International Research Journal of Education and Technology (IRJEdT)
║ ISSN: 2581-7795, Impact Factor 5.007 ║
║ Peer-Reviewed | Open Access | International Journal ║
║ Volume: 09 | Issue: 09 | Sep- 2026 | www.irjweb.com ║

AI Student Mentorship Platform: Next-Generation Generative AI Frameworks for Personalized Academic Guidance and Course Orchestration

Dr. ANTONY CYNTHIA[1],Srihari.D [2],Roshini.N.S [3],Irfan.S [4]

Assistant professor [1] ,Student of BCA[2,3,4]

Department of Computer Applications, Sri Krishna Arts And Science College, Coimbatore

Abstract

Traditional educational mentorship models face persistent structural limitations, including constrained mentor availability, lack of scalability, geographical boundaries, and high operational costs [1]. The emergence of Large Language Models (LLMs) and Generative Artificial Intelligence (GenAI) offers a paradigm shift toward scalable, personalized, and real-time academic guidance [2]. This paper presents a comprehensive research and implementation blueprint for an AI Student Mentorship Platform—an intelligent, multimodal educational system engineered to assist students through dynamic mentorship, automated course content orchestration, real-time query resolution, and adaptive voice/visual interaction [1, 2].

Synthesizing methodologies from contemporary generative tutoring systems and professional development frameworks, this platform integrates OpenAI's GPT-3.5-Turbo/GPT-4 for natural language dialog and course creation, DALL-E 3 for contextual image generation, and Whisper for automated speech-to-text and text-to-voice interaction [2]. We outline the end-to-end system architecture, prompt engineering guardrails, database structures, and a five-stage implementation methodology (Define, Design, Develop, Deliver, Monitor & Maintain) [1, 2]. Furthermore, we establish a robust evaluation framework spanning quantitative learning indicators and qualitative rapport metrics while providing an in-depth analysis of ethical considerations, including algorithmic bias, data privacy, model hallucinations, and human-in-the-loop governance [1, 2].

Keywords: AI-Powered Mentorship, Generative AI in Education, Large Language Models (LLMs), Natural Language Understanding, Adaptive Learning, Multimodal Course Orchestration, Automated Speech Recognition.

International Research Journal of Education and Technology
(IRJEdT)
‖ ISSN: 2581-7795, Impact Factor 5.007 ‖
‖ Peer-Reviewed | Open Access | International Journal ‖
‖ Volume: 09 | Issue: 09 | Sep- 2026 | www.irjweb.com ‖

1. Introduction

1.1 Context and Evolution of AI in Higher Education

The global higher education landscape has expanded exponentially, with market projections estimating growth from USD 95.19 billion in 2022 to USD 210.06 billion by 2030 at a Compound Annual Growth Rate (CAGR) of 10.4% [2]. Concurrently, modern higher education faces unprecedented demand for personalized, accessible, and high-quality instructional environments [2].

Traditional academic mentoring plays a foundational role in student development [1]. It provides:

- **Skill Enhancement:** Assisting students in identifying strengths, remediating learning gaps, and acquiring domain expertise [1].
- **Career and Academic Guidance:** Steering students through curriculum selection, research opportunities, and long-term career choices [1].
- **Confidence Reinforcement:** Offering psychological safety, constructive feedback, and stress mitigation [1].
- **Work-Life Balance:** Helping learners manage academic workloads, time prioritization, and personal well-being [1].

Despite these clear benefits, conventional human-to-human mentorship schemes suffer from severe scalability bottlenecks [1]. Human faculty and professional mentors encounter strict time constraints, sub-optimal mentor-mentee matching, geographical limitations, and difficulties in systematically monitoring long-term learning outcomes [1].

The advent of Generative AI (GenAI) and Large Language Models (LLMs) since late 2022 offers a viable technological framework to overcome these historical constraints [2]. By leveraging advanced Natural Language Understanding (NLU), speech synthesis, and multimodal generation, AI-powered platforms can deliver continuous, context-aware, and highly scalable academic support [1, 2].

Table 1. Comparative Analysis of Traditional vs. AI-Powered Mentorship Models

Dimension	Traditional Mentorship	AI Student Mentorship Platform
Availability & Access	Constrained scheduling & office hours	24/7/365 real-time ubiquitous access
Geographical Reach	Location-bound / physical presence needed	Global cloud accessibility across devices
Scalability	Difficult to scale linearly per student	Infinite, parallel linear scalability
Feedback Dynamics	Variable and subjective evaluation	Standardized, data-driven, instant path

Cost Structure	High human resource and operational cost	Cost-effective, automated execution

1.2 Statement of the Problem

While modern e-learning systems (e.g., MOOCs, Learning Management Systems) provide access to instructional materials, they frequently lack adaptive support mechanisms, conversational depth, and dynamic content generation [2]. Students operating in asynchronous or remote learning environments often experience isolation, leading to lower course completion rates and unaddressed conceptual misunderstandings [1, 2].

Conversely, existing basic chatbots rely on static decision trees that struggle to grasp complex context, process variable human queries, or maintain long-term memory of a student's learning trajectory [1, 2]. There is a critical need for an integrated, multi-tenant platform that unifies AI mentorship with automated course creation, context-bounded question answering, and multimodal interaction engines [1, 2].

1.3 Research Objectives and Scope

This study formulates an end-to-end framework for an AI Student Mentorship Platform. The primary objectives are:

- **1. System Architecture Design:** Construct a flexible, multi-layered architecture connecting student/educator frontends to backend LLM engines, vector databases, and multi-modal synthesis tools [2].
- **2. Automated Course Orchestration:** Establish an AI-driven workflow enabling educators to dynamically generate curricula, chapter outlines, audio lectures, and illustrative visual aids using LLMs, DALL-E 3, and Whisper [2].
- **3. Context-Bounded AI Mentorship:** Formulate system-level prompt engineering architectures that enforce domain-bounded dialog, preventing hallucinations while serving as a continuous interactive assistant [2].
- **4. Implementation & Evaluation Protocols:** Outline a comprehensive 5-stage deployment methodology and dual-perspective (quantitative/qualitative) evaluation metrics [1].
- **5. Ethical Guardrail Definition:** Analyze data privacy, algorithmic bias, model explainability, and strategies for preserving human-in-the-loop oversight [1].

2. Literature Review and Theoretical Foundations

2.1 Intelligent Tutoring Systems (ITS) and AI in Education (AIED)

Research in Artificial Intelligence in Education (AIED) traces back several decades, progressing from basic Computer-Assisted Instruction (CAI) to dynamic Intelligent Tutoring Systems (ITS) [2]. Historical analyses covering AIED research from 1993 to 2020 emphasize that early systems were limited by expert rules and rigid psychometric Item Response Theory (IRT) models [2].

Table 2. Paradigm Shift in Artificial Intelligence in Education (AIED)

International Research Journal of Education and Technology (IRJEdT)
‖ ISSN: 2581-7795, Impact Factor 5.007 ‖
‖ Peer-Reviewed | Open Access | International Journal ‖
‖ Volume: 09 | Issue: 09 | Sep- 2026 | www.irjweb.com ‖

Era & Technological Paradigm	Core Mechanism	Capabilities & Constraints
1980s: Rule-Based CAI	Static decision trees & predefined branching rules.	Low flexibility; no natural conversation.
2000s: Early ITS & IRT	Adaptive knowledge modeling & score-based pacing.	Limited semantic depth & content creation.

While early ITS platforms successfully adapted learning pace based on test scores, they lacked conversational fluidity, semantic depth, and creative content synthesis [2]. Contemporary machine learning and deep learning methodologies have significantly augmented student modeling, enabling real-time performance tracking and dynamic curriculum adjustment [2].

2.2 Generative AI and Multimodal Large Language Models

The technological baseline shifted dramatically with Transformer architectures, self-attention mechanisms, and large-scale pre-trained models [2]. Tools such as OpenAI's GPT series, Gemini, and open-source alternatives possess unprecedented capabilities in text summarization, reasoning, and context-driven response generation [2].

In educational contexts, LLMs serve a dual purpose [2]:

- **For Educators:** Automating routine instructional tasks, including syllabus design, lesson planning, assignment generation, and visual presentation drafting [2].
- **For Students:** Delivering instant, customized explanations, adapting tone to learner knowledge levels, converting raw text into auditory formats, and supporting interactive dialog [1, 2].

Table 3. Multimodal Generative AI Engine Integration Spectrum

Modality	Core AI Model / Engine	Educational Application
Text & Dialogue	GPT-3.5-Turbo / GPT-4	Mentorship Chat, Chapter Logic, Outline Synthesis
Image & Graphics	DALL-E 3 / Stable Diffusion	Contextual Graphics, Illustrative Charts, Visual Assets
Audio & Speech	OpenAI Whisper API	Text-to-Voice Lecture Audio, Voice Queries
Code & Analytics	Python / Code Interpreter	Interactive Problem Solving, Data Analysis

International Research Journal of Education and Technology (IRJEdT)
‖ ISSN: 2581-7795, Impact Factor 5.007 ‖
‖ Peer-Reviewed | Open Access | International Journal ‖
‖ Volume: 09 | Issue: 09 | Sep- 2026 | www.irjweb.com ‖

2.3 Synthesis of Technical and Human Factors

Recent empirical studies evaluating student Willingness to Accept (WTA) AI-assisted learning environments highlight that adoption is governed by performance expectancy, effort expectancy, and perceived functional or psychological risks [2]. Students demonstrate a strong willingness to utilize AI tools when outputs are accurate, user interfaces are intuitive, and interactions feel natural [2]. However, concerns regarding data privacy, loss of authentic human connection, and algorithmic bias remain significant friction points that demand structural safeguards [1].

3. System Architecture and Conceptual Design

The proposed AI Student Mentorship Platform is engineered as a cloud-native, modular system composed of three main layers: the Interface Layer, the Orchestration & Core Intelligence Engine, and the Data Management Layer [2].

Table 4. Multi-Layer System Architecture Breakdown

Architecture Layer	Sub-System Component	Functional Responsibility
1. User Interface Layer	Student Dashboard & Client UI	Audio playback, text chat, course progress indicators, 'Ask AI Tutor' side-panel.
1. User Interface Layer	Tutor / Educator Console	Subject creation controls, chapter structure setup, automated visual/audio review.
2. Orchestration Engine	Conversational Engine	GPT-3.5-Turbo / GPT-4 for dynamic dialog processing, context parsing, and intent classification.
2. Orchestration Engine	Speech & Visual Engine	Whisper API for speech-to-text/audio synthesis; DALL-E 3 for chapter graphics.
2. Orchestration Engine	Guardrail Engine	Enforces strict chapter boundary prompts and prevents off-topic model drift.

3. Data Layer	Relational DBMS	Stores user accounts, course metadata, activity logs, and interaction analytics.
3. Data Layer	Vector Knowledge Base	Stores vector embeddings of textbooks and references for Retrieval-Augmented Generation (RAG).

3.1 User Interface Layer

The frontend provides dedicated views tailored to student and educator workflows [2]:

- **Student Interface:** Displays active enrollments, course completion rates, interactive lecture screens with voice playback controls, and a context-aware 'Ask AI Tutor' side-panel [2].
- **Educator/Tutor Interface:** Enables subject creation, dynamic chapter structure setup, automated visual thumbnail generation, and content validation controls before publication [2].

3.2 Orchestration and AI Intelligence Engine

At the core of the platform is an AI orchestration pipeline that interfaces with external models via secure API frameworks [2]:

- **1. Conversational Engine (NLU & Reasoning):** Uses GPT-3.5-Turbo or GPT-4 for dynamic dialogue processing, context parsing, and intent classification [1, 2].
- **2. Speech Synthesis Engine:** Integrates OpenAI Whisper API to convert text transcripts into clear, natural voice output for auditory learning [2].
- **3. Visual Engine:** Leverages DALL-E 3 to interpret course chapter themes and render high-resolution illustrative visual content [2].

3.3 Data Management Layer

The data architecture combines relational and vector-based engines:

- **Relational DBMS:** Stores structured user records, course metadata, activity logs, and system metrics [2].
- **Vector Database & Knowledge Base:** Stores mathematical vector embeddings of textbook content, reference PDFs, and lecture scripts, enabling fast semantic retrieval (Retrieval-Augmented Generation - RAG) [2].

4. Implementation Framework and Course Orchestration Workflow

4.1 The Five-Stage Implementation Lifecycle

To ensure seamless institutional integration, platform execution follows a structured five-stage framework: Define, Design, Develop, Deliver, and Monitor & Maintain [1].

International Research Journal of Education and Technology (IRJEdT)
║ ISSN: 2581-7795, Impact Factor 5.007 ║
║ Peer-Reviewed | Open Access | International Journal ║
║ Volume: 09 | Issue: 09 | Sep- 2026 | www.irjweb.com ║

Table 5. Five-Stage Implementation Lifecycle Methodology

Phase	Key Activities & Deliverables
Stage 1: DEFINE	Outline project scope, student needs, institutional KPIs, and key stakeholder boundaries [1].
Stage 2: DESIGN	Configure prompt engineering structures, conversational UX, and backend system architecture [1, 2].
Stage 3: DEVELOP	Curate vector knowledge bases, establish API integrations, and enforce end-to-end security/RBAC [1].
Stage 4: DELIVER	Deploy pilot programs across academic cohorts, collect usability metrics, and refine models [1].
Stage 5: MONITOR	Continuously track response latency, audit model outputs for bias/accuracy, and update guardrails [1].

4.2 Multi-Phase Course Creation Workflow

The platform features an automated course generation lifecycle that pairs educator oversight with AI generation [2]:

- **Phase 1: Subject Request & Initialization:** An educator initiates course creation by providing a topic title (e.g., 'Artificial Intelligence vs. Machine Learning'), specifying chapter count (2–10 chapters), and uploading optional references [2].
- **Phase 2: Automated Server Processing:** The server triggers API calls to the LLM to generate outline and lesson scripts, calls DALL-E 3 for chapter visual graphics, and calls Whisper for audio lecture streams [2].
- **Phase 3: Review and Publishing:** Generated assets are indexed into the DBMS/Vector store. The educator reviews materials, makes edits if needed, and approves the module for distribution [2].

4.3 Prompt Engineering Design and Guardrails

To prevent context leakage, model drift, and off-topic responses during student mentorship sessions, strict system-level prompt boundaries are implemented [2].

System Configuration for Chapter-Bounded Q&A Bot

The following API prompt layout constrains the AI tutor to respond strictly within the scope of assigned chapter text [2]:

International Research Journal of Education and Technology
(IRJEdT)
‖ ISSN: 2581-7795, Impact Factor 5.007 ‖
‖ Peer-Reviewed | Open Access | International Journal ‖
‖ Volume: 09 | Issue: 09 | Sep- 2026 | www.irjweb.com ‖

```
{
  "model": "gpt-3.5-turbo",
  "temperature": 0.2,
  "messages": [
    {
      "role": "system",
      "content": "You are an official AI Tutor in an academic LMS. Users will ask questions
regarding a specific subject chapter. The chapter text is defined as follows: <chapter_text>.
Strict Policy: You must ONLY answer questions directly related to this chapter text. If a user
query is off-topic or unverified by the text, politely decline to answer and guide them back to
the course material."
    },
    {
      "role": "user",
      "content": "<student_query>"
    }
  ]
}
```

Prompt Configuration for Sequential Chapter Generation

To format structured course modules cleanly, the system utilizes deterministic instruct prompts [2]:

```
{
  "model": "gpt-3.5-turbo-instruct-0914",
  "prompt": "You are an automated curriculum design engine. Generate exactly <chapter_count>
chapter titles for a course titled '<course_topic>'. Format output strictly as a numbered list
(1., 2., 3.). Provide concise title strings only. Do not include introductory prose or
explanations.",
  "max_tokens": 250,
  "temperature": 0.3
}
```

5. Quantitative and Qualitative Evaluation Metrics

Evaluating an AI Student Mentorship Platform requires a multi-dimensional strategy that measures operational efficacy, academic outcome improvements, and user satisfaction [1].

5.1 Mathematical Formulations for Quantitative Tracking

1. Skill Acquisition Improvement Index (SAI)

To quantify academic learning gain resulting from platform engagement, pre-assessment score (S_pre) and post-assessment score (S_post) are tracked across user cohorts:

Skill Acquisition Improvement Index (SAI) Formula

$SAI = [(S\_post - S\_pre) / (S\_max - S\_pre)] \times 100$

International Research Journal of Education and Technology (IRJEdT)
║ ISSN: 2581-7795, Impact Factor 5.007 ║
║ Peer-Reviewed | Open Access | International Journal ║
║ Volume: 09 | Issue: 09 | Sep- 2026 | www.irjweb.com ║

Where S_max represents the maximum attainable assessment score, S_pre is the initial benchmark score, and S_post is the score achieved after AI mentorship engagement.

2. Student Engagement Factor (SEF)

Platform engagement intensity is determined by aggregating user session duration (D_i in minutes), login frequency (F_i), and interactive queries submitted (Q_i) over time period T:

Student Engagement Factor (SEF) Formula

$$SEF = \sum [\, \alpha \cdot D_i + \beta \cdot F_i + \gamma \cdot Q_i \,] \quad (\text{for } i = 1 \text{ to } N)$$

Where α, β, γ represent scaling parameters normalizing weights across session metrics.

5.2 Comparative Metric Analysis

Table 6. Comprehensive Platform Evaluation Benchmarks

Metric Domain	Assessment Parameter	Target Benchmark	Measurement Strategy
System Efficacy [1]	Context Adherence Rate	> 98.5%	Automated verification auditing output accuracy against source documents [1, 2].
System Efficacy [2]	Audio Processing Latency	< 1.5 sec	Measuring round-trip time for Whisper API text-to-speech audio streaming [2].

6. Implementation Scenarios and Use Cases

The modular AI Student Mentorship Platform adapts to diverse academic environments [1]:

6.1 Higher Education Institutions

International Research Journal of Education and Technology
(IRJEdT)
‖ ISSN: 2581-7795, Impact Factor 5.007 ‖
‖ Peer-Reviewed | Open Access | International Journal ‖
‖ Volume: 09 | Issue: 09 | Sep- 2026 | www.irjweb.com ‖

Universities can integrate the AI platform within existing learning ecosystems to provide continuous academic support [1]. The platform acts as an automated Teaching Assistant (TA), addressing student queries outside standard faculty office hours, explaining complex concepts, and guiding students through literature review structures [1, 2].

6.2 Distance and Remote Learning Environments

In remote or hybrid environments, student disengagement presents a primary challenge [1]. The platform counters this through dynamic multi-modal content generation [2]. By combining voice lectures, illustrative graphical summaries, and real-time interactive Q&A, remote learners maintain active engagement with course materials [2].

6.3 Professional Skill Development & Upskilling

Beyond traditional academic settings, commercial organizations can leverage the platform for employee onboarding and technical skill acquisition [1]. The platform assesses individual skill profiles, identifies competency gaps, and automatically generates tailored learning modules aligned with institutional goals [1].

7. Ethical Considerations, Safeguards, and Implementation Challenges

Deploying AI systems in educational environments introduces critical ethical, technical, and regulatory considerations that must be actively managed [1, 2].

7.1 Data Privacy, Security, and Governance

Educational platforms handle sensitive student information, including academic records, performance assessments, and dialogue logs [1, 2].

- **Encryption Standards:** Systems must enforce AES-256 bit encryption for data at rest and TLS 1.3 protocol standards for data in transit [1].
- **Regulatory Compliance:** Platform data workflows must adhere to established data protection standards, such as the General Data Protection Regulation (GDPR) and the Family Educational Rights and Privacy Act (FERPA) [1].
- **Data Anonymization:** Personally Identifiable Information (PII) must be stripped prior to transmitting prompt inputs to third-party model APIs [1].

7.2 Mitigating Algorithmic Bias and Ensuring Fairness

Language models trained on broad public web datasets risk inheriting underlying societal biases, unverified assumptions, or exclusionary language [1, 2].

- **Dataset Auditing:** Core reference embeddings must undergo content filtering to ensure diverse, non-discriminatory representation across diverse user populations [1].
- **Fairness Testing:** System validation must systematically test outputs across demographic subgroups to verify equity in response tone, recommendation accuracy, and instructional quality [1].

7.3 Maintaining Human-in-the-Loop (HITL) Governance

A foundational principle of educational AI deployment is that automated systems should enhance and supplement human mentorship rather than replace human faculty [1].

> **CORE GOVERNANCE PRINCIPLE**
> *"Artificial Intelligence in educational environments should serve as a force multiplier for human educators, offloading routine administrative workflows while preserving authentic human empathy, mentorship, and critical ethical leadership." [1]*

- **Editorial Oversight:** Educators maintain final review authority over generated course outlines, visual assets, and lecture scripts prior to student distribution [2].
- **Escalation Triggers:** When the AI engine detects high student distress, repeated conceptual failure, or complex emotional needs via sentiment analysis, the interaction is automatically escalated to human academic advisors [1].

7.4 Model Hallucinations and Explainability

LLMs are inherently probabilistic engines capable of generating plausible yet factually incorrect assertions (hallucinations) [2].

- **Retrieval-Augmented Generation (RAG):** Grounding model generation directly in vetted, institutional knowledge bases restricts response parameters to verified reference texts [2].
- **Citation & Provenance Tracking:** AI mentorship responses must incorporate clear references and source citations pointing back to official textbook chapters or course materials [2].

8. Conclusion and Future Work

8.1 Summary of Contributions

This paper has introduced a holistic framework for an AI Student Mentorship Platform, demonstrating how modern generative technologies can redefine academic guidance and curriculum orchestration [1, 2]. By unifying LLM conversational capabilities with multimodal visual generation (DALL-E 3) and automated speech synthesis (Whisper API), the proposed system offers a dynamic, 24/7 learning environment that bridges traditional mentorship gaps [1, 2].

Key framework contributions include:

- **1. Architectural Integration:** A three-tiered scalable layout linking student frontends to multimodal AI engines and secure storage systems [2].
- **2. Automated Course Generation:** A 3-phase course creation workflow enabling fast content generation paired with educator validation [2].
- **3. Rigorous Guardrails:** System-level prompt designs that bind AI responses strictly to assigned academic chapters [2].

- **4. Evaluation & Ethics Framework:** Dual quantitative/qualitative assessment metrics coupled with clear data privacy, fairness, and human-in-the-loop safeguards [1].

8.2 Directions for Future Research

As generative models evolve, several avenues offer promising opportunities for further system expansion:

- **Real-Time Emotional Intelligence & Avatar Mentors:** Integrating facial emotion recognition and photorealistic embodied virtual avatars to enhance social presence and empathy during mentorship sessions [1, 2].
- **Adaptive Student Knowledge Mapping:** Utilizing dynamic Bayesian Knowledge Tracing (BKT) alongside LLM capabilities to dynamically alter course difficulty based on real-time assessment performance [2].
- **Edge-Deployable Open-Source Models:** Transitioning core intelligence engines toward lightweight, fine-tuned open-source LLMs deployed locally within university infrastructure to maximize data sovereignty and minimize API dependency costs [1, 2].

References

[1] Bagai, R., & Mane, V. (2023). Designing an AI-Powered Mentorship Platform for Professional Development: Opportunities and Challenges. International Journal of Computer Trends and Technology (IJCTT), 71(4), 108–114. https://doi.org/10.14445/22312803/JCTT-V71I4P114

[2] Banjade, S., Patel, H., & Pokhrel, S. (2024). Empowering Education by Developing and Evaluating Generative AI-Powered Tutoring System for Enhanced Student Learning. Journal of Artificial Intelligence and Capsule Networks (JAICN), 6(3), 278–298. https://doi.org/10.36548/jaicn.2024.3.003

[3] Alam, A. (2021). Should robots replace teachers? Mobilisation of AI and learning analytics in education. In Proceedings of the 2021 International Conference on Advances in Computing, Communication, and Control (ICAC3), 1–12. IEEE.

[4] Chine, D. R., Gupta, S., & Koedinger, K. R. (2022). Personalized Learning2: A Human Mentoring and AI Tutoring Platform Ensuring Equity. Proceedings of the Ninth ACM Conference on Learning at Scale.

[5] Köbis, L., & Mehner, C. (2021). Ethical Questions Raised by AI-Supported Mentoring in Higher Education. Frontiers in Artificial Intelligence, 4, Article 624050. https://doi.org/10.3389/frai.2021.624050

[6] McIntosh, T. R., Susnjak, T., Liu, T., Watters, P., & Halgamuge, M. N. (2023). From Google Gemini to OpenAI Q*: A Survey of Reshaping the Generative Artificial Intelligence Research Landscape. arXiv preprint arXiv:2312.10868.

[7] Minn, S. (2022). AI-assisted knowledge assessment techniques for adaptive learning environments. Computers and Education: Artificial Intelligence, 3, Article 100050.

[8] Nye, B. D., et al. (2021). Feasibility and Usability of MentorPal, a Framework for Rapid Development of Virtual Mentors. Journal of Research on Technology in Education, 53(1), 21–43.

[9] Pokhrel, S., Ganesan, S., Akther, T., & Karunarathne, L. (2024). Building Customized Chatbots for Document Summarization and Question Answering using Large Language Models using a Framework with OpenAI, LangChain, and Streamlit. Journal of Information Technology and Digital World, 6(1), 70–86.