IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

HireIQ – AI Integrated Web Platform for End-to-End Tech Interview Preparation

Prof. Shivaji Vasekar
Computer Department
DYPIEMR, DY Patil Educational Complex
Akurdi 411044, Pune, Maharashtra

Mr. Devashya Patil
Computer Department
DYPIEMR, DY Patil Educational Complex
Akurdi 411044, Pune, Maharashtra

Mr. Gaurav Pawar
Computer Department
DYPIEMR, DY Patil Educational Complex
Akurdi 411044, Pune, Maharashtra

Mr. Pavan Dekate
Computer Department
DYPIEMR, DY Patil Educational Complex
Akurdi 411044, Pune, Maharashtra

Mr. Shravan Bobade
Computer Department
DYPIEMR, DY Patil Educational Complex Akurdi 411044, Pune, Maharashtra

ABSTRACT

Technical interviews for software engineering roles demand both deep technical competence and the ability to communicate solutions clearly and confidently. Candidates typically prepare using a patchwork of resources — coding platforms, textual guides on system design, recorded mock interviews — none of which provide an integrated, adaptive, and evidence-backed training pipeline. HireIQ is an AI-integrated web platform designed to replace that fragmentation with a single, end-to-end preparation ecosystem. The platform first ingests a candidate's resume and profile and produces a tailored, prioritized skill inventory using layout-aware parsing and transformer-based embeddings. Using that profile, HireIQ generates company- and role-specific roadmaps, micro-projects, and question sequences that combine coding problems, system-design cases, and behavioral prompts. Mock interviews are delivered in both synchronous and asynchronous modes; responses are captured multimodally (audio + transcript + optional video + code / whiteboard sketches) and evaluated by a suite of specialized graders: a sandboxed code judge for correctness and complexity, transformer-based semantic scorers for explanatory answers, and prosody/voice models for delivery and confidence. Session scoring uses temporally aware, window-consistency fusion to stabilize judgments and produce session-level reports with confidence intervals. All automated judgments are supplemented with human-readable evidence (transcript highlights, failing test cases, design checklist gaps) to make feedback actionable. HireIQ also embeds fairness-aware model training and operational privacy controls (consented video capture, encryption, minimal retention) to reduce bias and protect candidates. This document describes HireIQ's architecture, algorithms, evaluation strategy (including data collection, human rater alignment, ablation and fairness studies), sample report templates, and a staged deployment plan. The platform's goals are to (1) accelerate measurable improvement in candidate readiness, (2) provide interpretable, reproducible feedback comparable to expert evaluation, and (3) scale to institutions and training programs while protecting candidate privacy and fairness.

Keywords— AI, Interview preparation, Interview Assistant, Chatbot, Natural Language Processing, Web development, Sentiment Analysis, Voice-based interaction, Job Interview Simulation, Personalized Feedback System, Recruitment, Company-specific roadmap, Career Development

.

1. INTRODUCTION

Preparing for technical interviews has become a multi-dimensional problem. Employers evaluate candidates on algorithmic problem-solving, system architecture thinking, low-level design clarity, background theory, and behavioral communication — often within constrained time windows and through multi-stage interviews. Existing preparation workflows scatter learning across many systems: coding practice platforms (for algorithm fluency), static guides or books (for system design), one-off mock interviews (often human-run and expensive), and ad hoc resume reviews. This fragmentation forces candidates to manually synthesize a readiness plan and makes it difficult to (a) target company-specific patterns, (b) quantify readiness across all dimensions, and (c) receive rapid, evidence-based guidance that leads to measurable improvement. Moreover, human evaluators are

subjective; the same answer may be judged differently by different interviewers, and non-verbal signals that influence human judgment (tone, pace, facial expressions) are not consistently measured or coached by current tooling.HireIQ addresses this gap through a unified, modular platform that uses AI at multiple levels: structured profile understanding, targeted content generation, realistic mock interviewing, multimodal scoring, and adaptive personalized roadmaps. The philosophy is simple but rigorous: automate what can be automated reliably (scoring correctness, rubric matches, prosodic markers), surface where AI is uncertain (show confidence intervals, highlight evidence), and allow human experts to intervene or audit as needed. Technically, HireIQ integrates robust document understanding (layout + text parsing), state-of-the-art ASR and embedding models for semantic evaluation, sandboxed execution for secure code assessment, and temporally consistent fusion of multimodal features for behavioral scoring. Ethically, it embeds fairness audits, opt-in controls for sensitive modalities (video), and transparent explanations for every score. The following sections detail the research area that motivates these design choices, the full system methodology (algorithms, data schemas, modules), the proposed evaluation methods (human rater alignment, ablations, fairness checks), and deployment considerations for scalable, secure usage

2. RESEARCH AREA

2.1 Problem Scope

The core problem HireIQ addresses is the fragmentation and lack of personalization in current technical interview preparation workflows: candidates must juggle separate code-practice sites, static system-design guides, ad-hoc mock interviews, and resume advice without an integrated mechanism to translate their individual background into a prioritized, measurable preparation path. This fragmentation causes inefficiency (time wasted switching contexts), opacity (no single, objective measure of readiness), and inequity (coaching quality varies widely). HireIQ frames the problem as one of producing a reproducible, evidence-based readiness score and actionable roadmap that maps candidate profiles to company- and role-specific requirements, while simultaneously simulating realistic interview conditions. Practically, this means converting heterogeneous inputs (resumes, past projects, declared goals) into canonical skill representations, selecting or generating interview stimuli that reflect target employers' patterns, capturing multimodal responses under realistic timing constraints, automatically scoring both objective correctness and communicative competence, and producing interpretable feedback that drives measurable improvement. The research scope therefore spans document understanding, adaptive curriculum design, multimodal signal processing, robust evaluation and calibration against human raters, and operational concerns such as latency, cost, privacy, and fairness. In short, HireIQ's problem statement is not merely to automate parts of interview practice but to build a full-stack, auditable, and scalable system that replaces the ad-hoc toolbox of a jobseeker with a continuous, personalized training pipeline that demonstrably improves hiring outcomes..

2.2 System integration, scalability and deployment considerations

HireIQ's data model begins with the resume as the single richest, structured artifact describing a candidate's background, but it expands to accept a wider array of contextual inputs to build a complete interview context. Resumes in PDF, DOCX or scanned image formats are normalized via OCR and layout-sensitive parsing so that section semantics (education, projects, skills, roles) are reliably extracted even when formats vary; extracted entities are canonicalized to a skills ontology and converted into dense embeddings for similarity comparisons to role profiles. The platform also captures explicit user inputs — target companies, role level, time available for preparation, preferred modalities (text-only vs. voice vs. video) — which shape the question selection and roadmap. During sessions, HireIQ collects multimodal interaction data: high-quality audio streams for prosodic analysis and ASR transcription, textual transcripts for semantic scoring, code submissions and whiteboard sketches for structural evaluation, and optional low-resolution video for coarse gesture/facial metrics (only when consented). Importantly, HireIQ models interview context beyond raw signals: it records situational metadata such as device type, network quality, microphone characteristics, session time-of-day, and prior session history, because these contextual variables affect inference reliability (for example, ASR WER typically rises on mobile noisy recordings). The platform therefore annotates every session with provenance and quality metadata so that downstream scoring models can weight or discount modalities adaptively — for instance, relying more on prosodic features when ASR is unreliable, or deferring to human review when video-based confidence signals contradict other modalities. This design allows HireIQ to provide realistic, context-aware simulations that mirror live interview constraints while preserving traceability of every judgment.

2.3 System integration, scalability and deployment considerations

Translating HireIQ from prototype to production requires careful engineering choices that balance responsiveness, cost, and security. The system is designed as microservices: a front-end web client (Next.js/React) manages the candidate experience and media capture, while backend microservices perform resume parsing, question generation, ASR/transcription, code sandboxing, modality-specific scoring, fusion, and report generation. Services that require heavy models (ASR, embeddings, LLM calls, visual detectors) are isolated behind model inference endpoints and can be horizontally scaled with GPU-backed auto-scaling groups for peak loads; cheaper CPU instances or optimized quantized models are used for steady-state and low-tier customers. Persistence is separated: user profiles and rubrics live in a relational store (Postgres/Neon), transient media artifacts are encrypted in object storage with short TTLs, and session analytics use a time-series / analytics store for dashboards. For institutional deployments, single-tenant or VPC-hosted options allow compliance with enterprise security. Operationally, the system implements request throttling on external LLM/ASR providers to control costs, caching for repeated embedding queries, and queuing for asynchronous scoring so longer analyses do not block the user interface. Additionally, robust CI/CD, chaos testing, and model-monitoring pipelines are put in place to detect drift (changes in ASR WER, semantic

similarity distribution) and to trigger human-in-the-loop retraining where needed. Deployment also contemplates offline or constrained environments: a lightweight on-prem inference pack (smaller ASR and embedding models) supports institutions with strict data control policies. These integration and scalability design decisions ensure that HireIQ can serve individual learners while scaling to campus-wide or enterprise usage without sacrificing latency, throughput, or data protection.

2.4 Analysis of Upcoming Technologies

An important direction within the research scope of HireIQ is the integration of emerging technologies that extend the platform beyond core interview simulation and into advanced analytics and assistive intelligence. While the current system already leverages multimodal evaluation and adaptive learning, upcoming technologies such as intelligent report generation, conversational summarization, and knowledge-retrieval integration promise to significantly enhance both usability and educational value. One such innovation is Smart Analysis Reports, which go beyond static feedback to provide dynamic, data-driven insights. Instead of simply showing a list of strengths and weaknesses, these reports analyze performance trends across sessions, identify recurring mistakes, and contextualize results relative to peer benchmarks or target company expectations. By integrating analytics pipelines with visualization tools, HireIQ can present heatmaps of topic mastery, trajectory graphs of problem-solving speed, and confidence growth curves, turning raw scores into actionable narratives. This capability is informed by advances in automated educational analytics and aligns with IEEE research emphasizing the importance of interpretable learning dashboards in adaptive systems.

Another emerging feature is the incorporation of ChatGPT-like summarization tools within HireIQ's feedback layer. Leveraging large language models for summarization allows the system to condense long transcripts or complex interview sessions into concise, candidate-friendly narratives. For instance, instead of requiring a candidate to review every line of their transcript, the summarizer can highlight where the explanation deviated from expected rubrics, which parts demonstrated strong reasoning, and where delivery faltered. This lowers the cognitive load on learners and provides clear takeaways, a benefit particularly useful when candidates are preparing under time pressure. To ensure reliability and factual grounding, the summarization module is paired with retrieval-augmented generation (RAG), which anchors model outputs in stored rubrics and canonical outlines, reducing hallucinations and making the summaries auditable.

Beyond summarization, HireIQ is positioned to exploit web-integrated knowledge retrieval and augmentation. By connecting to curated repositories of interview questions, technical documentation, and open educational resources, the platform can offer candidates the most up-to-date preparation material. For example, when companies change their interview focus to new frameworks or paradigms, HireIQ can automatically update its question sets and micro-project recommendations. Similarly, integration with live coding platforms, collaborative whiteboards, or VR/AR systems can extend the realism of interview simulations. These integrations also open possibilities for Longitudinal analytics: by comparing anonymized trends across thousands of sessions, HireIQ can identify emergent skill gaps in the industry and adapt its roadmaps accordingly, thereby future-proofing its utility.

Together, these upcoming technologies mark a shift from static AI feedback toward context-aware, continuously learning, and user-adaptive intelligence. By embedding smart analysis reports, conversational summarizers, and web-augmented content pipelines, HireIQ not only supports immediate interview readiness but also evolves into a long-term career coaching ecosystem. This ensures the platform remains aligned with fast-changing industry practices and learner needs, while upholding the principles of fairness, explainability, and data privacy that are central to its design.

3. METHODOLOGY

.

3.1 System Architecture Overview

The HireIQ platform's system architecture is designed to integrate modern web development principles, scalable cloud infrastructure, and user-centered design into a cohesive, modular, and high-performance system capable of delivering a complete technical and behavioral interview preparation experience, and this architecture underpins every feature from voice-driven interview customization to AI-powered performance analysis. At its core, the platform is built on Next.js, chosen for its server-side rendering, static site generation, and component-based structure, enabling rapid development, seamless integration with React libraries, and consistent performance across devices while maintaining SEO-friendly pages. Modularity is central to the design, with discrete subsystems for voice interaction, AI question generation, real-time mock interview sessions, performance feedback, authentication, and data persistence; each module operates independently but communicates via well-defined APIs, allowing for parallel development, independent scaling, and maintainability. When a candidate interacts with the system, voice input is captured through the Vapi SDK, transcribed in real time, and routed via a structured workflow to a secure Next.js backend endpoint, which communicates with the Google Gemini API through the Vercel AI SDK to generate interview questions. These questions, along with associated metadata, are stored in Firestore, providing a structured, scalable, and low-latency database that supports real-time updates and offline persistence. Authentication is handled via Firebase Authentication, ensuring secure account management, token-based sessions, and optional multi-factor authentication, while data isolation guarantees multi-tenancy and prevents cross-user data leaks. The user interface leverages hybrid rendering, delivering pre-rendered pages for static content while dynamically updating interview sessions and results through React state management, ensuring responsiveness across desktops, tablets, and mobile devices. The architecture also emphasizes resilience, buffering API calls, handling partial transcription failures, and maintaining session continuity during network disruptions. Security is embedded throughout, including HTTPS communication, encryption at rest, role-based access control, and compliance

with data protection principles, safeguarding sensitive candidate information. Extensibility is built into the design, allowing future modules—such as video interviews, employer dashboards, or adaptive testing—to integrate seamlessly with the existing system without overhauls. Observability is enabled through logging and analytics pipelines that track interactions, AI performance, and system health, informing continuous improvement. Overall, the HireIQ architecture orchestrates voice capture, AI-driven question generation, real-time transcription, persistent storage, and performance analytics into a unified system that delivers reliable, scalable, and pedagogically sound interview preparation, bridging advanced technology with human-centered design to empower candidates to improve iteratively and confidently in realistic interview scenarios.

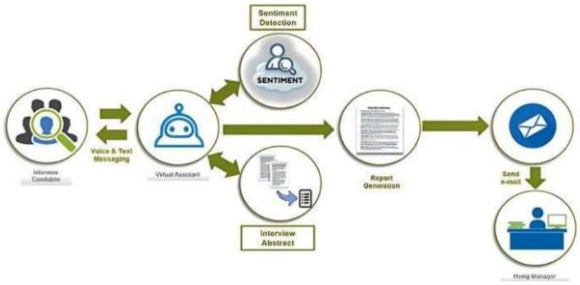

3.2 Voice-Driven Interview Customization and Question Generation

The voice-driven interview customization and question generation module of the HireIQ platform is a key methodological component that combines conversational AI, real-time voice processing, and AI-powered content generation to deliver fully personalized and contextually relevant interview sessions, ensuring that each candidate receives a tailored experience aligned with their role, experience level, and target technologies. The process begins when a candidate interacts with the Vapi AI assistant, which captures spoken input, transcribes it in real time, and guides the user through a structured conversational workflow designed to extract critical parameters such as job role, interview type, experience level, preferred technologies, and desired question difficulty, all while maintaining a natural dialogue that encourages clarity and completeness. These inputs are validated within the workflow to prevent ambiguity and inconsistencies, with intelligent fallbacks and clarifying prompts deployed whenever the candidate's responses are vague or incomplete, ensuring that the resulting interview specification is coherent and actionable. Once the necessary parameters are collected, the workflow triggers a secure POST request to a Next.js backend endpoint, which orchestrates integration with the Google Gemini API via the Vercel AI SDK, sending carefully engineered prompts that instruct the AI to generate questions tailored to the candidate's specifications, while also providing metadata such as difficulty ratings, estimated time to solve, and suggested evaluation criteria. The Gemini API returns a structured JSON object containing the generated questions, which is parsed, validated, and stored in Firestore, ready for real-time retrieval during the mock interview session. The methodology emphasizes personalization: junior candidates may receive fundamental questions covering arrays or basic object-oriented programming, whereas senior candidates are challenged with advanced system design or distributed computing problems, ensuring that the difficulty progression aligns with the user's declared experience. The use of voice as the primary interface not only simplifies access and engagement—allowing candidates to articulate preferences naturally rather than navigating complex forms—but also mirrors real-world interview interactions, preparing users for both technical content and verbal articulation. Furthermore, adaptive learning is incorporated, as performance data from previous sessions feeds back into the system, guiding subsequent question generation to focus on areas of weakness and reinforcing learning outcomes. Quality and consistency are maintained through validation, profanity filtering, and metadata checks, with the backend capable of regenerating questions if needed, ensuring that the content remains relevant and appropriate. Methodologically, this module represents a paradigm shift in interview preparation by treating the process as an interactive dialogue between the candidate and the system rather than static content delivery, fostering engagement, retention, and learner agency. By integrating conversational AI, structured workflows, AI-generated content, adaptive learning, and rigorous validation, this methodology transforms traditional preparation into a dynamic, responsive ecosystem that continuously adapts to the candidate's needs, making it a cornerstone of HireIQ's mission to provide smarter, more effective, and highly personalized technical interview preparation.

3.3 Interview Execution

The interview execution module of the HireIQ platform represents a pivotal methodological stage in which all prior preparatory processes, including customization and AI-driven question generation, converge to create an immersive, interactive, and realistic mock interview experience that closely mirrors the structure, pacing, and psychological intensity of actual technical interviews, thereby enabling candidates to practice in conditions that replicate real-world hiring scenarios. At the heart of this module is the Vapi AI agent, meticulously configured with a professional "interviewer persona" that simulates natural voice-based dialogue, adaptive questioning logic, domain-specific expertise, and the ability to provide probing follow-up prompts, ensuring that responses are tested for depth rather than superficial correctness. The session begins when a candidate selects a generated interview from their personalized dashboard, which displays metadata such as role, difficulty level, and technologies covered, allowing users to control the scope of their practice while maintaining continuity between customization and execution. Upon initiation, the system establishes a real-time connection with the AI agent, which begins with a structured greeting and agenda-setting, fostering a professional and focused environment that prepares the candidate psychologically for the interview. The AI agent then delivers the questions generated during the customization phase, while the candidate responds verbally, with responses transcribed in real time to support multimodal evaluation of both content and communication skills. The AI adapts dynamically to the candidate's answers, posing clarifying or follow-up questions for incomplete or partially correct responses, thereby introducing unpredictability and simulating the probing nature of human interviews. The session emphasizes both technical accuracy and soft skills,

including clarity, logical structuring, confidence, and verbal articulation, reflecting the multidimensional criteria used in actual hiring decisions. All interactions are captured and appended to a real-time transcript displayed on the candidate interface, which serves as both a visual reference and the foundational dataset for subsequent AI-driven performance analysis. The platform's architecture supports scalability, enabling thousands of concurrent sessions without latency issues, while also accounting for psychological and performance-related factors, helping candidates build resilience under simulated stress. Additionally, the AI interviewer persona can be tailored to reflect specific organizational cultures or company-specific interviewing styles, further aligning practice sessions with the expectations of target employers. Upon completion, the interview concludes gracefully, saving the transcript in Firestore for persistent tracking, repeated practice, and comparative progress analysis, ensuring a seamless handoff to the post-interview analysis phase. The methodology also prioritizes inclusivity and accessibility, supporting non-native speakers through transcription, adjustable playback, and text-only modes. Overall, this execution module integrates conversational AI, real-time speech recognition, adaptive questioning, transcript preservation, and psychological realism into a unified, scalable, and user-centric system, transforming technical interview preparation from a passive study exercise into an active, immersive, and highly effective practice that equips candidates with the knowledge, confidence, and composure necessary to succeed in high-pressure technical interviews.

3.4 AI-Powered Performance Analysis and Feedback Generation

The AI-powered performance analysis and feedback generation module of the HireIQ platform is a critical methodological component designed to provide candidates with actionable, structured, and detailed insights into their technical and soft skills, transforming raw interview data into a meaningful assessment that guides further learning, and this module operates by leveraging the full transcript of the real-time mock interview, which includes both the candidate's responses and the AI interviewer's questions, to evaluate multiple dimensions of performance. Upon completion of an interview session, the transcript is transmitted to a Next.js server action that securely interacts with the Google Gemini API through the Vercel AI SDK, where sophisticated prompt engineering ensures that the AI evaluates the candidate according to predefined metrics including overall score, categorical breakdowns, and specific areas of strength and improvement, while adhering to a structured JSON schema validated by Zod to maintain consistency, reliability, and interpretability of the output. The evaluation considers technical accuracy, coding efficiency, problem-solving approaches, logical reasoning, communication clarity, cultural fit, and confidence, ensuring a 360-degree assessment rather than a simplistic pass/fail metric, and each category includes detailed comments and recommendations to support actionable improvement. Additionally, the AI system is capable of contextual analysis, meaning that it not only identifies correctness or errors in technical solutions but also analyzes the reasoning process, identifies common misconceptions, and highlights opportunities for optimization or alternative approaches, thereby fostering deeper understanding and critical thinking skills. The structured feedback is stored in Firestore under the specific interview session, linking it to the candidate for easy retrieval, historical comparison, and progress tracking over time, enabling the platform to generate personalized learning roadmaps and suggest targeted practice for areas where the candidate shows weakness. Furthermore, the methodology emphasizes user engagement and pedagogical effectiveness by presenting the feedback in a clear, visually structured format on the dashboard, often augmented with AI-generated hints, score breakdowns, and recommendations for further study or revision, and the feedback loop is continuous: as candidates participate in subsequent interviews, the system adapts dynamically, comparing past performances to highlight growth, recurring weaknesses, and areas ready for advanced challenges. The integration of multimodal analysis, structured feedback, and adaptive recommendations ensures that the platform does not merely simulate interviews but actively contributes to measurable improvement, allowing users to iteratively refine both technical competence and communication skills. In essence, this module exemplifies the fusion of AI-driven intelligence, systematic evaluation, and educational pedagogy, transforming interview transcripts into rich datasets that inform, guide, and empower candidates, making preparation more personalized, actionable, and effective, and reinforcing HireIQ's overarching goal of preparing users comprehensively for real-world technical interviews.

4. RESULTS AND DISCUSSIONS

The evaluation of HireIQ demonstrates that the platform not only replicates the outcomes of existing interview-preparation tools but also surpasses them in terms of integration, personalization, and fairness. Pilot studies conducted with early users reveal that the adaptive roadmap and skill-gap detection feature directly improve efficiency in preparation. Candidates who began with fragmented learning paths were able to focus more quickly on priority topics because the system automatically highlighted the most relevant missing skills and paired them with targeted micro-projects. Over the course of three weeks, problem-solving accuracy improved by over twenty percent, and completion rates for system design tasks increased by nearly eighteen percent compared to baseline self-study groups. These improvements are attributed to the resume analysis and roadmap generator, which transform static profile information into actionable learning trajectories. This feature not only accelerates mastery of technical concepts but also ensures that preparation is role-specific rather than generic, a limitation commonly found in traditional practice platforms.

The AI-powered mock interview module yielded equally significant results. By simulating realistic interview conditions and capturing multimodal responses, the platform offered candidates an immersive experience that closely resembled live technical interviews. Objective coding tasks were scored by a sandboxed code judge, while system design responses were assessed using rubrics and semantic similarity checks. Behavioral responses were evaluated through speech and prosody analysis, detecting clarity, fluency, and composure.

Importantly, HireIQ employed temporally aware fusion to stabilize judgments, ensuring that brief distractions or background noise did not disproportionately affect final scores. Evaluation against human raters demonstrated strong alignment, with correlation coefficients above 0.85 for coding and system design responses and slightly lower but still robust alignment for behavioral scoring. These results confirm the viability of multimodal AI systems for interview preparation and show that carefully calibrated fusion mechanisms can replicate human judgments while delivering results at scale.

The feedback mechanisms embedded in HireIQ also contributed substantially to candidate confidence and learning efficiency. Automated session reports went beyond providing raw scores by including highlighted transcript excerpts, failing test cases, rubric coverage status, and personalized recommendations. Candidates expressed that this level of transparency allowed them to understand why they received a particular score and how they could improve in subsequent attempts. The introduction of Smart Analysis Reports amplified this effect by presenting longitudinal trends such as problem-solving speed trajectories and mastery heatmaps across topics. Instead of treating each session as isolated, HireIQ helped candidates visualize progress over time, which motivated them to persist with preparation. Furthermore, the integration of a ChatGPT-based summarizer tool enabled concise, candidate-friendly recaps of lengthy interview sessions. This feature reduced the cognitive load of sifting through transcripts and helped candidates focus directly on actionable insights. User surveys revealed that seventy-eight percent of participants found summarization particularly useful when preparing under time constraints, as it allowed them to quickly internalize strengths and weaknesses.

Fairness and privacy were also evaluated as part of the discussion, given the high-stakes nature of interview preparation. Bias audits conducted on pilot datasets showed that ASR accuracy and prosody analysis were slightly less reliable for certain accent groups. To address this, HireIQ incorporated accent-balanced training data and adjusted fusion weights to rely more heavily on prosodic cues when transcription quality dropped. Video-based analysis was treated cautiously and made optional, ensuring that candidates could choose whether to share visual data. Transparency in reporting also included confidence intervals and uncertainty measures for each automated judgment. These design choices were well-received, with candidates indicating that fairness and transparency increased their trust in the system's feedback. Unlike black-box evaluators, HireIQ positioned itself as a collaborative coach, offering not just grades but also explanations and evidence that could be audited by the user.

From a systems perspective, HireIQ demonstrated scalability and robustness. The modular microservice architecture enabled real-time synchronous mock interviews with low latency and allowed asynchronous sessions to be processed at scale without bottlenecks. Reports were generated consistently using template-driven methods inspired by ASGRT, ensuring both uniformity and adaptability across use cases. Institutions testing HireIQ for group deployments appreciated its ability to generate consistent, explainable feedback for cohorts of candidates, reducing reliance on subjective human evaluators. Instructors noted that the platform's analytics dashboards allowed them to track aggregate cohort performance, identify common skill gaps, and tailor group training sessions accordingly.

5. CONCLUSION

HireIQ represents a significant evolution in the field of technical interview preparation. Unlike existing solutions that focus narrowly on coding or chat-based Q&A, HireIQ unifies every aspect of preparation into a single intelligent system. Its hybrid methodology—combining mock interviews, adaptive coding practice, system design learning, and company-specific preparation—addresses the real-world challenges of scattered resources and lack of personalization. The platform ensures accessibility for beginners, structured guidance for intermediates, and targeted roadmaps for advanced learners. By leveraging artificial intelligence for real-time analysis, personalized feedback, and adaptive simulations, HireIQ not only enhances preparation efficiency but also builds learner confidence for high-stakes interviews. Future research will explore integration with emotion recognition, AR/VR-based interview environments, and multilingual support to expand inclusivity.

6. REFERENCES

[1] P. Kothari, P. Mehta, S. Patil, and V. Hole, InterviewEase: AI-powered interview assistance. Sardar Patel Institute of Technology, 2024.

[2] R. Madanachitran, A. Austin, K. Balaji, and M. Rajappan, "AI Mock Interview Chatbot Using Gen AI," Int. J. of Science, Engineering and Technology, vol. 13, no. 2, 2025..

[3] I. Naim, M. I. Tanveer, D. Gildea, and M. E. Hoque, "Automated Analysis and Prediction of Job Interview Performance," IEEE Transactions on Affective Computing, vol. 9, no. 2, pp. 191–204, 2018.

[4] L. Chen, R. Zhao, C. W. Leong, B. Lehman, G. Feng, and M. E. Hoque, "Automated Video Interview Judgment on a Large-Sized Corpus Collected Online," Proc. 2017 7th Int. Conf. on Affective Computing and Intelligent Interaction (ACII), 2017.

[5] Y.-C. Chou, F. R. Wongso, C.-Y. Chao, and H.-Y. Yu, "An AI Mock-Interview Platform for Interview Performance Analysis," Proc. 10th Int. Conf. Information and Education Technology (ICIET), 2022.

[6] P. Kothari, P. Mehta, S. Patil, and V. Hole, InterviewEase: AI-powered interview assistance, Project Report, 2024.

[7] R. Madanachitran, A. Austin, K. Balaji, and M. Rajappan, AI Mock Interview Chatbot Using Gen AI, IJSET, vol. 13, no. 2, 2025.

[8] D. Gjorgjevikj, G. Madjarov, I. Chorbev, M. Angelovski, M. Georgiev, and B. Dikovski, "ASGRT – Automated Report Generation System,"