An Accessibility-First Generative AI Chatbot to Support Inclusive MOOC Discovery and Early Onboarding

OPEN PRAXIS

RESEARCH ARTICLES

INTERNATIONAL COUNCIL FOR OPEN AND DISTANCE EDUCATION

SALWA MRAYHI
MOHAMED KOUTHEAIR KHRIBI
MOHAMED JEMNI

*Author affiliations can be found in the back matter of this article

ABSTRACT

Massive Open Online Courses (MOOCs) have expanded access to education globally but remain largely inaccessible to learners with disabilities and those facing multilingual barriers. Persistent issues such as insufficient captioning, poor screen reader compatibility, and limited adaptive support constrain equity in participation. This study introduces RECMOOC4All Chatbot, the user-facing interaction interface of the RECMOOC4All MOOC aggregator, designed to reduce friction during course discovery and early onboarding by embedding accessibility at the interaction layer through multimodal (text/voice) support, multilingual mediation, and accessibility-aware prompting based on available course signals. The prototype integrates speech-to-text (STT), optional text-to-speech (TTS), ensemble language identification, neural machine translation, and intent-aware routing across discovery, assistance, and feedback workflows. An exploratory live pilot with 15 MOOC learners (including six reporting disability-related needs) assessed feasibility and early utility for search/refinement, clarification, first-step guidance, and feedback. Results indicate reliable runtime behavior for the targeted tasks: intent routing agreement reached 92.1%, STT achieved a 9.2%-word error rate on a controlled speech set, and language identification precision was 95.82% on a multilingual benchmark; translation fidelity for domain queries was supported by semantic-similarity evaluation (F1 = 0.8962). An offline retrieval comparison showed a 38% precision gain for accessibility-tagged course queries relative to a TF-IDF baseline. Participants reported high perceived usefulness and clarity, and interaction telemetry suggested reduced friction during early discovery across diverse profiles. Findings should be interpreted as feasibility evidence from a small, heterogeneous pilot without a control condition; broader effectiveness and long-term outcomes require larger controlled studies and depend on the completeness of underlying course metadata.

CORRESPONDING AUTHOR:

Salwa Mrayhi

LaTICE Research Laboratory, National Higher Engineering School of Tunis (ENSIT), University of Tunis, Tunisia

salwa.mrayhi@ieee.org

KEYWORDS:

MOOC accessibility; inclusive open education; conversational artificial intelligence; generative ai in education; multimodal interaction; multilingual mediation; universal design for learning (udl); accessibility-first design; MOOC discovery and onboarding; open and distance learning

TO CITE THIS ARTICLE:

Mrayhi, S., Khribi, M. K., & Jemni, M. (2026). An Accessibility-First Generative AI Chatbot to Support Inclusive MOOC Discovery and Early Onboarding. *Open Praxis*, *18*(3), pp. 441–461. DOI: https://doi.org/10.55982/openpraxis.18.3.1026

INTRODUCTION

Massive Open Online Courses (MOOCs) have expanded educational participation by reducing geographic, temporal, and financial constraints (Kilgore & Al-Freih, 2017). Yet access remains uneven for learners with disabilities and for learners who rely on multilingual support, particularly during course discovery, onboarding, and early navigation, where interaction barriers can prevent meaningful participation before learning activities begin. The World Health Organization estimates that more than 1.3 billion people, approximately 16% of the global population, live with disabilities (WHO, 2021), reinforcing the urgency of inclusive digital learning aligned with international rights-based commitments (e.g., UNCRPD, Article 9) and widely adopted web accessibility guidance (e.g., WCAG 2.2 as applied to web content and interfaces). Empirical studies continue to document persistent MOOC accessibility gaps, including limited assistive technology compatibility, inconsistent accessibility metadata, and insufficient support for diverse sensory, cognitive, and linguistic needs (Farrow et al., 2024; Mrayhi et al., 2024). In practice, video-dominant instructional designs, inconsistent captioning, missing alternative text, and screen-reader friction remain common sources of exclusion (Hamideh Kerdar et al., 2024).

Beyond technical accessibility, many MOOC environments provide static content delivery and generic help functions with limited personalization or adaptive support, making it difficult to operationalize Universal Design for Learning (UDL) in ways that are transparent and evaluable. UDL requires multiple means of engagement, representation, and action/expression, supported by explicit design logic that links learner needs to concrete interaction supports (e.g., scaffolds, alternative modalities, stepwise guidance, and self-regulation supports) (Meyer et al., 2014). Where such logic is absent or implicit, learners who require alternative engagement routes or cognitive supports may experience higher cognitive load and reduced autonomy, which are associated with disengagement and dropout risk in online learning contexts (Wang et al., 2023).

Recent advances in Artificial Intelligence (AI), particularly Generative AI (GenAI), have increased interest in conversational interfaces that can provide real-time assistance, multimodal interaction, and multilingual mediation (Guettala et al., 2024; Holmes et al., 2022; Khong et al., 2024). However, current educational chatbot deployments (e.g., Khanmigo, Coursera AI, and earlier agent-based teaching assistants such as Jill Watson) are typically framed around tutoring, personalization, or engagement, while accessibility is often treated as peripheral or inconsistently integrated across courses and platforms (Caelen & Blete, 2024; Goel & Polepeddi, 2018; Rao & Lin, 2024). As a result, the evidence base remains limited on accessibility-first GenAI chatbots evaluated as user-facing interfaces for MOOC discovery and onboarding, with transparent reporting of their design logic, ethical safeguards, and operational constraints.

This study addresses this gap by presenting the RECMOOC4All Chatbot, a GenAI-powered conversational system designed with accessibility as a primary design constraint. The chatbot integrates GPT-4o-mini for intent routing and response planning, Whisper for speech-to-text, and OPUS-MT for multilingual translation, and is orchestrated via LangGraph/LangChain to support modular conversational workflows. In this paper the chatbot is examined specifically as the interaction interface within the broader RECMOOC4All ecosystem (MOOC aggregator, learner model, recommender, and accessibility checker), and claims are limited to its role in mediating discovery and onboarding tasks rather than asserting platform-level accessibility compliance. Consistent with this scope, non-chatbot modules are described only insofar as they inform the chatbot's interaction design and evaluation.

A methodological emphasis is placed on research traceability and reporting clarity. This work reports an exploratory evaluative pilot study of a live prototype. It investigates feasibility (can the system operate reliably under realistic conditions?) and generates early evidence of utility (does it help participants' complete discovery and onboarding tasks through text and voice across languages?), rather than establishing generalizable effects. The aim of the study is to examine whether an accessibility-first GenAI chatbot can reduce interaction friction during MOOC discovery and early navigation while supporting multimodal and multilingual access under pilot conditions. The following research questions were formulated to guide the inquiry: (RQ1) How reliably and how quickly does the chatbot support core discovery and navigation tasks across text and voice, as indicated by intent handling, speech-to-text performance, and response latency? (RQ2) How robust is multilingual mediation for MOOC queries, as reflected in the adequacy of translations for domain-specific intents and entities? (RQ3) To what extent does the chatbot enable accessibility-aware course discovery by surfacing accessibility-relevant attributes using available metadata and/or audit outputs? (RQ4) What pilot evidence from user feedback and interaction telemetry indicates changes in access and engagement for diverse learners, including participants reporting disabilities?

To ensure construct consistency across sections, key outcomes are operationalized as follows: access is defined as successful completion of core discovery and navigation tasks via a preferred modality (text or voice) with appropriate interaction supports; engagement is operationalized through telemetry proxies (e.g., session duration and depth of interaction) rather than learning gains; achievement refers to task completion within scenarios rather than course performance; and persistence/retention denotes continued system use beyond an initial interaction, interpreted cautiously given the short pilot window. WCAG 2.2 is used as a design and evaluation reference for interaction screens and discovery workflows, consistent with the definition of conformance as a property of web content and interfaces. Accordingly, the chatbot's design follows WCAG 2.2 principles and aims to facilitate access to WCAG-aligned content, but it does not itself confer conformance on external course materials.

The remainder of this paper synthesizes related work on MOOC accessibility and educational chatbots, then presents the methodology and scenario-based evaluation used to assess the chatbot interface. Results are reported by research question, followed by a discussion that contextualizes findings against prior literature, articulates feasibility and scalability considerations (including external service dependence), and acknowledges limitations and future directions for inclusive GenAI-supported learning environments.

LITERATURE REVIEW

This section synthesizes research relevant to accessibility-first conversational support for MOOC discovery and early navigation across three domains: (1) accessibility barriers in MOOCs, with emphasis on discovery-layer signals and the role of accessibility metadata; (2) GenAI affordances that enable multimodal and multilingual mediation; and (3) evidence from educational GenAI chatbots, focusing on what is reported and evaluated regarding accessibility- and UDL-relevant interaction supports. This synthesis motivates a study focus on the interaction layer (search, onboarding, and early navigation), where barriers can prevent participation before learning activities begin.

MOOC ACCESSIBILITY BARRIERS AND DISCOVERY-LAYER CONSTRAINTS

MOOCs have expanded participation opportunities, yet access remains uneven for learners who rely on assistive technologies or alternative sensory, cognitive, or linguistic supports (Bozkurt et al., 2024; Farrow et al., 2024; Hamideh Kerdar et al., 2024; Ingavélez-Guerra et al., 2023; Królak et al., 2022). Platform-level accounts commonly describe variability across courses and content types within the same provider rather than treating accessibility as a uniform platform property, reflecting differences in authoring practices, partner institutions, and media formats (Rodriguez-Ascaso et al., 2024).

Two barriers recur across recent analyses. First, video lectures remain a dominant format, yet captions and transcripts are inconsistently available or vary in quality; automated captions may degrade under real-world conditions, reducing reliability for deaf and hard-of-hearing learners and for learners who benefit from text alternatives (Burgstahler, 2023; Lee, 2024; Risi, 2025). Visual learning resources also frequently lack equivalent alternatives (e.g., meaningful alternative text or audio description), limiting usability for blind and low-vision learners. Second, accessibility is often difficult to identify during course selection. When platforms do not support structured filtering by accessibility-relevant features (e.g., captioning, transcripts, language options, screen-reader compatibility indicators), learners may need to manually inspect courses to determine fit an added burden in low-bandwidth contexts and for multilingual users (Aguis, 2024; UNESCO, 2021). Overall, barriers are documented at both content level (missing alternatives) and interaction level (weak accessibility signals during search and early navigation).

Table 1 summarizes what platform documentation and recent evaluations describe about accessibility targets, auditing, and recurring gaps (2024–early 2025), emphasizing reported/observed provisions rather than platform-wide conclusions (Coursera Support, 2025; EdX, 2025; FutureLearn, 2025; Support-Khan Academy, 2025). Across the reviewed platforms, captions are frequently present but inconsistently implemented; alternatives for visuals are uneven and often depend on course teams; and real-time transcription is not described as a standard feature in the sources reviewed. These patterns motivate support earlier in the learner pathway (discovery and early navigation), not only course-by-course remediation.

PLATFORM	REPORTED ACCESSIBILITY TARGET AND AUDIT STATUS	CAPTIONING AND TRANSCRIPTS (REPORTED/OBSERVED)	ALTERNATIVES FOR VISUALS (REPORTED/OBSERVED)	REAL-TIME TRANSCRIPTION	RECURRING GAPS REPORTED IN EVALUATIONS
Coursera (*Coursera Support*, 2025)	Reports a WCAG 2.2 Level AA target; biannual independent audits reported	Yes, for many videos; quality varies by course	Limited; often instructor-dependent	Not reported as a standard feature	Accessibility metadata inconsistent; variable screen-reader experience; course-level implementation uneven
edX (EdX, 2025)	Reports a WCAG 2.1 Level AA target; transition to WCAG 2.2 varies	Yes; captions provided, but upload/review practices vary	Limited; varies by course	Not reported as a standard feature	Instructor-dependent implementation; uneven automation/verification
Khan Academy (*Support-Khan Academy*, 2025)	Reports partial coverage of accessibility requirements; gaps acknowledged	Yes; auto-captions via YouTube for some content; accuracy varies	Minimal; limited coverage reported	Not reported as a standard feature	Limited built-in speech interaction; variable assistive-technology compatibility
FutureLearn (FutureLearn, 2025)	Reports a WCAG 2.2 Level AA target; improvements ongoing	Yes; captions on many videos; quality varies	Limited; not uniform	Not reported as a standard feature	Limited multilingual/cognitive supports; dependence on partner institutions

Table 1 Accessibility-related provisions reported for major MOOC platforms (2024–2025).

GENAI AFFORDANCES FOR MULTIMODAL AND MULTILINGUAL ACCESS

Natural Language Processing (NLP) has advanced from symbolic systems to transformer-based architectures that support context-sensitive generation and multilingual processing, enabling conversational interaction at scale (Goel & Polepeddi, 2018; Kuddus, 2022). When integrated with speech and translation technologies, LLM-based GenAI can support interaction-level accessibility through speech-to-text (STT), text-to-speech (TTS), multilingual mediation, and structured explanation/summarization (Holmes et al., 2022; Khong et al., 2024; Mrayhi et al., 2025a). These capabilities are particularly relevant to MOOC discovery and onboarding, where learners may require modality alternatives and language support to understand requirements, navigate interfaces, and identify courses that match accessibility needs.

Educational AI has long emphasized personalization through tutoring systems and recommender algorithms, but accessibility-specific outcomes are less consistently evaluated (Rane, 2023). GenAI extends the design space by enabling assistive interaction supports that can reduce friction during early navigation (e.g., speech-enabled assistance; stepwise conversational scaffolding) (Grover, 2024; Kumar & Nagar, 2024). At the same time, the inclusivity of GenAI is constrained by ethical and technical factors, including bias, uneven performance across languages, and underrepresentation of low-resource languages (Educause, 2025). These constraints increase the importance of explicit scope statements and transparent reporting of what is implemented and evaluated in accessibility-oriented GenAI systems.

EDUCATIONAL GENAI CHATBOTS: REPORTED SUPPORTS AND RECURRENT LIMITATIONS

GenAI chatbots are increasingly deployed to support tutoring, navigation, and engagement, yet accessibility-oriented design and evaluation are inconsistently specified. Targeted studies indicate that chatbots can support accessibility-relevant interaction (e.g., speech output for navigation support; structured conversational scaffolding for comprehension) (Grover, 2024; Kumar & Nagar, 2024). However, published accounts often provide limited detail on multimodal coverage (speech input/output across contexts), multilingual robustness, and whether discovery-layer accessibility attributes are incorporated into guidance (Nataraj et al., 2024; Sarangam, 2024).

Mainstream deployments illustrate similar constraints. Khanmigo offers speech output and multilingual support, but reported coverage does not necessarily extend across platform contexts and may provide limited speech input and disability-sensitive adaptation (Caelen & Blete, 2024; *Khan Academy*, 2025). Jill Watson has been presented as a scalable virtual teaching assistant for courseware-grounded Q&A in LMS/forum contexts; published accounts foreground teaching-presence support and automation rather than accessibility-oriented interaction supports as evaluation targets (Goel & Polepeddi, 2018; Kakar et al., 2024). Coursera's AI features emphasize recommendation and learner assistance, yet published descriptions do not consistently indicate that accessibility metadata or disability-responsive supports are incorporated into interaction flows (*Coursera Support*, 2025; Rao & Lin, 2024). Duolingo's chatbot incorporates speech and feedback supports within language-learning tasks (Jiang et al., 2024), but domain specificity limits transferability to cross-disciplinary MOOC discovery and navigation (Ouyang et al., 2024).

Because accessibility outcomes cannot be attributed to a chatbot in isolation, Table 2 summarizes these systems using capability-oriented descriptors (what is reported as an interaction support and what is emphasized in evaluation), rather than implying platform-wide accessibility claims. Across studies and deployments, three gaps recur: (1) accessibility-oriented discovery remains limited when accessibility attributes are not consistently surfaced during search and onboarding; (2) assistive interaction capabilities (speech input/output, multilingual mediation, structured scaffolding) are implemented unevenly or described at a high level; and (3) UDL is referenced more often than it is operationalized through reported design logic mapping learner needs to checkpoints and triggerable supports.

SYSTEM	PRIMARY FUNCTION (REPORTED)	ACCESSIBILITY-RELEVANT INTERACTION SUPPORTS REPORTED	MULTILINGUAL SUPPORT REPORTED	UDL MAPPING OR CHECKPOINT LOGIC REPORTED	LIMITATIONS
Khanmigo (Support-Khan Academy, 2025)	Tutoring/guidance	TTS for chatbot responses (coverage varies by context)	Partial	Not reported	Limited STT; partial content coverage (Caelen & Blete, 2024)
Jill Watson (Goel & Polepeddi, 2018)	Forum/LMS Q&A	Not reported as accessibility-oriented support	Not reported	Not reported	Forum-bound; accessibility supports not reported
Coursera AI (Coursera Support, 2025)	Assistance/ recommendation	Captions; indirect UX support	Varies	Not reported	Accessibility-aware discovery not explicit
Duolingo AI (Ouyang et al., 2024)	Language learning dialogue/feedback	TTS, visual aids, gamified interaction	Yes	Partial	Domain-specific; limited transfer to MOOC discovery

Table 2 Comparison of GenAI Chatbots and Accessibility Features.

POSITIONING AN ACCESSIBILITY-FIRST CHATBOT FOR DISCOVERY AND EARLY NAVIGATION

This literature indicates that improving access in MOOCs requires attention to interaction friction during discovery and early navigation, in addition to content-level remediation. It also shows that existing chatbots rarely foreground accessibility as an evaluation target or document how interaction supports operationalize UDL beyond general claims. The RECMOOC4All Chatbot is therefore positioned as an accessibility-first interaction interface that targets discovery and early navigation through multimodal interaction (speech-to-text and text-to-speech), multilingual mediation, and accessibility-aware discovery when accessibility attributes are available via metadata and/or audit outputs. This scope supports a focused evaluation of interaction-layer accessibility supports while acknowledging that content-level accessibility depends on course materials and provider implementation.

METHODOLOGY

This paper presents the design and implementation of the RECMOOC4All GenAI-powered chatbot and reports an exploratory pilot evaluation of a live prototype. To evaluate the RECMOOC4All chatbot for MOOC discovery and early navigation, we followed the Cross-Industry Standard Process for Data Mining (CRISP-DM) (Chapman et al., 2000) as an iterative framework from requirements elicitation to data understanding, data preparation, modeling, and evaluation. Interaction design was guided by UDL principles and UNCRPD-informed inclusion goals, and evaluation reports evidence from selected WCAG 2.2-relevant checks applied to the chatbot's user-facing flows.

BUSINESS UNDERSTANDING

Learners with visual, cognitive/neurodiverse, motor and multilingual needs encounter barriers during MOOC discovery and onboarding, including limited multimodal options, inaccessible navigation pathways and high cognitive load in early decision-making. RECMOOC4All positions the chatbot as the primary interaction interface within an ecosystem that also includes a MOOC aggregator, learner modeling, a recommendation engine and an accessibility checker (Figure 1). Within this scope, the chatbot serves two functions:

1. Interaction and guidance: Supporting discovery-oriented dialogue (course search, clarification, onboarding help and first-step guidance) through text and voice modalities, with multilingual capabilities.
2. Signal capture for personalization: Logging declared preferences and interaction signals (modality choice, reformulations, navigation loops) and passing them to the learner-modeling module to refine recommendations and subsequent interactions.

This delineation clarifies that the chatbot manages the conversational layer, while personalization and accessibility assessments are handled by connected services.

SYSTEM ARCHITECTURE

The RECMOOC4All ecosystem comprises five coordinated modules (Figure 1):

- GenAI conversational chatbot (interaction layer): manages real-time dialogue, supports discovery for non-registered users and personalization for registered users via backend services, and orchestrates dialogue routing.
- MOOC aggregator: consolidates course metadata and attaches available accessibility-related attributes (e.g., subtitle/caption indicators when available, language, and content descriptors).
- Learner modeling: maintains and updates adaptive profiles using declared preferences and interaction telemetry.
- Recommender system: generates recommendations using content-based retrieval (e.g., TF-IDF) and ranking strategies appropriate to the available signals, then filters or annotates outputs based on accessibility constraints and available indicators.
- Accessibility checker and remediation (decision support): performs selected WCAG-relevant checks and produces structured outputs (e.g., indicators and remediation hints) that can be surfaced to the chatbot to support informed discovery decisions.

The chatbot integrates GPT-4o-mini for intent routing and response planning, Whisper for speech-to-text, and OPUS-MT for multilingual translation, orchestrated via LangGraph/LangChain to support modular conversational workflows. Pilot sessions were conducted on a live prototype (not a simulated environment), and system logs captured latency, timeouts/failures, and fallback behavior to characterize feasibility under realistic network conditions.

Open Praxis
DOI: 10.55982/openpraxis.18.3.1026

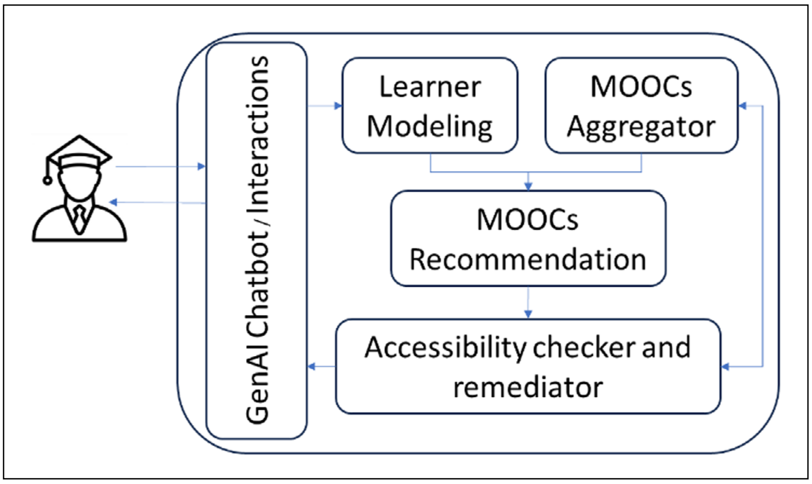

Figure 1 Overview of the RECMOOC4ALL system architecture.

Runtime workflow and orchestration

Figure 2 illustrates the state-aware routing workflow: language detection, optional STT for audio input, optional affect/engagement components if enabled, query classification, and dispatch to task-specific chains (recommendation, assistance, feedback, quiz or roadmap), with optional TTS for voice output. State transitions are logged to support traceability of routing and adaptation decisions.

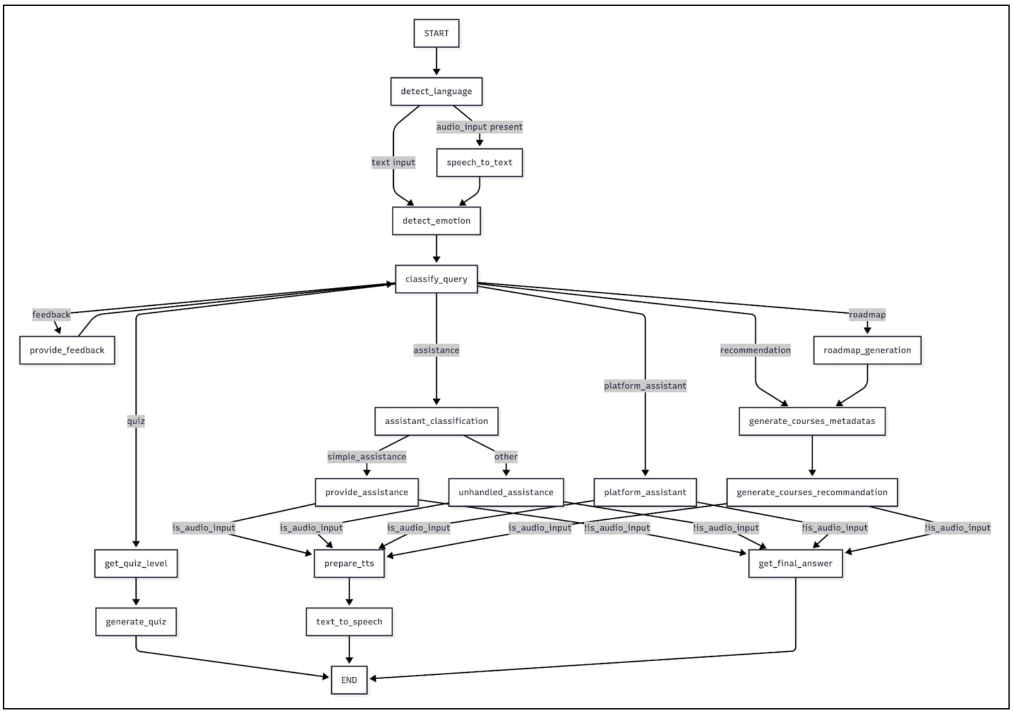

Figure 2 Global architecture of the RECMOOC4All chatbot subsystem.

Accessibility mapping and adaptation policy

Accessibility needs were operationalized as interaction supports during discovery and early navigation via a rules-based policy. Table 3 summarizes how each need activates specific features.

Table 3 Accessibility mapping in RECMOOC4All.

ACCESSIBILITY NEED	SYSTEM FEATURE
Visual Impairment	TTS output; screen-reader-compatible; voice input via Whisper STT
Cognitive Impairment	Simplified language; stepwise guidance; reduced information density (low-density formatting)
Multilingual Access	Real-time Neural Machine Translation (NMT) (M2M-100, OPUS-MT) with glossary control for accessibility and pedagogical terminology
Motor Impairment	Keyboard-first interaction flows; voice input via STT as an alternative to pointer-dependent interaction

Supports are activated through: (i) declared preferences (e.g., audio first; simplified output), (ii) interaction signals (e.g., repeated reformulations; prolonged clarification loops; sustained keyboard-only use) and (iii) context cues (e.g., detected language mismatch). When accessibility information is missing, the system preserves uncertainty (e.g., "unknown/not specified") rather than inferring accessibility properties. All decisions and logs are sent to the learner model for ongoing refinement.

Study objectives

Objectives were scoped to discovery/onboarding and mapped to the CRISP-DM phases:

- Improve discovery and navigation (business & data understanding): Enable natural-language search and guided clarification; surface available accessibility signals to inform selection decisions.
- Support personalization (data understanding & modeling): Capture multimodal preferences and interaction traces to inform learner profiles and refine recommendations.
- Reduce interaction friction (evaluation & deployment): Use explicit and implicit feedback to identify breakdowns (e.g., loops/reformulations) and iteratively improve dialogue scaffolding and modality selection.
- Enable multilingual access (data preparation & deployment): Facilitate language detection and translation so learners can interact in their preferred language while preserving critical meaning.

DATASETS AND DATA PREPARATION

Following CRISP-DM, the chatbot pipeline used multimodal, multilingual and accessibility-tagged resources supporting three functions: intent routing, assistance dialogue and MOOC discovery prompts (Table 4). Training and benchmarking resources were derived from publicly accessible sources (e.g., provider support materials and open multilingual benchmarks); no proprietary platform logs were used for training. Pilot interaction traces were collected under consent and used only for evaluation and telemetry-driven refinement.

Table 4 Overview of integrated datasets.

DATASET/SOURCE	DATA TYPE	PURPOSE
MOOC Q&A pairs (Coursera, edX, Udemy)	Text (QA pairs)	Base corpus for intent modeling and accessibility-aware response patterns
Paraphrase-augmented Q&A (GPT-2/GPT-3.5 + BERT checks)	Text	Increase lexical robustness while controlling meaning drift
Dual-labeled classification set (n = 14,082)	Text + labels	Train/validate intent routing and accessibility-aware dispatch
Assistance dataset (n = 2,597)	Text (query–response pairs)	Dialogue patterns for early navigation support
Synthetic speech set (TTS → Whisper STT transcripts)	Audio → text (+ confidence/timestamps)	Validate multimodal routing and voice interaction under controlled conditions
Translation resources (M2M-100, OPUS-MT; optional Google Translate)	Multilingual text	Real-time multilingual mediation with glossary control for domain terminology
Language-ID benchmark (Zarajamshaid; public)	Text	Benchmark language identification before translation/routing
Translation validation corpus (OPUS-100; public)	Text (parallel translations)	Quality checks for translation adequacy
Pilot interaction & feedback logs (consented)	Text + telemetry/metadata	Evaluation and feedback-driven refinement

MOOC Q&A corpus and augmentation

A corpus of 2,526 MOOC-related Q&A items was compiled from publicly accessible support/help materials of selected providers (Coursera, edX, Udemy). After filtering for relevance, linguistic clarity and redundancy, 613 Q&A pairs were retained. To improve robustness to phrasing variation, an offline paraphrase augmentation step produced 1,520 instances, with consistency checks to limit meaning drift. Each instance was labeled for (i) intent (discovery/recommendation, technical assistance, feedback/reporting) and (ii) accessibility need (visual, auditory, cognitive, or motor). Augmentation was used offline; online routing relied on lightweight GPT-4o-mini chains.

Dual-labeled classification dataset

A 14,082-row classification dataset followed the same intent + accessibility labeling schema. Labeling combined expert annotation on a seed subset with rule-based expansion driven by lexical complexity indicators, interaction cues, and modality markers. Inter-rater reliability on the manually labeled subset was assessed using Cohen's Kappa ($\kappa = 0.86$), indicating strong agreement (Kappas, 2013). These labels support state-aware routing (recommendation/assistance/feedback) in LangGraph.

Assistance dataset

A parallel assistance dataset of 2,597 query–response pairs was constructed from structured scenarios aligned with early navigation needs (onboarding, clarification, troubleshooting). Items were filtered for semantic fidelity post-translation, STT/TTS compatibility and cognitive load constraints.

Multilingual resources and language identification

All textual datasets were translated into 100+ languages using M2M-100 and OPUS-MT, with glossary-based control for accessibility terms. Translation quality was verified using semantic-similarity measures and OPUS-aligned corpora (e.g., OPUS-100). Language detection used an ensemble (LangDetect, LangID, Polyglot) benchmarked on the Zarajamshaid dataset (17,000 Wikipedia entries) with an accuracy of 95.82 % (Hugging Face, 2025).

Synthetic speech set for voice interaction

To validate voice-based interaction, a synthetic speech set was created by converting representative queries to audio (TTS), resampling to 16 kHz, normalizing, and transcribing via Whisper (retaining confidence scores and timestamps). These transcripts were processed through the same routing pipeline as text inputs. Possible degradation in low-resource languages or noisy conditions is regarded as an operational risk (Liu et al., 2024).

Pilot feedback logs

During pilot sessions, feedback was captured via RESTful APIs. Logged signals included explicit feedback (ratings, satisfaction) and implicit signals (reformulations, navigation loops, session drop-offs). Telemetry was used to characterize feasibility (runtime behavior) and interaction friction, and was forwarded to the learner-modeling module for iterative refinement.

These datasets enable multimodal and multilingual discovery while preserving traceability of routing decisions and adaptation triggers within the prototype scope.

DATA PREPARATION

A reproducible pipeline transformed raw inputs into standardized, accessibility-tagged records (Figure 3). The pipeline sequences (1) missing-field identification; (2) cross-referencing external public sources; (3) consolidating provider metadata; (4) normalizing fields; and (5) validating interaction-layer accessibility indicators. The following steps ensured data readiness:

- Normalization and harmonization: unify schemas across text, transcripts and feedback logs; standardize encodings, timestamps and language codes; resample audio to 16 kHz mono.
- Accessibility-aware enrichment: add structured indicators used by the interaction policy (e.g., readability proxies; response density/chunking controls; STT/TTS readiness; keyboard-first pathways).

- Validation and remediation hints: apply rule-based checks to interaction-layer metadata and selected WCAG-relevant indicators (e.g., subtitle/transcript flags; keyboard-path readiness) alongside UDL-informed readability/structure indicators. Missing or inconsistent attributes trigger structured remediation hints (e.g., request transcript; reduce density; provide keyboard-first guidance).
- Missing/uncertain data handling: impute accessibility-critical fields only when defensible; otherwise mark as unknown to avoid overstated assumptions.
- Deduplication and quality control: remove duplicates/corrupted entries; apply conservative imputations for non-accessibility fields while preserving accessibility tags.

This preparation ensures each record can be routed by intent and will activate the appropriate interaction supports (e.g., simplified responses, voice-first workflows, multilingual mediation) under consistent and auditable rules.

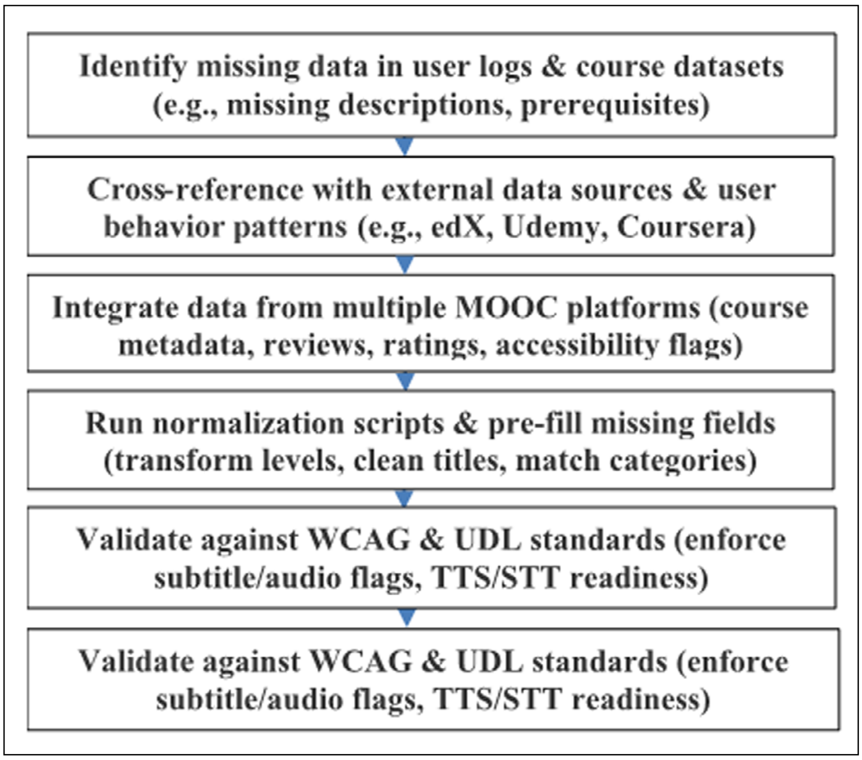

Figure 3 Data refinement proces.

PILOT STUDY DESIGN

Participants and context

A feasibility pilot was conducted with 15 participants (8 women, 7 men) recruited through the Avaxia accessibility initiative. Six participants (40%) self-reported disability-related needs (visual impairments, n = 3; cognitive processing difficulties, n = 2; combined visual–linguistic profile, n = 1). Participants interacted with the live prototype using their preferred modality (text and/or voice) and preferred language. Sessions were conducted in a live setting, and the system recorded interaction telemetry alongside post-session feedback. The pilot was designed to assess feasibility and early utility for discovery and onboarding tasks, rather than to produce statistically generalizable effectiveness estimates.

Procedure and tasks

Participants completed a scenario-based session covering early-stage MOOC interactions: (1) course search and refinement, (2) clarification requests (e.g., prerequisites, pace, subtitles/captions), (3) onboarding/help requests, (4) preference expression (e.g., audio-first, simplified output, language choice) and (5) feedback submission. Sessions were run on a live prototype and telemetry was captured for feasibility analysis and iterative refinement.

MEASURES AND INSTRUMENTS

System and interaction measures

System logs captured: (i) intent-routing accuracy (match between expected task intent and routed chain), (ii) end-to-end latency (request-to-response time) including timeouts/failures and fallback activations, (iii) multimodal behavior (text/voice selection, switches, repeat attempts), (iv) interaction friction (reformulations, clarification loops, early drop-offs), and (v) multilingual pipeline events (language-ID outcomes, translation activations, and verification flags).

User-reported measures

After each session, participants completed short Likert-scale items (plain language; available in the participant's preferred language) assessing perceived usefulness, clarity, ease of navigation, and perceived accessibility support. Two open-ended questions captured perceived strengths, friction points, and improvement suggestions.

Qualitative feedback

Open-ended responses were analyzed using a hybrid deductive–inductive thematic approach. An a priori framework (friction points; clarity/cognitive load; modality usefulness; multilingual clarity; accessibility-signal usefulness) guided first-cycle coding; inductive subcodes were added as needed. Two coders independently coded an initial subset (n_1 = 9/30, 30%) to refine the codebook, then reliability was computed on a held-out subset (n_2 = 6/30, 20%) using Cohen's κ at the parent-code level (presence/absence per response; multi-coding permitted). Disagreements were resolved through codebook-based discussion and consensus, and refinements were logged.

Data analysis

Telemetry and Likert responses were summarized descriptively (counts and means/medians with dispersion). Subgroup comparisons (e.g., by modality) were exploratory and reported without population-level claims. Qualitative themes were finalized through iterative grouping and coherence checks, with illustrative de-identified excerpts used to support interpretation.

MODELING AND EVALUATION

Speech-to-Text model selection

Whisper was integrated for multilingual STT with standardized preprocessing. Transcripts included confidence scores and timestamps, enabling the system to seek clarification when transcriptions were uncertain. Evaluation was based on the synthetic speech set; potential degradation in noisy or low-resource language conditions was treated as an operational risk.

Intent routing and dialogue orchestration

User requests were routed into three intents (recommendation/discovery, assistance, and feedback) via a lightweight classifier. A state-aware orchestration layer dispatched requests to the appropriate chain and logged state transitions for traceability.

Retrieval/recommendation support

For discovery prompts, the system combined lexical and semantic retrieval over the MOOC metadata store and surfaced results annotated with available accessibility signals (e.g., subtitles/captions). Where signals were absent or uncertain, the system preserved uncertainty rather than inferring accessibility properties.

Multilingual mediation

An ensemble of language detectors identified the input language. When needed, NMT mediated between the user's language and the internal processing language, with glossary controls to stabilize terminology. Translation quality was validated using semantic similarity checks and targeted back-translation for high-risk queries (e.g., accessibility-specific requests).

RUNTIME AND DEPLOYMENT CONSIDERATIONS

The prototype integrates third-party AI components (STT, NMT, and LLM-based routing/assistance). To assess feasibility under realistic conditions, pilot sessions were conducted on a live prototype and telemetry captured latency, timeouts/failures, and fallback behavior (e.g., retry, simplified response, or alternative modality suggestion). These measures informed scalability discussion by grounding performance claims in observed runtime behavior rather than assuming ideal network conditions.

ETHICS, PRIVACY, AND DATA GOVERNANCE

This feasibility pilot was conducted as a minimal-risk usability study. Participants were recruited on a voluntary basis and provided informed consent after receiving a plain-language summary of (i) what data would be collected (text queries, optional voice input, system telemetry and feedback), (ii) the purpose of the pilot (testing early-stage MOOC discovery and onboarding), (iii) storage and retention practices, and (iv) their right to withdraw at any time without penalty. Only the data necessary for assessing interaction feasibility were collected; no video, facial imagery or other biometric signals were recorded. Interaction logs were pseudonymized prior to analysis and stored on access-restricted servers. No formal institutional research ethics committee (REC/IRB) review was available or required for this type of minimal-risk pilot; therefore, the study adhered to prevailing research integrity and data-protection standards. Any sharing of derived materials (e.g. de-identified excerpts, aggregated telemetry or evaluation scripts) is restricted to datasets from which individuals cannot be re-identified and will occur only with participant consent and appropriate data-protection safeguards.

DATA AVAILABILITY

All third-party datasets and pre-trained models used for development and benchmarking (such as public MOOC support corpora and open multilingual resources) are publicly available. Study-generated interaction traces are not made public because they may contain sensitive or personally identifying content. Anonymized excerpts or aggregated summaries can be provided on reasonable request, subject to participant consent and compliance with data-protection requirements.

RESULTS

This section reports the empirical and technical outcomes of the RECMOOC4All GenAI-powered chatbot pilot evaluation, focusing on MOOC discovery and early onboarding (search/refinement, clarification, first-step guidance, and feedback). Consistent with the exploratory design, findings are presented as feasibility and early utility evidence from a live prototype rather than population-level effectiveness claims. Building on the architecture described earlier (Figure 2), results examine how the chatbot's modular workflow supports inclusive interaction across modalities and languages. Grounded in UDL and using WCAG 2.2 as a reference for the chatbot's user-facing interaction flows, the system leverages multimodal and multilingual interfaces to reduce friction for learners with visual impairments, cognitive challenges, and multilingual needs. Results are organized by research question: (i) routing/voice and runtime feasibility (RQ1); (ii) multilingual mediation adequacy (RQ2); (iii) accessibility-aware discovery behavior (RQ3); and (iv) pilot utility evidence from user feedback and telemetry (RQ4).

RQ1: RELIABILITY AND RUNTIME FEASIBILITY ACROSS TEXT AND VOICE

Using the GPT-4o-mini routing component integrated via LangChain/LangGraph, intent detection achieved 92.1 % agreement with expected scenario intents. Whisper (tiny) transcribed voice queries with a Word Error Rate (WER) of 9.2 % on the controlled speech set described in the Methodology. When transcription confidence was low, the system prompted users for clarification.

System logs recorded mean end-to-end response times of 823 ms for text queries and 1.24 s for voice queries (including STT). Timeouts and fallbacks (e.g., retry prompts, simplified responses, modality-switch suggestions) were logged as feasibility indicators. Together, routing accuracy, STT performance, and runtime latency indicate that the system can reliably and quickly support discovery/onboarding tasks across text and voice. Building on these modality-level findings, we next examine how well the chatbot mediates across languages.

RQ2: ROBUSTNESS OF MULTILINGUAL MEDIATION

The ensemble language detector (LangDetect, LangID, Polyglot) achieved 95.82 % precision on the Zarajamshaid benchmark, enabling dependable language identification. Translation fidelity was evaluated with BERTScore on domain queries containing MOOC and accessibility terminology. OPUS-MT outperformed M2M-100 (F1 = 0.8962 vs. 0.8155); therefore, OPUS-MT was chosen as the default translation layer for the pilot's main languages (Arabic, French, Urdu), with M2M-100 retained for fallback. Glossary constraints stabilized key accessibility and pedagogical terms (e.g., captions, screen reader, prerequisites). These results indicate that multilingual mediation preserves domain meaning for MOOC queries in the evaluated settings. Since effective translation alone is not enough to ensure equitable access, we next evaluate whether the chatbot can surface relevant accessibility information during course discovery.

RQ3: ACCESSIBILITY-AWARE COURSE DISCOVERY USING AVAILABLE SIGNALS

During discovery dialogues, the chatbot surfaced accessibility-related attributes when they existed in course metadata or structured outputs (e.g., language, caption/subtitle indicators, modality descriptors). When signals were absent, the system explicitly displayed "unknown/not specified" and asked clarifying questions (e.g., "Do you require captions?" or "Is keyboard-only navigation essential?"). WCAG 2.2 is treated as a design reference for the chatbot's interaction flows rather than as a blanket conformance claim; thus, the interaction layer surfaces available signals and preserves uncertainty when attributes are missing. This behavior shows that the system can support accessibility-aware discovery at the interaction layer. To understand how users experience these features and to gauge pilot-level utility, we are now transitioning to user feedback and telemetry.

RQ4: PILOT UTILITY EVIDENCE FROM USER FEEDBACK AND TELEMETRY

Post-session Likert scores averaged 4.4/5, indicating positive perceived usefulness and usability. Accessibility-specific reactions differed by profile: participants with visual impairments valued the voice/TTS interface and screen-reader-friendly navigation; those with cognitive processing challenges highlighted simplified outputs and reduced information density; multilingual participants emphasized clarity and stability of translated responses.

Open-ended responses were thematically analyzed as described in the Methodology. The resulting themes (friction points, clarity/cognitive load, modality usefulness, multilingual clarity, accessibility-signal usefulness) were coded with substantial inter-rater reliability ($\kappa = 0.81$).

Telemetry provided additional pilot-level insights: relative to an internal text-only baseline under the same tasks, session duration increased by about 25%, bookmarking by about 15%, and session completion among participants with disability-related needs by about 31%. These deltas are exploratory and should be interpreted cautiously given the small sample and short observation window, but they are consistent with reduced friction during early discovery and onboarding.

Taken together, self-reports, coded feedback, and interaction proxies provide early evidence of improved access-related task completion and engagement behavior across diverse learner profiles. No video, facial, or biometric data were collected; therefore, no affective-tracking outcomes are reported.

SCENARIO-BASED INTERACTION CASES

Scenario-based sessions illustrate how the system behaves across varied accessibility needs:

Scenario 1: Sighted learner (English text)

A neurotypical learner asked "*I want to learn SQL*"; the chatbot parsed the intent and delivered targeted course recommendations with low latency (Figure 4).

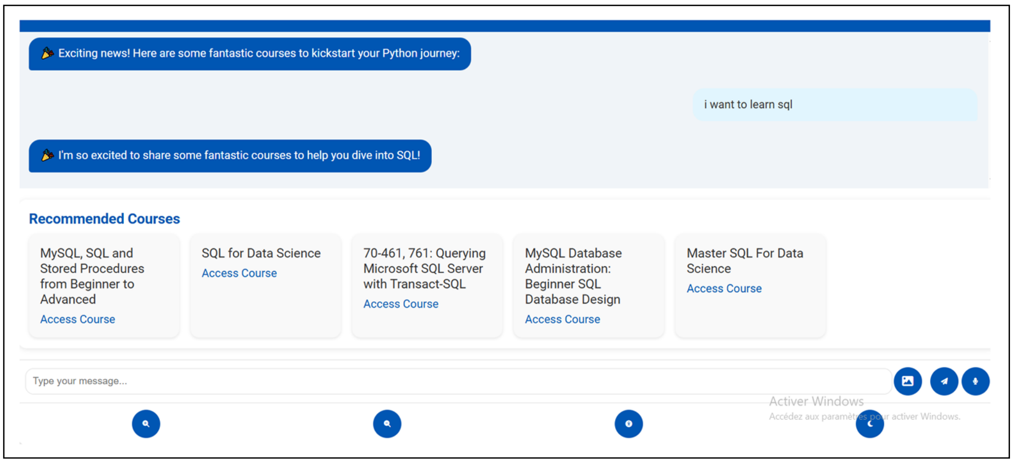

Figure 4 SQL course recommendations for a sighted learner using English text input.

Scenario 2: Visually impaired learner (Arabic voice)

A blind learner interacted via Arabic speech; after STT processing, the chatbot returned concise Arabic responses, surfaced available caption/subtitle indicators, and offered a voice-first, keyboard-compatible flow. The user reported improved independent navigation compared to prior experiences (Figure 5).

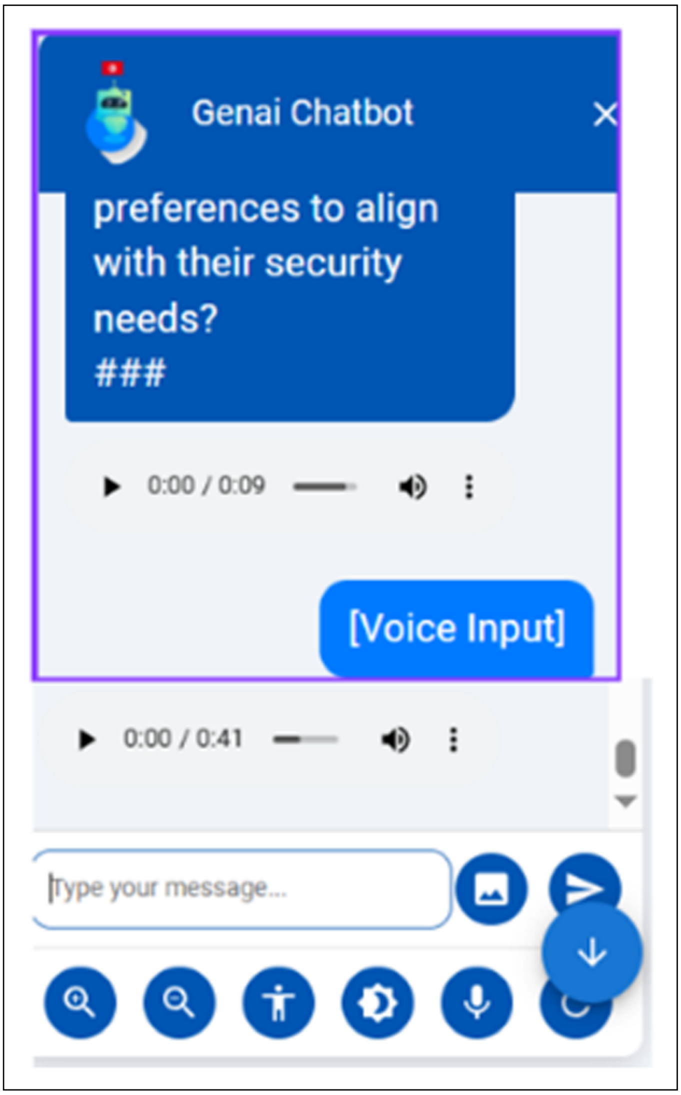

Figure 5 Arabic voice interaction with TTS output for a visually impaired learner.

Scenario 3: Learner with cognitive processing challenges (simplified UI)

- An adult learner with executive-function limitations used a simplified text interface; simplified language, reduced information density, and controlled pacing enabled successful search and bookmarking of courses without assistance (Figure 6).

DOI: 10.55982/openpraxis.18.3.1026

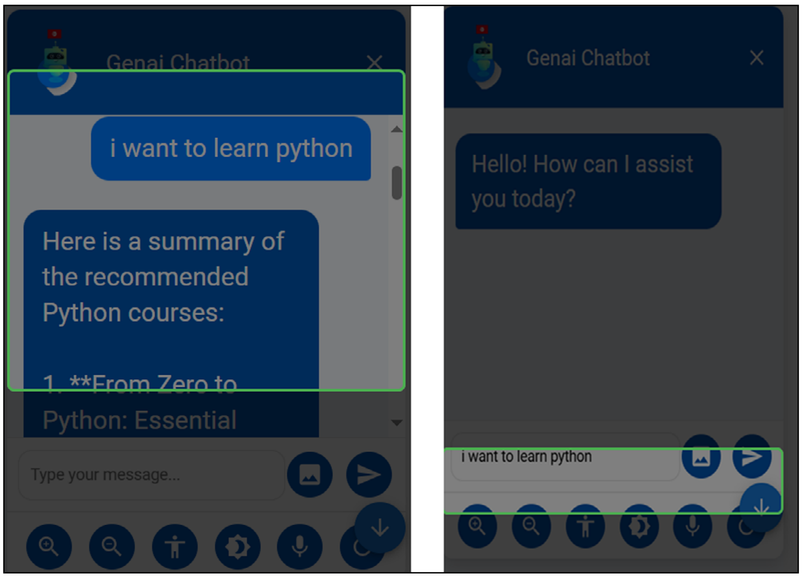

Figure 6 Simplified UI interaction for a learner with cognitive challenges.

Scenario 4: Multilingual learner with combined needs (Arabic voice + translation)

A learner with visual impairment and mild learning difficulties used Arabic voice input; language detection and glossary-constrained translation maintained stable terminology, while the interface adapted outputs to both visual and cognitive accessibility requirements. The learner felt understood and included (Figure 7).

Figure 7 Multilingual voice input with adapted output for a learner with combined needs.

Scenario 5: First-time MOOC learner (English onboarding)

A novice user asked how to create an account; the chatbot delivered clear, step-by-step instructions that enabled independent account creation (Figure 8).

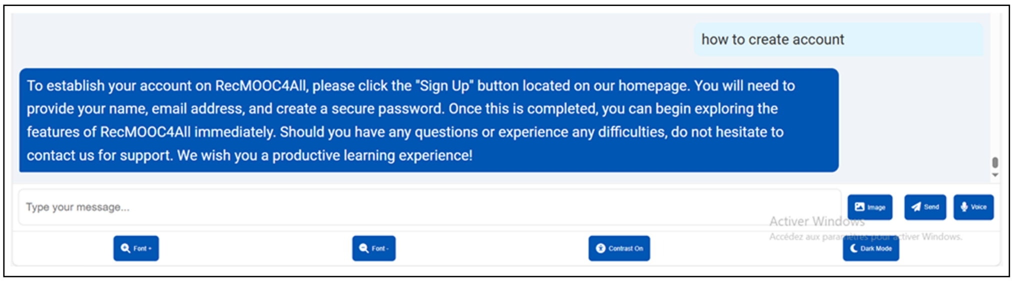

Figure 8 Account creation guidance via GenAI chatbot.

These scenarios show that the chatbot can accommodate varied modality preferences and accessibility profiles through voice and text, multilingual mediation, simplified phrasing, and structured guidance, all within an inclusive interaction design.

PERSONALISATION AND ADAPTIVITY

During pilot sessions, the chatbot adapted outputs based on declared preferences and interaction signals (e.g., modality choice, reformulations, clarification loops). Observed behaviors included chunked summaries and simplified phrasing, modality-aware scaffolding, and accessibility-aware annotation of discovery results when signals were available. Offline comparison of retrieval configurations showed a 38% precision gain for accessibility-tagged queries relative to a TF-IDF-only baseline in offline evaluation, supporting the retrieval design choice.

ACCESSIBILITY SIGNAL MONITORING

The integrated accessibility checker generated course-level accessibility indicators and remediation hints from available metadata (e.g., caption/subtitle flags, modality tags, transcript/alt-text fields when present) and lightweight page parsing where feasible. In the pilot recommendation set, 87.5% of courses contained sufficient signals to be flagged as "indicators present"; this does not imply full WCAG conformance of the course experience. When signals were missing or ambiguous, courses were labelled "unknown/not specified", deprioritized for accessibility-constrained queries, or flagged for follow-up.

DISCUSSION

This pilot study examined the feasibility of embedding a multimodal, multilingual GenAI chatbot into a MOOC aggregator with accessibility as a guiding principle. By centering accessibility rather than retrofitting it, the design responds to longstanding calls for UDL-aligned platforms that adapt to diverse learner profiles. The results demonstrate that a modular GenAI pipeline can reliably support discovery and onboarding across text and voice (92.1 % routing accuracy; 9.2 % WER) and handle multilingual inputs with high fidelity (language-ID precision ≈95.8 %; OPUS-MT F1 ≈0.8962). These outcomes are consistent with (Kumar & Nagar, 2024), who found that speech-enabled interfaces enhance autonomy for visually impaired learners, and with (Kennedy et al., 2023), who noted the importance of high-quality translation for under-served languages.

Behavioral proxies such as increased time-on-task, more bookmarking, and higher session completion for participants reporting disabilities suggest that accessibility-aware adaptations may support deeper engagement. This aligns with (Ding et al., 2024), who identified session persistence as a proxy for inclusive learning. However, these improvements are derived from a small, heterogeneous sample and lack a control group, so they should be interpreted as indicative rather than definitive.

POSITIONING AGAINST EXISTING GENAI CHATBOTS

Most existing educational chatbots (e.g., Khanmigo, Jill Watson) prioritize tutoring and general personalization, providing limited and inconsistent accessibility features (Caelen & Blete, 2024; Goel & Polepeddi, 2018). RECMOOC4All differs by integrating STT, TTS, NMT, simplified phrasing, and accessibility metadata into a unified interaction layer. This aggregator-level design may reduce dependence on individual instructors and improve consistency of accessibility signals across courses. To contextualize this distinction, Table 5 compares key interaction-layer features across several educational GenAI systems. The comparison emphasizes descriptive differences rather than formal conformance.

FEATURE	PRIOR SYSTEMS (KHANMIGO, JILL WATSON, COURSERA AI, DUOLINGO AI)	RECMOOC4ALL CHATBOT
Accessibility support (interaction layer)	Partial (basic captions, limited TTS/STT)	Accessibility-first interaction supports (STT/TTS options, simplified/chunked responses, keyboard-first flows) informed by WCAG 2.2/UDL design references
Primary GenAI role	Tutoring/personalization, assistance, content support	Discovery/onboarding assistance with accessibility-aware prompts and surfacing of available accessibility signals
Adaptation logic	Often opaque or platform-dependent	Rule-based adaptation triggers and logged state transitions (prototype)
Integration Scope	Single platform/ecosystem	Aggregator-level interface connecting to metadata, recommender, and checker outputs
Multilingual Support	Varies; often limited to major languages	NMT-mediated interaction (OPUS-MT/M2M-100) with glossary constraints for key accessibility terms

Table 5 Descriptive comparison of educational GenAI chatbots and RECMOOC4All chatbot.

Note. Comparison is descriptive and based on publicly documented feature sets; "WCAG/UDL" refers to design references for the chatbot interface, not conformance of external course content.

ARCHITECTURAL SCALABILITY AND MODULARITY

The RECMOOC4All prototype uses a modular architecture, built around LangGraph orchestration, LangChain prompts, ChromaDB semantic retrieval and SQL-Alchemy learner modelling, to support horizontal scalability and real-time personalization (Holmes et al., 2022). By decoupling components, it permits substitution of translation, STT or recommendation modules without interrupting service, which is essential for long-term evolution and potential cross-platform deployment. This modularity also enables continuous refinement: logged state transitions and feedback loops allow the system to adapt to emergent user needs.

As Schmitt notes, feedback-aware architectures enhance personalization by interpreting implicit user behavior, such as navigation loops and query reformulations. The present pilot validates the feasibility of this approach but does not yet test scalability across multiple providers (Schmitt, 2020); future work should examine how component substitution and feedback loops perform at larger scale and in diverse MOOC ecosystems.

LIMITATIONS

Several factors limit the generalizability of this pilot. The sample (n = 15) is small and heterogeneous, preventing robust statistical comparisons across subgroups. The absence of a randomized control condition precludes causal attribution of engagement gains. Language support is incomplete: OPUS-MT and M2M-100 cover Arabic, French and Urdu well, but low-resource languages and dialects are insufficiently supported. Speech-recognition accuracy declines in noisy or accented conditions, and the evaluation focused only on discovery and onboarding tasks. Future research should therefore avoid overgeneralizing these preliminary findings.

FUTURE DIRECTIONS

To enhance inclusivity, future work should expand language coverage and improve robustness for dialects and noisy inputs. Sign-language support could aid Deaf learners, as suggested by (Yarrow et al., 2023), , while emotion-aware pacing may personalize interaction for diverse cognitive profiles. Larger, controlled, multi-site studies such as those advocated by (Tlili, 2024), are needed to evaluate long-term learning outcomes and equity impacts. Integrating domain-specific glossaries and paraphrasing models may assist learners with limited literacy or cognitive constraints. Finally, any expansion should be accompanied by transparent privacy safeguards and independent accessibility audits.

In sum, this pilot provides preliminary evidence that an accessibility-first, modular GenAI chatbot can enhance discovery and early navigation in MOOCs. The approach complements existing literature (Mrayhi et al., 2025b; Meyer et al., 2014; Tlili et al., 2025) by demonstrating a practical implementation of UDL principles within a Generative AI framework. However, its broader effectiveness and scalability remain to be validated through larger and more diverse evaluations.

CONCLUSION

This study assessed the feasibility of deploying the RECMOOC4All chatbot as the user-facing, accessibility-first interaction interface of a MOOC aggregator for course discovery and early onboarding. In a live pilot, the prototype supported text/voice interaction and multilingual mediation, and the evaluation provided convergent pilot evidence (scenario cases, retrieval precision gains for accessibility-tagged queries, and accessibility-checker outputs) that the interaction layer can surface available accessibility signals and reduce early navigation friction.

These findings should be interpreted as feasibility and early-utility evidence rather than population-level effectiveness. The pilot sample was small (n = 15) and heterogeneous, with no control group; therefore, observed telemetry deltas cannot be generalized or attributed causally. Moreover, accessibility-aware guidance is contingent on the completeness and reliability of course metadata and/or checker outputs and should not be interpreted as evidence of full WCAG conformance of external course experiences.

Overall, the pilot suggests that an accessibility-first generative AI strategy can support more inclusive MOOC discovery and onboarding. Validating this approach at larger scale and through more controlled evaluations, and through collaboration among key stakeholders, will be essential to achieve reliable, sustainable impact.

DATA ACCESSIBILITY STATEMENT

The datasets used and/or analyzed during the current study are available from the corresponding author on reasonable request.

SUSTAINABLE DEVELOPMENT GOALS (SDGs)

This study is linked to the following SDG(s): Quality education (SDG 4); Reduced inequalities (SDG 10).

ETHICS AND CONSENT

Ethical approval was not obtained because this work was conducted as a minimal-risk feasibility/usability pilot, and no formal REC/IRB review was available or required for this type of study. Participants were recruited voluntarily and provided informed consent using a plain-language (accessible/multilingual) summary describing the data collected (text queries, optional voice input, telemetry, and feedback), the purpose of the pilot, storage/retention practices, and the right to withdraw without penalty. Data collection followed data-minimization principles; no video, facial imagery, or biometric signals were recorded, and logs were pseudonymised and stored on access-restricted servers.

ACKNOWLEDGEMENTS

We gratefully acknowledge Avaxia for staff support and resources that were crucial to the project.

COMPETING INTERESTS

The authors have no competing interests to declare.

AUTHOR CONTRIBUTIONS (CRediT)

Salwa Mrayhi: Conceptualization; Methodology; Investigation; Software; Resources; Data curation; Formal analysis; Validation; Visualization; Writing-original draft; Writing-review & editing. Mohamed Koutheair Khribi: Conceptualization; Methodology; Formal analysis; Validation; Supervision; Writing-review & editing. Mohamed Jemni: Conceptualization; Methodology; Supervision; Project administration; Resources; Validation; Writing-review & editing. All authors have read and agreed to the published version of the manuscript.

AUTHOR NOTES

Based on *Academic Integrity and Transparency in AI-assisted Research and Specification Framework* (Bozkurt, 2024), the authors acknowledge the use of ChatGPT (OpenAI, GPT-5 Thinking, accessed September 2025) to support writing and refinement of this manuscript. All AI-assisted content was reviewed, critically edited, and validated by the authors to ensure academic rigor and adherence to ethical standards. The authors also assessed and mitigated potential biases in AI-generated output. The final content, interpretations, and conclusions are the sole responsibility of the authors.

AUTHOR AFFILIATIONS

Salwa Mrayhi orcid.org/0009-0000-2515-8304
LaTICE Research Laboratory, National Higher Engineering School of Tunis (ENSIT), University of Tunis, Tunisia

Mohamed Koutheair Khribi orcid.org/0000-0003-0901-0267
LaTICE Research Laboratory, National Higher Engineering School of Tunis (ENSIT), University of Tunis, Tunisia

Mohamed Jemni orcid.org/0000-0001-8841-5224
LaTICE Research Laboratory, National Higher Engineering School of Tunis (ENSIT), University of Tunis, Tunisia

REFERENCES

Aguis, J. (2024). Inclusive education for all: A qualitative exploration of fostering social justice and supporting students with special educational needs in the learning environment. *MCAST Journal of Applied Research & Practice, 7*(3), 20–39. https://doi.org/10.5604/01.3001.0054.2570

Bozkurt, A. (2024). GenAI et al.: Cocreation, authorship, ownership, academic ethics and integrity in a time of generative AI. *Open Praxis, 16*(1), 1–10. https://doi.org/10.55982/openpraxis.16.1.654

Bozkurt, A., Xiao, J., Farrow, R., Bai, J. Y. H., Nerantzi, C., Moore, S., Dron, J., Stracke, C. M., Singh, L., & **Crompton, H.** (2024). The manifesto for teaching and learning in a time of generative AI: A critical collective stance to better navigate the future. *Open Praxis, 16*(4), 487–513. https://doi.org/10.55982/openpraxis.16.4.777

Burgstahler, S. (2023). Designing inclusive formal and informal online learning: What do instructors need to know? In *Guide to digital accessibility* (pp. 113–128). Routledge. https://doi.org/10.4324/9781003438922-12

Caelen, O., & **Blete, M.-A.** (2024). *Developing apps with GPT-4 and ChatGPT: Build intelligent chatbots, content generators, and more*. O'Reilly Media.

Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., & **Wirth, R.** (2000). *CRISP-DM 1.0: Step-by-step data mining guide*. https://www.the-modeling-agency.com/crisp-dm.pdf

Coursera. (2025). *Coursera Support*. Retrieved December 24, 2025, from https://www.coursera.support/s/?language=en_US

Ding, Y., Cho, S.-J., Dong, C., Luo, Z., & **Li, H.** (2024). The involvement of academic and emotional support for sustainable use of MOOCs. *Behavioral Sciences, 14*(6), 461. https://doi.org/10.3390/bs14060461

EDUCAUSE. (2015). Using design thinking in higher education. *EDUCAUSE Review*. https://er.educause.edu/articles/2015/1/using-design-thinking-in-higher-education

edX. (2025). *edX*. Retrieved December 24, 2025, from https://www.edx.org/

Farrow, R., Iniesto, F., Pitt, R., Weller, M., & **Bossu, C.** (2024). Innovation with open educational resources: An integrative review of drivers, barriers and enablers. *Journal of Open, Distance, and Digital Education, 1*(2), 1–41.

FutureLearn. (2025). *FutureLearn Help Centre*. Retrieved December 24, 2025, from https://futurelearn.zendesk.com/hc/en-us

Goel, A. K., & **Polepeddi, L.** (2018). Jill Watson: A virtual teaching assistant for online education. In *Learning engineering for online education* (pp. 120–143). Routledge. https://doi.org/10.4324/9781351186193-7

Grover. (2024). Next-generation education: The impact of generative AI on learning. *Journal of Informatics Education and Research, 4*(2).

Guettala, M., Bourekkache, S., Kazar, O., Harous, S., & **others.** (2024). Generative artificial intelligence in education: Advancing adaptive and personalized learning. *Acta Informatica Pragensia, 13*(3), 460–489. https://doi.org/10.18267/j.aip.235

Hamideh Kerdar, S., Bächler, L., & **Kirchhoff, B. M.** (2024). The accessibility of digital technologies for people with visual impairment and blindness: A scoping review. *Discover Computing, 27*(1), 24. https://doi.org/10.1007/s10791-024-09460-7

Holmes, W., et al. (2022). State of the art and practice in AI in education. *European Journal of Education, 57*(4), 542–570. https://doi.org/10.1111/ejed.12533

Hugging Face. (2025). *Language identification dataset (papluca/language-identification)*. Retrieved December 24, 2025, from https://huggingface.co/datasets/papluca/language-identification

Ingavélez-Guerra, P., Otón-Tortosa, S., Hilera-González, J., & **Sánchez-Gordón, M.** (2023). The use of accessibility metadata in e-learning environments: A systematic literature review. *Universal Access in the Information Society, 22*(2), 445–461. https://doi.org/10.1007/s10209-021-00851-x

Jiang, X., Peters, R., Plonsky, L., & **Pajak, B.** (2024). The effectiveness of Duolingo English courses in developing reading and listening proficiency. *CALICO Journal, 41*(3), 249–272. https://doi.org/10.1558/cj.26704

Kakar, S., Maiti, P., Taneja, K., Nandula, A., Nguyen, G., Zhao, A., Nandan, V., & **Goel, A.** (2024). Jill Watson: Scaling and deploying an AI conversational agent in online classrooms. In *Proceedings of the International Conference on Intelligent Tutoring Systems* (pp. 78–90). https://doi.org/10.1007/978-3-031-63028-6_7

Kappas, A. (2013). Social regulation of emotion: Messy layers. *Frontiers in Psychology, 4*, Article 51. https://doi.org/10.3389/fpsyg.2013.00051

Kennedy, D., Halim, E., Condrobimo, A. R., Syamsuar, D., & **Ferdianto, F.** (2023). Overcoming language barriers in MOOCs with artificial intelligence: An AI-based approach for multilingual education. In *2023 Eighth International Conference on Informatics and Computing (ICIC)* (pp. 1–6). https://doi.org/10.1109/ICIC60109.2023.10381900

Khan Academy. (2025). *Does Khanmigo have a read-aloud or text-to-speech feature?* Retrieved December 24, 2025, from https://support.khanacademy.org/hc/en-us/articles/23772334788365-Does-Khanmigo-have-a-read-aloud-or-Text-to-speech-feature

Khong, P., Holmes, D., Masoudian, B., Lund, G. C., & **Garwood, S.** (2024). Lecture capture, transcripts, and captioning in US colleges of osteopathic medicine: Descriptive cross-sectional survey. *Medical Science Educator, 35*, 625–628. https://doi.org/10.1007/s40670-024-02224-4

Kilgore, W., & **Al-Freih, M.** (2017). MOOCs as an innovative pedagogical design laboratory. *International Journal on Innovations in Online Education, 1*(1). https://doi.org/10.1615/IntJInnovOnlineEdu.2016015210

Królak, A., Zając, P., & **Zając, P.** (2022). Analysis of the accessibility of selected massive open online courses (MOOCs) for users with disabilities. *Universal Access in the Information Society*, 1–12. https://doi.org/10.1007/s10209-022-00927-2

Kuddus, K. (2022). Artificial intelligence in language learning: Practices and prospects. In *Advanced analytics and deep learning models* (pp. 1–17). https://doi.org/10.1002/9781119792437.ch1

Kumar, A., & **Nagar, D. K.** (2024). AI-based language translation and interpretation services: Improving accessibility for visually impaired students. In *Transforming learning: The power of educational* (p. 178).

Lee, T. E. (2024). Exploring naturalness in MOOC video lecture MT subtitles: A case study. *TransLogos Translation Studies Journal, 7*(1), 41–61. https://doi.org/10.29228/transLogos.64

Liu, J., Li, S., & **Dong, Q.** (2024). Collaboration with generative artificial intelligence: An exploratory study based on learning analytics. *Journal of Educational Computing Research*. Advance online publication. https://doi.org/10.1177/07356331241242441

Meyer, A., Rose, D. H., & **Gordon, D.** (2014). *Universal Design for Learning: Theory and practice*. CASTProfessional Publishing. http://udltheorypractice.cast.org/

Mrayhi, S., et al. (2024). Designing future education for all: Principles and frameworks. In *Envisioning the future of education through design* (pp. 147–177). Springer. https://doi.org/10.1007/978-981-97-0076-9_7

Mrayhi, S., et al. (2025a). A generative AI-powered chatbot for enhancing accessibility and personalized learning in MOOCs. https://doi.org/10.1109/ICALT64023.2025.00059

Mrayhi, S., Khribi, M. K., Jemni, M., & **Mrayhi, S.,** et al. (2025b). RECMOOC4ALL: Bridging the accessibility gap in MOOCs. In *2025 IEEE International Conference on Advanced Learning Technologies (ICALT)* (pp. 196–198). https://doi.org/10.1109/ICALT64023.2025.00061

Nataraj, V., Liao, W.-H., Chang, Y.-S., Chiang, C.-Y., Lin, C.-Y., Lin, Y.-A., & **Day, M.-Y.** (2024). Generative AI in multimodal cross-lingual dialogue system for inclusive communication support. In *2024 IEEE International Conference on Information Reuse and Integration for Data Science (IRI)* (pp. 204–209). https://doi.org/10.1109/IRI62200.2024.00051

Ouyang, Z., Jiang, Y., & **Liu, H.** (2024). The effects of Duolingo, an AI-integrated technology, on EFL learners' willingness to communicate and engagement in online classes. *International Review of Research in Open and Distributed Learning, 25*(3), 97–115. https://doi.org/10.19173/irrodl.v25i3.7677

Rane, N. (2023). Chatbot-enhanced teaching and learning: Implementation strategies, challenges, and the role of ChatGPT in education. https://doi.org/10.2139/ssrn.4603204

Rao, J., & **Lin, J.** (2024). RAMO: Retrieval-augmented generation for enhancing MOOCs recommendations. *arXiv*. https://arxiv.org/abs/2407.04925

Risi, M. (2025). Accessible images: AI-generated alt text in educational publishing: Key criteria for creating alt text to enhance accessibility in educational publishing. *Logos, 36*(1), 50–69. https://doi.org/10.1163/18784712-20240039

Rodriguez-Ascaso, A., Molanes-López, E. M., Pérez-Martín, J., & **Letón, E.** (2024). Performance of students with different accessibility needs and preferences in "Design for All" MOOCs. *PLOS ONE, 19*(3), e0299090. https://doi.org/10.1371/journal.pone.0299090

Sarangam, A. (2024). Enhancing accessibility through inclusive design and artificial intelligence: A comprehensive analysis based on WCAG guidelines. *North American Journal of Engineering Research, 5*(3).

Schmitt, É. (2020). *Big data: An art of decision making*. John Wiley & Sons. https://doi.org/10.1002/9781119777014

Support—Khan Academy. (2025). *Khan Academy support*. Retrieved December 24, 2025, from https://support.khanacademy.org/hc/fr

Tlili, A. (2024). Towards next-generation user interfaces: Chinese perspective of implementing artificial intelligence (AI) to support people with disabilities. *Nafath, 9*(26). https://doi.org/10.54455/MCN2602

Tlili, A., Saqer, K., Salha, S., & **Huang, R.** (2025). Investigating the effect of artificial intelligence in education (AIEd) on learning achievement: A meta-analysis and research synthesis. *Information Development*. https://doi.org/10.1177/02666669241304407

UNESCO, U. I. (2021). *Welcoming learners with disabilities in quality learning environments: A tool to support countries in moving towards inclusive education*. https://unesdoc.unesco.org/ark:/48223/pf0000380256

Wang, H., Huang, T., Zhao, Y., & **Hu, S.** (2023). The impact of dashboard feedback type on learning effectiveness, focusing on learner differences. *Sustainability*, *15*(5), 4474. https://doi.org/10.3390/su15054474

World Health Organization. (2021). *Disability and health*. https://www.who.int/news-room/fact-sheets/detail/disability-and-health

Yarrow, N., Pynnonen, L., Song, C., Bhardwaj, R., & **Spiezio, M.** (2023). *Use of assistive education technologies to support children with visual and hearing difficulties in the East Asia and Pacific region*. https://doi.org/10.1596/40510

TO CITE THIS ARTICLE:
Mrayhi, S., Khribi, M. K., & Jemni, M. (2026). An Accessibility-First Generative AI Chatbot to Support Inclusive MOOC Discovery and Early Onboarding. *Open Praxis*, *18*(3), pp. 441–461. DOI: https://doi.org/10.55982/openpraxis.18.3.1026

Submitted: 25 September 2025
Accepted: 29 December 2025
Published: 04 August 2026