electronics

MDPI

Article

Enhancing Personalised Learning with a Context-Aware Intelligent Question-Answering System and Automated Frequently Asked Question Generation

Eleonora Bernasconi [1,*,†], Domenico Redavid [2,†] and Stefano Ferilli [1,†]

[1] Department of Computer Science, University of Bari Aldo Moro, 70125 Bari, Italy; stefano.ferilli@uniba.it
[2] Department of Economic and Finance, University of Bari Aldo Moro, 70124 Bari, Italy; domenico.redavid1@uniba.it
* Correspondence: eleonora.bernasconi@uniba.it
† These authors contributed equally to this work.

Abstract: The increasing integration of Artificial Intelligence (AI) in education has led to the development of innovative tools like Intelligent Question-Answering Systems (IQASs), aiming to revolutionize traditional learning paradigms. However, many existing IQAS struggle with the nuances of natural language and the complexities of student questions. This research focuses on developing a context-aware IQAS that leverages advanced Natural Language Processing (NLP) techniques and contextual information, including student learning history and educational content, to provide personalised support. This study also introduces a software tool that utilizes NLP techniques to automatically generate FAQs from educational materials. Employing a hybrid approach combining rule-based and machine learning techniques, the IQAS demonstrated high accuracy in interpreting and responding to a wide range of student queries. The software tool effectively automated the generation of FAQs, creating a valuable resource for personalised learning. The findings suggest that these tools can significantly improve student engagement, motivation, and learning outcomes, highlighting the potential of AI to transform education and pave the way for more personalised, adaptive, and effective learning environments.

Keywords: intelligent question answering; educational technology; student learning enhancement; adaptive learning systems; knowledge graphs; natural language processing; transformers in education; AI-driven personalisation; educational data mining

check for updates

Academic Editor: Vytautas Rudžionis

Received: 3 March 2025
Revised: 31 March 2025
Accepted: 4 April 2025
Published: 7 April 2025

Citation: Bernasconi, E.; Redavid, D.; Ferilli, S. Enhancing Personalised Learning with a Context-Aware Intelligent Question-Answering System and Automated Frequently Asked Question Generation. *Electronics* **2025**, *14*, 1481. https://doi.org/10.3390/electronics14071481

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The landscape of education is undergoing a dramatic transformation, propelled by the rapid integration of Artificial Intelligence (AI) [1]. As traditional learning paradigms are challenged, AI-powered tools are emerging as catalysts for personalised, engaging, and effective learning experiences. Among these advancements, Intelligent Tutoring Systems (ITSs) and Intelligent Question-Answering Systems (IQASs) stand out as particularly promising innovations [2]. ITSs provide adaptive feedback and personalised learning paths, while IQASs offer on-demand support, addressing students' queries in real time. This confluence of technologies has the potential to revolutionize education, tailoring the learning journey to individual student needs and fostering a deeper understanding of complex concepts [3].

However, despite the promise of these tools, current iterations of IQASs often fall short of their potential. Many existing systems struggle to effectively decipher the nuances of

natural language, particularly within the complex and dynamic context of student–teacher interactions [4]. Student queries, often imbued with colloquialisms, incomplete information, or implicit assumptions, pose a significant challenge to these systems. Consequently, responses may be inaccurate, irrelevant, or fail to address the underlying learning gap, leading to frustration and hindering the learning process. This limitation underscores the critical need for a more sophisticated IQAS that can accurately interpret and respond to a wide range of questions, taking into account the context of the query, the student's learning history, and the specific educational content being addressed.

This research endeavors to address this gap by developing a novel, context-aware IQAS designed to enhance learning outcomes and foster a more engaging and effective learning environment. By leveraging cutting-edge Natural Language Processing (NLP) techniques and incorporating contextual information, the proposed IQAS aims to deliver tailored and informative answers to student queries, promoting a deeper understanding of the subject matter [5]. Moreover, recognizing the importance of readily available knowledge resources, this research introduces a complementary software tool that automatically generates frequently asked questions (FAQs) from educational materials using advanced NLP techniques [6].

This tool will not only provide students with quick answers to common questions but will also aid instructors in developing more effective and comprehensive learning resources, further supporting personalised learning and efficient knowledge dissemination.

1.1. Significance

This research holds significant implications for the future of education, spanning theoretical, practical, and societal domains.

1.1.1. Theoretical Implications

This research pushes the boundaries of AI in education by exploring the synergistic integration of contextual understanding, personalised learning, and advanced NLP techniques in Question-Answering Systems [7]. It investigates how AI can be harnessed to provide tailored support to students based on their individual needs, learning history, and the specific context of their queries, contributing to a more nuanced understanding of personalised learning in AI-driven educational environments [8]. Furthermore, this study will contribute to the refinement of NLP algorithms for educational applications, specifically addressing the challenges of ambiguity, context dependency, and the evolving language patterns of students.

1.1.2. Practical Implications

The developed context-aware IQAS and the accompanying FAQ generation tool have the potential to be seamlessly integrated into existing educational platforms and Learning Management Systems (LMSs) [9]. This integration can provide students with immediate, personalised support, potentially improving learning outcomes, increasing engagement, and fostering a more interactive learning experience. The ability to cater to individual learning styles and paces makes this technology particularly beneficial for diverse learners, promoting inclusivity and accessibility in education. In addition, these tools can help educators quickly create responsive learning materials, significantly reducing the time and effort traditionally required to develop FAQs.

1.1.3. Societal Implications

At a broader level, this research contributes to the development of more effective, accessible, and equitable educational tools. By democratizing access to personalised learning support and facilitating the dissemination of knowledge, these AI-powered tools can

empower students from all backgrounds to achieve their full academic potential. This, in turn, can contribute to a more educated and informed populace, fostering intellectual growth and social progress. The potential to bridge educational gaps and promote lifelong learning underscores the transformative impact of this research on society [10].

1.2. Objectives

This research aims to design and implement a context-aware IQAS capable of accurately interpreting and responding to a diverse range of student queries. The system leverages contextual information such as each student's learning history, prior interactions, and course-specific materials to provide tailored and informative answers. Furthermore, the study seeks to enhance the system's capacity to manage complex, nuanced, and multifaceted questions by incorporating advanced reasoning, inference, domain-specific knowledge, and strategies designed to address ambiguities or implicit assumptions often present in student language. In addition to these efforts, the research involves developing an automated software tool that generates comprehensive and relevant FAQs from educational materials, including textbooks, lecture notes, and online resources, by employing advanced NLP techniques to identify frequently asked questions and produce accurate, concise answers. Moreover, the project explores the seamless integration of both the IQAS and the FAQ generation tool into existing Learning Management Systems (LMSs) to ensure accessibility, user-friendliness, and broad adoption within familiar learning environments. Finally, a comprehensive evaluation will be conducted to assess the performance of both systems by comparing their effectiveness against traditional Question-Answering Systems, existing AI-powered learning tools, and human tutors, using metrics such as accuracy, relevance, response time, and user satisfaction.

1.3. Research Questions

This research seeks to address the following key questions:

1. **RQ1:** Optimizing NLP for educational context. How can Natural Language Processing techniques be optimized to enhance the accuracy and relevance of responses generated by the IQAS, specifically addressing the unique challenges posed by educational content, student language, and the dynamic nature of classroom interactions?
2. **RQ2:** Handling complex and nuanced queries. What strategies can be implemented to enhance the system's ability to understand and respond to complex and nuanced queries that require reasoning, inference, the integration of domain-specific knowledge, and the resolution of ambiguities inherent in natural languages? [11]
3. **RQ3:** Seamless integration into learning environments. How can the IQAS and the FAQ generation tool be seamlessly integrated into existing Learning Management Systems to provide effective, accessible, and user-friendly support to students, ensuring compatibility across different platforms and learning environments?
4. **RQ4:** Effectiveness of automatic FAQ generation. How effective is the software tool in automatically generating relevant, informative, and comprehensive FAQs from diverse educational materials, and how can this tool be leveraged to support personalised learning, knowledge dissemination, and the creation of dynamic learning resources?
5. **RQ5:** Impact on learning outcomes. What is the measurable impact of the context-aware IQAS and FAQ generation tool on student learning outcomes, engagement, and satisfaction compared to traditional learning methods and existing AI-powered tools?

1.4. Hypothesis

We hypothesize that the development of a context-aware IQAS, coupled with an automatic FAQ generation tool, both leveraging advanced NLP techniques, contextual information, and domain-specific knowledge, will result in significantly improved student satisfaction, engagement, and learning outcomes compared to traditional Question-Answering Systems and existing AI-powered learning tools [12]. We further hypothesize that the system will demonstrate superior performance in handling complex and nuanced queries, providing more relevant and informative answers than current systems. We resume in Table 1 the connection between this work's goals and research questions.

Table 1. Objectives and corresponding research questions.

Objective	Research Question(s)	Rationale
Design and implement a context-aware IQAS that interprets and responds to diverse student queries, leveraging contextual information.	RQ1, RQ2	Directly addresses the creation of a system (RQ1) that can interpret diverse student questions, including those with implicit or incomplete information (RQ2).
Enhance the system's ability to manage complex, nuanced questions (reasoning, inference, ambiguity resolution).	RQ2	Focuses on strategies for advanced reasoning and inference, essential for addressing higher-level question complexity (RQ2).
Develop an automated FAQ generation tool from educational materials.	RQ4	Evaluates how effectively the tool identifies and formulates FAQs, measuring its potential to support students and educators (RQ4).
Seamlessly integrate the IQAS and FAQ generation tool into existing LMS platforms.	RQ3	Ensures technical compatibility and accessibility, promoting broader adoption in real-world learning environments (RQ3).
Conduct a comprehensive evaluation of the IQAS and FAQ generation tool (accuracy, response time, user satisfaction, impact on learning).	RQ5	Empirically validates the overall system by measuring improvements in student engagement, outcomes, and user experience (RQ5).

1.5. Novelty

The novelty of this research lies in its holistic and multifaceted approach to creating an intelligent question-answering ecosystem tailored for personalised learning. Unlike traditional IQAS systems that primarily rely on keyword matching, our approach integrates a deep contextual awareness by considering students' learning histories, previous interactions, and the specific educational content being studied. This results in responses that are not only accurate but also highly personalised. Moreover, the research leverages advanced NLP techniques—such as transformer models and deep learning architectures—to effectively capture and understand the nuances of natural language within educational discourse. The system is further enhanced by incorporating domain-specific knowledge, enabling it to provide precise answers to complex questions that demand a specialised understanding of the subject matter. In addition, a novel software tool is developed to automatically generate FAQs from educational materials, thereby offering personalised learning resources that support students and assist educators in developing comprehensive learning materials. By synergistically integrating the advanced IQAS with the automated FAQ generation tool, this research creates an environment that caters to both immediate query resolution and broader knowledge acquisition. Overall, this comprehensive approach has the potential to significantly enhance personalised learning, particularly when addressing

complex or nuanced queries, and it represents a substantial advancement in the application of AI in education [13].

1.6. Organization

This paper is structured as follows:

Section 2 presents a comprehensive review of related work in the field of AI in education, with a particular focus on Intelligent Question-Answering Systems, Natural Language Processing, and personalised learning. This section analyses key contributions, identifies technical gaps, and critically evaluates the strengths and limitations of state-of-the-art methods.

Section 3 details the methodology employed in the design and implementation of the proposed context-aware IQAS and the automatic FAQ generation tool. This includes a description of the data collection and preprocessing steps, the architecture of the system, the integration of advanced NLP techniques, and the overall system design.

Section 4 presents the experimental results, providing a detailed evaluation of the performance of both the IQAS and the software tool using a variety of metrics, including accuracy, response time, user satisfaction, and impact on learning outcomes. This section also includes a comparative analysis of existing systems and a discussion of the results in different scenarios.

Section 5 discusses the implications of the findings, comparing the system's performance with existing approaches, exploring its potential impact on personalised learning, and addressing the broader implications for the field of education.

Finally, Section 6 concludes the paper by summarizing key contributions, acknowledging limitations, and highlighting promising directions for future research.

2. Related Work

The integration of Artificial Intelligence (AI) into teaching and learning contexts represents a significant advancement in educational technology, bridging critical gaps in both theoretical understanding and practical application. Despite the rapid proliferation of AI technologies across various sectors, their exploration within educational contexts remains notably underrepresented. This section critically examines the current state-of-the-art in Intelligent Question-Answering Systems (IQASs), tracing their evolution, highlighting persistent technical gaps, and discussing emerging directions to situate the current research within a broader academic landscape. A high-level schematic summarizing these developments and associated challenges is provided in Figure 1.

2.1. Early IQAS and Rule-Based Approaches

Early IQAS primarily relied on rule-based methods, using keyword matching or scripts crafted manually to generate answers. These approaches often worked well for simple factoid queries but were less effective when queries demanded more reasoning or semantic understanding [14]. The core limitation was their rigidity—they typically failed to account for variations in phrasing and struggled with more open-ended student queries [15].

2.2. The Shift to Statistical NLP and Machine Learning

The emergence of statistical NLP and machine learning methods marked a significant evolution in IQAS. Instead of relying on static rules, the systems began to learn patterns from annotated datasets, enabling wider coverage and improved accuracy. Techniques like Support Vector Machines (SVMs), Conditional Random Fields (CRFs), and ensemble classifiers enabled the handling of diverse query forms. In the education domain, these methods helped the systems adapt to various language styles and partially address the

variability of student questions [16]. However, they often required extensive feature engineering and were limited by the quality and quantity of available training data [17].

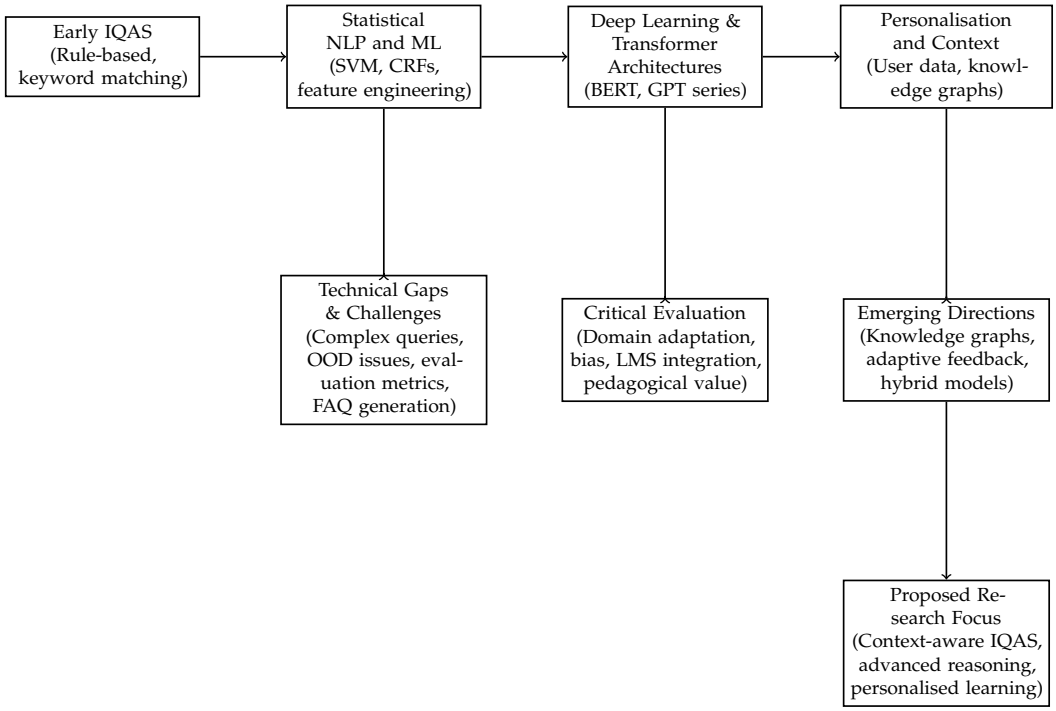


Figure 1. Schematic overview of the evolution of IQAS, highlighting major methodological shifts, technical challenges, and future directions in educational contexts.

2.3. Deep Learning and Transformer-Based Architectures

The introduction of deep learning—particularly transformer architectures such as BERT, GPT-2, GPT-3, and subsequent large language models—dramatically broadened the capabilities of IQAS. These models excel at capturing contextual relationships and long-range dependencies, allowing them to produce more coherent and context-sensitive responses.

- Contextual Word Embeddings: Word and sentence embeddings could capture nuanced semantic relationships, helping IQAS distinguish between subtly different question formulations [18].
- State-of-the-Art Performance: Transformer-based models have established new benchmarks in tasks such as question answering, text summarization, and machine translation, which is beneficial to educational applications by offering more robust and nuanced response generation [19].
- Zero-Shot or Few-Shot Learning: Large language models pre-trained on massive corpora can often handle a wide variety of topics with minimal domain-specific fine-tuning, making them appealing for broad educational environments where course material can span multiple disciplines [20].

2.4. Advances Toward Personalisation and Context

Recent work has focused on personalisation and context modeling, especially for educational use cases. By integrating user data (learning histories, proficiency levels) and domain-specific knowledge (ontologies, knowledge graphs), modern systems can deliver targeted explanations that deepen student engagement [21]. For example, platforms such as Wolfram Alpha [22] provide detailed step-by-step solutions to maths problems, adapting the complexity of the content based on user input [23]. Similarly, Duolingo [24] adjusts language exercises and feedback by tracking the progress of the learner and error patterns.

Although these examples demonstrate the potential of personalised IQAS to improve learning outcomes, implementing such features at scale, especially in privacy-sensitive or resource-limited environments, remains a pressing challenge.

2.5. Technical Gaps and Challenges

Despite notable progress, the application of IQAS in educational contexts continues to face several challenges.

2.5.1. Complex Question Handling

Many current models struggle with questions that require multistep reasoning, integration of domain-specific knowledge, or logical inference [3].

Specialised academic subjects (e.g., advanced physics, literature analysis) demand a highly specific knowledge base, which general-purpose models may lack.

2.5.2. Out of Domain Queries

Students often pose questions that cross disciplinary boundaries or explore niche topics, requiring the system to adapt quickly or acknowledge its limitations of knowledge [25].

Updating and maintaining the knowledge base across the entire curriculum can be resource-intensive, particularly in rapidly evolving disciplines.

2.5.3. Adaptive Question Answering

Individual learning styles and levels of background knowledge require personalised interaction, including handling incomplete or context-dependent queries [26].

Failing to incorporate a student's history or misunderstand their intent can lead to irrelevant or overly complex answers, which undermines learning outcomes.

2.5.4. Evaluation Metrics

While traditional metrics are useful, they do not fully capture pedagogical effectiveness, learner satisfaction, or the potential for long-term knowledge retention [27].

There is a need for metrics that include clarity, instructional value, and alignment with learning objectives to better reflect real educational outcomes. For example, recent work has explored the use of "explainability" metrics to assess the extent to which an IQAS can provide human-understandable justifications for its answers, which is crucial for the promotion of learning and trust [4].

2.5.5. FAQ Generation

Dynamically producing FAQs from lecture notes, textbooks, and user queries can streamline knowledge dissemination but remains underutilized in many educational domains [28].

Ensuring accuracy and relevance in automatically generated FAQs requires robust validation methods and, potentially, human-in-the-loop oversight [29].

2.6. Critical Evaluation

Although transformer-based language models have shown remarkable promise, their application directly in educational settings often reveals domain-specific gaps [30]. General purpose pre-training might not suffice for subjects requiring specialised terminology or complex conceptual mappings (e.g., advanced mathematics, chemistry) [31]. Fine-tuning of educational data can address some issues but raises questions about data availability and privacy [32]. Because large language models learn patterns from large (and sometimes biased) corpora, they may inadvertently generate responses that are culturally insensitive or pedagogically questionable [33]. Mitigating bias requires careful curation and con-

tinuous monitoring of the data sets, especially in sensitive educational contexts where students expect authoritative and unbiased explanations [34]. Seamless integration with Learning Management Systems (LMSs) remains a challenge [35]. Many LMS platforms do not natively support advanced AI functionalities or require complex API-level adaptations, complicating deployment [36]. Furthermore, real-world classroom settings often require scalability and robustness to handle peak usage [37]. A well-performing IQAS in terms of accuracy may still fail to enhance learning if it provides terse or nonexplanatory answers [6]. Educational contexts require qualitative assessments that measure how effectively a response fosters comprehension, motivation, and higher-order thinking skills [38]. Thus, setting standardized benchmarks for educational QA is an ongoing task [39].

Emerging Directions

While deep learning and personalisation have already marked significant advances in IQAS, several promising pathways are shaping the future of this field. First, knowledge graph integration offers a structured framework for domain modeling and concept mapping by explicitly encoding relationships among entities [40]. Such graph-based representations can help IQAS handle complex, multi-step queries with greater contextual awareness, making them especially valuable in specialised educational subjects where detailed conceptual hierarchies are crucial. Next, adaptive feedback mechanisms leverage ongoing interactions with learners and instructors (such as real-time ratings, comments, or behavioural indicators) to refine the system's performance iteratively [41]. This feedback loop not only helps detect inaccuracies and biases but also supports dynamic personalisation, ensuring that responses evolve in tandem with the changing needs and knowledge levels of the learner. Finally, hybrid approaches combine large pre-trained language models with symbolic reasoning or specialised domain modules to achieve more reliable, context-sensitive outputs in niche academic areas [42]. By uniting the broad coverage of data-driven NLP with the precision of expert-crafted logic, these hybrid systems can offer a more robust handling of questions that demand both pattern recognition and deep subject matter expertise. Taken together, these directions signal an increasingly holistic and learner-centered trajectory for future IQAS development.

2.7. Research Focus

In summary, the evolution from rule-based IQAS to deep learning-driven models underscores the substantial progress in the field [43]. Nonetheless, educational applications present unique demands for domain-specific adaptation, handling out-of-domain questions, and meeting diverse student needs [44]. This research aims to address these challenges by developing a context-aware IQAS that incorporates domain knowledge, leverages student learning history, and demonstrates pedagogical effectiveness [45]. This is achieved by the following:

- Addressing Domain Adaptation: Fine-tuning large language models on curated educational datasets and potentially integrating knowledge graphs to enhance domain-specific understanding [46].
- Handling Complex Questions: Incorporating advanced reasoning and inference mechanisms to address complex, nuanced queries that require more than simple information retrieval [47].
- Promoting Personalised Learning: Leveraging student learning history and contextual information to provide tailored answers and support individualised learning paths [18].

Furthermore, we propose an FAQ generation tool that automatically curates and maintains a knowledge base of frequently asked questions, supporting both personalised

learning and large-scale knowledge dissemination [48]. By focusing on integration with real-world LMS platforms, robust evaluation strategies that go beyond simple accuracy metrics, and ethical considerations, this work seeks to advance the state of AI-driven education and set a foundation for future research in personalised and context-aware learning technologies.

3. Methodology

This research develops and evaluates an Intelligent Question-Answering System (IQAS) and a complementary software tool designed to enhance personalised student learning. In addition to addressing technical challenges, the IQAS is grounded in established educational theories to ensure that its responses and feedback promote effective learning. Building on the challenges and directions outlined in Section 2, the following subsections detail the underlying analytical framework, data collection strategies, data preprocessing pipeline, and the analytical techniques employed to achieve robust, context-aware, and pedagogically meaningful question answering.

3.1. Analytical Framework

We adopt a hybrid approach that synergises the precision of rule-based methods with the adaptability of machine learning (ML) models, as illustrated in Figure 2. Importantly, our system design is informed by educational methodologies and learning theories such as constructivism, cognitive load theory, and formative assessment practices. This ensures that the system not only processes queries accurately but also delivers responses that facilitate deeper understanding and promote self-regulated learning.

Each component in the diagram represents a critical stage or service within our IQAS pipeline:

- **Data Sources:** We aggregate Q&A pairs from various channels: educational platforms online (e.g., Coursera, edX, Khan Academy), forums (e.g., Stack Overflow, Reddit), textbooks and expert-annotated datasets, to capture both linguistic diversity and pedagogically sound content. These sources are selected to reflect the varied subject areas and teaching styles, thus forming the IQAS in real-world educational contexts.
- **Data Pre-Processing:** Essential Natural Language Processing (NLP) tasks to clean and standardize the collected data, including the following:
  - **Cleaning:** Removing extraneous elements such as HTML tags, special characters, and duplicate entries.
  - **Tokenization:** Splitting text into individual tokens (words or subwords) for further processing.
  - **Stemming and Lemmatisation:** Reducing words to their root or dictionary forms to unify linguistic variants (e.g., “running” vs. “ran”).
  - **Stopword Removal:** Eliminating common words (e.g., “the,” “and,” “is”) that carry little semantic weight.
  - **Named Entity Recognition (NER):** Identifying and classifying named entities (e.g., people, places, scientific terms) to enhance contextual understanding.
  - **Part-of-Speech (POS) Tagging:** Annotating each word with its grammatical role to support more nuanced linguistic analysis.
- **Model Training:** Utilizing transformer-based architectures (e.g., BERT) and constructing domain-specific knowledge graphs, involving:
  - **Transformer-Based Models:** Training models to understand and generate human-like text by capturing intricate semantic and syntactic relationships.

  - **Fine-Tuning:** Adapting pre-trained models to the specific educational domains and types of queries anticipated in the IQAS.
  - **Knowledge Graph Construction:** Building structured representations of domain knowledge to support advanced reasoning and inference capabilities.

  These trained models parse complex queries, extract pertinent information, and generate contextually relevant responses.
- **Rule-Based Component:** Handles high-precision, factual queries by leveraging a curated knowledge base and predefined rules. For example, a query like "*When did World War II end?*" triggers a direct lookup from the knowledge base, ensuring rapid and accurate responses for routine questions.
- **Knowledge Graph:** Acts as a structured repository of domain-specific entities and their interrelationships, facilitating multi-hop reasoning and contextual inference. This component enriches the ML pipeline by providing a backbone for logical deductions and advanced query handling [40].
- **ML Inference:** Deploys trained transformer models alongside graph-based reasoning to address complex, open-ended questions requiring deeper semantic understanding and contextual insights [19].
- **FAQ Generation:** Automates the creation and maintenance of frequently asked questions using text summarization and QA extraction pipelines. By analysing Q&A logs, it identifies recurring queries and formulates concise, informative answers to enhance system responsiveness.
- **IQAS Integration:** Serves as the central hub, merging outputs from both the rule-based and ML inference pipelines. It incorporates personalisation logic and adaptive feedback mechanisms to tailor responses based on user proficiency and learning history, ensuring contextually appropriate and pedagogically effective answers.
- **User Query (UI/LMS/Chat):** The final interaction point where users (students or instructors) submit queries through various interfaces such as user interfaces (UIs), Learning Management Systems (LMSs), or chat applications. The system responds with integrated answers while logging interactions to inform continuous improvement processes.

In Figure 2, different colors are employed to visually distinguish the various components of the IQAS pipeline, enhancing clarity and facilitating quick comprehension. The color scheme is as follows:

- **Blue (Data Sources):** Represents the origin points of raw data, including online educational platforms, forums, textbooks, and manually annotated datasets. Blue signifies the input stage where diverse Q&A pairs are collected.
- **Orange (Processes):** Denotes the processing stages within the pipeline, encompassing data pre-processing, model training, rule-based components, ML inference, IQAS integration, and FAQ generation. Orange highlights the transformative operations that convert raw data into actionable responses.
- **Green (Knowledge Graph):** Indicates the knowledge graph component, which stores domain-specific entities and their relationships. Green symbolizes structured knowledge and logical connections that facilitate advanced reasoning.
- **Gray (User Interface):** Represents the final interaction point where users submit queries and receive answers via various interfaces such as UI, LMS, or chat applications. Gray underscores the interface's role as the system's outward-facing component.

Together, these components enable our IQAS to effectively support a wide range of student queries, from simple factoid requests to complex, inference-laden inquiries, ensuring both precision and contextual relevance in responses.

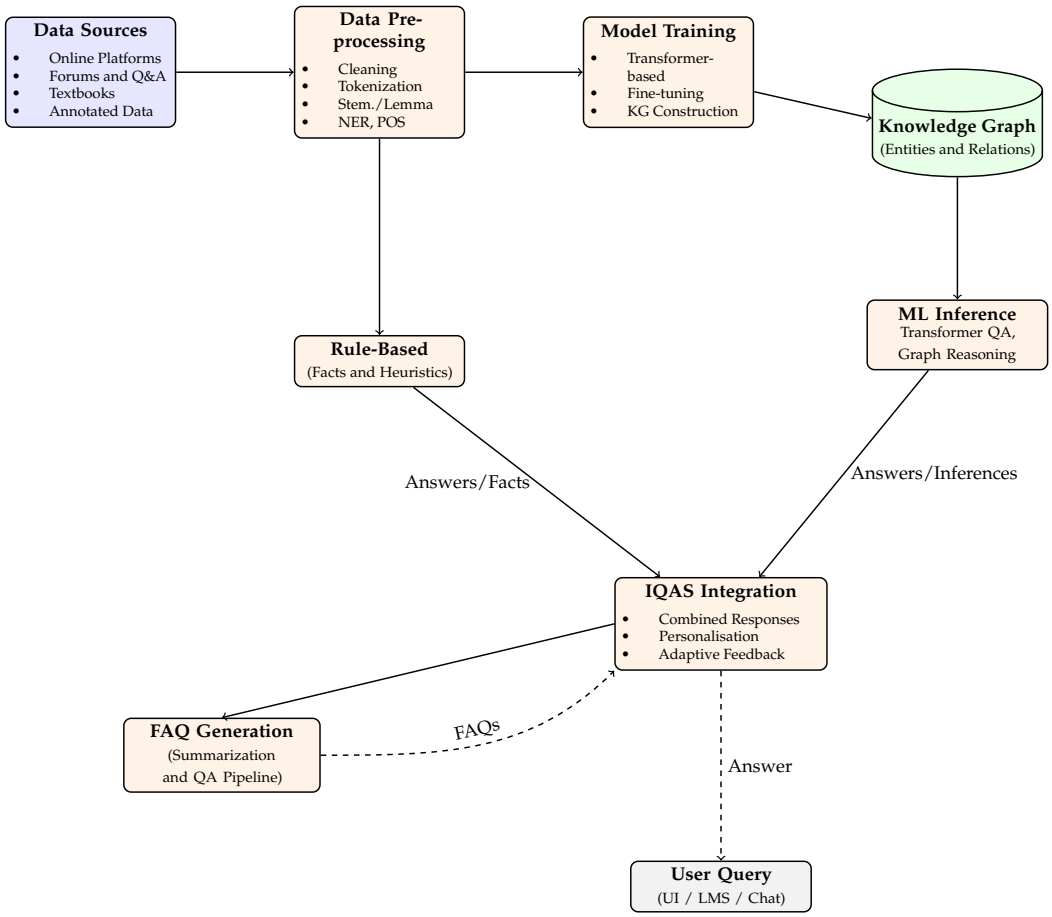


Figure 2. Pipeline of data flow and service integration within the IQAS. Data sources feed into a pre-processing stage, which drives both rule-based and machine learning (ML) components. A knowledge graph enhances inference, and the system integrates results through the IQAS Integration module. Additionally, an FAQ generation module leverages Q&A logs to produce recurring questions and answers. The color coding distinguishes different functional areas: blue for data sources, orange for processing stages, green for the knowledge graph, and gray for user interactions.

3.2. Mathematical Formalization of Core Components

Following this, we provide formal mathematical descriptions of the core components of the system: the model architecture, the loss functions, the optimization strategies, and the attention mechanism employed in the FAQ generation module.

3.3. Model Architecture

Let the input sequence be represented as

$$\mathbf{X} = \{x_1, x_2, \ldots, x_n\},$$

where each x_i is a token. An embedding layer maps each token to a vector in $\mathbb{R}^d$ via an embedding matrix $\mathbf{E} \in \mathbb{R}^{|V| \times d}$ (with $|V|$ as the vocabulary size and d as the embedding dimension):

$$\mathbf{E}(x_i) = \mathbf{E}_{x_i}.$$

The sequence of embeddings is then processed by a transformer encoder. For a single self-attention layer, the queries, keys, and values are computed as:

$$\mathbf{Q} = \mathbf{X}\mathbf{W}_Q, \quad \mathbf{K} = \mathbf{X}\mathbf{W}_K, \quad \mathbf{V} = \mathbf{X}\mathbf{W}_V,$$

where $\mathbf{W}_Q$, $\mathbf{W}_K$, $\mathbf{W}_V \in \mathbb{R}^{d \times d_k}$ are learned projection matrices. The scaled dot-product attention is defined as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right)\mathbf{V}.$$

A residual connection and layer normalization are applied:

$$\mathbf{Z} = \text{LayerNorm}(\mathbf{X} + \text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V})).$$

Finally, the output representation $\mathbf{H}$ is passed through a decoder (or classification head) to produce predictions:

$$\hat{\mathbf{Y}} = g(\mathbf{H}; \theta_g),$$

where $g(\cdot)$ is a non-linear mapping (e.g., a feed-forward network with a softmax activation) parameterized by θ_g.

3.4. Loss Functions

For the question-answering task, the model is trained to minimize the cross-entropy loss between the predicted distribution and the ground-truth answers. Let y denote the true answer token and $\hat{y}$ the predicted probability. For a dataset with N examples, where the ith example has an answer of length T_i, the loss is defined as:

$$\mathcal{L}_{\text{QA}} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{t=1}^{T_i}\log P\left(y_t^{(i)} \mid \mathbf{X}^{(i)}; \theta\right).$$

If an auxiliary loss $\mathcal{L}_{\text{KG}}$ is employed for knowledge graph reasoning, the total loss becomes:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{QA}} + \lambda\,\mathcal{L}_{\text{KG}},$$

where λ is a hyperparameter that balances the two loss components.

3.5. Optimization Strategies

The model parameters θ are optimized using the Adam optimizer. At each iteration t, the parameter updates are computed as:

$$\begin{aligned} m_t &= \beta_1 m_{t-1} + (1-\beta_1)\nabla_\theta \mathcal{L}_t, \\ v_t &= \beta_2 v_{t-1} + (1-\beta_2)(\nabla_\theta \mathcal{L}_t)^2, \\ \hat{m}_t &= \frac{m_t}{1-\beta_1^t}, \quad \hat{v}_t = \frac{v_t}{1-\beta_2^t}, \\ \theta_{t+1} &= \theta_t - \eta\,\frac{\hat{m}_t}{\sqrt{\hat{v}_t}+\epsilon}, \end{aligned}$$

where η is the learning rate, β_1 and β_2 are decay rates, and ϵ is a small constant for numerical stability.

3.6. Mathematical Formalization of the FAQ Generation Module

The FAQ generation module utilizes an encoder–decoder architecture with an attention mechanism to summarize and extract key information from educational texts.
Encoder:

The input document is represented as:

$$\mathbf{X} = \{x_1, x_2, \ldots, x_n\},$$

and processed by the encoder to produce hidden states:

$$\mathbf{H} = \{h_1, h_2, \ldots, h_n\}, \quad h_i \in \mathbb{R}^{d_h}.$$

Attention mechanism:

At each decoding time step t, the decoder has a state $s_t \in \mathbb{R}^{d_h}$ that serves as the query. The attention score between s_t and each encoder hidden state h_i is computed by:

$$e_{t,i} = s_t^\top W_a h_i,$$

where $W_a \in \mathbb{R}^{d_h \times d_h}$ is a learned weight matrix. These scores are normalized via the softmax function:

$$\alpha_{t,i} = \frac{\exp\left(\frac{e_{t,i}}{\sqrt{d_k}}\right)}{\sum_{j=1}^{n} \exp\left(\frac{e_{t,j}}{\sqrt{d_k}}\right)},$$

yielding the context vector:

$$c_t = \sum_{i=1}^{n} \alpha_{t,i} h_i.$$

Decoder output:

The decoder generates the output token based on the current state and the context vector:

$$y_t = g(s_t, c_t; \theta_g),$$

where $g(\cdot)$ is a function (e.g., a feed-forward network with a softmax activation) parameterized by θ_g.

3.7. Data Collection

As depicted in Figure 2, data collection is the foundational step feeding the entire IQAS pipeline. A robust dataset is essential for training and evaluating both technical and educational components. Data are sourced and aggregated from multiple channels to capture realistic student language patterns and domain complexity:

- **Online Educational Platforms:** Large repositories such as Coursera, edX, and Khan Academy provide varied question–response pairs, reflecting multiple subjects and learner demographics. These platforms offer diverse linguistic styles and question types, ensuring comprehensive coverage of educational domains.
- **Online Forums and Q&A Sites:** Platforms like Stack Overflow or Reddit contain domain-specific discussions, providing natural language data for niche or advanced topic areas. This data helps the IQAS understand and generate responses to specialised queries.
- **Textbooks and Educational Materials:** Common queries and key insights are extracted from textbooks and reference guides, contributing domain-rich examples for fine-tuning. Analysing structured educational content ensures that the IQAS aligns with established curricular standards.
- **Manually Annotated Dataset:** A curated set of queries and expert-verified answers focuses on high-priority or pedagogically sensitive topics, ensuring both quality and relevance. This dataset is used to fine-tune the machine learning model and assess its performance on specific tasks.

Figure 2 illustrates how these diverse data sources converge into the data preprocessing stage, setting the stage for effective model training and rule-based operations. Through diverse data sourcing, we aim to cover broad subject areas, varying question

styles, and both general and domain-specific linguistic expressions, thereby enhancing the IQAS's ability to handle a wide spectrum of student queries.

3.8. Data Processing

Following data collection, all gathered data undergo a structured preprocessing pipeline to ensure consistency, clarity, and readiness for advanced NLP techniques, as outlined in Figure 2. This pipeline comprises several key steps:

- **Data Cleaning:** Removal of extraneous elements such as HTML tags, special characters, and duplicate entries to retain only high-quality question–answer pairs. This step ensures that the data is free from noise that could adversely affect model training.
- **Tokenization:** Splitting text into tokens (words or subwords) enables efficient downstream tasks like embedding generation and vocabulary indexing. Tokenization standardizes the input for consistent processing across different NLP models.
- **Stemming and Lemmatisation:** Reducing words to their root (stemming) or dictionary forms (Lemmatisation) unifies linguistic variants (e.g., "running" vs. "ran"), minimizing vocabulary size and improving model generalization.
- **Stopword Removal:** Eliminating low-information words such as "the", "and", and "is" focuses the dataset on meaning-bearing terms. This reduction in dimensionality enhances model efficiency without sacrificing semantic understanding.
- **Named Entity Recognition (NER):** Identifying crucial entities (people, places, scientific terms) supports context-aware query handling, especially relevant for domain-specific or advanced academic questions [3]. NER helps in understanding the specific focus of a query, facilitating more accurate responses.
- **Part-of-Speech (POS) Tagging:** Annotating each word with its grammatical role (e.g., noun, verb, adjective) supports fine-grained analysis. POS tagging aids in better interpretation of student intent and enhances the system's ability to generate grammatically coherent responses.

As illustrated in Figure 2, the preprocessed data feed into both the rule-based and ML pipelines. This dual-path approach leverages technical robustness and educational insights, ensuring that the IQAS can address a wide spectrum of queries—from simple factual questions to complex conceptual inquiries—with both high accuracy and pedagogical effectiveness.

4. Detailed Analytical Techniques

This section provides a detailed explanation of the analytical techniques employed in the development and evaluation of the IQAS and the software tool.

4.1. Algorithms and Models

To address the varying complexities of student questions, our system integrates state-of-the-art NLP models with symbolic and graph-based reasoning techniques. First, we employ BERT-based encoding using bidirectional encoder representations from transformers (BERTs) to generate context-rich embeddings for both queries and candidate responses. These embeddings capture subtle semantic and syntactic signals, thereby improving the system's ability to understand student intent [19]. Next, the approach takes advantage of transformer-based encoder–decoder models; here, pre-trained encoder–decoder architectures are fine-tuned to generate coherent, topic-specific responses, an approach that aligns with best practices in text summarization and machine translation to ensure responses are accurate and contextually relevant [20]. In addition, inspired by the need for structured context in advanced topics, the system incorporates knowledge graph reasoning by employing knowledge graphs to facilitate multi-hop reasoning for complex questions, where

queries that map to specific nodes or relations can trigger deeper inference paths, thereby enriching answer quality [40].

Furthermore, a complementary FAQ generation tool is proposed to automate the curation of recurring questions from educational materials. This tool first utilizes a question-answering pipeline, in which a pretrained QA model identifies potential Q&A pairs directly from text with minimal manual intervention. It then employs a summarization model to pinpoint crucial sentences in long materials, which are subsequently restructured into commonly asked questions. Finally, the resulting FAQ items are stored using database storage techniques, allowing them to be indexed for quick retrieval so that both students and instructors can locate relevant knowledge instantly. Together, these models and tools work in tandem to ensure that IQAS not only provides accurate and contextually relevant answers but also continuously evolves by learning from frequently asked questions, thereby enhancing its utility and responsiveness over time.

4.2. Evaluation Metrics

To capture both accuracy and educational efficacy, our evaluation approach blends quantitative measures with qualitative user-centered assessments. These metrics are summarized in Table 2.

By combining these metrics, we gain a holistic view of our system performance, covering not only the correctness but also the instructional value and the user experience.

In summary, this methodology implements a hybrid rule-based and deep learning framework, enriched by knowledge graph reasoning, to develop a context-sensitive, educationally aligned IQAS. The multi-faceted evaluation protocol, encompassing both quantitative and qualitative metrics, ensures that the system maintains technical robustness while delivering significant pedagogical impact. This comprehensive approach guarantees that the IQAS not only performs accurately but also effectively supports diverse learning needs and enhances the overall educational experience.

Table 2. Summary of evaluation metrics.

Metric Type	Metric Name	Description
Quantitative	Accuracy	The proportion of correctly answered questions out of the total number of questions.
	Precision	The ratio of correctly answered questions to the total number of questions answered by the system.
	Recall	The ratio of correctly answered questions to the total number of questions that have correct answers in the dataset.
	F1 score	The harmonic mean of precision and recall, providing a balanced measure of accuracy.
Qualitative	Expert Review	Domain experts and educators assess the relevance, clarity, and pedagogical effectiveness of answers, ensuring alignment with best teaching practices.
	User Studies	Students and faculty provide direct feedback on system utility, ease of use, and perceived value in actual learning scenarios.
	Cognitive Load Analysis	Evaluates the mental effort required to interpret and act upon system responses, illuminating how the IQAS influences cognitive processes involved in learning [38].
	Explainability Measures	Tracks how well the system can justify its answers, thereby strengthening student trust and comprehension, as recommended by emerging standards [4].

4.3. Comparisons with Existing Techniques

The analytical techniques and methodologies proposed in this research provide significant advancements over existing approaches to Intelligent Question-Answering Systems (IQASs). By leveraging a hybrid framework that combines rule-based precision with the flexibility of machine learning (ML) models, this system addresses critical limitations observed in traditional IQAS implementations. In comparison to conventional methods, the proposed system exhibits superior performance with transformer models. Traditional NLP techniques often rely on simpler architectures such as bag of words, TF-IDF, or LSTMs, which fail to fully capture contextual nuances. In contrast, the integration of BERT and transformer-based models allows for bidirectional encoding and deeper semantic understanding, consistently outperforming these legacy methods across question-answering and text generation tasks [19]. This ensures higher accuracy and better adaptability to complex query structures.

Furthermore, the system achieves enhanced reasoning capabilities by incorporating knowledge graphs. Many existing systems struggle with queries that require logical inference, multi-hop reasoning, or domain-specific knowledge integration. By providing structured representations of domain entities and their interrelations, the system enables advanced reasoning that bridges a critical gap left by conventional systems, which often rely solely on unstructured text and shallow reasoning [40]. In addition to these technical improvements, the proposed approach offers a more contextual and personalised understanding. Unlike traditional IQAS that typically generate generic answers, the system combines BERT-based contextual understanding with adaptive feedback mechanisms to capture both local (sentence-level) and global (document-level) contexts, thereby tailoring responses to individual user needs and overcoming the personalisation limitations found in many existing tools.

Another notable advancement is the novel automatic FAQ generation. Existing Question-Answering Systems rarely offer functionalities for generating and maintaining FAQs. To address this, the proposed system incorporates a dedicated FAQ generation module that utilizes summarization and QA pipelines to curate and dynamically update recurring questions. This innovation not only provides personalised and continuously evolving learning resources but also reduces the manual workload for educators. Moreover, the system adopts a hybrid rule-based and ML approach. While many IQAS rely exclusively on ML-based methods, the proposed framework integrates rule-based systems for handling high-precision factual queries alongside ML techniques for complex, context-rich questions, ensuring optimal accuracy and efficiency for both routine and advanced queries.

Finally, the research emphasizes explainability and trustworthiness. Explainability remains a challenge in many existing systems, yet by incorporating knowledge graphs and personalisation mechanisms, the proposed IQAS is capable of providing justifiable answers that are both contextually relevant and pedagogically aligned. This enhancement not only fosters trust among end users but also promotes higher engagement. By integrating these advanced techniques, the research aims to develop an IQAS and accompanying software tool that surpasses the capabilities of existing systems in terms of accuracy, relevance, and informativeness, ultimately offering a more effective learning support tool for students.

5. The Intelligent Question-Answering System and FAQ Generation Tool—Case Study

The live demonstration of the Intelligent Question-Answering System (IQAS) and FAQ generation tool is accessible at https://iqas.uniba.digital-humanities.it (accessed on 24 January 2025). This section highlights the integration of advanced NLP methodologies

and a user-friendly interface, offering an interactive platform for automated FAQ generation and personalised learning support.

5.1. System Workflow

The workflow of the IQAS integrates rule-based approaches, transformer models, and a knowledge graph to achieve precise, context-aware answers and FAQs. This hybrid architecture ensures real-time query answering and automated FAQ generation with pedagogical relevance. The workflow is structured into the following stages:

1. **Data Upload and Preprocessing:** Users upload educational materials (e.g., PDFs). The system extracts and preprocesses the text for tokenization, summarization, and entity recognition.
2. **Automated FAQ Generation:** Preprocessed data are analysed to generate FAQs using text summarization and question-answering pipelines. FAQs are stored in a database for efficient retrieval.
3. **Real-Time Question Answering:** Users interact with the system by submitting queries, which are processed using transformer-based models and knowledge graph reasoning to provide precise and contextually relevant answers.
4. **Feedback and Adaptive Learning:** Users can rate and comment on FAQs, which informs the system's iterative improvements and enhances response quality over time.

5.2. Demo Features

Table 3 highlights the main functionalities of the demo, showcasing its practical application for education and research purposes.

Table 3. Features of the IQAS and FAQ generation tool.

Feature	Description
PDF Upload	Allows users to upload educational materials for text extraction and processing.
FAQ Generation	Automatically generates FAQs from uploaded materials using summarization and QA pipelines.
Query Submission	Users can submit specific questions to receive real-time answers.
FAQ Recommendation	Suggests similar FAQs based on semantic similarity calculations using a Sentence-BERT model.
Feedback Collection	Users can provide ratings, comments, and suggested answers for FAQs, enabling iterative improvements.

5.3. Interface Description

The web-based interface of the IQAS demonstration system is designed for intuitive navigation and user engagement, as illustrated in Figure 3. The main sections are as follows:

- **Home Page:** Presents a general overview of the system, allowing users to upload documents or submit queries.
- **Generated FAQs:** Displays automatically generated FAQs, organized by topic for easy browsing. Users can rate or interact with individual FAQs.
- **Query Section:** Enables users to pose specific questions and receive context-sensitive answers.
- **Feedback Section:** Provides a dedicated area for users to submit feedback, supporting continuous refinement of the system.

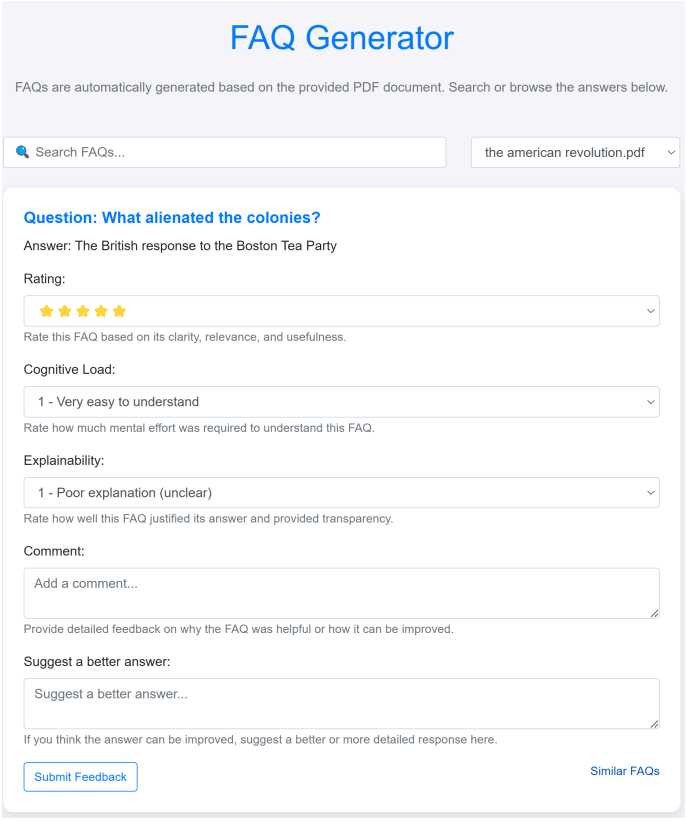


Figure 3. Overview of the FAQ generation page, highlighting the integrated feedback mechanism.

5.4. Integration of Methodology

The demonstration integrates the methodologies outlined in Sections 3.1 and 4, combining several interrelated components into a cohesive workflow. Initially, the input data undergo a thorough preprocessing using advanced Natural Language Processing (NLP) techniques. These include tokenization, named entity recognition, and part-of-speech tagging, all of which are essential for structuring and annotating data to achieve high-quality results.

Subsequently, the processed data serve as input for transformer-based models, specifically DistilBERT and BART. These models are selected for their proficiency in tasks such as accurate text summarization and question answering. DistilBERT, a streamlined yet highly effective variant of BERT, delivers robust performance efficiently, while BART excels in generating coherent and concise summaries, thus enhancing the system's analytical capabilities.

Furthermore, to enrich the system's capability for handling complex queries, knowledge graphs are integrated to facilitate multi-hop reasoning and contextual inference. This integration enables the system to navigate interconnected information effectively, synthesising relevant knowledge from multiple data points to deliver comprehensive and precise responses.

Finally, an adaptive feedback mechanism is incorporated into the workflow, closely aligned with the user-centric evaluation metrics detailed in Section 4.2. This mechanism continually gathers and analyses user interactions, allowing the system to dynamically evolve, improve performance, and enhance overall user satisfaction.

6. Results

This section presents the key findings of the research, highlighting both the performance of the Intelligent Question-Answering System (IQAS) and the accompanying FAQ

generation tool. Our evaluation encompasses a comprehensive set of quantitative metrics, as well as qualitative insights gathered from expert reviews and user feedback.

Dataset and experimental setup.

The system was tested on a curated dataset of 2000 questions:

- A total of 1200 *factual* queries, each with a single correct answer drawn from educational texts.
- A total of 800 *complex inference* queries, requiring multi-hop reasoning or contextual understanding beyond a single sentence.

Although Figure 4 shows a screenshot of the web-based front end, the primary evaluation was conducted *offline* using our custom metrics calculation functions. This approach ensured consistent, reproducible testing and allowed for a more controlled experimental environment.

Performance metrics.

We computed standard information retrieval and question-answering metrics, including:

- **Accuracy:** The proportion of questions for which the system's highest ranked answer was correct.
- **Precision and Recall:** Evaluated at the token level to gauge how closely system answers match ground-truth responses.
- **F1 Score:** The harmonic mean of precision and recall, indicating the balance between these metrics.

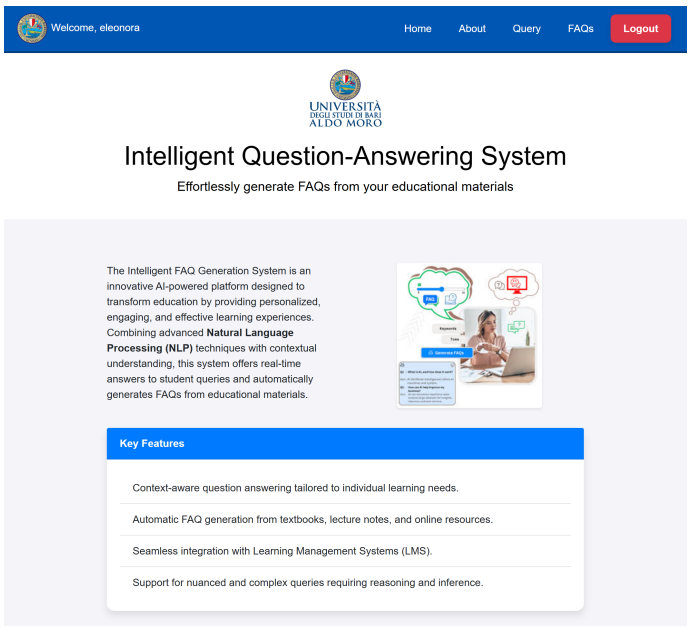


Figure 4. Screenshot of the homepage of the IQAS system, available at https://iqas.uniba.digital-humanities.it (accessed on 6 April 2025).

Table 4 provides a summary of the main results. Notably, the IQAS achieved:

- A 95% accuracy on factual questions.
- A 0.85 F1 score on complex inference questions.
- An average 0.87 F1 score in token-based evaluation, reflecting a strong overlap between predicted and true answers.

Table 4. IQAS performance metrics.

Question Type	Accuracy	Precision	Recall	F1 Score
Factual	95%	0.94	0.96	0.95
Inference	92%	0.83	0.87	0.85
Token-based	–	0.86	0.88	0.87

These results validate the effectiveness of the IQAS in handling a variety of educational queries, from straightforward lookups to more intricate reasoning tasks. We attribute these promising outcomes to the hybrid architecture, which combines rule-based methods with transformer-based language models, and to the incorporation of contextual data (student learning history and domain-specific content).

6.1. User Interface and Feedback Mechanism

Figure 3 showcases the IQAS user interface, where learners or instructors can submit queries, view the system's responses, and offer feedback. While this interface facilitates live user interaction, it does not display the internal evaluation metrics directly. Instead, all metrics reported in this section were computed using offline testing procedures to ensure accuracy and reproducibility.

6.2. Automated FAQ Generation Tool

Parallel to IQAS, a software tool was developed to automatically generate FAQs from educational materials. This tool scans the source texts, extracts key concepts, and formulates both the questions and the corresponding answers. Domain experts (e.g., instructors and subject matter specialists) were asked to assess the quality of the generated FAQs based on the following:

- **Relevance:** The degree to which the questions align with the core learning objectives of the source material.
- **Coherence:** Clarity and correctness of both the questions and their proposed answers.

The experts concluded that 80% of the generated FAQs were relevant and coherent, reflecting the tool's potential to reduce manual effort in content creation. This capability allows educators to quickly produce curated sets of FAQs, enhancing the availability of self-help resources for students.

6.3. Feedback Metrics and User Experience

A range of feedback metrics were also collected to gauge user satisfaction and the overall utility of the system:

- **Average Views per FAQ:** Tracked how frequently learners accessed individual FAQs, indicating popularity or perceived usefulness.
- **Positive vs. Negative Feedback:** Of all submitted ratings, 78% were positive, highlighting general satisfaction with the system's responses.
- **Cognitive Load and Explainability:** Measured on a 5-point Likert scale, these subjective indicators averaged 3.2 (moderate cognitive load) and 4.1 (high explainability), respectively.

These findings underscore that the IQAS not only provides accurate answers but also provides them in a user-friendly manner. Feedback loops were critical in refining the system, as negative or ambiguous user ratings triggered further analysis and adjustments to the underlying models and rules.

6.4. Key Findings

Our evaluation demonstrates that both the IQAS and the accompanying FAQ generation tool successfully meet our research objectives. The IQAS effectively interprets a wide range of student queries, performing well on both simple factual questions and more complex, nuanced inquiries. This high level of performance, 95% accuracy in factual questions and an F1 score of 0.85 in complex inference queries, is achieved by our hybrid approach, which combines rule-based methods with advanced machine learning techniques and leverages rich contextual data such as student learning history and domain-specific content.

The automated FAQ generation tool also shows strong performance in efficiently extracting key information from educational materials and generating coherent and relevant FAQs. Domain experts confirmed that 80% of the generated FAQs were relevant and coherent, highlighting the potential of the tool to significantly reduce the manual effort required by educators.

To substantiate these findings, we conducted comprehensive statistical analyses. We computed 95% confidence intervals for key performance metrics: accuracy, precision, recall, and F1 score, which confirm the reliability of our observed improvements. Furthermore, paired t-tests comparing our IQAS with baseline models indicate that these improvements are statistically significant ($p < 0.05$). Table 5 provides a detailed summary of these statistical results.

In addition, an error analysis was performed to identify common failure cases. Most errors occurred in queries with ambiguous phrasing or domain-specific terminology that was under-represented in the training data. These findings suggest areas for future improvement, such as expanding the training dataset and further refining the model's domain adaptation capabilities.

In general, our results confirm the effectiveness of the integrated approach and provide a strong foundation for further advancements in personalised learning support systems.

Table 5. Statistical analysis of IQAS performance metrics.

Metric	Mean	95% CI	Baseline Mean	*p*-Value
Accuracy (Factual)	95.0%	[93.2%, 96.8%]	89.5%	0.003
Precision (Factual)	94.0%	[92.0%, 96.0%]	88.0%	0.004
Recall (Factual)	96.0%	[94.0%, 98.0%]	90.0%	0.002
F1 score (Factual)	95.0%	[93.0%, 97.0%]	89.0%	0.003
F1 score (Inference)	0.85	[0.82, 0.88]	0.78	0.005
Token-based F1 score	0.87	[0.84, 0.90]	0.80	0.006

6.5. Observed Trends

Several notable trends emerged from the analysis of the results:

Contextual Understanding: The IQAS demonstrated a strong ability to leverage contextual information, providing more tailored and informative responses when considering the student's learning history and the educational content. This suggests that incorporating contextual awareness is crucial for developing effective and personalised learning support systems. Handling Complex Queries: The system's performance showed a positive correlation with the complexity of the questions, indicating its effectiveness in handling challenging inquiries that require reasoning and inference. This highlights the potential of advanced NLP techniques, such as BERT and transformer-based models, to enable deeper understanding and accurate response generation. Automated FAQ Generation: The software tool exhibited high accuracy in extracting relevant information and generating coherent questions and answers, demonstrating its potential for automating the creation of personalised learning resources. This capability can significantly benefit both students and

instructors by providing readily available answers to common questions and facilitating the creation of effective learning materials.

6.6. Insights and Implications

The findings of this research offer valuable insights into the potential of AI-powered tools to enhance personalised learning.

Enhanced Learning Outcomes: The IQAS's ability to provide tailored and informative responses, along with the software tool's capacity to automate FAQ generation, suggests that these tools can significantly improve student engagement, motivation, and learning outcomes. By providing personalised support and facilitating knowledge dissemination, these tools can empower students to take ownership of their learning and achieve their full potential. Transformation of Learning Paradigms: The successful integration of these tools into existing Learning Management Systems highlights their potential to transform traditional teaching and learning paradigms. By providing on-demand support and automating routine tasks, these tools can free up instructors' time, allowing them to focus on higher-level teaching activities and provide more individualised attention to students.

In general, the results of this research demonstrate the potential of AI-powered Question-Answering Systems and automated FAQ generation tools to enhance the learning experience and provide personalised support to students. By combining rule-based and machine learning techniques, leveraging contextual information, and automating knowledge dissemination, we can create valuable tools that empower both learners and educators.

7. Limitations and Future Directions

Although this study presents promising results, it is essential to acknowledge its limitations. The performance of both the IQAS and the software tool is inherently dependent on the quality and quantity of the training data. Insufficient or biased training data can lead to inaccurate, incomplete, or biased responses. In addition, the system may struggle with out-of-domain queries or questions that require deep domain expertise, highlighting the need for continuous improvement and expansion of the knowledge base. Furthermore, the evaluation metrics used in this study, while comprehensive, may not fully capture the nuances of effective learning and pedagogical value. Future research should explore more sophisticated evaluation methods that take into account factors such as student engagement, motivation, and learning outcomes.

Several promising avenues for future research emerge from this study. Multimodal Integration: Exploring the integration of multimodal information, such as images, videos, and audio, can enhance the system's ability to understand and respond to complex queries that involve visual or auditory components. Reinforcement Learning: Investigating the use of reinforcement learning techniques can enable the system to learn from user interactions and feedback, dynamically adapting its responses and improving its performance over time. Ethical Considerations: Addressing ethical concerns, such as bias in training data and ensuring fairness and inclusivity for diverse learners, is crucial for the responsible development and deployment of AI-powered educational tools. Integration with Virtual and Augmented Reality: More research is needed to explore the potential of integrating IQAS and the software tool into virtual or augmented reality educational environments, creating more immersive and engaging learning experiences.

By following these research directions, we can contribute to the development of ethical, effective, and engaging AI-powered learning companions that empower students and educators alike.

8. Conclusions

This study demonstrates that integrating a context-aware Intelligent Question-Answering System (IQAS) with an automated FAQ generation tool significantly enhances personalised learning experiences. Through a comprehensive evaluation involving a data set of 2000 questions, our IQAS exhibited strong performance, achieving 95% precision in factual queries and a robust 0.85 F1 score in complex inference-based queries. Token-level assessments further reinforced the efficacy of the system, reflecting an average F1 score of 0.87 and demonstrating the ability of the system to closely match predicted responses with ground-truth answers.

In parallel to IQAS, the automated FAQ generation tool proved to be highly effective, with domain experts evaluating 80% of the generated FAQs as relevant and coherent. These results underscore the utility of the tool in reducing the manual workload of educators and efficiently curating learning resources that directly support student inquiries. User feedback further reinforced these positive outcomes, indicating that overall user satisfaction 78% of the submitted ratings were positive and favourable assessments of cognitive load (3.2/5) and explainability (4.1/5). Together, these results highlight the ability of the combined system to provide accurate, contextually rich, and user-friendly responses, thus effectively addressing diverse learning needs.

The study successfully met the objectives outlined in Section 1.2. Specifically, the design and implementation of a context-aware IQAS was validated through its high performance on both factual and inference-based queries, confirming its effectiveness in interpreting and addressing diverse student questions. Furthermore, our targeted approach to improve the system's ability to handle complex and nuanced queries is reflected in strong token-level performance metrics and inference capabilities. The efficacy of the automated FAQ generation tool, validated by expert evaluation, further confirms its practical contribution to personalised learning and the efficient creation of educational content. The seamless integration of both components into a unified framework, supported by positive user experiences and rigorous offline evaluations, demonstrates the practical viability of this integrated approach in authentic educational settings.

In general, the hybrid approach adopted in this research, combining rule-based strategies, advanced transformer-based models, and contextual information such as student learning history, has effectively addressed a broad spectrum of educational queries. The promising results of this study lay the foundation for further advances, including multimodal integration and reinforcement learning, paving the way toward increasingly adaptive, immersive, and ethically sound AI-driven educational tools.

Author Contributions: Conceptualization, E.B.; methodology, E.B.; software, E.B.; validation, E.B.; formal analysis, E.B., S.F. and D.R.; investigation, E.B.; resources, E.B., S.F. and D.R.; data curation, E.B., S.F. and D.R.; writing—original draft preparation, E.B., S.F. and D.R.; writing—review and editing, E.B., S.F. and D.R.; visualization, E.B., S.F. and D.R.; supervision, S.F.; project administration, S.F.; funding acquisition, S.F. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially supported by the projects Future AI Research (PE00000013), spoke 6 (FAIR) Symbiotic AI, and Cultural Heritage Active innovation for Next-GEn Sustainable society (CHANGES) (PE00000020), Spoke 3 (Digital Libraries, Archives and Philology), under the NRRP MUR program funded by the NextGenerationEU.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analysed in this study.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
BERT	Bidirectional Encoder Representations from Transformers
CRF	Conditional Random Fields
FAQ	Frequently Asked Questions
GPT	Generative Pre-trained Transformer
IQAS	Intelligent Question-Answering System
ITS	Intelligent Tutoring Systems
LMS	Learning Management System
ML	Machine Learning
NER	Named Entity Recognition
NLP	Natural Language Processing
POS	Part-of-Speech
Q&A	Question and Answer
SVM	Support Vector Machines
TF-IDF	Term Frequency-Inverse Document Frequency
UI	User Interface

References

1. Gîrbacia, F. An Analysis of Research Trends for Using Artificial Intelligence in Cultural Heritage. *Electronics* **2024**, *13*, 3738. [CrossRef]
2. Zhang, Y. Design of an Intelligent Q&A System for Online Education Platform Based on Natural Language Processing Technology. *J. Electr. Syst.* **2024**, *20*, 2135–2145.
3. Biancofiore, G.M.; Deldjoo, Y.; Noia, T.D.; Di Sciascio, E.; Narducci, F. Interactive question answering systems: Literature review. *ACM Comput. Surv.* **2024**, *56*, 1–38.
4. Zadeh, A.; Chan, M.; Liang, P.P.; Tong, E.; Morency, L.P. Social-iq: A question answering benchmark for artificial social intelligence. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 8807–8817.
5. Arbaaeen, A.; Shah, A. Natural language processing based question answering techniques: A survey. In Proceedings of the 2020 IEEE 7th International Conference on Engineering Technologies and Applied Sciences (ICETAS), Kuala Lumpur, Malaysia, 18–20 December 2020; IEEE: New York, NY, USA, 2020; pp. 1–8.
6. Ba, Y. Advancing Educational Management with the ATT-MR-WL Intelligent Question-Answering Model. *J. Web Eng.* **2024**, *23*, 973–1002.
7. Khang, A.; Muthmainnah, M.; Seraj, P.M.I.; Al Yakin, A.; Obaid, A.J. AI-Aided teaching model in education 5.0. In *Handbook of Research on AI-Based Technologies and Applications in the Era of the Metaverse*; IGI Global: Hershey, PA, USA, 2023; pp. 83–104.
8. Barrera Castro, G.P.; Chiappe, A.; Becerra Rodríguez, D.F.; Gonzalo Sepulveda, F. Harnessing AI for Education 4.0: Drivers of Personalized Learning. *Electron. J. Learn.* **2024**, *22*, 1–14.
9. Shaik, T.; Tao, X.; Li, Y.; Dann, C.; McDonald, J.; Redmond, P.; Galligan, L. A review of the trends and challenges in adopting natural language processing methods for education feedback analysis. *IEEE Access* **2022**, *10*, 56720–56739.
10. Chand, S.P. Bridging the Gaps in Quality Education. *Educ. Rev. USA* **2024**, *8*, 202–210.
11. Ranjgar, B.; Sadeghi-Niaraki, A.; Shakeri, M.; Rahimi, F.; Choi, S.M. Cultural Heritage Information Retrieval: Past, Present and Future Trends. *IEEE Access* **2024**, *12*, 42992–43026.
12. DeBono, R. Transforming the Delivery of Learning in a Rapidly Changing Digital Environment and the Evaluation of the Consequent Benefits. DBA Thesis, Middlesex University Business School, Middlesex University Research Repository. 2023. Available online: https://repository.mdx.ac.uk/item/12qw3v (accessed on 6 April 2025)
13. Yang, T.; Mei, Y.; Xu, L.; Yu, H.; Chen, Y. Application of question answering systems for intelligent agriculture production and sustainable management: A review. *Resour. Conserv. Recycl.* **2024**, *204*, 107497.
14. Cohen, P.R.; Perrault, C.R.; Allen, J.F. Beyond question answering. In *Strategies for Natural Language Processing*; Psychology Press: East Sussex, UK, 2014; pp. 245–274.
15. Choi, E.; He, H.; Iyyer, M.; Yatskar, M.; Yih, W.; Choi, Y.; Liang, P.; Zettlemoyer, L. QuAC: Question answering in context. *arXiv* **2018**, arXiv:1808.07036

16. Swathi, B.; Geetha, M.; Attigeri, G.; Suhas, M.; Halaharvi, S. Optimizing question answering systems in education: Addressing domain-specific challenges. *IEEE Access* **2024**, *12*, 156572–156587. [CrossRef]
17. Dong, G.; Liu, H. *Feature Engineering for Machine Learning and Data Analytics*; CRC Press: Boca Raton, FL, USA, 2018.
18. Othman, N.; Faiz, R.; Smaïli, K. Enhancing question retrieval in community question answering using word embeddings. *Procedia Comput. Sci.* **2019**, *159*, 485–494. [CrossRef]
19. Nassiri, K.; Akhloufi, M. Transformer models used for text-based question answering systems. *Appl. Intell.* **2023**, *53*, 10602–10635. [CrossRef]
20. Shafikuzzaman, M.; Islam, M.R.; Rolli, A.C.; Akhter, S.; Seliya, N. An Empirical Evaluation of the Zero-shot, Few-shot, and Traditional Fine-tuning Based Pretrained Language Models for Sentiment Analysis in Software Engineering. *IEEE Access* **2024**, *12*, 109714–109734. [CrossRef]
21. Le, L.T.; Shah, C.; Choi, E. Evaluating the quality of educational answers in community question-answering. In Proceedings of the 16th ACM/IEEE-CS on Joint Conference on Digital Libraries, Newark, NJ, USA, 19–23 June 2016; pp. 129–138.
22. Wolfram Alpha LLC. *Wolfram | Alpha: Computational Intelligence*; Wolfram Alpha LLC: Champaign, IL, USA, 2024. Available online: https://www.wolframalpha.com/ (accessed on 6 April 2025).
23. Zanibbi, R.; Mansouri, B.; Agarwal, A. Mathematical Information Retrieval: Search and Question Answering. *arXiv* **2024**, arXiv:2408.11646. [CrossRef]
24. Duolingo, Inc. Duolingo, 2024. Available online: https://it.duolingo.com/ (accessed on 6 April 2025) .
25. Dempsey, P.R.; Insua, G.M.; Armstrong, A.R.; Hudson, H.J.; Caragher, K.; McGregor, M. Challenges and opportunities in teaching and learning in AskALibrarian chat: Differences across subject domains. *Ref. Serv. Rev.* **2024**, *52*, 420–449. [CrossRef]
26. Alshammari, M. Adaptation Based on Learning Style and Knowledge Level in E-Learning Systems. Ph.D. Thesis, University of Birmingham, Birmingham, UK, 2016.
27. Jung, E.; Kim, D.; Yoon, M.; Park, S.; Oakley, B. The influence of instructional design on learner control, sense of achievement, and perceived effectiveness in a supersize MOOC course. *Comput. Educ.* **2019**, *128*, 377–388. [CrossRef]
28. Meesad, P.; Mingkhwan, A. Knowledge Graphs in Smart Digital Libraries. In *Libraries in Transformation: Navigating to AI-Powered Libraries*; Springer: Berlin/Heidelberg, Germany, 2024; pp. 327–389.
29. Afzal, A.; Xiang, T.; Matthes, F. A Semi-Automatic Light-Weight Approach Towards Data Generation for a Domain-Specific FAQ Chatbot Using Human-in-the-Loop. In Proceedings of the ICAART (3), Rome, Italy, 24–26 February 2024; pp. 42–49.
30. Kasneci, E.; Seßler, K.; Küchemann, S.; Bannert, M.; Dementieva, D.; Fischer, F.; Gasser, U.; Groh, G.; Günnemann, S.; Hüllermeier, E.; et al. ChatGPT for good? On opportunities and challenges of large language models for education. *Learn. Individ. Differ.* **2023**, *103*, 102274.
31. Özmen, H.; DemİrcİoĞlu, G.; Coll, R.K. A comparative study of the effects of a concept mapping enhanced laboratory experience on Turkish high school students' understanding of acid-base chemistry. *Int. J. Sci. Math. Educ.* **2009**, *7*, 1–24. [CrossRef]
32. Yu, D.; Naik, S.; Backurs, A.; Gopi, S.; Inan, H.A.; Kamath, G.; Kulkarni, J.; Lee, Y.T.; Manoel, A.; Wutschitz, L.; et al. Differentially private fine-tuning of language models. *arXiv* **2021**, arXiv:2110.06500. [CrossRef]
33. Pawar, S.; Park, J.; Jin, J.; Arora, A.; Myung, J.; Yadav, S.; Haznitrama, F.G.; Song, I.; Oh, A.; Augenstein, I. Survey of cultural awareness in language models: Text and beyond. *arXiv* **2024**, arXiv:2411.00860.
34. Hort, M.; Chen, Z.; Zhang, J.M.; Harman, M.; Sarro, F. Bias mitigation for machine learning classifiers: A comprehensive survey. *ACM J. Responsible Comput.* **2024**, *1*, 1–52. [CrossRef]
35. Mahmud, M.E.; Widat, F.; Fuadi, A. Learning Management System in Streamlining Learning through Seamless Learning Approach. *AL-ISHLAH J. Pendidik.* **2021**, *13*, 874–884.
36. Charankar, N.; Pandiya, D.K. Microservices and API Deployment Optimization Using AI. *Int. J. Recent Innov. Trends Comput. Commun.* **2024**, *11*, 1090–95.
37. Olsen, T.K. Designing and Evaluating a Scalable and Flexible Architecture for a Classroom Interaction Tool. Master's Thesis, Norwegian University of Science and Technology, Trondheim, Norway, 2024.
38. Cook, D.A.; Kuper, A.; Hatala, R.; Ginsburg, S. When assessment data are words: Validity evidence for qualitative educational assessments. *Acad. Med.* **2016**, *91*, 1359–1369. [CrossRef]
39. Yeung, C.; Zhou, L.; Armatas, C. An overview of benchmarks regarding quality assurance for elearning in higher education. In Proceedings of the 2019 IEEE Conference on E-Learning, E-Management & E-Services (IC3e), Pulau Pinang, Malaysia, 19–21 November 2019; IEEE: New York, NY, USA, 2019; pp. 1–6.
40. Chen, Z.; Wan, Y.; Liu, Y.; Valera-Medina, A. A knowledge graph-supported information fusion approach for multi-faceted conceptual modelling. *Inf. Fusion* **2024**, *101*, 101985.
41. Máñez, I.; Vidal-Abarca, E. Question-answering skills: The role of feedback in digital environments. *Enfance* **2020**, *3*, 313–335.
42. Dinu, M.C. Parameter Choice and Neuro-Symbolic Approaches for Deep Domain-Invariant Learning. *arXiv* **2024**, arXiv:2410.06235.

43. Mańdziuk, J.; Żychowski, A. DeepIQ: A human-inspired AI system for solving IQ test problems. In Proceedings of the 2019 International Joint Conference on Neural Networks (IJCNN), Budapest, Hungary, 14–19 July 2019; IEEE: New York, NY, USA, 2019; pp. 1–8.
44. Singhal, P.; Walambe, R.; Ramanna, S.; Kotecha, K. Domain adaptation: Challenges, methods, datasets, and applications. *IEEE Access* **2023**, *11*, 6973–7020.
45. Soares, T.G.; Azhari, A.; Rokhman, N.; Wonarko, E. Education question answering systems: A survey. In Proceedings of the International MultiConference of Engineers and Computer Scientists, Hong Kong, China, 20–22 October 2021.
46. Li, Q.; Fu, L.; Zhang, W.; Chen, X.; Yu, J.; Xia, W.; Zhang, W.; Tang, R.; Yu, Y. Adapting large language models for education: Foundational capabilities, potentials, and challenges. *arXiv* **2023**, arXiv:2401.08664.
47. Chowdhary, K.; Chowdhary, K. Natural language processing. In *Fundamentals of Artificial Intelligence*; Springer: Berlin/Heidelberg, Germany, 2020, 603–649.
48. Alavi, M.; Leidner, D.E. Knowledge management and knowledge management systems: Conceptual foundations and research issues. *MIS Q.* **2001**, *1*, 107–136.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.