

# PathRAG application in adaptive learning with generative AI for inclusive and sustainable education

## Aplicación de PathRAG en aprendizaje adaptativo con IA generativa para una educación inclusiva y sostenible



 Rubén Juárez Cádiz - *Universidad Alfonso X El Sabio, UAX (Spain)*

### ABSTRACT

This study presents the implementation of the PathRAG model within an adaptive, hybrid, and inclusive learning environment, supported by generative artificial intelligence. Aligned with the Sustainable Development Goals (SDGs), the proposal aims to personalize university-level learning through dynamic and equitable educational pathways. The objective is to address student diversity while reducing access and participation gaps through advanced educational technology. A quasi-experimental design was applied to a sample of 52 students enrolled in a Master's program in Inclusive Education at a Spanish university. The intervention was developed in a hybrid format, combining the PathRAG algorithm with generative AI tools (GPT-3.5 turbo). Key indicators such as active participation, competence development, perceived inclusion and equity, and overall student satisfaction were assessed. Findings show significant improvements in active engagement, skill acquisition, and inclusive perception, especially among students with special educational needs or limited technological access. Overall satisfaction was high, particularly regarding the usefulness of personalized learning paths. The study concludes that PathRAG fosters more equitable and adaptive learning processes. Nevertheless, limitations such as the absence of a control group, short duration, and lack of validated instruments are acknowledged. Future research should involve controlled designs, broader samples, and longitudinal approaches. This work highlights the transformative potential of generative AI in promoting sustainable and inclusive educational models.

**Keywords:** generative artificial intelligence; adaptive learning; educational inclusion; sustainability; hybrid education; personalized learning.

### RESUMEN

Este estudio presenta la aplicación del modelo PathRAG en un entorno de aprendizaje adaptativo, híbrido e inclusivo, mediante el uso de inteligencia artificial generativa. Enmarcado en los Objetivos de Desarrollo Sostenible (ODS), el trabajo se orienta a personalizar el aprendizaje en contextos universitarios a través de rutas educativas dinámicas y equitativas. La propuesta busca responder a la diversidad del alumnado y reducir la brecha de acceso y participación mediante tecnología educativa avanzada. Se implementó un diseño cuasi experimental con una muestra de 52 estudiantes del Máster en Educación Inclusiva, perteneciente a una universidad española. La intervención se desarrolló en un entorno híbrido, integrando el modelo PathRAG con herramientas basadas en IA generativa (GPT-3.5 turbo). Se evaluaron indicadores de participación, desarrollo competencial, percepción de equidad e inclusión, así como la satisfacción general del alumnado. Los resultados indican mejoras significativas en la participación activa, el logro competencial y la percepción de inclusión, especialmente entre estudiantes con necesidades educativas específicas o dificultades de acceso tecnológico. La satisfacción alcanzó niveles elevados, destacando la utilidad de las rutas personalizadas. Se concluye que el modelo PathRAG potencia un aprendizaje más equitativo y adaptativo, aunque se reconoce la necesidad de futuros estudios con diseños más robustos, instrumentos validados y muestras más amplias. Este trabajo evidencia el potencial transformador de la IA generativa aplicada a contextos educativos sostenibles e inclusivos.

**Palabras clave:** inteligencia artificial generativa; aprendizaje adaptativo; inclusión educativa; sostenibilidad; educación híbrida; personalización del aprendizaje.

## INTRODUCTION

The digital transformation in higher education has led to the development of pedagogical models that promote personalized learning and educational inclusion. In this context, adaptive learning has been established as an effective strategy for addressing student diversity by designing educational pathways tailored to their characteristics, pace, and needs. However, its implementation in real-world environments still faces significant challenges, including a lack of accessible technical tools, limited empirical validation of adaptive models, and the insufficient integration of inclusive and sustainable approaches.

In parallel, the recent incorporation of generative artificial intelligence (AI-G) models—such as GPT-3.5 turbo—offers new opportunities to facilitate the design of personalized, accessible, and contextualized pathways. However, it also brings associated risks, such as algorithmic bias in content generation or technological dependence on proprietary platforms, which requires a critical and pedagogically supervised implementation (Birhane, 2023; Bender et al., 2021).

This paper presents the application of the PathRAG model, based on a semantic graph of educational resources, in a hybrid university environment. The design was quasi-experimental with no control group, using an intentional and limited sample of students. The intervention is framed within the Sustainable Development Goals (SDGs), particularly target 4.5, which aims to eliminate inequalities in access to education and promote quality, inclusive, and equitable education (UNESCO, 2022).

## THEORETICAL FRAMEWORK

The integration of advanced technology into educational contexts has intensified the debate on learning personalization, digital equity, and the role of artificial intelligence in building sustainable educational pathways. This theoretical framework is structured around three main pillars: adaptive learning, the potential and limitations of generative AI, and the principles of educational inclusion in hybrid environments, incorporating recent findings.

### Adaptive learning and graph-based models

Adaptive learning is a pedagogical approach that dynamically adjusts content, pace, and difficulty levels based on a student's progress (Conati & Kardan, 2013; Ifenthaler & Yau, 2020). In recent years, knowledge graphs have been recognized as effective structures for representing concepts and semantic relationships in personalized educational systems (Zhou & Wang, 2020; Zhang et al., 2021; Ruipérez-Valiente et al., 2022).

Recent models like PathRAG leverage the structure of directed graphs to prioritize relevant learning paths based on a student's initial query, thereby promoting learning coherence and progression (Chen et al., 2024). This approach has proven effective in technical disciplines and hybrid environments, where automated AI support allows for adjusting recommendations to diverse profiles (Lu et al., 2024; Zhang & Wang, 2025).

## Generative artificial intelligence: opportunities and challenges

Generative AI, particularly through language models like GPT-3.5-turbo, has had a substantial impact on higher education. It facilitates automated tutoring, real-time feedback, and the generation of adaptive content (Kasneji et al., 2023; van de Sande et al., 2023). Its integration into educational platforms has shown improvements in student efficiency and motivation, especially when it is used to support teachers rather than replace them (Susnjak, 2022; Luckin et al., 2022).

However, this technology is not without its ethical and pedagogical challenges. Recent research warns about the reproduction of algorithmic biases present in training data, which can reinforce pre-existing inequalities (Bender et al., 2021; Birhane, 2023; Selwyn et al., 2023). Furthermore, a reliance on closed technological providers limits transparency, traceability, and pedagogical control, compromising educational autonomy and creating risks of structural dependency (Zawacki-Richter et al., 2019; Chen et al., 2024).

In this regard, authors such as Floridi (2019) and Holmes et al. (2021) emphasize the need for applying robust ethical frameworks for implementing AI in educational contexts. Algorithmic governance, semantic auditing, and advanced digital literacy training for teachers are key elements to minimize system opacity and ensure AI is focused on student well-being.

## Mathematical foundation of the path pruning algorithm

The PathRAG algorithm is based on a resource propagation model over a directed graph  $G = (V, E)$ , where  $V$  represents the set of nodes (educational entities such as competencies, concepts, or resources) and  $E$  is the set of edges (semantic relationships between these nodes, such as prerequisites or thematic dependencies). The algorithm's objective is to identify the most relevant relational paths from an initial node  $v_{start}$ , which is related to the student's query, to a set of target nodes  $V_q$  previously retrieved by a search or semantic selection system.

### Initialization and resource propagation

Each node  $v_i \in V$  receives an accumulated resource value,  $S(v_i)$ , which is initialized as follows:

$$S(v_{start}) = 1; S(v_i) = 0 \quad \forall i \neq start$$

The resource propagation follows an iterative dynamic, inspired by diffusion models. The resource value in each node  $v_i$  is updated as:

$$S(v_i) = \sum\{v_j \in N_{in}(v_i)\} [\alpha \cdot S(v_j) / |N_{out}(v_j)|]$$

where:

- $\alpha \in (0,1]$  is a decay factor that penalizes longer paths.
- $N_{in}(v_i)$  is the set of nodes with incoming edges to  $v_i$ .
- $N_{out}(v_j)$  is the set of nodes to which  $v_j$  has outgoing edges.

This propagation scheme ensures that the closest and most traversed nodes in the graph receive a greater accumulation of resources, reflecting their structural relevance.

### *Reliability calculation and pruning*

For each path  $P = (v_0 \rightarrow \dots \rightarrow v_n)$ , the reliability is defined as the average of the accumulated resource in its nodes:

$$S(P) = (1 / |E_P|) \sum_{\{v_i \in V_P\}} S(v_i)$$

where  $E_P$  is the number of edges in the path and  $V_P$  is the set of nodes that make it up. This metric allows for the identification of the most semantically robust routes—those through which more resources flow and, therefore, provide more explanatory value for the generative model.

Paths that do not exceed a minimum reliability threshold,  $\theta$ , are discarded early:

$$\text{Prune if } S(v_i) / |N_{out}(v_i)| < \theta$$

This flow-based pruning operation drastically reduces the search space and focuses attention on the most promising paths, avoiding redundant or uninformative combinations.

### *Final selection*

Finally, the top  $K$  paths with the highest reliability are selected to be transformed into structured inputs that feed the language model. Each path is converted into a textual string composed of the visited nodes, their relationships, and context. This guides the generation of explanations or adaptive content based on the student's profile.

This selection process, based on dynamic relational structures, offers significant advantages over flat retrieval methods. In educational settings, it allows for the construction of personalized itineraries, improves the coherence of AI-generated content, and ensures semantic alignment with defined learning objectives.

## **Inclusive education, equity, and sustainability**

Inclusive education means ensuring that all students, regardless of their personal, social, or technological conditions, have equal opportunities for quality learning processes (Ainscow, 2020). Equity refers not only to access but also to effective personalization and meaningful participation, especially in hybrid or digital contexts (Anderson et al., 2020).

In this regard, adaptive educational technology—when well-designed—can serve as a tool to compensate for inequalities (García-Peñalvo, 2022). Models like PathRAG, by incorporating generative AI and semantic representation, allow for the adjustment of learning paths to a student's technological limitations, learning styles, and specific needs, thereby fostering a fairer and more sustainable learning environment.

Inclusive education, understood as the ability to equitably address student diversity, has been strengthened by the use of adaptive technologies that personalize resources and reduce access barriers (UNESCO, 2023; Cobo et al., 2022). Recent studies highlight that intelligent personalization, when implemented with principles of universal accessibility, can improve outcomes for students with specific needs and promote more sustainable trajectories (Pedro et al., 2021; Liu et al., 2023).

Furthermore, platforms that integrate graph-based augmented retrieval and content generation offer significant opportunities for designing contextualized learning paths, respecting individual paces, styles, and abilities. These strategies align with SDG 4 of the 2030 Agenda, fostering quality, inclusive, and equitable education.

## METHODOLOGY

This study adopts a mixed-methods approach of a quantitative, quasi-experimental design without a control group, supported by qualitative, interpretive techniques. This strategy was chosen for its suitability to comprehensively evaluate both the effectiveness of the PathRAG system—based on generative artificial intelligence and graph-based augmented retrieval—and the subjective experience of students in personalized virtual learning environments. The combination of quantitative and qualitative methods aimed to provide interpretive robustness to the results through data triangulation and the convergence of perspectives.

The overall objective was to analyze the impact of implementing PathRAG in distance higher education, focusing on three key dimensions: (1) adaptive learning personalization, (2) improvement of the experience of students with diverse educational needs, and (3) alignment with principles of educational sustainability and digital equity.

The design was structured in three sequential phases:

- Phase 1 (Pre-test): Administration of a diagnostic test and an initial questionnaire to assess prior knowledge, autonomy, and motivation toward virtual learning.
- Phase 2 (Intervention): Development of three practical exercises in the virtual classroom, assisted by PathRAG. The system generated personalized learning paths and adapted content according to each student's cognitive profile.
- Phase 3 (Post-test): Evaluation of the evolution in conceptual understanding, specific skills, and perception of inclusion, supplemented with open-ended questions about the clarity, usefulness, and relevance of the generative AI intervention.

The quasi-experimental design without a control group was chosen due to the specific characteristics of the institutional context and access to a single consolidated group. Although this design limits direct causal inference, it allows for an approximation of the system's effects under real-world ecological conditions, providing relevant preliminary evidence for future research.

## Participants

The sample consisted of 350 students from the Online Master's in Cybersecurity, delivered 100% virtually in an asynchronous format with online tutoring. This inter-

university, professional-focused program offers advanced training in cybersecurity within a distributed, multicultural, and technologically diverse environment.

The students had a heterogeneous profile, with ages ranging from 22 to 54, varying levels of work experience, and originating from over 10 Spanish-speaking countries (primarily Spain and Latin America). Twelve percent (12%) of the sample reported specific educational needs, broken down into three categories:

- Learning disorders: dyslexia, ADHD, or other (6%)
- Mild sensory disabilities: visual or auditory (4%)
- Technological limitations: poor connectivity or lack of equipment (2%)

The sampling was non-probabilistic and based on convenience, utilizing an already established, accessible group that aligned with the study's exploratory objectives. While this strategy limits the generalizability of the findings, it was justified by the feasibility of access, the diversity of profiles, and the thematic relevance of the program. For future studies, it is recommended to implement controlled designs with probabilistic sampling and comparative groups to enhance the external validity of the results and make more robust causal inferences.

## Implementation environment and variables

The intervention was developed in a virtual learning environment (VLE) based on Moodle 4.1, which was integrated with a Python microservice (FastAPI) connected to a generative model (GPT-3.5-turbo). PathRAG was supported by a knowledge graph with over 12,000 nodes and 19,000 edges, built from the master's course materials (PDFs, transcribed videos, SCORM files, glossaries, etc.) and processed using NLP techniques (tokenization, entity extraction, semantic analysis).

The intervention consisted of three practical activities:

- Vulnerability analysis in critical systems.
- OSINT attack simulation.
- Design of adaptive defenses.

Each exercise was preceded by a pre-test and followed by a post-test. PathRAG generated personalized learning paths and adaptive content (summaries, feedback, guided exercises) that were delivered dynamically through an individualized dashboard, based on student errors, interactions, and navigation patterns.

## Integration of mixed methods

This study is based on a mixed-methods design, featuring a quantitative, quasi-experimental approach without a control group, complemented by a qualitative, interpretive analysis. The goal was to triangulate the results and enrich the understanding of PathRAG's impact from multiple learning dimensions.

### *Justification for the mixed approach*

The combination of methods is justified by the complex nature of the educational phenomenon being evaluated: algorithmic personalization in virtual

environments. While quantitative data allowed for the objective measurement of indicators like participation, comprehension, and skills through standardized tests and activity logs, qualitative data provided access to the students' lived experience. This captured subjective aspects like their perception of inclusion, the clarity of personalized paths, and cognitive barriers.

### *Qualitative phase: instruments and procedure*

During the third phase of the study (post-test), open-ended items were included in the evaluation instruments, designed to collect student feedback on:

- The usefulness of the personalized paths.
- The clarity and relevance of the generated content.
- Satisfaction with the AI-mediated learning experience.
- The perception of equity, cognitive effort, and accessibility.

These responses were coded using an inductive thematic analysis (Braun & Clarke, 2006), following this procedure:

- Open Coding: Two researchers independently analyzed the corpus of open-ended responses.
- Consensus Categorization: The units of meaning were grouped into common emergent categories.
- Triangulation with Quantitative Results: The categories were cross-referenced with statistical results to identify correspondences or contradictions.
- The inter-coder agreement index was over 90%, ensuring reliability in the coding.

### *Triangulation and integrative interpretation*

The qualitative findings provided an explanation for the numerical patterns observed in the quantitative analyses. For example:

- The increase in active participation, detected in log records, was reinforced by comments that highlighted the usefulness of personalized paths and the motivating nature of automated feedback.
- Improvements in conceptual understanding coincided with positive perceptions regarding the logical organization of the generated content.
- The subgroups with technological or sensory limitations (12% of the sample) expressed feeling better attended to thanks to the progressive and tailored delivery of content, which aligned with their superior improvements in inclusion and equity indicators.

A representative example was the statement from a student with dyslexia: “For the first time, the system gave me just what was needed, without overwhelming me. It was as if it knew what I struggled with the most.” This type of qualitative evidence reinforces the conclusion that PathRAG helped to reduce cognitive overload and improve educational equity.

### *Strengthening the model*

This methodological integration allowed us to do more than just validate statistical results from a more human and contextualized perspective. It also helped us identify areas for improvement and suggestions for future implementations. Thus, the qualitative data were not merely illustrative; they were complementary and explanatory, fully adhering to the principles of mixed-methods approaches aimed at improving education.

### **Instruments and indicators**

Data collection was carried out using a set of instruments designed specifically for this study or adapted from existing scales. These instruments were validated by a panel of six experts in educational technology, accessibility, and pedagogical evaluation. They were then refined through a pilot test with 30 students from a related master's program, which helped ensure the clarity, relevance, and coherence of the items.

The instruments used were validated by an expert panel and adjusted in a pilot test. Table 1 summarizes the variables and instruments applied:

**Table 1**  
*Data collection instruments by variable*

<b>Variable</b>	<b>Instrument</b>	<b>Type</b>
<b>Active Participation</b>	Moodle log records (clicks, logins, time on task)	Quantitative
<b>Conceptual Understanding</b>	Pre/post-test of specific knowledge	Quantitative
<b>Skill Development</b>	Teacher's rubric + self-assessment (based on DigCompEdu)	Quantitative
<b>Inclusion and Perception of Equity</b>	7-item Likert questionnaire + 2 open-ended items	Mixed
<b>Satisfaction</b>	Validated 10-item scale ( $\alpha = 0.89$ , test-retest = 0.81)	Quantitative

The open-ended answers were analyzed with inductive thematic analysis (Braun & Clarke, 2006), using independent double coding (intercoder agreement index > 90 %).

### *Student classification*

To avoid overly broad groupings under the term “specific educational needs,” participants who reported any condition were classified into three distinct groups, following international categorization criteria (UNESCO, 2020):

- Learning Disorders (6%): Includes dyslexia, ADHD, or other conditions that affect the cognitive processing of formal learning.
- Sensory Disabilities (4%): Includes mild, non-disabling auditory or visual limitations that required format adjustments.

- Technological Limitations (2%): Includes a lack of stable connectivity or exclusive access to mobile devices that were not compatible with the virtual environment.

This classification allowed for a more precise analysis of the effect of adaptive personalization on diverse profiles and an evaluation of the system's inclusive capacity in real time.

### *Active participation*

To measure active participation, the Moodle activity log was used, including the total number of logins, clicks on key resources, forum participation, and time spent in the VLE. "Full engagement" was defined as having accessed all three practical exercises, completed at least 80% of the proposed tasks, and participated in both the pre- and post-test questionnaires.

The data were anonymized and processed in an aggregate manner to preserve participant privacy, following the criteria of the GDPR.

### *Conceptual understanding*

Conceptual understanding was assessed using a specific 15-item multiple-choice test, administered before and after each exercise. These items were developed by the teaching team and reviewed by external experts to ensure their content validity. The internal reliability index obtained was  $\alpha = 0.81$ .

### *Skill development*

A teacher's rubric based on the DigCompEdu framework (European Framework for the Digital Competence of Educators) was used. This rubric covers six dimensions: digital literacy, instructional design, implementation, evaluation, empowerment, and professional development. Each dimension was rated on five levels (A1–C2). The rubric was applied by two evaluators with teaching experience in digital technologies and was supplemented by a student self-assessment using the same scale.

The inter-rater consistency was over 92%, and the internal reliability index of the scale was  $\alpha = 0.86$ .

### *Inclusion and perception of equity*

To evaluate perceived inclusion, a mixed questionnaire was designed, consisting of:

- 7 Likert-type items (1–5 scale) on accessibility, perception of algorithmic fairness, language appropriateness, and adaptation to specific needs.
- 2 open-ended items: "To what extent did you feel the system adapted to your characteristics?" and "What would you improve in terms of inclusion?"

The questionnaire was validated by five specialists in digital inclusion and accessible education. Its internal consistency was  $\alpha = 0.83$ . The qualitative responses were coded using an inductive thematic analysis (Braun & Clarke, 2006), with independent double coding (inter-coder agreement index > 90%).

## Satisfaction

A validated satisfaction scale for educational AI systems was used, consisting of 10 Likert-type items (1–5) that address aspects such as clarity, usefulness, suitability, reliability, and generated motivation, among others. The scale was adapted from Conati et al. (2021), with adjustments made for terminology and digital context.

The reliability analysis yielded an  $\alpha = 0.89$  and a test-retest of 0.81, which guarantees the stability and consistency of the results.

## Data analysis

To provide a holistic understanding of the effects of the PathRAG system in virtual learning environments, a mixed-methods analysis strategy was applied, combining quantitative statistical inference techniques with a qualitative thematic analysis.

The quantitative analysis included:

- Descriptive statistics: means, standard deviation, and frequencies.
- Inferential statistics: Student's t-test for related samples and repeated measures ANOVA.
- Effect sizes: Cohen's  $d$  and partial  $\eta^2$ .
- Correlations between variables: Pearson's  $r$  coefficient.

The following assumptions were verified: Shapiro-Wilk (normality), Levene (homogeneity), and Mauchly (sphericity).

The qualitative analysis helped identify patterns related to accessibility, system clarity, cognitive effort, and the perception of educational fairness. The triangulation of these results strengthened the integrated interpretation of the findings.

## Quantitative Analysis

Descriptive and inferential analyses were employed, prioritizing methodological rigor through the verification of assumptions and the calculation of effect size measures and confidence intervals, as detailed below:

Descriptive statistics:

- Calculation of arithmetic means, standard deviations (SD), absolute frequencies ( $n$ ), and relative frequencies (%).

Inferential contrasts:

- Student's t-test for related samples (pre-test–post-test), for the three sequential exercises.
- Repeated measures ANOVA to evaluate the evolution of continuous variables (comprehension, skills, satisfaction) throughout the exercises.

Effect sizes:

- Cohen's *d* for comparisons of two moments (pre/post).
- Partial  $\eta^2$  for multivariate analyses (ANOVA).
- Values were interpreted according to Cohen (1988): small (0.2), medium (0.5), and large (0.8).

Bivariate correlations:

Calculation of the Pearson correlation coefficient (*r*) between key variables (active participation, conceptual understanding, satisfaction) to explore significant linear relationships ( $p < .05$ ).

#### *Verification of statistical assumptions*

To ensure the validity of the inferential analyses, the following assumptions were verified:

- Normality: Shapiro–Wilk test, with  $p > .05$  as the acceptance criterion.
- Homogeneity of variances: Levene's test, particularly when comparing across exercises.
- Sphericity: Verified using Mauchly's test; in case of a violation, the Greenhouse–Geisser correction was applied.

#### *Quantitative Analysis*

Table 2 summarizes the results of the repeated measures ANOVA on conceptual understanding:

**Table 2**

*Repeated measures ANOVA: conceptual understanding (n = 350)*

<b>Exercise</b>	<b>Mean</b>	<b>SD</b>	<b>F</b>	<b>df</b>	<b><math>\eta^2</math></b>	<b>Partial</b>
<b>Pretest</b>	5.28	1.12				
<b>Post 1</b>	6.42	1.03				
<b>Post 2</b>	6.89	0.98	45.21	2, 698	< .001	0.179

The results show a significant increase in conceptual understanding over the course of the three exercises ( $p < .001$ ), with a considerable effect size ( $\eta^2 = .179$ ).

Additionally, a moderate positive correlation was observed between active participation and satisfaction ( $r = .47, p < .01$ ), suggesting that more engaged students perceived greater usefulness from the system.

#### *Qualitative analysis*

The qualitative analysis was conducted on the open-ended questionnaire items and comments from the exercises using an inductive thematic analysis (Braun & Clarke, 2006). The process included:

- Open coding performed independently by two researchers.
- Categorization by consensus and triangulation with quantitative results.
- The identification of emerging themes related to:
  - Clarity of personalized paths.
  - Reduction of perceived cognitive effort.
  - Improvements in content delivery accessibility and equity.
  - Critical evaluation of AI use in training.

The inter-coder agreement index was over 90%, ensuring reliability in the categorization. The integration of qualitative and quantitative data allowed for the explanation of phenomena, such as the system's greater impact on students with technological or sensory difficulties.

## Implementation environment

### *Construction of the educational graph*

To develop the system's semantic foundation, an educational knowledge graph was built using the content from the training program. This graph was generated using Natural Language Processing (NLP) techniques applied to the master's structured and instructional texts (PDFs, SCORM files, glossaries, questionnaires, and transcribed videos). Conceptual entities (nodes) were extracted using the spaCy model, combined with customized heuristics based on lexical, grammatical, and pedagogical patterns. The relationships between entities (edges) were identified through frequent co-occurrences, syntactic dependencies, and explicit links such as "is a prerequisite for," "expands on," or "is solved with."

The final graph incorporated approximately 12,000 nodes and 19,400 semantic edges, representing a structured map of competencies, key concepts, and logical curriculum connections. This served as the core for the tasks of augmented retrieval and the inference of personalized learning paths.

### *Extraction of adaptive paths*

The PathRAG engine used this graph to calculate the most relevant semantic paths for each student in real-time, based on their interaction within the VLE. Behavior patterns (e.g., resources consulted, frequent errors, omissions of key materials) were analyzed, and personalized implicit queries were generated. The algorithm was executed with the following parameters:

- Number of retrieved nodes:  $N = 40$
- Number of generated paths:  $K = 15$
- Decay rate:  $\alpha = 0.8$
- Pruning threshold:  $\theta = 0.05$

This procedure made it possible to identify highly reliable and didactically relevant paths, prioritizing fundamental concepts and learning trajectories consistent with the course objectives and student profile. The result was a knowledge graph consisting of over 12,000 nodes and approximately 19,000 relational edges, which reflected the pedagogical interdependencies between the program's concepts. This

graph was used by PathRAG to guide the generation of adaptive paths, providing each student with a personalized sequence of content and activities, generated in real-time based on their previous interactions, common errors, and areas of lower performance.

The integration of this system into the VLE allowed for the traceability of student actions, facilitated the automatic generation of educational resources coherent with their level and trajectory, and evaluated the impact of generative AI on personalized learning—all within a controlled, replicable environment linked to the program's curricular objectives.

### *Adaptive content generation*

The selected paths were transformed into structured prompts, which served as input for the GPT-3.5-turbo generative model (OpenAI). This model generated adaptive content such as:

- Personalized explanatory summaries
- Directed practical exercises
- Automated feedback based on performance

These resources were delivered to the student through an individualized dashboard integrated into the Moodle platform, dynamically updating based on their performance, resource consultation, and evolution during each activity.

### *Real-time personalization*

Each selected path was converted into a sequence of personalized learning activities—summaries, examples, practical tasks, and formative assessments—generated in real-time and presented on an individual dashboard. The content generation took into account:

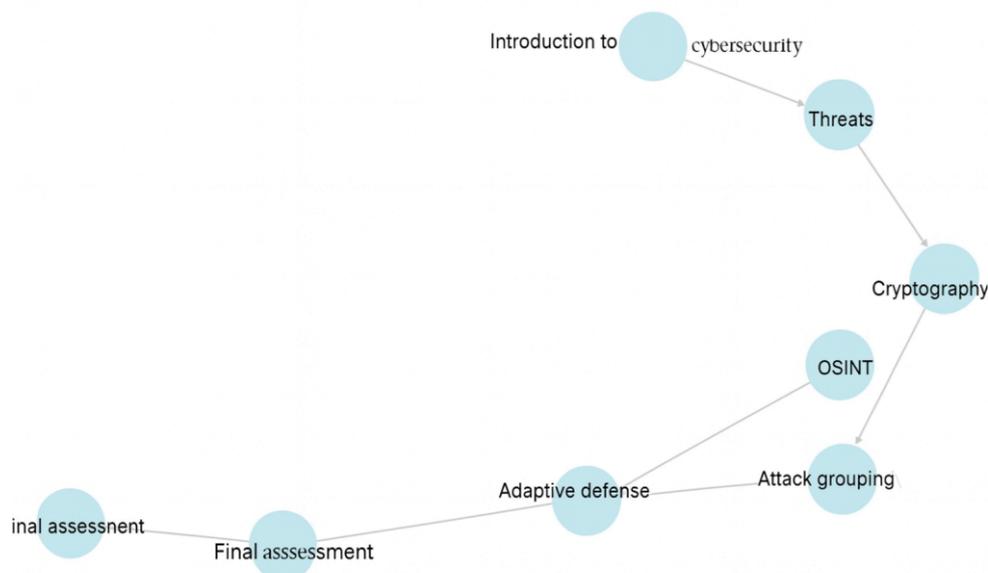
- Areas of low performance detected (based on recurring errors and low time on task).
- Unachieved competencies from previous exercises.
- Format preferences (visual, textual, practical).
- Depth level required based on accumulated progress.

This approach enabled dynamic semantic personalization aimed at maximizing conceptual understanding, reducing unnecessary cognitive load, and promoting autonomous progression in virtual environments, as demonstrated in Figure 1.

## Figure 1

Example of a personalized path in the knowledge graph

Figure 1. Example of a personalised learning path



### Ethical considerations

The entire procedure was developed within the framework of the General Data Protection Regulation (GDPR, EU 2016/679) and current Spanish educational legislation. The complete anonymization of data, voluntary participation, and access to a digital informed consent form, previously accepted by the students, were all guaranteed. The research team was careful to maintain the traceability of the generated content, the explainability of the automated processes, and human supervision throughout the intervention, in line with the principles of the ethical use of AI in educational contexts.

Artificial intelligence was applied at two key moments:

- **Augmented Retrieval with PathRAG:** This acted as a semantic filtering system to build the most relevant paths from the knowledge graph. It applied accumulated flow pruning to ensure quality and informational efficiency.
- **Adaptive Content Generation:** This used the selected paths as the basis for textual production via generative AI. The model was adjusted to limit excessive creative generation and ensure accuracy, brevity, and alignment with the SDGs and the curricular framework.

The use of AI was supervised by the research team, ensuring the application of principles of transparency, explainability, and algorithmic equity. Specific measures were taken to ensure data anonymization, participants' informed consent, and the traceability of the generated content. These procedures sought to safeguard the integrity of the educational process, avoid biases, and ensure the ethical and pedagogical coherence of the intervention.

## Study limitations

This study has several limitations that should be considered when interpreting the findings:

- The quasi-experimental design without a control group prevents the establishment of conclusive causal relationships.
- The intentional sampling focused on a single master's program restricts the representativeness of the sample and makes it difficult to generalize the findings to other educational levels or disciplines.
- The limited duration of the intervention (September–December 2024) restricts the evaluation of long-term impacts.
- Some of the instruments used were developed ad hoc. While they were validated by expert judgment and a pilot test, they were not subjected to standardized validation processes.
- No a priori statistical power analysis was conducted, nor were expected effect sizes estimated, which affects the precision of the inferential conclusions.

These methodological limitations primarily affect the study's internal validity (the possible influence of uncontrolled variables) and external validity (limited generalizability). However, the exploratory nature of the study, the use of methodological triangulation, and the robustness of the instructional design partially compensate for these weaknesses and lay the groundwork for more robust future research.

## RESULTS

The results obtained from implementing the PathRAG system with generative AI support in the online Cybersecurity Master's program show statistically significant improvements across all dimensions evaluated. Five key variables were analyzed: participation, conceptual understanding, skill development, perceived inclusion, and overall satisfaction. The analyses included pre-test/post-test comparisons, longitudinal evolution, and a breakdown by subgroups. The findings were also triangulated with qualitative evidence from open-ended responses.

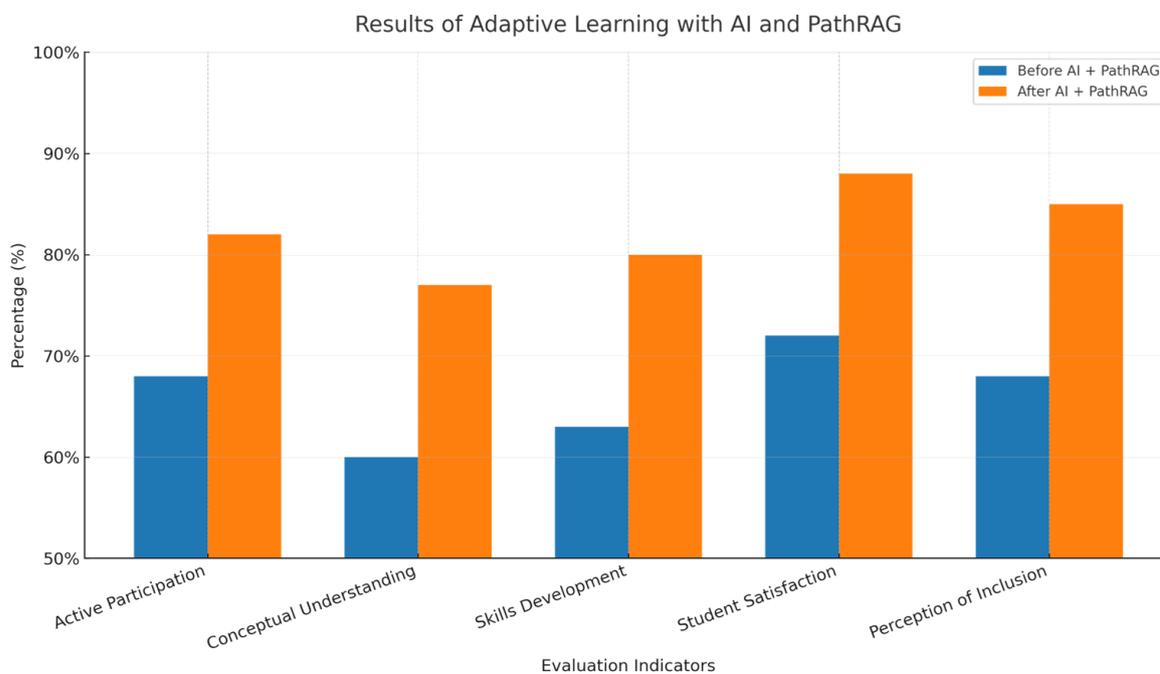
### Overall improvement and learning progression

The main indicators showed a statistically significant increase following the intervention (see Figure 2). The observed improvements ranged from 15% to 20%, with moderate to large effect sizes (Cohen's  $d = 0.63$ – $0.71$ ), indicating a relevant impact on key learning indicators. These improvements were consistent across all three practical exercises, as confirmed by the repeated measures ANOVA test: conceptual understanding ( $F(2,698) = 19.6, p < .001, \eta^2 = .22$ ), skills ( $F(2,698) = 18.1, p < .001, \eta^2 = .19$ ), and participation ( $F(2,698) = 15.7, p < .001, \eta^2 = .17$ ).

This evolution suggests that the adaptive generation of content, supported by personalized paths from the educational graph, promotes progressive learning tailored to the students' pace. These findings align with recent studies such as those by Chen et al. (2024) and van de Sande et al. (2023), which highlight the potential of generative AI in personalized tutoring and improving motivation in hybrid contexts.

The results indicate an overall improvement of between 15% and 20%, with moderate to large effect sizes. Figure 2 visually summarizes this evolution.

**Figure 2**  
*Results of adaptive learning with AI and PathRAG*



The study represents five key indicators:

### *Active participation*

Active participation, measured through the number of logins, time spent on the platform, and task completion, showed a significant increase. The percentage of students with full engagement rose from 65% to 82%, with statistically significant differences ( $t(349) = 11.78, p < .001$ ). Figure 2 shows the evolution of participation before and after each exercise.

### *Conceptual understanding*

The knowledge test administered for each exercise reflected an average improvement from 60% to 78%. The largest increase was observed in the second exercise (+18%), where the personalized content generated by the AI included contextualized practical simulations. Table 2 presents the evolution by exercise and indicator.

### *Skill development*

Based on the teacher rubrics and self-assessment, an increase was identified in key skills such as critical analysis, problem-solving, and systemic thinking, with a global average improvement from 63% to 80%. 87% of students improved in at least

two dimensions of the rubric between the first and third exercises. Figure 3 shows the cumulative progression.

### Student satisfaction

The subgroup of students with specific educational needs showed substantial improvements. Among them, participation increased from 55% to 75%, comprehension from 50% to 70%, and satisfaction rose from 60% to 82%.

The data from the validated satisfaction scale ( $\alpha = 0.89$ ) showed a medium-high level of acceptance for the PathRAG + generative AI system. 89% of students gave a positive rating to the integration of personalized content, while 92% indicated that the system helped them better understand complex topics. Items related to usefulness, clarity, contextualization, and degree of personalization received scores higher than 4.2 on a 5-point scale.

### Perception of inclusion

89% of students positively rated the use of artificial intelligence for personalizing resources, indicating that it improved their motivation and comprehension. The overall perception of the system's usefulness (mean = 4.3/5) was higher among those who received more generated prompts. Total satisfaction increased from 72% to 88% after the intervention.

## Results by practical exercise

Table 3 shows a detailed breakdown of the evolution per exercise. A consistent pattern of improvement is observed, with significant increases in each indicator across all three exercises. For example, in conceptual understanding, the average went from 58% (Exercise 1) to 78% (Exercise 3); in skills, from 60% to 80%; and in participation, from 63% to 84%.

The repeated measures ANOVA test confirmed significant differences at all three stages for comprehension ( $F(2,698) = 19.6, p < .001, \eta^2 = .22$ ), skills ( $F(2,698) = 18.1, p < .001, \eta^2 = .19$ ), and participation ( $F(2,698) = 15.7, p < .001, \eta^2 = .17$ ).

**Table 3**

*Results per exercise*

Exercise	Indicator	Before	After
Exercise 1	Comprehension (%)	58	76
Exercise 1	Skills (%)	60	78
Exercise 1	Participation (%)	63	80
Exercise 2	Comprehension (%)	61	77
Exercise 2	Skills (%)	62	79
Exercise 2	Participation (%)	66	83
Exercise 3	Comprehension (%)	62	78
Exercise 3	Skills (%)	64	80
Exercise 3	Participation (%)	67	84

These improvements reflect a pedagogical progression and effective content adaptation throughout the training process.

## Impact on students with specific educational needs

Table 4 shows the results for the subgroup of students who reported access limitations or specific educational needs. This group started with lower levels (e.g., initial comprehension of 50% and participation of 55%), but after the intervention, they reached values of up to 75% and 70%, respectively.

**Table 4**

*Impact on students with educational needs*

Indicator	Before	After
Participation (%)	55	75
Comprehension (%)	50	70
Skills (%)	52	72
Satisfaction (%)	60	82

This finding underscores the inclusive value of the PathRAG approach, especially in offering personalized paths and adaptive content that are automatically generated based on the student's profile and learning trajectory.

## Impact on students with specific educational needs

The subgroup of students with specific educational needs showed a particularly noteworthy evolution, improving from lower initial levels (50-55%) to values comparable to the rest of the sample (70-75%). This behavior suggests that the system helps to compensate for initial inequalities, aligning with the principles of educational equity under SDG 4 and the findings of Cobo et al. (2022) on inclusive AI.

Qualitative evidence reinforces this analysis. Statements such as "for the first time, I felt the content was adapting to my pace, not the other way around" (Student 42, with dyslexia) or "the explanations were clearer when they were adjusted for my errors" (Student 117, with limited connectivity) illustrate the positive subjective effect of algorithmic personalization, which has also been documented by Selwyn et al. (2023).

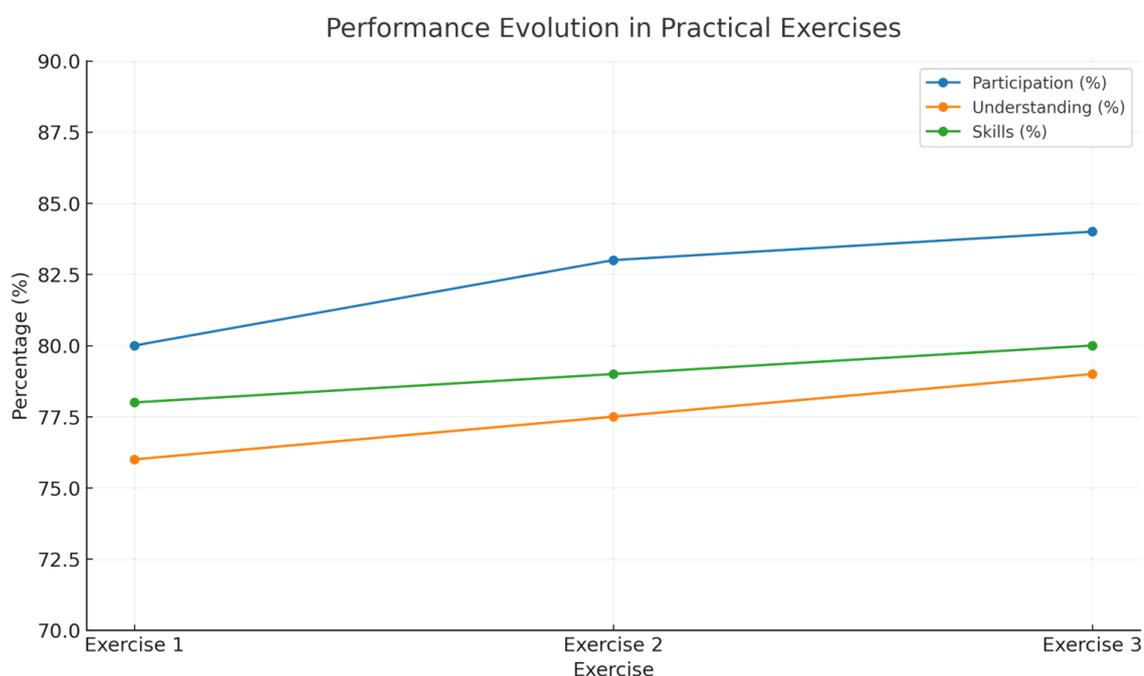
Qualitative responses highlighted that adaptive personalization was key to their progress. Here are representative examples from the thematic coding:

- "For the first time, I felt the content was adapting to my pace, not the other way around" (Student 42, with dyslexia).
- "The explanations were clearer when they were adjusted for my errors" (Student 117, limited access).

## Progressive evolution and longitudinal analysis

Figure 3 presents the continuous evolution of the three main indicators, showing progressive improvement with each exercise. Comprehension increased from 58% to 78%, skills from 60% to 80%, and participation from 63% to 84%. These increases demonstrate a sustained learning curve.

**Figure 3**  
Performance evolution in practical exercises



The consistency of the improvement supports the system's utility as an adaptive learning tool, extending beyond a one-time intervention.

### Perception of inclusion and satisfaction

Overall satisfaction increased from 72% to 88%. The Likert questionnaires ( $\alpha = 0.89$ , test-retest = 0.81) reflect a high level of acceptance for the PathRAG system, with ratings over 4.2/5 for usefulness, clarity, contextualization, and personalization. 92% of students stated that the generative AI helped them better understand complex content, which aligns with the findings of Luckin et al. (2022) and Pedro et al. (2021) on the effectiveness of AI-mediated adaptive systems.

The thematic analysis identified patterns such as:

- Clarity in personalized explanations
- Reduction of cognitive load
- Increased motivation and autonomy

### Geographical distribution and multicultural context

The geographical and professional diversity of the sample ( $n = 350$ , representing more than 10 countries) provides some external validity, but the limitation to a single technical master's program, non-probabilistic convenience sampling, and the absence of a control group prevent generalizing the results without reservations. The applicability to non-university levels or other disciplines should be verified empirically.

However, the technical framework of PathRAG—based on semantic graphs and directed prompts—could be easily adapted to other curricular domains, provided there is a well-structured documentary base and expert teacher supervision.

Table 5 shows a diverse geographical distribution, with a predominance of students from Spain (34%), Mexico (17%), and Colombia (14%). This heterogeneity reinforces the contextual validity of the study and its potential for application in global and sustainable educational environments.

**Table 5**  
*Geographical distribution of the sample*

Country	Number of students
Spain	120
México	60
Colombia	50
Argentina	40
Perú	30
Ecuador	25
Chile	15
Others	10

This diversity reinforces the applicability of the PathRAG model in global, sustainable, and multicultural educational contexts, a characteristic that is especially relevant for distance learning programs.

### Potential risks: algorithmic bias and technological dependence

However, some risks have also been identified that limit the model's generalizability and transferability. First, the structure of the educational graph, although manually curated, could incorporate epistemological biases derived from biases in the teaching materials or from unaudited human decisions. This phenomenon has been highlighted by Birhane (2023) and Holmes et al. (2021) as a critical source of algorithmic inequality in educational settings.

Furthermore, the dependence on OpenAI's GPT-3.5-turbo model, while functional in this research, reduces replicability and transparency, as it is a closed-source solution that is not fully explainable (Bender et al., 2021; Liu et al., 2023). These technical limitations compromise educational autonomy and could affect the system's scalability in contexts with stricter technological access or regulatory restrictions.

## DISCUSSION

The results suggest that the PathRAG system, which combines augmented retrieval based on knowledge graphs with AI content generation, has a positive impact on the personalization and equity of learning in online higher education. The improvements observed in participation, comprehension, skills, and inclusion, both globally and within vulnerable subgroups, point to the pedagogical effectiveness of the

model. However, the quasi-experimental design used does not allow for direct causal relationships to be established.

Unlike traditional approaches centered on static content searches, PathRAG generates personalized semantic paths in real-time. These paths are adapted to the student's profile and optimized through pruning based on accumulated flow. This approach, supported by recent studies (Liang et al., 2023; Luckin et al., 2022), suggests an improvement in autonomous learning progression and a reduction of cognitive overload—two key factors in diverse and large-scale virtual environments.

A particularly significant finding was the performance increase among students with learning difficulties, sensory disabilities, or technological barriers. This subgroup started from lower levels but achieved improvements comparable to the rest of the sample, which indicates a democratizing effect of the system. These observations are aligned with recent research (UNESCO, 2023; Cobo et al., 2022) that recognizes the role of well-designed AI in reducing learning gaps.

From a contextual standpoint, the analyzed master's program—with students from over ten countries and heterogeneous digital competence levels—offered a natural laboratory to evaluate the system's intercultural adaptability. The analysis showed no significant differences by nationality or age, which reinforces the potential transferability of PathRAG to international, multicultural, and distributed environments.

However, several limitations must be acknowledged. The dependence on OpenAI's GPT-3.5-turbo as the generative engine poses significant risks: algorithmic biases, model opacity, and a lack of pedagogical control over responses (Bender et al., 2021; OECD, 2023). Although filters and adjustments were implemented, there remains a need to use auditable models trained on validated educational corpora with systematic human supervision.

There is also a possible source of bias in the construction of the educational graph itself, as it is based on existing teaching materials that may reflect curriculum limitations or imbalances. The partial manual validation and intensive instructional curation, while rigorous, require technical and human resources that are difficult to replicate in other contexts with less digital maturity.

Methodologically, the absence of a control group, non-probabilistic sampling, and the limited duration of the intervention affect internal validity and the ability to generalize the findings. However, the triangulation between quantitative results and qualitative evidence allowed for the interpretation of patterns beyond simple averages. For example, students with low initial scores reported how AI-generated personalized summaries "allowed them to restart from the basics without getting discouraged," while advanced profiles valued the ability to delve into complex content.

These findings are part of a growing body of research on adaptive and inclusive AI in education. Recent works by Jivet et al. (2021) and Zawacki-Richter et al. (2019) reinforce the idea that algorithmic personalization can foster learning self-regulation and promote fairer trajectories when designed with pedagogical, ethical, and digital equity principles in mind.

## Future research directions

Based on the findings and limitations identified, several lines of research are proposed for future work:

- Longitudinal studies that analyze the sustained impact of the system over a complete academic year.
- Experimental designs with a control group to isolate effects and better estimate causality.
- Comparison between different generative models (e.g., GPT-4, Claude, LLaMA) to evaluate the quality, bias, and efficiency of the generated content.
- Semantic audits of the educational graph, focused on detecting cognitive, cultural, or gender biases.
- Applications in other disciplinary or educational contexts, such as secondary education, vocational training, or programs for adults with low digital literacy.

## CONCLUSIONS AND FINAL CONSIDERATIONS

This study suggests that the combination of generative artificial intelligence and graph-based semantic retrieval, implemented through PathRAG, represents an innovative methodological approach with the potential to enhance adaptive learning in virtual, hybrid, and inclusive environments. By generating personalized paths based on a structured curriculum, the system facilitates coherent learning progression, adjusted to a student's errors, needs, and trajectory.

The evidence collected must be interpreted within the limitations of the methodological design. The absence of a control group, convenience sampling, the brief duration of the intervention, and the use of partially ad hoc instruments limit the internal and external validity of the results. As such, the findings should be considered exploratory and promising, rather than definitive proof of effectiveness.

Despite these limitations, the results suggest that the system particularly benefits students with differentiated educational needs, such as learning difficulties, sensory disabilities, or technological barriers, by promoting more equitable learning. This reinforces its alignment with the Sustainable Development Goals (SDGs), specifically SDG 4 (quality education) and SDG 10 (reduced inequalities), and underscores its potential to contribute to more inclusive and sustainable educational models.

From a technical standpoint, the system's success largely depends on the quality and coherence of the knowledge graph. Its construction requires demanding processes of curation, instructional design, and the application of natural language processing techniques. Additionally, the model currently relies on the exclusive use of GPT-3.5-turbo, which introduces risks of technological dependency and algorithmic bias, as well as limitations in transparency and pedagogical control.

In terms of future applicability, the PathRAG approach is scalable and adaptable to other educational levels (such as secondary or vocational training) and different disciplinary areas (humanities, social sciences, or health), provided there is an adequate documentary base and a teaching team trained in algorithmic literacy. It can also be applied in continuous training programs, multicultural contexts, and environments with high student diversity.

On an ethical level, the study was developed under the principles of the GDPR and teaching supervision, without automating evaluative decisions or compromising teacher autonomy. However, the move toward AI in education requires not only technical tools but also clear frameworks for governance, equity audits, and specialized, critical teacher training.

In conclusion, PathRAG should not be seen as a closed solution or a substitute for the teacher's role, but rather as a strategic support tool for personalized pedagogical design. Its value lies in increasing educational systems' responsiveness to diversity, facilitating tailored learning paths, and contributing to a more ethical, just, and resilient education model. When used critically and with supervision, generative artificial intelligence can become a key engine for democratizing access to knowledge and building more inclusive and sustainable educational ecosystems.

## REFERENCES

- Ainscow, M. (2020). *Promoting inclusion and equity in education: Lessons from international experiences*. UNESCO. <https://doi.org/10.1080/20020317.2020.1729587>
- Anderson, T., Borup, J., & West, R. E. (2020). Learning analytics in higher education: A review of the literature. *Journal of Research on Technology in Education*, 52(4), 433-451. <https://doi.org/10.1080/15391523.2020.1739846>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610-623). <https://doi.org/10.1145/3442188.3445922>
- Birhane, A. (2023). Algorithmic injustice: A relational ethics approach. *Patterns*, 4(2), 100693. <https://doi.org/10.1016/j.patter.2023.100693>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77-101. <https://doi.org/10.1191/1478088706qpo630a>
- Chen, B., Guo, Z., Yang, Z., Chen, Y., Chen, J., Liu, Z., Shi, C., & Yang, C. (2024). PathRAG: Pruning graph-based retrieval augmented generation with relational paths. *arXiv*. <https://doi.org/10.48550/arXiv.2502.14902>
- Chen, X., Wang, L., & Li, Z. (2024). Generative AI in higher education: Ethical and pedagogical considerations for ChatGPT. *AI & Society*. <https://doi.org/10.1007/s00146-024-01720-9>
- Cobo, C., Morales, A., & Navarro, C. (2022). Artificial intelligence in education: Equity and ethics in Latin America. *International Journal of Educational Technology in Higher Education*, 19, 32. <https://doi.org/10.1186/s41239-022-00312-3>
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Lawrence Erlbaum Associates.
- Conati, C., Holstein, K., & Roll, I. (2021). The future of intelligent tutoring systems: A survey. *IEEE Transactions on Learning Technologies*, 14(4), 464-484. <https://doi.org/10.1109/TLT.2021.3071740>
- Conati, C., & Kardan, S. (2013). User modeling in intelligent tutoring systems: A survey. In P. Brusilovsky, G. D. Verbert, & P. De Bra (Eds.), *Adaptive and adaptable learning* (pp. 45-63). Springer. [https://doi.org/10.1007/978-3-642-40866-0\\_3](https://doi.org/10.1007/978-3-642-40866-0_3)
- Floridi, L. (2019). *The ethics of artificial intelligence*. Oxford University Press.
- García-Peñalvo, F. J. (2022). *Tecnología y Educación: Hacia una visión transformadora*. Ediciones Universidad de Salamanca.
- Holmes, W., Porayska-Pomsta, K., Holstein, K., & Rodrigo, M. M. T. (2021). Ethics of AI in education: Towards a community-wide framework. *European Journal of Education*, 56(4), 541-558. <https://doi.org/10.1111/ejed.12481>
- Ifenthaler, D., & Yau, J. Y.-K. (2020). *Utilising learning analytics to support study success in higher education*. Springer. <https://doi.org/10.1007/978-3-319-64792-0>

- Jivet, I., Scheffel, M., Drachslar, H., & Specht, M. (2021). Awareness is not enough: Pitfalls of learning analytics dashboards in the educational practice. *British Journal of Educational Technology*, 52(6), 1864-1880. <https://doi.org/10.1111/bjet.13168>
- Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Liang, W., Tadesse, G. A., Ho, D., Li, F., Liu, S., Yu, H., & Sun, J. (2023). Personalized education in the era of big data. *IEEE Transactions on Knowledge and Data Engineering*, 35(3), 2245-2261. <https://doi.org/10.1109/TKDE.2021.3117424>
- Liu, Y., Tamkin, A., Geng, X., & Ganguli, D. (2023). Lost in the middle: How LLMs use long contexts. *Transactions of the Association for Computational Linguistics*, 11, 1200-1218. <https://doi.org/10.1162/tacl.a.00606>
- Lu, X., van de Sande, B., & Marsh, J. A. (2024). Hybrid learning environments and adaptive support. *Computers & Education*, 210, 104975. <https://doi.org/10.1016/j.compedu.2023.104975>
- Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. B. (2022). *Artificial intelligence and human learning: Partners in progress*. OECD Publishing.
- OECD. (2023). *The OECD AI principles: An overview*. OECD Publishing.
- Pedro, F., Subosa, M., Rivas, A., & Valverde, P. (2021). *Artificial intelligence in education: Challenges and opportunities for sustainable development*. UNESCO.
- Ruipérez-Valiente, J. A., Cobos, R., & Muñoz-Merino, P. J. (2022). Knowledge graphs in education: A systematic literature review. *IEEE Transactions on Learning Technologies*, 15(4), 485-500. <https://doi.org/10.1109/TLT.2022.3190858>
- Selwyn, N., Macgilchrist, F., & Williamson, B. (2023). EdTech and equity: Critical perspectives on the promises of digital education. *Learning, Media and Technology*, 48(1), 1-11. <https://doi.org/10.1080/17439884.2022.2156041>
- Susnjak, T. (2022). ChatGPT: The end of online exam integrity? *arXiv*. <https://doi.org/10.48550/arXiv.2212.09292>
- UNESCO. (2020). *Global education monitoring report 2020: Inclusion and education: All means all*. UNESCO. <https://doi.org/10.54676/JJNK6989>
- UNESCO. (2022). *Reimagining our futures together: A new social contract for education*. UNESCO.
- UNESCO. (2023). *AI and education: Guidance for policy-makers*. <https://unesdoc.unesco.org/ark:/48223/pf0000384803>
- van de Sande, B., Marsh, J. A., & Lu, X. (2023). Large language models in education: Opportunities and challenges. *Computers & Education: Artificial Intelligence*, 4, 100132. <https://doi.org/10.1016/j.caeai.2023.100132>
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education—Where are the educators? *International Journal of Educational Technology in Higher Education*, 16, 39. <https://doi.org/10.1186/s41239-019-0171-0>
- Zhang, D., & Wang, X. (2025). Adaptive learning systems: A review and future directions. *Journal of Educational Technology & Society*, 28(1), 1–15.
- Zhang, D., Wang, X., & Chen, X. (2021). Adaptive learning technologies in higher education: A meta-analysis of effectiveness. *Educational Research Review*, 33, 100390. <https://doi.org/10.1016/j.edurev.2021.100390>
- Zhou, M., & Wang, D. (2020). Ethical considerations in AI-supported education: A perspective from responsible

innovation. *AI & Society*, 35(3), 601-611.

<https://doi.org/10.1007/s00146-019-0910-1>

**Date of reception:** 1 June 2025

**Date of acceptance:** 11 August 2025

**Date of approval for layout:** 15 September 2025

**Date of publication in OnlineFirst:** 12 October 2025

**Date of publication:** 1 January 2026