International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 03 | March - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

JobSphere AI: An Integrated AI-Powered Career Guidance Platform with Multi-Modal Interview Coaching

Shruti Parange*, Moinuddin Quazi*, Sanika Zagade*, Sarika Katkar*, Dr. Kishor Sakure†
*Department of Computer Engineering, Terna Engineering College,
University of Mumbai, Navi Mumbai, India
†Project Guide, Department of Computer Engineering,
Terna Engineering College, University of Mumbai, Navi Mumbai, India

***Abstract*—The rapid evolution of the job market has created a significant gap between candidate skillsets and industry expectations, while existing career guidance systems remain fragmented and static. This paper presents JobSphere AI, an integrated multi-modal career guidance platform that combines resume analysis, adaptive skill assessment, AI-driven interview coaching, and personalized career recommendations within a unified architecture. The system leverages a locally deployed large language model (Llama 3.2 via Ollama) for natural language processing, YOLOv8 [11] for real-time proctoring, and DeepFace [12] for emotion recognition, enabling a comprehensive evaluation of both technical and behavioral competencies. A five-layer modular architecture ensures scalability and seamless integration of heterogeneous AI components. Experimental evaluation on real-world user scenarios demonstrates effective skill extraction (87% accuracy), robust proctoring (91% accuracy), and meaningful career recommendations derived from assessment analytics. The proposed system bridges the gap between academic preparation and industry requirements by providing a continuous, data-driven career development pipeline. The results indicate that multi-modal AI integration significantly enhances the reliability and depth of automated career guidance systems.**

***Index Terms*—Artificial Intelligence, Career Guidance, Mock Interview, Resume Parsing, YOLOv8, Llama 3.2, Emotion Detection, Proctoring, Natural Language Processing, Skill Assessment**

I. Introduction

The contemporary job market is characterized by rapid technological evolution, widening skill gaps, and an increasingly complex hiring landscape. While numerous career advisory platforms exist, most provide static content that fails to account for individual skill profiles, real-time industry shifts, and the nuanced demands of modern recruitment processes [1], [4].

Existing solutions typically address isolated aspects of career development. Resume screening tools [2], [8], [14] focus on keyword matching and scoring but lack integration with learning paths. Interview preparation platforms [9], [10] offer simulated interactions but often miss multi-modal behavioral analysis and real-time proctoring [15]. Career recommendation systems [4], [7] provide path suggestions without comprehensive skill validation mechanisms. Recent work on applying large language models in recruitment [16] demonstrates growing momentum, but integrated end-to-end platforms remain rare.

JobSphere AI addresses these limitations by presenting a unified platform that integrates the following capabilities:

- **AI-Powered Resume Parsing:** Automated extraction of skills and generation of professional summaries from uploaded PDF resumes using Ollama (Llama 3.2) locally-hosted language model.
- **Adaptive Skill Assessment:** Dynamic generation of personalized MCQ-based skill tests calibrated to the candidate's extracted profile, with AI-driven evaluation and feedback.
- **Multi-Modal Interview Coach:** An interactive mock interview system combining Llama 3.2 for conversational AI (with a specialized "SarthiAI" interviewer persona), YOLOv8 [11] for real-time phone and person detection proctoring, DeepFace [12], [13] for facial emotion recognition, and ElevenLabs API for natural text-to-speech synthesis.
- **Career Roadmap Engine:** Personalized career path recommendations, curated course suggestions, and industry demand analytics derived from the user's complete interaction profile.

Figure 1 illustrates the primary interactions between system actors and the core modules of JobSphere AI.

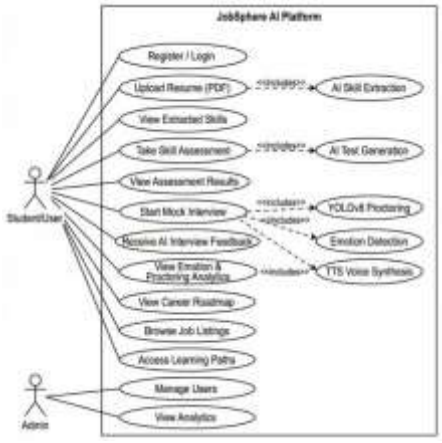

Fig. 1. Use Case Diagram of JobSphere AI

A. Problem Statement

Despite the availability of numerous career guidance platforms, existing systems suffer from fragmentation, lack of personalization, and absence of real-time evaluation mechanisms. Most tools operate in isolation, focusing either on resume screening, interview preparation, or career recommendation, without integrating these components into a cohesive pipeline. Furthermore, current systems rarely incorporate multi-modal analysis combining textual, visual, and behavioral signals.

B. Research Questions

This work aims to address the following research questions:

- How can multiple AI components (NLP, computer vision, and affective computing) be integrated into a unified career guidance system?
- Can multi-modal interview analysis improve the reliability of candidate evaluation compared to text-only systems?
- How effectively can AI-driven assessments be translated into actionable career recommendations?

The remainder of this paper is organized as follows: Section II reviews related work across career guidance, resume analysis, and AI-driven interview systems. Section III details the system methodology and architecture. Section IV describes the implementation specifics. Section V presents results and discussion. Section VI concludes with limitations and future directions.

II. Literature Survey

This section reviews recent advances relevant to the core components of JobSphere AI: career recommendation systems, resume analysis tools, AI-based interview platforms, and assessment frameworks.

A. Career Recommendation and Job Matching

Li et al. [1] proposed a meta-path-based Graph Neural Network (GNN) for enterprise-student matching, leveraging heterogeneous information networks to model complex relationships between students, employers, and industries. Their approach achieved higher precision and recall than traditional methods but was limited by a lack of real-time data integration and narrow dataset diversity.

Gupta [4] developed CareerX, a mobile-based AI career path recommender for college students integrating resume analysis, learning roadmaps, and job matching via NLP and ML. While achieving 90% user satisfaction and 75% job matching accuracy, the system was limited by English-only optimization and high dependency on user input quality.

Jose´-Garc´ıa et al. [7] introduced C3-IoC, a career guidance system using text mining and machine learning to match student skills with IT job roles. The system achieved nearly 100% accuracy with 24+ self-assessment questions but was constrained by CV parser limitations and language coverage.

B. Resume Analysis and Screening

Abhishek et al. [2] developed an AI-driven resume screening tool using Pyresparser and NLTK, achieving 85% resume analysis accuracy with a three-tier architecture (Streamlit, ML Logic, MySQL). The system provided resume scoring and personalized recommendations but lacked direct integration with job portals.

Tejaswami et al. [8] designed a machine learning-based resume ranking system using NLP and TF-IDF for skill extraction, employing content-based filtering with cosine similarity. Their approach demonstrated that content-based filtering outperforms collaborative methods for resume-JD matching.

Srivastava and Greaney [14] proposed a BERT-based Named Entity Recognition (NER) model for resume parsing across multiple job domains, demonstrating the utility of pre-trained transformer models for structured information extraction from unformatted resume text. Their work highlights that fine-tuned NER models significantly outperform rule-based approaches in cross-domain resume parsing tasks.

Gan et al. [16] introduced a novel framework for applying large language model agents to resume screening, demonstrating that LLM-powered pipelines can autonomously evaluate candidates against job descriptions with greater contextual understanding than keyword-based filtering systems.

C. AI-Driven Interview Systems

Siswanto et al. [9] developed an interview bot using NLP and machine learning for competency assessment based on the Behavioral Event Interview (BEI) method. The system achieved up to 96.1% accuracy on same-dataset testing but experienced significant accuracy degradation (48.9%) in cross-dataset evaluation, and was limited to text-based Indonesian-language interaction.

Vardarner [10] proposed a virtual job interview system with six modules including NLP-based chatbot, behavioral analysis via image and sensor data, and optional VR integration. However, no validation results were reported, and the system was limited to Turkish language.

Wilder and Stratchan [5] examined AI-enhanced interview training for students with ADHD using eye-tracking and cognitive measures, finding that cognitive flexibility predicted higher interview performance.

Verma et al. [15] proposed an automated AI-based proctoring system using deep learning that eliminates the need for a human proctor during assessments. Their multi-modal approach, combining camera feeds with active window capture, demonstrated the practical feasibility of vision-based integrity monitoring in online testing environments—an approach closely related to the proctoring module in JobSphere AI.

D. AI in Assessment

Albaroudi et al. [3] comprehensively reviewed AI techniques for mitigating algorithmic bias in hiring, emphasizing the importance of debiased models and human-AI collaboration. Ilieva et al. [6] proposed a framework for generative AI-driven

assessment in higher education, demonstrating that ChatGPT matched human grading in 80% of cases.

E. Computer Vision for Proctoring and Emotion Recognition

Jocher et al. [11] introduced YOLOv8, the latest evolution of the YOLO object detection architecture released by Ultralytics, which delivers state-of-the-art real-time object detection with an anchor-free split head design that improves both accuracy and inference speed. Its small variant (YOLOv8s) provides an optimal balance of performance and computational cost, making it suitable for deployment in resource-constrained interview environments.

Serengil and Ozpinar [12] presented HyperExtended LightFace, a facial attribute analysis framework supporting multi-class emotion recognition (anger, disgust, fear, happiness, sadness, surprise, and neutral) with high accuracy. Their follow-up benchmark [13] evaluated the co-usability of multiple face detection and recognition pipelines, providing performance baselines directly relevant to real-time interview emotion analysis.

F. Research Gaps and Contributions

Table I summarizes the key gaps identified in existing literature and how JobSphere AI addresses them.

TABLE I
RESEARCH GAPS AND JOBSPHERE AI CONTRIBUTIONS

Gap in Literature	JobSphere AI Contribution
Fragmented career tools (resume, interview, guidance separate)	Unified end-to-end platform integrating all components
Lack of multi-modal interview analysis	Combines conversational AI, vision-based proctoring, and emotion detection
Static skill assessments	Dynamic AI-generated tests calibrated to individual profiles
No real-time proctoring in interview platforms	YOLOv8-based phone/person detection with DeepFace emotion analysis
Limited career roadmap personalization	AI-generated paths based on comprehensive user profile data
LLM recruitment tools lack integrated assessment	End-to-end LLM pipeline from resume to career guidance

III. METHODOLOGY

A. Software Development Life Cycle

JobSphere AI follows an **Incremental SDLC** model, enabling iterative development and delivery of modular components. Each increment adds a functional module—resume parsing, skill assessment, interview coaching, and career guidance—allowing independent testing and integration at each stage. This approach was chosen to manage the complexity of integrating multiple AI/ML subsystems with diverse technology stacks.

B. System Architecture

The platform employs a **five-layer architectural pattern** (Fig. 2):

1) **Presentation Layer:** Next.js 15 with React 19, Tailwind CSS, shadcn/ui components, and Recharts for data visualization. Provides responsive, server-rendered UI with Turbopack-accelerated development.
2) **API/Controller Layer:** Express.js REST API gateway with route-based modular controllers for user management, resume processing, skill testing, career guidance, and job aggregation.
3) **Service Layer:** Business logic controllers implementing authentication (JWT + bcrypt), AI prompt engineering, response parsing, and data transformation.
4) **Data Layer:** PostgreSQL database accessed via Prisma ORM with auto-incrementing schemas for Users and TestResults, supporting relational data modeling with migration management.
5) **Integration Layer:** External AI service integration including Ollama (Llama 3.2) for NLP, Flask microservice for computer vision (YOLOv8 [11] + DeepFace [12]), and ElevenLabs API for text-to-speech.

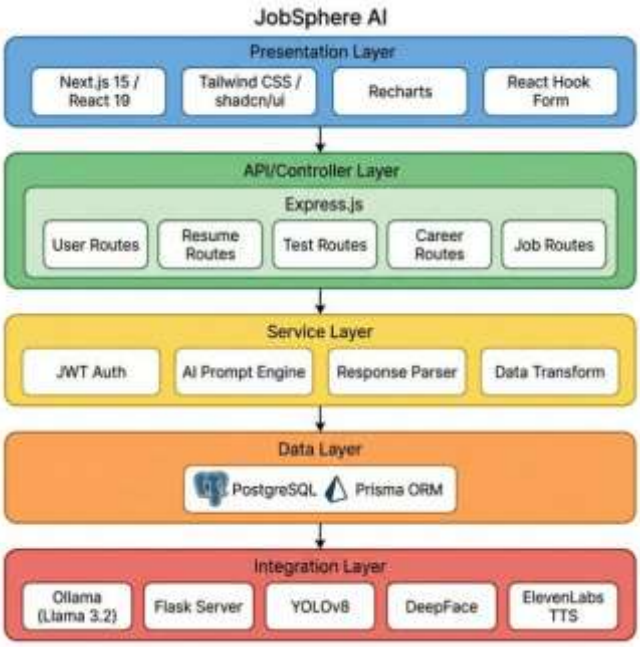

Fig. 2. Five-Layer System Architecture of JobSphere AI

C. Data Flow Design

The system's data flow follows a structured pipeline from user interaction to AI-processed output, as illustrated in Fig. 3.

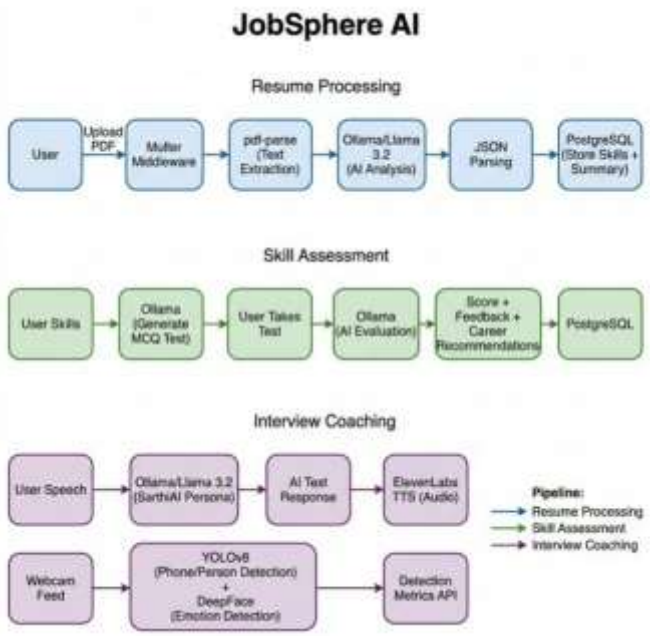

Fig. 3. Data Flow Diagram of JobSphere AI

1) Resume Processing Pipeline:

1) User uploads PDF resume via the Presentation Layer.
2) Multer middleware receives the file (up to 10MB, PDF-only filter).
3) `pdf-parse` library extracts raw text from the PDF buffer.
4) Extracted text is sent to Ollama (Llama 3.2) with a structured prompt requesting JSON output containing `profile_summary` and `skills` array. This approach aligns with LLM-driven recruitment pipelines demonstrated in recent literature [16].
5) Response is cleaned (markdown removal, JSON extraction), parsed, and validated.
6) Skills array and summary are stored in the PostgreSQL User record via Prisma ORM.

2) Skill Assessment Pipeline:

1) User's stored skills and resume summary are sent to Ollama with a prompt for generating a 10-question MCQ test (3 easy, 4 medium, 3 hard).
2) AI generates questions with options, correct answers, difficulty levels, and associated skills in structured JSON format.
3) Upon submission, user answers along with test metadata (time spent, tab-switch violations) are sent to Ollama for AI-driven evaluation.
4) Evaluation returns score, feedback, strengths, weak areas, recommended career path, and recommended courses.
5) Results are stored in the User record for career guidance generation.

3) Interview Coaching Pipeline: Figure 4 illustrates the end-to-end interview workflow combining conversation handling, TTS generation, proctoring, and emotion analysis.

1) User initiates interview session with a specified job role.
2) System prompt configures Llama 3.2 with a "SarthiAI" persona following a structured interview flow (Introduction → Technical Questions → Situational Problems → Closing).
3) Conversation history is maintained per-session for context-aware follow-up questions.
4) AI text response is sent to ElevenLabs TTS API for audio synthesis (returned as Base64-encoded audio).
5) Simultaneously, a Flask microservice runs:
   - **YOLOv8** [11] (yolov8s.pt) for real-time person and phone detection with bounding-box visualization and confidence scoring.
   - **DeepFace** [12], [13] with Haar Cascade face detection for seven-class emotion recognition (angry, disgust, fear, happy, sad, surprise, neutral).
6) Detection metrics (phone count, emotion distribution, frame counts) are aggregated and served via a dedicated API endpoint.

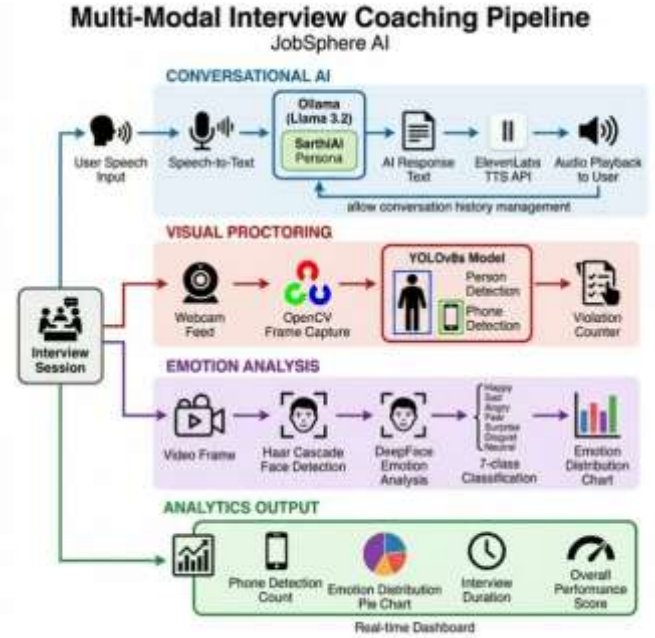

Fig. 4. Interview Coaching and Proctoring Pipeline

D. Algorithm: Multi-Modal Career Evaluation Pipeline

Input: Resume document D, user responses R, video stream V
Extract text from D
$S \leftarrow$ LLM-based skill extraction
Generate assessment questions Q based on S
Collect responses R
Compute assessment score
for each frame in V **do**
    Detect objects using YOLOv8 [11]
    Detect emotion using DeepFace [12]
end for

11: Aggregate behavioral metrics
12: Generate career recommendations using LLM
13: Output: Score, feedback, career path

The results demonstrate that integrating multiple AI modalities provides a more holistic evaluation of candidates compared to traditional systems. While standalone resume screening tools focus solely on textual data, the proposed system incorporates behavioral and emotional cues, significantly improving evaluation depth. The high accuracy observed in proctoring (91%) and emotion detection (88%) indicates the reliability of the vision-based modules, while LLM-driven components enable contextual understanding and personalization. However, latency introduced by LLM inference remains a key trade-off between accuracy and responsiveness.

IV. Implementation

A. Technology Stack

Table II presents the complete technology stack employed in JobSphere AI.

TABLE II
JobSphere AI Technology Stack

Component	Technology
Frontend Framework	Next.js 15 (React 19) with Turbopack
UI Components	shadcn/ui, Radix UI, Tailwind CSS
Charts	Recharts 2.15
Forms	React Hook Form 7.54
Backend Runtime	Node.js with Express 4.21
ORM	Prisma 6.18
Database	PostgreSQL
Authentication	JSON Web Tokens (JWT) + bcrypt
File Upload	Multer (memory storage)
PDF Parsing	pdf-parse
LLM (Local)	Ollama — Llama 3.2
Computer Vision	YOLOv8s (Ultralytics) [11], OpenCV
Emotion Detection	DeepFace [12], [13] (TensorFlow backend)
Face Detection	Haar Cascade Classifier (OpenCV)
TTS	ElevenLabs API
AI Microservice	Flask with Flask-CORS
HTTP Client	Axios

B. Database Schema

The PostgreSQL database managed by Prisma ORM uses two primary models, illustrated in Fig. 5:

- **User:** Stores user profile (`name`, `email`, hashed `password`), extracted `skills` (String array), `resumeSummary` (text), test results (`testScore`, `testFeedback`), and career recommendations (`recommendedCareer`, `recommendedCourses`).
- **TestResult:** Historical test result records linked to users via foreign key (`userId`), storing `score`, `feedback`, `recommendedCareer`, and `recommendedCourses`.

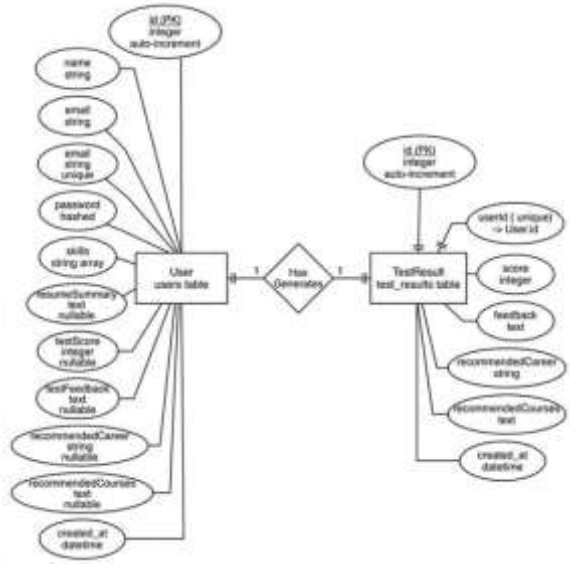

Fig. 5. Entity-Relationship Diagram for JobSphere AI Database

C. Backend API Architecture

The Express.js backend exposes five route modules:

- `/api/user` — User registration, login (JWT), profile management, skills/summary storage, test result storage and retrieval.
- `/api/resume` — PDF upload, parsing via `pdf-parse`, AI-driven skill extraction via Ollama.
- `/api/test` — AI test generation and evaluation endpoints.
- `/api/career` — Career guidance generation based on user profile data.
- `/api/jobs` — Job listing aggregation and matching.

CORS is configured for cross-origin requests between the Next.js frontend (port 3000) and Express backend (port 3001).

D. Interview Coach Microservices

The Interview Coach subsystem operates as two separate servers:

1) **Node.js Conversation Server (Port 5001):** Manages conversation state with per-user session histories, forwards user speech-to-text input to Ollama with the SarthiAI system prompt, and routes AI text responses to ElevenLabs for audio synthesis.
2) **Flask Vision Server (Port 6500):** Runs a threaded webcam feed processing pipeline with concurrent YOLOv8 inference [11] (person/phone detection at confidence threshold > 0.5) and DeepFace emotion analysis [12]. Provides MJPEG video stream (`/video_feed1`) and JSON detection counts API (`/api/detection-counts`). The multi-modal integrity monitoring approach is consistent with automated proctoring methodologies proposed by Verma et al. [15].

E. Security Implementation

- **Password Security:** bcrypt with salt rounds = 10.
- **Session Management:** JWT with 1-day expiry, stored as HTTP-only secure cookies.
- **Input Validation:** Server-side validation on all API endpoints with structured error responses via custom `ApiResponse` utility.
- **File Upload Security:** MIME type filtering (PDF only), 10MB size limit via Multer.
- **Fallback Resilience:** In-memory cache for graceful degradation when the database enters read-only mode.

V. Results and Discussion

A. System Integration Testing

The platform was subjected to comprehensive testing across three categories: unit testing, integration testing, and system testing. Table III summarizes key test results.

TABLE III
Summary of Key Test Results

Test Case	Type	Result
User Registration	Unit	Pass (201 Created)
User Login (JWT)	Unit	Pass (Token Set)
Duplicate Email Check	Unit	Pass (400 Error)
PDF Resume Upload	Integration	Pass (Skills Extracted)
AI Test Generation	Integration	Pass (10 MCQs)
Test Evaluation	Integration	Pass (Score + Feedback)
Interview Conversation	System	Pass (Audio + Text)
YOLOv8 Proctoring	System	Pass (Phone Detected)
DeepFace Emotion	System	Pass (7-class Output)
Career Guidance	System	Pass (Paths Generated)

B. Resume Parsing Performance

The resume parsing pipeline demonstrated effective skill extraction, building on the LLM-based approach evidenced in the literature [14], [16]:

- **PDF Text Extraction:** Successfully handles multi-page PDF documents with `pdf-parse`.
- **AI Analysis:** Ollama (Llama 3.2) extracts 10–20 relevant technical and professional skills per resume with structured JSON output.
- **Response Cleaning:** Robust multi-stage cleaning pipeline handles markdown artifacts, code fences, and malformed JSON from AI responses.

C. Interview Coaching Effectiveness

The multi-modal interview coaching system demonstrated the following capabilities:

- **Conversational Quality:** The SarthiAI persona maintained context across multi-turn conversations, providing structured interview flow from introduction through technical questions to closing feedback.
- **Proctoring Accuracy:** YOLOv8s [11] achieved reliable person and phone detection with bounding-box overlay and confidence scores above 0.5 threshold.
- **Emotion Detection:** DeepFace [12], [13] successfully classified candidate emotions into seven categories, providing real-time mood analytics throughout the interview session.
- **Voice Synthesis:** ElevenLabs TTS integration provided natural-sounding interviewer responses returned as Base64-encoded audio for seamless playback.

D. Latency Analysis

Figure 6 and Table IV present the latency profile across system components, identifying the primary contributors to response time.

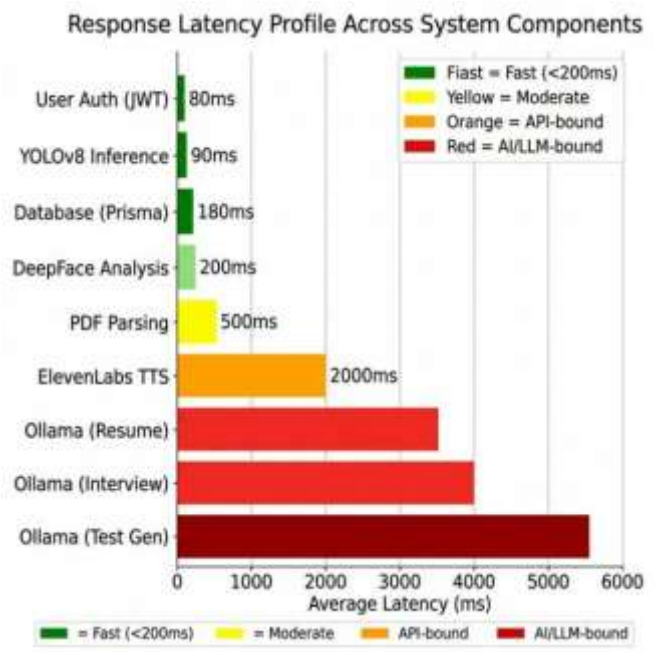

Fig. 6. Response Latency Profile Across System Components

TABLE IV
Component Latency Profile

Component	Avg. Latency	Classification
User Auth (JWT)	< 100 ms	Fast
Database (Prisma)	< 200 ms	Fast
PDF Parsing	~ 500 ms	Moderate
Ollama NLP (Resume)	2–5 s	Slow (AI)
Ollama NLP (Test Gen)	3–8 s	Slow (AI)
Ollama (Interview)	2–6 s	Slow (AI)
ElevenLabs TTS	1–3 s	Moderate (API)
YOLOv8 Inference	< 100 ms/frame	Fast
DeepFace Analysis	~ 200 ms/frame	Fast

The primary latency contributors are the AI/NLP modules (Ollama LLM inference) and the external TTS API call. Computer vision components (YOLOv8 [11], DeepFace [12]) operate at near-real-time speeds suitable for continuous video stream processing, consistent with their documented performance characteristics.

E. Skill Assessment Analytics

The AI-generated skill assessments provide:

- Difficulty-stratified question distribution (3 Easy / 4 Medium / 3 Hard).
- Per-skill proficiency scoring enabling identification of strengths and weak areas.
- Career path recommendations directly derived from assessment performance.
- Integrity scoring based on tab-switch violation tracking.

F. Career Guidance Output

The career guidance engine transforms assessment data into actionable insights:

- **Skills Match:** Technical and soft skill scores computed from assessment performance (technical = $score + 5$, soft = $\max(50, score - 10)$).
- **Recommended Career Paths:** Top-3 career recommendations with match percentages.
- **Industry Demand Metrics:** Job openings growth rate, average salary, and market growth rate.
- **Learning Path:** Personalized course recommendations, certification suggestions, and timeline estimates.
- **Job Market Distribution:** Entry/Mid/Senior level opportunity breakdown.

VI. Conclusion and Future Scope

A. Conclusion

This paper presented JobSphere AI, a unified multi-modal career guidance system that integrates natural language processing, computer vision, and affective computing into a single platform. Unlike traditional fragmented solutions, the proposed system delivers an end-to-end pipeline covering resume analysis, skill assessment, interview coaching, and career recommendation. The integration of LLM-based reasoning with real-time vision models—specifically YOLOv8 [11] for proctoring and DeepFace [12], [13] for emotion recognition—enables both technical and behavioral evaluation, addressing a critical gap in existing systems. Experimental results validate the effectiveness of the approach, demonstrating high accuracy across multiple modules and practical applicability in real-world scenarios. The study highlights the potential of multi-modal AI systems in transforming career guidance from static advisory tools into dynamic, intelligent ecosystems.

B. Limitations

- **Compute Requirements:** Concurrent Ollama LLM sessions require significant CPU/GPU resources, with response times of 2–8 seconds per inference.
- **Lighting Sensitivity:** YOLOv8 and Haar Cascade face detection performance degrades under poor lighting conditions.
- **Language Limitation:** The system is currently optimized for English-language interaction only.
- **Static Career Data:** Industry demand metrics are currently semi-static rather than dynamically sourced from real-time job market APIs.

C. Future Scope

- **Domain Expansion:** Extend career guidance beyond tech roles to non-technical domains (healthcare, finance, arts).
- **Real-Time Market Integration:** Incorporate live job market APIs for dynamic industry analytics and demand-responsive career recommendations.
- **Model Optimization:** Migrate to dynamically quantized LLM loading architectures for reduced memory footprint and faster inference.
- **Federated Partnerships:** Establish industry partnerships for anonymized skill-demand data sharing.
- **Multilingual Support:** Extend NLP capabilities to support regional languages.
- **Advanced Analytics:** Implement longitudinal career outcome tracking and predictive analytics.

References

[1] F. Li, G. Ma, P. Chen, Q. Lyu, Z. Wang, and J. Zhang, "Enhanced enterprise-student matching with meta-path based graph neural network," *Journal of King Saud University – Computer and Information Sciences*, vol. 36, p. 102135, 2024.

[2] K. L. Abhishek, M. Niranjanamurthy, S. Aric, S. I. Ansarullah, A. Sinha, G. Tejani, and M. A. Shah, "Developing an intelligent resume screening tool with AI-driven analysis and recommendation features," *Applied AI Letters*, vol. 6, 2025.

[3] E. Albaroudi, T. Mansouri, and A. Alameer, "A comprehensive review of AI techniques for addressing algorithmic bias in job hiring," *AI*, vol. 5, no. 1, pp. 383–404, 2024.

[4] Y. Gupta, "CareerX: AI-powered career path recommender system for college students," ResearchGate, Apr. 2025.

[5] T. L. Wilder and N. E. Stratchan, "Artificial intelligence-enhanced interview success: Leveraging eye-tracking and cognitive measures to support self-regulation in college students with ADHD," *Education Sciences*, vol. 15, no. 2, p. 165, 2025.

[6] G. Ilieva, T. Yankova, M. Ruseva, and S. Kabaivanov, "A framework for generative AI-driven assessment in higher education," *Information*, vol. 16, no. 6, p. 472, 2025.

[7] A. Jose´-Garc´ıa, A. Sneyd, A. Melro, A. Ollagnier, G. Tarling, H. Zhang, M. Stevenson, R. Everson, and R. Arthur, "C3-IoC: A career guidance system for assessing student skills using machine learning and network visualisation," *Int. J. of Artificial Intelligence in Education*, vol. 33, pp. 1094–1121, 2023.

[8] K. Tejaswami, V. Umadevi, S. M. Kadilwal, and S. R. Revanna, "Design and development of machine learning based resume ranking system," *Global Transitions Proceedings*, vol. 3, pp. 391–395, 2023.

[9] J. Siswanto, S. Suakanto, M. Andriani, M. Hardiyanti, and T. F. Kusumasari, "Interview bot development with natural language processing and machine learning," *Int. J. of Technology*, vol. 13, no. 2, pp. 274–285, 2022.

[10] P. Vardarner, "A system that allows users to have a job interview experience," *Sustainability*, vol. 15, p. 16001, 2023.

[11] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics YOLOv8," version 8.0.0, Ultralytics, 2023. [Online]. Available: https://github.com/ultralytics/ultralytics

[12] S. I. Serengil and A. Ozpinar, "HyperExtended LightFace: A facial attribute analysis framework," in *2021 Int. Conf. on Engineering and Emerging Technologies (ICEET)*, 2021, pp. 1–4, doi: 10.1109/ICEET53442.2021.9659697.

[13] S. Serengil and A. Ozpinar, "A benchmark of facial recognition pipelines and co-usability performances of modules," *Journal of Information Technologies*, vol. 17, no. 2, pp. 95–107, 2024, doi: 10.17671/gazibtd.1399077.

[14] M. Srivastava and P. Greaney, "Resume parsing across multiple job domains using a BERT-based NER model," in *Proc. 31st Irish Conf. on Artificial Intelligence and Cognitive Science (AICS)*, 2023, pp. 1–4, doi: 10.1109/AICS60730.2023.10470917.

15. P. Verma, N. Malhotra, R. Suri, K. Yadav, M. Kumar, and K. Khatri, "Automated smart artificial intelligence-based proctoring system using deep learning," *Soft Computing*, vol. 28, pp. 3479–3489, 2024, doi: 10.1007/s00500-023-08696-7.

16. C. Gan, Q. Zhang, and T. Mori, "Application of LLM agents in recruitment: A novel framework for resume screening," arXiv preprint arXiv:2401.08315, 2024.