generate explanations that are both interpretable and pedagogically meaningful.

At the same time, the findings reveal that different methods excel based on specific goals and data characteristics. Generative methods offer strong utility and catalog coverage, while PGPR is valuable for diversity and serendipity. For explainability-focused contexts, CAFE balances familiarity in explanations by focusing on popular entities, whereas PEARLM emerges with its diverse explanation types. Overall, KG reasoning methods not only enhance explainability but also serve as effective alternatives. This observation connects naturally to the conversational interaction strategies explored in Chapter 4 and provides the foundation for Chapter 6, where large language models are employed to translate structured reasoning paths into context-aware natural language explanations. Together, these chapters position path-based reasoning as a critical grounding mechanism that enables faithful explainability while supporting more accessible and adaptive explanation delivery.

Chapter 6

Methods for Path-based Natural Language Explanation

Research Highlights

- *Path-based grounding* constrains explanation generation.
- LLM can generate *fluent* explanations when grounded in KG paths.
- Explanation quality varies across *models*, *contexts*, and *tone* settings.
- Motivational explanations tone favors *engagement*.
- Professional explanations enhance *faithfulness* and *transparency*.

6.1 Introduction

The increasing reliance on AI-driven support systems in lifelong learning environments intensifies the need for explanations that are not only correct but also understandable, engaging, and aligned with learners' cognitive and contextual needs. In educational recommendation scenarios, explanations play a critical role in helping learners interpret the system output, and make informed decisions [158] [52].

In the preceding chapters, we investigated path-based recommendation methods, demonstrating how structured reasoning paths over knowledge graphs can support faithful and transparent recommendations. These path-based methods ensure that explanations are grounded in explicit semantic relations, thereby addressing core requirements of transparency and accountability. Despite their advantages, path-based explanations are typically expressed through symbolic representations or template-based textual realizations, which directly verbalize reasoning paths. While such explanations preserve faithfulness to the underlying decision logic, they often lack linguistic flexibility, personalization, and adaptability. As a result, explanations may appear overly technical or insufficiently address users' information needs in lifelong learning settings [159].

Recent advances in large language models (LLMs) offer promising opportunities to address these limitations. LLMs are capable of generating fluent, context-aware natural language and can adapt explanations to different tones, levels of detail, and learning contexts [160] [161]. When applied to explanation generation, they have the potential to transform structured reasoning outputs into user-friendly narratives that improve comprehension and engagement. However, the use of LLMs in explainable decision support systems introduces critical challenges. In particular, unconstrained generation may lead to hallucinated content or explanations that are persuasive but unfaithful to the actual reasoning process, thereby undermining trust and transparency [162].

In this chapter, we investigate knowledge-aware explanation generation using large language models for educational decision support, focusing on how LLMs can be integrated with path-based recommendation methods without sacrificing faithfulness. Building directly on the structured reasoning paths produced by the methods presented in Chapter 5, we design a generation pipeline in which LLMs are explicitly grounded in knowledge graph paths, learner context, and recommendation metadata. This grounding enables explanations that remain consistent with the underlying decision logic while benefiting from the expressive capabilities of modern language models.

The contribution of this chapter is two-fold:

- A knowledge-aware pipeline for explanation generation, in which LLMs transform path-based reasoning outputs into natural language explanations.
- A large-scale automatic evaluation of LLM-generated explanations, assessing multiple explanation quality and dimensions.

The remainder of this chapter is structured as follows. Section 6.2 describes the proposed methodology, including the grounding strategy, prompt design, and explanation generation process. Section 6.3 presents the experimental results from evaluations. Finally, section 6.4 discusses the implications of the findings for knowledge-aware explainable decision support in lifelong learning.

6.2 Methodology

This chapter builds upon the structured reasoning outputs produced by the path-based recommendation methods introduced in Chapter 5. The methodology is designed to transform faithful, knowledge-graph-grounded reasoning paths into natural language explanations that are accessible to learners, while preserving consistency with the decision logic. To this end, the pipeline consists of two main stages: (i) path generation via a path-reasoning method and (ii) natural language explanation generation.

6.2.1 Path Generation

Path generation constitutes the first stage of the explanation pipeline and provides the structured input required for the natural language explanation generation. In this stage, reasoning paths are extracted from educational knowledge graphs using path-based recommendation methods. These paths explicitly connect learners to recommended educational resources through sequences of semantically meaningful entities and relations [20]. For instance, a generated reasoning path such as $user\_w$ enrolled_in $course\_x$ covers $concept\_y$ prerequisite_for $course\_z$ captures a link between a learner's past activity to a new recommendation. This path can be verbalized as the explanation: "$course\_z$ is recommended to you because it builds on concept $concept\_y$, which you covered in $course\_x$. Such paths provide an explicit and faithful trace of the reasoning process underlying the recommendation and serve as the grounding input for subsequent natural language explanation generation.

Data Preparation

The input to the path generation process consists of educational knowledge graphs constructed using a learner-centered ontology (for instance, those proposed in Chapter 3). These graphs represent learners, educational resources, learning objectives, subjects, and their associated semantic relationships. Before model training, the data undergo a preprocessing phase to ensure consistency, quality, and suitability for path-based reasoning. This phase includes the removal of infrequent entities and relations, the normalization of relation types, and the partitioning of the dataset into training, validation, and test sets, following standard practices in knowledge graph-based recommendation systems. Special attention is devoted to preserving the semantic integrity of prerequisite relations and learner-resource interactions, as these relations are essential for generating pedagogically meaningful and interpretable reasoning paths.

Model Training

Path-based recommendation models are trained to identify informative reasoning paths that connect learners to candidate educational resources (such as those described in Chapter 5). Depending on the adopted reasoning paradigm, the models learn to explore the underlying knowledge graph by optimizing objective functions that balance recommendation accuracy with the relevance and interpretability of the generated paths. Through this process, the models acquire the ability to capture both structural patterns in the graph and pedagogical relationships among learners, concepts, and materials.

Model Inference

At inference time, the trained path reasoning model generates one or more reasoning paths for each learner-resource pair. Each path is represented as an ordered sequence of entities and relations, forming a structured trace that explains how a specific recommendation is derived. These paths serve a dual purpose: they provide a justification for the recommendation and act as a grounding signal for the subsequent explanation generation stage. To promote explanation diversity and robustness, up to the top ten reasoning paths are usually generated for each recommendation. The resulting paths are then filtered according to criteria such as semantic coherence and relevance to the learner's profile, ensuring that only meaningful paths are retained for explanation generation.

6.2.2 Explanation Generation

The second stage of the methodology focuses on transforming reasoning paths into natural language explanations using LLMs [163]. This stage is designed to preserve faithfulness by grounding generation in the paths produced during the previous stage.

Prompt Design

Prompt design plays an important role in aligning the expressive capabilities of LLM with the requirements of faithful explanation generation. Prompts are constructed to explicitly include the reasoning path, relevant entity labels, relation descriptions, and contextual information about the recommended resource. Rather than allowing free-form generation, prompts constrain the model to base its explanation solely on the provided path information. Instructional constraints specify explanation length, style, and output format, discouraging unsupported claims and ensuring consistency across generated explanations.

In addition to grounding constraints, prompts control explanation tone and learning context, which are treated as experimental variables. Two tones are considered: professional and motivational. Professional explanations adopt a formal, goal-oriented style and are generated only for formal learning and learning-for-credits (L4C) contexts. Motivational explanations adopt a more engaging style and are generated across formal and non-formal settings, as well as learning-for-credits (L4C) and learning-for-pleasure (L4P) contexts.

Model Prompting

Using the designed prompts, LLMs are queried to generate explanations corresponding to individual reasoning paths. In this work, explanation generation is performed using open-weight language models executed through the Ollama runtime, including Gemma, LLaMA, Qwen, and Mistral. The models process the structured path information and produce natural language explanations that articulate the rationale behind a recommendation. To ensure consistency across explanations, generation parameters such as temperature and maximum output length are controlled. The resulting explanations are stored alongside their originating paths. A representative prompt template used is reported in Figure 6.1.

```
 You are an educational course recommendation assistant.
Your task is to write a brief, clear explanation for why a student should take a specific course.

Write a two-sentence explanation for why a student should take the course "{course}".

Information:
- The student has previously studied:  {learner_history}
- Relation type between these courses:  {reason_type}
- Tone:  {tone}
- Context:  {context}

Instructions:
- Write only two natural, fluent sentences.
- Interpret the meaning of the relation type yourself and express it naturally.
- Do not mention the relation name explicitly.
- Adapt the explanation to the specified tone and context.
- Do not include reasoning steps, planning, or any additional text.
- Output only the final explanation.
```

Fig. 6.1: Prompt template for explanation generation. The template constrains LLMs to generate short, grounded explanations based on structured reasoning paths, while explicitly controlling explanation tone and learning context.

Explanation Evaluation

The quality of the generated explanations is assessed through a property-based evaluation framework, designed to capture multiple dimensions of explanation effectiveness relevant to lifelong learning decision support [164]. To this end, we define eleven explanation properties, inspired by prior work on explainable recommender systems and explainable artificial intelligence [165], and adapted to the educational domain. These properties are evaluated automatically using large language models acting as evaluators [166], following a structured and controlled prompting protocol. Specifically, the selected properties cover complementary aspects and are grouped into five conceptual categories:

- **Support** to capture how explanations assist in making informed choices:
  - *Effectiveness*: the extent to which the explanation helps a learner decide whether to choose the recommended course. For brevity, we report only the evaluation prompt for effectiveness Figure 6.2, as all other property prompts share the same structure and differ only in the properties criterion.
  - *Decision Speed*: how quickly the explanation enables understanding.
  - *Confidence*: how the explanation increases learner's confidence.
- **Engagement** to reflect motivational and affective aspects:
  - *Motivation*: how strongly the explanation encourages interest.
  - *Satisfaction*: the perceived completeness and helpfulness of the explanation.
- **Transparency** related to explainability and user agency:
  - *Transparency*: how clearly the explanation reveals the reasoning.

```
 Consider that you are a student enrolled in courses on a MOOCs platform.
You are shown a recommended course along with an explanation.
You also have access to information about the course and your past interaction.

Use the explanation, the course information, and the past interaction provided below to evaluate
how effective the explanation is in helping you decide whether to choose the recommended course.

Choose one effectiveness category:
- excellent:  extremely effective and strongly motivates course choice
- very good:  highly effective and clearly helpful
- good:  reasonably effective and somewhat helpful
- moderate:  partially helpful but not strong
- poor:  ineffective or unhelpful in supporting the decision

Return ONLY this JSON object (no extra text):
{
"label":  "<excellent | very good | good | moderate | poor>",
"justification":  "<one sentence explaining why this label fits>"
}

Explanation (to be evaluated):  {explanation}
Course information (context of the course):  {course_info}
Past interaction (your past interactions):  {user_history}
```

Fig. 6.2: Prompt template used for automatic evaluation of explanation effectiveness. The prompt instructs the language model to assess explanation quality from a learner perspective and to return a structured judgment for a single explanation property.

  - *Correction Ease*: how easily a learner could identify which aspects of their profile influenced the recommendation and could be adjusted if needed.

- **Expressiveness** related to linguistic and semantic quality:
  - *Readability*: the clarity and ease of understanding of the explanation text.
  - *Informativeness*: the extent to which the explanation provides useful details.
- **Pertinence** to ensure alignment with the underlying data and reasoning:
  - *Faithfulness*: the accuracy of the explanation with respect to the course information and learner history, including the absence of fabricated content.
  - *Relevance*: the degree to which the explanation meaningfully connects the recommendation to the learner's past interactions, goals, or interests.

Each explanation is evaluated independently for each property using a five-level ordinal scale: excellent, very good, good, moderate, and poor. For every property, a dedicated evaluation prompt is constructed, providing the evaluator with the generated explanation, the recommended course information, the learner's past interaction history, and a precise definition of the target property and the evaluation criteria. The evaluator is instructed to return a structured JSON object containing the assigned label and a short textual justification. This protocol enforces consistency across evaluations and enables systematic comparison across explanation generation strategies and language models.

LLMs are employed as evaluators to enable fine-grained assessment across a large number of explanations. In this work, the evaluation is conducted using LLaMA 3.1

(8B), executed through the Ollama runtime, which is used consistently across all evaluation tasks. Each property is evaluated independently, ensuring that judgments remain focused on a single quality dimension at a time. To reduce variability, generation parameters are fixed and evaluation prompts are standardized across all experiments. While LLM-based evaluation does not replace human judgment, it provides a practical and reproducible means to analyze explanation behavior at scale. In the context of this chapter, the evaluation framework is used comparatively, focusing on relative differences between models, prompting strategies, and explanation configurations, rather than absolute quality claims most reliable models for explanation generation, though their strengths vary across properties, motivating more fine-grained analysis.

Table 6.1: RQ1: Overall Explanation Quality Benchmark. Average and standard deviation across explanation properties for COCO.

Property	Gemma3		LLaMA3.2		Mistral		Qwen2.5	
	Avg	Std	Avg	Std	Avg	Std	Avg	Std
Confidence Boost	0.745	0.019	0.761	0.023	0.696	0.034	**0.778**	0.034
Correlation Ease	0.754	0.037	0.816	0.042	**0.877**	0.036	0.783	0.029
Context Relevance	0.597	0.040	**0.619**	0.047	0.584	0.025	0.585	0.008
Decision Speed	0.796	0.013	0.808	0.007	0.782	0.008	**0.815**	0.010
Effectiveness	0.669	0.015	0.668	0.019	0.654	0.016	**0.680**	0.022
Faithfulness	0.586	0.017	0.584	0.028	0.586	0.010	0.565	0.006
Informativeness	0.896	0.027	**0.907**	0.026	0.877	0.036	0.891	0.032
Motivation	**0.806**	0.018	0.788	0.005	0.778	0.028	0.778	0.014
Readability	0.781	0.004	**0.792**	0.006	0.778	0.009	0.789	0.005
Satisfaction	0.684	0.028	0.713	0.028	0.654	0.045	**0.720**	0.029
Transparency	0.687	0.029	**0.730**	0.026	0.678	0.038	0.725	0.031
Total (Avg)	0.727	0.023	**0.744**	<u>0.023</u>	0.722	0.026	<u>0.737</u>	**0.020**

Table 6.2: RQ1: Overall Explanation Quality Benchmark. Average and standard deviation across explanation properties for MOOCCube.

Property	Gemma3		LLaMA3.2		Mistral		Qwen2.5	
	Avg	Std	Avg	Std	Avg	Std	Avg	Std
Confidence Boost	0.784	0.018	**0.808**	0.024	0.784	0.028	0.780	0.018
Correlation Ease	0.800	0.011	**0.871**	0.022	0.814	0.017	0.838	0.018
Context Relevance	0.717	0.033	0.707	0.030	0.723	0.038	**0.705**	0.028
Decision Speed	**0.823**	0.016	0.817	0.011	0.808	0.007	0.802	0.012
Effectiveness	0.628	0.014	0.617	0.014	0.634	0.013	**0.639**	0.011
Faithfulness	0.647	0.036	0.655	0.018	**0.660**	0.029	0.665	0.015
Informativeness	**0.919**	0.026	0.923	0.022	0.903	0.030	0.912	0.012
Motivation	**0.793**	0.024	0.772	0.007	0.783	0.029	0.753	0.017
Readability	**0.809**	0.017	0.798	0.001	0.794	0.007	0.800	0.007
Satisfaction	**0.737**	0.019	0.734	0.019	0.716	0.024	0.733	0.018
Transparency	0.746	0.027	**0.769**	0.020	0.744	0.030	0.754	0.015
Total (Avg)	0.764	0.022	<u>0.767</u>	<u>0.017</u>	0.767	0.022	**0.771**	**0.015**

Table 6.3: RQ1: Overall Explanation Quality Benchmark. Average and standard deviation across explanation properties for MOOPer.

Property	Gemma3		LLaMA3.2		Mistral		Qwen2.5	
	Avg	Std	Avg	Std	Avg	Std	Avg	Std
Confidence Boost	0.876	0.020	0.861	0.012	0.876	0.028	**0.879**	0.017
Correlation Ease	0.856	0.025	0.901	0.024	0.859	0.016	**0.894**	0.021
Context Relevance	**0.884**	0.085	0.887	0.031	0.883	0.011	0.906	0.023
Decision Speed	**0.842**	0.020	0.826	0.012	0.834	0.028	0.823	0.011
Effectiveness	0.691	0.014	0.672	0.018	0.690	0.020	**0.716**	0.026
Faithfulness	**0.770**	0.064	0.745	0.018	0.770	0.039	0.789	0.015
Informativeness	**0.965**	0.022	0.952	0.019	0.935	0.024	0.973	0.011
Motivation	0.832	0.028	0.805	0.015	**0.836**	0.047	0.808	0.017
Readability	**0.816**	0.025	0.791	0.004	0.796	0.013	0.801	0.004
Satisfaction	0.794	0.020	0.780	0.016	0.763	0.020	**0.805**	0.013
Transparency	0.773	0.030	0.786	0.014	0.765	0.021	**0.796**	0.012
Total (Avg)	0.828	0.031	0.830	0.017	0.830	0.026	**0.844**	**0.016**

Table 6.4: RQ1: Overall Explanation Quality Benchmark. Average and standard deviation across explanation properties for PEEK.

Property	Gemma3		LLaMA3.2		Mistral		Qwen2.5	
	Avg	Std	Avg	Std	Avg	Std	Avg	Std
Confidence Boost	0.722	0.015	0.718	0.034	0.712	0.021	**0.734**	0.020
Correlation Ease	0.754	0.017	**0.768**	0.026	0.736	0.012	0.747	0.014
Context Relevance	**0.568**	0.032	0.557	0.024	0.558	0.046	0.570	0.049
Decision Speed	0.793	0.015	0.779	0.014	0.780	0.010	**0.782**	0.026
Effectiveness	**0.600**	0.011	0.561	0.021	0.597	0.008	0.580	0.015
Faithfulness	**0.605**	0.025	0.504	0.027	0.532	0.030	0.538	0.040
Informativeness	**0.897**	0.021	0.854	0.034	0.848	0.034	0.870	0.017
Motivation	0.751	0.020	0.722	0.007	**0.739**	0.036	0.729	0.029
Readability	**0.810**	0.017	0.782	0.008	0.791	0.005	0.796	0.015
Satisfaction	**0.675**	0.019	0.628	0.035	0.648	0.024	0.638	0.021
Transparency	0.700	0.030	0.688	0.035	0.689	0.035	**0.711**	0.012
Total (Avg)	0.716	**0.021**	0.678	0.024	0.693	0.024	**0.718**	0.022

6.3 Experimental Results

This section presents the experimental results of the explanation generation and evaluation pipeline described in Section 6.2. The analysis is structured around three research questions, each addressing a distinct aspect of explanation quality in LLM-based explanation generation for lifelong learning decision support. Results are reported across four educational datasets and four LLMs, using the 11 properties introduced earlier.

6.3.1 RQ1: Overall Explanation Quality Benchmark

This research question investigates how different LLMs perform across explanation properties and datasets, and which models are most effective for specific explanation dimensions. Tables 6.1- 6.4 report the overall performance (Avg ± Std) across the explanation properties for the COCO, MOOCCube, MOOPer, and PEEK datasets, respectively.

Results reveal that, across all datasets, no single model consistently dominates all properties, highlighting the multi-dimensional nature of explanation quality. However, LLaMA 3.2 and Qwen 2.5 consistently achieve higher overall mean scores across datasets, with Qwen 2.5 obtaining the best total mean on MOOCCube (0.771), MOOPer (0.844), and PEEK (0.718), and competitive performance on COCO (0.737). The property-level analysis reveals that Informativeness and Readability achieve the highest scores across all models and datasets, indicating that LLMs are particularly effective at generating detailed and linguistically clear explanations when grounded in structured paths. However, Faithfulness and Context Relevance exhibit comparatively lower scores, especially on sparser datasets such as PEEK, confirming the challenge of maintaining strict grounding in learner history and knowledge graph content. Dataset characteristics also influence results. MOOPer, which provides richer and more structured learner-resource relations, yields consistently higher scores across most properties, whereas PEEK presents lower overall performance, particularly for faithfulness-related dimensions. These findings suggest that explanation quality is strongly dependent on both the language model and the underlying data structure. Nevertheless, Qwen 2.5 and LLaMA 3.2 emerge as the most viable models for the task of natural language explanation on top of paths.

6.3.2 RQ2: Professional vs. Motivational Explanation Contexts

In this second analysis, we examine how explanation quality differs between professional and motivational tone, focusing on the two best-performing models identified in RQ1. Tables 6.5- 6.8 present a comparative analysis between professional and motivational explanations for LLaMA 3.2 and Qwen 2.5 across the four datasets.

Table 6.5: RQ2: Professional vs Motivational Explanation Quality. Average and standard deviation across explanation properties for COCO.

Property	LLaMA3.2				Qwen2.5			
	Professional		Motivational		Professional		Motivational	
	Avg	Std	Avg	Std	Avg	Std	Avg	Std
Confidence Boost	0.783	0.004	0.750	0.020	**0.810**	0.001	**0.763**	0.031
Correlation Ease	**0.861**	0.001	**0.794**	0.034	0.813	0.006	0.768	0.023
Context Relevance	**0.682**	0.002	**0.587**	0.018	0.593	0.004	0.581	0.006
Decision Speed	0.811	0.004	0.806	0.008	**0.816**	0.001	**0.815**	0.014
Effectiveness	0.691	0.001	0.656	0.013	**0.700**	0.002	**0.670**	0.020
Faithfulness	**0.619**	0.007	0.566	0.014	0.561	0.006	**0.567**	0.005
Informativeness	**0.937**	0.004	**0.893**	0.020	0.921	0.002	0.876	0.030
Motivation	**0.782**	0.003	**0.792**	0.003	0.775	0.001	0.779	0.020
Readability	**0.798**	0.000	**0.788**	0.006	0.795	0.000	0.786	0.003
Satisfaction	0.745	0.005	0.697	0.019	**0.748**	0.000	**0.705**	0.022
Transparency	0.760	0.001	0.714	0.019	**0.760**	0.003	**0.708**	0.025
Total (Avg)	**0.788**	0.003	0.750	0.018	<u>0.781</u>	0.002	0.745	0.021

Table 6.6: RQ2: Professional vs Motivational Explanation Quality. Average and standard deviation across explanation properties for MOOCube.

Property	LLaMA3.2				Qwen2.5			
	Professional		Motivational		Professional		Motivational	
	Avg	Std	Avg	Std	Avg	Std	Avg	Std
Confidence Boost	**0.829**	0.002	**0.798**	0.028	0.770	0.013	0.785	0.018
Correlation Ease	**0.884**	0.006	**0.865**	0.024	0.822	0.020	0.846	0.021
Context Relevance	**0.731**	0.001	0.695	0.032	0.682	0.003	**0.717**	0.031
Decision Speed	**0.826**	0.005	**0.813**	0.011	0.790	0.002	0.809	0.010
Effectiveness	0.624	0.001	0.613	0.014	**0.638**	0.010	**0.640**	0.011
Faithfulness	**0.670**	0.001	0.647	0.018	0.661	0.005	**0.667**	0.017
Informativeness	**0.937**	0.004	**0.915**	0.024	0.907	0.002	0.914	0.013
Motivation	**0.767**	0.003	**0.775**	0.008	0.731	0.010	0.764	0.008
Readability	**0.798**	0.001	0.798	0.001	0.797	0.001	**0.802**	0.007
Satisfaction	**0.747**	0.001	**0.728**	0.022	0.727	0.001	0.736	0.008
Transparency	**0.783**	0.001	**0.762**	0.019	0.770	0.002	0.756	0.015
Total (Avg)	**0.800**	0.002	0.783	0.021	0.781	0.006	0.785	0.017

Table 6.7: RQ2: Professional vs Motivational Explanation Quality. Average and standard deviation across explanation properties for MOOPer.

Property	LLaMA3.2				Qwen2.5			
	Professional		Motivational		Professional		Motivational	
	Avg	Std	Avg	Std	Avg	Std	Avg	Std
Confidence Boost	**0.873**	0.008	0.856	0.009	0.863	0.003	**0.888**	0.014
Correlation Ease	**0.913**	0.025	0.895	0.028	0.874	0.006	**0.905**	0.021
Context Relevance	**0.922**	0.021	0.870	0.020	0.886	0.001	**0.916**	0.024
Decision Speed	**0.816**	0.006	**0.831**	0.011	0.810	0.003	0.830	0.007
Effectiveness	0.682	0.013	0.667	0.018	**0.714**	0.015	**0.718**	0.032
Faithfulness	0.765	0.006	0.735	0.016	**0.780**	0.009	**0.794**	0.014
Informativeness	**0.973**	0.004	0.942	0.016	0.962	0.003	**0.979**	0.010
Motivation	0.795	0.006	0.810	0.012	**0.788**	0.011	**0.819**	0.009
Readability	0.791	0.001	0.791	0.002	**0.798**	0.001	**0.802**	0.002
Satisfaction	0.786	0.014	0.777	0.017	**0.792**	0.008	**0.812**	0.014
Transparency	**0.799**	0.004	0.779	0.011	0.795	0.001	**0.797**	0.016
Total (Avg)	0.830	0.011	0.808	0.016	0.829	0.006	**0.840**	0.017

Results show that professional tone outperform motivational ones across most properties, particularly for Transparency, Informativeness, Faithfulness, and Correlation Ease. In particular, Motivational explanations tone, while often scoring slightly lower in transparency and faithfulness, show competitive or higher scores for Motivation-related properties, especially in datasets with richer learner profiles such as MOOPer. This suggests that motivational framing can enhance engagement-oriented properties, albeit at the cost of stricter grounding. Overall, we uncovered a trade-off between explanatory rigor and motivational expressiveness: professional contexts favoring transparency and faithfulness, while motivational contexts emphasizing engagement-related dimensions.

Table 6.8: RQ2: Professional vs Motivational Explanation Quality. Average and standard deviation across explanation properties for PEEK.

Property	LLaMA3.2				Qwen2.5			
	Professional		Motivational		Professional		Motivational	
	Avg	Std	Avg	Std	Avg	Std	Avg	Std
Confidence Boost	**0.753**	0.014	0.700	0.027	0.717	0.018	**0.743**	0.018
Correlation Ease	**0.783**	0.010	**0.761**	0.027	0.736	0.020	0.753	0.010
Context Relevance	**0.569**	0.013	0.551	0.029	0.545	0.005	**0.583**	0.059
Decision Speed	**0.796**	0.013	0.771	0.005	0.760	0.004	**0.794**	0.025
Effectiveness	**0.578**	0.004	0.553	0.023	0.573	0.005	**0.584**	0.014
Faithfulness	**0.540**	0.003	0.487	0.018	0.532	0.011	**0.542**	0.045
Informativeness	**0.887**	0.012	0.838	0.026	0.861	0.001	**0.875**	0.016
Motivation	**0.721**	0.006	0.722	0.007	0.688	0.012	**0.750**	0.019
Readability	0.782	0.003	0.782	0.009	**0.791**	0.002	**0.799**	0.017
Satisfaction	**0.657**	0.002	0.613	0.032	0.633	0.008	**0.641**	0.013
Transparency	**0.721**	0.005	0.673	0.026	0.717	0.004	**0.708**	0.013
Total (Avg)	<u>0.728</u>	0.008	0.679	0.023	0.714	0.008	**0.729**	0.023

Table 6.9: RQ3: Explanation Quality over Motivational Contexts. Average and standard deviation across explanation properties for COCO.

Property	Formal	L4C	Non-formal	L4P
Confidence Boost	0.802	0.780	0.718	0.750
Correlation Ease	0.795	0.777	0.730	0.769
Context Relevance	0.578	0.590	0.572	0.584
Decision Speed	0.832	0.821	0.799	0.809
Effectiveness	0.692	0.682	0.637	0.667
Faithfulness	0.563	0.568	0.563	0.575
Informativeness	0.914	0.888	0.833	0.870
Motivation	0.802	0.783	0.753	0.778
Readability	0.789	0.786	0.782	0.786
Satisfaction	0.736	0.717	0.665	0.703
Transparency	0.739	0.717	0.676	0.701
Total (Avg)	<u>0.749</u>	0.738	0.703	**0.754**

6.3.3 RQ3: Impact of Learning Contexts on Motivational Tone

In this third and final analysis, we explore how different learning contexts—formal, non-formal, learning-for-credits (L4C), and learning-for-pleasure (L4P) affects motivational explanations for the best-performing model per dataset. Tables 6.9- 6.12 demonstrates results for the top-performing model on each dataset, focusing exclusively on motivational explanations. Across datasets, formal and L4C contexts consistently achieve higher scores for Faithfulness, Transparency, and Effectiveness, while non-formal and L4P contexts tend to favor Motivation and Satisfaction-related properties. For instance, on MOOPer and MOOCCube, formal contexts yield the highest total mean scores, reflecting the benefit of structured learning objectives and clearer learner intent. Conversely, non-formal contexts show improvements in Motivation and Readability, particularly in datasets with diverse learner interactions. While, PEEK exhibits the largest

Table 6.10: RQ3: Explanation Quality over Motivational Contexts. Average and standard deviation across explanation properties for MOOCube.

Property	Formal	Non-formal	L4C	L4P
Confidence Boost	0.804	0.780	0.793	0.761
Correlation Ease	0.858	0.837	0.854	0.835
Context Relevance	0.690	0.756	0.703	0.717
Decision Speed	0.808	0.822	0.803	0.802
Effectiveness	0.653	0.624	0.643	0.641
Faithfulness	0.642	0.688	0.663	0.673
Informativeness	0.921	0.929	0.907	0.899
Motivation	0.757	0.776	0.763	0.760
Readability	0.796	0.813	0.797	0.800
Satisfaction	0.745	0.730	0.740	0.727
Transparency	0.776	0.735	0.761	0.751
Total (Avg)	<u>0.768</u>	**0.772**	0.766	0.761

Table 6.11: RQ3: Explanation Quality over Motivational Contexts. Average and standard deviation across explanation properties for MOOPer.

Property	Formal	Non-formal	L4C	L4P
Confidence Boost	0.896	0.869	0.907	0.879
Correlation Ease	0.916	0.872	0.927	0.905
Context Relevance	0.937	0.881	0.937	0.908
Decision Speed	0.820	0.837	0.831	0.833
Effectiveness	0.731	0.668	0.729	0.742
Faithfulness	0.797	0.785	0.780	0.812
Informativeness	0.981	0.987	0.982	0.967
Motivation	0.813	0.817	0.814	0.831
Readability	0.800	0.804	0.800	0.804
Satisfaction	0.826	0.789	0.815	0.817
Transparency	0.808	0.771	0.804	0.805
Total (Avg)	<u>0.848</u>	0.830	0.848	**0.852**

variability across contexts, with non-formal explanations outperforming formal ones in several properties. This behavior reflects the dataset's heterogeneous nature, where motivational framing may compensate for limited historical grounding. In summary, we found that learning context significantly influences explanation quality, and that no single explanation style is optimal across all scenarios. Instead, context-aware explanation strategies are required to balance faithfulness, transparency, and learner engagement.

6.4 Discussion and Implications

The LLMs can effectively enhance the accessibility and expressiveness of explanations in knowledge-aware educational decision support systems, when their generation process is explicitly grounded in structured reasoning paths. This chapter demonstrates how path-based explanations derived from educational knowledge graphs can be trans-

Table 6.12: RQ3: Explanation Quality over Motivational Contexts. Average and standard deviation across explanation properties for PEEK.

Property	Formal	Non-formal	L4C	L4P
Confidence Boost	0.737	0.748	0.768	0.720
Correlation Ease	0.744	0.757	0.768	0.742
Context Relevance	0.542	0.678	0.552	0.559
Decision Speed	0.769	0.830	0.790	0.786
Effectiveness	0.570	0.603	0.592	0.571
Faithfulness	0.489	0.612	0.528	0.537
Informativeness	0.868	0.901	0.875	0.857
Motivation	0.676	0.699	0.725	0.755
Readability	0.785	0.825	0.794	0.793
Satisfaction	0.631	0.656	0.654	0.624
Transparency	0.714	0.694	0.725	0.698
Total (Avg)	0.684	**0.727**	<u>0.726</u>	0.716

formed into fluent and context-sensitive natural language without sacrificing alignment with the underlying decision logic. The results highlight which explanation properties are most reliably supported by current language models, and how explanation quality varies across models, learning contexts, and tone settings.

However, LLMs improve the linguistic realization of explanations, they do not replace the foundational role of knowledge representation and reasoning. Explanation faithfulness and contextual relevance remain strongly dependent on the quality of the ontology design and the reasoning paths introduced in the earlier chapters. This confirms that explainability in lifelong learning decision support must be addressed holistically, spanning representation, reasoning, interaction, and communication.

With this chapter, the thesis completes its investigation of knowledge-aware methods for explainable decision support in lifelong learning, progressing from learner-centered ontology design *Chapter* 3, to conversational interaction *Chapter* 4, to path-based reasoning *Chapter* 5, and finally to natural language explanation generation. The final chapter synthesizes these contributions, reflects on their broader implications, and outlines future directions toward trustworthy, user-aligned, and explainable AI systems.

Chapter 7

Conclusions

In this thesis, we investigated the design, evaluation, and application of knowledge-aware methods for explainable decision support in lifelong learning. The research focused on how structured knowledge representations, reasoning mechanisms, and large language models can be integrated to support transparent, trustworthy, and user-aligned educational decision-making. To this end, we designed, analyzed, and evaluated a set of methods spanning ontology design, conversational interaction, path-based recommendation, and natural language explanation generation.

7.1 Contribution Summary

The research conducted in this thesis provides the following main contributions:

- We introduced a learner-centered educational ontology (and four resulting knowledge graphs) that explicitly models learners, competencies, objectives, and resources to support explainability by construction. The ontology enables semantically grounded reasoning paths that make outputs transparent at the knowledge representation level.
- We designed a conversational interface for educational knowledge graph exploration, leveraging large language models to mediate between users and structured knowledge. This interface lowers cognitive and technical barriers, enabling interactive and explainable dialogue over complex educational knowledge graphs.
- We investigated path-based recommendation methods for educational decision support, demonstrating how structured reasoning paths can provide faithful and interpretable reasoning for recommendations. These methods establish a clear link between learner history, educational structure, and recommended learning actions.
- We proposed a knowledge-aware natural language explanation generation pipeline that combines path-based reasoning with large language models to produce fluent, context-sensitive natural language explanations. The approach preserves faithfulness to the reasoning logic while improving explanation accessibility and expressiveness.

It is worth noting that the contributions of this thesis collectively show that explainable decision support in lifelong learning cannot be addressed by isolated techniques. Instead, explainability emerges from the integration of knowledge modeling, reasoning, interaction, and communication. While some components were primarily evaluated in controlled experimental settings, the proposed methods are designed with real-world educational platforms in mind, and their limitations are discussed below.

7.2 Limitations and Open Research Issues

While each chapter has provided several insights specific to the knowledge modeling, decision support, and interaction, the contributions presented in this thesis allows us to make also broader conclusions that are commonly shared and, subsequently, derive important limitations that motivate future research along these lines:

- The effectiveness of knowledge-aware decision support systems is strongly dependent on the availability and quality of structured educational data. Although multiple datasets and knowledge graphs were used in this thesis, limitations in data completeness, sparsity, and heterogeneity remain a challenge for large-scale deployment.

- Ontology-driven representations enable explainability by construction, but they require careful design choices and domain expertise. Modeling learner goals, competencies, and educational relations remains a complex process, and suboptimal modeling decisions may limit the interpretability of downstream reasoning and explanations.

- Path-based recommendation methods provide faithful and interpretable reasoning, yet their performance is sensitive to the structure of the underlying knowledge graph. Sparse learner interactions or shallow semantic relations can constrain the relevance of extracted reasoning paths and therefore limit the decision-support effectiveness.

- Large language models improve the accessibility and fluency of explanations, but their behavior remains difficult to fully control. Despite grounding mechanisms, ensuring consistent faithfulness and avoiding subtle deviations from underlying reasoning paths remains an open issue, particularly across diverse datasets and contexts.

- Automatic, property-based evaluation enables scalable assessment of explanation quality, but it cannot fully capture human perception and cognitive impact. The absence of direct user feedback limits the extent to which explanation effectiveness can be assessed from a learner-centered perspective.

- The integration of multiple components (i.e., ontology design, reasoning models, conversational interfaces, and language models) introduces system-level complexity. Coordinating these components while maintaining transparency, robustness, and efficiency remains a non-trivial challenge for real-world educational platforms.

7.3 Future Work Directions

Motivated by the identified limitations, the research presented in this thesis advances knowledge-aware methods for explainable decision support in lifelong learning, but still several research directions remain open and deserve further investigation:

- **Scalable Educational Knowledge Graphs**. While this thesis demonstrates the effectiveness of learner-centered ontologies and educational knowledge graphs, future work should focus on constructing larger, continuously evolving knowledge graphs that integrate data from multiple platforms and institutions. This would enable broader coverage of learning domains and support more diverse learning scenarios.
- **Generalization of Path-Based Reasoning Methods**. Path-based recommendation methods show strong explainability properties but remain sensitive to graph structure and data sparsity. Future research should investigate techniques that improve generalization across datasets, learning contexts, and educational domains, including hybrid reasoning strategies that combine symbolic paths with learned representations.
- **Context-Aware Explanation Strategies**. Results from this thesis indicate that explanation effectiveness varies across explanation tone and contexts. Future work should explore adaptive explanation mechanisms that dynamically adjust explanation style, depth, and framing based on learner goals, context, and interaction history.
- **Human-Centered Evaluation of Explanations**. Although this thesis adopts scalable, offline evaluation methods, future research should complement assessment with user studies involving learners and educators. Such evaluations would provide deeper insights into the long-term learning outcomes associated with different strategies.
- **Ethical and Governance Aspects**. As knowledge-aware decision support systems become increasingly integrated into lifelong learning platforms, future work should address ethical considerations related to fairness, accountability, transparency, and data governance. Developing evaluation frameworks and design principles that operationalize these properties remains a key research challenge.

The rapid evolution of artificial intelligence in education suggests that knowledge-aware and explainable decision support systems will play an increasingly central role in lifelong learning. By extending the methods proposed in this thesis along the directions outlined above, future research can further enhance the reliability, trustworthiness, and societal impact of AI-driven educational systems.

References

[1] World Economic Forum, "The Future of Jobs Report 2025," Jan. 2025. Accessed: 2026-01-16.

[2] G. F. McKenna, "Theory and practice of online learning, 2nd revised edition, terry anderson, fathi elloumi. AU press (2008), ISBN: 1897425082," *Int. J. Inf. Manag.*, vol. 29, no. 6, pp. 503–504, 2009.

[3] B. Means, Y. Toyama, R. Murphy, and M. Baki, "The effectiveness of online and blended learning: A meta-analysis of the empirical literature," *Teachers college record*, vol. 115, no. 3, pp. 1–47, 2013.

[4] AboutChromebooks, "Google for education user statistics," 2026. [Online]. Available: `https://www.aboutchromebooks.com/google-for-education-user-statistics/`. [Accessed: Jan. 8, 2026].

[5] Coursera, "Coursera impact report 2025," 2025. [Online]. Available: `https://www.coursera.org/explore/learner-outcomes`. [Accessed: Jan. 8, 2026].

[6] Udemy, "Udemy annual report 2024," 2024. [Online]. Available: `https://about.udemy.com/press-releases/udemy-annual-learning-skills-trends-report/`. [Accessed: Jan. 8, 2026].

[7] S. Skelcher, D. Yang, J. Trespalacios, and C. Snelson, "Connecting online students to their higher learning institution," *Distance Education*, vol. 41, no. 1, pp. 128–147, 2020.

[8] T. Liu, Q. Wu, L. Chang, and T. Gu, "A review of deep learning-based recommender system in e-learning environments," *Artif. Intell. Rev.*, vol. 55, no. 8, pp. 5953–5980, 2022.

[9] A. Agarwal, D. S. Mishra, and S. V. Kolekar, "Knowledge-based recommendation system using semantic web rules based on learning styles for moocs," *Cogent Engineering*, vol. 9, no. 1, p. 2022568, 2022.

[10] S.-C. Kong, S.-M. Korte, S. Burton, P. Keskitalo, T. Turunen, D. Smith, L. Wang, J. C.-K. Lee, and M. C. Beaton, "Artificial intelligence (ai) literacy–an argument for ai literacy in education," *Innovations in education and teaching international*, vol. 62, no. 2, pp. 477–483, 2025.

[11] S. Mumtaz, J. Carmichael, M. Weiss, and A. Nimon-Peters, "Ethical use of artificial intelligence based tools in higher education: are future business leaders ready?," *Educ. Inf. Technol.*, vol. 30, no. 6, pp. 7293–7319, 2025.

[12] J. Jiao, S. Afroogh, Y. Xu, and C. Phillips, "Navigating LLM ethics: advancements, challenges, and future directions," *AI Ethics*, vol. 5, no. 6, pp. 5795–5819, 2025.

[13] H. Zhu, Y. Sun, and J. Yang, "Towards responsible artificial intelligence in education: a systematic review on identifying and mitigating ethical risks," *Humanities and Social Sciences Communications*, vol. 12, no. 1, pp. 1–14, 2025.

[14] N. Afreen, G. Balloccu, L. Boratto, G. Fenu, F. M. Malloci, M. Marras, and A. G. Martis, "Can path-based explainable recommendation methods based on knowledge graphs generalize for personalized education?," in *Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*, pp. 273–278, 2025.

[15] S. Sharma, P. Mittal, M. Kumar, and V. Bhardwaj, "The role of large language models in personalized learning: a systematic review of educational impact," *Discover Sustainability*, vol. 6, no. 1, pp. 1–24, 2025.

[16] A. L. Paul, "Improving explainability in large language models (llms)," 2025.

[17] M. Giannakos, R. Azevedo, P. Brusilovsky, M. Cukurova, Y. A. Dimitriadis, D. H. Leo, S. Järvelä, M. Mavrikis, and B. Rienties, "The promise and challenges of generative AI in education," *Behav. Inf. Technol.*, vol. 44, no. 11, pp. 2518–2544, 2025.

[18] I. Jurenka, M. Kunesch, K. R. McKee, D. Gillick, S. Zhu, S. Wiltberger, S. M. Phal, K. Hermann, D. Kasenberg, A. Bhoopchand, *et al.*, "Towards responsible development of generative ai for education: An evaluation-driven approach," *arXiv preprint arXiv:2407.12687*, 2024.

[19] K. Qu, K. C. Li, B. T. Wong, M. M. Wu, and M. Liu, "A survey of knowledge graph approaches and applications in education," *Electronics*, vol. 13, no. 13, p. 2537, 2024.

[20] S. Zhao, "Personalized learning path recommendation strategy based on subject knowledge graph," in *Proceedings of the 2024 3rd International Conference on Artificial Intelligence and Education*, pp. 398–402, 2024.

[21] W. Wu, S. Yang, F. Tian, X. Wei, and Y. Pan, "Knowledge graph enhanced recommendation for semantic structure construction and feedback in online learning," *Physical Communication*, vol. 69, p. 102590, 2025.

[22] A. Hanif, X. Zhang, and S. Wood, "A survey on explainable artificial intelligence techniques and challenges," in *2021 IEEE 25th International Enterprise Distributed Object Computing Workshop (EDOCW)*, pp. 81–89, 2021.

[23] F. G. Febrinanto, F. Xia, K. Moore, C. Thapa, and C. Aggarwal, "Graph lifelong learning: A survey," *IEEE Computational Intelligence Magazine*, vol. 18, no. 1, pp. 32–51, 2023.

[24] B. Zhao, J. Sun, B. Xu, X. Lu, Y. Li, J. Yu, M. Liu, T. Zhang, Q. Chen, H. Li, L. Hou, and J. Li, "EDUKG: a heterogeneous sustainable K-12 educational knowledge graph," *CoRR*, vol. abs/2210.12228, 2022.

[25] A. Li, Y. Li, and X. Gao, "Personalized learning path recommendation based on knowledge graphs: A survey," *Electronics*, vol. 15, no. 1, p. 238, 2026.

[26] H. Zhou, R. Bamler, C. M. Wu, and Tejero-Cantero, "Predictive, scalable and interpretable knowledge tracing on structured domains," 2024. [Online]. Available: `https://arxiv.org/abs/2403.13179`.

[27] H. Abu-Rasheed, C. Weber, and M. Fathi, "Knowledge graphs as context sources for llm-based explanations of learning recommendations," in *2024 IEEE Global Engineering Education Conference (EDUCON)*, pp. 1–5, IEEE, 2024.

[28] G. Jhajj, X. Zhang, J. R. Gustafson, F. Lin, and M. P.-C. Lin, "Educational knowledge graph creation and augmentation via llms," in *Generative Intelligence and Intelligent Tutoring Systems: 20th International Conference, ITS 2024, Thessaloniki, Greece, June 10–13, 2024, Proceedings, Part II*, (Berlin, Heidelberg), p. 292–304, Springer-Verlag, 2024.

[29] P. Chen, Y. Lu, V. W. Zheng, X. Chen, and X. Li, "An automatic knowledge graph construction system for k-12 education," in *Proceedings of the fifth annual ACM conference on learning at scale*, pp. 1–4, 2018.

[30] F. Dang, J. Tang, K. Pang, T. Wang, S. Li, and X. Li, "Constructing an educational knowledge graph with concepts linked to wikipedia," *J. Comput. Sci. Technol.*, vol. 36, no. 5, pp. 1200–1211, 2021.

[31] P. Chen, Y. Lu, V. W. Zheng, X. Chen, and B. Yang, "Knowedu: A system to construct knowledge graph for education," *IEEE Access*, vol. 6, pp. 31553–31563, 2018.

[32] F. Biessmann and V. Treu, "A turing test for transparency," *arXiv preprint arXiv:2106.11394*, 2021.

[33] Y. Li, K. Liu, R. Satapathy, S. Wang, and E. Cambria, "Recent developments in recommender systems: A survey," *IEEE Computational Intelligence Magazine*, vol. 19, no. 2, pp. 78–95, 2024.

[34] Q. Guan, X. Cheng, F. Xiao, Z. Li, C. He, L. Fang, G. Chen, Z. Gong, and W. Luo, "Explainable exercise recommendation with knowledge graph," *Neural Networks*, vol. 183, p. 106954, 2025.

[35] M. C. Urdaneta-Ponte, A. Mendez-Zorrilla, and I. Oleagordia-Ruiz, "Recommendation systems for education: Systematic review," *Electronics*, vol. 10, no. 14, p. 1611, 2021.

[36] Q. Bin, M. F. Zuhairi, and J. Morcos, "A comprehensive study on personalized learning recommendation in e-learning system," *IEEE Access*, vol. 12, pp. 100446–100482, 2024.

[37] J. Bobadilla, F. Ortega, A. Hernando, and A. Gutiérrez, "Recommender systems survey," *Knowledge-Based Systems*, vol. 46, pp. 109–132, 2013.

[38] H. Drachsler, K. Verbert, O. C. Santos, and N. Manouselis, "Learning with personalized recommender systems: A psychological view," *Computers in Human Behavior*, vol. 28, no. 1, pp. 207–216, 2012.

[39] S. Reddy, I. Labutov, and T. Joachims, "Latent skill embedding for personalized lesson sequence recommendation," *arXiv preprint arXiv:1602.07029*, 2016.

[40] X. Chen, J. Shen, W. Xia, J. Jin, Y. Song, W. Zhang, W. Liu, M. Zhu, R. Tang, K. Dong, D. Xia, and Y. Yu, "Set-to-sequence ranking-based concept-aware learning path recommendation," in *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023* (B. Williams, Y. Chen, and J. Neville, eds.), pp. 5027–5035, AAAI Press, 2023.

[41] Y. Zhou, H. Wang, J. He, and H. Wang, "From intrinsic to counterfactual: On the explainability of contextualized recommender systems," *arXiv preprint arXiv:2110.14844*, 2021.

[42] Q. Guo, F. Zhuang, C. Qin, H. Zhu, X. Xie, H. Xiong, and Q. He, "A survey on knowledge graph-based recommender systems," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 8, pp. 3549–3568, 2020.

[43] Y. Gao, Y. Li, Y. Lin, H. Gao, and L. Khan, "Deep learning on knowledge graph for recommender system: A survey," *CoRR*, vol. abs/2004.00387, 2020.

[44] K. Qu, K. C. Li, B. T. M. Wong, M. M. F. Wu, and M. Liu, "A survey of knowledge graph approaches and applications in education," *Electronics*, vol. 13, no. 13, p. 2537, 2024.

[45] X. Wang, X. He, Y. Cao, M. Liu, and T.-S. Chua, "Kgat: Knowledge graph attention network for recommendation," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, pp. 950–958, 2019.

[46] H. Wang, M. Zhao, X. Xie, W. Li, and M. Guo, "Knowledge graph convolutional networks for recommender systems," in *The world wide web conference*, pp. 3307–3313, 2019.

[47] e. Li, "Bilateral knowledge graph enhanced online course recommendation," *Information Systems*, vol. 107, p. 102000, 2022.

[48] Y. Xian, Z. Fu, S. Muthukrishnan, G. de Melo, and Y. Zhang, "Reinforcement knowledge graph reasoning for explainable recommendation," in *Proc. of SIGIR 2019*, pp. 285–294, ACM, 2019.

[49] Y. Xian, Z. Fu, H. Zhao, Y. Ge, X. Chen, Q. Huang, S. Geng, Z. Qin, G. de Melo, S. Muthukrishnan, and Y. Zhang, "CAFE: coarse-to-fine neural symbolic reasoning for explainable recommendation," in *Proc. of CIKM 2020*, pp. 1645–1654, ACM, 2020.

[50] C.-Y. Tai, L.-Y. Huang, C.-K. Huang, and L.-W. Ku, "User-centric path reasoning towards explainable recommendation," in *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '21, (New York, NY, USA), p. 879–889, Association for Computing Machinery, 2021.

[51] X. Qin, "Kg-xqrec: Knowledge graph-based explainable question recommendation," in *Proceedings of the 2025 International Conference on Artificial Intelligence and Educational Systems*, pp. 433–437, 2025.

[52] Y. Feldman-Maggor, M. Cukurova, C. Kent, and G. Alexandron, "The impact of explainable AI on teachers' trust and acceptance of AI edtech recommendations: The power of domain-specific explanations," *Int. J. Artif. Intell. Educ.*, vol. 35, no. 5, pp. 2889–2922, 2025.

[53] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, 2019.

[54] N. Tintarev and J. Masthoff, "Evaluating the effectiveness of explanations for recommender systems," *User Modeling and User-Adapted Interaction*, vol. 22, pp. 399–439, 2012.

[55] A. Samih, A. Adadi, and M. Berrada, "Towards a knowledge based explainable recommender systems," in *BDIoT'19: The 4th International Conference On Big Data and Internet of Things, Rabat, Morocco, October 23-24, 2019* (M. Lazaar, C. Duvallet, M. A. Achhab, O. Mahboub, and H. Silkan, eds.), pp. 21:1–21:5, ACM, 2019.

[56] J. Barria-Pineda, K. Akhuseyinoglu, S. Želem undefinedelap, P. Brusilovsky, A. K. Milicevic, and M. Ivanovic, "Explainable recommendations in a personalized programming practice system," in *Artificial Intelligence in Education: 22nd International Conference, AIED 2021, Utrecht, The Netherlands, June 14–18, 2021, Proceedings, Part I*, (Berlin, Heidelberg), p. 64–76, Springer-Verlag, 2021.

[57] J. Barria-Pineda and P. Brusilovsky, "Explaining educational recommendations through a concept-level knowledge visualization," in *Companion Proceedings of the 24th International Conference on Intelligent User Interfaces*, IUI '19 Companion, (New York, NY, USA), p. 103–104, Association for Computing Machinery, 2019.

[58] W. Song, Z. Duan, Z. Yang, H. Zhu, M. Zhang, and J. Tang, "Ekar: an explainable method for knowledge aware recommendation," *arXiv preprint arXiv:1906.09506*, 2019.

[59] A. Mangold, J. Zietz, S. Weinhold, and S. Pannasch, "On the design and evaluation of human-centered explainable AI systems: A systematic review and taxonomy," *CoRR*, vol. abs/2510.12201, 2025.

[60] G. Cabour, A. Morales, É. Ledoux, and S. Bassetto, "Towards an explanation space to align humans and explainable-ai teamwork," *CoRR*, vol. abs/2106.01503, 2021.

[61] S. Hong and W. Park, "Developing user-centered system design guidelines for explainable ai: a systematic literature review," *Artificial Intelligence Review*, vol. 58, p. 386, 2025.

[62] A. Wijekoon, N. Wiratunga, D. Corsar, K. Martin, I. Nkisi-Orji, B. Díaz-Agudo, and D. Bridge, "Xeq scale for evaluating xai experience quality," *arXiv preprint*, 2024.

[63] F. M. Cau and L. D. Spano, "Exploring the impact of explainable ai and cognitive capabilities on users' decisions," *User Modeling and User-Adapted Interaction*, vol. 36, no. 3, 2026.

[64] W. Jin, J. Fan, D. Gromala, P. Pasquier, X. Li, and G. Hamarneh, "Transcending XAI algorithm boundaries through end-user-inspired design," *CoRR*, vol. abs/2208.08739, 2022.

[65] T. Laugel, M.-J. Lesot, C. Marsala, X. Renard, and M. Detyniecki, "The dangers of post-hoc interpretability: Unjustified counterfactual explanations," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2801–2807, 2019.

[66] W. J. Yeo, R. Satapathy, and E. Cambria, "Towards faithful natural language explanations: A study using activation patching in large language models," in *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, EMNLP 2025, Suzhou, China, November 4-9, 2025* (C. Christodoulopoulos, T. Chakraborty, C. Rose, and V. Peng, eds.), pp. 10425–10447, Association for Computational Linguistics, 2025.

[67] A. Madsen, S. Chandar, and S. Reddy, "Are self-explanations from large language models faithful?," in *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024* (L. Ku, A. Martins, and V. Srikumar, eds.), Findings of ACL, pp. 295–337, Association for Computational Linguistics, 2024.

[68] K. Randl, J. Pavlopoulos, A. Henriksson, and T. Lindgren, "Mind the gap: from plausible to valid self-explanations in large language models," *Machine Learning*, vol. 114, p. 220, 2025.

[69] C. Agarwal, S. H. Tanneru, and H. Lakkaraju, "Faithfulness vs. plausibility: On the (un) reliability of explanations from large language models, 2024," *URL https://arxiv.org/abs/2402.04614*, 2024.

[70] H. Wang and K. Shu, "Explainable claim verification via knowledge-grounded reasoning with large language models," *arXiv preprint arXiv:2310.05253*, 2023.

[71] Z. Chen, J. Chen, Y. Chen, H. Yu, A. K. Singh, and M. Sra, "Lmexplainer: Grounding knowledge and explaining language models," *arXiv preprint arXiv:2303.16537*, 2023.

[72] H. Bian, "Llm-empowered knowledge graph construction: A survey," *arXiv preprint arXiv:2510.20345*, 2025.

[73] F. L. da Silva, B. K. Slodkowski, K. K. A. da Silva, and S. C. Cazella, "A systematic literature review on educational recommender systems for teaching and learning: research trends, limitations and opportunities," *Educ. Inf. Technol.*, vol. 28, no. 3, pp. 3289–3328, 2023.

[74] J. Barria-Pineda, "Exploring the need for transparency in educational recommender systems," in *Proc. of ACM UMAP 2020*, pp. 376–379, ACM, 2020.

[75] Y. Zhang and X. Chen, "Explainable recommendation: A survey and new perspectives," *Found. Trends Inf. Retr.*, vol. 14, no. 1, pp. 1–101, 2020.

[76] N. Tintarev and J. Masthoff, "A survey of explanations in recommender systems," in *Proc. of ICDE 2007*, pp. 801–810, IEEE Computer Society, 2007.

[77] B. Goodman and S. R. Flaxman, "European union regulations on algorithmic decision-making and a "right to explanation"," *AI Mag.*, vol. 38, no. 3, pp. 50–57, 2017.

[78] T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," *Artif. Intell.*, vol. 267, pp. 1–38, 2019.

[79] L. Ehrlinger and W. Wöß, "Towards a definition of knowledge graphs," in *Proc. of SEMANTiCS 2016*, vol. 1695 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2016.

[80] D. Liu, J. Lian, Z. Liu, X. Wang, G. Sun, and X. Xie, "Reinforced anchor knowledge graph generation for news recommendation reasoning," in *Proc. of KDD 2021*, pp. 1055–1065, ACM, 2021.

[81] G. Balloccu, L. Boratto, G. Fenu, and M. Marras, "Reinforcement recommendation reasoning through knowledge graphs for explanation path quality," *Knowl. Based Syst.*, vol. 260, p. 110098, 2023.

[82] S. Geng, Z. Fu, J. Tan, Y. Ge, G. de Melo, and Y. Zhang, "Path language modeling over knowledge graphsfor explainable recommendation," in *Proc. of WWW 2022*, pp. 946–955, ACM, 2022.

[83] G. Balloccu, L. Boratto, C. Cancedda, G. Fenu, and M. Marras, "Faithful path language modelling for explainable recommendation over knowledge graph," *CoRR*, vol. abs/2310.16452, 2023.

[84] N. Guarino, D. Oberle, and S. Staab, "What is an ontology?," in *Handbook on Ontologies*, International Handbooks on Information Systems, pp. 1–17, Springer, 2009.

[85] S. Ji, S. Pan, E. Cambria, P. Marttinen, and P. S. Yu, "A survey on knowledge graphs: Representation, acquisition, and applications," *IEEE Trans. Neural Networks Learn. Syst.*, vol. 33, no. 2, pp. 494–514, 2022.

[86] J. Frej, N. Shah, M. Knezevic, T. Nazaretsky, and T. Käser, "Finding paths for explainable MOOC recommendation: A learner perspective," *CoRR*, vol. abs/2312.10082, 2023.

[87] L. Wu, Q. Liu, W. Zhou, G. Mao, J. Huang, and H. Huang, "A semantic web-based recommendation framework of educational resources in e-learning," *Technol. Knowl. Learn.*, vol. 25, no. 4, pp. 811–833, 2020.

[88] E. Ilkou, H. Abu-Rasheed, M. Tavakoli, S. Hakimov, G. Kismihók, S. Auer, and W. Nejdl, "Educor: An educational and career-oriented recommendation ontology," in *Proc. of ISWC 2021*, vol. 12922, pp. 546–562, Springer, 2021.

[89] N. W. Rahayu, R. Ferdiana, and S. S. Kusumawardani, "A systematic review of ontology use in e-learning recommender system," *Com. Ed. Art. Int.*, vol. 3, p. 100047, 2022.

[90] K. Stancin, P. Poscic, and D. Jaksic, "Ontologies in education - state of the art," *Educ. Inf. Technol.*, vol. 25, no. 6, pp. 5301–5320, 2020.

[91] H. X. Li, G. Appleby, C. D. Brumar, R. Chang, and A. Suh, "Knowledge graphs in practice: Characterizing their users, challenges, and visualization opportunities," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 1, pp. 584–594, 2024.

[92] B. A. Chaushi, B. Selimi, A. Chaushi, and M. Apostolova, "Explainable artificial intelligence in education: A comprehensive review," in *Proc. of xAI 2023*, vol. 1902, pp. 48–71, Springer, 2023.

[93] D. Dessì, G. Fenu, M. Marras, and D. R. Recupero, "COCO: semantic-enriched collection of online courses at scale with experimental use cases," in *Proc. of the World Conference on Information Systems and Technologies (WorldCIST)*, vol. 746, pp. 1386–1396, 2018.

[94] J. Yu, G. Luo, T. Xiao, Q. Zhong, Y. Wang, W. Feng, J. Luo, C. Wang, L. Hou, J. Li, *et al.*, "Mooccube: a large-scale data repository for nlp applications in moocs," in *Proc. of ACL 2020*, pp. 3135–3142, 2020.

[95] K. Liu, X. Zhao, J. Tang, W. Zeng, J. Liao, F. Tian, Q. Zheng, J. Huang, and A. Dai, "Mooper: a large-scale dataset of practice-oriented online learning," in *Proc. of the China Conference on Knowledge Graph and Semantic Computing: Knowledge Graph Empowers New Infrastructure Construction (CCKS 2021)*, pp. 281–287, Springer, 2021.

[96] S. Bulathwela, M. Pérez-Ortiz, E. Novak, E. Yilmaz, and J. Shawe-Taylor, "PEEK: A large dataset of learner engagement with educational videos," *CoRR*, vol. abs/2109.03154, 2021.

[97] A. Iliadis, A. Acker, W. Stevens, and S. B. Kavakli, "One schema to rule them all: How schema. org models the world of search," *Journal of the Association for Information Science and Technology*, 2022.

[98] N. Hubert, A. Brun, and D. Monticolo, "New ontology and knowledge graph for university curriculum recommendation," in *Proceedings of the ISWC 2022 Posters, Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 21st International Semantic Web Conference (ISWC 2022), Virtual Conference, Hangzhou, China, October 23-27, 2022* (A. Dimou, A. Haller, A. L. Gentile, and P. Ristoski, eds.), CEUR Workshop Proceedings, CEUR-WS.org, 2022.

[99] N. Tintarev and J. Masthoff, "A survey of explanations in recommender systems," in *Proc. of ICDEW 2007*, pp. 801–810, IEEE, 2007.

[100] A. Haleem, M. Javaid, M. A. Qadri, and R. Suman, "Understanding the role of digital technologies in education: A review," *Sustainable Operations and Computers*, vol. 3, pp. 275–285, 2022.

[101] D. Dessì, G. Fenu, M. Marras, and D. R. Recupero, "Leveraging cognitive computing for multi-class classification of e-learning videos," in *Proc. of The Semantic Web Satellite Events (ESWC 2017)*, vol. 10577, pp. 21–25, Springer, 2017.

[102] A. A. Abdulquadir, O. Abdulsalam, R. Mcmahon, and A. Prabhakar, "Usability study of a learning management system (lms)," in *Proc. of the Annual Conf. on Info. Tech. Ed. (ITE 2023)*, pp. 150–156, 2023.

[103] G. Fenu and M. Marras, "Controlling user access to cloud-connected mobile applications by means of biometrics," *IEEE Cloud Comput.*, vol. 5, no. 4, pp. 47–57, 2018.

[104] I. P. G. P. Duta, M. R. Febriansyah, M. S. Anggreainy, *et al.*, "Effectiveness of lms in online learning by analyzing its usability and features," in *Proc. of the International Conference on Computer Science and Artificial Intelligence (ICCSAI 2021)*, vol. 1, pp. 56–61, IEEE, 2021.

[105] O. Talbi and A. Ouared, "Goal-oriented student motivation in learning analytics: How can a requirements-driven approach help?," *Education and Information Technologies*, vol. 27, no. 9, pp. 12083–12121, 2022.

[106] I. Maslov and S. Nikou, "Usability and ux of learning management systems: an eye-tracking approach," in *Proc. of the IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC 2020)*, pp. 1–9, IEEE, 2020.

[107] P. Mejia-Domenzain, M. Marras, C. Giang, and T. Käser, "Identifying and comparing multi-dimensional student profiles across flipped classrooms," in *Proc. of the 23rd International Conference on Artificial Intelligence in Education (AIED 2022)*, vol. 13355, pp. 90–102, Springer, 2022.

[108] K. Wiley, Y. Dimitriadis, and M. Linn, "A human-centred learning analytics approach for developing contextually scalable k-12 teacher dashboards," *British Journal of Educational Technology*, vol. 55, no. 3, pp. 845–885, 2024.

[109] V. Swamy, S. Du, M. Marras, and T. Käser, "Trusting the explainers: Teacher validation of explainable artificial intelligence for course design," in *Proc. of the 13th International Learning Analytics and Knowledge Conference (LAK 2023)*, pp. 345–356, ACM, 2023.

[110] D. Ifenthaler, D. Gibson, D. Prasse, A. Shimada, and M. Yamada, "Putting learning back into learning analytics: Actions for policy makers, researchers, and practitioners," *Educational Technology Research and Development*, vol. 69, pp. 2131–2150, 2021.

[111] D. Dessì, M. Dragoni, G. Fenu, M. Marras, and D. R. Recupero, "Evaluating neural word embeddings created from online course reviews for sentiment analysis," in *Proc. of the 34th ACM/SIGAPP Symposium on Applied Computing (SAC 2019)*, pp. 2124–2127, ACM, 2019.

[112] K. S. Kaswan, J. S. Dhatterwal, and R. P. Ojha, "Ai in personalized learning," in *Advances in Technological Innovations in Higher Education*, pp. 103–117, CRC Press, 2024.

[113] A. Lieb and T. Goel, "Student interaction with newtbot: An llm-as-tutor chatbot for secondary physics education," in *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pp. 1–8, 2024.

[114] A. Ifelebuegu, "Rethinking online assessment strategies: Authenticity versus ai chatbot intervention," *J. of App. Learn. and Teach.*, vol. 6, no. 2, 2023.

[115] M. S. Ali, F. Azam, A. Safdar, and M. W. Anwar, "Intelligent agents in educational institutions: Nedbot-nlp-based chatbot for administrative support using dialogflow," in *Proc. of the IEEE International Conference on Agents (ICA 2022)*, pp. 30–35, IEEE, 2022.

[116] L. N. Paschoal, A. L. Krassmann, F. B. Nunes, M. M. de Oliveira, M. Bercht, E. F. Barbosa, and S. d. R. S. de Souza, "A systematic identification of pedagogical conversational agents," in *Proc. of Frontiers in Ed. Conf. (FIE 2020)*, pp. 1–9, IEEE, 2020.

[117] J. Qadir, "Engineering education in the era of chatgpt: Promise and pitfalls of generative ai for education," in *Proc. of the IEEE Global Engineering Education Conference (EDUCON 2023)*, pp. 1–9, IEEE, 2023.

[118] X. Zou, "A survey on application of knowledge graph," in *Journal of Physics: Conference Series*, vol. 1487, p. 012016, IOP Publishing, 2020.

[119] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, and X. Wu, "Unifying large language models and knowledge graphs: A roadmap," *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 7, pp. 3580–3599, 2024.

[120] F. Huang, Y. Deng, C. Zhang, M. Guo, K. Zhan, S. Sun, J. Jiang, Z. Sun, and X. Wu, "KOSA: KO enhanced salary analytics based on knowledge graph and LLM capabilities," in *Proc. of the IEEE International Conference on Data Mining (ICDM 2023)*, pp. 499–505, IEEE, 2023.

[121] D. Edge, H. Trinh, N. Cheng, J. Bradley, A. Chao, A. Mody, S. Truitt, and J. Larson, "From local to global: A graph RAG approach to query-focused summarization," *CoRR*, vol. abs/2404.16130, 2024.

[122] N. Afreen, G. Balloccu, L. Boratto, G. Fenu, F. M. Malloci, M. Marras, and A. G. Martis, "Learner-centered ontology for explainable educational recommendation," in *Proc. of the ACM Conf. on User Modeling, Adaptation and Personalization (UMAP 2024)*, pp. 1–4, ACM, 2024.

[123] N. Hubert, A. Brun, and D. Monticolo, "New ontology and knowledge graph for university curriculum recommendation," in *Proc. of the International Semantic Web Conference (ISWC 2022)*, 2022.

[124] G. Balloccu, L. Boratto, G. Fenu, and M. Marras, "Post processing recommender systems with knowledge graphs for recency, popularity, and diversity of explanations," in *Proc. of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)* (E. Amigó, P. Castells, J. Gonzalo, B. Carterette, J. S. Culpepper, and G. Kazai, eds.), pp. 646–656, 2022.

[125] G. Balloccu, L. Boratto, C. Cancedda, G. Fenu, and M. Marras, "Knowledge is power, understanding is impact: Utility and beyond goals, explanation quality, and fairness in path reasoning recommendation," in *Prco. of the European Conference on Information Retrieval, (ECIR)*, vol. 13982, pp. 3–19, 2023.

[126] Y. Yang, C. Huang, L. Xia, and C. Li, "Knowledge graph contrastive learning for recommendation," in *Proc. of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pp. 1434–1443, 2022.

[127] M. Marras and G. Fenu, "Characterizing learner's comments and rating behavior in online course platforms at scale," in *Advances in Software Engineering, Education, and e-Learning: Proc. from FECS'20, FCS'20, SERP'20, and EEE'20*, pp. 781–791, 2021.

[128] Y. Lin, F. Lin, W. Zeng, J. Xiahou, L. Li, P. Wu, Y. Liu, and C. Miao, "Hierarchical reinforcement learning with dynamic recurrent mechanism for course recommendation," *Knowl. Based Syst.*, vol. 244, p. 108546, 2022.

[129] Y. Lin, S. Feng, F. Lin, J. Xiahou, and W. Zeng, "Multi-scale reinforced profile for personalized recommendation with deep neural networks in moocs," *Appl. Soft Comput.*, vol. 148, p. 110905, 2023.

[130] Z. Liang, L. Mu, J. Chen, and Q. Xie, "Graph path fusion and reinforcement reasoning for recommendation in moocs," *Educ. Inf. Technol.*, vol. 28, no. 1, pp. 525–545, 2023.

[131] Y. Lin, W. Zhang, F. Lin, W. Zeng, X. Zhou, and P. Wu, "Knowledge-aware reasoning with self-supervised reinforcement learning for explainable recommendation in moocs," *Neural Comput. Appl.*, vol. 36, no. 8, pp. 4115–4132, 2024.

[132] J. Frej, N. Shah, M. Knezevic, T. Nazaretsky, and T. Käser, "Finding paths for explainable MOOC recommendation: A learner perspective," in *Proc. of the 14th Learning Analytics and Knowledge Conference (LAK)*, pp. 426–437, 2024.

[133] N. Afreen, G. Balloccu, L. Boratto, G. Fenu, and M. Marras, "Towards explainable educational recommendation through path reasoning methods," in *Proc. of the Italian Information Retrieval Workshop (IIR)*, vol. 3448, pp. 131–136, CEUR-WS.org, 2023.

[134] S. S. Khanal, P. W. C. Prasad, A. Alsadoon, and A. Maag, "A systematic review: machine learning based recommendation systems for e-learning," *Educ. Inf. Technol.*, vol. 25, no. 4, pp. 2635–2664, 2020.

[135] H. Ko, S. Lee, Y. Park, and A. Choi, "A survey of recommendation systems: recommendation models, techniques, and application fields," *Electronics*, vol. 11, no. 1, p. 141, 2022.

[136] J. Yu, G. Luo, T. Xiao, Q. Zhong, Y. Wang, W. Feng, J. Luo, C. Wang, L. Hou, J. Li, Z. Liu, and J. Tang, "Mooccube: A large-scale data repository for NLP applications in moocs," in *Proc. of the Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, pp. 3135–3142, 2020.

[137] S. Bulathwela, M. Perez-Ortiz, E. Novak, E. Yilmaz, and J. Shawe-Taylor, "Peek: A large dataset of learner engagement with educational videos," *arXiv preprint arXiv:2109.03154*, 2021.

[138] W. Yin, J. Hay, and D. Roth, "Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach," in *Proc. of the Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 3912–3921, 2019.

[139] A. Bordes, N. Usunier, A. García-Durán, J. Weston, and O. Yakhnenko, "Translating embeddings for modeling multi-relational data," in *Proc. of the Annual Conference on Neural Information Processing Systems (NIPS)*, pp. 2787–2795, 2013.

[140] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. Riedl, "Evaluating collaborative filtering recommender systems," *ACM Trans. Inf. Syst.*, vol. 22, no. 1, pp. 5–53, 2004.

[141] C. Chen, M. Zhang, Y. Zhang, Y. Liu, and S. Ma, "Efficient neural matrix factorization without sampling for recommendation," *ACM Trans. Inf. Syst.*, vol. 38, no. 2, pp. 14:1–14:28, 2020.

[142] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme, "BPR: bayesian personalized ranking from implicit feedback," in *Proc. of the Conference on Uncertainty in Artificial Intelligence (UAI)* (J. A. Bilmes and A. Y. Ng, eds.), pp. 452–461, 2009.

[143] X. Wang, X. He, M. Wang, F. Feng, and T.-S. Chua, "Neural graph collaborative filtering," in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR'19, (New York, NY, USA), p. 165–174, Association for Computing Machinery, 2019.

[144] X. Wang, H. Jin, A. Zhang, X. He, T. Xu, and T.-S. Chua, "Disentangled graph collaborative filtering," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '20, (New York, NY, USA), p. 1001–1010, Association for Computing Machinery, 2020.

[145] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang, "Lightgcn: Simplifying and powering graph convolution network for recommendation," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '20, (New York, NY, USA), p. 639–648, Association for Computing Machinery, 2020.

[146] J. Ma, C. Zhou, P. Cui, H. Yang, and W. Zhu, *Learning disentangled representations for recommendation.* Red Hook, NY, USA: Curran Associates Inc., 2019.

[147] Q. Ai, V. Azizi, X. Chen, and Y. Zhang, "Learning heterogeneous knowledge base embeddings for explainable recommendation," *Algorithms*, vol. 11, no. 9, p. 137, 2018.

[148] F. Zhang, N. J. Yuan, D. Lian, X. Xie, and W.-Y. Ma, "Collaborative knowledge base embedding for recommender systems," in *Proc. of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 353–362, 2016.

[149] H. Wang, F. Zhang, M. Zhao, W. Li, X. Xie, and M. Guo, "Multi-task feature learning for knowledge graph enhanced recommendation," in *The World Wide Web Conference*, WWW '19, (New York, NY, USA), p. 2000–2010, Association for Computing Machinery, 2019.

[150] Y. Cao, X. Wang, X. He, Z. Hu, and T.-S. Chua, "Unifying knowledge graph learning and recommendation: Towards a better understanding of user preferences," in *The World Wide Web Conference*, WWW '19, (New York, NY, USA), p. 151–161, Association for Computing Machinery, 2019.

[151] H. Wang, F. Zhang, J. Wang, M. Zhao, W. Li, X. Xie, and M. Guo, "Ripplenet: Propagating user preferences on the knowledge graph for recommender systems," in *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, CIKM '18, (New York, NY, USA), p. 417–426, Association for Computing Machinery, 2018.

[152] H. Wang, M. Zhao, X. Xie, W. Li, and M. Guo, "Knowledge graph convolutional networks for recommender systems," in *The World Wide Web Conference*, WWW '19, (New York, NY, USA), p. 3307–3313, Association for Computing Machinery, 2019.

[153] H. Wang, F. Zhang, M. Zhang, J. Leskovec, M. Zhao, W. Li, and Z. Wang, "Knowledge-aware graph neural networks with label smoothness regularization for recommender systems," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, (New York, NY, USA), p. 968–977, Association for Computing Machinery, 2019.

[154] Y. Wang, L. Wang, Y. Li, D. He, and T. Liu, "A theoretical analysis of NDCG type ranking measures," in *Proc. of the Annual Conference on Learning Theory (COLT)* (S. Shalev-Shwartz and I. Steinwart, eds.), vol. 30, pp. 25–54, 2013.

[155] N. Craswell, *Mean Reciprocal Rank*, pp. 1703–1703. Boston, MA: Springer US, 2009.

[156] A. Bellogín, P. Castells, and I. Cantador, "Precision-oriented evaluation of recommender systems: an algorithmic comparison," in *Proc. of the ACM Conference on Recommender Systems (RecSys)* (B. Mobasher, R. D. Burke, D. Jannach, and G. Adomavicius, eds.), pp. 333–336, 2011.

[157] M. Kaminskas and D. Bridge, "Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems," *ACM Trans. Interact. Intell. Syst.*, vol. 7, no. 1, pp. 2:1–2:42, 2017.

[158] J. D. Pinto, L. Paquette, V. Swamy, T. Käser, Q. Liu, and L. Cohausz, "Preface to the proceedings of the human-centric explainable ai in education workshop (hexed 2024)," in *Joint Proceedings of the Human-Centric eXplainable AI in Education and the Leveraging Large Language Models for Next Generation Educational Technologies Workshops (HEXED-L3MNGET 2024)*, vol. 3840, 2024.

[159] H. Abu-Rasheed, C. Weber, and M. Fathi, "Experimental interface for multimodal and large language model based explanations of educational recommender systems," *CoRR*, vol. abs/2402.07910, 2024.

[160] X. Zhang, Y. Li, J. Wang, B. Sun, W. Ma, P. Sun, and M. Zhang, "Large language models as evaluators for recommendation explanations," in *Proceedings of the 18th ACM Conference on Recommender Systems, RecSys 2024, Bari, Italy, October 14-18, 2024* (T. D. Noia, P. Lops, T. Joachims, K. Verbert, P. Castells, Z. Dong, and B. London, eds.), pp. 33–42, ACM, 2024.

[161] K. Matton, R. O. Ness, J. V. Guttag, and E. Kiciman, "Walk the talk? measuring the faithfulness of large language model explanations," *CoRR*, vol. abs/2504.14150, 2025.

[162] K. Randl, J. Pavlopoulos, A. Henriksson, and T. Lindgren, "Mind the gap: from plausible to valid self-explanations in large language models," *Mach. Learn.*, vol. 114, no. 10, p. 220, 2025.

[163] S. Lubos, T. N. T. Tran, A. Felfernig, S. Polat Erdeniz, and V.-M. Le, "Llm-generated explanations for recommender systems," in *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*, pp. 276–285, 2024.

[164] V. Adlakha, P. BehnamGhader, X. H. Lù, N. Meade, and S. Reddy, "Evaluating correctness and faithfulness of instruction-following models for question answering," *Trans. Assoc. Comput. Linguistics*, vol. 12, pp. 681–699, 2024.

[165] G. Balloccu, L. Boratto, C. Cancedda, G. Fenu, and M. Marras, "Faithful path language modelling for explainable recommendation over knowledge graph," *arXiv preprint arXiv:2310.16452*, 2023.

[166] H. Li, Q. Dong, J. Chen, H. Su, Y. Zhou, Q. Ai, Z. Ye, and Y. Liu, "Llms-as-judges: a comprehensive survey on llm-based evaluation methods," *arXiv preprint arXiv:2412.05579*, 2024.

This thesis was produced while attending the PhD programme in *Mathematics and Computer Science* at the University of Cagliari, Cycle XXXVIII, with the support of a scholarship financed by the Ministerial Decree no. 351 of 9th April 2022, based on the NRRP - funded by the European Union - NextGenerationEU - Mission 4 “Education and Research”, Component 1 “Enhancement of the offer of educational services: from nurseries to universities” *or Investment 4.1 “Extension of the number of research doctorates and innovative doctorates for public administration and cultural heritage”.*