



Università telematica delle  
Camere di Commercio Italiane

**PHD PROGRAMME IN  
BIG DATA AND ARTIFICIAL INTELLIGENCE**  
*Curriculum "Artificial Intelligence for Industry 5.0 and the Circular Economy"*  
38th Cycle

*PhD Dissertation in  
AI-Driven Learning Assistants for Reflective, Adaptive, and  
Inclusive Education*

Stefano D'Urso

**Programme Coordinator**

Prof.ssa Barbara Martini

**Supervisor**

Prof. Filippo Sciarrone

**Co-Supervisor**

Prof. Andrea Mazzitelli

Academic Year 2024/ 2025

# Table of Contents

Table of Contents.....	2
List of Acronyms.....	6
1 Introduction.....	1
1.1 Context and Motivation .....	1
1.2 Problem Statement .....	2
1.3 Research Aim and Objectives .....	3
1.4 Original Contribution and Impact of the Study .....	4
Research Path.....	6
1.5 .....	6
1.5.1 Phase I – Systematic Literature Review (SLR). .....	7
1.5.2 Phase II – Early Prototyping: Adaptive and Collaborative Intelligence.....	7
1.5.3 Phase III – Multimodal and Action-Oriented Learning.....	8
1.5.4 Phase IV – Inclusive and Reflective Intelligence. ....	8
1.5.5 Integration and Validation. ....	8
1.6 Significance of the Study .....	9
1.7 Thesis Structure.....	10
2 Scope.....	11
2.1 Research Focus and Boundaries .....	11
2.2 Research Questions (RQs) .....	11
2.3 Conceptual Framework.....	12
2.4 Methodological Boundaries .....	13
2.4.1 Research Design.....	13
2.4.2 Data Collection and Evaluation .....	14
2.4.3 Scope and Limitations.....	15
2.5 Integration of the Research Path .....	15
2.5.1 Evolution of the Research.....	15

2.5.2	Synthesis and Coherence .....	17
2.5.3	Research Coherence and Continuity .....	17
2.6	Delimitations and Assumptions .....	17
2.6.1	Delimitations .....	18
2.6.2	Assumptions.....	18
2.6.3	Risks and Critical Reflections.....	19
2.7	Expected Contributions.....	20
2.7.1	Scientific Contributions .....	20
2.7.2	Technological Contributions.....	20
2.7.3	Educational and Societal Contributions.....	21
2.7.4	Overall Impact.....	21
3	Methodology .....	23
3.1	Methodological Orientation .....	23
3.2	Systematic Literature Review .....	23
3.3	Design Based Research Framework .....	24
3.4	Data Collection and Analysis.....	25
3.5	Ethical and Methodological Considerations .....	25
3.6	Methodological Limitations.....	26
4	Results.....	27
4.1	Introduction.....	27
	Results of the Systematic Literature Review .....	28
4.2	.....	28
4.2.1	Purpose and Scope .....	28
4.2.2	Search Strategy and Corpus .....	28
4.2.3	Classification and Trends.....	29
4.2.4	Interpretation.....	29
4.3	Early Prototyping: Adaptive and Collaborative Intelligence .....	30
4.3.1	Objective and Context.....	30

4.3.2	EduPort Prototype .....	31
4.3.3	Boulez Prototype.....	32
4.3.4	Industrial Collaboration (ReSkilla).....	33
4.3.5	Outcomes .....	35
4.4	Multimodal and Action-Oriented Learning .....	35
4.4.1	Objective .....	35
4.4.2	AI4Architect Prototype .....	36
4.4.3	Intent-Based LLM Framework .....	37
4.4.4	Methodological Impact .....	38
4.5	Inclusive and Reflective Intelligence .....	39
4.5.1	Introduction.....	39
4.5.2	Targeted Literature Review on AI, Dyslexia, and Inclusive Learning .....	40
4.6	AI4LA: Fine-Tuned Conversational Assistant for Dyslexic Students.....	47
4.6.1	System Architecture and Functionalities .....	48
4.6.2	Fine-Tuning Process .....	49
4.6.3	Evaluation and Results.....	50
4.6.4	Contribution .....	50
4.7	DyslexIA: Readability-Aware and Reflective Learning Environment.....	50
4.7.1	System Overview .....	52
4.7.2	Conversation Engine .....	54
4.7.3	Learning Modes .....	55
4.7.4	Context Management .....	60
4.7.5	Knowledge Visualization.....	62
4.7.6	Readability and Accessibility Layer .....	65
4.7.7	Evaluation .....	67
4.8	Complementary Research Threads .....	70
4.8.1	EEG-Based Cognitive Framework.....	70
4.8.2	Criterion: Rubric-Aligned Evaluation of Open-Ended Responses .....	71

5	Discussion .....	73
5.1	Reflectivity: From Dialogue to Metacognition .....	73
5.2	Adaptivity: Context, Language, and Cognitive State .....	75
5.3	Inclusivity: Technology as a Bridge for Diversity .....	78
5.4	Assessment and Fairness: The Human-in-the-Loop Approach .....	80
5.5	Ethical and Theoretical Implications .....	83
5.6	Limitations and Future Directions .....	86
6	Conclusions and Next Steps.....	90
	Bibliography .....	92
	Acknowledgements .....	102
	Appendix A – Example of Concept-Map Generation and Merging.....	103
	Appendix B – Examples of Behavioural Fine-Tuning Conversations .....	105
	Appendix C – Tutor Mode Concept Coverage and Reflective Transition .....	107
	Appendix D – Prototype Access Information.....	109

## List of Acronyms

<b>Acronym</b>	<b>Meaning</b>
AI	Artificial Intelligence
AID-LA	AI Driven-Learning Assistant
AID-RLF	AI Driven Reflective Learning Framework
AIED	Artificial Intelligence in Education
CLT	Cognitive Load Theory
CM	Concept Map
DBR	Design-Based Research
GenAI	Generative Artificial Intelligence
ICAP	Interactive-Constructive-Active-Passive
ITS	Intelligent Tutoring Systems
IVH	Improved Vee Heuristic
LA	Learning Analytics
LLM	Large Language Models
ML	Machine Learning
MOOC	Massive Open Online Course
RAG	Retrieval Augmented Generation
RQ	Research Question
SLD	Specific Learning Disorder
SLR	Systematic Literature Review
UDL	Universal Design for Learning

# 1 Introduction

## 1.1 Context and Motivation

Over the past decade, the digital transformation of education has profoundly reshaped how knowledge is produced, distributed, and experienced. The diffusion of online courses(Coursera Inc., 2025; edX Inc., 2025; FutureLearn Ltd., 2025; Udemy Inc., 2025), virtual classrooms(Google LLC, 2025; Microsoft Corporation, 2025; Zoom Video Communications Inc., 2025), and learning platforms(Blackboard Inc., 2025; Instructure, Inc., n.d.; Kajabi LLC, 2025; Moodle Pty Ltd, n.d.; Teachable, Inc., n.d.), accelerated by the COVID-19 pandemic(Kara, 2021), has revealed both the potential and the fragility of digital education. Students today expect learning environments that are flexible, adaptive, and accessible anytime and anywhere(Almansor & Hussain, 2019; Peng et al., 2019). At the same time, educators face growing challenges in providing individualized guidance and feedback to increasingly heterogeneous cohorts.

In this scenario, Artificial Intelligence (AI) has emerged as a key enabling technology. From early Intelligent Tutoring Systems (ITS) to large-scale Learning Analytics (LA) platforms, AI has been progressively integrated into teaching and assessment processes.

A decisive breakthrough occurred with the introduction of the Transformers(Vaswani et al., 2017) architecture, which revolutionized natural language processing by enabling large-scale parallel training and contextual attention mechanisms. This development paved the way for Generative AI (GenAI) and Large Language Models (LLMs) such as OpenAI's ChatGPT(OpenAI, Inc., n.d.-b), Anthropic's Claude(Anthropic PBC, n.d.), Google's Gemini(Google LLC, n.d.), and Meta's LLaMA(Meta Platforms, Inc., n.d.), marking a qualitative leap in the field. These models support natural-language dialogue, multimodal reasoning, and contextual understanding, opening new avenues for human-machine interaction in education. Their ability to adapt explanations, generate examples, and assist in knowledge construction positions them as potential cognitive partners rather than mere automation tools.

At the policy level, this work is closely aligned with the vision of Industry 5.0, which places the worker at the centre of production processes and aims to generate benefits for industry, workers, and society as a whole. In this framework, the European Commission identifies two strategic challenges for the coming years in the field of education and training: (i) adopting a human-centric approach to digital technologies, including artificial intelligence, and (ii) up-skilling and re-skilling the European workforce, with particular emphasis on digital competences(European Commission, n.d.).

These priorities compel all stakeholders involved in education to rethink teaching and learning models, market strategies, and the design of courses and learning materials, which increasingly must be multimedia, personalized, and socially grounded, while explicitly including learners with disabilities and supporting them in developing their competences.

This evolution aligns with the vision of Learning 5.0, where AI complements, not replaces, the educator, acting as a cognitive assistant that personalizes learning experiences and promotes reflection and self-regulation. The new paradigm places emphasis on inclusion, empathy, and human–AI collaboration, extending beyond performance optimization toward human-centred intelligence. However, while LLMs demonstrate remarkable linguistic fluency, their integration into educational practice still raises critical questions of pedagogical grounding, interpretability, and ethical transparency.

## **1.2 Problem Statement**

Despite rapid progress, current AI-based learning systems exhibit notable limitations. Many ITS architectures remain rule-based or domain-specific, lacking the flexibility required for open, interdisciplinary contexts. Generative models, on the other hand, often produce accurate yet pedagogically shallow responses, focusing on information retrieval rather than cognitive engagement.

Drawing from the state-of-the-art analysis, my research defined four strategic areas of innovation in Artificial Intelligence in Education (AIED):

1. *Personalization of learning*: adaptive delivery that accounts for learner’s prior knowledge, learning pace, and motivation.

2. *Accessibility and inclusion*: multimodal interfaces that accommodate sensory, linguistic, or cognitive diversity.
3. *Reflectivity and metacognition*: mechanisms that encourage learners to articulate reasoning, self-assessment, and critical thinking.
4. *Automation of feedback and evaluation*: AI-driven analytics capable of producing timely, explainable, and fair assessments.

Existing approaches typically address one or two of these dimensions in isolation only. For instance, adaptive learning engines rely on performance metrics but seldom support self-reflection; accessibility tools improve usability but rarely connect to conceptual understanding. Conversely, many generative-AI chatbots provide fluent answers without grounding them in verified or pedagogically structured content.

Therefore, there is a clear research gap in designing integrated AI-Driven Learning Assistants (AID-LAs) that combine reflectivity, adaptivity, and inclusivity within a single, coherent architecture. Bridging this gap requires rethinking how generative models interact with educational content, user profiles, and cognitive frameworks.

### **1.3 Research Aim and Objectives**

The overarching goal of my research is to define, implement, and validate a model of AID-LA capable of supporting learners and teachers across diverse educational settings through reflective, adaptive, and inclusive interaction.

To achieve this goal, this research formulated the following objectives:

1. *Conceptualization*: develop an AI-Driven Reflective Learning Framework (AID-RLF), grounded in cognitive and pedagogical theory, defining the core functions of an intelligent educational assistant:
  - *Adaptive Delivery*: dynamic adjustment of content, style, and complexity.
  - *Reflective Scaffolding*: promotion of self-questioning, reasoning, and self-assessment.
  - *Inclusive Interaction*: accessibility through multimodal and readability-aware design.
  - *Feedback and Guidance*: contextual, personalized feedback based on the learner's progress.

2. *Prototyping and experimentation*: design multiple working prototypes integrating LLMs, Retrieval-Augmented Generation (RAG)(Gao et al., 2023; Y. Guo et al., 2024; Lewis et al., 2020), and fine-tuning techniques; validate their usability in real educational environments.
3. *Integration of evidence-based pedagogy*: combine AI methods with theoretical frameworks such as the Cognitive Load Theory (CLT)(Paas et al., 2003; Van Merriënboer & Sweller, 2005), Interactive-Constructive-Active-Passive (ICAP)(Chi & Wylie, 2014) model, Improved Vee Heuristic (IVH)(Novak, 2010; Novak & Gowin, 1984) and Universal Design for Learning (UDL)(Coffman & Draper, 2022; Fornauf & Erickson, 2020) to ensure cognitive coherence.
4. *Empirical validation*: evaluate system performance and learner perception through mixed qualitative-quantitative methods, including readability metrics, interaction logs, and user feedback.

#### **1.4 Original Contribution and Impact of the Study**

This dissertation contributes a unified and empirically grounded model for AID-LAs that integrate reflectivity, adaptivity, and inclusivity within a single pedagogical and computational framework. While existing studies in AIED, LAs, and accessibility technologies tend to address these dimensions separately, this work demonstrates, through theory, system design, and experimental validation, that they can be coherently combined.

The main original contribution is the formulation of the AID-RLF, a framework that operationalises cognitive and pedagogical principles (CLT, ICAP, IVH, UDL) into design constraints for conversational, multimodal, and context-grounded AI systems. This shifts the role of AI from a content-delivery tool to an epistemic mediator that supports learners in constructing meaning, externalising reasoning, and monitoring understanding.

A second contribution lies in the methodological distinction between behavioural fine-tuning and retrieval-based grounding, which is demonstrated experimentally across different prototypes. The thesis shows that fine-tuning is effective for regulating linguistic accessibility, empathy, and cognitive load (particularly for dyslexic learners) while RAG mechanisms are essential for factual

alignment and contextual coherence. This separation clarifies two roles that are often conflated in current AI educational systems.

The research also provides a complete architectural implementation of reflective tutoring, realised in DyslexIA through Tutor Mode and Reflective Mode. These modes guide learners from initial comprehension to deeper metacognitive engagement, supported by knowledge visualisation tools (concept maps and mind maps) that externalise understanding and make conceptual structures manipulable. This contributes to a growing but still underdeveloped line of inquiry on reflective AI systems capable of supporting higher-order cognition.

A further original contribution is the development of readability-aware and accessibility-oriented GenAI systems for dyslexic learners. The targeted literature review reveals that current research on GenAI for dyslexia is sparse, fragmented, and rarely grounded in educational theory. DyslexIA is the first system to combine bilingual behavioural fine-tuning, RAG-based contextualisation, adaptive summarisation, multimodal accessibility tools, and reflective scaffolding in a single environment. The empirical results confirm measurable improvements in readability, comprehension, and cognitive manageability for dyslexic students. Finally, the dissertation extends these principles to assessment through Criterium, a rubric-aligned, explainable, human-in-the-loop evaluation system for open-ended responses. This tool demonstrates how transparency and pedagogical grounding can be preserved in AI-supported assessment, addressing concerns about fairness and automation bias.

Taken together, these contributions offer a meaningful advance for the literature. They provide (i) a theoretical model for reflective and inclusive GenAI-mediated learning, (ii) methodological clarity on the interplay between fine-tuning, retrieval, and adaptive generation, (iii) empirical evidence that accessibility-oriented GenAI benefits both dyslexic and non-dyslexic learners, and (iv) a replicable approach for explainable, rubric-based AI evaluation. The research thus advances the current understanding of how AI can support human learning not by replacing teachers or simplifying tasks, but by facilitating comprehension, reflection, and equitable participation across diverse learning needs

## 1.5 Research Path

The research unfolded through a sequence of interconnected phases, each contributing a building block to the development of AID-LA, a model for Reflective, Adaptive, and Inclusive Education enriched, but never replaced, by Generative AI. This trajectory is illustrated in Figure 1 - PhD Research Timeline, which provides an overview of the major projects, thematic areas, and their temporal relationships within the PhD journey.

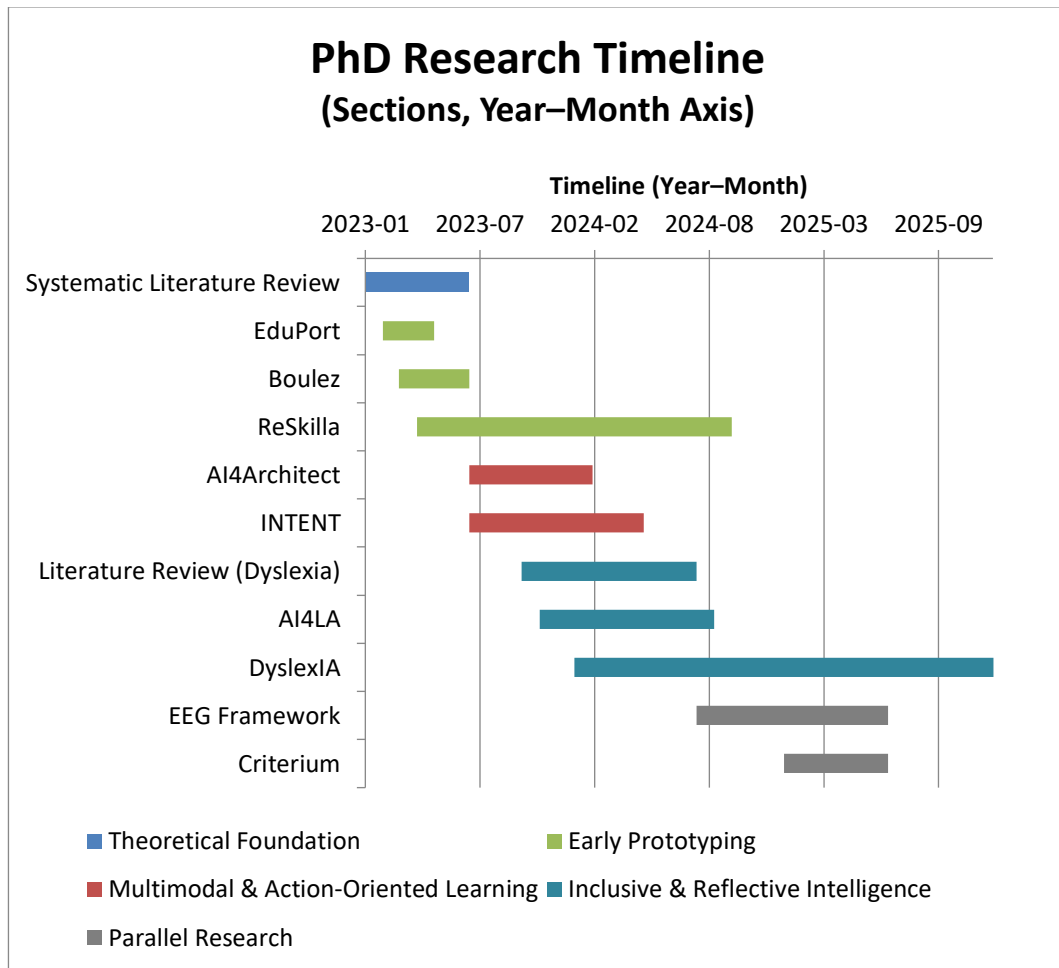


Figure 1 - PhD Research Timeline

Over the three-year period, the work evolved from theoretical grounding toward multimodal experimentation, inclusive AI design, and parallel methodological innovation. Each phase translated conceptual insights into operative prototypes, allowing theoretical assumptions to be iteratively tested, validated, and refined. The resulting framework progressively consolidated into a unified vision: GenAI as an augmentative layer that supports human teaching and

learning processes while preserving interpretability, autonomy, and pedagogical intent.

The following subsections examine these phases in detail: what each contributed, how they informed subsequent developments, and how they collectively shaped the proposed AID-LA model.

### **1.5.1 Phase I – Systematic Literature Review (SLR).**

The first phase established the theoretical foundation of the work through a systematic analysis of the state of the art in AIED, LA, and accessibility technologies. Conducted with the PRISMA (Page et al., 2021; Takkouche & Norman, 2011) methodology, the review mapped the evolution of educational AI from 2020 to 2023, identifying key research gaps in adaptivity, inclusivity, and reflection within existing intelligent-tutoring approaches. These findings provided the conceptual background for the subsequent Design-Based Research (T. Anderson & Shattuck, 2012; Reimann, 2010) (DBR) cycles.

### **1.5.2 Phase II – Early Prototyping: Adaptive and Collaborative Intelligence.**

Building on this foundation, the second phase explored the potential of chatbot (Adamopoulou & Moussiades, 2020)-based learning environments through two preliminary prototypes. EduPort (Cuneo et al., 2023), an exploratory multi-tiered conversational assistant for Massive Open Online Courses (MOOCs), tested the integration of GenAI within course materials, while Boulez (D’Urso et al., 2023) introduced a federated-learning architecture that enabled collaboration among distributed educational chatbots while preserving data privacy. Several developments of these early systems were carried out in collaboration with the Italian company ReSkilla (ReSkilla s.r.l., n.d.), which provided a proprietary e-learning platform for integrating and testing conversational assistants in real-world scenarios. This industrial partnership enabled the practical validation of GenAI integration, user interaction design, and agile development strategies within a professional learning context. Together, these activities laid the technical groundwork for adaptive and collaborative learning assistants.

### **1.5.3 Phase III – Multimodal and Action-Oriented Learning.**

The third phase expanded the framework into multimodal and procedural domains. AI4Architect(Abrusci et al., 2023) combined textual and visual generation to foster ideation and reflective design reasoning, while the Intent-based LLM Framework(D’Urso, Martini, et al., 2024) advanced system configurability using Retrieval-Augmented Generation (RAG)(D’Urso & Sciarrone, 2024), Function Calling(OpenAI, Inc., n.d.-c), and Assistant API(OpenAI, Inc., n.d.-a) mechanisms (currently deprecated). This phase marked a methodological turning point, transforming the assistants from reactive conversational agents into interactive and task-oriented collaborators. The collaboration with ReSkilla continued in this stage, supporting applied experimentation with RAG pipelines, fine-tuning methods, and the development of summarization and concept-mapping features for adaptive digital learning environments.

### **1.5.4 Phase IV – Inclusive and Reflective Intelligence.**

The final research phase concentrated on accessibility and inclusion. AI4LA applied fine-tuned LLMs to support students with Specific Learning Disorders (SLDs)(American et al., 2013; Organization, 2019), in particular dyslexic students, through adaptive text simplification and contextual feedback. These efforts culminated in DyslexIA (D’Urso & Sciarrone, 2025b), a comprehensive bilingual environment integrating readability-aware fine-tuning, multi-level summarization, and reflective scaffolding grounded in the ICAP model. Parallel research threads reinforced the framework from methodological and applied perspectives:

- an EEG-based cognitive framework(D’Urso, Luongo, et al., 2024) linking indices such as Focus, Engagement, and Flow to adaptive feedback mechanisms.
- Criterium(D’Urso & Sciarrone, 2025a), a rubric-aligned evaluation system extending the reflective paradigm to assessment through transparent, explainable scoring of open-ended responses.

### **1.5.5 Integration and Validation.**

The integration and validation activities are currently ongoing and unfold across multiple contexts that extend the experimental work conducted in earlier stages. Building on the developments already carried out in collaboration with

ReSkilla, where the prototypes were first embedded into a real operational setting for applied testing, the company is now continuing to put the results of the research into practice within its private training pathways, enabling further iterative refinement with adult learners.

A second validation pathway involves DyslexIA, which is currently being tested with secondary-school students and in two study centres, under the supervision of educational tutors. This context allows the evaluation of accessibility-enhanced generation, readability-oriented fine-tuning, and reflective scaffolding in inclusive learning environments.

In parallel, an international pilot has begun at the Tilburg School of Economics and Finance(Tilburg University, n.d.) (The Netherlands) where the integration has been tested on multiple LMS environments (Canvas(Instructure, Inc., n.d.), Moodle(Moodle Pty Ltd, n.d.), and Teachable(Teachable, Inc., n.d.)). A gradual adoption is planned within the new LMS platform that Tilburg will deploy next year.

Overall, these ongoing initiatives illustrate that integration and validation are not a single final stage but a distributed, evolving process in which the adaptive, inclusive, and reflective principles of the AID-LA framework are progressively tested, consolidated, and expanded across diverse real-world settings.

## **1.6 Significance of the Study**

The significance of this work lies in its integration of technological innovation with cognitive and pedagogical theory. Whereas most AIED research isolates technical optimization from educational design, this thesis demonstrates a coherent model connecting GenAI to evidence-based learning principles.

The proposed framework contributes to three research frontiers:

- *Theoretical*: it formalizes the AID-RLF, articulating how reflectivity, adaptivity, and inclusivity can coexist within generative-AI systems.
- *Technological*: it advances architectures for LLM-based educational assistants, integrating RAG, fine-tuning, and accessibility-aware natural-language generation.
- *Practical*: it provides validated prototypes and open methodologies applicable to real-world e-learning platforms (Canvas, Moodle, Teachable),

thereby bridging research and educational practice and contributing to the emerging vision of Learning 5.0, where artificial intelligence empowers human cognition through adaptive, inclusive, and reflective interaction.

## 1.7 Thesis Structure

The remainder of this dissertation is organized into six chapters.

- Chapter 2: **Scope**  
Defines the boundaries of the research, formulates the RQs, and presents the conceptual and methodological framework guiding the study. It also clarifies delimitations, assumptions, and the expected theoretical, technological, and educational contributions.
- Chapter 3: **Methodology**  
Describes the methodological orientation of the research, combining SLR and DBR. It details data-collection strategies, analytical procedures, ethical considerations, and methodological limitations.
- Chapter 4: **Results**  
Presents the core experimental and developmental outcomes of the thesis across the four research phases. It includes the SLR findings, the early prototypes on adaptive and collaborative intelligence, multimodal and action-oriented systems, and the inclusive, readability-aware assistants. Complementary research threads, EEG-based cognitive modelling and rubric-aligned evaluation, are also reported.
- Chapter 5: **Discussion**  
Interprets the results in relation to the RQs and theoretical frameworks. It articulates the contributions of the work across four dimensions, reflectivity, adaptivity, inclusivity, and human-in-the-loop assessment, while discussing broader pedagogical, cognitive, and ethical implications.
- Chapter 6: **Conclusions and Next Steps**  
Summarizes the thesis contributions, highlights its limitations, and outlines future directions for research and development of AI-Driven Learning Assistants in educational contexts.

## 2 Scope

### 2.1 Research Focus and Boundaries

This thesis investigates the design and validation of AID-LAs that foster reflective, adaptive, and inclusive learning in higher-education and digital-learning environments.

The study lies at the intersection of AIED, Human–Computer Interaction, and Cognitive Psychology, integrating GenAI architectures with established pedagogical frameworks (CLT, ICAP, IVH, UDL).

Three complementary dimensions delineate the research space:

- *Reflectivity*: activating metacognition and self-questioning through conversational scaffolds and Vee-based reflective prompts.
- *Adaptivity*: dynamically adjusting explanations, readability, and feedback through RAG, fine-tuning, and cognitive-state analytics.
- *Inclusivity*: enabling accessibility for diverse learners through multimodal interfaces, readability-aware design, and assistive features for dyslexic students.

The assistants are conceived as complements to teachers, augmenting human feedback rather than automating evaluation or affective mediation.

### 2.2 Research Questions (RQs)

Building on the conceptual framework introduced in Introduction, this section defines the RQs that guide the investigation.

The study explores how GenAI, particularly LLMs, can be embedded within pedagogically grounded architectures to enhance learning processes that are reflective, adaptive, and inclusive.

The overarching question is therefore:

**RQ0:** *How can AID-LAs based on LLMs enhance reflectivity, adaptivity, and inclusivity in digital education?*

From this, four specific RQs are derived:

- *RQ1: What conceptual and architectural principles define an AID-LA that integrates cognitive scaffolding with GenAI?*

- *RQ2: How can RAG, fine-tuning, and readability metrics be combined to create adaptive and context-aware educational dialogue?*
- *RQ3: In what ways can LLM-based systems be designed to ensure accessibility and engagement for learners with dyslexia and other SLDs?*
- *RQ4: To what extent can such assistants contribute to transparent, explainable, and fair feedback or evaluation processes for open-ended learning tasks?*

These RQs define the analytical dimensions of the study and frame the construction of the theoretical model that underpins the thesis. The next section introduces the AID-RLF, which integrates these dimensions into a coherent set of design principles.

### 2.3 Conceptual Framework

The conceptual basis of this research is articulated through the AID-RLF, a model that connects pedagogical theory with computational implementation. It translates the three guiding dimensions identified in the thesis, that is reflectivity, adaptivity, and inclusivity, into design principles for AID-LAs, complemented by a cross-cutting focus on feedback and evaluation. This framework performs a dual role. First, it provides a theoretical lens through which to interpret the educational value of AI systems, grounding technological innovation in cognitive and instructional theory. Second, it serves as a design blueprint for the experimental prototypes described in subsequent chapters, ensuring that each implementation adheres to a coherent pedagogical rationale.

*Table 1 - Pedagogical Foundations and AI Implementations within the AID-RLF*

<b>Dimension</b>	<b>Pedagogical / Cognitive Foundations</b>	<b>AI Implementation Strategy</b>
Reflectivity	IVH; ICAP Model	Conversational scaffolding, reflective questioning, and metacognitive prompting integrated in dialogue design
Adaptivity	CLT; LA; Self-Regulated Learning	RAG, fine-tuning, and readability calibration for dynamic content delivery

Inclusivity	UDL; Assistive Technologies for SLDs	Multimodal interfaces, text simplification, and personalized accessibility features (e.g., speech-to-text, font customization)
Feedback & Evaluation	Formative Assessment; Explainable AI; Rubric-based Evaluation	LLM-based feedback generation, rubric-aligned scoring (Criterion), and transparent explanation of model reasoning

Within this structure, reflectivity encourages learners to articulate reasoning and monitor understanding; adaptivity ensures that explanations and feedback adjust to context and learner profile; and inclusivity guarantees that interaction remains accessible to students with diverse cognitive or linguistic needs. Feedback and evaluation connect the learner-facing dimension with the teacher-facing one, promoting transparency and pedagogical coherence.

By integrating these four pillars, the AID-RLF provides the theoretical backbone of the thesis, guiding both the development and the interpretation of all subsequent prototypes.

## 2.4 Methodological Boundaries

The methodological approach adopted in this research follows a design-based and iterative paradigm, aligning technological experimentation with pedagogical inquiry. The overall process integrates qualitative and quantitative methods within successive design–implementation–evaluation cycles, reflecting the dual nature of the thesis as both theoretical and applied.

### 2.4.1 Research Design

The study employs DBR as its overarching methodological framework. DBR is well suited to educational innovation because it combines the development of practical solutions with the generation of theoretical knowledge through iterative refinement in authentic contexts.

Each research cycle comprised the following phases:

- *Design and modelling*: translating pedagogical principles into computational architectures (e.g., adaptive retrieval, reflective prompting).

- *Implementation and testing*: developing and deploying functional prototypes in controlled or semi-real learning environments.
- *Evaluation and reflection*: analysing interaction data, readability indices, and qualitative feedback to refine both the system and the underlying model.

Operationally, the DBR framework was implemented through Agile(Fowler et al., 2001; Kamat, 2012; Salza et al., 2019)-inspired development cycles and the PDCA(J. C. Anderson et al., 1994; Moen et al., 2006) (Plan–Do–Check–Act) logic of continuous improvement. Agile principles guided the short, iterative sprints of prototype development, emphasizing flexibility, user feedback, and incremental enhancement. The PDCA cycle provided the meta-structure for each iteration:

- *Plan*: defining objectives, metrics, and expected learning outcomes.
- *Do*: implementing and testing the system in a pilot setting.
- *Check*: evaluating usability, accessibility, and cognitive impact.
- *Act*: integrating findings to improve both design and theory.

Together, these approaches ensured methodological coherence between research-level iteration (DBR) and development-level iteration (Agile/PDCA), maintaining alignment between pedagogical goals, technological artefacts, and empirical validation.

#### **2.4.2 Data Collection and Evaluation**

Data were collected from pilot experiments involving higher-education and digital-learning contexts, typically through:

- *Interaction logs* (dialogue content, coherence metrics, response accuracy).
- *Quantitative indices* such as readability measures (Gulpease (Lucisano et al., 1988), Flesch(Talburt, 1986), LIX(J. Anderson, 1983)) and timing data.
- *Qualitative evidence* from user interviews, educator feedback, and observational notes.

This mixed-methods approach provided a multidimensional understanding of system performance, combining usability and accessibility metrics with cognitive and pedagogical indicators.

### **2.4.3 Scope and Limitations**

The empirical validation was conducted across a limited set of pilot environments, including a secondary school, two study centres supporting learners with tutors, a university-level design course involved in the AI4Architect evaluation, the private e-learning platform ReSkilla, and an early-stage university context at Tilburg. These settings involve small cohorts and are used primarily for feasibility-oriented testing rather than large-scale generalization. While the methodology validates the feasibility of the proposed framework, it remains exploratory in scale:

- Large-scale longitudinal or cross-institutional deployments are beyond the current scope.
- Ethical, privacy, and transparency aspects are embedded in the design of the systems but are not subjected to formal compliance audits.
- Neurophysiological investigations (e.g., the EEG-based cognitive framework discussed separately in the next Chapters) are considered complementary methodological explorations rather than part of the central evaluation protocol.

These boundaries ensure a balance between technical experimentation and pedagogical relevance, allowing the thesis to validate the educational and design principles of AID-LAs within realistic research constraints.

## **2.5 Integration of the Research Path**

The research path underlying this thesis was structured as a coherent sequence of studies that progressively advanced the conceptual and technical maturity of AID-LAs. Each phase contributed a specific layer of evidence, technical, cognitive, or pedagogical, while remaining aligned with the overarching goal of enhancing reflectivity, adaptivity, and inclusivity in digital education.

### **2.5.1 Evolution of the Research**

The research evolved through a sequence of overlapping design-based iterations in which theoretical insights, technical developments, and empirical pilots progressively informed one another. Rather than unfolding as a linear pipeline, the trajectory alternated moments of conceptual consolidation through cycles of

prototyping and field experimentation, gradually shaping the AID-RLF as a unified framework for reflective, adaptive, and inclusive AI-mediated learning.

The initial phase consisted of a SLR, which identified key gaps in adaptivity, inclusivity, and reflective scaffolding across existing AIED systems. These findings guided the development of the first exploratory prototypes, EduPort and Boulez, which investigated content-aligned support and collaborative intelligence in distributed learning scenarios. Early versions of these systems were also tested within ReSkilla's private training environment, enabling preliminary application of GenAI-based assistants in authentic learning settings.

A second group of studies extended the research into multimodal and task-oriented interactions. AI4Architect and the Intent-based LLM Framework introduced capabilities such as RAG, Function Calling, fine-tuning, and action-oriented dialogue, marking a shift from static Q&A support to procedural, co-creative, and system-configurative assistance.

The third phase deepened the focus on accessibility and reflectivity. AI4LA explored fine-tuned linguistic adaptation for learners with dyslexia, while DyslexIA integrated readability-aware generation, multilevel summarization, and ICAP-aligned scaffolding into a unified bilingual environment. In parallel, Criterium and the EEG-based cognitive framework broadened the scope of the research toward explainable assessment and cognitive-state-aware adaptivity.

Current work extends these contributions into ongoing multi-site pilots. These include applications in professional training contexts, evaluations with secondary-school learners supported by educational tutors, and a collaborative pilot in a university setting. Together, these environments provide authentic conditions for iterative refinement and ecological validation of the assistants.

Taken together, these developments do not constitute separate, isolated phases but a cumulative working path in which prototypes, empirical observations, and theoretical constructs continuously inform one another. This evolution has converged toward the AID-RLF, which synthesizes these strands into a coherent pedagogical and technical model.

### **2.5.2 Synthesis and Coherence**

While each prototype addressed specific RQs, their cumulative design reflects an integrated methodological logic:

- *From adaptivity to collaboration*: ensuring that AI assistance remains contextually aligned and pedagogically reliable.
- *From creativity to inclusivity*: expanding generative capacities toward accessible, multimodal, and human-centred learning experiences.
- *From learning support to reflective evaluation*: embedding transparency and metacognition in both learner and teacher interfaces.

This integration process demonstrates how successive experimental layers converged toward the AID-RLF introduced earlier. Each system iteration not only validated a distinct aspect of the framework but also informed the refinement of its theoretical principles, creating a reciprocal relationship between design practice and conceptual advancement.

### **2.5.3 Research Coherence and Continuity**

The final synthesis achieved in this thesis represents the consolidation of all intermediate results into a unified model. The prototypes and methodological studies form an interconnected ecosystem rather than isolated projects, collectively addressing the four RQs outlined earlier. This cumulative approach illustrates a deliberate evolution, from foundational experimentation to integrated theory, confirming the coherence of the proposed paradigm of AID-LAs for reflective, adaptive, and inclusive education.

## **2.6 Delimitations and Assumptions**

To clarify the interpretive boundaries of the research, this section outlines the principal delimitations, the choices that intentionally narrow the focus of the study, and the assumptions that underpin its methodological and theoretical orientation.

These delimitations ensure coherence with the thesis objective: to conceptualize, design, and validate AID-LAs that are reflective, adaptive, and inclusive within formal educational settings.

### 2.6.1 Delimitations

- *Scope of Interaction*: the systems developed in this research are limited to textual and multimodal dialogue. Embodied or immersive interfaces (e.g., VR, AR, or haptic environments) are excluded, as the focus remains on language-based interaction and its cognitive and pedagogical implications.
- *Educational Context*: the empirical validation took place across a limited number of pilot environments: one secondary school, two study centres where learners work with tutors, a university-level design course used for testing AI4Architect, the private training platform ReSkilla, and an initial university context at Tilburg. The scope of validation reflects the pilot nature of the research and is intentionally limited to contexts where iterative testing is feasible.
- *Evaluation Focus*: evaluation emphasizes usability, accessibility, readability, and cognitive impact. Commercial scalability, market adoption, or cost–benefit analyses are deliberately excluded.
- *Ethical and Legal Boundaries*: the research integrates principles of ethical design (transparency, bias mitigation, privacy by design) at the architectural level, but does not undertake formal compliance audits or external certification.
- *Language and Cultural Setting*: prototypes were developed and tested primarily in English and Italian, reflecting the linguistic contexts of participating institutions. Broader multilingual or cross-cultural generalization is beyond the present scope.
- *Temporal Scale*: the evaluation considers short to mid-term learning interactions and does not extend to long-term retention, transfer, or behavioural change, which are proposed directions for future

### 2.6.2 Assumptions

- *AI as Epistemic Mediator and Cognitive Amplifier*: LLMs are conceived not as autonomous teaching agents but as epistemic mediators—tools that facilitate reasoning, reflection, and conceptual understanding through dialogue. Their role is to augment human cognition, supporting sense-

making and metacognition, rather than to automate pedagogical decision-making or replace the teacher's evaluative judgment. This perspective situates AI as a co-executive cognitive partner that externalizes part of the learner's reasoning process while keeping the human agent in control of interpretation and meaning.

- *Pedagogical Validity*: it is assumed that integrating cognitive and instructional theories, such as CLT, the ICAP model, and the IVH, provides a sound pedagogical basis for AI-driven learning systems.
- *Generative Reliability*: although LLMs can produce variable outputs, the combination of RAG, fine-tuning, and readability control is assumed to ensure sufficient reliability for educational use.
- *User Agency and Oversight*: learners and educators are regarded as active agents within the human–AI interaction loop. The systems are designed to support, not replace, the teacher's evaluative and instructional role.
- *Ethical Intent and Transparency*: all prototypes assume the responsible use of data and AI, with transparency toward users about the system's artificial nature and operational logic.

### 2.6.3 Risks and Critical Reflections

While the thesis focuses on the constructive integration of AI in education, it also acknowledges emerging epistemic and behavioural risks.

- *Epistemic Risk (Epistemia)*(Loru et al., 2025): the phenomenon of Epistemia describes a collapse of epistemic diversity and critical discourse when individuals rely on algorithmically mediated information flows. In educational contexts, excessive dependence on AI-generated content may reinforce confirmation bias and reduce exposure to divergent viewpoints. The proposed framework mitigates this by embedding reflective scaffolds that prompt learners to question, compare, and verify generated information rather than accepting it passively.
- *Cognitive Laziness*(Fan et al., 2025) and *Delegated Thinking*: AI assistance can inadvertently foster intellectual passivity or “delegated cognition,” where learners defer reasoning to the system. This risk is addressed through the design of active learning prompts and metacognitive engagement loops

that encourage students to interpret, justify, and re-elaborate AI responses rather than merely consuming them.

By explicitly recognizing these risks, the thesis frames AI not only as a technological opportunity but also as a pedagogical responsibility, emphasizing the need for reflective use, human oversight, and epistemic plurality in AI-mediated education.

## **2.7 Expected Contributions**

Within the boundaries defined above, this thesis aims to make a set of scientific, technological, and educational contributions to the field of AIED.

Each contribution emerges from the integration of theoretical reflection and experimental validation across the prototypes and studies described in previous sections.

### **2.7.1 Scientific Contributions**

- *A validated theoretical framework*: the thesis introduces the AID-RLF, which systematically links cognitive and pedagogical theories, such as the ICAP model, CLT, and the IVH, with computational mechanisms for adaptivity, reflectivity, and inclusivity. This framework contributes to the theoretical foundations of AIED by formalizing how GenAI can serve as a cognitive and epistemic mediator rather than a content-delivery tool.
- *Integration of reflection, adaptation, and accessibility*: the work shows that these three educational dimensions, often addressed separately in literature, can coexist within a unified model supported by LLMs and GenAI. This integrative perspective expands current understandings of how AI can support higher-order learning and self-regulated cognition.

### **2.7.2 Technological Contributions**

- *Operational prototypes*: the research delivers a sequence of working systems, EduPort, Boulez, AI4Architect, AI4LA, DyslexIA, and Criterium, each implementing specific facets of the framework, from adaptive retrieval and multimodal creativity to accessibility and explainable evaluation.

- *Readability-aware fine-tuning and adaptive summarization*: the fine-tuning strategies and summarization pipelines developed for DyslexIA represent novel approaches for aligning LLM-generated language with cognitive accessibility requirements, particularly for learners with dyslexia.
- *Human-centred architecture design*: the integration of LLMs with RAG, Function Calling, and intent translation demonstrates a replicable architecture for AI systems that remain interpretable, teacher-supervised, and aligned with pedagogical goals.

### **2.7.3 Educational and Societal Contributions**

- *Empirical evidence of reflective and inclusive learning support*: through pilot studies in secondary schools and higher-education contexts, the thesis provides qualitative and quantitative data on how AI-driven assistants can enhance student engagement, comprehension, and self-awareness.
- *Methodological transferability*: the DBR approach, operationalized through Agile and PDCA cycles, offers a methodological template for future educational-technology projects seeking to balance innovation with theoretical grounding.
- *Guidelines for responsible AI integration in education*: the work formulates principles for designing AI systems that promote epistemic plurality, transparency, and human agency, explicitly addressing the risks of Epistemia and AI-induced cognitive laziness identified in Risks and Critical Reflections.

### **2.7.4 Overall Impact**

This thesis integrates insights from cognitive theory, AI methodologies, and iterative experimentation to outline a coherent approach to AI-driven learning support. Rather than positioning technology as a substitute for human guidance, the work emphasizes its role in complementing understanding, reflection, and accessibility within diverse learning contexts.

The AID-RLF proposed in this research offers a conceptual basis that can inform the design of AI-Driven Learning Assistants with reflective, adaptive, and

inclusive characteristics. Ongoing pilots across different educational settings will continue to refine and evaluate the framework's applicability and effectiveness.

### **3 Methodology**

The methodological design of this research evolved in parallel with its conceptual development. From the outset, the study aimed not only to observe the rise of GenAI in education but also to actively explore how such technologies could be embedded in pedagogically coherent learning environments. Consequently, the research adopted a mixed methodological stance: empirical, iterative, and design-based, where theoretical inquiry and system development informed one another in successive cycles.

#### **3.1 Methodological Orientation**

The overall approach follows the logic of DBR, a methodology particularly suited for educational innovation that unfolds through iterative design, implementation, analysis, and reflection in authentic contexts.

Unlike experimental paradigms that isolate variables, DBR allows the co-evolution of theory and artifact: every prototype becomes both an experimental tool and a unit of analysis. Within this framework, the research process was structured in multiple phases—beginning with a systematic mapping of the scientific landscape and progressing toward iterative design and validation of AID-LAs.

#### **3.2 Systematic Literature Review**

The methodological foundation of the research was established through a SLR conducted according to the PRISMA guidelines.

The objective of this review was to construct a comprehensive map of the evolution of AIED in the post-transformer era, with specific attention to chatbot-based learning systems, generative models (GPT/BERT), and accessibility-oriented AI. Searches were carried out between 2020 and 2023 across four major academic repositories—IEEE Xplore(IEEE, 2025), ACM Digital Library(Association for Computing Machinery, 2025), Springer Link(Springer Nature, 2025), and MDPI Journals—using standardized Boolean queries (learning, education, chatbot, GPT, BERT). After identification, screening, and eligibility checks, approximately 400 papers were retained for qualitative synthesis.

The PRISMA process revealed a methodological and technological turning point in late 2022 with the introduction of ChatGPT, marking a shift from domain-specific tutoring systems to general-purpose generative assistants.

The analysis exposed three persistent gaps in the literature:

- the lack of reflective scaffolding in AI-driven learning systems;
- the fragmentation between adaptive and inclusive approaches;
- the limited pedagogical grounding of generative models.

These findings provided the conceptual and empirical baseline for the experimental work that followed, informing the iterative design and validation cycles of the DBR framework.

### **3.3 Design Based Research Framework**

Building on the findings of the review, the study adopted DBR as its overarching methodological approach. DBR provided the structure for translating theoretical insights into iterative cycles of design, implementation, and reflection in authentic learning environments. Operationally, these cycles were informed by Agile development principles and the PDCA logic of continuous refinement, as outlined earlier. This dual perspective ensured methodological coherence between research-level iteration and development-level iteration, maintaining alignment between pedagogical intent, technical artefacts, and empirical validation. Within this DBR framework, the four developmental phases introduced in the Introduction were realized as successive research cycles. Each cycle addressed a different dimension of the model—adaptive and collaborative learning, multimodal and procedural interaction, inclusive and reflective intelligence, and ongoing integration and validation. Collaborations with external partners provided authentic contexts for early experimentation and cross-platform testing, supporting the ecological validity of the approach without requiring any fixed institutional or technological configuration. As is typical of DBR, the iterative evolution of the prototypes prompted successive refinements of the theoretical grounding. While the initial review offered a broad map of AIED and generative technologies, the emergence of accessibility and inclusion as central design requirements—particularly during the development of AI4LA and DyslexIA—necessitated an additional thematic review. This second review, conducted through a structured

PRISMA-inspired process and later published as a short literature review, examined the intersection of AI, dyslexia, and education. Its findings revealed a striking scarcity of generative-AI systems designed for learners with SLDs, highlighting gaps in accessibility, readability, and reflective scaffolding. These insights directly informed the inclusive-intelligence phase of the DBR cycle and guided the pedagogical and technical choices underpinning the DyslexIA environment.

### **3.4 Data Collection and Analysis**

Data collection was multi-layered and adapted to the nature of each prototype. It combined:

- *Quantitative data*: interaction logs, readability indices (Gulpease, Flesch, LIX), accuracy and coherence metrics, usage analytics.
- *Qualitative data*: educator and learner feedback, semi-structured interviews, and content analysis of reflective responses.
- *Comparative data*: cross-evaluation of LLM configurations, fine-tuning parameters, and retrieval strategies.

Analytical methods included descriptive statistics for performance indicators, thematic coding for qualitative data, and triangulation across data sources to ensure validity. Data collected across both industrial and academic settings informed micro-level analyses (usability, engagement, accessibility) and macro-level synthesis (overall coherence of the AID-RLF).

### **3.5 Ethical and Methodological Considerations**

All experimental activities adhered to principles of research integrity, informed consent, and data protection. User data were anonymized, and no personally identifiable information was stored. Special attention was devoted to inclusion and accessibility, particularly in studies involving dyslexic participants. When commercial APIs (e.g., OpenAI, Azure Cognitive Services) were employed, data handling and retention policies were verified for compliance with the General Data Protection Regulation (GDPR(*General Data Protection Regulation*, 2016)).

Throughout the research, transparency was maintained regarding the capabilities and limitations of generative models, with human oversight emphasized in all evaluative processes.

### **3.6 Methodological Limitations**

While the adopted methodology ensured flexibility and ecological validity, several limitations remain:

- Empirical validations were conducted on pilot-scale cohorts, limiting large-scale generalization.
- Dependence on proprietary LLM APIs restricted full control over model updates and internal data management.
- The neurocognitive (EEG) component was exploratory, providing methodological insight but not yet enabling real-time adaptive feedback.
- Finally, long-term educational outcomes, such as knowledge retention or transfer, fell outside the temporal scope of this research.

## 4 Results

### 4.1 Introduction

This chapter presents the outcomes of the research, following the evolutionary path of the proposed AID-LA Framework from its conceptual foundations to its empirical validation.

Each section corresponds to one of the major stages identified in Introduction, showing how successive DBR cycles progressively integrated adaptivity, multimodality, reflection, and inclusion within a coherent model of human–AI collaboration in education.

The first part reports the results of the SLR, which established the theoretical background and revealed critical gaps in adaptivity, inclusivity, and reflective scaffolding across existing Intelligent Tutoring and LA systems.

The subsequent section examines the development of adaptive and collaborative intelligence, encompassing the EduPort and Boulez prototypes. These systems laid the technological foundations for retrieval-augmented dialogue and federated cooperation among chatbots, demonstrating how GenAI could enhance content alignment and peer-like interaction in learning environments.

The following section expands into multimodal and action-oriented learning, through AI4Architect and the Intent-based LLM Architecture. Here, textual and visual generation were integrated with function-calling and RAG pipelines, transforming chatbots from reactive assistants into active co-designers and interpreters of complex tasks.

The next section addresses inclusive and reflective intelligence, encompassing AI4LA and DyslexIA, which operationalized accessibility through readability-aware fine-tuning, multi-level summarization, and reflective scaffolding grounded in cognitive and pedagogical theory.

Finally, two complementary research lines consolidated the methodological depth of the framework:

- the EEG-based Cognitive Framework, linking neural indicators such as focus and engagement to adaptive feedback

- Criterium, which extended the reflective paradigm to assessment through transparent, rubric-aligned evaluation.

Taken together, these contributions trace a coherent progression, from conceptual analysis to inclusive implementation, demonstrating the feasibility of an AID-LA that mediates between technological performance and pedagogical intent.

## 4.2 Results of the Systematic Literature Review

### 4.2.1 Purpose and Scope

The first phase of the research aimed to establish a comprehensive understanding of AIED in the transformer-driven era (2020–2023). To this end, a SLR was conducted according to the PRISMA methodology, with the objective of mapping the evolution of chatbot-based learning systems, LLMs, and accessibility-oriented educational AI published between 2020 and 2023. The review sought to identify conceptual and technological gaps (particularly concerning adaptivity, reflectivity, and inclusivity) that would later guide the design of the experimental prototypes and the consolidation of the AID-RLF.

### 4.2.2 Search Strategy and Corpus

The literature search was performed across the main scientific repositories, using standardized Boolean queries to capture the intersection between education, conversational AI, and transformer-based models.

*Table 2 - Search Strategy and Corpus*

Database	Search period	Query / Keywords	Results
IEEE Xplore	2020 – 2023	learning AND education AND (chat OR bot OR chatbot) AND (GPT OR BERT)	≈ 315 records
ACM Digital Library	2020 – 2023	same terms, E-publication date filter	≈ 381 records
Springer Link	2021 – 2023	Multimodal interfaces, text simplification, and personalized education AND learning AND chatbot AND (BERT OR GPT) features (e.g., speech-to-text, font customization)	≈ 794 records

MDPI Journals	2018 – 2023	(education, chatbot, BERT, learning) and (education, chatbot, GPT, learning)	≈ 86 records
---------------	-------------	--	--------------

After duplicate removal and relevance screening, approximately 400 papers were retained for qualitative analysis and thematic classification.

### 4.2.3 Classification and Trends

The selected corpus was organized into four thematic clusters:

- Adaptive and Intelligent Tutoring Systems (ITS / LLM-based chatbots): studies developing tutoring agents capable of personalized or retrieval-augmented dialogue.
- GenAI for Learning Feedback and Assessment: works leveraging LLMs to generate formative feedback, grade open-ended responses, or synthesize educational content.
- Accessibility and Inclusive AI: research addressing SLDs, multimodal interfaces, and adaptive readability for learners with dyslexia.
- Cognitive and Reflective Learning Support: approaches emphasizing metacognition, self-explanation, or reflective scaffolding within AI-mediated learning.

The quantitative analysis revealed several patterns:

- A sharp increase in publications after 2022, coinciding with the release of ChatGPT, marked a paradigm shift from rule-based or domain-specific AI to generative, conversational systems.
- Only a small subset (<10%) of papers explicitly addressed inclusion or accessibility, confirming a persistent research gap later targeted by AI4LA and DyslexIA.
- A growing number of studies employed retrieval-augmented or fine-tuned transformer pipelines but rarely incorporated pedagogical frameworks such as ICAP or UDL.

### 4.2.4 Interpretation

The analysis provided the conceptual grounding for the entire thesis. It highlighted three urgent needs in contemporary AIED research

- Integration of reflective scaffolding within generative dialogue systems, enabling learners to articulate reasoning and monitor understanding.
- Development of adaptive mechanisms that move beyond surface content delivery to achieve cognitive alignment and personalization.
- Embedding of inclusivity and accessibility-aware design directly into AI-driven assistants rather than treating them as external add-ons.

These insights informed the subsequent DBR phase and oriented the design of the experimental prototypes—EduPort, Boulez, and AI4Architect—which operationalized adaptive, collaborative, and reflective intelligence within the proposed framework.

### **4.3 Early Prototyping: Adaptive and Collaborative Intelligence**

#### **4.3.1 Objective and Context**

Following the insights that emerged from the literature review, this stage of the research marked the transition from conceptual exploration to experimental validation. The overarching objective was to begin testing concrete scenarios in which AID-LAs could operate, assessing both their technical feasibility and their pedagogical relevance. Within the framework of DBR, this phase corresponded to the first iterative cycle where early prototypes were conceived as exploratory tools to examine how generative models could be embedded within authentic educational contexts.

Three complementary lines of experimentation were pursued.

The first, EduPort, focused on adaptive learning support in Massive Open Online Courses (MOOCs), using ChatGPT as a conversational engine to test retrieval-augmented dialogue grounded in course materials. Its goal was to understand how generative models behave when constrained by curricular content and how their responses could be evaluated for accuracy and alignment.

The second, Boulez, addressed the problem of scalability and collaboration by proposing a federated framework for inter-institutional cooperation among educational chatbots. In this conceptual model, each chatbot represented a local learning node, coordinated through a shared orchestration layer designed to preserve privacy while enabling knowledge exchange.

Finally, the partnership with ReSkilla offered a concrete setting in which early concepts could be applied and evaluated within an operational e-learning ecosystem. This environment allowed the team to examine implementation constraints, learner experience, and organizational factors tied to the introduction of AI-driven assistance.

Taken together, these initiatives defined the Early Prototyping stage of the research: a formative phase aimed at bridging theory and practice, clarifying design principles for adaptive and collaborative intelligence, and establishing the methodological foundation for the subsequent development of multimodal and inclusive systems.

### **4.3.2 EduPort Prototype**

The EduPort prototype constituted the first practical step toward testing GenAI within educational interaction. The system was conceived to address a recurring problem in large university courses, limited teacher availability for answering individual questions, by providing a conversational interface capable of supporting students through natural-language dialogue. While inspired by the scale and autonomy typical of MOOCs, EduPort was designed as a general-purpose Q&A environment adaptable to standard university contexts.

Technically, the platform implemented a three-tiered response pipeline:

- retrieval of existing Frequently Asked Questions provided by instructors;
- automatic synthesis of answers through the BART transformer model trained on course-specific documents;
- fallback generation via ChatGPT (GPT-3.5-turbo) for open-ended queries requiring broader contextualization.

The experiment aimed primarily to understand how generative models behave when grounded in curricular content, and how their output can balance fidelity and fluency. Two RQs guided the evaluation:

- How close are the system's answers to the official learning material?
- How can generative output be constrained to preserve alignment without losing conversational quality?

The results confirmed the feasibility of integrating LLM-based dialogue within existing course infrastructures and exposed critical trade-offs. Corpus-

grounded responses ensured higher factual accuracy but tended to sound mechanical, whereas fully generative answers were engaging yet risked semantic drift beyond the syllabus. These findings emphasized the need for retrieval-augmented pipelines and transparent grounding mechanisms—principles that later informed the development of Boulez and the ReSkilla integration.

### **4.3.3 Boulez Prototype**

Where EduPort explored adaptivity at the single-course level, Boulez extended the investigation toward inter-system collaboration. The project introduced a conceptual and architectural framework for federated learning among educational chatbots, envisioning a distributed ecosystem in which institutional agents could cooperate while preserving autonomy and data privacy. The model took inspiration from the orchestration of orchestral performance, hence the name Boulez, where individual instruments contribute to a shared outcome under the coordination of a conductor. In this framework, each chatbot instance represented a local learning node, capable of answering students' queries and collecting interaction data within its own course or institutional domain. A central Orchestrator component mediated communication among these nodes, selecting and aggregating the most relevant completions provided by participating systems. The architecture incorporated auxiliary modules such as a Trust Manager, assigning reliability scores to nodes based on performance and feedback, and an Internal Storage layer maintaining metadata and interaction histories. Together, these elements constituted a prototype protocol for standardized cooperation among educational agents. Unlike traditional centralized LA, Boulez proposed a federated paradigm in which model improvements and contextual knowledge could be shared across nodes without exposing raw student data. This approach anticipated current trends in privacy-preserving AI and collaborative intelligence, offering a blueprint for future large-scale educational networks. The implementation remained primarily theoretical, supported by simulation scenarios rather than full deployment, yet it clarified the methodological requirements for interoperability, trust calibration, and orchestration in distributed learning environments.

From a research perspective, Boulez expanded the notion of adaptivity from individual learner–AI interaction to collective intelligence across learning

communities. It demonstrated how principles of federated learning could be translated into the educational domain to balance personalization with privacy, and autonomy with coordination. The insights gained here later informed the applied experiments carried out with ReSkilla, where interoperability, modularity, and data governance became essential design criteria.

#### **4.3.4 Industrial Collaboration (ReSkilla)**

The collaboration with ReSkilla, an Italian company specializing in digital learning and professional training, represented the transition from conceptual prototyping to applied experimentation. This partnership enabled the integration of AID-LAs into a production-level e-learning platform, testing their technical feasibility, pedagogical relevance, and organizational impact in real-world conditions.

The main goal of the collaboration was to examine how GenAI could enhance the learner experience within an established digital ecosystem. Specifically, the study sought to (i) implement conversational assistance for personalized learning support, (ii) evaluate RAG as a mechanism for content contextualization, and (iii) define strategies for scalable integration of AI models under GDPR and institutional constraints.

ReSkilla provided the operational context: a cloud-based platform used for corporate and academic training, offering a rich testbed of content repositories, user interaction data, and pre-existing analytics modules. The integration followed an agile and iterative methodology, structured around two main cycles.

In the first cycle, the focus was on compatibility analysis and functional prototyping. The existing ReSkilla architecture was examined to identify interfaces for chatbot integration, ensuring interoperability with the platform's learning modules and content management system. Two prototypes—web and mobile—were developed to demonstrate the chatbot's interaction capabilities, using OpenAI APIs for conversational processing and RAG pipelines for document-grounded responses. This stage also included a comparative study of proprietary versus open-source LLMs, assessing differences in performance, customization, and cost-effectiveness. The analysis highlighted the trade-off between the high contextual

accuracy of proprietary APIs and the greater transparency and controllability of open frameworks.

The second cycle concentrated on refinement and optimization. The system's retrieval engine was expanded to handle large course corpora, introducing dynamic chunking and adaptive summarization to improve response precision. A dedicated teacher mode was added, allowing educators to generate hints, summaries, and contextual feedback aligned with learning objectives. On the engineering side, extensive refactoring improved modularity, reliability, and testing coverage, while the back-end infrastructure was reorganized for scalability and fault tolerance.

The experimentation followed the principles of DBR, combining technological iterations with empirical feedback from users. Testing sessions involved both students and instructors, who interacted with the assistant through structured scenarios representing typical learning tasks. Data were collected on three dimensions:

- technical performance (latency, retrieval accuracy, and response coherence);
- pedagogical utility (perceived usefulness, clarity, and engagement);
- ethical and legal compliance (privacy handling and transparency).

Feedback indicated that learners valued the system's ability to provide contextualized and timely explanations, while instructors appreciated its potential to automate routine clarification requests. However, the tests also revealed challenges related to overgeneralization in generative responses and the need for continuous teacher supervision—reinforcing the thesis's emphasis on human-centred AI in education.

The ReSkilla collaboration validated the feasibility of embedding AI-driven assistants into professional learning platforms and demonstrated measurable gains in personalization and accessibility. It also produced a set of transferable design principles:

- retrieval and generation must remain tightly coupled to ensure pedagogical grounding.
- conversational models should be modular and explainable to maintain user trust.

- iterative co-design between educators and developers is essential for meaningful adoption.

This experience provided the first real-world validation of the proposed framework, bridging the conceptual work of EduPort and Boulez with the subsequent research on inclusivity and reflection (AI4LA and DyslexIA). Beyond its scientific results, the collaboration also served as a technological transfer initiative, laying the groundwork for the future DyslexIA s.r.l. academic spin-off dedicated to AI for inclusive learning.

#### **4.3.5 Outcomes**

The ReSkilla experimentation directly contributed to addressing RQ1 and RQ2, validating the architectural and pedagogical principles of an AI-driven reflective learning assistant. It demonstrated that adaptive mechanisms based on RAG and fine-tuning can effectively personalize dialogue while preserving contextual grounding (RQ2), and that these systems can be coherently designed within a human–AI collaborative model that supports instructional reflection rather than automation (RQ1). This stage thus marked the empirical consolidation of adaptive and collaborative intelligence within the broader framework of AI-driven education.

### **4.4 Multimodal and Action-Oriented Learning**

#### **4.4.1 Objective**

This stage of the research extended the early prototypes into multimodal and procedural domains, exploring how GenAI could support not only adaptive dialogue but also creative reasoning and operational understanding.

The objective was twofold:

- to investigate the role of fine-tuned models in mediating complex, design-oriented learning tasks.
- and to examine RAG and Function Calling as complementary strategies for dynamic, goal-directed interaction.

Two systems were developed to address these aims: AI4Architect, focused on creativity and reflection in design education, and the Intent-Based LLM

Framework, aimed at operational learning scenarios requiring precise, explainable, and executable reasoning.

#### **4.4.2 AI4Architect Prototype**

The AI4Architect system was conceived to test the potential of fine-tuned language models in supporting creative ideation and reflective dialogue in the design domain. The system integrated ChatGPT and DALL·E within a web-based interface that allowed students to post design-related queries, receive descriptive answers, and generate visual representations of concepts derived from the text.

The underlying hypothesis was that fine-tuning a generative model on domain-specific dialogues could improve contextual understanding, coherence, and stylistic adequacy of the system's responses.

To test this hypothesis, the model was fine-tuned on a corpus of question–answer pairs drawn from architecture-related learning materials and curated instructor interactions. The dataset followed a two-stage preprocessing pipeline: (i) domain adaptation through textual segmentation and prompt–completion formatting; (ii) linguistic and stylistic alignment with design terminology and reflective phrasing.

Fine-tuning was carried out through OpenAI APIs, yielding a customized model variant specialized in the architectural lexicon and pedagogical tone.

A controlled experiment was conducted with a small cohort of undergraduate design students enrolled in a studio course. Participants used AI4Architect to generate textual explanations and corresponding concept images for their projects. Feedback indicated that the assistant improved idea formulation and conceptual clarity, stimulating reflection and dialogue around design choices.

However, detailed analysis of the responses revealed critical limitations. While fine-tuning enhanced stylistic coherence and discipline-specific vocabulary, it did not improve factual grounding or retrieval accuracy. The model tended to reuse trained patterns rather than access contextual course materials, showing that fine-tuning alone is unsuitable for dynamic content retrieval or Q&A tasks requiring up-to-date knowledge.

This finding represented a turning point in the research trajectory: it clarified the distinct roles of fine-tuning and retrieval-based strategies within the AI-Driven

Learning Framework. Fine-tuning proved effective for behavioural adaptation—for example, controlling tone, empathy, or pedagogical style—but insufficient for ensuring content validity or contextual precision. These insights directly informed the subsequent development of retrieval-augmented and function-based architectures.

### 4.4.3 Intent-Based LLM Framework

Building on the limitations identified through fine-tuning alone, the Intent-Based LLM Framework was developed to investigate how different GenAI components—Prompt Engineering, RAG, Function Calling, and the Assistant API—could be combined into a coherent architecture for intent interpretation and action generation. Originally explored in the domain of network configuration, the framework enabled a systematic comparison of alternative configurations and provided a controlled environment for evaluating the strengths and weaknesses of each component.

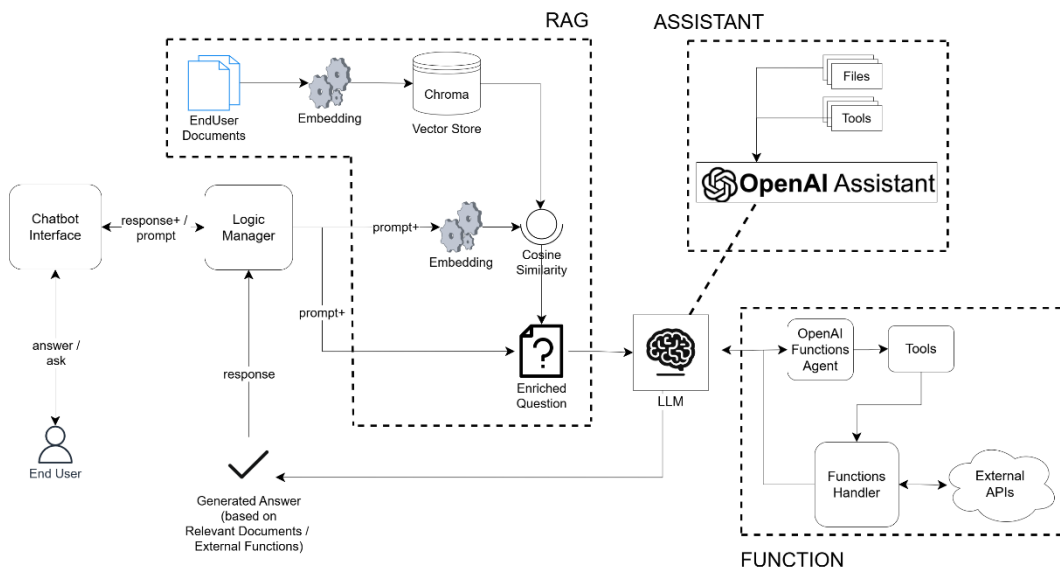


Figure 2 - Intent-Based LLM Framework Architecture

The system architecture in Figure 2 - Intent-Based LLM Framework Architecture consisted of four key layers:

- a *Prompt Engineering* layer responsible for structuring user instructions and generating enriched prompts;

- a *RAG pipeline* used to ground the model’s responses in relevant documentation through embedding, vector storage, and similarity-based retrieval;
- a *Function Calling* mechanism for translating linguistic intents into parameterized actions or executable solutions; and
- the *Assistant API* (now deprecated), which coordinated persistent conversation threads and enabled tool and file invocation within a unified interface.

Through this architecture, user inputs were first interpreted as candidate intents, enriched with retrieved contextual information, and then processed by the LLM to generate either explanatory responses or operational outputs (e.g., configuration scripts). The modularity of the system made it possible to test different combinations of components, such as RAG-only, Function Calling-only, or hybrid configurations, across multiple scenarios.

The experimental evaluation showed that the hybrid architectures integrating RAG and Function Calling achieved higher interpretability, accuracy, and robustness than fine-tuned models. RAG ensured contextual grounding, while Function Calling provided reliable action execution, reducing hallucination and increasing reproducibility. Moreover, the framework demonstrated that procedural reasoning, such as step-by-step interaction, intent clarification, and parameter verification, could be systematically supported through structured prompts and modular tool invocation.

While the application domain was technical, the findings were pedagogically relevant: the system provided a model of human–AI co-action where learners could iteratively refine intents, inspect generated reasoning, and receive immediate feedback. These insights informed the subsequent integration of retrieval-based and action-oriented capabilities within the broader AID-LA framework.

#### **4.4.4 Methodological Impact**

The research conducted in this stage consolidated the transition from reactive conversation to goal-oriented collaboration.

AI4Architect confirmed that fine-tuning enhances stylistic and behavioural alignment but lacks contextual grounding, whereas the Intent-Based LLM

Framework demonstrated that retrieval and function-based architectures are better suited for educational contexts requiring accuracy, transparency, and adaptivity.

Together, these systems established the methodological and technical foundation for the reflective and inclusive learning assistants developed in the next stage, AI4LA and DyslexIA.

## **4.5 Inclusive and Reflective Intelligence**

This final stage of the research trajectory concentrates on the integration of adaptivity, accessibility, and reflective engagement into a coherent model of inclusive AI-driven learning assistance. The previous phases established the technical foundations for building conversational systems grounded in course materials, enhanced with multimodal capabilities, and able to scaffold higher-order reasoning. Yet an essential dimension remained insufficiently addressed: the needs of learners with SLDs, and in particular dyslexia, whose challenges make them especially sensitive to linguistic complexity, cognitive load, and poorly structured learning environments.

The overarching objective of this phase was therefore to investigate how GenAI could be employed to design learning assistants that support comprehension and metacognition while remaining accessible to dyslexic learners. Achieving this required not only architectural innovation but also a clear understanding of the existing scientific landscape. For this purpose, a targeted interdisciplinary literature review was conducted to map how AI has been used to support dyslexic learners in educational contexts and to identify the research gaps that would guide the development of AI4LA and DyslexIA.

### **4.5.1 Introduction**

Dyslexia represents the most prevalent SLDs, affecting between 5% and 15% of the population and significantly influencing educational trajectories in both traditional and digital environments. Its impact extends beyond decoding difficulties: dyslexic learners often experience reduced reading fluency, slower processing speed, and challenges in sustaining attention on dense or poorly structured text. These characteristics expose them to a high risk of cognitive

overload, especially in learning settings that rely heavily on written explanations, long-form content, and autonomous study.

At the same time, the broader movement toward Learning 5.0 emphasizes human-centred, inclusive, and reflective learning, in which AI acts not as a replacement for the teacher but as a cognitive partner capable of supporting comprehension, self-regulation, and deeper engagement. For dyslexic learners, this vision demands systems that can adaptively simplify content, structure explanations, and prompt active reflection, functions that GenAI is uniquely positioned to enable when pedagogically grounded.

Before designing such systems, however, it was necessary to determine whether existing AI solutions already addressed these needs. This motivated a dedicated literature review focused explicitly on the intersection of AI, dyslexia, and education, complementing the broader AIED review presented earlier in the dissertation.

#### **4.5.2 Targeted Literature Review on AI, Dyslexia, and Inclusive Learning**

To inform the inclusive-intelligence phase of this research, a comprehensive interdisciplinary literature review was conducted to examine how Artificial Intelligence, particularly machine learning (ML) and emerging GenAI, has been applied to support students with dyslexia in educational contexts. This targeted review complemented the broader AIED mapping (Systematic Literature Review) by focusing exclusively on the intersection of AI, dyslexia, and education, as documented in the extended systematic review manuscript.

##### **Methodological Scope and PRISMA Workflow**

The review followed the PRISMA 2020 protocol and analysed literature published between 2018 and 2024, corresponding to the post-transformer era. The search query "artificial intelligence AND dyslexia AND education" was applied across six major repositories (IEEE, ACM, Springer Link, Elsevier/Scopus(Elsevier, 2025), Google Scholar(Google, 2025), ArXiv(Cornell University, 2025)).

- 582 records were initially retrieved.
- 74 duplicates were removed.

- 508 records underwent screening.
  - 412 studies met eligibility criteria, which included:
    - o focus on dyslexia,
    - o educational relevance,
    - o AI-based methods,
    - o English language,
    - o publication years 2018–2024.
  - 72 studies remained after full-text eligibility assessment (Figure 3)
- This final subset formed the analytical corpus for the review.

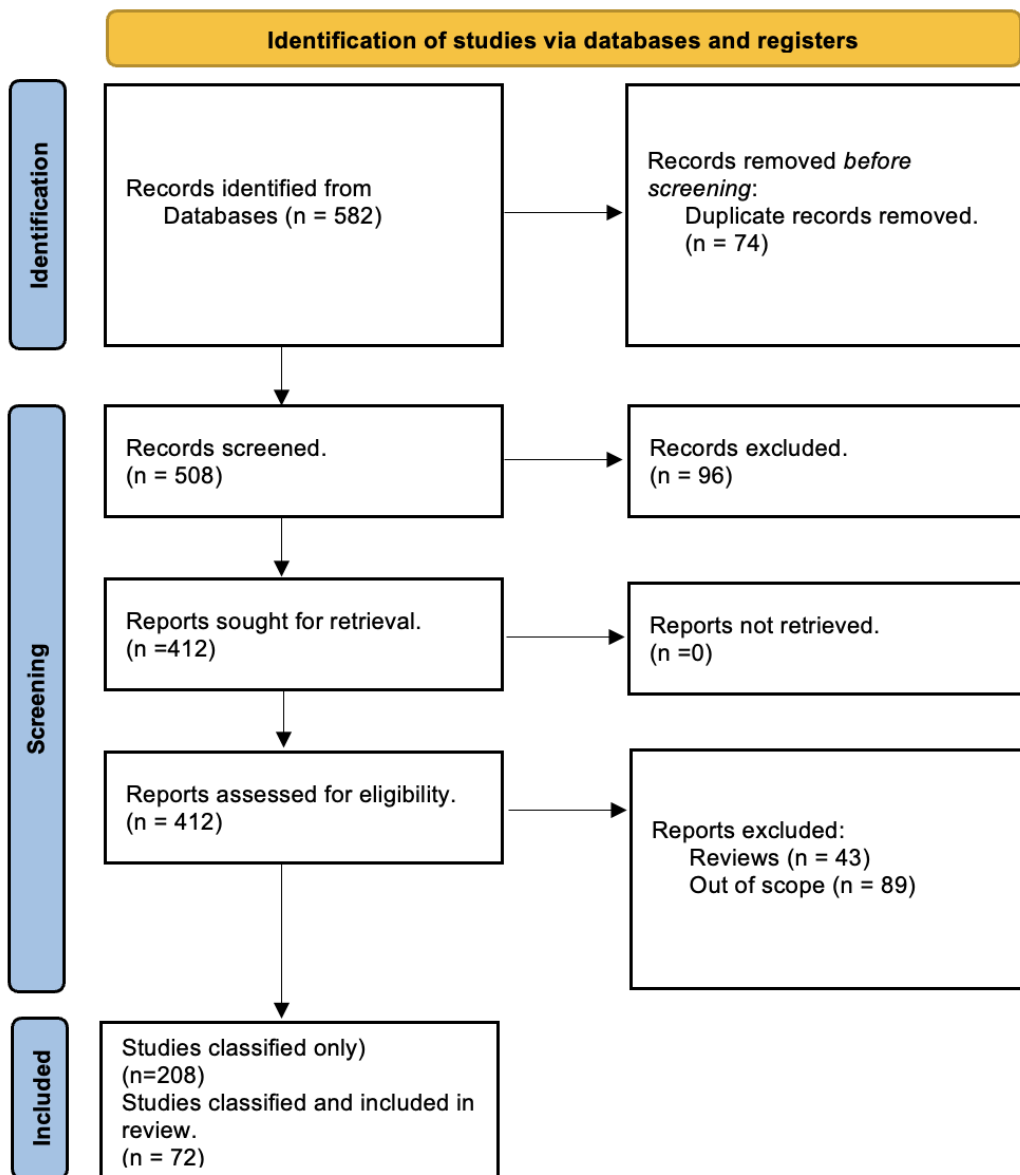


Figure 3 - The whole PRISMA workflow followed in the study

## Semantic Categorization and Interdisciplinary Structure

The 412 screened studies were categorized using a seven-category semantic framework created by the authors:

Table 3 - Categorization

Category	Description
Research	Theoretical, conceptual, or methodological discussions
Help	Assistive tools supporting learning or daily tasks
Medical	Neurological, cognitive, or physiological investigations
Review	Surveys and synthesis work
Education	Classroom experiments or pedagogical applications
Game	Gamified or serious-game learning interventions
Detection	Diagnostic or screening systems

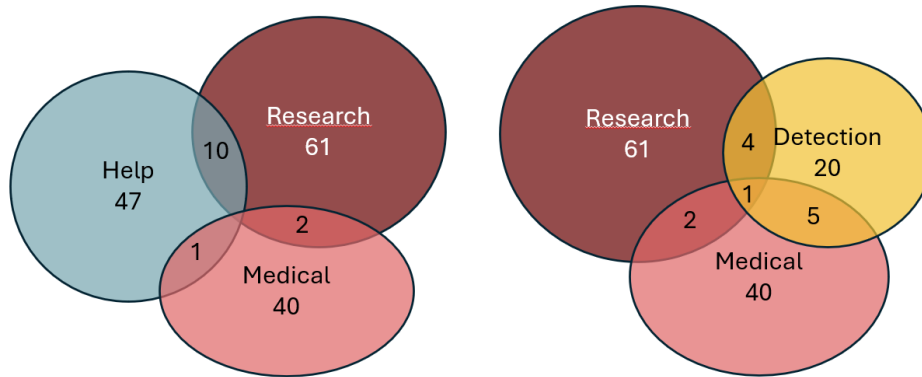


Figure 4 - The intersections among the categories: Research-Help-Medical (left) and the categories: Research-Medical-Detection (right), having respectively zero and one paper in common among the sets. The number under the name of the category is for the number of papers belonging to that category only

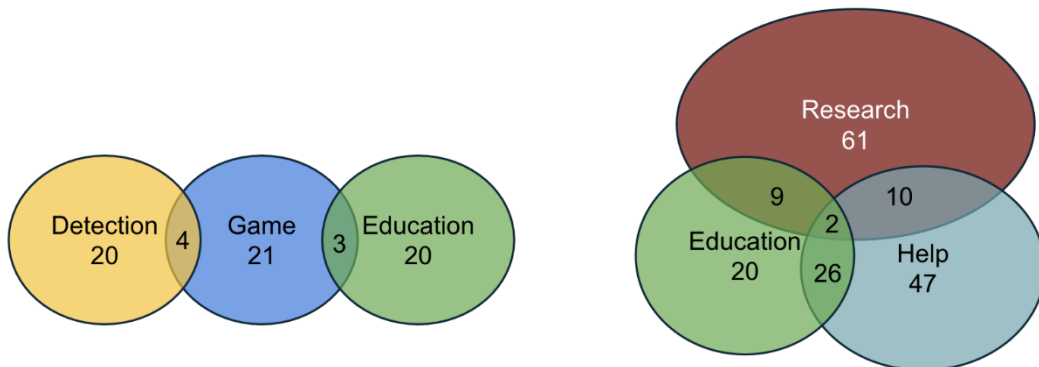


Figure 5 - The intersections among the categories: Game-Education-Detection (left) and the categories: Education-Help-Research (right), having respectively zero and two papers in common among the sets. The number under the name of the category is for the number of papers belonging to that category only

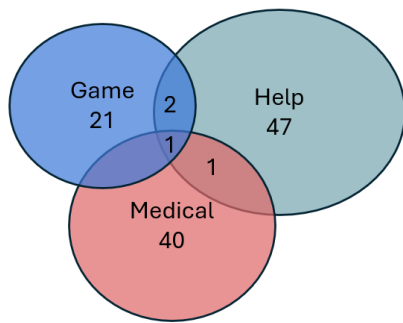


Figure 6 - The intersections among the categories: Game-Help-Medical, having one paper in common. The number under the name of the category is for the number of papers belonging to that category only

Many studies fell into multiple categories, revealing the interdisciplinary nature of work on dyslexia. Intersection analyses in Figure 4, Figure 5 and Figure 6 demonstrated that:

- Some intersections (e.g., Research–Help–Medical) contained 0 or 1 study, indicating disciplinary fragmentation.
- Others (e.g., Education–Help–Research) contained multiple studies, representing richer interdisciplinary engagement.

The most educationally relevant cluster was the Help–Education intersection.

### Quantitative Distribution Across Categories

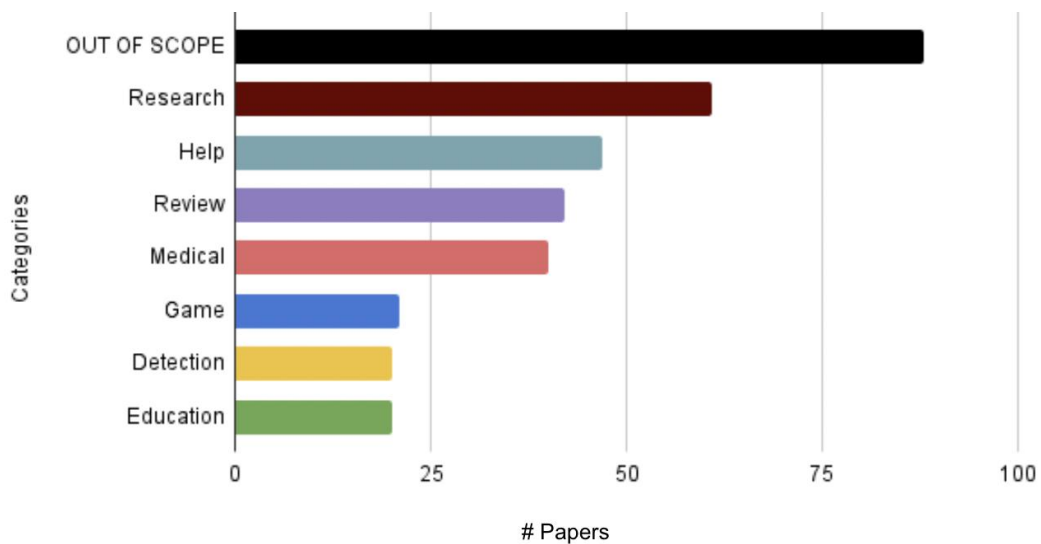


Figure 7 - The distribution of the retrieved papers belonging to one category only (n = 341)

The single-category distribution in Figure 7 showed that:

- 258 papers came from Springer,
- 221 were conference papers,

- 140 journal articles,
- and only 6 book chapters and 2 books, indicating a field dominated by emerging, not yet consolidated research.

The classification revealed:

- A large cluster of detection studies (~35% of relevant papers)
- A significant group of assistive systems (~30%)
- A smaller but growing number of educational applications, especially after 2022
- A very small number of GenAI studies, nearly all appearing only in 2024

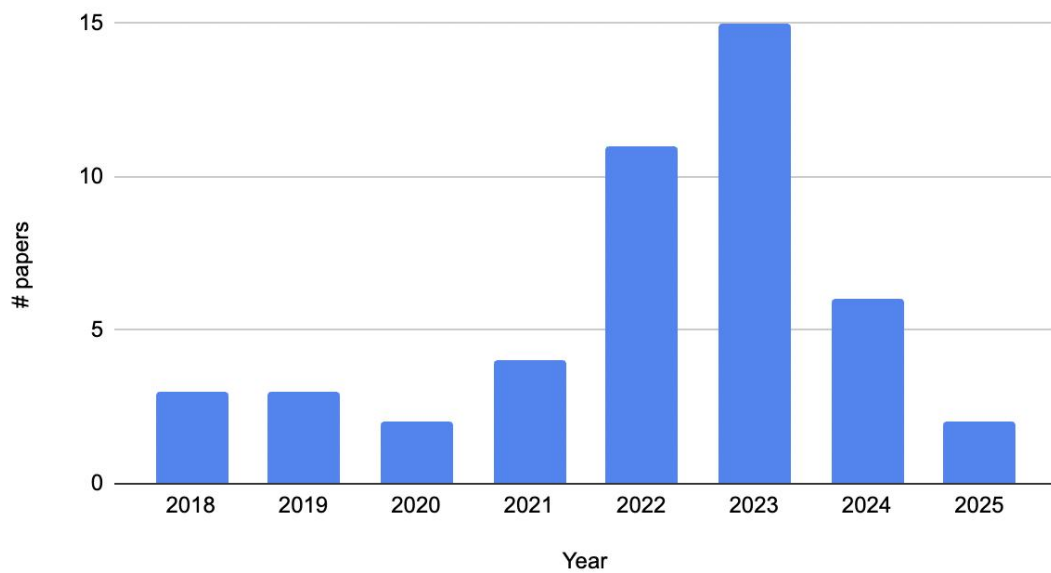


Figure 8 - The Other Papers distribution during time.

### **Analysis of the Help–Education Intersection**

The Help–Education cluster was the most directly relevant to this dissertation, as it contained systems designed to support learning rather than diagnosis.

#### **Machine Learning Studies**

These studies addressed:

- Personalization (e.g., an ML-driven student model for learners with disabilities(Benmarrakchi et al., 2018))
- Domain-specific tutoring, such as music(Della Ventura, 2019) or numeracy(Tayal, 2021)

- Engagement enhancement through ML-mediated tools(Shardlow et al., 2022)
- Adaptive practice with assistive pens(J. V. Guo et al., 2023)
- Reinforcement learning in robotics to support dyslexic learners(Mcvey et al., 2022)
- Recommendation systems to match students with appropriate assistive technologies(Morciano et al., 2024; Sharif & Elmedany, 2022; Zingoni et al., 2024)

These ML systems were technically innovative but often:

- lacked grounding in pedagogical theory,
- were not evaluated longitudinally,
- or targeted micro-tasks (letters, numbers, phonological exercises) rather than extended learning processes.

### **GenAI Studies**

A distinct subset of studies in the review focused on the recent emergence of GenAI to support dyslexic learners. Notably, all identified GenAI-based works were published in 2024, underscoring how recent this research direction is within the broader field . Despite their novelty, these contributions can be grouped into three thematic areas, reflecting different application goals and levels of pedagogical integration.

#### **Intelligent Chatbots.**

Several studies proposed conversational agents based on large language models to assist dyslexic learners through interactive dialogue, text simplification, or multimodal explanations. These systems typically leverage models such as GPT-3.5 or GPT-4 to generate adaptive responses, reading support, or contextualized explanations(De Marco et al., 2024; Li, 2024; Subramani et al., 2024). Some prototypes integrate multimodal elements, including synthesized voice or simple avatar-based interfaces, offering an immersive and more accessible interaction format. Although promising, these systems remain in an exploratory phase, with limited empirical validation in authentic learning scenarios.

### **Teacher Training.**

A second line of work examined how GenAI can assist educators in acquiring inclusive teaching strategies or understanding the instructional needs of students with dyslexia (Gilbert et al., 2023; Khazanchi & Khazanchi, 2024). These contributions focus on scenarios where teachers use LLM-based assistants to generate adapted materials, redesign assessments, or explore differentiated instruction strategies. While not learner-facing, such tools highlight how GenAI could indirectly contribute to dyslexia support by empowering educators. However, the literature suggests that teacher-oriented GenAI tools also remain preliminary and require deeper integration with established pedagogical frameworks.

### **Exploratory Studies.**

A third cluster of studies investigated the potential of GenAI in dyslexia-related contexts without proposing full learning systems. These works include conceptual analyses of how ChatGPT might support dyslexic users (Botchu et al., 2024), GenAI-enabled email-writing assistance for adults with reading difficulties (Goodman et al., 2024), preliminary investigations into GenAI's impact on learner creativity (Habib et al., 2024) [59], and broader reflections on GenAI in STEM or special education scenarios (Grover, 2024; Mayilyan, 2024). These exploratory papers do not yet provide structured interventions or validated learning tools, but they collectively signal growing academic interest in GenAI as a component of future inclusive learning environments.

Across all three clusters, GenAI-based contributions share two characteristics:

- they demonstrate substantial potential for adaptive support, linguistic simplification, and interactive learning.
- they remain technically and pedagogically immature, often lacking grounding in cognitive or instructional theory, and rarely evaluated with dyslexic learners in real educational contexts.

### **Synthesis of Trends and Research Gaps**

Across the 72 detailed studies, several patterns emerged:

- AI for dyslexia is dominated by diagnostic work, not learning support.

- Assistive tools focus on compensation, not conceptual understanding.
- Very few systems engage with metacognition, reflection, or conceptual scaffolding.
- GenAI is underrepresented, with tools appearing only recently and rarely tested.
- Interdisciplinary fragmentation (education vs. ML vs. HCI) limits integrative solutions.
- Long-term evaluations are almost entirely absent, weakening evidence of impact.
- No system integrates readability-aware GenAI, adaptive summarization, and structured reflective scaffolding in a single environment.

### **Relevance to This Dissertation**

The targeted review revealed a clear and significant research void:

- There is no GenAI-based learning assistant designed specifically for dyslexic learners that integrates readability-aware generation, chunked summarization, concept mapping, reflective questioning, contextual retrieval, bilingual interaction, and pedagogical grounding (CLT, ICAP, UDL).
- This gap directly motivated the design of the two inclusive systems developed in this thesis:
  - AI4LA – A first proof-of-concept demonstrating the feasibility of accessibility-aware generative assistance.
  - DyslexIA – A comprehensive bilingual platform integrating GenAI, RAG, readability modelling, multimodality, and reflective scaffolding.

These systems respond precisely to the deficiencies identified in the literature, establishing a novel direction in the field of inclusive AI-driven learning.

### **4.6 AI4LA: Fine-Tuned Conversational Assistant for Dyslexic Students**

The AI4LA (Artificial Intelligence for Learning Assistant) prototype marked the first complete implementation of the inclusive dimension of the research. It was conceived as a web-based intelligent chatbot designed to assist dyslexic learners

during study activities by combining fine-tuned generative dialogue, RAG, and visual knowledge representation.

#### **4.6.1 System Architecture and Functionalities**

The architecture consists of three main components (Figure 9 - The Fine-Tuned + RAG pipeline embedded in the AI4LA system.):

- *Conversational Engine*: a fine-tuned GPT-4 model optimized for short, clear, and supportive dialogue adapted to dyslexic learners' reading profiles.
- *RAG*: implemented through a ChromaDB(Chroma, n.d.) vector database for semantic retrieval of user-uploaded study materials.
- *Visualization Layer*: dynamically generates Concept Maps (CMs) based on the ongoing learner–assistant dialogue, supporting organization and retention of ideas.

Learners interact through a simple chat interface where uploaded learning materials are indexed and contextualized. The system produces simplified explanations, topic summaries, and visual maps reflecting the evolving conversation.

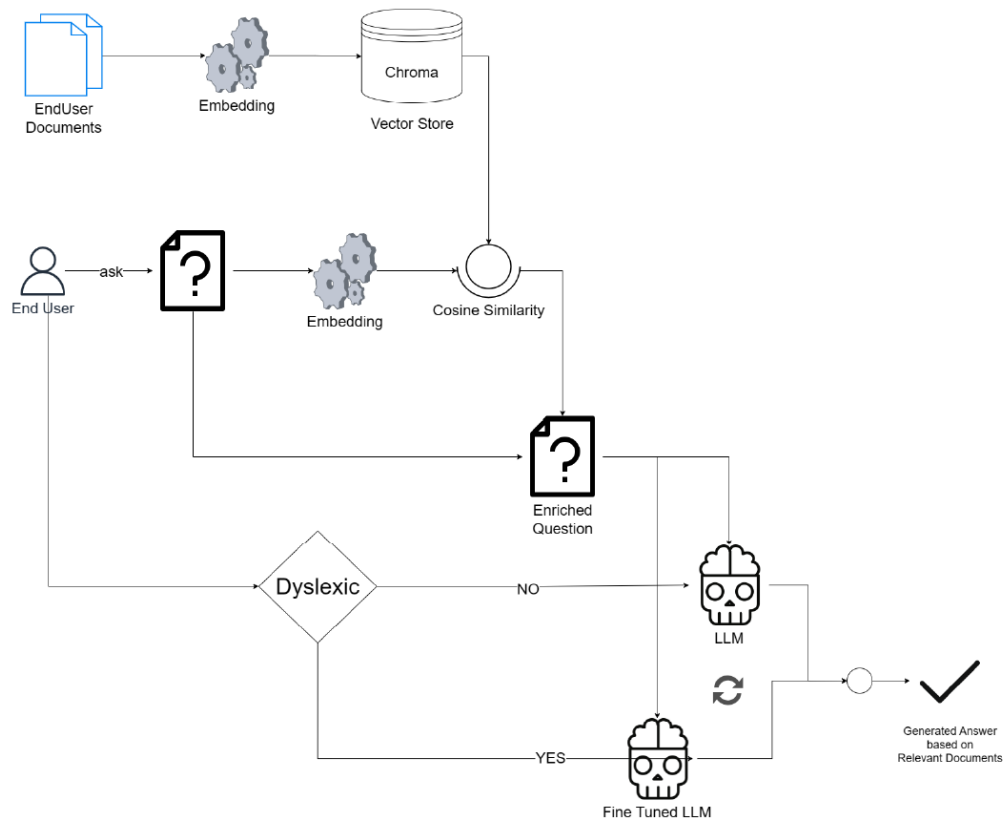


Figure 9 - The Fine-Tuned + RAG pipeline embedded in the AI4LA system.

#### 4.6.2 Fine-Tuning Process

A specialized training dataset of 80 simulated dialogues representing typical interactions with dyslexic learners was developed with the guidance of a cognitive psychologist.

Each conversation included:

- a system instruction defining role and communication style.
- multiple (user query, assistant response) pairs based on real learning tasks.

The fine-tuning focused on:

- simplifying syntax and vocabulary,
- maintaining positive and empathetic tone, and
- structuring responses around concise conceptual units.

This process produced a model capable of clear, motivational feedback and accessible paraphrasing key features for learners with dyslexia.

### 4.6.3 Evaluation and Results

A pilot study involving 23 secondary-school students diagnosed with dyslexia evaluated the system's usability, engagement, and educational impact.

Results were overwhelmingly positive:

- Mean usability score: 4.4 / 5,
- Engagement score: 4.36 / 5,
- Effectiveness score: 4.56 / 5.

Participants emphasized the assistant's ease of communication, supportive tone, and clarity of explanations. Teachers reported increased student autonomy and improved comprehension of core texts.

However, qualitative analysis revealed that as expected from the SLR findings indicating that fine-tuning generally improves stylistic alignment but not factual grounding. The model performed well in behavioural alignment, producing empathetic and accessible answers, but occasionally diverged from factual sources. This highlighted the need for RAG to ensure contextual accuracy.

### 4.6.4 Contribution

AI4LA demonstrated the pedagogical value of fine-tuned LLMs in inclusive learning environments and established a methodological distinction crucial to the rest of the research:

- Fine-tuning is best suited for behavioural and stylistic adaptation (tone, empathy, linguistic accessibility).
- Retrieval-based mechanisms such as RAG are essential for content access and factual grounding.

These insights directly inspired the next iteration of the work, DyslexIA, which integrates fine-tuning, RAG, and reflective scaffolding within a readability-aware architecture.

## 4.7 DyslexIA: Readability-Aware and Reflective Learning Environment

DyslexIA represents the most advanced and comprehensive system developed in this research, embodying the full integration of accessibility, adaptivity, and reflective engagement within a single learning environment. Building upon the insights gained from AI4LA, particularly the distinction between

behavioural fine-tuning and factual retrieval, DyslexIA was conceived as a bilingual (Italian/English) platform capable of delivering contextually grounded, readability-aware, and pedagogically aligned interaction to learners with dyslexia.

A fully functional prototype of DyslexIA, corresponding to the version described in this chapter, is available online at <https://dyslexiapp.replit.app/>. Access credentials have been provided separately to the referees and are included in Appendix D – Prototype Access Information.

Whereas AI4LA demonstrated that fine-tuned conversational behaviour could substantially improve linguistic accessibility and learner engagement, its reliance on fine-tuning alone revealed intrinsic limitations in factual consistency and contextual grounding. DyslexIA extends this work by unifying three technological and pedagogical pillars:

- A readability-optimized conversational engine informed by behavioural fine-tuning
- A retrieval-augmented architecture anchored in a persistent library of learner-provided materials.
- A dual visualisation layer enabling structured knowledge construction through concept maps and expressive synthesis through mind maps.

Together, these elements form a system designed not only to simplify language, but also to support comprehension, metacognition, and conceptual integration.

The system's design was guided by the objective of ensuring that learners with dyslexia could access study materials in a form that aligned with their cognitive profiles while still encouraging active and reflective engagement. To this end, DyslexIA incorporates multiple interaction levels, distinct learning modes, multimodal accessibility mechanisms (e.g., text-to-speech, speech-to-text, dyslexia-friendly fonts and layouts), and a pipeline capable of transforming raw documents, such as PDFs, slides, transcripts, or OCR-extracted images, into structured summaries, salient concept lists, and knowledge visualisations.

DyslexIA thereby embodies the core principles of the AID-RLF introduced in this dissertation:

- *Inclusivity*, through readability-aware generation and accessible interfaces.

- *Adaptivity*, through dynamic RAG-enabled contextualisation and interaction-level modulation.
- *Reflectivity*, through structured questioning and the integration of knowledge visualisation tools that support deeper modes of learning.

The following subsections describe the system in detail, beginning with an overview of its architectural structure and continuing with dedicated sections on the conversational engine, learning modes, context management, knowledge visualisation, and the readability and accessibility layer. Evaluation results and implications are then presented to illustrate the educational significance and practical impact of the system.

#### 4.7.1 System Overview

DyslexIA is implemented as a modular and extensible architecture in which accessibility-aware language generation, contextual retrieval, and visual knowledge construction operate as coordinated subsystems. The design follows a multilayer logic: each layer is responsible for a distinct aspect of the learning experience, yet all contribute to a unified workflow that transforms raw educational materials into accessible explanations, structured representations, and reflective prompts.

At a high level, the architecture is composed of five primary components:

- *Conversation Engine*: the central generative module responsible for producing responses. It incorporates behavioural fine-tuning specifically crafted for dyslexic learners and supports four interaction levels (Full, Standard, Smart+, Smart). Through these levels, the engine modulates verbosity, structure, conceptual density, and linguistic style, enabling the system to adapt dynamically to user needs and task demands.
- *Learning Modes*: a pedagogical layer that orchestrates how the assistant interacts with the learner. DyslexIA implements two complementary modes:
  - *Tutor Mode*, which focuses on comprehension and semantic coverage.
  - *Reflective Mode*, which promotes deeper understanding through epistemic questioning.

These modes operate independently from the interaction levels, allowing the system to blend linguistic accessibility with metacognitive scaffolding.

- *Context Management*: the subsystem responsible for handling learner-provided materials. Documents uploaded to the Library, whether textual files, PDFs, YouTube transcripts, or images processed via OCR, are automatically chunked, summarized, and embedded into a Qdrant (Qdrant Technology GmbH, 2025) vector database. During a chat session, the system retrieves the most relevant chunks and assembles contextual information for the generative model, ensuring factual grounding and alignment with the user's study content.
- *Knowledge Visualization*: a dual mechanism for externalizing knowledge.
  - The CM generator produces analytical, relation-based maps that reflect the conceptual structure of the source materials or emerging dialogue.
  - The Mind Map engine synthesizes information into expressive, associative layouts that support memory and personalization.

Together, these tools provide structural and interpretive perspectives on the content.

- *Readability and Accessibility Layer*: a set of linguistic and interface-level adaptations designed to reduce cognitive load for dyslexic learners. This includes fine-tuned generation patterns (short sentences, active voice, controlled information density), as well as multimodal aids such as text-to-speech, speech-to-text, and customizable fonts, spacing, and colour schemes.

These components interact through a two-phase operational model:

- *Offline phase* (Library ingestion): documents are transformed into structured, indexable, and pedagogically meaningful representations—summaries, chunks, salient concepts, embeddings.
- *Online phase* (Learning interaction): the assistant integrates behavioural fine-tuning, conversational control, retrieval-augmented generation, and visualization to provide accessible explanations and reflective guidance.

Through this layered and tightly integrated architecture, DyslexIA achieves a form of adaptivity that operates simultaneously at the linguistic, contextual, and cognitive levels.

#### 4.7.2 Conversation Engine

The Conversation Engine represents the generative core of DyslexIA and is responsible for producing linguistically accessible, contextually grounded, and pedagogically aligned responses throughout the interaction with the learner. It serves as the central orchestrator of the system's communicative behaviour, integrating behavioural fine-tuning, interaction-level modulation, contextual awareness, and structural constraints on language generation.

At the foundation of the engine lie two large language models with complementary roles. The primary model, GPT-4.1, is used for tasks demanding conceptual depth, long-range coherence, or reflective questioning, while GPT-4.1-mini supports lightweight exchanges and rapid clarifications. This dual-model approach allows DyslexIA to balance computational efficiency with linguistic expressiveness, enabling smooth transitions between elaborated explanations and succinct guidance. A key characteristic of the conversation engine is its interaction-level hierarchy, consisting of four progressively constrained communication profiles: Full, Standard, Smart+, and Smart:

- *Full Mode* provides detailed, highly structured answers, suitable for learners seeking comprehensive explanations or engaging with complex material.
- *Standard Mode* reduces verbosity while maintaining clarity and coverage, offering a balanced middle ground between detail and accessibility.
- *Smart Mode* and *Smart+ Mode* implement behavioural fine-tuning specifically optimized for dyslexic learners. These levels constrain sentence length, lexical complexity, and information density while enhancing coherence, explicit structuring, and pacing. Smart+ further incorporates reflective cues and implicit scaffolding patterns that are helpful for learners who require additional cognitive support.

Examples of fine-tuning conversation pairs used to train Smart and Smart+ behavioural modes are provided in Appendix B – Examples of Behavioural Fine-

Tuning Conversations, illustrating the controlled sentence length, pacing, lexical simplicity, and scaffolding patterns specifically optimised for dyslexic learners.

Across all levels, the conversation engine applies core constraints that promote linguistic accessibility, short sentences, active voice, explicit transitions, and predictable structuring. A detailed discussion of these readability principles is provided in Readability and Accessibility Layer. Internal prompting templates also regulate how the model handles lists, examples, explanations, and checks for understanding, ensuring consistency across responses. In parallel, the engine incorporates conversation-flow control mechanisms that adapt generation to the learner's behaviour. These include dynamic adjustment of verbosity, insertion of micro-summaries in longer outputs, and real-time modulation of cognitive load based on detected question complexity. When contextual material is available, the engine incorporates retrieved chunks into its generative process, ensuring that explanations remain anchored to the learner's documents rather than relying solely on latent knowledge. Finally, the engine is responsible for maintaining interaction continuity, managing state variables such as the learner's selected Library item, the current conversational goal, and the active learning mode. Through this orchestration, the engine provides a coherent experience that can shift fluidly between explanation, guidance, simplification, and context-dependent elaboration. The conversation engine acts as the linguistic and behavioural substrate upon which the entire DyslexIA experience is built. It enables the system to communicate in a manner that is supportive, predictable, and cognitively appropriate for dyslexic learners, while retaining the flexibility needed to accommodate diverse study scenarios.

### **4.7.3 Learning Modes**

While the conversation engine defines the system's linguistic behaviour, DyslexIA's learning modes determine how the assistant engages with the learner from a pedagogical standpoint. These modes structure the interaction into coherent phases of comprehension, exploration, and reflection, enabling the system to respond not only with accessible language but also with appropriate cognitive scaffolding. DyslexIA introduces two complementary modes—Tutor Mode and

Reflective Mode—each designed to address different stages of the learning process and different levels of cognitive engagement.

### **Tutor Mode**

Tutor Mode is activated when learners seek clarification, explore uploaded material, or express uncertainty about core concepts. Its primary function is to support comprehension and semantic coverage.

Within this mode, the system behaves as an adaptive tutor that:

- Identifies key concepts from the contextual materials retrieved through RAG.
- Presents simplified explanations aligned with the learner’s interaction level (Full, Standard, Smart+, Smart).
- Monitors coverage of the document’s conceptual space, detecting omissions or misconceptions.
- Generates targeted understanding-check questions, prompting the learner to paraphrase, exemplify, or relate information in their own words.
- Adapts pacing, ensuring that explanations are delivered gradually, with clear transitions and opportunities for learner input.

Tutor Mode follows a progression logic: as the learner engages with content, the system evaluates whether a sufficient proportion of salient concepts has been addressed. When the system detects that the learner has gained adequate comprehension—either through exploration or explicit confirmation—it transitions toward more reflective forms of guidance.

A key feature of Tutor Mode is its concept-coverage pipeline, which guides the system’s behaviour during comprehension-oriented interactions. When Tutor Mode is activated, DyslexIA extracts the set of salient concepts associated with the learner’s selected document or contextual segment. Each concept is assigned a salience score, reflecting its relevance, recurrence, and semantic centrality within the material. This produces an ordered list of concepts—from the most structurally important to secondary or peripheral ones.

Using this list, the system initiates a first-cycle coverage phase. It generates an accessible explanation that incorporates the highest-salience concepts, ensuring

that learners immediately encounter the core meaning of the text. Alongside this explanation, the assistant provides two types of pedagogical suggestions:

- In-depth prompts, which encourage the learner to explore nuances or implications of the concepts already covered.
- Exploration prompts, which direct attention toward salient concepts not yet addressed.

As the dialogue progresses, DyslexIA continuously monitors which concepts have been explicitly mentioned, paraphrased, or discussed by the learner. This coverage tracker is updated turn by turn. When new concepts are introduced, the system adapts subsequent explanations and suggestions to maintain a coherent progression through the conceptual space of the document.

This mechanism continues until the learner has covered a sufficient proportion of the salient concepts, either all of them or a predefined threshold (typically operationalised as a critical mass of high-salience and medium-salience items). Once comprehension is stabilised, the system prepares a transition toward a higher level of cognitive engagement: Reflective Mode. A short illustrative example of the concept-coverage pipeline and its transition toward Reflective Mode is presented in Appendix C – Tutor Mode Concept Coverage and Reflective Transition.

### **Reflective Mode**

Reflective Mode represents the highest cognitive tier in DyslexIA and is designed to promote deep understanding, metacognition, and meaning making. Its operation draws inspiration from epistemic frameworks such as the Vee heuristic and learner-engagement models resembling ICAP, but it is implemented in a manner tailored to the conversational medium. This transition aligns with the principles of the IVH, which emphasises the shift from surface understanding to conceptual integration once the learner has stabilised the core meaning of the content. In IVH terms, Reflective Mode is activated when the learner is ready to move from record and transform operations toward integrate and reflect operations. The reflective question is therefore not a generic dialogic move, but a deliberate prompt designed to consolidate the learner's conceptual structure and make the underlying relationships, assumptions, and meanings explicit.

In this mode, the system:

- **Generates reflective prompts** that encourage the learner to analyse relationships, assumptions, justifications, and implications behind core ideas.
- **Connects concepts**, guiding the learner to identify causal links, hierarchical structures, or contrasting perspectives.
- **Activates metacognitive awareness**, prompting learners to articulate why a concept is important, how it relates to prior knowledge, or how it might be applied.
- **Revisits knowledge visually**, often linking reflective dialogue with updates to the CM or prompting synthesis that later informs the Mind Map.

Reflective Mode does not replace explanation; instead, it complements it by shifting the learner from consuming information toward interacting with it. Reflective questions typically begin when the learner demonstrates a degree of stability in comprehension—for example, after completing exploration of a significant portion of the document or through demonstrated confidence in prior turns. The mode is sensitive to cognitive load: when learners appear overwhelmed or confused, the system reduces the intensity or frequency of reflective prompts, reverting temporarily to Tutor Mode scaffolding. Reflective dialogue in DyslexIA is not a linear sequence of questions, but a contingent process shaped by the learner's response. After each reflective question, the system evaluates the student's answer in terms of relevance, completeness, and conceptual alignment with the salient ideas previously explored. This lightweight assessment determines the next pedagogical move. When the response is judged sufficiently accurate or elaborated, the assistant may propose one of three reflective continuations:

- Deepening the same theme, inviting the learner to explore implications, assumptions, or broader meanings.
- Addressing a new reflective question, selected among the salient concepts that have not yet been examined at a deeper level.
- Generating a personalised mind map, synthesising the knowledge constructed during the reflective exchange and externalising the learner's conceptual structure.

If the response appears partial, vague, or misaligned, the system does not advance to a new reflective step but instead provides clarification or supportive guidance, temporarily returning to a more scaffolded form of dialogue. This adaptive behaviour ensures that reflection remains productive, proportionate to the learner's readiness, and consistent with the IVH principles of stabilising understanding before moving to higher-order integration.

This adaptive progression also corresponds to the upper tiers of the ICAP framework. Whereas Tutor Mode primarily supports Active and Constructive engagement, prompting learners to process, paraphrase, and reorganise information, Reflective Mode is designed to elicit Interactive behaviours. The reflective question requires the learner to integrate concepts, articulate relationships, and generate inferences, all of which are characteristic of ICAP's highest levels of cognitive engagement. By ensuring that Reflective Mode is triggered only after a stable conceptual foundation has been established, DyslexIA promotes a transition from constructive processing to genuinely interactive meaning-making.

Reflective Mode operationalises this shift through a controlled mechanism of *guided reflective questioning*. Once the system identifies sufficient conceptual stability, it generates *one* open-ended reflective question at a time, drawing directly on the salient concepts just explored. To avoid redundancy, the assistant keeps track of previously asked questions and explicitly excludes them from subsequent reflective prompts. The internal template driving this behaviour follows the pattern: *“Formulate only one open-ended reflection question. The question should help the student think critically, personally, and deeply about the concepts studied, without repeating any previously asked questions.”*

This approach ensures that reflection remains focused, cognitively appropriate, and incrementally progressive: each reflective question encourages the learner to re-elaborate, connect, or evaluate the concepts encountered during Tutor Mode without introducing excessive cognitive load.

### **Mode Coordination**

Both modes operate independently of the interaction levels (Full, Standard, Smart+, Smart), meaning that DyslexIA can deliver reflective guidance in

simplified language or provide detailed tutoring in more elaborate forms depending on the chosen profile. This separation of concerns—linguistic behaviour vs. pedagogical function—enables the system to accommodate diverse cognitive profiles and study strategies. In combination, Tutor Mode and Reflective Mode guide learners through the dual progression from comprehension → integration → reflection, allowing DyslexIA to function not merely as a supportive tool but as a facilitator of structured and meaningful learning.

#### **4.7.4 Context Management**

The Context Management subsystem constitutes the backbone of DyslexIA’s ability to operate on learner-provided study materials in a reliable, scalable, and pedagogically meaningful way. Its function is to transform heterogeneous inputs, text documents, slides, transcripts, and images, into structured, indexable, and retrievable resources that can be leveraged during interaction. This subsystem enables the assistant to provide grounded explanations, maintain factual accuracy, and support the semantic coherence of tutoring and reflective dialogue.

##### **Ingestion Pipeline**

At the centre of Context Management there is the **Library**, a persistent repository in which learners create items and associate them with one or more sources. Supported inputs include TXT and PDF documents, YouTube video URLs (converted into transcripts), and images processed through OCR pipelines to extract textual content. Once materials are added, DyslexIA performs a multi-stage ingestion workflow that standardizes the content and prepares it for retrieval.

The ingestion phase begins with text extraction and normalization, ensuring that the material is converted into a consistent internal representation. The system then applies semantic chunking, dividing the content into coherent segments optimized for cognitive processing and downstream retrieval. Chunking decisions are informed by text structure, semantic cohesion, and token constraints, producing units that are neither excessively fragmented nor too large for effective retrieval.

##### **Summaries and Markdown Rendering**

Each document undergoes a dual summarization process, consisting of an initial orientation summary and a more comprehensive map-reduce summary. These summaries provide learners with quick insight into the document's structure, support Tutor Mode during early exploration, and act as fallback resources when only partial content is available.

At this stage, DyslexIA also generates structured Markdown versions of the ingested material. Two variants are produced:

- A standard Markdown format that preserves the logical structure of the source.
- A Simplified Markdown format optimised for readability.

The Simplified version applies targeted accessibility strategies: bold highlighting of key concepts, predictable formatting, active-voice rewrites, reduced information density, and optimisation for Gulpease and Flesch-type indices. Learners may freely edit these Markdown outputs, reorganise sections, adjust language, or add images, supporting personalised study workflows and promoting ownership of the material. Following summarization and Markdown rendering, the system performs concept extraction, identifying salient terms and phrases that form a conceptual scaffold for both retrieval and downstream pedagogical operations.

### **Retrieval and Contextual Operation**

Processed chunks are embedded into dense vector representations and stored in Qdrant, a high-performance vector database that maintains a separate collection for each Library item. This ensures that retrieval remains tightly scoped to the selected study context, preserving relevance and interpretability. Semantic search and metadata filtering enable the system to retrieve not only similar chunks but also those suited to the learner's interaction history or study mode. During a chat session, the subsystem orchestrates RAG. When a Library item is active, relevant chunks are ranked based on semantic similarity, diversity, and contextual coverage. The system assembles a token-aware context package, including excerpts, summaries, and conceptual cues, which is passed to the conversation engine to ground its generative output. This selective retrieval mitigates hallucinations and ensures alignment with the learner's actual material.

DyslexIA distinguishes between two operational states:

- *Pure Chat Mode*, in which the assistant relies solely on behavioural fine-tuning and pretrained knowledge.
- *Contextual Mode*, in which all generation is grounded in the indexed material.

This separation ensures behavioural clarity and prevents unintended mixing of internal knowledge with user-provided content.

Through this structured workflow, the *Context Management* subsystem provides the foundation for DyslexIA’s adaptivity and contextual fidelity. It transforms disparate educational resources into coherent, retrievable knowledge structures that the conversational engine and learning modes can leverage to deliver accessible and meaningful support. The next section, Knowledge Visualization, describes how this contextual information is further externalized and enriched through visual tools.

#### **4.7.5 Knowledge Visualization**

The *Knowledge Visualization* subsystem constitutes one of the defining elements of DyslexIA, enabling the transformation of textual information and conversational exchanges into structured visual representations that support comprehension, memory, and metacognitive reflection. This subsystem is composed of two complementary components: the *CM generator* and the *Mind Map engine*, each fulfilling distinct cognitive functions while contributing to the learner’s overall understanding of the material.

##### **CM Generator**

The CM generator is responsible for producing analytical, structure-oriented representations of knowledge. It operates by synthesizing information extracted from either the learner’s contextual materials or the ongoing dialogue, thereby offering visual scaffolds that reflect the conceptual architecture of the content under study.

The process begins with deep concept extraction. Units of meaning (keywords, phrases, or domain-specific entities) are identified and assigned salience scores based on relevance, recurrence, and structural importance derived from the

ingestion pipeline. Saliency scores guide the selection of nodes to ensure that concept maps remain focused and avoid cognitive overload.

Next, the system detects semantic relationships among these concepts. DyslexIA leverages syntactic cues, embedding-space proximity, and higher-level discourse features (such as cause-effect, part-whole, or temporal links) to infer meaningful edges. These relationships are expressed as labelled links, enabling maps to convey not only conceptual proximity but also the nature of the connection.

DyslexIA supports both contextual and cumulative concept maps. Contextual maps are generated in response to individual learner queries or specific passages within the document, offering localized structural clarity. Cumulative maps aggregate concepts and relationships across multiple turns or across entire documents, building a progressively richer conceptual network that mirrors the learner's engagement trajectory.

To preserve clarity and avoid semantic drift, the map generator applies merging and harmonization strategies, ensuring that similar or identical concepts are unified, redundant edges are pruned, and conflicting relations are resolved. These mechanisms allow maps to remain coherent even as new information emerges during extended study sessions.

Learners may also personalize the visual output by adjusting colours, icons, or node shapes, as well as by adding images associated with concepts. This customization supports memory encoding and personal relevance, both of which are particularly helpful for students with dyslexia.

Concept maps play an important role in DyslexIA's pedagogical flow. In Tutor Mode, they assist the learner in maintaining a clear view of the content's structure and enable the system to detect coverage gaps. In Reflective Mode, the maps act as resources for deeper questioning, such as comparing relations, hypothesizing connections, or identifying conceptual hierarchies. A complete example of this process is presented in Appendix A – Example of Concept-Map Generation and Merging, where a short interaction on quantum-computing topics (qubit, superposition, and entanglement) demonstrates how DyslexIA integrates information extracted from multiple responses to generate a unified and coherent

concept map. The example illustrates the merging and harmonisation mechanisms described in this section.

### **Mind Map Engine**

Complementing the analytic clarity of concept maps, the Mind Map engine provides a more expressive and associative form of visualization. Whereas concept maps emphasize precise relationships and structured knowledge, mind maps highlight the learner's synthesis of ideas, encouraging flexibility and personalized meaning-making. Mind maps are typically generated after the learner has engaged with multiple explanations or completed a significant portion of a document. The engine aggregates textual summaries, concept clusters, and recurring themes, arranging them in a radial layout centred on the main topic. This layout supports intuitive, non-linear exploration and is particularly compatible with dyslexic learners' tendency to benefit from spatial and visual organization. The system extracts high-level themes and sub-themes, distilling the conceptual landscape into nodes that radiate outward. Each branch may include keywords, brief explanations, or user-added images, allowing learners to re-encounter the content in a condensed and visually appealing format. Mind maps excel at promoting recall and self-expression. They provide a visual synthesis that learners can annotate, reorganize, or expand. This aligns with the goal of fostering active involvement, as students are encouraged to interact with the representation rather than passively consume it. The mind map generator also supports DyslexIA's reflective processes. At the end of a study cycle, the assistant may invite the learner to examine connections between branches, identify missing elements, or articulate how different parts of the map relate to broader themes. These reflective prompts help close the learning loop by integrating comprehension, association, and personal interpretation.

### **Role in the Learning Workflow**

Together, concept maps and mind maps contribute to DyslexIA's ability to externalize mental models and support metacognitive engagement. They serve as visual anchors that complement textual explanations, provide navigational cues, and offer multiple entry points for reviewing and internalizing information.

By integrating both analytical and expressive visualization modes, DyslexIA ensures that learners not only understand individual concepts but also develop a coherent representation of their relationships and broader meaning. This dual-visualization strategy is particularly beneficial for learners with dyslexia, for whom structured visual support can significantly enhance comprehension, integration, and recall.

#### **4.7.6 Readability and Accessibility Layer**

The *Readability* and *Accessibility Layer* constitute the dimension of DyslexIA that most directly addresses the linguistic, perceptual, and cognitive needs of learners with dyslexia. While other subsystems handle retrieval, generative behaviour, or knowledge structuring, this layer ensures that the system’s outputs and interface remain accessible, sustainable, and aligned with the principles of inclusive design. It operates both at the level of language generation and at the level of multimodal interaction, providing a continuous scaffold for comprehension.

##### **Readability-Aware Language Generation**

A defining feature of DyslexIA is its use of behavioural fine-tuning explicitly aimed at producing text that is easier to read, process, and retain for dyslexic learners. Unlike domain-specific fine-tuning, which augments content knowledge, the fine-tuning used in DyslexIA optimizes how the system communicates. The training data—comprising dialogues designed according to linguistic and cognitive criteria—shapes the model’s behaviour in Smart and Smart+ modes, enforcing characteristics such as:

- Short, clear sentences, reducing working-memory demands.
- Active voice constructions, improving syntactic transparency.
- Controlled information density, limiting the amount of new material introduced per sentence or paragraph.
- Explicit discourse markers (“first...”, “in summary...”, “this means that...”), supporting coherence and predictability.
- Linear conceptual structure, avoiding unnecessary digressions or embedded clauses.

The outputs generated in Smart and Smart+ modes differ from those produced by generic LLMs not only in simplicity but in consistency: learners can rely on predictable linguistic structure, which reduces the cognitive effort associated with deciphering sentence forms and identifying conceptual relationships. The system also adapts these principles bilingually, ensuring that both Italian and English outputs follow the same readability strategies even though the languages differ in morphology, syntax, and typical sentence rhythm. Although Italian and English differ significantly in morphology, syntax, and typical sentence rhythm, the behavioural objective remains consistent across both languages. DyslexIA therefore applies language-specific readability criteria (e.g., Gulpease index for Italian and Flesch Reading Ease for English) within a unified set of generative constraints aimed at reducing cognitive load and increasing predictability.

### **Multimodal Accessibility Features**

Beyond textual accessibility, DyslexIA incorporates a wide range of multimodal features that support learners with diverse perceptual and cognitive profiles:

- *Text-to-Speech (TTS)*: allows learners to listen to generated content with adjustable voice speed and prosody, improving access for those who benefit from auditory reinforcement.
- *Speech-to-Text (STT)*: provides an alternative input modality for learners who find typing effortful or cognitively taxing.
- *Dyslexia-friendly fonts and layouts*: including adjustable line spacing, character spacing, font size, and colour schemes that reduce visual stress and improve legibility.
- *Colour personalization*: enabling users to configure background–foreground combinations that suit their perceptual comfort.
- *Image inclusion*: supporting scenarios in which visual elements can enhance comprehension, whether added by the learner or used in concept and mind maps.
- *Markdown structuring*: ensuring that explanations, lists, definitions, and examples follow highly predictable formatting patterns that support scanning and chunking.

These accessibility tools ensure that DyslexIA remains usable in varied educational settings, including secondary school, university study, and individual learning, while reducing the barriers that often characterize text-heavy learning environments.

### **Synergy Between Linguistic and Multimodal Support**

The effectiveness of this layer arises not from isolated features but from their integration. The linguistic simplification of Smart/Smart+ modes aligns with the structured formatting facilitated by Markdown; the reflective questions generated in Reflective Mode can be read aloud through TTS; concept and mind maps can be colour-customized to support perceptual clarity. This multimodal synergy allows DyslexIA to adapt to the learner's preferred mode of engagement while maintaining coherence across all forms of output. The Readability and Accessibility Layer thus complete the system's inclusive architecture, ensuring that the cognitive benefits of retrieval, structured tutoring, and conceptual visualisation remain accessible to learners with dyslexia. This alignment between linguistic form, interface affordances, and pedagogical design is essential to DyslexIA's objective of supporting comprehension, reducing cognitive load, and promoting sustained engagement.

#### **4.7.7 Evaluation**

The evaluation of DyslexIA focused on four dimensions aligned with the system's design goals: (1) readability and linguistic accessibility, (2) contextual fidelity in retrieval-augmented generation, (3) contribution of visualisations to comprehension and monitoring, and (4) usability and perceived support. The assessment combined system-level metrics (interaction logs, retrieval traces, readability indicators) with qualitative feedback from dyslexic learners and teachers.

The system was tested in two study centres supporting learners with dyslexia, involving approximately 25 middle- and high-school students and their educators. Additional, non-continulative sessions were carried out with high-school classes and special-education tutors. These settings provided authentic study conditions and allowed us to observe how DyslexIA was used during real learning activities.

Quantitative analyses focused primarily on readability gains and structural properties of the generated text, while qualitative observations captured learners' and teachers' perceptions of clarity, cognitive load, and the usefulness of visual and accessibility features. Together, these complementary perspectives provide a coherent account of DyslexIA's effectiveness across the four evaluation dimensions.

### **Readability and Linguistic Output**

Outputs generated in Smart and Smart+ modes were compared with baseline responses from an untuned LLM, revealing clear and measurable improvements in linguistic accessibility. Average readability values increased substantially: in Italian, the Gulpease index rose by 10–15 points on the same set of prompts, while English outputs showed parallel gains in Flesch-based readability. These improvements stem from a consistent reduction in sentence length and structural complexity, together with a marked increase in the use of active-voice constructions. Across samples, the proportion of active-voice sentences, initially around 40% in non-adapted outputs, reached 80–100% in the simplified modes.

Lexical dispersion also decreased, with responses focusing more tightly on high-frequency, unambiguous terms. These effects were stable across documents and topics, confirming that fine-tuning shaped a systematic behavioural profile rather than producing isolated stylistic changes. Learners consistently reported that the simplified and predictable structure made explanations easier to follow and less cognitively demanding over extended study periods.

### **Contextual Fidelity and RAG Behaviour**

The behaviour of the retrieval subsystem was analysed through chunk-level logs and systematic inspection of the explanations produced in Contextual Mode. Across sessions, retrieved passages consistently aligned with the learner's active document segments, and semantic drift remained minimal. Quantitatively, when the same prompts were answered with and without retrieval, the number of concepts introduced in the response decreased by over 60% in contextual mode, indicating a significant reduction in generative dispersion.

This contraction of the conceptual space reflects the intended behaviour of the RAG pipeline: the assistant prioritises salient information from the retrieved chunks rather than expanding into loosely related topics. As a result, explanations tend to be more compact, thematically coherent, and easier to trace back to the source material. Educators involved in the study confirmed that the assistant remained reliably anchored to the document, avoiding speculative or off-topic content typically observed in generic LLM behaviour.

### **Knowledge Visualization**

Interaction traces showed that both concept maps and mind maps were used strategically by learners at different stages of the study process. Concept maps were accessed most frequently during early comprehension, supporting the construction of a global view of the material, while mind maps were predominantly used at the end of sessions for synthesis and recall.

Quantitatively, the visualisations benefited from the increased linguistic focus observed in the simplified modes: as the number of concepts introduced in responses decreased (often by more than 50% compared to non-adapted outputs) the resulting maps contained fewer nodes and clearer relational paths. This reduction in visual complexity contributed to more readable and cognitively manageable structures, which learners reported as easier to interpret and modify.

The extensive customisation capabilities of the visualisation layer further strengthened this effect. Learners frequently adjusted colours, node positions, and annotations, turning the generated maps into personalised cognitive artefacts. Educators noted that these customisations provided meaningful insight into each learner's interpretation and areas of emphasis. No major usability issues emerged in the generation, editing, or navigation of visualisations.

### **Accessibility Tools**

Multimodal features were consistently used as complementary supports:

- Text-to-speech was employed for longer outputs or under fatigue.
- Speech-to-text reduced effort for students who struggled with typing.
- Font, spacing, and colour customisation were frequently adjusted to individual comfort.

These tools functioned reliably and were integrated naturally into study workflows, suggesting good alignment with the target population's needs.

### **Perceived Support and Overall Usability**

Although no formal rating scales were administered, usage logs and qualitative observations provided clear indications of strong usability and perceived support. Across the two study centres involved in the evaluation, nearly all learners engaged with DyslexIA's multimodal tools: text-to-speech was activated in a substantial proportion of long responses, and speech-to-text was regularly used by students who experienced typing fatigue. Font and spacing adjustments were applied in the majority of sessions, demonstrating active personalisation rather than passive feature availability.

Learners consistently described the explanations as clear and predictable, and several reported a noticeable reduction in cognitive effort when approaching dense or unfamiliar material. Educators likewise emphasised the system's coherence and pacing, noting that DyslexIA often acted as a bridging tool that supported communication and facilitated study mediation. The combined evidence suggests that the usability benefits were not incidental but directly linked to the accessibility-oriented design of the system.

## **4.8 Complementary Research Threads**

While DyslexIA represents the most mature instantiation of the AID-RLF, two additional research threads developed over the course of the doctoral project contributed essential methodological and theoretical depth: an EEG-based cognitive framework for exploring neuro-adaptive learning support, and Criterium, a rubric-aligned evaluation system designed for transparent and fair assessment of open-ended responses. Although distinct from DyslexIA in purpose and architecture, both lines of work reinforced the dissertation's overarching themes of adaptivity, reflectivity, and inclusivity.

### **4.8.1 EEG-Based Cognitive Framework**

This line of research explored whether lightweight EEG devices and machine learning models could extract cognitive indices, such as attention, fatigue, and

engagement, that may eventually inform adaptive behaviours in AI-driven learning systems. Two complementary investigations were conducted.

The first (D’Urso, Luongo, et al., 2024) developed a methodological pipeline for collecting EEG data during learning tasks, preprocessing signals, and inferring interpretable cognitive indicators through supervised models. The study demonstrated that even consumer-grade EEG headsets can yield meaningful patterns when combined with structured learning protocols. Importantly, it introduced a taxonomy of cognitive indices relevant to educational adaptation (e.g., sustained attention, mental effort), together with design constraints for real-time use in digital environments. The second investigation (D’Urso et al., 2025), examined whether neural signals collected during reasoning tasks could predict students’ performance and reveal latent cognitive processes. By applying attention-based models to temporal EEG features, the study showed promising associations between neural activity patterns and abstract reasoning ability, suggesting that EEG can provide complementary information beyond behavioural performance alone. Although not integrated into DyslexIA, these studies contribute three key insights to the broader research programme:

- *Feasibility*: EEG-derived cognitive indices can be extracted reliably in ecological, non-laboratory settings using accessible hardware.
- *Interpretability*: Certain EEG patterns correlate with task difficulty, engagement, and reasoning performance, offering signals that could support adaptive pacing or scaffolding.
- *Future integration*: Neurocognitive sensing may eventually form a fourth adaptive channel (alongside linguistic, contextual, and representational adaptivity) enabling AI assistants to respond to cognitive state in real time.

These findings remain exploratory but outline a pathway toward multimodal adaptivity, where physiological signals complement interaction logs and linguistic cues to refine the learning experience.

#### **4.8.2 Criterium: Rubric-Aligned Evaluation of Open-Ended Responses**

The Criterium system represents a second complementary research thread focused on the evaluation of open-text student responses; it operationalizes a

human-in-the-loop, rubric-based evaluation pipeline designed to support—rather than automate—teacher judgment.

Unlike DyslexIA, which centres on tutoring, retrieval, and reflective engagement, Criterium addresses the challenge of producing transparent, pedagogically grounded, and explainable assessments of student writing. The system embeds three architectural principles:

- *Rubric alignment*: Teacher-defined minimum and advanced criteria, model answers, and weighting schemes serve as explicit anchors for evaluation. The LLM does not improvise criteria; it interprets student responses through the lens of teacher-provided standards.
- *Explainable scoring*: The model generates structured rationales that connect elements of the student response to specific rubric criteria, highlighting evidence, omissions, and inferred misconceptions. This fosters accountability and interpretability for both teachers and learners.
- *Teacher oversight*: Criterium’s outputs are advisory. Teachers review, adjust, or reject AI-generated suggestions, ensuring that human expertise remains central while AI reduces cognitive load and ensures consistency.

Empirical studies conducted with educators demonstrated that the system:

- improved inter-rater consistency,
- reduced evaluation time,
- facilitated more uniform application of criteria, and
- promoted reflective dialogue between teachers and students about expectations and performance.

As a complementary research thread, Criterium extends the AID-RLF by showing how principles of reflectivity, transparency, and human-centred design can be translated into the domain of assessment. It also reinforces the broader position of the dissertation: AI should support human reasoning, not replace it, particularly in high-stakes or subjective evaluative contexts.

## 5 Discussion

The systems and prototypes presented in the previous chapter collectively converge toward a unified vision: **the development of AID-LAs that are not only technically advanced but also pedagogically grounded**. This chapter interprets those findings through the three foundational dimensions of the proposed framework—reflectivity, adaptivity, and inclusivity—and examines how the results contribute to ongoing debates in AIED.

In addition to addressing these pedagogical dimensions, the chapter discusses the implications of the research for assessment, ethical considerations, and the broader theoretical landscape of AI-supported learning. Each section corresponds to one of the research questions and offers an evidence-based interpretation of how the systems developed throughout the dissertation collectively advance reflective, adaptive, and inclusive educational technologies.

### 5.1 Reflectivity: From Dialogue to Metacognition

This section addresses *RQ1: What conceptual and architectural principles define an AI-driven reflective learning assistant that integrates cognitive scaffolding with GenAI?*

Across my research direction, reflectivity emerged not as an add-on capability but as a fundamental pedagogical function shaping the design of each successive system. From early explorations with AI4Architect to the fully realized reflective workflows in DyslexIA, the findings demonstrate that generative models can be configured to move beyond answer delivery toward supporting higher-order thinking processes. A key insight is that reflective learning with conversational AI is enabled by architectural decisions, not by generic LLM behaviour. Three principles were consistently validated across prototypes.

#### **Reflection Emerges Through Structured Dialogue, Not Monologic Explanation**

Traditional tutoring systems centre on providing clarifications or answers, typically following a unidirectional model of support. In contrast, the systems

developed in this dissertation leverage dialogue as a reflective space in which the learner is prompted to articulate reasoning, compare ideas, and connect concepts.

In AI4Architect, this was first observed through interactions where learners were encouraged to justify or refine design decisions. In DyslexIA, the conversational engine was explicitly shaped to sustain reflective engagement through:

- prompts that require the learner to elaborate or rephrase.
- questions that probe assumptions or relationships between concepts.
- gradual shifts from explanation to metacognitive inquiry.

This conversational scaffolding transforms the chatbot from an information source into a mediator of thought, prompting learners to externalize and refine their evolving understanding. This reflective scaffolding also mitigates the risk of passive or uncritical engagement by prompting learners to articulate and justify ideas, thereby countering tendencies toward over-reliance on AI explanations.

### **Reflective Scaffolding Requires Pedagogical Models Operationalized in System Logic**

Two pedagogical frameworks guided the design: an adaptation of the IVH and principles related to ICAP. Their integration was not conceptual but architectural:

- **Tutor Mode** addressed Active engagement by focusing on comprehension and coverage.
- **Reflective Mode** promoted Constructive and Interactive engagement by guiding learners to examine relations, purposes, and implications.

This demonstrates that actionable reflection in AI systems must be coded into system behaviour, through structured prompting templates, mode transitions, and dialogue control, not left to the emergent properties of the language model.

DyslexIA shows that LLMs can instantiate epistemic scaffolds by prompting learners to consider why information matters, how ideas connect, or what assumptions underlie a concept. Reflection becomes an intentional function of the system rather than a by-product of open conversation.

### **Externalization Tools Amplify Metacognition**

Reflection deepens when learners can perceive the structure of their own understanding. Concept maps and mind maps served as external artifacts enabling learners to inspect the relationships, hierarchies, and thematic connections implicit in their study materials or dialogue.

These artifacts made reasoning visible and manipulable, supporting reflection in three ways:

- *Structural reflection*: concept maps highlight how ideas relate and where gaps exist.
- *Integrative reflection*: mind maps synthesize information, encouraging learners to articulate global meaning.
- *Dialogic reflection*: visual representations informed subsequent conversational turns, allowing the assistant to ask questions grounded in the learner's evolving conceptual model.

This demonstrates that reflective AI assistance benefits from a hybrid design that integrates linguistic scaffolding with visual epistemic tools. The findings across the prototypes confirm that an AI-driven reflective assistant must integrate three architectural elements:

- Dialogue-based scaffolding that shifts interaction from explanation to inquiry.
- Pedagogical models operationalized in system logic, enabling structured transitions into reflective questioning.
- Knowledge externalization tools, such as concept and mind maps, that make conceptual structure explicit and manipulable.

Together, these components allow conversational AI to function as a metacognitive partner, supporting learners in the construction, evaluation, and refinement of their own understanding.

## **5.2 Adaptivity: Context, Language, and Cognitive State**

This section addresses *RQ2: How can RAG, fine-tuning, and readability metrics be combined to create adaptive, context-aware educational dialogue?*

Across the systems developed in this research, adaptivity emerged as the central technological and methodological principle enabling AI-driven assistants to adjust explanations, linguistic style, and conversational strategies according to the

learner's needs and context. The findings reveal that adaptivity in educational AI is not a single mechanism, but a multi-layered process involving factual grounding, behavioural shaping, linguistic control, and—at the frontier—neurocognitive responsiveness.

### **Contextual Adaptivity Through RAG**

The introduction of RAG was a turning point in the evolution from early prototypes to mature systems. EduPort first demonstrated the value of contextual retrieval: LLM answers became more relevant, more aligned with course content, and less prone to hallucination when anchored in teacher-provided materials.

AI4LA further exposed the distinction between:

- fine-tuning, which governs how the model communicates,
- retrieval, which governs what the model communicates.

Evaluation showed that fine-tuning alone improved linguistic accessibility but could not guarantee alignment with factual content. DyslexIA resolved this by integrating Qdrant-based RAG as a stable architectural component. This ensured that explanations reflected user-provided documents, PDFs, transcripts, OCR-extracted images, and preserved curricular coherence regardless of topic or instructional level.

Across all systems, the findings show that contextual adaptivity requires retrieval, not just model prompting or parameter tweaking.

### **Linguistic Adaptivity Through Behavioural Fine-Tuning**

From AI4LA onward, fine-tuning was used not as domain training but as behavioural conditioning, shaping tone, pacing, empathy, and syntactic form. DyslexIA refined this approach through Smart+ and Smart modes, which applied linguistic constraints grounded in cognitive psychology and readability theory.

The evaluation demonstrated that these modes consistently produced:

- shorter sentences,
- simplified structure,
- reduced lexical dispersion,
- predictable discourse markers,
- accessible formulations in both Italian and English.

This establishes that educational adaptivity must operate at the level of linguistic form, enabling learners, especially neurodivergent learners, to process, interpret, and retain knowledge without excessive cognitive load.

In contrast to traditional adaptive learning systems, which focus on task difficulty or sequencing, DyslexIA shows that language itself becomes a medium of adaptivity.

### **Representational Adaptivity Through Concept Maps and Mind Maps**

The Knowledge Visualization subsystem adds a third dimension to adaptivity: the ability to adapt representations to the learner's stage and comprehension level.

- Concept maps provide structural clarity during comprehension.
- Mind maps support synthesis and expressive reasoning.

These tools externalize conceptual relationships that the model infers from source documents and dialogue, enabling the system to adapt tutoring and reflective questioning to the learner's emergent understanding.

Representational adaptivity complements linguistic and contextual adaptivity, offering a multimodal pathway for learners to build and refine mental models.

### **Neurocognitive Adaptivity: The Exploratory EEG Framework**

The exploratory EEG-based cognitive framework introduced a fourth dimension: physiological adaptivity. Although preliminary, the EEG studies suggested that cognitive indices focus, fatigue, engagement can be inferred from brain activity using lightweight headsets.

The methodological insight is significant: AI systems can potentially adjust linguistic complexity, pacing, or reflective depth in real time based on neurocognitive signals.

While not implemented in DyslexIA, these findings point toward a future class of AID-LAs that integrate:

- contextual retrieval,
- behavioural fine-tuning,
- readability control,

- neurocognitive sensing.

This positions adaptivity as an evolving capability spanning multiple modalities.

### **5.3 Inclusivity: Technology as a Bridge for Diversity**

This section addresses *RQ3: In what ways can LLM-based systems be designed to ensure accessibility and engagement for learners with dyslexia and other SLDs?*

Inclusivity served as a guiding thread throughout the entire research trajectory, shaping design choices from the earliest prototypes to the fully developed DyslexIA platform. Across all systems, the findings consistently show that inclusivity in AI-driven learning environments does not consist merely of accommodating deficits, but rather of removing barriers and creating multimodal entry points for diverse learners. The systems developed in this work demonstrate how GenAI can act as a bridge, linking linguistic clarity, emotional safety, cognitive sustainability, and equitable participation.

#### **Accessibility Beyond Simplification**

A central insight of the research is that inclusivity requires far more than text simplification. While simplified language is essential for dyslexic learners, true inclusivity emerged when the systems provided multiple pathways for accessing content, including:

- text-to-speech for auditory reinforcement,
- speech-to-text for reducing writing effort,
- adjustable fonts, colour schemes, and spacing,
- concept maps for structural clarity,
- mind maps for expressive synthesis.

By integrating these modalities into cohesive environments, AI4LA and DyslexIA reduced the fragmentation typical of traditional assistive tools. Learners no longer needed to switch between external software for reading, writing, visualization, and comprehension; the assistant unified these functions within a single learning flow.

Teachers reported that this integration “normalised” accessibility, providing support without stigmatizing learners or altering curricular expectations.

### **Linguistic Inclusivity Through Behavioural Fine-Tuning**

AI4LA and DyslexIA demonstrated that LLMs can be fine-tuned not to increase their domain knowledge, but to align their communicative behaviour with the needs of dyslexic learners. This includes:

- predictable sentence rhythm,
- reduced syntactic complexity,
- clearer transitions between ideas,
- empathetic tone and motivational cues,
- careful pacing aligned with cognitive load management.

The Smart and Smart+ modes embedded in DyslexIA show that linguistic inclusivity can be systematically engineered through fine-tuned behaviour, ensuring that accessible communication is not incidental but structurally guaranteed.

This avoids the inconsistency seen in generic LLMs, which may fluctuate unpredictably between simple and complex formulations.

### **Inclusive Learning Through Reflective Engagement**

Inclusivity extends beyond allowing learners to “access” content; it also encompasses enabling them to participate meaningfully in knowledge building.

Reflective Mode in DyslexIA operationalized this principle by prompting learners to:

- articulate reasoning,
- explore relationships between ideas,
- connect current learning with prior knowledge,
- consider assumptions or implications,
- visually inspect conceptual structures.

These reflective behaviours, often underdeveloped in dyslexic learners due to cognitive load barriers, were facilitated through structured questioning and concept/mind map visualizations.

By distributing cognitive effort across modalities—linguistic, visual, conversational—the system supported deeper and more confident engagement with complex material.

### **Inclusivity as Universal Design**

Although the systems were designed with dyslexic students in mind, evaluations revealed that many features like readable language, clear structure, visual reasoning tools, multimodal interfaces provided benefits to all learners.

This reflects the principle of UDL: accessibility is not a special intervention but a design strategy that improves the learning environment for everyone. Clear explanations, structured maps, reflective prompts, and multimodal input/output support broader student populations—whether multilingual learners, students with attention differences, or learners navigating cognitively demanding content.

The systems thus demonstrate that inclusivity and academic excellence are mutually reinforcing, not competing objectives.

Across the prototypes, the findings show that LLM-based systems can ensure accessibility and engagement for learners with dyslexia when they integrate:

- Readability-aware fine-tuning that ensures consistent linguistic accessibility.
- Multimodal interfaces (TTS, STT, customizable visuals) that lower perceptual and cognitive barriers.
- Knowledge visualization tools that make structure and meaning explicit.
- Reflective scaffolds that empower learners to participate in higher-order thinking.

Taken together, these elements transform inclusion from a compensatory mechanism into a universal pedagogical strategy, enabling AI-driven systems to bridge cognitive diversity and support equitable learning.

### **5.4 Assessment and Fairness: The Human-in-the-Loop Approach**

This section addresses *RQ4: To what extent can AI-driven assistants contribute to transparent, explainable, and fair evaluation of open-ended learning tasks?*

While the primary focus of this dissertation is on reflective, adaptive, and inclusive learning, the complementary research on Criterium extends these principles into the domain of assessment. Evaluation is a critical component of the learning cycle, yet it is also one of the most labour-intensive and cognitively demanding tasks for educators. The findings from the Criterium system demonstrate that GenAI, when designed within a human-in-the-loop paradigm, can meaningfully augment the evaluation of open-ended responses without compromising fairness, interpretability, or teacher authority.

### **Rubric-Aligned Scoring as a Foundation for Fairness**

Criterium confirmed that alignment with explicit teacher rubrics is essential for trustworthy AI-assisted assessment. Instead of substituting human judgment, the system uses:

- teacher-defined minimum and advanced criteria,
- model-answer exemplars,
- structured scoring formulas, and
- text similarity retrieval when relevant

to guide the evaluation process.

The LLM acts not as an autonomous grader, but as a scoring assistant that applies rubric principles consistently. This reduces variability among graders, limits subjective bias, and enhances transparency by explicitly referencing the pedagogical standards on which judgments are based.

### **The Human-in-the-Loop Model Preserves Teacher Agency**

A crucial finding is that educators remained the epistemic authority in all evaluation workflows. The system supports teachers by:

- producing draft evaluations aligned with the rubric,
- generating concise justifications for each criterion,
- highlighting matched concepts and missing elements,
- offering suggested scores rather than definitive ones.

Teachers validated, revised, or rejected these suggestions, ensuring that AI served as an analytical aid, not a replacement. This approach directly responds to

ethical concerns about automation in assessment: grading is never outsourced to the model but enhanced and streamlined through AI-assisted analysis.

### **Explainability Through Structured Rationales**

Criterion provided evidence that LLMs can produce interpretable scoring rationales by mapping student responses to rubric elements and articulating reasoning in clear, structured prose.

These rationales supported fairness in three ways:

- *Transparency*: students and teachers can see why a response received a particular score.
- *Equity*: decisions are grounded in explicit pedagogical standards, not opaque model behaviour.
- *Consistency*: rubric-aligned reasoning reduces intra- and inter-grader variability.

Explainability is essential in educational contexts, and the results show that well-designed prompting and retrieval can make LLM scoring rationales both intelligible and pedagogically meaningful.

### **Complementarity With Reflective and Inclusive Learning**

Beyond grading, the Criterion system reinforces themes from earlier parts of the research:

- Its reflective scoring prompts mirror DyslexIA's reflective scaffolding, encouraging learners to understand why certain criteria matter.
- Its focus on fairness and transparency parallels the inclusivity focus of AI4LA and DyslexIA, ensuring that assessment does not introduce barriers for neurodiverse learners.
- Its human-in-the-loop workflow aligns with the broader thesis position: AI should augment cognition and teacher judgment, not supplant them.

Thus, while Criterion addresses a different stage of the learning process, it aligns with the overall AID-RLF.

The findings show that AI-driven evaluation systems can contribute to transparency, explainability, and fairness when designed with:

- Explicit rubric alignment, ensuring that evaluation criteria remain pedagogically grounded.
- Explainable scoring rationales, enabling accountability and interpretability.
- A human-in-the-loop workflow, ensuring teacher oversight and preventing automation bias.
- Support for reflective assessment, reinforcing learning rather than merely producing scores.

In this way, AI-supported assessment becomes an extension of inclusive and reflective learning rather than a parallel or competing process.

## **5.5 Ethical and Theoretical Implications**

The systems and findings presented in this dissertation carry significant ethical and theoretical implications for the future of AI in education. Although the technical evaluations in Chapter 4 examined performance, accessibility, and user experience, the broader meaning of these results becomes apparent only when interpreting them in light of long-standing debates regarding agency, transparency, autonomy, and the role of technology in learning. This section synthesizes these broader implications without introducing further empirical claims, positioning the contributions of this thesis within contemporary discussions in AIED.

### **Ethics Beyond Compliance: Toward Human-Centred AI**

Throughout the research, ethical considerations were embedded not merely as procedural requirements, such as GDPR compliance in data processing or transparency in model behaviour, but as design principles. The systems developed here were intentionally crafted to augment, rather than replace, human judgment and cognitive processes. This human-centred orientation manifests in several ways:

- Teacher authority is preserved, particularly in assessment scenarios such as Criterium, where educators retain the final decision-making role.
- Learner agency is supported, as reflective scaffolding encourages students to articulate reasoning and make conceptual connections autonomously.
- Transparency is maintained, especially through RAG, which grounds explanations in identifiable source materials rather than opaque model inference.

These design decisions collectively counter the tendency of some AI systems to obscure reasoning processes or promote passive consumption. Instead, the assistants developed in this thesis foreground interpretability, justification, and learner control. At the same time, the findings foreground specific epistemic risks that design must explicitly guard against. One is the danger of epistemic drift—sometimes referred to as Epistemia—where learners gradually detach their beliefs from evidence due to over-reliance on fluent but weakly grounded explanations. Another is cognitive laziness, in which students outsource interpretive or sense-making processes to the system. The architectural principles adopted in this thesis—retrieval grounding, reflective prompts, and visual externalization—were intentionally designed to counter these tendencies, ensuring that AI augments rather than replaces cognitive effort.

### **Theoretical Shift: From Automation to Augmentation**

The research contributes to a growing theoretical shift in AIED: reimagining AI not as an automated instructor, but as a collaborator in the learning process. Traditional tutoring systems often aim to approximate expert teaching behaviour, emphasizing task completion and error correction. In contrast, the systems developed here demonstrate that GenAI can function as a cognitive partner, supporting:

- externalization of thought (via conversation and visualization),
- reflective reasoning (via Vee-inspired scaffolds),
- linguistic simplification for cognitive access,
- multimodal comprehension and expression.

This positions AI assistants not as replacements for instruction, but as epistemic mediators that help learners construct, negotiate, and refine understanding. In theoretical terms, this contributes to the conceptualization of AI as a facilitator of distributed cognition, where the locus of learning is shared between human and artificial agents.

### **Rethinking Adaptivity: From Item-Level Personalization to Cognitive Ecology**

The systems in this thesis illustrate a transition from conventional adaptivity (e.g., adjusting task difficulty) toward a richer model grounded in what might be

termed cognitive ecology, the idea that learning emerges from interactions among linguistic, visual, contextual, and affective processes.

Three contributions stand out:

- Linguistic adaptivity through fine-tuned language generation tailored to dyslexic learners.
- Contextual adaptivity via RAG-based grounding in authentic study materials.
- Neurocognitive adaptivity, explored in the EEG framework, which hints at future systems that adjust dynamically to learner cognitive states.

This multi-layered adaptivity suggests that future learning systems may function not as static tutors but as dynamic, context-sensitive ecosystems.

### **Inclusivity as a Universal Principle of System Design**

Results from AI4LA and DyslexIA show that accessibility-oriented features, originally designed for dyslexic students, proved beneficial to a broader population. This reinforces the universal design perspective: systems built for inclusivity tend to enhance learning for everyone, not only those with specific needs. The implication is theoretical as well as ethical: accessibility is not a peripheral concern but a core design principle that reshapes how AI systems structure explanations, representation formats, and interactions. Inclusivity becomes a lens through which the entire learning ecosystem is constructed, rather than a secondary accommodation.

### **Implications for Assessment and Accountability**

The work on Criterium demonstrates how transparent and human-aligned scoring pipelines can redefine fairness in AI-assisted assessment. Theoretical contributions include:

- A model of explainable evaluation, grounded in teacher rubrics.
- A rejection of black-box scoring, ensuring that automation never eclipses human reasoning.
- A reaffirmation of the educator's central role, even in technologically advanced settings.

This supports a conception of AI-enabled assessment as a collaborative process, where automation assists but does not supplant human expertise.

Taken together, the ethical and theoretical implications of this work point toward a model of educational AI grounded in collaborative intelligence, where systems support human reasoning, respect agency, and promote equitable participation. The assistants developed in this dissertation, reflective, adaptive, and inclusive, demonstrate how AI technologies can act not as authorities but as partners in the co-construction of knowledge.

## **5.6 Limitations and Future Directions**

Although the systems and studies presented in this dissertation provide substantial evidence for the potential of reflective, adaptive, and inclusive AID-LAs, several limitations remain. These limitations do not diminish the contributions of the research; rather, they outline the boundaries within which the results should be interpreted and identify promising directions for future work.

### **Scope and Scale of Empirical Validation**

The empirical studies conducted in this research, whether evaluating AI4LA, DyslexIA, or the exploratory EEG-based framework, were performed primarily in small-scale and context-specific settings.

While these studies yielded meaningful insights, including high levels of usability and accessibility, they were limited in:

- sample size,
- geographic and institutional diversity,
- duration of exposure to the systems,
- breadth of curricular domains represented.

Future research should involve longitudinal, multi-site evaluations to assess learning gains, retention, metacognitive development, and behavioural changes over time. Such studies would contribute to the generalizability and external validity of the framework.

### **Dependence on Proprietary LLMs and APIs**

The systems developed in this dissertation rely on proprietary large language models and cloud-based APIs. While these tools enabled rapid prototyping and state-of-the-art performance, they introduce several constraints:

- limited transparency into model internals,
- potential cost barriers for large-scale deployment,
- challenges for reproducibility and open research practices,
- exposure to model updates outside the system designer's control.

Future work may explore open-source or locally deployable models, or hybrid architectures in which sensitive components (e.g., summarization, chunking, readability enforcement) are handled on-device while generative operations are optionally offloaded to cloud services.

### **Limitations of Behavioural Fine-Tuning**

While behavioural fine-tuning successfully enhanced linguistic accessibility and empathy, it remains inherently sensitive to:

- the quality and representativeness of the training data,
- conversational drift in longer interactions,
- the risk of over-simplification when dealing with advanced or abstract concepts.

As LLM technology evolves, future research may explore fine-tuning strategies that incorporate linguistic metrics, reinforcement learning from dyslexic users, or adaptive prompt tuning, enabling finer-grained control over readability and discourse structures.

### **Early Stage of Neuro-Adaptive Integration**

The EEG-based cognitive framework represented an exploratory proof-of-concept showing the feasibility of using physiological indicators such as attention, fatigue, or engagement, to inform adaptive feedback. However, this line of work remains in its infancy. Limitations include:

- the variability of consumer-grade EEG accuracy,
- the small scale of the experiments,
- limited integration between EEG signals and generative modelling,
- the need for robust ethical protocols for neuro-data collection.

Future directions include integrating multimodal signals (e.g., EEG, eye tracking, interaction logs), developing adaptive policies that respond to cognitive states in real time, and conducting controlled studies to assess the pedagogical impact of neuro-adaptive feedback.

### **Integration into Institutional Learning Ecosystems**

The systems developed here were tested primarily as standalone prototypes. Their integration into institutional LMS environments such as Canvas, Moodle, or Teachable, is an open challenge. Such integration raises practical questions about:

- interoperability,
- data governance and privacy,
- teacher-facing analytics,
- scalability in real-world classroom contexts,
- institutional trust and adoption.

Future research should explore deployment at scale, including secure interfaces for rubric imports (Criterion), course-level context ingestion (DyslexIA), and automated activity logging for LA.

### **Toward Comprehensive Theoretical Consolidation**

While the framework proposed in this dissertation unifies reflectivity, adaptivity, and inclusivity, further theoretical consolidation may explore:

- the role of cognitive load theory in generative dialogue design,
- the alignment between concept maps and meta-representational competence,
- the interplay between linguistic simplicity and conceptual precision,
- the positioning of reflective scaffolding within epistemic cognition models.

Such work may lead to a broader theoretical model for AI-mediated learning grounded in empirical findings and cross-disciplinary perspectives.

### **Future Research Agenda**

Based on the limitations identified above, future research will evolve along three main trajectories:

- *Scalability and Systemic Adoption*: embedding reflective and inclusive AI-driven assistants directly into LMS platforms, enabling end-to-end workflows that integrate tutoring, context ingestion, assessment, and analytics.
- *Longitudinal and Cross-Linguistic Studies*: evaluating DyslexIA and related systems across languages, cultural contexts, and educational levels to assess universality, sustainability, and long-term impact.
- *Neuro-Adaptive Learning Systems*: investigating how cognitive indices derived from EEG or other physiological sensors can inform real-time adaptation in generative dialogue, potentially opening paths toward multimodal personalization.

Collectively, these limitations and directions outline a road map for future research that deepens and extends the AID-RLF presented in this dissertation. They highlight the potential for scalable, ethically grounded, and empirically validated systems capable of transforming reflective, adaptive, and inclusive education through human–AI collaboration.

## 6 Conclusions and Next Steps

This dissertation examined how Artificial Intelligence, and in particular Large Language Models, can be shaped and regulated to support reflective, adaptive, and inclusive learning. Across the different prototypes developed during the research—spanning retrieval-based systems, multimodal assistants, fine-tuned dialogue models, reflective scaffolding tools, and rubric-aligned evaluators—a coherent vision gradually emerged. AI-driven learning assistants are most effective not when they attempt to automate teaching, but when they act as mediators: systems that accompany learners, help them make sense of complexity, encourage metacognition, and reduce accessibility barriers.

The work conducted demonstrates that generative AI can move beyond the narrow role of question-answering and become a partner in meaning-making. This shift required aligning technical architectures with cognitive and pedagogical principles, integrating RAG pipelines with readability-aware generation, modelling interaction around the ICAP spectrum, structuring reflection through the IVH, and ensuring that evaluation remains transparent and explainable. In doing so, the dissertation proposed the AID-RLF, which brings together reflectivity, adaptivity, and inclusivity as mutually reinforcing dimensions rather than as isolated features.

The development of systems such as EduPort, Boulez, AI4Architect, AI4LA, DyslexIA, Criterium and the EEG-based cognitive framework illustrates how this vision can be translated into concrete implementations. Each prototype explores a different facet of learning support—from contextual grounding to multimodal creativity, from behavioural fine-tuning to cognitive-state modelling—and each one contributed evidence that the integration of AI and pedagogy can be deliberate, incremental, and theoretically grounded. The international validation pathways, including the collaboration with Tilburg University and the deployment in schools and study centres, show that these systems are not only technically feasible but also meaningful in real educational settings.

Despite these advances, several limitations remain. The empirical studies were necessarily circumscribed, involving small cohorts and short interaction periods. Long-term impacts on learning, reflection, and motivation still need to be investigated, as do issues of cultural and linguistic generalisation. Reliance on

proprietary models poses questions of sustainability and transparency, while neurocognitive components are still at an exploratory stage and require careful ethical consideration. At the same time, the increasing pervasiveness of AI in education demands attention to risks such as epistemic shortcuts, superficial engagement, and the possible erosion of intellectual effort. These concerns underscore the importance of designing systems that stimulate reasoning rather than replacing it.

The perspectives that open from this work point toward several avenues of development. One concerns the consolidation of AI-driven assistants within institutional ecosystems, supported by teacher dashboards, monitoring tools, and integration with LMS infrastructures. Another lies in longitudinal studies capable of capturing how reflective and accessibility-oriented AI affects learners over time, particularly in terms of transfer, self-regulation, and metacognitive awareness. A further direction involves strengthening neuro-adaptive components, if they remain ethically robust and pedagogically justified. Finally, the emergence of AI-mediated learning requires clearer guidelines and governance frameworks, so that automation enhances—rather than diminishes—human agency.

This research demonstrates that the educational value of AI does not stem from automation, but from mediation—from systems that help learners interpret complexity, articulate their reasoning, and engage more deeply with knowledge. The prototypes developed throughout the dissertation illustrate how such an approach can be operationalised in practice, without compromising depth, agency, or the centrality of human dialogue. Future work should continue to refine this trajectory, designing AI that strengthens the intellectual practices underlying meaningful learning rather than replacing them.

## Bibliography

- Abrusci, L., Dabaghi, K., D'Urso, S., & Sciarrone, F. (2023). AI4Architect: An Intelligent Help System to Support Students in the Design Domain. *Methodologies and Intelligent Systems for Technology Enhanced Learning, 13th International Conference*, 65–72.
- Adamopoulou, E., & Moussiades, L. (2020). An overview of chatbot technology. *IFIP International Conference on Artificial Intelligence Applications and Innovations*, 373–383.
- Almansor, E. H., & Hussain, F. K. (2019). Survey on intelligent chatbots: State-of-the-art and future research directions. *Conference on Complex, Intelligent, and Software Intensive Systems*, 534–543.
- American, Psychiatric, Association, DSM-5, Task, & Force. (2013). *Diagnostic and statistical manual of mental disorders: DSM-5<sup>TM</sup> (5th ed.)*. American Psychiatric Publishing Inc.
- Anderson, J. (1983). Lix and rix: Variations on a little-known readability index. *Journal of Reading*, 26(6), 490–496.
- Anderson, J. C., Rungtusanatham, M., & Schroeder, R. G. (1994). A theory of quality management underlying the Deming management method. *Academy of Management Review*, 19(3), 472–509.
- Anderson, T., & Shattuck, J. (2012). Design-based research: A decade of progress in education research? *Educational Researcher*, 41(1), 16–25.
- Anthropic PBC. (n.d.). *Claude* [Computer software]. <https://claude.ai/>
- Association for Computing Machinery. (2025). *ACM Digital Library*. <https://dl.acm.org/>

- Benmarrakchi, F., Ouherrou, N., Elhammoumi, O., & El Kafi, J. (2018). An innovative approach to involve students with learning disabilities in intelligent learning systems. *International Conference on Advanced Intelligent Systems for Sustainable Development*, 39–50.
- Blackboard Inc. (2025). *Blackboard Learn*. <https://www.blackboard.com>
- Botchu, B., Iyengar, K. P., & Botchu, R. (2024). Can ChatGPT empower people with dyslexia? *Disability and Rehabilitation: Assistive Technology*, 19(5), 2131–2132.
- Chi, M. T., & Wylie, R. (2014). The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational Psychologist*, 49(4), 219–243.
- Chroma. (n.d.). *ChromaDB*. <https://www.trychroma.com/>
- Coffman, S., & Draper, C. (2022). Universal design for learning in higher education: A concept analysis. *Teaching and Learning in Nursing*, 17(1), 36–41.
- Cornell University. (2025). *arXiv e-Print Archive*. <https://arxiv.org/>
- Coursera Inc. (2025). *Coursera*. <https://www.coursera.org>
- Cuneo, A., D’Urso, S., Sciarrone, F., & Temperini, M. (2023). EduPort: An Intelligent Learning System Based on Deep Learning-A Preliminary Evaluation. *International Symposium on Emerging Technologies for Education*, 79–93.
- De Marco, V., Sciarrone, F., & Temperini, M. (2024). TutorChat: A Chatbot for the Support to Dyslexic Learner’s activity through Generative AI. *Proceedings*

*of the 24th IEEE International Conference on Advanced Learning Technologies, ICALT 2024, (in press).*

Della Ventura, M. (2019). Exploring the impact of artificial intelligence in music education to enhance the Dyslexic student's skills. *Learning Technology for Education Challenges: 8th International Workshop, LTEC 2019, Zamora, Spain, July 15–18, 2019, Proceedings 8*, 14–22.

D'Urso, S., Luongo, R., & Sciarrone, F. (2024). Enhancing educational outcomes through eeg-based cognitive indices and supervised machine learning: A methodological framework. *2024 28th International Conference Information Visualisation (IV)*, 1–6.

D'Urso, S., Martini, B., & Sciarrone, F. (2024). A Novel LLM Architecture for Intelligent System Configuration. *2024 28th International Conference Information Visualisation (IV)*, 326–331.

D'Urso, S., Pan, Z., Cristea, A. I., & Sciarrone, F. (2025). From Test Scores to Neural Spikes: Predicting Students' Abstract Reasoning Ability Using EEG with Attention-Based Models. *International Conference on Artificial Intelligence in Education*, 352–367.

D'Urso, S., & Sciarrone, F. (2024). AI4LA: An Intelligent Chatbot for Supporting Students with Dyslexia, Based on Generative AI. *International Conference on Intelligent Tutoring Systems*, 369–377.

D'Urso, S., & Sciarrone, F. (2025a). Criterium: Assisting Teachers in Fair and Consistent Grading of Open-Ended Questions. *International Workshop on Smart Education and Training Environments (SETE 2025), in Conjunction with ICWL*.

- D'Urso, S., & Sciarrone, F. (2025b). DyslexIA: Empowering Dyslexic Students with Accessible Learning through AI. *International Conference on Web-Based Learning (ICWL 2025)*.
- D'Urso, S., Sciarrone, F., & Temperini, M. (2023). Boulez: A Chatbot-Based Federated Learning System for Distance Learning. *2023 27th International Conference Information Visualisation (IV)*, 210–215.
- edX Inc. (2025). *edX*. <https://www.edx.org>
- Elsevier. (2025). *Scopus Document Search*. <https://www.scopus.com/>
- European Commission. (n.d.). *Industry 5.0*. [https://research-and-innovation.ec.europa.eu/research-area/industrial-research-and-innovation/industry-50\\_en](https://research-and-innovation.ec.europa.eu/research-area/industrial-research-and-innovation/industry-50_en)
- Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, *56*(2), 489–530.
- Fornauf, B. S., & Erickson, J. D. (2020). Toward an inclusive pedagogy through universal design for learning in higher education: A review of the literature. *Journal of Postsecondary Education and Disability*, *33*(2), 183–199.
- Fowler, M., Highsmith, J., & others. (2001). The agile manifesto. *Software Development*, *9*(8), 28–35.
- FutureLearn Ltd. (2025). *FutureLearn*. <https://www.futurelearn.com>
- Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, H., & Wang, H. (2023). Retrieval-augmented generation for large language models: A survey. *arXiv Preprint arXiv:2312.10997*, *2*(1).

- General Data Protection Regulation*. (2016, April 27). European Parliament and Council of the European Union. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- Gilbert, B., Stubblefield, J., Qualls, J., Huang, X., Pait, A., Yanowitz, K., Hays, A., Richmond, E., Parker, L., & Washington, T. (2023). Dyslexia and AI: the use of artificial intelligence to identify and create font to improve reading ability of individuals with Dyslexia. *Society for Information Technology & Teacher Education International Conference*, 856–865.
- Goodman, S. M., Buehler, E., Clary, P., Coenen, A., Donsbach, A., Horne, T. N., Lahav, M., MacDonald, R., Michaels, R. B., Narayanan, A., Pushkarna, M., Riley, J., Santana, A., Shi, L., Sweeney, R., Weaver, P., Yuan, A., & Morris, M. R. (2024). LaMPost: AI Writing Assistance for Adults with Dyslexia Using Large Language Models. *Commun. ACM*.
- Google. (2025). *Google Scholar*. <https://scholar.google.com/>
- Google LLC. (n.d.). *Gemini* [Computer software]. <https://deepmind.google/>
- Google LLC. (2025). *Google Classroom*. <https://classroom.google.com>
- Grover, D. (2024). Next-Generation Education: The Impact of Generative AI on Learning. *Journal of Informatics Education and Research*, 4(2).
- Guo, J. V., Francis, P. A., & Tan, G. (2023). D-Pen—Personalised Practice with Assistive Pen for Young Dyslexic Learners Using Machine Learning. *IRC-SET 2022: Proceedings of the 8th IRC Conference on Science, Engineering and Technology, August 2022, Singapore*, 209–243.
- Guo, Y., Qiu, W., Leroy, G., Wang, S., & Cohen, T. (2024). Retrieval augmentation of large language models for lay language generation. *Journal of*

*Biomedical Informatics*, 149, 104580.  
<https://doi.org/10.1016/j.jbi.2023.104580>

Habib, S., Vogel, T., Anli, X., & Thorne, E. (2024). How does generative artificial intelligence impact student creativity? *Journal of Creativity*, 34(1), 100072.

IEEE. (2025). *IEEE Xplore Digital Library*. <https://ieeexplore.ieee.org/>

Instructure, Inc. (n.d.). *Canvas LMS – Learning Management System*.  
<https://www.instructure.com/canvas>

Kajabi LLC. (2025). *Kajabi*. <https://www.kajabi.com>

Kamat, V. (2012). Agile manifesto in higher education. *2012 IEEE Fourth International Conference on Technology for Education*, 231–232.

Kara, A. (2021). COVID-19 PANDEMIC AND POSSIBLE TRENDS INTO THE FUTURE OF HIGHER EDUCATION: A REVIEW. *Journal of Education and Educational Development*, 8(1).

Khazanchi, R., & Khazanchi, P. (2024). Generative AI to Improve Special Education Teacher Preparation for Inclusive Classrooms. In M. Searson, E. Langran, & J. Trumble (Eds), *Exploring New Horizons: Generative Artificial Intelligence and Teacher Education* (pp. 159–177). AACE.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., & others. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474.

Li, N. (2024). GENERATIVEREADER: EXPLORING THE USE OF GENERATIVE AI IN ENHANCING READING EXPERIENCES FOR CHILDREN WITH DYSLEXIA. *INTED2024 Proceedings*, 8023–8027.

- Loru, E., Nudo, J., Di Marco, N., Santirocchi, A., Atzeni, R., Cinelli, M., Cestari, V., Rossi-Arnaud, C., & Quattrocioni, W. (2025). The simulation of judgment in LLMs. *Proceedings of the National Academy of Sciences*, *122*(42), e2518443122.
- Lucisano, P., Piemontese, M. E., & others. (1988). Gulpease: Una formula per la predizione della leggibilità di testi in lingua italiana. *Scuola e Città*, 110–124.
- Mayilyan, A. (2024). INTEGRATING GENERATIVE AI IN STEM EDUCATION: A PATHWAY TO INCLUSIVE AND ENHANCED LEARNING. *Education in the 21st Century*, *11*(1), 93–101.
- Mcvey, S.-M., Chew, E., & Carroll, F. (2022). Educational Robotics and Dyslexia: Investigating How Reinforcement Learning in Robotics Can Be Used to Help Support Students with Dyslexia. *International Conference on Technological Ecosystems for Enhancing Multiculturality*, 43–49.
- Meta Platforms, Inc. (n.d.). *LLama* [Computer software]. <https://www.llama.com/>
- Microsoft Corporation. (2025). *Microsoft Teams*. <https://www.microsoft.com/en/microsoft-teams>
- Moen, R., Norman, C., & others. (2006). *Evolution of the PDCA cycle*.
- Moodle Pty Ltd. (n.d.). *Moodle – Open-source Learning Platform*. <https://moodle.org/>
- Morciano, G., Llergo, J. M. A., Zingoni, A., Bolívar, E. Y., Taborri, J., & Calabrò, G. (2024). Use of recommendation models to provide support to dyslexic students. *Expert Systems with Applications*, *249*, 123738.

- Novak, J. D. (2010). *Learning, creating, and using knowledge: Concept maps as facilitative tools in schools and corporations*. Routledge.
- Novak, J. D., & Gowin, D. B. (1984). *Learning how to learn*. Cambridge University Press.
- OpenAI, Inc. (n.d.-a). *Assistant API*.  
<https://platform.openai.com/docs/assistants/tools>
- OpenAI, Inc. (n.d.-b). *ChatGPT* [Computer software]. <https://chat.openai.com/>
- OpenAI, Inc. (n.d.-c). *Function Calling*.  
<https://platform.openai.com/docs/guides/function-calling>
- Organization, W. H. (2019). *International Classification of Diseases, 11th Revision (ICD-11)*. World Health Organization. <https://icd.who.int/>
- Paas, F., Renkl, A., & Sweller, J. (2003). Cognitive load theory and instructional design: Recent developments. *Educational Psychologist*, 38(1), 1–4.
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., & others. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *Bmj*, 372.
- Peng, H., Ma, S., & Spector, J. M. (2019). Personalized adaptive learning: An emerging pedagogical approach enabled by a smart learning environment. *Smart Learning Environments*, 6(1), 1–14.
- Qdrant Technology GmbH. (2025). *Qdrant Vector Database*. <https://qdrant.tech>
- Reimann, P. (2010). Design-based research. In *Methodological choice and design: Scholarship, policy and practice in social and educational research* (pp. 37–50). Springer.

- ReSkilla s.r.l. (n.d.). *ReSkilla – Digital Learning Solutions*.  
<https://www.reskilla.com/>
- Salza, P., Musmarra, P., & Ferrucci, F. (2019). Agile methodologies in education: A review. *Agile and Lean Concepts for Teaching and Learning*, 25–45.
- Shardlow, M., Sellar, S., & Rousell, D. (2022). Collaborative augmentation and simplification of text (CoAST): Pedagogical applications of natural language processing in digital learning environments. *Learning Environments Research*, 1–23.
- Sharif, M. S., & Elmedany, W. (2022). A Proposed Machine Learning Based Approach to Support Students with Learning Difficulties in The Post-Pandemic Norm. *2022 IEEE Global Engineering Education Conference (EDUCON)*, 1988–1993.
- Springer Nature. (2025). *SpringerLink*. <https://link.springer.com/>
- Subramani, S., Jerish S, A., C, S., M, S., T, M., & M, M. (2024). Learn Buddy – Transforming Education with Generative AI. *2024 International Conference on Computing and Data Science (ICCDs)*, 1–5.  
<https://doi.org/10.1109/ICCDs60734.2024.10560378>
- Takkouche, B., & Norman, G. (2011). PRISMA statement. *Epidemiology*, 22(1), 128.
- Talburt, J. (1986). The Flesch index: An easily programmable readability analysis algorithm. *Proceedings of the 4th Annual International Conference on Systems Documentation*, 114–122.

- Tayal, A. (2021). An Interactive Alphabet and Number learning system using OpenCV and CNN. *2021 Asian Conference on Innovation in Technology (ASIANCON)*, 1–4.
- Teachable, Inc. (n.d.). *Teachable – Online Course Creation Platform*.  
<https://teachable.com/>
- Tilburg University. (n.d.). *Tilburg School of Economics and Finance*.  
<https://www.tilburguniversity.edu/>
- Udemy Inc. (2025). *Udemy*. <https://www.udemy.com>
- Van Merriënboer, J. J., & Sweller, J. (2005). Cognitive load theory and complex learning: Recent developments and future directions. *Educational Psychology Review*, *17*(2), 147–177.
- Vaswani, A., Shazeer, N. M., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is All you Need. *Neural Information Processing Systems*.  
<https://api.semanticscholar.org/CorpusID:13756489>
- Zingoni, A., Taborri, J., & Calabrò, G. (2024). A machine learning-based classification model to support university students with dyslexia with personalized tools and strategies. *Scientific Reports*, *14*(1), 273.
- Zoom Video Communications Inc. (2025). *Zoom Video Communications*.  
<https://zoom.us>

## **Acknowledgements**

I would like to express my deepest gratitude to my family, whose support has sustained me throughout this entire journey. My warmest thanks go to my wife, Silvia, who has both supported and tolerated me—even when it felt as if I was constantly heading into unexplored territory. Her patience and presence made this work possible.

I am also grateful to my parents, who have always supported my learning choices and taught me that learning is what truly matters.

My sincere thanks go to the ReSkilla team, who welcomed me as one of their own. I am also thankful to Prof. Marco Da Rin for giving me the opportunity to work within the academic environment of Tilburg University—an experience I loved so much that it became, for me, a genuine “I love Tilburg” moment.

Finally, and most importantly, I would like to thank Università Mercatorum and my tutor, Prof. Filippo Sciarrone, whose encouragement and guidance have been essential throughout this adventure. Without him, I would not be here today, and I would have missed many of the wonderful experiences of these three years.

Thank you all.

## Appendix A – Example of Concept-Map Generation and Merging

This appendix presents a full example of the concept-map generation and merging process performed by DyslexIA based on multiple assistant responses within a single learner session. The aim is to show in a transparent manner how the system extracts concepts, identifies relations, and produces a unified representation from distinct explanatory passages.

The example concerns a brief exchange on quantum-computing concepts. The learner first asks: *Tell me what a qubit is*, and subsequently requests clarifications on *superposition* and *entanglement*. The assistant's responses provide definitional elements and semantic relations that DyslexIA uses as input for the construction of a cumulative concept map.

The following transcript reports the interaction in full, documenting the linguistic material from which the system derived concepts and relational links.

Student:

*Tell me about what a qubit is.*

Assistant:

*A qubit is the basic unit of information in quantum computing. Unlike a classical bit, which can be either 0 or 1, a qubit can be in the state 0, in the state 1, or in a combination of both states at the same time, a phenomenon known as superposition. When measured, the qubit collapses to either 0 or 1. Qubits can also be entangled, meaning that the state of one qubit depends on the state of another, even if they are far apart. These properties allow quantum computers to solve certain problems much faster than classical computers.*

Student:

*What is superposition and what about entanglement?*

Assistant's reply:

*Superposition means that a qubit can be in multiple states simultaneously, with specific probabilities. When measured, this combined state collapses to either 0 or 1. Entanglement is a special connection between two or more qubits: knowing*

*the state of one entangled qubit immediately determines the state of the other, regardless of the distance between them. Together, superposition and entanglement give quantum computers their unique computational capabilities.*

From these responses, DyslexIA extracted the main concepts (qubit, classical bit, quantum state, superposition, measurement, collapse, entanglement, quantum correlation, computational advantage) and identified the semantic relations connecting them (for example: «a qubit can be in superposition», «measurement causes state collapse», «entanglement creates correlations between qubits», «superposition and entanglement enable computational advantage»).

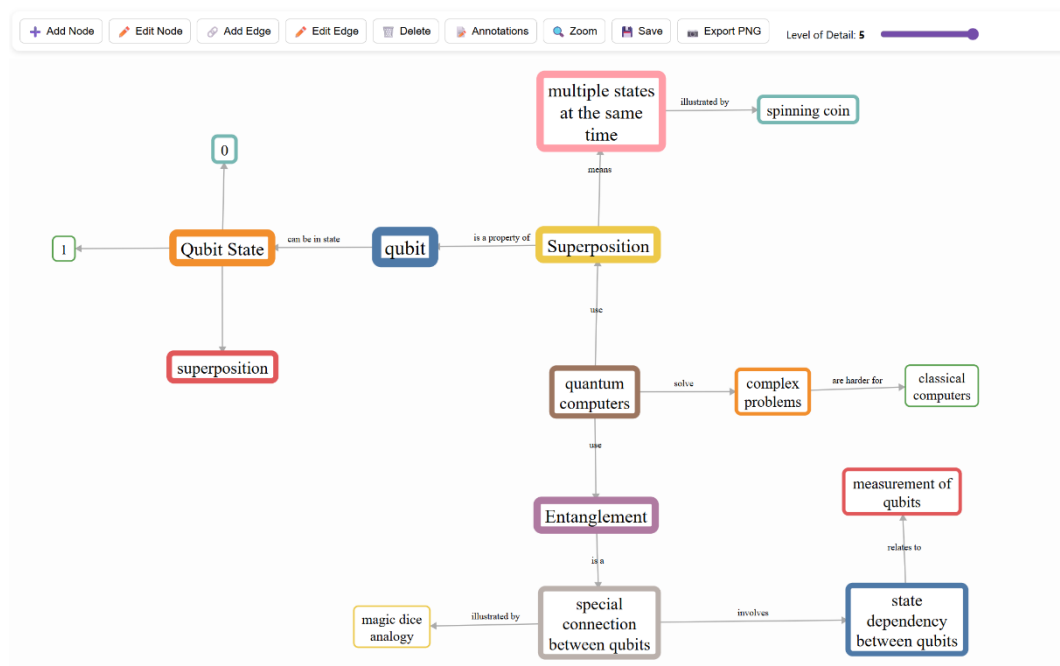


Figure 10 - Merged Concept Map from the conversation

The Figure 10 - Merged Concept Map from the conversation visualises how DyslexIA merged information from the two responses. The system harmonised overlapping nodes, eliminated duplicated relations, and preserved a coherent structure. This example illustrates the internal functioning of the concept-map generator and demonstrates the behaviour of the merging and harmonisation pipeline described in the main text.

## Appendix B – Examples of Behavioural Fine-Tuning Conversations

This appendix presents a set of sample conversation pairs used during the behavioural fine-tuning of the Smart and Smart+ modes of DyslexIA. These examples illustrate the stylistic constraints applied to optimise readability for dyslexic learners, including reduced sentence length, controlled lexical complexity, low information density, explicit structuring, and supportive conversational pacing. Each row reproduces a message sequence from the original fine-tuning dataset in JSONL format.

```
{
  "messages": [
    {"role": "user", "content": "What is an operating system?"},
    {"role": "assistant", "content": "It is a software that controls computer hardware and software. It also gives services to computer programs. It works as a link between the user and the computer hardware. Examples are Windows, macOS, and Linux. Does it make sense?"},
    {"role": "user", "content": "Tell me about its functions."},
    {"role": "assistant", "content": "Sure! A computer needs an operating system to work. It does a few key things: 1) It runs programs. 2) It gives memory to each task. 3) It helps open and save files. Does that make sense?"},
    {"role": "user", "content": "Yes, they do. Thanks for the explanation!"},
    {"role": "assistant", "content": "You're welcome! If you have more questions about tech or any other topic, I am here to help. Enjoy exploring!"}
  ],
},
{
  "messages": [
    {"role": "user", "content": "Tell me about the concept of variable in programming."},
    {"role": "assistant", "content": "A variable in programming is like a container or a box that stores data. It can hold values like numbers, text, or more complex data. The value in a variable can change or vary, hence the name. Does that make sense?"},
    {"role": "user", "content": "How do they work?"},
    {"role": "assistant", "content": "Variables are declared with a specific data type, like 'int' for integers or 'string' for text. Once declared, you can assign a value. For instance, in Python, x = 5 means that the variable 'x' now holds the value 5. Does that make sense?"},
    {"role": "user", "content": "Yes, it's clearer now. Thanks!"},
    {"role": "assistant", "content": "You're welcome! If you dive deeper into programming or have other questions, feel free to ask. Happy coding!"}
  ],
},
{
```

```

    "messages": [
      {"role": "user", "content": "What is meant by philology?"},
      {"role": "assistant", "content": "Philology is the study of
language in old written sources. It mixes linguistics, text study,
and history to understand old texts and their cultural context. In
short, it studies old texts and how languages change. Does that
make sense?"},
      {"role": "user", "content": "How is it different from
linguistics?"},
      {"role": "assistant", "content": "Good question! Both study
language, but linguistics is a wider field. It studies structure,
meaning, and use in general. Philology focuses on old texts. It
studies how language changes and the culture around those texts.
You can think of it as a mix of linguistics and history. Does that
make sense?"},
      {"role": "user", "content": "Yes, that clears things up. Thank
you!"},
      {"role": "assistant", "content": "You're welcome! If you have
more questions about humanities, languages, or any other topic, I'm
here to help. Stay curious!"}
    ]
  }

```

These examples demonstrate the stylistic constraints implemented during fine-tuning: controlled sentence length, active voice, explicit sequencing, reduced ambiguity, slow pacing, and supportive conversational closures. Smart+ fine-tuning extends this pattern by introducing reflective cues, metacognitive prompts, and gentle self-evaluation scaffolds.

## Appendix C – Tutor Mode Concept Coverage and Reflective Transition

### Concept Extraction and Saliency (Example)

When Tutor Mode is activated, DyslexIA extracts the main concepts associated with the selected document and assigns each a saliency score based on relevance and recurrence.

A brief illustrative excerpt is shown below:

```
[
  { "concept": "transformation", "saliency": 5 },
  { "concept": "connection", "saliency": 5 },
  { "concept": "journey of discovery", "saliency": 5 },
  { "concept": "stories", "saliency": 5 },
  { "concept": "friendship", "saliency": 5 },
  { "concept": "courage", "saliency": 4 },
  { "concept": "adventure", "saliency": 4 }
]
```

This list guides the system’s progression through the conceptual space of the material.

### First-Cycle Coverage

After extracting and ranking salient concepts, the system initiates the first coverage cycle. It produces an accessible explanation that integrates the highest-saliency concepts and, immediately after, generates explicit concept-related questions.

These questions fall into two categories:

- Elaboration questions, referring to the concepts already covered in the explanation;
- Exploration questions, inviting the learner to move toward salient concepts not yet addressed.

A minimal example drawn from the excerpt above:

- Elaboration questions (concepts already covered):
  - *“How does transformation differ between John’s and George’s stories?”*
  - *“In what ways does the journey of discovery shape their views on life?”*

- *“How do the concepts of connection and stories contribute to their transformations?”*
- Exploration questions (concepts not yet covered):
  - *“What role does courage play in overcoming challenges?”*
  - *“Why is nature preservation important for future generations?”*
  - *“How does connection influence our daily lives?”*

In the interface, these suggestions appear as predefined buttons to support quick navigation through the conceptual space while maintaining coherence with the source material.

### **Coverage Tracking**

As the interaction unfolds, DyslexIA keeps track of which concepts have been covered.

Covered after Turn 1:  
["transformation", "connection", "stories"]

Covered after Turn 2:  
["transformation", "connection", "stories", "friendship"]

Remaining high-salience:  
["journey of discovery"]

This ensures coherent and progressive coverage throughout Tutor Mode.

### **Transition to Reflective Mode**

Once a sufficient proportion of salient concepts has been addressed, the system shifts from comprehension to reflection. Reflective Mode generates a single open-ended question designed to stimulate critical and personal thinking.

*“How does the idea of transformation influence the way you interpret the characters’ journey?”*

## **Appendix D – Prototype Access Information**

The online DyslexIA prototype can be accessed using the following credentials (for referees only):

URL: <https://dyslexiapp.replit.app/>

Username: demo\_user

Password: demo\_password



Università telematica delle  
Camere di Commercio Italiane