

BABEȘ BOLYAI UNIVERSITY, CLUJ NAPOCA, ROMÂNIA
FACULTY OF MATHEMATICS AND COMPUTER SCIENCE

Virtual Assistants, Ontologies and AI: A Waltz in Education Innovation

– PhD Thesis –

PhD student: **Ioana-Alexandra Todericiu**
Scientific supervisor: Prof. Laura Dioșan

2025

Abstract

The time for virtual assistants is now. For more than half a decade, Artificial Intelligence (AI) has been making its way into the spotlight, and today, we are all taking notice. AI has delivered its promise. In November 2022, ChatGPT took the stage by storm. The application, which reached 1 million users in just 5 days, made it clear that virtual assistants are not just a once-believed far-off future but a concrete reality. This thesis, which began two years prior to the "GPT phenomenon", has gained even greater relevance in light of the current technological landscape.

AI-powered tools are transforming industries, and education is no exception to the the rule. Through current innovations in virtual assistants, personalized, accessible, and student-centered learning can be achieved. The only limit is our imagination.

This thesis aims to provide a novel approach that bridges the fields of virtual assistants and education in a cohesive and structured manner. A comprehensive literature review is conducted to explore the theoretical concepts surrounding virtual assistants, including their taxonomy and applications. Subsequently, existing efforts aimed at formalizing the connection between virtual assistants and the educational landscape are examined. Given that this field is relatively new and less explored from a formal perspective, this thesis seeks to contribute by offering a robust, platform-agnostic ontology designed for the development of educational tools based on virtual assistants. The power of a unified blueprint lies in its versatility and the unique ways it can be applied. To illustrate this, the thesis presents three concrete solutions, selected based on empirical studies that uncovered students needs that could be addressed via virtual assistants.

In conclusion, this thesis aims to lay a foundation for a future in which learning is enhanced through the use of intelligent tools. By embracing the possibilities of AI and reimagining how education can be delivered, we move closer to a reality where no student is left behind, and every question finds an answer. The future of learning is here—virtual, accessible, and only a voice command away.

Contents

1	Introduction	1
1.1	Context and Motivation	1
1.2	Original Contribution	2
1.3	Challenges and Opportunities	3
1.4	List of Publications	5
2	Theoretical Foundations: Ontologies, Virtual Assistants and Smart Speakers	7
3	Proposed approaches	9
3.1	Virtual Assistants Ontology	9
3.2	Ontology Instances	12
3.2.1	Quiz Skill - Alexa	12
3.2.2	University Assistant Skill - Alexa	13
3.2.3	University Assistant Skill - MS Copilot	13
3.2.4	Mentorship for Female Students Skill - MS Copilot	14
3.3	A Comparative Framework for Virtual Assistants	15
3.4	Validation	16
3.4.1	Validation Methodology	16
3.4.2	Pilot Validation	16
4	Conclusions and Future work	17

Chapter 1

Introduction

1.1 Context and Motivation

In education, a field rooted in tradition yet increasingly influenced by innovation, the need for personalized, adaptive, and accessible learning tools has never been more pressing. Traditional systems often struggle to meet the unique needs of individual learners, deliver timely and constructive feedback, or provide equitable access for all students. Technology, particularly the one anchored in artificial intelligence, offers immense potential to address these challenges [HT24], and its adoption in education has been piecemeal. The integration of ontologies and virtual assistants (VAs) presents a promising yet underexplored avenue for transforming the learning landscape.

Ontologies provide a structured framework for organizing and representing knowledge [VCGO23], ensuring consistency, adaptability, and clarity. Virtual assistants, on the other hand, excel at creating interactive, dynamic, and engaging environments by leveraging natural language processing and conversational AI [MVP21]. Combining these technologies could redefine education—making it more tailored to individual learning styles, responsive to diverse needs, and inclusive of underrepresented groups. This thesis is driven by the belief that technology should empower education to adapt to the learners it serves.

While the potential of these tools is evident, there are significant gaps in how they have been applied [AAAA25, LGM23]. Existing research often explores ontologies and VAs as separate domains, focusing on their individual strengths rather than their combined capabilities. Ontologies have been used to enhance interoperability and formalize knowledge across systems [HB20], and VAs have demonstrated their utility in improving user engagement and access [Vij24, ISSD20]. However, the intersection of these technologies, particularly in education, remains underdeveloped. There is a lack of formalized frameworks that integrate ontologies with VAs to create personalized, adaptive, and scalable solutions [EGPCB19]. Additionally, little attention has been given to validating such frameworks in diverse educational settings or exploring their role in addressing challenges like accessibility and inclusivity [TLCCLM18].

This thesis aims to bridge these gaps by proposing an ontology-driven framework that enhances the capabilities of virtual assistants for education. The research focuses on tailoring ontologies to the educational domain, defining the relationships and attributes that underpin

personalized learning experiences. It also maps these theoretical constructs to practical use cases, such as creating quiz-based knowledge assessments [SDMGH⁺22], providing administrative support, and delivering personalized study recommendations. The proposed solution is validated through empirical studies, assessing its impact on accessibility, scalability, and overall learning outcomes [SPJ20].

By addressing these gaps, the thesis seeks to provide a new lens through which education can leverage the power of technology. The anticipated impact is twofold: first, creating a more engaging and effective learning experience for students; and second, offering a blueprint for researchers and developers to explore the integration of ontologies and VAs in other domains. This work is a step toward a future where education is as dynamic, responsive, and inclusive as the world it prepares students to navigate.

1.2 Original Contribution

This work presents a structured framework that integrates ontology-driven design with the adaptive and interactive capabilities of virtual assistants, specifically tailored for educational environments. By leveraging AI, natural language processing, and scalable cloud technologies, the research proposes innovative solutions that address challenges such as personalized learning, instant feedback, and enhanced accessibility.

The exploration on this topics started 5 years ago, right before the global pandemic, in 2020, and it gained momentum as the gaps it was addressing were becoming hard to ignore. As students refrained from human interaction, this thesis explored available technologies that mimicked the dialog exchange typically experienced between individuals. The aim was to identify a solution that could potentially work for all, reaching even users with disabilities who were traditionally under-served by conventional software. The thesis pursued the development of a customized solution, tailored to address the dynamic needs of students.

The initial exploration began with Alexa, culminating in the publication of the first research paper on the topic [ST20]. This early work helped define the trajectory that would guide the subsequent stages of this thesis. During an academic conference, a key question was raised: how might this approach expand to other technologies and diverse educational needs? In essence, how could others adapt the underlying concept to different technology stacks and specific requirements? This perspective proved eye-opening. Beyond the pandemic context, additional opportunities were identified where virtual assistants could contribute meaningfully. In the following years, the work centered on developing a unique ontology designed to serve as a blueprint capable of answering three fundamental questions:

- What concepts specific for Virtual Assistants need to be implemented in order to provide a solution that tackle a specific need?
- How technical and non-technical aspects of a solution relate and influence each other when it comes to integrating virtual assistants in education?

- Can the ontology be validated on a set of unique use-cases and work for it, regardless of the technology stack and solution scope?

The ontology serves as a guiding framework, while virtual assistants—especially those designed for the educational environment—are the primary focus of our research. While formalizing the ontology, we explored possible instances and their values, from administrative information retrieval (classes schedules, announcements, news) to knowledge verification and validation (gamification of quizzes). The ontology steps and guidelines were followed, and during this process, the classes, attributes, and relationships of the ontology were further refined.

Starting from the premise that the target audience should be vast, diverse, and inclusive, different categories of students and their specific educational needs were also studied. Two empirical studies were conducted [TSV21, MATZ22] to gain a better understanding of the intended users and to define their personas. The first study [MATZ22] examined the factors influencing women’s persistence and attrition in computer science education and careers, with a focus on Romanian higher education institutions. It revealed critical barriers and motivators that inform strategic interventions to address gender disparities and sustain academic interest. The second study [TSV21] investigated students’ perceptions of their involvement in the learning process, highlighting how active participation, such as generating questions and collaborative learning, impacts motivation and knowledge retention. This approach provided a deeper understanding of the pedagogical methods that resonate with students’ learning dynamics. Together, these studies offer a comprehensive view of the diverse challenges and opportunities in creating more inclusive and effective educational practices.

1.3 Challenges and Opportunities

This thesis aims to address the following research questions:

RQ1: What are the key elements and methodologies required for the formalization of virtual assistants in education?

RQ2: What instances of virtual assistants, with a focus on education, demonstrate the practical benefits of formalized architectural solutions?

RQ3: How can the proposed solutions be systematically evaluated to measure their performance?

RQ4: How does the proposed solution extend and integrate existing approaches in knowledge representation, learning, and accessibility?

These research questions collectively address the identified gaps in the literature by exploring the foundational methodologies for formalizing virtual assistant solutions (RQ1), demonstrating

their real-world applicability in education (RQ2), and systematically evaluating their effectiveness (RQ3). Additionally, by comparing the proposed solutions with existing approaches (RQ4), this work provides a comprehensive framework for advancing virtual assistant development in education. The answers to these questions are detailed throughout the thesis: foundational methodologies are discussed in Section 4.1, practical use cases in Section 4.2, evaluation metrics in Section 4.3, and comparative analysis in Chapter 3. Together, these findings contribute to bridging the gap between theoretical innovation and practical implementation in virtual assistant applications for education.

1.4 List of Publications

Below is the list of publications derived from the research conducted during the PhD program.

1. **Ioana Alexandra Todericiu**, Camelia Șerban, Laura Dioșan. "Towards Accessibility in Education through Smart Speakers: An Ontology-Based Approach." *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 25th International Conference KES2021, Procedia Computer Science*, vol. 192, pp. 883-892, 2021/1/1. DOI: <https://doi.org/10.1016/j.procs.2021.08.091>.
Paper Type: Regular Paper
Conference Category: Conference of Category B
Score: 4
Citations: 11
2. **Ioana Alexandra Todericiu**, Camelia Șerban, Andreea Vescan. "Students' Perception on the Impact of Their Involvement in the Learning Process: An Empirical Study." In *Proceedings of the 3rd International Workshop on Education through Advanced Software Engineering and Artificial Intelligence*, pp. 39-46, 2021/8/23. DOI: <https://doi.org/10.1145/3472673.3473964>.
Paper Type: Regular Paper
Conference Category: Workshop of an A conference*
Score: 6
Citation: 4
3. Simona Motogna, Lenuta Alboae, **Ioana Alexandra Todericiu**, Cristina Zaharia. "Retaining Women in Computer Science: The Good, the Bad and the Ugly Sides." In *Proceedings of the Third Workshop on Gender Equality, Diversity, and Inclusion in Software Engineering*, pp. 35-42, 2022/5/19. DOI: <https://doi.org/10.1145/3524501.3527598>.
Paper Type: Regular Paper
Conference Category: Workshop of an A conference*
Score: 3
Citations: 14
4. **Ioana Alexandra Todericiu**, Pop Mihai Daniel, Camelia Șerban, Laura Dioșan. "Quiz-Ifying Education: Exploring the Power of Virtual Assistants." *16th International Conference on Computer Supported Education*, 2024. DOI: <https://doi.org/10.5220/0012722400003693>.
Paper Type: Short Paper
Conference Category: Conference of Category B
Score: 4/3
5. **Ioana Alexandra Todericiu**, Laura Dioșan, Camelia Serban. "Alexa and Copilot: A Tale of Two Assistants" *17th International Conference on Agents and Artificial*

Intelligence, 2025. DOI: TBD.

Paper Type: Short Paper

Conference Category: Conference of Category B

Score: 8/3

6. **Ioana Alexandra Todericiu.** "Virtual Assistants: A Review of the Next Frontier in AI Interaction" *Acta Universitatis Sapientiae*, 2025. DOI: TBD.

Paper Type: Journal Paper

Journal Category: Journal D (Indexed Web Of Science)

Score: 1

Total Points: 18

Chapter 2

Theoretical Foundations: Ontologies, Virtual Assistants and Smart Speakers

The field of ontologies and abstractization in the context of virtual assistants in higher education is a rapidly evolving field. According to Wiederhold and Thomsen (2019) [WT19], current research in this area has focused on developing ontologies that can be used to represent knowledge and skills in various domains, such as course content and student information.

Ontologies, within the realm of computer and information sciences, are formal representations of a set of concepts and the relationships between them within a particular domain [Gru93]. They serve as structured frameworks that enable the sharing and reuse of knowledge across different applications and groups. By providing a common vocabulary and a set of agreed-upon definitions, ontologies facilitate interoperability between systems and improve communication between stakeholders [FSZ23].

From a formal perspective, an ontology defines a set of representational primitives with which to model a domain of knowledge or discourse [Gru09]. Each element which constitutes the formal representation of an ontology, such as the classes, objects, instances, inference rules contribute towards a generalization of different topics. They offer a common language for describing the architecture, and the logic in any area.

Virtual assistants (VAs) are an essential subset of Question Answering Systems (QASs), evolving from simple information retrieval tools into sophisticated interactive systems [BDDN⁺24].

A classification of virtual assistants, originally proposed in [Tod25], is presented, based on interaction types, task specificity, deployment methods, and decision-making frameworks. These categories provide a comprehensive framework for understanding the diverse functionalities and applications of VAs in various domains.

Smart speakers are functioning based on virtual assistants, which represent the brain or core of their "judgment". In a simplistic manner, virtual assistants can be perceived as a software mechanism that intercepts the command of the user (speech recognition and natural language understanding), maps it to a request, and, when needed, calls the service that is responsible for handling that action, based on which it formulates a proper response to the user [BLS⁺18].

Despite the widespread integration of Virtual Assistants (VAs) and ontologies into various

domains, including education, a significant gap persists in their systematic formalization and application. Existing systems tend to focus either on the development of standalone ontologies or the implementation of VAs without addressing the interplay between these two fields. This disconnect limits their potential to create transformative, scalable solutions, particularly in the field of education, where personalized, adaptive, and accessible learning environments are urgently needed.

Chapter 3

Proposed approaches

A novel approach is proposed in Chapter 3, introducing an ontology-based framework that bridges the gap between virtual assistants and the academic domain. This framework formalizes the interaction between intelligent systems and educational processes, focusing on addressing students' and educators' needs. By leveraging virtual assistants like Amazon Alexa and Microsoft Copilot, the solution integrates advanced AI capabilities, such as natural language processing and cloud-based services, to provide tailored skills for knowledge evaluation, retrieval, and suggestion. These skills enable personalized learning, streamline academic workflows, and enhance accessibility for diverse users. The technical architecture combines ontology-driven design with scalable cloud services to deliver an adaptive, efficient, and interactive solution that reshapes how students and educators engage with educational systems. A validation methodology and an initial pilot are described at the end of the chapter.

3.1 Virtual Assistants Ontology

The development of the ontology began with a study of best practices and methodologies from various sources [NM01, SLPS21]. Although no universally mandatory guideline exists for creating a high-quality ontology, several recommended phases have been identified to support the systematic organization of the domain. The initial phase focuses on defining the scope and understanding the knowledge domain.

The scope of an ontology is to set boundaries and refine the knowledge domain relevant for it. In respect to this, we propose the expectations that the ontology must fulfill as a list of the competencies [BLR19] as follows:

- S1:** Set a bridge between virtual assistants and the academic sphere.
- S2:** Define the users that facilitate from the connection bridge between virtual assistants and university.
- S3:** Demystify the technical architecture needed for a solution based on a skill for virtual assistants oriented towards the university.
- S4:** Clearly identify what needs the virtual speaker is responding to.

Table 3.1 presents several main concepts from domain knowledge that are common among other existing ontologies, with some focusing exclusively on virtual assistants, others on education, and a few addressing both areas. The first three listed papers [Had13, AAA⁺21, ZAzJS19] from the below table focus on an university ontology, whereas the last two [MMM19, WACY17] focus on an ontology based on virtual assistants. In respect of the topics presented in Table 3.1, the proposed ontology is comprised of classes and sub-classes created based on the presented topics that all the listed papers have in common.

Study	Academic Staff	Student	Quiz	Intelligence/ Knowledge	Conversational Actions	Virtual Assistant
<i>Hadjar (2013)</i> [Had13]	✓	✓				
<i>Alrehaili et al. (2021)</i> [AAA ⁺ 21]	✓	✓	✓	✓		
<i>Zeebaree et al. (2019)</i> [ZAzJS19]	✓	✓				
<i>Maier et al. (2019)</i> [MMM19]				✓	✓	✓
<i>Wessel et al. (2017)</i> [WACY17]					✓	✓

Table 3.1: Common Patterns in Domain Knowledge in Literature

For the development phase of the ontology, the main tool used was Protege [Mus15, GMF⁺03]. Protege is notorious for being the go-to software tool when it comes to ontologies [ZZZ12], offering a wide variety of features that ease not only the development phase, but also the validation. Given the complexity of the structure, snapshots alone were deemed insufficient to fully capture the design of the proposed ontology.

The visualization of ontologies is crucial for effective management, browsing, and understanding [KHL⁺07], but can present significant challenges for large and complex structures [TH06]. In the search for an appropriate software solution for ontology visualization, the OWLViz plugin [CSM09] offered by Protégé was initially evaluated; however, its outdated design and limited user-friendliness were identified as major limitations. Consequently, the interactive web-based application WebVOWL¹ [LLMN15, WLA18] was adopted, offering a more convenient and user-friendly environment. WebVOWL supports ontology development directly within the platform and provides various features for graphical representation, including adjustable granularity, dynamic text box sizing, zoom controls, and others. Nevertheless, despite the advantages of WebVOWL’s graphical interface, visualizing a highly complex ontology continues to pose challenges in terms of clarity.

Instances can be regarded as individuals or concrete representation of the ontology. The instances of an ontology help us map the formal definition of a system to real scenarios. These instances offer an overview of the architecture of our proposed solution: 3 skills that are custom-made for students, that addresses basic needs: keeping up with the administrative and organizational updates, testing their knowledge by taking quizzes and receiving suggestions of what they should study and look into.

We define the consistency of an ontology as the absence of ambiguity, that may lead to poor static assertions and relationships.

In order to check for the consistency validation, the HermiT reasoner was used [SMH08]. Reasoner is a visual representation of inconsistent behavior in the ontology (such as disjointness

¹<https://service.tib.eu/webvowl/>

that causes inconsistencies - which is one of the most common axiom that causes this is happen. In order for an inconsistency to occur, there must be at least 2 axioms.). It displays in red the errors from the lower to the upper level class. In order to avoid possible propagation of errors, reasoner was used in all the phases of design and development process. HermiT reasoner comes as a built-in plugin in Protege, and can be triggered directly from it.

In order to address the scope, we will use the classes and the relationships that link them as follows:

In order to address the scope, we will use the classes and the relationships that link them as follows:

S1: Set a bridge between virtual assistants and the academic sphere.

S1 Explained: The bridge between virtual assistants and the academic sphere can be regarded as a custom *virtual assistant skill* created for a university, to which we can refer to as *university skill*.

S2: Define the users that benefit from the connection bridge between virtual assistants and university.

S2 Explained: The users benefiting from the university skill are the *students*, as well as the *academic staff*. Students are the direct beneficiaries, as the skill answers to their need, and the academic staff benefit indirectly from it by having an interactive tool to support their teaching processes.

S3: Demystify the technical architecture needed for a solution based on a skill for virtual assistants oriented towards the university.

S3 Explained: A skill, more specifically our custom skill called *University Skill*, is developed in *Skill Development Platform*, which is a system offered by the virtual assistant provider in order to add new capabilities (skills) to the VA. In order to formulate a response, a cloud function needs to be called. That means that we need to decide which *cloud* provider to leverage, from Google Cloud, Microsoft Azure to Amazon Web Services and others. These providers offer services that respond to our design needs, more precisely: a *historical log service* which logs the conversational actions and possible failures, a *serverless compute service* that is responsible for constructing a *response/prompt* and can manipulate *database or cloud storage* services if needed.

S4: Clearly identify what needs the virtual assistant is responding to.

S4 Explained: The most important need that the virtual assistant is responding to (by using a university skill) is the need for *knowledge*. It does that by offering organizational knowledge that helps students navigate through their daily academic life as well as professional knowledge that helps them prepare better. It also comes in handy for academic staff, as it helps *understanding needs* and supports the *teaching process*.

3.2 Ontology Instances

This section presents several practical implementations of ontology-based virtual assistants developed for educational use. Each subsection explores a distinct skill or assistant, detailing its purpose, target users, technical architecture, and conversational flow. The implementations vary in platform (Alexa [LRK⁺19] vs. Microsoft Copilot Studio [Str24]) and educational focus (quizzes, university services, mentorship), offering a comparative perspective on how ontological modeling can support adaptive, voice-enabled educational tools.

3.2.1 Quiz Skill - Alexa

The proposed approach, "Quiz for Uni", is a skill developed for Alexa, designed to augment personalized learning through interactive quizzes. The primary goal of the skill is to provide students with a stimulating learning tool that adapts to their individual needs and preferences. We will outline the key stages involved in the development process, including the selection of technologies, the modeling of the conversational flow, and the implementation details.

We selected Alexa as our interface due to its wide user base and the Alexa Developer Console, which streamlines skill development. This console offers comprehensive tools for creating, testing, and deploying interactive voice applications, making it ideal for our quiz system. Additionally, the integration between Alexa and AWS is seamless as both are Amazon products. This synergy allows for efficient and coherent development, with AWS providing robust backend services like Lambda and DynamoDB, ensuring scalability, reliability, and ease of management for 'Quiz for Uni'.

This solution is developed for second-year students in Software Engineering (SE), with a focus on enhancing their understanding of the Java programming language and object-oriented programming (OOP) concepts. At this stage of their education, students are transitioning from basic programming skills to more complex concepts and practices in SE. The Alexa-based interactive quiz system is tailored to meet the educational needs of this learning phase.

A preliminary pilot validation was conducted to assess the skill, presented in greater detail in section 3.4.2. The pilot study involved six volunteers, including three current third-year Computer Science students and three alumni, selected on a voluntary basis. Data collection was carried out through a structured questionnaire combining Likert-scale ratings and open-ended questions, completed after a two-week interaction with the Alexa quiz skill. Quantitative data were analyzed using descriptive statistics, while qualitative responses were thematically analyzed via open coding [BC21]. Results showed high user satisfaction (average 8.6/10), strong interest in future use (9.33/10), and positive feedback on usability, with suggestions for improvements.

The choice of this skill was inspired by previous work [TSV21], that studied the impact of quizzes on students academic performance. The study was conducted on 83 students. The purpose of the study was to investigate how a complementary Project-Based Learning method—specifically, the “Students Generating Questions” (SGQ) strategy—impacts student engagement and learning outcomes in Software Engineering courses at Babeş-Bolyai University.

The methodology combined a mixed-methods design, which includes a structured survey (employing Likert scales, ranking questions, and open-ended items) administered to both undergraduate and master’s students, with rigorous quantitative analysis and qualitative thematic coding using open coding techniques to evaluate the collected data. The results indicate that the SGQ method helped students better regulate their learning, improve their understanding of course concepts, reduce exam anxiety, and strengthen collaboration among peers.

3.2.2 University Assistant Skill - Alexa

The "University Assistant Skill - Alexa" functions as a novel instrument aimed at improving the daily academic experiences of students. This ability utilizes Alexa’s voice interaction features to provide smooth communication between users and their educational environment. The main aim is to streamline logistical academic responsibilities, including schedule management, class notifications, and communication in virtual learning settings. Integrating Alexa with popular platforms like the official university website, Microsoft Teams and Outlook enables students to rapidly retrieve essential information, including class schedules, assignments, and future meetings, with straightforward voice requests.

The Amazon ecosystem was selected for this skill because of its support for voice-controlled applications and its simple connection with AWS services. AWS Lambda guarantees that the backend of the "University Assistant Skill" is scalable and responsive, efficiently processing voice inputs and providing precise responses in real time. DynamoDB offers a dependable and adaptable database solution for managing fluctuating schedules and resources. Microsoft Power Automate connects voice commands to Microsoft Office 365 services, facilitating efficient workflows for tasks like posting updates on Teams or summarizing emails. This combination of technology guarantees that the "University Assistant Skill" provides a unified and effective experience, customized to address the dynamic requirements of modern education.

3.2.3 University Assistant Skill - MS Copilot

The "University Assistant Skill - MS Copilot" aims to deliver a fluid and adaptable solution for academic activities by utilizing features of Microsoft Copilot Studio [Che24]. It employs a probabilistic methodology, informed by large language models (LLMs) [ZZL⁺23], to provide dynamic and context-sensitive interactions, in contrast to deterministic virtual assistants such as Alexa. The skill is linked with the Microsoft environment, providing powerful capabilities such real-time collaboration, contextual replies, and improved flexibility to various linguistic and academic contexts. Students can employ the skill for tasks such as classes management, accessing academic information, and collaboration.

Microsoft Copilot Studio was selected as the basis for this skill due to its connectivity with Microsoft Services and the use of generative algorithms. The probabilistic methodology allows the assistant to enhance responses according to user behavior, fostering a progressively evolving learning experience. The integration of retrieval-augmented generation for retrieving information from many sources boosts its capabilities, allowing users to engage with both

structured and unstructured data effortlessly. The platform used for development (MS Copilot Studio) facilitates deployment across several platforms, including Microsoft Teams, custom applications, chat applications (Whatsapp, MS Teams), and web-based interfaces, thereby assuring accessibility and scalability.

3.2.4 Mentorship for Female Students Skill - MS Copilot

The initiative, "Mentorship for Female Students," aims to identify and tackle the specific issues encountered by female students in computer science, as outlined in the paper "Retaining Women in Computer Science: The Good, the Bad, and the Ugly Sides" [MATZ22]. This study performed a thorough examination of factors affecting women's retention in computer science studies, employing data from 125 MSc students, 31 PhD students, and 46 female academic personnel throughout Romania. It identified obstacles, including lack of mentorship, insufficient representation, and restricted access to role models. The study employed a comprehensive, longitudinal methodology that integrates both quantitative surveys and qualitative open coding of free-text responses. The design, which evolved through iterative refinement based on pilot feedback, included structured questionnaires, and rigorous validation procedures (involving independent coding and consensus-building) to ensure the reliability and validity of our findings.

Quantitative responses provided essential insights into the representation of mentorship among female students in computer science. Among MSc students, 91 female students indicated they did not benefit from mentorship programs, whereas only 34 benefited, reflecting a mere 27.2% participation rate. Furthermore, merely 5.6% of MSc students had access to programs explicitly designed to support women (7 out of 125 participants). Alarming, hardly 2.4% reported receiving assistance from teachers. When questioned about their decision to discontinue their studies, 15.2% cited the lack of a mentor at their university as the reason.

Among PhD students, 51.7% (15 of 31 participants) availed themselves of mentorship programs, whereas 93.5% reported a lack of access to women-specific institutional support programs (29 of 31 participants). Among postdoctoral researchers, 21.3% (10 of 47 participants) had mentors, while merely 2.1% engaged with women-focused programs (1 of 47 participants).

These findings underscore the urgent necessity for accessible and scalable mentorship programs, especially for underrepresented demographics in the industry. The primary objective is to facilitate an efficient mentorship pairing process. Upon identifying a female student in need of assistance —either through an explicit request for mentorship help or because the underlying model determines that mentorship is the optimal response — the Copilot initiates a dialogue to determine her particular requirements and preferences. The skill subsequently cross-references this information with a designated database of mentors, organized by areas of expertise, and launches a customized mentorship introduction via email. The email serves as an introduction between the mentor and the mentee. This automation improves the accessibility and effectiveness of mentorship programs, in accordance with the study's proposal to establish institutional mechanisms that more effectively support female students.

3.3 A Comparative Framework for Virtual Assistants

Although Amazon Alexa excels in its primary language, issues with pronunciation and comprehension may limit its efficacy. In contrast, Microsoft Copilot Studio advantages from its underlying design that accommodates several languages, along with the adaptability provided by GPT, facilitating communication in languages not explicitly enumerated, but its performance is not always excellent. As educational institutions increasingly cater to different populations, the capacity to deliver appropriate language support will be essential for the efficacy of virtual assistants in improving the learning experience.

Upon finalizing the development of a skill for Amazon Alexa, the final product is the Alexa skill, which can be deployed on the Alexa platform. Developers must utilize the Alexa Developer Console to publish a skill, allowing for comprehensive testing to verify that its functionality and user experience adhere to their criteria. Upon completion of testing, developers submit the skill for certification. This certification procedure guarantees that the skill adheres to Amazon’s standards for privacy, security, and usability[CA23]. Upon successful certification, the skill becomes available to users on many Alexa-enabled devices, including Echo speakers, smartphones, and tablets. This deployment enables students to engage with the skill through voice commands, facilitating convenient access to information and help. A student can inquire, “Alexa, what classes do I have today?” and obtain prompt, tailored responses derived from the skill’s programming and data integration. Moreover, Copilot solutions can be implemented on either a demonstration website, provided by Microsoft, or a custom website. In addition to conventional applications, Copilot Studio facilitates deployment across several platforms, such as Slack, custom mobile apps, Telegram, and Direct Line Speech.

Amazon Alexa generally attains an average response time of 2.5 seconds for simple queries that roughly correspond with established intentions. When users conform to the prescribed language, the skill can provide results very instantaneously. Nevertheless, if the inquiry diverges from the specified intents, the response time may extend to 3–4 seconds while the system try to digest the input and align it with the closest intent, query the appropriate dataset and formulate a response. Microsoft Copilot Studio typically attains an average response time of 2–3 seconds for basic orders, which is comparable to the performance of Amazon Alexa. For intricate inquiries necessitating comprehensive analysis or contextual comprehension, response times may extend to 4-5 seconds. The diversity in response time is mostly due to the utilization of GPT (Generative Pre-trained Transformer) technology, enabling Copilot to comprehend a wider array of user inputs. Although GPT improves the assistant’s capacity to produce nuanced and contextually appropriate responses, this probabilistic characteristic may result in extended processing durations as the model evaluates and adjusts to user behavior [RCD+20]. Therefore, the abundance of interactions provided by Copilot results in occasional delays, especially when addressing complex or multifaceted inquiries.

When users stick to with established intents, Alexa attains a 100% accuracy rate for inquiries. This deterministic approach ensures dependable responses, as the system is engineered to provide accurate answers when user commands correspond with the designated intents. Nevertheless, if the inquiries are reformulated in a way that deviates from the established

expressions for each intent—specific terms that the assistant is trained to identify—the accuracy rate may diminish. Microsoft Copilot Studio reveals its versatility and promise in educational environments, while also highlighting key areas for improvement, especially in hallucinations, precision and contextual comprehension. Confronting these problems will be crucial for optimizing the assistant’s efficacy and improving the educational experience for students.

3.4 Validation

3.4.1 Validation Methodology

To evaluate the proposed virtual-assistant “skills,” we designed a flexible, three-pronged validation framework that can be reused across different educational use-cases. Participants—recruited from the same year and specialization to keep academic context uniform—received clear usage instructions and interacted with the assistant in Microsoft Teams (or on a smart speaker) for two–three weeks. Data were gathered through (1) a post-trial questionnaire capturing satisfaction, usability, and perceived learning gains; (2) internal telemetry logs recording real-world command frequency, duration, and flow; and (3) optional external engagement tools such as an IDE time-tracker to see whether VA use translates into deeper coding activity. Analysis combined descriptive statistics and thematic coding of open-ended answers with correlation of subjective ratings against objective telemetry, yielding a holistic picture of effectiveness, usability, and adoption feasibility while remaining GDPR-compliant and easily replicable for future studies.

3.4.2 Pilot Validation

A first exploratory pilot focused on the quiz skill, supplying six volunteers (three current CS students, three alumni) with pre-configured Alexa Echo devices for two weeks. Participants completed the structured questionnaire only, producing early but encouraging evidence: enjoyment averaged 4.5/5, technical performance 4/5, overall satisfaction 8.6/10, and likelihood of future use 9.3/10. Qualitative feedback praised the game-like, self-navigating flow and its ability to motivate further study, while calling for customizable quiz length/difficulty and better speech-recognition for non-native accents. Primary threats to validity stem from the small, self-selected sample and the absence of objective behavioral metrics; future larger-scale studies will add telemetry and IDE time-tracking to triangulate findings and strengthen generalizability.

Chapter 4

Conclusions and Future work

The research journey began with a goal that evolved and took on different shapes and forms over time. Initially, the destination was supposed to be the technical intricacies behind the creation of virtual assistants for the educational sector. It was soon clear that this static goal is not enough in order to capture the dynamic nature of both the technical landscape and the educational system. Challenges and changes started to arise, from COVID-19, from the launch of ChatGPT. A singular approach would simply not do. But what if that approach would be robust enough to be able to extent to new solutions, opportunities and to the unknown future?

The ontology perspective offered a unique appeal and turned out to be the missing piece of the puzzle. After comprehensive research of the intersection between virtual assistants and education, the surprise did not disappoint and arrived just right. Even though both fields have gained significant momentum individually, their combination remains scarce. The field was clearly a new one, waiting to be discovered. Therefore, this works build on top of existing work, with a strong theoretical foundation. On top of it, the thesis introduces a novel taxonomy for virtual assistant and smart speakers, that paves the way to a better understanding of the concepts that will play a key role in the ontological formalization.

The backbone of the study is represented by the unique ontology that we developed. It can be regarded as a ceremony between the opportunities that the educational sector holds and the promises that the virtual assistants deliver. The ontology covers a vigorous collection of classes, relationships, properties and instances.

Instances are the highlight of the formalization, as they prove how the ontology can be applied on a wide range of use cases, with different technologies and technical stacks. Behind the instances are full fledged empirical studies and prove their value and represent the motivation behind the specific needs they are addressing. Using the ontology, we follow closely the technical landscape, the conversational flow and the particularities of each solution. This thesis does not subscribe to one technical spectrum, being platform-agnostic, as it follows different technologies and SaaS providers, proving that the ontology is not restrained by a specific technology or business need.

To further expand on the how the impact can be measured, we provide a robust methodology, generic enough so that it can be applied to any of the virtual assistants based application

described in this thesis. The methodology provides information on participants selection, data collection and data analysis. Based on the methodology, an initial pilot validation was described. In the pilot validation phase, a carefully designed questionnaire gathered subjective feedback—assessing aspects such as enjoyment, technical performance, overall satisfaction, and likelihood of future use, on a small group of users. Despite the limited sample size, the pilot phase yielded valuable insights that confirm the effectiveness of the ontology-driven framework and its potential for scalability.

Although this thesis represents an important milestone, the work in this area remains ongoing. This work serves as a foundation for advancing new technical endeavors. Looking optimistically towards the future, the following areas of research are worth mentioning:

- The piloting of the ontology’s instances is the natural next step, which will enable the validation and refinement of its practical applications. The piloting of the ontology’s instances is the natural next step, which will enable the validation and refinement of its practical applications. This phase will help assess its usability, scalability, and effectiveness in real-world scenarios, paving the way for further advancements and innovations. This will inspire the enhancements of the ontology, making it more robust, complex and future-ready.
- The ontology enhancement is a topic that can be refined every day. Therefore, a future ramification of the existing work oughts to find ways in which this evolution of the current ontology can be done dynamically, leveraging machine learning algorithms to adapt to changing educational paradigms, students needs and technical advancements in real time.
- Expanding the ontology to accommodate a wider variety of use-cases and technical complexities can tap new opportunities. A richer set of classes, properties and relations translates into more applicability and guidance. This scaling can be done horizontally, to expand to new technologies and business needs, but at the same time, there is value in vertical scaling, with an emphasize or greater granularity and details.
- AI Agents are poised to be the virtual assistants of tomorrow, and the future work aims to align with and advance this trajectory. Understanding their potential to enhance user experiences, improve decision-making processes and support across a wide range of domain is the first step in the exploration journey. Their technical intricacies, new design paradigms they introduce, as well as integration with different systems are some potential area of study, that later may translate in an enhanced novel ontology, a pillar for a detailed examination into this area.

As virtual assistants continue to push the boundaries of innovation, creativity, and productivity, the future will be defined by our relentless pursuit of ways to harness this technology to our advantage — a vision this thesis strives to realize.

Bibliography

- [AAA⁺21] Nada Alrehaili, Muhammad Aslam, Dimah Alahmadi, Dina Alrehaili, Muhammad Asif, and Muhammad Malik. Ontology-Based Smart System to Automate Higher Education Activities. *Complexity*, 2021:1–20, 08 2021.
- [AAAA25] Safa Albo Abdullah and Ahmed Al-Azawei. Predicting online learners’ performance through ontologies: A systematic literature review. *The International Review of Research in Open and Distributed Learning*, 26:16–37, 02 2025.
- [BC21] V. Braun and V. Clarke. *Thematic Analysis: A Practical Guide*. SAGE Publications, 2021.
- [BDDN⁺24] Giovanni Biancofiore, Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, and Fedelucio Narducci. Interactive Question Answering Systems: Literature Review. *ACM Computing Surveys*, 56, 04 2024.
- [BLR19] Maricela Bravo, Hoyos-Reyes Luis, and José Reyes. Methodology for ontology design and construction. *Contaduría y Administración*, 64:134, 03 2019.
- [BLS⁺18] Frank Bentley, Chris Luvogt, Max Silverman, Rushani Wirasinghe, Brooke White, and Danielle Lottridge. Understanding the Long-Term Use of Smart Speaker Assistants. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2(3), September 2018.
- [CA23] Sudip Chakraborty and Sreeramana Aithal. Let Us Create an Alexa Skill for Our IoT Device Inside the AWS Cloud. *International Journal of Case Studies in Business, IT, and Education*, pages 214–225, 05 2023.
- [Che24] Wei-Yu Chen. Intelligent tutor: Leveraging chatgpt and microsoft copilot studio to deliver a generative ai student support and feedback system within teams. *arXiv preprint arXiv:2405.13024*, 2024.
- [CSM09] Nadia Catenazzi, Lorenzo Sommaruga, and Riccardo Mazza. User-Friendly Ontology Editing and Visualization Tools: The OWLeasyViz Approach. In *13th International Conference Information Visualisation*, pages 283–288, 07 2009.
- [EGPCB19] Lenin Erazo-Garzón, Andrés Patiño, Priscila Cedillo, and Alexandra Bermeo. CALMS: A Context-Aware Learning Mobile System Based on Ontologies. In *2019 Sixth International Conference on eDemocracy & eGovernment*, pages 84–91, 2019.

- [FSZ23] Karim Farghaly, Ranjith K. Soman, and Shanjing Alexander Zhou. The evolution of ontology in AEC: A two-decade synthesis, application domains, and future directions. *Journal of Industrial Information Integration*, 36:100519, 2023.
- [GMF⁺03] John H Gennari, Mark A Musen, Ray W Ferguson, William E Grosso, Monica Crubézy, Henrik Eriksson, Natalya F Noy, and Samson W Tu. The evolution of Protégé: an environment for knowledge-based systems development. *International Journal of Human-Computer Studies*, 58(1):89–123, 2003.
- [Gru93] Thomas R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.
- [Gru09] Tom Gruber. *Ontology*, pages 1963–1965. Springer US, Boston, MA, 2009.
- [Had13] Karim Hadjar. University Ontology: case study Ahlia University. *Semantic Web: Implications for Technologies and Business Practices*, pages 50–54, 07 2013.
- [HB20] Martina Husáková and Vladimír Bureš. Formal Ontologies in Information Systems Development: A Systematic Review. *Information*, 11(2), 2020.
- [HT24] Ryan Hare and Ying Tang. Ontology-driven Reinforcement Learning for Personalized Student Support. *2024 IEEE International Conference on Systems, Man, and Cybernetics*, abs/2407.10332:2555–2560, 2024.
- [ISSD20] Vinayak Iyer, Kshitij Shah, Sahil Sheth, and Kailas Devadkar. Virtual assistant for the visually impaired. In *2020 5th International Conference on Communication and Electronics Systems*, pages 1057–1062, 2020.
- [KHL⁺07] Akrivi Katifori, Constantin Halatsis, George Lepouras, Costas Vassilakis, and Eugenia Giannopoulou. Ontology Visualization Methods—a Survey. *ACM Comput. Surv.*, 39(4):10–es, nov 2007.
- [LGM23] Lasha Labadze, Maya Grigolia, and Lela Machaidze. Role of ai chatbots in education: systematic literature review. *International Journal of Educational Technology in Higher Education*, 20:1–17, 10 2023.
- [LLMN15] Steffen Lohmann, Vincent Link, Eduard Marbach, and Stefan Negru. WebVOWL: Web-based Visualization of Ontologies. In *Knowledge Engineering and Knowledge Management: EKAW 2014 Satellite Events, VISUAL, EKM1, and ARCOE-Logic, Linköping, Sweden, November 24-28, 2014. Revised Selected Papers.*, volume 8982, pages 154–158, 04 2015.
- [LRK⁺19] Irene Lopatovska, Katrina Rink, Ian Knight, Kieran Raines, Kevin Cosenza, Harriet Williams, Perachya Sorsche, David Hirsch, Qi Li, and Adrianna Martinez. Talk to me: Exploring user interactions with the Amazon Alexa. *Journal of Librarianship and Information Science*, 51(4):984–997, 2019.
- [MATZ22] Simona Motogna, Lenuța Alboai, Ioana Alexandra Todericiu, and Catrinel Zaharia. Retaining women in computer science: the good, the bad and the ugly sides. In *Proceedings of the Third Workshop on Gender Equality, Diversity, and Inclusion in Software*

- Engineering*, GE@ICSE '22, page 35–42, New York, NY, USA, 2022. Association for Computing Machinery.
- [MMM19] Torsten Maier, Jessica Menold, and Christopher McComb. Towards an Ontology of Cognitive Assistants. *Proceedings of the Design Society: International Conference on Engineering Design*, 1:2637–2646, 07 2019.
 - [Mus15] Mark A. Musen. The Protégé Project: A Look Back and a Look Forward. *AI Matters*, 1(4):4–12, jun 2015.
 - [MVP21] Ann Mathew, Rohini .V, and Joy Paulose. NLP-based personal learning assistant for school education. *International Journal of Electrical and Computer Engineering*, 11:4522–4530, 10 2021.
 - [NM01] N. Noy and Deborah McGuinness. Ontology Development 101: A Guide to Creating Your First Ontology. *Knowledge Systems Laboratory*, 32, 01 2001.
 - [RCD⁺20] Daniele Rogora, Antonio Carzaniga, Amer Diwan, Matthias Hauswirth, and Robert Soulé. Analyzing system performance with probabilistic performance annotations. In *Proceedings of the Fifteenth European Conference on Computer Systems*, pages 1–14, 04 2020.
 - [SDMGH⁺22] Asya Stoyanova-Doycheva, Sebiha Madanska, Maria Grancharova-Hristova, Todorka Glushkova, and Georgi Cholakov. Development of Ontologies in Different Domains for a Test Generation Environment. In *8th International Conference on Higher Education Advances*, pages 925–933, 06 2022.
 - [SLPS21] Alexander Smirnov, Tatiana Levashova, Andrew Ponomarev, and Nikolay Shilov. Methodology for multi-aspect ontology development. In Uchitha Jayawickrama, Pavlos Delias, María Teresa Escobar, and Jason Papathanasiou, editors, *Decision Support Systems XI: Decision Support Systems, Analytics and Technologies in Response to Global Crisis Management*, pages 97–109, Cham, 2021. Springer International Publishing.
 - [SMH08] Rob Shearer, Boris Motik, and Ian Horrocks. HermiT: A highly-efficient OWL reasoner. In *Proceedings of the Fifth Web Ontology Language Experiences and Directions Workshop on OWL: Experiences and Directions, collocated with the 7th International Semantic Web Conference (ISWC-2008), Karlsruhe, Germany, October 26-27, 2008*, volume 432, 01 2008.
 - [SPJ20] Kristian Stancin, Patrizia Pošćić, and Danijela Jaksic. Ontologies in education – state of the art. *Education and Information Technologies*, 25:5301–5320, 11 2020.
 - [ST20] Camelia Serban and Ioana Alexandra Todericiu. Alexa, What classes do I have today? The use of Artificial Intelligence via Smart Speakers in Education. *Procedia Computer Science*, 176:2849–2857, 2020. Knowledge-Based and Intelligent Information and Engineering Systems: Proceedings of the 24th International Conference 2020.
 - [Str24] Jess Stratton. An introduction to microsoft copilot. In *Copilot for Microsoft 365: Harness the Power of Generative AI in the Microsoft Apps You Use Every Day*, pages 19–35. Springer, 2024.

- [TH06] Yannis Tzitzikas and Jean-Luc Hainaut. On the Visualization of Large-Sized Ontologies. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, AVI '06, page 99–102, New York, NY, USA, 2006. Association for Computing Machinery.
- [TLCCLM18] Mariela Tapia-Leon, Abdon Carrera, Janneth Chicaiza, and Sergio Luján-Mora. Application of ontologies in higher education: A systematic mapping study. In *2018 IEEE Global Engineering Education Conference*, pages 1344–1353, 04 2018.
- [Tod25] Ioana Alexandra Todericiu. Virtual assistants: A review of the next frontier in ai interaction. *Acta Universitatis Sapientiae*, 2025.
- [TSV21] Ioana Todericiu, Camelia Serban, and Andreea Vescan. Students perception on the impact of their involvement in the learning process: an empirical study. In *Proceedings of the 3rd International Workshop on Education through Advanced Software Engineering and Artificial Intelligence*, EASEAI 2021, page 39–46, New York, NY, USA, 2021. Association for Computing Machinery.
- [VCGO23] William Villegas-Ch and Joselin García-Ortiz. Enhancing Learning Personalization in Educational Environments through Ontology-Based Knowledge Representation. *Computers*, 12(10):1–19, 2023.
- [Vij24] M. Vijayakumar. A Study on Chatbots and Virtual Assistants in Customer Engagement: A Review. *International Journal of Engineering and Management Research*, 14:204–208, 03 2024.
- [WACY17] Michael Wessel, Girish Acharya, James Carpenter, and Min Yin. OntoVPA - An Ontology-Based Dialogue Management System for Virtual Personal Assistants. In *Advanced Social Interaction with Agents : 8th International Workshop on Spoken Dialog Systems*, pages 219–233. Springer International Publishing, 2017.
- [WLA18] Vitalis Wiens, Steffen Lohmann, and S. Auer. WebVOWL Editor: Device-Independent Visual Ontology Modeling. In *International Workshop on the Semantic Web*, volume 2180, 2018.
- [WT19] Stefan Wiederhold and Christian Thomsen. A Survey of Personalized Learning Support Systems. In *Proceedings of the 12th International Conference on Educational Data Mining, 2019, Montreal, Canada, June 19-22, 2019*, pages 3–14, 2019.
- [ZAJS19] Subhi Zeebaree, Adel Al-zebari, Karwan Jacksi, and Ali Selamat. Designing an Ontology of E-learning system for Duhok Polytechnic University Using Protégé OWL Tool. *Journal of Advanced Research in Dynamical and Control Systems*, 11:5, 01 2019.
- [ZZL⁺23] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 1(2), 2023.
- [ZZZ12] Huiqun Zhao, Shikan Zhang, and Junbao Zhao. Research of Using Protégé to Build Ontology. In *2012 IEEE 11th International Conference on Computer and Information Science*, pages 697–700, 2012.