

Leveraging Generative AI and Audio-Visual Cloning to Democratise Digital Learning

Matthew Montebello, Keith Azzopardi, Gabriel Borg, Karl Cini, Dylan Seychell, and Vanessa Camilleri

AI Department, Faculty of ICT, University of Malta, Malta

matthew.montebello@um.edu.mt

Abstract: This paper explores the transformative potential of Generative Artificial Intelligence (GenAI) and audio-visual cloning technologies in reshaping digital education, grounded in the ongoing project ACCLAIMED (Artificial Intelligence Content Cloning of Language-Agnostic Media for Education Democratisation). As education systems increasingly rely on digital platforms, the imperative to ensure accessibility, inclusion, and ethical integrity becomes more pronounced. ACCLAIMED introduces a novel triadic framework, namely, Course Generator, Guardrail, and AI-Renderer, to address e-learning challenges. The Course Generator collaborates with educators to produce comprehensive, pedagogically sound multilingual content. The Guardrail ensures human oversight, reinforcing societal norms, factual accuracy, and ethical alignment. The AI-Renderer transforms materials into realistic, human-like audio-visual formats, delivering engaging, culturally sensitive learning experiences. We discuss how ACCLAIMED advances the state-of-the-art by surpassing conventional AI tutors and adaptive learning systems through deeper pedagogical integration and ethical AI moderation. A key feature is its ability to deliver high-quality content across multiple languages, removing linguistic barriers and fostering educational equity especially for underserved or non-English-speaking populations. The paper also critically addresses broader issues: ethical concerns around AI-generated content, privacy and data protection (GDPR, EU AI Act), and digital sovereignty. Consideration is given to how such innovations can bridge or deepen the digital divide depending on their responsible and inclusive deployment. Ultimately, this paper calls for a reconceptualisation of digital learning, not merely as content delivery but as an inclusive, ethical, and adaptive ecosystem. It positions ACCLAIMED as a forward-looking blueprint for educational technologies prioritising innovation, equity, and societal impact.

Keywords: Generative AI, Audio-visual cloning, Digital learning, Ethical AI.

1. Introduction

Digital education is undergoing a profound transformation, driven by advances in artificial intelligence (AI), natural language processing, and multimedia generation (Arias Ortiz, et al., 2025). The integration of AI technologies in educational contexts has demonstrated significant potential to address longstanding challenges in learning delivery, with applications ranging from intelligent tutoring systems to automated content generation. Natural language processing technologies have emerged as particularly powerful tools for creating personalized learning experiences, enabling adaptive content delivery and multilingual support that caters to diverse learning styles and linguistic backgrounds (Tamilarasan, Selvaraj, & Buckingham, 2025). The proliferation of multimedia integration in e-learning environments has further enhanced educational accessibility and engagement. Modern educational platforms increasingly leverage interactive videos, simulations, and dynamic content to transform passive learning into active participation, addressing various learning modalities and improving knowledge retention. These technological advances align with contemporary pedagogical frameworks that emphasize learner-centered approaches and inclusive educational practices (MOEE, 2012). Despite the proliferation of e-learning platforms and technological innovations, significant barriers remain that prevent equitable access to quality digital education (Vishnu, Tengli, Ramadas, Sathyan, & Bhatt, 2024). Linguistic exclusion represents one of the most pervasive challenges, with the digital language divide creating disparities between dominant languages and low-resourced languages in terms of digital content availability and technological support (Okolo & Tano, 2024). This disparity particularly affects speakers of indigenous and minority languages, who often find themselves excluded from the benefits of digital technologies and educational resources. Cultural adaptation deficits further compound accessibility challenges in digital learning environments (Zhao, Wang, Li, Zhou, & Li, 2021). Many mainstream e-learning platforms fail to adequately address the unique learning styles, preferences, and cultural sensitivities of diverse populations, leading to reduced effectiveness and engagement among learners from different cultural backgrounds (Viberg, Jivet, & Scheffel, 2022). The lack of culturally responsive educational content not only diminishes learning outcomes but also perpetuates educational inequities by failing to recognize and validate diverse cultural knowledge systems (Milheim, 2015). Ethical concerns surrounding the deployment of AI-generated content in educational settings have emerged as another critical barrier (Nguyen, 2025). The complexity of AI systems often makes it difficult to understand decision-making processes, leading to concerns about transparency, accountability, and potential bias in educational applications (EC, 2025). These ethical considerations are particularly significant given the formative nature of educational experiences and their long-term impact on learners' development and

opportunities (EC, 2025). These challenges are particularly acute for underserved, non-English-speaking, or resource-constrained populations, exacerbating the global digital divide (Okolo & Tano, 2024). Rural communities, low-income families, and historically marginalized groups face compounded disadvantages in accessing digital education, including limited internet connectivity, inadequate technological infrastructure, and insufficient digital literacy skills (Raihan, et al., 2025). The digital divide not only affects immediate educational outcomes but also perpetuates existing inequalities by limiting access to opportunities for skill development and career advancement (Bentley, Naughtin, & McGrath, 2024). Research indicates that students from historically and systematically excluded backgrounds, including English language learners and students with disabilities, often rely more heavily on basic assistive tools rather than advanced educational technologies, further widening the achievement gap (Assefa, Gebremeskel, Moges, Tilwani, & Azmera, 2025). The pandemic has particularly highlighted these disparities, as only educational systems with robust digital capabilities were able to maintain continuity of learning, leaving millions of students behind (Sun, Tang, & Zhao, 2024).

The rest of the paper is organised as follows. The next section will go through any related work that encompasses the different facets of this study, while Section 3 introduces ACCLAIMED (Artificial Intelligence Content Cloning of Language-Agnostic Media for Education Democratisation), a novel framework leveraging generative AI and audio-visual cloning to democratise digital learning. The next section goes through the methodology we adopted during the ACCLAIMED project followed by the early results and findings in our initial implementations, in the following section. A conclusions and future work section completes our paper as we close with numerous recommendations and best practices especially in relation to important issues like ethical concerns, content, privacy and data protection.

2. Related Work

2.1 Generative AI in Education

Generative AI has rapidly emerged as a transformative force in education, demonstrating significant promise in areas such as automated content creation, assessment, and personalised tutoring (Qian, 2025). The advent of large language models (LLMs), including GPT-4 and Google Gemini, has enabled the rapid and scalable generation of high-quality textual learning materials, question banks, and adaptive feedback systems (Fenta, 2025). These models are capable of producing contextually relevant, curriculum-aligned content that can be tailored to individual learner needs, thereby supporting differentiated instruction and fostering learner engagement (Yunjo, Hyun, & Shadarra, 2025). Beyond text, advances in generative models, such as diffusion models and neural rendering, have facilitated the creation of synthetic media, including images, animations, and even interactive simulations (Kumar, 2025). These innovations enable the development of rich, multimodal educational experiences that transcend the limitations of traditional textbooks and static e-learning modules. However, despite these advances, most generative AI systems remain predominantly focused on English-language content and Western-centric educational paradigms (Agarwal, Naaman, & Vashistha, 2025). There is a notable lack of mechanisms for deep pedagogical integration, cultural adaptation, or robust ethical oversight. As a result, existing solutions often fail to address the needs of diverse learner populations or to ensure alignment with local curricular standards and values (Nyaaba, Wright, & Choi, 2024).

2.2 Audio-Visual Cloning

Audio-visual cloning technologies represent another frontier in educational innovation, enabling the synthesis of realistic human-like avatars, voices, and gestures through AI-driven techniques (Ingle et al., 2024). Platforms such as Synthesia, HyGen and D-ID have demonstrated the technical feasibility of generating video lectures and virtual instructors that closely mimic human appearance and delivery. These tools hold significant potential for scaling high-quality instruction and making learning more engaging and accessible, particularly in remote and asynchronous settings. Nevertheless, the adoption of audio-visual cloning in education is not without challenges (Chheang, et al., 2024). Many current implementations lack sufficient alignment with established educational standards or fail to incorporate cultural and contextual nuances critical for effective pedagogy (Baker, Smith, & Anissa, 2022). The risk of producing content that is culturally insensitive or pedagogically inappropriate underscores the need for frameworks that integrate educational expertise and local context into the design and deployment of AI-driven media.

2.3 Ethical and Legal Considerations

The deployment of AI technologies in educational contexts raises a host of ethical and legal concerns, including issues of algorithmic bias, misinformation, learner privacy, and regulatory compliance (Williamson & Eynon, 2020). AI systems, if not properly monitored, may inadvertently perpetuate or even amplify existing biases, leading to unfair or discriminatory outcomes for certain groups of learners (Kotek, Dockum, & Sun, 2023). The

risk of misinformation where AI-generated content is factually incorrect or misleading poses further challenges, particularly in high-stakes educational environments (Locke & Hodgdon, 2025). Privacy and data protection are paramount, especially given the sensitive nature of educational data and the increasing regulatory scrutiny exemplified by frameworks such as the General Data Protection Regulation (GDPR) and the EU AI Act. Many current AI-powered educational platforms operate with limited transparency, offering few mechanisms for human-in-the-loop oversight or for learners and educators to understand, challenge, or intervene in AI-driven decisions (Liao & Wortman-Vaughan, 2023). This lack of transparent guardrails not only undermines trust but also raises significant questions about accountability and the responsible use of AI in shaping educational futures.

3. The ACCLAIMED Framework

3.1 Overview

ACCLAIMED is structured around three core modules (see Figure 1):

1. Course Generator: Collaborates with educators to produce pedagogically robust content, assisting in both planning and content creation.
2. AI-Renderer: Translates textual content and transforms it into realistic audio-visual formats, adapting to cultural and linguistic contexts.
3. Guardrail: Flags concerning content for human review, ensuring factual accuracy, ethical alignment, and regulatory compliance.

3.2 Course Generator

The Course Generator component within the ACCLAIMED framework facilitates a structured, educator-guided authoring process that blends human pedagogical insight with AI-assisted automation to generate high-quality, scalable educational content. The process begins with educators uploading their instructional resources into a centralised database, from which the AI agent can retrieve relevant materials on demand using Retrieval-Augmented Generation (RAG) methods (Lee, 2025). This initial step ensures that content generation remains grounded in expert-sourced materials. Subsequently, educators engage in a co-creative dialogue with the AI to formulate a syllabus, delineating the key topics to be addressed throughout the course. Following this, educators define essential course parameters (such as intended duration, lesson length and frequency, and target audience) which the AI uses to construct a scheme of work. This scheme outlines the structure of lessons and assessments, and is subject to iterative review by the educator to ensure that it lines up with their individual style and goals. Each lesson is then individually developed through a collaborative process wherein the AI generates a detailed lesson plan based on Gagné's Nine Events of Instruction, which is validated and adjusted through feedback from both the educator and the Guardrail component described in Section 3.4. The educator may select which activities to include based on their personal style and intended audience, as recommended in Fan et al., (2024). Building on this foundation, the lesson script is automatically generated by the Lesson Generator and undergoes further scrutiny and revision by both human and AI-driven checks. The finalised script is passed to the AI Renderer (Section 3.3), which produces a video lecture that simulates human delivery in the target selection of languages, allowing the educator to review and approve the generated media. In parallel, instructional slides are automatically produced from the lesson plan's defined sections. These slides are reviewed by the Guardrail component and integrated into Google Slides via the API, enabling further manual refinement by educators within a familiar interface. The Assessment Generator creates multiple-choice quizzes to evaluate student understanding in a formative manner, exportable to platforms such as Kahoot and Google Forms, granting educators flexibility in tailoring and deploying formative assessments.




Figure 1: ACCLAIMED Architecture

3.3 AI Renderer

The AI Renderer receives the content generated and validated by the educator in their native language, and translates it into the chosen languages, before making the necessary calls to third-party services to create engaging content that is accessible by diverse learner populations. Multilingual support is achieved via translation models and cross-lingual embeddings, ensuring content is accessible to diverse learner populations. The translated content is first checked by the guardrail, and then makes use of the following third-party services to generate the deliverables as shown in Figure 1.

3.3.1 Video Avatar Generation Platform

The approved and translated script is used to generate a “talking head” video of the educator delivering the content. The output is a video lecture reproduced in the desired language set, that can be uploaded to the Learning Management System (LMS) of choice. Heygen was selected as the platform of choice for the proof of concept of ACCLAIMED as it seemed to offer the highest quality content at competitive prices upon an initial evaluation. A comprehensive comparison of providers and open source implementations will be discussed later in Section 3.5 as part of planned future updates.

3.3.2 Presentation Creation Platform

The lecture slides designed by the Course Generator are uploaded to a presentation creation platform, in this case Google Slides. This was chosen as it is widely adopted and easily accessible since a dedicated license is not required to use the platform. The integration is built via APIs made available by Google, which enables users to pass in the ID of a template presentation to use for styling. This was considered a desirable feature as it allows

educators to align the generated content with their existing course material, enabling seamless incorporation of the new content into their workflows. This feature will also allow institutions to standardise content styling across courses and educators if they desire, and any changes made to the underlying template will be propagated to the generated material automatically.

3.3.3 Quiz Creation Platform (Kahoot/Google forms)

The multiple-choice quizzes created by the Course Generator are similarly uploaded to a quiz creation platform that allows the educator to make any final modifications to the content, as well as distribute the quiz to students. To demonstrate the extensibility of the ACCLAIMED platform, two providers are integrated for educators to choose from, Kahoot and Google Forms. The gamified interface of Kahoot has proven to be effective at increasing engagement in various domains (Wang & Tahir, 2020), and allows the educator to carry out the quiz in class or distribute as an assignment. On the other hand, Google Forms might be a favourable integration for academics who might consider Kahoot as too informal for a higher education institution, or are more familiar with the Google Workspace ecosystem. Both platforms offer extensive reporting capabilities, allowing educators to assess student understanding of the concepts covered in class. By reducing the authoring burden, the ACCLAIMED platform allows educators in resource-constrained environments to reap the benefits of formative assessment with ease. This can provide near-real-time feedback to educators while the course is in progress, allowing them to adapt their teaching practices and plans accordingly. As discussed later on, we intend to look into tighter integrations with these platforms to pull in performance analytics and incorporate it directly in the content authoring loop, providing recommendations when designing/updating the course. The modular design of the AI Renderer is easily extensible to make use of any other provider for each of these use cases. This was considered a key feature when architecting the ACCLAIMED framework, so that the project can be adapted to the suite of tools in use at a particular academic institution. This will circumvent the issue of having to purchase new licenses and re-training staff. The project leans heavily into third-party services since educators are already familiar with the content authoring experience. This allows educators to fine tune the design and/or content using the comprehensive set of tools on the platforms. This combines the efficiency of AI content generation with the flexibility of proven tools.

3.4 Guardrail

The Course Generator communicates with the Guardrail component to run various checks in the content generation loop. Each of the checks listed in this section outputs a “probability of violation” and a validation error message, which is surfaced through the UI and requires educator validation when the probability exceeds a certain threshold. The checks are implemented using a mixture of existing open source platforms, namely Guardrails AI and Nvidia NeMo, as well as some custom implementations. The array of checks included:

- **Content Accuracy Check:** This check ensures that factual statements found in the text are supported by at least one authoritative source. Preference is given to sources from the approved list uploaded by the educator. However, these are also cross-checked against online references for consistency. This ensures that any claims made (for example, about the current state-of-the-art technique in a particular domain) are up to date, even if the content uploaded by the educator is out of date. Conflicts are escalated to the educator to resolve before proceeding. Combined with RAG, this check aims at minimising hallucinations, a known problem in LLMs (Huang, et al., 2025). Spell checks, grammar checks, and inconsistency checks are also run against the content to ensure linguistic coherence.
- **Pedagogical Alignment Check:** The objectives stated in the lesson plans are cross-checked against the syllabus, to ensure that there is adequate mapping between the generated content and the overall course goals. A separate check ensures that the objectives, discussions and activities are suitable for the course audience as specified by the educator in the initial course planning stages with the Course Generator. This aspect of the guardrail will evolve further.
- **Ethical and Legal Compliance Check:** All text output by the various components of the Course Generator, as well as any modifications made by the educator are checked to ensure that the content does not promote hate, harassment, extremism, political persuasion or self-harm. At the course design stage, the educator is asked whether they would like any medical or legal advice surfaced in the content to be blocked or flagged for verification. While it is expected that most educators will want to outright block such content, this opens the door for manual review in courses where discussions about financial and legal decisions is a core component.
- **Bias and Inclusion Check:** All content is checked to ensure that there are no gender, racial, cultural or ability stereotypes. This protects the educator from inclusion of micro-aggressions or exclusionary phrasing. A reading level check is also run to ensure that the content can be understood by the target

audience specified at the course level. Content that mentions violence, sexual intercourse or substance use is surfaced for educator review. The content is not blocked to allow for cases where a mature discussion of these subjects is required to meet the course objectives.

- **GDPR Compliance Check:** All content output by the various components of the Course Generator are checked to ensure that there is no leakage of Personally Identifiable Information (PII). While the platform does not directly deal with student data, this is added as a fail-safe against situations where the educator might (accidentally) upload material that includes student information such as grades or attendance records.
- **Multilingual Consistency Check:** Once the content has been translated by the AI renderer, a series of linguistic checks are run to ensure that the tone is consistent across all languages, and that there are no grammatical or spelling errors.
- **Intellectual Property Check:** The generated content is reviewed to ensure that adequate references are included where portions of source material is reproduced verbatim, and to ensure that all document snippets retrieved via RAG have been cited.

4. Results and Discussion

The development of ACCLAIMED has, to date, focused on rigorous internal testing and iterative proof-of-concept trials, providing a robust foundation for future large-scale deployment. This phase has been instrumental in evaluating the technical feasibility, pedagogical soundness, and practical usability of the ACCLAIMED framework. Initial results have shown that the technical performance and system stability of the beta version are very promising and encouraging. Multiple rounds of system trials were conducted to assess the performance of the Course Generator, Guardrail, and AI-Renderer modules. These trials involved generating multilingual course materials across a range of subjects and educational levels. The Course Generator demonstrated strong capabilities in producing coherent, curriculum-aligned content in English, Spanish, French, Italian and Maltese, with ongoing fine-tuning to improve output in lower-resourced languages like Maltese itself. The AI-Renderer successfully synthesized human-like audio-visual content, with avatars accurately reflecting linguistic and cultural nuances as validated by internal reviewers. System stability was evaluated through stress-testing and error-logging protocols. The framework maintained high uptime and responsiveness, with latency for content generation and rendering consistently within acceptable thresholds for educational applications. These results suggest that ACCLAIMED's modular architecture is technically robust and scalable, supporting future integration with external learning management systems. To ensure pedagogical validity, subject-matter experts and instructional designers participated in structured walkthroughs of the generated content. Feedback highlighted the system's strengths in producing logically structured lessons, appropriate scaffolding, and adaptive feedback mechanisms. Notably, the Guardrail module's human-in-the-loop oversight effectively identified and corrected factual inaccuracies, cultural mismatches, and inappropriate content, demonstrating the practical value of embedding ethical review within the workflow (Montebello, 2021). Qualitative feedback from internal testers indicated that the audio-visual materials enhanced learner engagement and comprehension, particularly for complex or abstract topics. These findings are consistent with broader research on AI-driven educational tools, which have shown improvements in student engagement and learning outcomes when adaptive and multimedia elements are incorporated (Mallia-Milanes & Montebello, 2021). Initial user experience trials involved the project team that is made up of numerous educators and learners that interacted with the prototype. The majority of participants found the system intuitive and the workflow for content co-creation straightforward. Educators appreciated the ability to review, edit, and approve AI-generated materials, reinforcing their role as pedagogical authorities. Learners responded positively to the diversity of avatars and languages, noting increased relatability and motivation. Iterative refinements continue to refine ACCLAIMED as they were based on user feedback that led to enhancements in navigation, customisation options, and accessibility features. These adjustments mirror findings from other AI education pilots, where continuous user-centered design has been critical for adoption and efficacy (Heung & Yim, 2024). While this initial internal testing phase yielded promising results, several limitations were identified. Language coverage provides particular challenges as further work is needed to improve accuracy and naturalness in underrepresented languages like Maltese. Cultural adaptation also has its challenges as ensuring deep cultural relevance remains a complex challenge, particularly for content intended for diverse global audiences. Finally, automated moderation through the Guardrail module presented some edge cases indicating it still requires manual intervention, underscoring the need for ongoing human oversight. These findings align with other recent proof-of-concept initiatives in educational AI, which have similarly highlighted the importance of iterative development, user feedback, and ethical safeguards in achieving practical and scalable solutions (DoE, 2024). The ACCLAIMED proof-of-concept demonstrates that generative AI and audio-visual cloning can be harnessed to produce high-quality, inclusive,

and ethically sound educational content. The modular, educator-centered design supports both scalability and adaptability, making ACCLAIMED a compelling candidate for broader pilot studies and eventual real-world deployment. These initial results reinforce the transformative potential of AI in education, particularly for addressing linguistic and cultural barriers and supporting underserved populations. However, the experience also underscores the necessity of embedding robust ethical, legal, and pedagogical oversight at every stage of development and deployment. Building on the success of internal trials, the next phase will involve pilot deployments in partnership with educational institutions, with a focus on diverse linguistic and cultural settings. Quantitative and qualitative data will be systematically collected to evaluate learner outcomes, educator workload, and system impact. Further research will also address the scalability of the Guardrail module and the ongoing refinement of language and cultural adaptation capabilities.

5. Conclusion and Future Work

ACCLAIMED represents a significant step forward in leveraging generative AI and audio-visual cloning for equitable, inclusive, and ethical digital learning. By foregrounding multilingualism, cultural adaptation, and robust ethical safeguards, it offers a scalable blueprint for the future of education technology. However, ongoing vigilance is required to ensure these technologies serve all learners, uphold privacy, and foster digital sovereignty. As the ACCLAIMED platform matures, several key areas have been identified to guide future research and development. These avenues aim to strengthen the platform's pedagogical robustness, enhance system autonomy and flexibility, and extend its adoption across varied educational contexts. Future work will look into integration with Learning Management Systems (LMS) to manage not only the distribution of generated content, but also to collect learner analytics and feedback. This integration would enable the coupling of ACCLAIMED with student modelling frameworks such as Knowledge Tracing and Cognitive Diagnosis Modelling, thereby creating an adaptive learning environment. Such an environment would support the dynamic adjustment of instructional content in response to individual learner needs, thus realising the project's ambition of personalisation at scale to improve accessibility. To further close the adaptive feedback loop, the alignment of assessment mechanisms with Bloom's Taxonomy will be deepened. Prior work has demonstrated the effectiveness of Bloom-aligned assessments in differentiating cognitive skill levels and informing instructional design. By enhancing the taxonomy mapping in ACCLAIMED's assessment generation tools, a more granular understanding of learner capabilities can be derived and used to guide the generation of tailored content that meets learners at their demonstrated level of mastery. To accommodate a broader range of educational and training contexts, especially those outside traditional academic environments, future iterations of the platform will incorporate alternative assessment techniques. This includes integrations with branching scenario platforms and virtual reality (VR) simulation environments to support Scenario-Based Learning (SBL). Such modalities are particularly relevant for vocational training, where experiential and performance-based evaluation methods are more pedagogically appropriate than conventional quiz-based assessments. The platform's export functionalities will be expanded to include Microsoft Forms and Microsoft PowerPoint, ensuring compatibility with institutional ecosystems that rely on Microsoft 365 infrastructure. This will enhance usability for educators in institutions where Office licenses are predominant, thereby lowering the barrier to adoption. To support institutional-level deployments, ACCLAIMED will be extended with capabilities for collaborative course creation. This will allow educators, including teaching assistants and co-instructors, to jointly author and refine course materials. The ability to support multi-author workflows is essential in contexts where course delivery is a distributed effort across teams, ensuring consistency, version control, and efficient division of labour in large-scale educational programs. Future versions could also allow for the sharing of educational content between courses, allowing educators to pull from their institutions content library and re-use to fit their needs. The current reliance on third-party AI rendering services presents limitations in terms of flexibility, cost, and control. A proposed future direction involves implementing the AI Renderer using recent open-source techniques for talking head generation. This would not only eliminate third-party dependencies but also allow fine-grained access control and local deployment. A qualitative study will be undertaken to compare the output quality of open-source renderers with that of proprietary solutions, incorporating metrics such as realism, engagement, and educator satisfaction. In parallel, the content generation pipeline will undergo two major comparative evaluations. First, a comprehensive benchmarking of Guardrail libraries will be conducted to assess their effectiveness in enforcing ethical, cultural, and pedagogical constraints. Second, the performance of various large language models (LLMs) will be compared in their ability to generate educational content. This comparison will be executed as a controlled study involving both educators and students, who will be asked to rate and rank content produced by different LLMs and by expert human authors. These evaluations will provide empirical grounding for model selection and configuration in future platform release. A series of pilot studies will be conducted in collaboration with educational institutions to assess both the technical and pedagogical

performance of each platform component, as well as the system as a whole. In addition to performance metrics, a user experience (UX) study will be carried out to evaluate the intuitiveness and usability of the educator-facing interface. Complementing these efforts, rigorous empirical studies will quantify the reduction in content authoring time achieved through ACCLAIMED. By measuring time savings in real-world instructional design tasks, the platform's value proposition for educators can be substantiated with evidence-based metrics.

Ethics declaration

No ethical clearance was needed for the research referred to in this paper.

AI declaration

A number of Generative engines were used responsibly to assist with language refinement, structure, and editing during the preparation of this paper. All content was critically reviewed, verified, and edited by the authors who employed Gemini, Co-Pilot, ChatGPT and Perplexity to ensure accuracy, originality, and strict compliance with academic integrity standards.

References

- Agarwal, D., Naaman, M., & Vashistha, A. (2025). AI Suggestions Homogenize Writing Toward Western Styles and Diminish Cultural Nuances. 2025 CHI Conference on Human Factors in Computing Systems (p. <http://dx.doi.org/10.1145/3706598.3713564>). ACM.
- Arias Ortiz, E., Castro Vergara, N., Forero Pabón, T., Della Nina Gambi, G., Giambruno, C., Pérez Alfaro, M., & Rodríguez Segura, D. (2025). AI and Education: Building the Future Through Digital Transformation. <https://doi.org/10.18235/0013500>.
- Assefa, Y., Gebremeskel, M., Moges, B., Tilwani, S., & Azmera, Y. (2025). Rethinking the digital divide and associated educational in(equity) in higher education in the context of developing countries: the social justice perspective. *International Journal of Information and Learning Technology*, 42(1), 15-32.
- Baker, T., Smith, L., & Anissa, N. (2022). The Ethics of Synthetic Media in Education: Cultural and Pedagogical Implications. *Journal of Educational Technology & Society*, 25(3), 45-59.
- Bentley, S., Naughtin, C., & McGrath, M. (2024). The digital divide in action: how experiences of digital technology shape future relationships with artificial intelligence. *AI Ethics*(<https://doi.org/10.1007/s43681-024-00452-3>), 901–915.
- Chheang, V., Sharmin, S., Marquez-Hernandez, R., Patel, M., Rajasekaran, D., Caulfield, G., . . . Barmaki, R. (2024). Towards Anatomy Education with Generative AI-based Virtual Assistants in Immersive Virtual Reality Environments. *Human-Computer Interaction* (p. <https://arxiv.org/abs/2306.17278>). Arxiv.
- DoE. (2024). Use Cases for Generative AI in Education. UK Department of Education.
- EC, E. C. (2025). Living Guidelines on the Responsible Use of Generative AI for Research. Brussels: EC.
- Fan, H., Chen, G., Wang, X., & Peng, Z. (2024). LessonPlanner: Assisting Novice Teachers to Prepare Pedagogy-Driven Lesson Plans with Large Language Models. 37th Annual ACM Symposium on User Interface Software and Technology, (pp. 1-20).
- Fenta, A. (2025). A review on enhancing education with AI: exploring the potential of ChatGPT, Bard, and generative AI. *Discover Education*, <https://doi.org/10.1007/s44217-025-00426-5>.
- Haniya, S., Montebello, M., Cope, W., & Tapping, R. (2018). Promoting Critical Clinical Thinking through e-Learning. 10th International Conference on Education and New Learning Technologies. Palma De Mallorca, Spain: EduLearn.
- Haniya, S., Tzirides, O., Montebello, M., Georgiadou, K., Cope, B., & Kalantzis, M. (2019). Maximizing learning potential with multimodality: A Case study. *World Journal of Educational Research*, 6(2).
- Heung, I., & Yim, Y. (2024). A critical review of teaching and learning artificial intelligence (AI) literacy: Developing an intelligence-based AI literacy framework for primary school education. *Computers and Education: Artificial Intelligence*, <https://doi.org/10.1016/j.caeai.2024.100319>.
- Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., . . . Liu, T. (2025). A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Transactions on Information Systems*, 43(2), 42:1-42:55.
- Ingle, A., Shaikh, M., Shelke, A., Thombre, B., & Wanjale, K. (2024). Generative AI Pipeline for Modeling Digital Humans for Interactive Learning. In: Asokan, R., Ruiz, D.P., Piramuthu, S. (eds) *Smart Data Intelligence. ICSMDI 2024. Algorithms for Intelli*. In R. Asokan, D. Ruiz, & S. (. Piramuthu, *Smart Data Intelligence. ICSMDI 2024*. (https://doi.org/10.1007/978-981-97-3191-6_33). Singapore: Algorithms for Intelligent Systems. Springer.
- Kotek, H., Dockum, R., & Sun, D. (2023). Gender bias and stereotypes in Large Language Models. *Computation and Language*, 10.1145/3582269.3615599.
- Kumar, P. (2025). Diffusion Models and Generative Artificial Intelligence: Frameworks, Applications and Challenges. *Archives of Computational Methods in Engineering*, <https://doi.org/10.1007/s11831-025-10266-z>.
- Lee, Y. (2025). Developing a computer-based tutor utilizing Generative Artificial Intelligence (GAI) and Retrieval-Augmented Generation (RAG). *Youngjin Lee in Education and Information Technologies*, 30(<https://doi.org/10.1007/s10639-024-13129-5>), 7841–7862.
- Liao, Q., & Wortman-Vaughan, J. (2023). AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap. *Human-Computer Interaction*, <https://arxiv.org/abs/2306.01941>.

- Locke, L., & Hodgdon, G. (2025). Gender bias in visual generative artificial intelligence systems and the socialization of AI. *AI & Society*, 40(<https://doi.org/10.1007/s00146-024-02129-1>), 2229–2236.
- Mallia-Milanes, M., & Montebello, M. (2021). Enhancing Educators Credibility in Intelligent Learning Environments. *The International Conference on Education and New Developments* (p. (in print)). Virtual: END.
- Milheim, K. L. (2015). Cultural Inclusivity in Online Learning. In R. W. (Ed.), *Student-Teacher Interaction in Online Learning Environments* (pp. 76-88). <https://doi.org/10.4018/978-1-4666-6461-6.ch004> : IGI Global.
- MOEE, M. o. (2012). *A National Curriculum Framework for All*. Malta: Salesian Press.
- Montebello, M. (2021). Smart Education through AI. 2nd International Conference of the Journal Scuola Democratica “Reinventing Education”, (pp. 161-171). VOL. 2, Learning with New Technologies, Equality and Inclusion.
- Nguyen, K. (2025). The Use of Generative AI Tools in Higher Education: Ethical and Pedagogical Principles. *Journal of Academic Ethics*, <https://doi.org/10.1007/s10805-025-09607-1>.
- Nyaaba, M., Wright, A., & Choi, G. (2024). Generative AI and Digital Neocolonialism in Global Education: Towards an Equitable Framework. *Computers and Society*, <https://arxiv.org/abs/2406.02966>.
- Okolo, C., & Tano, M. (2024, December 12). Closing the gap: A call for more inclusive language technologies. From <https://www.brookings.edu/articles/closing-the-gap-a-call-for-more-inclusive-language-technologies/>
- Qian, Y. (2025). Pedagogical Applications of Generative AI in Higher Education: A Systematic Review of the Field. *TechTrends*, <https://doi.org/10.1007/s11528-025-01100-1>.
- Raihan, M., Subroto, S., Chowdhury, N., Koch, K., Ruttan, E., & Turin, T. (2025). Dimensions and barriers for digital (in)equity and digital divide: a systematic integrative review. *Digital Transformation and Society*, 4(2), 111-127.
- Sun, Y., Tang, C., & Zhao, Z. (2024). Does online education magnify educational inequalities? Evidence from the COVID-19 pandemic. *China Economic Review*, 88.
- Tamilarasan, P., Selvaraj, V., & Buckingham, L. (2025). Impact of Deep Learning for Multilingual Natural Language Processing in Educational Applications. *Proceedings of International Conference on Recent Trends in Computing* (https://doi.org/10.1007/978-981-97-8836-1_33). ICRTC 2024. *Lecture Notes in Networks and Systems*, vol 885. Springer, Singapore.
- Viberg, O., Jivet, I., & Scheffel, M. (2022, December 19). Designing Culturally Aware Learning Analytics: A Value Sensitive Perspective. From Arxiv.org: <https://arxiv.org/abs/2212.09645>
- Vishnu, S., Tengli, M., Ramadas, S., Sathyan, A., & Bhatt, A. (2024). Bridging the Divide: Assessing Digital Infrastructure for Higher Education Online Learning. *TechTrends*, [https://doi.org/10.1007/s11528-024-00997-4\(68\)](https://doi.org/10.1007/s11528-024-00997-4(68)), pp. 1107–1116.
- Wang, A., & Tahir, R. (2020). The effect of using Kahoot! for learning – A literature review. *Computers & Education*, 149, <https://doi.org/10.1016/j.compedu.2020.103818>.
- Williamson, B., & Eynon, R. (2020). Historical threads, missing links, and future directions in AI in education. *Learning, Media and Technology*, 45(3), 223–235.
- Yunjo, A., Hyun, Y., & Shadarra, J. (2025). Investigating the higher education institutions’ guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. *International Journal of Educational Technology in Higher Education*, 22(1).
- Zhao, Y., Wang, N., Li, Y., Zhou, R., & Li, S. (2021). Do cultural differences affect users’ e-learning adoption? A meta-analysis. *British Journal of Educational Technology*, 52(1)(<https://doi.org/10.1111/bjet.13002>), 20–41.