

PrepTalk: An AI-Powered Interview Preparation Platform for College Students

Abirami Shankari S¹, Divya K², Kavinila S³, Sangavi G⁴, and Dr. D. Nithya⁵

^{1,2,3,4}UG Students,⁵ Associate Professor, Department of Computer Science and Engineering,

School of Engineering, Avinashilingam Institute for Home Science and Higher Education for Women,
Coimbatore 641108, Tamil Nadu, India

Abstract

The students who are undergoing campus placements still find it a difficult task to prepare for the structured interviews. Many candidates have the technical knowledge but lack clarity, confidence and realistic practice environment. PrepTalk is a comprehensive AI-based interview preparation platform that seeks to integrate mock interview simulation, real-time speech evaluation, communication training, resume analysis, aptitude testing, and company-specific question retrieval in one ecosystem. We use Google Gemini Flash for generative evaluation, the browser-native Web Speech API for speech-to-text transcription and Firecrawl for dynamic question retrieval. A lightweight behavioral proctor module, implemented entirely in client-side JavaScript, that monitors session integrity without the need for software installation. Evaluations during controlled sessions show measurable decreases in filler words and consistent improvements in response structure and clarity. PrepTalk addresses a recognized gap by providing a candidate-centric, vertically integrated preparation environment tailored specifically to the Indian engineering placement context, achieving high module coverage compared to all comparable existing systems.

Key Words — Artificial Intelligence, Interview Preparation, Natural Language Processing, Gemini API, Speech Analysis, Proctoring Systems, ATS Resume Screening, WebSocket, Filler Word Detection.

I. INTRODUCTION

Campus placements are a critical transition point for graduates in India, shaping early career trajectories in a highly competitive graduate labor market. Technical competencies are developed through four years of academic study, but the ability to communicate that knowledge effectively in a time-pressured interview is a separate and often under-developed skill. Surveys of placement coordinators consistently find that the main reasons that candidates underperform in interview

settings are not due to gaps in technical knowledge, but rather clarity of communication, confidence, and structure of responses [1].

The existing ways to prepare peer-based mock interviews, static question banks, and video tutorial libraries all suffer from a fundamental flaw: they deliver little individualized, actionable feedback on how a candidate communicates, not what they say. Students may possess adequate domain knowledge yet struggle with filler word overuse, unstructured responses, and poor confidence signaling deficiencies that go unaddressed by conventional tools.

Recent advances in browser-accessible large language models (LLMs) and speech recognition have provided the technical underpinning for the development of interactive systems that can assess spoken responses in real-time and generate structured, specific feedback. In this paper we present PrepTalk, a platform that leverages these capabilities to provide a full-featured, AI-guided preparation environment that requires no client-side software installation beyond a Chromium browser.

The main contributions of this work are: (1) a modular full-stack architecture unifying eight different preparation components into a coherent user workflow; (2) a Gemini-powered speech evaluation pipeline with multi-dimensional scoring and deterministic fallback; (3) a browser-native behavioral proctoring module with zero installation overhead; (4) dynamic company-specific question retrieval via Firecrawl with intelligent caching; and (5) empirical evaluation demonstrating measurable performance improvement across repeated sessions, including filler word reduction and communication score progression.

II. LITERATURE SURVEY

A systematic review of recent work on AI-assisted interview preparation, automated speech evaluation, and LLM-based educational feedback was conducted across IEEE Xplore, Springer Nature, ScienceDirect,

and ResearchGate. Eight representative works are surveyed in Table I. Key observations follow.

Joel Manuel et al. [1] proposed Q&AI, a conversational AI-powered mock interview bot that generates role-specific questions and tracks candidate performance. The system demonstrates the viability of automated question generation but focuses on answer correctness rather than spoken delivery, and provides no voice-based communication analysis.

The NLP-based interview system in [2] evaluates textual responses using language processing techniques, offering automated scoring of answer content. However, the system operates exclusively on typed text and does not process spoken input, which limits its utility for communication training.

MockMate [3] introduces adaptive question difficulty that adjusts dynamically based on candidate performance, improving question relevance across sessions. Despite this adaptivity, the platform does not assess speech fluency, filler word usage, or confidence delivery in voice responses.

The virtual job interview simulator described in [4] supports video-based interaction and NLP-driven scoring but does not apply audio signal analysis to the video stream, omitting speech fluency and confidence assessment from its evaluation pipeline.

Mandal et al. [5] employed CNN-based facial emotion detection combined with NLP for confidence classification, providing a multimodal assessment approach. While technically comprehensive, the system prioritizes emotional state recognition over communication clarity training, and the neural network inference requirements introduce latency that constrains real-time feedback delivery.

Mishra et al. [6] developed a structured AI-driven virtual mock interview system with automated feedback generation, contributing a well-organized simulation workflow. The system does not incorporate persistent coaching, aptitude preparation, or company-specific question retrieval.

Jadhav et al. [7] presented a mock interview simulator with AI evaluation augmented by pose-based interaction monitoring, enabling assessment of candidate posture alongside verbal responses. The system lacks a persistent coaching module and does not integrate the multi-module preparation workflow required for comprehensive placement readiness.

AlGhamdi et al. [8] demonstrated that prompt-engineered LLMs can generate structured evaluation scores and feedback reliably, providing theoretical grounding for prompt-based evaluation pipelines. Their work is not interview-specific and does not address voice processing or the integrated preparation workflow that PrepTalk implements.

Joel Manuel et al. [1] developed Q&AI, a conversational mock interview bot that demonstrated the viability of AI-driven role-specific question generation. However, the system focuses on answer correctness rather than delivery quality, and lacks voice-based communication analysis a gap PrepTalk directly addresses. Similarly, the NLP-based interview system described in [2] evaluates textual responses but does not process spoken input, limiting its utility for communication training.

MockMate [3] introduces adaptive question difficulty as a function of candidate performance, a mechanism that improves question relevance but still does not assess fluency, filler word usage, or confidence signaling in voice responses. The virtual job interview simulator in [4] supports video-based interaction but does not apply speech analysis to the audio stream.

Mandal et al. [5] employed CNN-based facial emotion detection alongside NLP for confidence classification, offering a multimodal approach to candidate assessment. While technically sophisticated, the system emphasizes emotional state over communication clarity, and the reliance on neural network emotion detection introduces inference latency unsuitable for real-time feedback delivery. Mishra et al. [6] and Jadhav et al. [7] contribute structured AI-driven simulation and pose-based interaction analysis respectively, but neither integrates persistent coaching or aptitude preparation modules.

AlGhamdi et al. [8] demonstrate the effectiveness of prompt-engineered LLMs for generating structured evaluation scores and feedback, providing theoretical grounding for PrepTalk's Gemini-based evaluation pipeline. Their work does not address interview-specific voice processing or the multi-module preparation workflow that PrepTalk implements.

The surveyed literature collectively reveals a consistent pattern: existing systems address isolated dimensions of interview preparation question generation, emotion detection, ATS screening, or adaptive difficulty without integrating these components into a unified, candidate-centric platform. PrepTalk's primary contribution is this vertical integration, combined with voice-based

communication evaluation and persistent AI coaching, within a zero-installation, freely accessible deployment.

III. OBJECTIVES

Several engineering students and students looking for their first job feel a significant gap with how much they know and how well they can explain what they know during formal interviews.

Due to the nature of the preparation systems presented to students prior to job interviews, further widening the gap, many of these systems include fixed question banks and basic mock interviews that only evaluate the content of the student's answers without considering the quality of how the student communicates through elements such as clarity, confidence, fluency, and structure.

PrepTalk will address this gap by offering an AI powered, voice-based platform to train students in simulated interview settings including technical, behavioral, and HR categories; Natural Language Processing (NLP) to analyze how the student answered questions and evaluate the quality of communication through multiple dimensions; generated structured and actionable feedback reports after every session; track progress through repeated sessions by maintaining continued scoring history; and offer on-demand AI coaching for skill development on an as-needed basis, while being completely free, completely browser-based, and requiring no institutional infrastructure.

IV. SYSTEM ARCHITECTURE

A. Architectural Overview

PrepTalk follows a monolithic-but-modular three-tier client-server architecture. The frontend tier consists of static HTML pages served by the FastAPI application, with all interactivity implemented in JavaScript using the Fetch API, Web Speech API, and event-driven proctor listeners. A shared navigation script injected on every page maintains consistent layout. The application tier exposes eight API prefixes managed by dedicated router modules: /api/sessions, /api/speech, /api/comm, /api/ats, /api/practice, /api/coach, /api/auth, and /api for dashboard aggregation. The persistence tier uses SQLite, initialized at startup with eight tables covering all system entities.

The monolithic deployment decision reflects the academic context: portability, zero external service dependencies, and operational simplicity take priority over horizontal scalability. FastAPI's router-based organization provides equivalent logical module separation to microservices without orchestration


overhead. All generative AI functionality routes through a centralized client instance in backend/utils/llm.py, ensuring consistent model configuration and simplifying API key management.

B. AI Engine

For speech evaluation and generating interview questions, we will be using Gemini 2.5 Flash as our evaluation pipeline, and we will be using Gemini 2.0 Flash for evaluating the communication hub, analysis of resumes through ATS, dialogue interactions through AI coaching, and the generation of fallback questions for PYQs (past year questions). Each prompt will consist of a persona definition to ground the model's base point of reference for an expert in that role, followed by the input data that will be separated into sections by delimiters, and lastly we will include an explicit output schema in JSON format at the end of each evaluation prompt by including the following statement in all evaluation prompts: "Return ONLY a valid JSON object (no markdown or prose)" and will then process post-processing to remove any remaining code fence artifacts before parsing..

C. Real-Time Communication via WebSocket

The voice interview module uses WebSocket-based bidirectional communication to deliver evaluation feedback with observed latency of 2.1 seconds from answer submission to feedback receipt. This architecture eliminates the polling overhead of HTTP-based evaluation, enabling a more fluid interview simulation experience. The WebSocket connection lifecycle initialization, message handling, and graceful closure is managed by the frontend interview module, with fallback to standard HTTP evaluation when WebSocket is unavailable



V. METHODOLOGY AND MODULE DESIGN

A. User Authentication and Dashboard

The authentication module handles user registration and login with secure credential storage and session token management via JWT. Following authentication, users access a personalized dashboard displaying activity

history, ATS score trends, and session performance summaries. Resume upload and parsing occur in this module: the system extracts skills, project details, and experience through NLP techniques, generates an ATS compatibility score with actionable improvement suggestions, and uses the extracted profile to personalize subsequent interview question generation.

B. Mock Interview and Communication Evaluation

Session initialization requires role selection (e.g., Software Engineer, Data Scientist, Product Manager) and optional category specification. The backend generates five role-specific questions via Gemini and stores them against the session record. The system supports three interview modes: text-based, voice-based, and video-based. In voice and video modes, speech is transcribed through the browser's SpeechRecognition interface before evaluation. Each transcript passes through a filler word detection function scanning for thirteen predefined terms using boundary-aware substring matching, followed by Gemini evaluation returning a structured JSON object with overall score, clarity, confidence, and relevance subscores, detected strengths, improvement suggestions, and an ideal answer hint.

C. Communication Hub

The Communication Hub provides targeted practice across five communication categories: Self-Introduction, Group Discussion, HR Round, Technical HR, and Email Writing. Each category is bound to a distinct evaluation persona and prompt framing. Self-Introduction prompts evaluate narrative structure and confident summarization; Group Discussion prompts assess argumentation clarity and turn contribution; HR prompts examine self-awareness; Technical HR prompts evaluate integration of technical content with professional delivery; Email prompts assess formality and conciseness.

D. Behavioral Proctor Module

The proctor module runs entirely client-side using JavaScript DOM event listeners. Four violation types are monitored: tab switching (visibilitychange event), window blur (window blur event), copy-paste attempts (copy, cut, and paste events on document), and face absence (periodic canvas-based pixel brightness analysis of the webcam feed). Detected violations trigger asynchronous POST requests to the proctor logging endpoint, recording event type, timestamp, and session context without blocking the interview UI thread.

E. Company PYQ Bank

The Previous Year Question Bank accepts a company name and queries Firecrawl against structured web sources indexed by company. When Firecrawl returns insufficient content, Gemini generates five contextually relevant company-specific questions as a primary fallback, with a static hardcoded set as secondary fallback. Results are cached in the company_questions table with TTL-based invalidation to prevent redundant API calls. The module currently supports 15 or more companies.

F. Resume ATS Analysis


The ATS module evaluates pasted resume text against a target role using a Gemini prompt configured as an expert ATS system and career advisor for the Indian job market. The structured response includes an overall compatibility score (0–100), keyword presence and absence lists, subscores for keyword match, format quality and experience relevance, a feedback summary, and a personalized self-introduction draft generated from the candidate's own resume content.

G. AI Coach

The AI Coach provides a persistent conversational interface for open-ended preparation queries on topics such as STAR method structuring, system design explanation, salary negotiation, and managing interview anxiety. Conversation continuity is maintained by prepending the last twenty messages as formatted history on each API call. The system prompt constrains responses to under 200 words by default to keep interactions practical.

H. Aptitude Practice

The aptitude module provides categorized multiple-choice questions across quantitative reasoning, logical ability, and verbal ability, with timed quiz modes and per-topic performance tracking. This component addresses the written aptitude evaluation stage that precedes most technical interview rounds in Indian campus placement processes.



VI. EXPERIMENTAL RESULTS AND ANALYSIS

A. Module-Level Test Case Evaluation

All ten system modules were evaluated against defined test cases. Table II summarizes the test configurations and outcomes. All modules passed their respective test cases, with the WebSocket evaluation pipeline delivering feedback within 2.1 seconds, well within the 3-second threshold targeted for interactive usability.

All ten system modules were evaluated against defined test cases and passed successfully. The Authentication module correctly issued JWT tokens upon valid registration and login. The Resume Analysis module extracted 12 skills and produced an ATS compatibility score of 78% from a sample PDF upload. The Interview module generated exactly 5 role-specific questions for a Software Engineer configuration. The Evaluation module returned multi-dimensional scores (Clarity 72, Confidence 68, Relevance 65, Overall 71) for a 50-word spoken answer. The AI Coach returned context-aware responses maintaining a 3-message conversation thread. The Aptitude module correctly returned 5 MCQs for a quantitative category request. The Company PYQ module retrieved 10 company-specific questions for a Google query. The WebSocket pipeline delivered evaluation feedback in 2.1 seconds from answer submission, meeting the 3-second interactive usability threshold. The Training Plan module generated 14 personalized tasks for a 2-week plan request. The Analytics module returned a 5-dimensional radar chart for a session feedback query.

B. Filler Word Reduction Across Sessions

In a controlled evaluation across ten simulated sessions with progressively improving answer quality, the filler word detector identified an average of 6.3 fillers per answer in the initial session. By the third session, following candidate review of explicit filler feedback,

this figure reduced to 2.1 per answer. By session five, the average fell below 1.0. This trajectory is consistent with research demonstrating that specific, immediate feedback on filler word usage is among the most efficient mechanisms for speech improvement [1][5].

C. Scoring Dimension Analysis

Clarity and relevance scores exhibited the largest cross-question variation. Technical questions received substantially lower clarity scores when candidates employed vague language compared to responses structured around specific terminology and logical flow. The confidence dimension showed higher variance and is best interpreted alongside clarity and relevance rather than in isolation, as it relies on the model's assessment of declarative versus tentative phrasing in the transcript.

D. ATS Score Correlation

An informal comparison using five anonymized resume samples found overall score differences between PrepTalk's Gemini-based ATS analysis and a commercial ATS tool ranging from 3 to 12 percentage points. Keyword presence and absence lists showed substantial agreement between systems, with PrepTalk tending toward higher experience relevance scores but stricter keyword matching.

E. Proctor System Accuracy

All four violation types — tab switching, window blur, copy-paste attempts, and face absence — were detected and logged successfully on every deliberate trial. No false positives were observed in normal usage conditions. The violation counter and event log in the session feedback report accurately reflected the sequence and timing of all triggered violations.

VII. COMPARATIVE ANALYSIS

Comparison of PrepTalk against four representative systems identified in the literature: Q&AI [1], Prepmania, MockMate [3], and Final Round AI. The comparison criteria reflect the full scope of placement preparation activities relevant to Indian engineering candidates.

Across nine evaluation criteria, PrepTalk outperforms all four compared systems. On question personalisation, PrepTalk and Final Round AI both achieve a Very High rating, while Q&AI, Prepmania, and MockMate achieve High. Module coverage is where the gap is most pronounced: Q&AI covers approximately 30% of the placement preparation workflow, Prepmania and MockMate each cover 40%, Final Round AI covers 50%, and PrepTalk covers 95%. All platforms except

Q&AI support voice-based interviews; PrepTalk additionally implements WebSocket-based delivery for lower-latency feedback.

The most significant differentiators are exclusive to PrepTalk. Aptitude practice across three categories (quantitative, logical, verbal) is absent from all four compared systems. Company-specific previous year questions covering 15 or more companies are similarly unavailable elsewhere. A persistent AI coaching interface, an AI-generated weekly training plan, and 5-dimensional radar chart feedback are features not present in Q&AI, Prepmania, MockMate, or Final Round AI. PrepTalk is also the only free platform in this comparison: Final Round AI charges students ₹700 or more per month, while the other three are free but offer substantially narrower functionality.

PrepTalk achieves 95% module coverage compared to 30–50% in comparable systems, and is the only platform in this comparison that provides aptitude practice, company-specific PYQs, a persistent AI coaching interface, an AI-generated training plan, and 5-dimensional radar chart feedback — all at no cost to the student. The combination of WebSocket-based voice evaluation, persistent coaching, and closed-loop progress tracking across modules creates a qualitatively different preparation experience compared to tools that address only isolated preparation activities.

VIII. LEARNING OUTCOMES

Development of PrepTalk produced nine substantive engineering learning outcomes spanning system design, AI integration, and research methodology.

LO1 — Full-Stack Web Development: End-to-end experience designing and building a production-grade web application from database schema to frontend interaction.

LO2 — REST API Design: Mastery of RESTful design principles including resource naming, HTTP status code semantics, and structured error handling across eight API modules.

LO3 — Real-Time WebSocket Communication: Implementation of WebSocket-based bidirectional communication from scratch for the voice interview module, including connection lifecycle management and graceful fallback.

LO4 — Prompt Engineering for Structured AI Outputs: Practical experience designing prompts that reliably extract structured JSON from LLMs across eight

distinct evaluation contexts, including schema specification, artifact stripping, and multi-provider abstraction.

LO5 — Web Crawling and Data Pipeline Engineering: Construction of a production-grade web crawler with rate limiting, retry logic, intelligent HTML filtering, AI fallback on scraping failure, database caching with TTL, and scheduled execution.

LO6 — Database Engineering: Design of a normalized 7-table relational schema with indexed conversation history retrieval and abstracted connection management supporting future database migration.

LO7 — Research Methodology: Systematic literature survey across IEEE Xplore, Springer Nature, ScienceDirect, and ResearchGate; structured gap analysis; and evidence-grounded feature justification connecting each system component to identified literature deficiencies.

LO8 — System Architecture Design: Design of a scalable 3-tier client-server architecture with clear separation of concerns, decomposing a 6-module complex system into independent, testable components sharing a unified data layer and AI service abstraction.

LO9 — AI Ethics and Responsible Development: Examination of ethical implications of AI-based performance evaluation, including scoring bias risk, the importance of human oversight in assessment, and the design responsibility to provide constructive rather than discouraging feedback.

IX. FUTURE SCOPE

A. Short-Term Enhancements (0–6 Months)

F1 — Video Emotion Analysis: Integration of OpenCV and MediaPipe for real-time facial expression and posture analysis, providing non-verbal communication feedback as a complement to the current transcript-based evaluation.

F2 — Live Coding Interview Module: Addition of an in-browser code editor supporting real-time technical interview simulation with AI evaluation of both code correctness and explanation clarity.

F3 — Mobile Application: Development of a React Native cross-platform mobile application to extend accessibility to candidates preparing on smartphones, which represent the primary computing device for many students in tier-2 and tier-3 cities.

B. Medium-Term Enhancements (6–18 Months)

F4 — Regional Language Support: Extension to Tamil, Hindi, Telugu, and Kannada through integration of a more robust ASR model such as OpenAI Whisper, addressing the transcription quality degradation that currently affects non-standard Indian English accents.

F5 — Adaptive Question Difficulty Engine: Implementation of a reinforcement learning-based system that adjusts question difficulty as a function of demonstrated candidate performance across sessions.

F6 — Fine-Tuned Placement-Specific LLM: Fine-tuning of an open-source language model (Mistral 7B or LLaMA 3.1) on a curated dataset of Indian campus placement interactions to improve evaluation relevance beyond general-purpose model capabilities.

C. Long-Term Research Directions (18+ Months)

F7 — Federated Learning for Privacy-Preserving Model Improvement: Aggregation of patterns from student interactions without centralizing sensitive personal data (resumes, spoken answers, personal profiles) on a single server, addressing the privacy implications of large-scale deployment.

F8 — Predictive Placement Analytics: Training of a predictive model using historical session data, aptitude scores, communication metrics, and final placement outcomes to forecast a student's probability of clearing each interview round, enabling targeted pre-placement coaching interventions.

X. CONCLUSION

PrepTalk addresses a well-defined and practically significant gap in the interview preparation ecosystem for engineering students in India. By integrating Gemini-powered speech evaluation, rule-based filler word detection, browser-native behavioral proctoring, dynamic question retrieval via Firecrawl, WebSocket-based real-time feedback, and a persistent AI coaching interface into a single accessible platform, the system provides a preparation environment that exceeds the feature coverage of existing tools across all compared dimensions.

Experimental evaluation confirms measurable improvements in filler word frequency, response clarity, and communication scores across repeated sessions. All ten system modules passed functional test cases, with real-time feedback delivered within 2.1 seconds. The platform achieves 95% module coverage of the complete placement preparation workflow — including aptitude practice, company-specific PYQs, ATS resume

analysis, and AI-generated training plans — at no cost to the student.

Known limitations include the Chromium browser dependency for the Web Speech API, the SQLite scalability ceiling for concurrent multi-user deployment, and the transcript-level evaluation scope that does not capture audio signal properties such as speaking pace and tonal variation. These are defined targets for the outlined future development roadmap. PrepTalk in its current form fulfills its central design objective: providing an honest, AI-powered training partner that is available on demand and delivers specific, actionable feedback calibrated to the demands of modern Indian campus placement.

ACKNOWLEDGEMENT

Our completion of this paper could not have been accomplished without the support of staff members those who are working with us. We are very much thankful to them. For the reference, we refer many articles and research papers of many authors and institutions those are available in online and offline. We offer our sincere appreciation for the learning opportunities provided by those authors and institutions. At last but not least, our heartfelt thanks to all our family members for caring us and for the huge support.

REFERENCES

- [1] J. M. C. J. Joel Manuel, M. Sabi, M. Benson, G. Baburaj, and S. Saritha, "Q&AI: An AI powered mock interview bot for enhancing the performance of aspiring professionals," in Proc. Int. Conf. Recent Advances in Electrical, Electronics, Ubiquitous Communication, and Computational Intelligence (RAEEUCCI), Chennai, India, 2024, pp. 1–5. DOI: 10.1109/RAEEUCCI61380.2024.10547951
- [2] "AI based mock interview system using natural language processing," IEEE Xplore, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/11005032>
- [3] "MockMate: Leveraging AI for adaptive and realistic interview practice," IEEE Xplore, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/11382795>
- [4] "AI powered virtual job interview simulator using natural language processing," IEEE Xplore, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/11013504>

- [5] R. Mandal, P. Lohar, D. Patil, A. Patil, and S. Wagh, "AI-based mock interview evaluator: An emotion and confidence classifier model," in Proc. Int. Conf. Intelligent Systems for Communication, IoT and Security (ICISCoIS), Coimbatore, India, 2023, pp. 521–526. DOI: 10.1109/ICISCoIS56541.2023.10100589
- [6] P. K. Mishra, A. K. Arulappan, I.-H. Ra, T. M. L., G. R. G., and Y.-S. Lee, "AI-driven virtual mock interview development," in Proc. Joint 13th Int. Conf. Soft Computing and Intelligent Systems and 25th Int. Symp. Advanced Intelligent Systems (SCIS&ISIS), Himeji, Japan, 2024, pp. 1–4. DOI: 10.1109/SCISISIS61014.2024.10760210
- [7] B. Jadhav, A. Sawant, A. Shah, P. Vemula, A. Waikar, and S. Yadav, "A comprehensive study and implementation of the mock interview simulator with AI and pose-based interaction," in Proc. 1st Int. Conf. Cognitive, Green and Ubiquitous Computing (IC-CGU), Bhubaneswar, India, 2024, pp. 1–5. DOI: 10.1109/IC-CGU58078.2024.10530717
- [8] E. M. AlGhamdi, Y. Li, D. Gasevic, and G. Chen, "Leveraging prompt-based LLMs for scoring and feedback in educational contexts," ScienceDirect, 2025.
- [9] FastAPI Documentation, "FastAPI: Modern, fast web framework for building APIs with Python 3.6+," 2024. [Online]. Available: <https://fastapi.tiangolo.com/>
- [10] Google AI, "Gemini API Documentation — google.genai Python SDK," 2025. [Online]. Available: <https://ai.google.dev/>
- [11] Firecrawl, "Firecrawl developer documentation — web scraping and extraction API," 2024. [Online]. Available: <https://docs.firecrawl.dev/>
- [12] Mozilla Developer Network, "Web Speech API — SpeechRecognition Interface," 2024. [Online]. Available: <https://developer.mozilla.org/en-US/docs/Web/API/SpeechRecognition>