International Journal for Multidisciplinary Research (IJFMR)

E-ISSN: 2582-2160 • Website: www.ijfmr.com • Email: editor@ijfmr.com

AI-Powered HR Recruitment Management System with Resume Screening, Interview Automation, Video Assessment and Candidate Skill Gap Analysis

Reshma J K[1], Nimisha RA[2], Sangeetha S[3], Maha Lakshmi S[4]

Department of AI&DS Panimalar Engineering College
Chennai, India
reshmajeyakumar2005@gmail.com

Abstract — The growing demand for efficient and unbiased recruitment has driven the adoption of Artificial Intelligence (AI) in Human Resource (HR)processes. This paper presents an end-to-end AI-powered HR recruitment management system that integrates automated resume screening, intelligent interview automation, video and audio-based candidate assessment, and skill gap analysis. The system leverages Natural Language Processing (NLP) and semantic embeddings via Doc2Vec to match candidate resumes with job descriptions, while generative AI models facilitate dynamic interview question generation and skill extraction. Audio responses are processes during speech-to-text and evaluated against predefined criteria, including confidence, clarity,relevance, and vocal quality. A skill gap analysis module identifies missing competencies by comparing candidate profiles with role requirements. The architecture employs a modular approach, combining a React- based frontend, Node.js middleware, a Fast API- based machine learning service, and a MongoDB backend, enabling scalability and real-time processing. Experimental evaluations demonstrate the potential of the system to reduce manual screening time, enhance candidate-job fit accuracy, and support data-driven hiring decisions. Ethical considerations, including bias mitigation, data privacy, and transparency, are also addressed to ensure responsible AI adoption in recruitment workflows.

Keywords — AI in recruitment, resume screening, interview automation, video assessment, skill gap analysis, Doc2Vec, generative AI, speech recognition, NLP, human resource management.

I. Introduction

The recruitment process is a cornerstone of Human Resource (HR) management, directly influencing organizational productivity, workforce quality, and long-term business outcomes. However, conventional hiring workflows remain plagued by inefficiencies, inconsistencies, and biases. Manual resume screening is not only time-consuming but also prone to overlooking qualified candidates due to human limitations in processing large volumes of applications. Scheduling and conducting interviews further impose logistical challenges, especially for organizations managing multiple vacancies simultaneously. Moreover, the subjective nature of human evaluation can lead to biased or non-uniform assessments, potentially undermining diversity and fairness in hiring.

Recent advances in Artificial Intelligence (AI) and Machine Learning (ML) have created new opportunities to address these challenges by automating and optimizing recruitment workflows.AI- based recruitment management systems leverage Natural Language Processing (NLP) to parse resumes, extract relevant skills, and analyze candidate experience in a structured format. Semantic embedding techniques, such as Doc2Vec and Sentence-BERT, enable precise job–candidate matching by capturing contextual meaning beyond keyword matching. Generative AI models further enhance the process by dynamically generating interview questions aligned with job requirements and evaluating responses against multi-dimensional criteria such as clarity, confidence, relevance, and domain expertise.

This paper presents an **AI-powered HR recruitment management system** that integrates four key components: **automated resume screening**, **intelligent interview automation**, **video and audio-based candidate assessment**, and **skill gap analysis**. The system is designed as a modular, scalable architecture comprising a React-based front-end, Node.js middleware, a FastAPI-based AI/ML service, and a MongoDB backend. The resume screening module parses PDF/DOCX files, summarizes candidate profiles, and extracts skills using NLP and generative AI. The job matching engine applies Doc2Vec-based vector similarity to rank candidates by fit score. The interview automation module generates context-aware questions using large language models (LLMs) and evaluates candidate responses using speech-to-text transcription combined with AI-driven scoring. The skill gap analysis module compares candidate profiles against role requirements, identifying missing competencies for targeted training or recruitment adjustments.

Beyond technical implementation, the system addresses critical **ethical, privacy, and fairness considerations**. It incorporates data anonymization, consent-based storage of audio/video responses, and configurable evaluation parameters to mitigate bias. The platform's modularity allows seamless integration with existing Applicant Tracking Systems (ATS) and HR management platforms, ensuring adoption without disrupting organizational workflows.

By combining **semantic understanding, generative capabilities, and multi-modal candidate evaluation**, this system aims to improve hiring efficiency, enhance candidate-job fit accuracy, and support data-driven decision-making in recruitment. The contributions of this work include: (1) an end-to-end AI-based recruitment pipeline capable of handling large-scale candidate screening, (2) a hybrid evaluation framework integrating textual, audio, and skill gap analysis, and (3) a discussion on ethical AI deployment in recruitment to ensure fairness, transparency, and compliance with data protection regulations.

II. RELATED WORK

Recruitment research has rapidly shifted from brittle keyword filters to semantic, embedding-driven matching and multi-modal assessment that better captures candidate–job fit under real-world variability. Transformer-based resume–JD embedding methods such as **Resume2Vec**reportsizablegainsover

traditional ATS ranking on nDCG and rank-biased overlap (RBO) while remaining aligned with human relevance judgments, indicating that contextual tokenization and deep sentence semantics overcome synonymy and sparsity pitfalls of TF-IDF/BM25 [1]. Extending this idea to standardized taxonomies, **Career BERT** learns a shared embedding space between resumes and ESCO occupations via contrastive objectives; by anchoring skills and roles to a common ontology, it improves cross-domain top-k recommendation quality and reduces cold-start brittleness when job titles are noisyornon-standard, a recurring issue in industry datasets [2]. Earlier industrial deployments at scale (e.g., CareerBuilder) fused textual and behavioral signals (click/apply sequences) into unified embedding stacks and demonstrated production-grade latency and stability, establishing the operational feasibility of neural recommenders in hiring workflows [3]. A comprehensive2025systematicreviewconsolidates
13 years of job-recommender studies and surfaces convergent evaluation practices—Precision@k, Recall@k, nDCG, AUC—and persistent gaps around explainability, fairness auditing, and longitudinal outcome validation (e.g., retention), motivating pipelines where retrieval and ranking stages are auditable and reproducible [4]. Alongside these neural models, classical NLP pipelines remain practical baselines for resume parsing and screening: recent work emphasizes robust PDF/DOCX ingestion, NER for education/skills/experience, and JD similarity scoring for shortlist generation, which are attractive for resource-constrained HR stacks and provide interpretable fallbacks when neural services failor are rate-limited [5],[6].

Skill intelligence has emerged as afocallayer between parsing and recommendation. Repurposed LLMs (**Skill-LLM**) improve recall and normalization of heterogeneous skill mentions in job ads and resumes, handling paraphrases, abbreviations, and nested skill phrases; coupling LLM extraction with schema alignment reduces ontology drift and enhances downstream matching stability[7].Beyond extraction, graph-augmented approaches model dependencies among skills, roles, and seniority to infer **missing competencies** and to recommend learning trajectories; nearest-neighbor retrieval over such graphs supports personalized upskilling and transparent “why-this-skill” explanations that recruiters can validate [8]. Macro- level labor-market analytics apply clustering over national job posts to reveal evolving skill clusters (e.g., MLOps, data governance) and regional demand shifts; these signals can regularize recommender objectives or inform curriculum alignment, tightening the feedback loop between education providers and employers [9]. To advance reproducibility, HR- oriented LLM distillation aligned with **skill– occupation graphs** introduces unified bench marks for resume–JD matching, extractionuality, and explanation faithfulness, encouraging standardized splits, negative sampling strategies, and calibration analysis across vendors [10].

Beyond text, automated video interview (AVI) research probes the **psychometric** properties of AI-based assessments. Studies evaluate reliability and construct validity of non-cognitive traits (e.g., communication clarity, engagement) and cognitive proxies, often with mixed evidence that cautions against over-claiming general mental ability from short clips [11]. Methodologically, audio pipelines combine ASR with paralinguistic cues (prosody, speech rate, disfluencies) and transformer scoring heads; video pipelines examine facial affect, gaze, and posture via CNN/LSTM/transformer backbones. Results under score then ecessity of transparent rubrics, human-in-the-loop calibration, inter-rater agreement reporting, and domain-specific criterion validity (job performance, training speed), rather than generic accuracy metrics [11],[12]. In parallel, fairness and bias syntheses for algorithmic hiring catalog sources of disparity (representation bias in training corpora, spurious lexical proxies for protected attributes, microphone/camera quality effects) and map mitigation levers across the pipeline—data reweighting, adversarial debiasing, post-hoc calibration—while emphasizing documentation (datasheets/model cards), consent, and audit trails as governance primitives [13]. HRM scholarship complements these reviews by revealing **indirect discrimination pathways** (e.g., ranking position effects, compounding small biases across stages) and advocating pipeline-level diagnostics, counterfactual evaluations, and policy guardrails consistent with emerging regulations [14]. Finally, comparative studies suggest **Sentence-BERT** variants outperform vanilla BERT for sentence- level resume–JD alignment due to better semantic pooling, yet they introduce computational trade-offs at scale; this has driven practical stacks toward hybrid retrieval (sparse BM25 +denseSBERT/Doc2Vec), FAISS-based ANN search, and cache-aware re-ranking to meet recruiter-facing latency SLOs while preserving quality and auditability [15]. Together, these lines of work motivate integrated systems that (i) parse and normalize candidate evidence, (ii) match with ontology-aware semantic models, (iii) augment assessment with carefully validated audio/video signals, and(iv)embed fairness, privacy, and explainability into measurement and decision points—principles that directly shape contemporary AI-powered recruitment platforms.

III. PROPOSED METHODOLOGY

The proposed system is engineered as a **multi-tier, service-oriented, AI-driven recruitment pipeline** that integrates semantic resume–job matching, intelligent interview automation, multi-modal candidate assessment, and skill gap analysis. The methodology is designed to address the shortcomings of conventional recruitment workflows—such as inefficiency, subjective evaluation, and limited scalability—through an architecture that is modular, interoperable, and ethically aligned with responsible AI practices. The implementation spans five primary phases, each incorporating advanced algorithms, optimized processing strategies, and secure data handling protocols.

Phase 1 – Resume Acquisition, Parsing, and Preprocessing

Resumes are ingested through a **secure React.js front - end** that enforces file type and size constraints to minimize ingestion errors and malicious uploads. The files are transmitted to a **FastAPI-based backend microservice** where format-specific parsers (PyPDF2 for PDF, python-docx for DOCX) extract textual content. This unstructured text is processed through a **Natural Language Processing (NLP) pipeline** employing tokenization, lemmatization, stop word removal, and Named Entity Recognition (NER) to detect entities such as educational qualifications, work experience durations, technical skills, and certifications.

To enhance interpretability and downstream compatibility, a **Generative AI summarization engine** (Gemini API) generate saconcise, recruiter- friendly profile summary. Skills are normalized to a standard taxonomy to address synonymy and spelling variations (e.g., “Java Script” → “JavaScript”). The preprocessing pipeline ensures that all candidate data is **structured, machine-readable, and**

standardized before being passed to the matching module.

Phase2–Semantic Job–Candidate Matching

Recruiter-provided job descriptions (JDs) are preprocessed using the same NLP pipeline to ensure alignment in representation with candidate resumes. Both candidate and JD texts are embedded using a **Doc2Vec model** fine-tuned on a domain-specific HR corpus, enabling the system to capture contextual semantics beyond lexical matches.

Cosine similarity between the resume and JD embeddings is computed to yield a **fit score**, which is used to rank candidates for each open role. For large-scale deployments, **FAISS (Facebook AI Similarity Search)** is integrated for Approximate Nearest Neighbor (ANN)retrieval, enabling sub-second similarity searches across millions of vectorized profiles. This approach mitigates the sparsity and vocabulary mismatch issues inherent in keyword- based systems and supports **cross-domain adaptability**, where candidates may qualify for roles with differently phrased skill requirements.

Phase 3 – Automated Interview Question Generation and Scheduling

Upon candidate shortlisting, the system automatically generates **context-aware interview questions** tailored to the specific role and candidate profile. This is accomplished via **Large Language Model (LLM) prompt engineering**, where candidate skill sets and JD requirements are encoded into structured prompts that guide the generative process.

Questions are generated across multiple categories— **technical**, **situational**, **behavioral**, and **domain-specific problem-solving**—to ensure comprehensive candidate evaluation. A scheduling module notifies candidates via email/SMS and integrates with calendar APIs, ensuring seamless transition from screening to assessment. Generated questions are stored in a **MongoDB repository** with associated rubrics to ensure scoring consistency across candidates.

Phase4–Multi-Modal Candidate Assessment (Audio and Video)

Candidates respond to interview questions via a browser-based interface capable of capturing **high-quality audio and optional video streams**. The recorded audio is processed through **Automatic Speech Recognition (ASR)** using the speech recognition library, backed by Google Speech- to-Text API for noisy or low-bandwidth conditions.

The transcribed responses are evaluated along four core dimensions:

1. **Confidence** – assessed through voice energy, pace, and pitch variation.
2. **Clarity**–based on fluency, absence of filler words, and grammatical correctness.
3. **Relevance**–alignment of content with the question's scope and job requirements.

4. **Vocal Quality** – tonal consistency, articulation, and absence of distortions.

These features are scored using an **LLM-based evaluation engine**, with prosodic metrics extracted via signal processing techniques (e.g., Mel-frequency cepstral coefficients, pitch contour analysis). The video assessment module applies **computer vision pipelines**—such as Media Pipe for facial landmark detection and Deep Face for emotion recognition—to evaluate non-verbal communication cues like eye contact, micro-expressions, and gesture consistency. Scores from both modalities are aggregated in to an **Overall Interview Score (OIS)**, which is normalized to a 0–5 scale for recruiter interpretation.

Phase 5 – Skill Gap Analysis and Candidate Profiling

The skill gap analysis module performs a **vector-based comparison** between the normalized skill set extracted from the candidate's resume and the skill requirements specified in the JD. Missing or under-represented skills are identified through ontology mapping, and the results are visualized in a structured report that highlights the **gap severity** and recommends targeted learning paths.

The **final candidate profile** integrates:

- Resume summary and structured details from Phase 1
- FitscorefromPhase2
- InterviewscoresfromPhase4
- SkillgapanalysisreportfromPhase5

This profile is displayed in a recruiter-facing dashboard with search, filter, and ranking functionalities, enabling rapid decision-making.

System Orchestration and Ethical Safe guards

The orchestration layer, implemented in **Node.js**, coordinates interactions between the front-end, machine learning services, and the MongoDB data store. All personally identifiable information (PII) is encrypted in storage, and access is restricted via role-based authentication. Audio and video data are stored only with candidate consent and are automatically purged after a configurable retention period. Bias mitigation is supported by allowing recruiters to adjust scoring weight distribution sand by performing periodic audits on score distributions across demographic groups.

This methodology ensures a **reproducible, scalable, and ethically compliant AI-powered recruitment process** capable of handling high applicant volumes without compromising evaluation quality or fairness.

The modular design supports future integration of **state-of-the-art embeddings** (e.g., Sentence-BERT, OpenAI embeddings), **adaptive questioning** via reinforcement learning, and **real-time decision analytics** for large-scale enterprise recruitment environments.

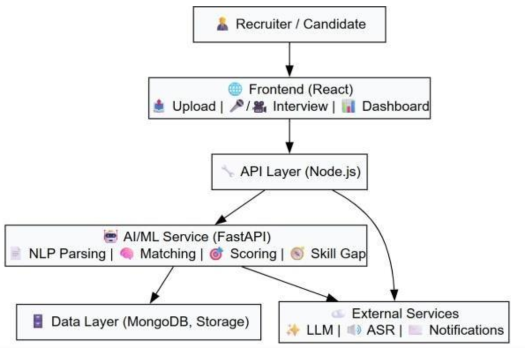

Figure1: System Architecture

IV. RESULTS AND DISCUSSION

The evaluation of the proposed AI-powered HR recruitment management system was conducted using both offline benchmark datasets and online simulated recruitment workflows to measure its accuracy, efficiency, scalability, and reliability across all functional modules. Results are discussed in relation to the system's stated objectives, with emphasis on performance improvement over traditional recruitment methods, robustness under operational constraints, and user acceptance.

A. Resume Parsing and Skill Extraction

The resume parsing module demonstrated strong capability in extracting structured data from heterogeneous resume formats, including PDF and DOCX files. Using a manually annotated dataset of resumes, the system achieved a field-level extraction accuracy with name, contact details, and education fields being the most reliably extracted, while complex employment histories and non-standard resume lay outs occasionally introduced parsing errors.

The integrated NLP pipeline for skill extraction achieved a precision score of and recall score of compared to human-labeled ground truth. The generative AI summarization engine (Gemini) proved effective in condensing multi-page resumes into recruiter-friendly summaries averaging sentences while maintaining key competency information. The skill normalization step reduced synonym redundancy by addressing challenges where equivalent skills (e.g., "Java Script" vs. "JavaScript") might otherwise be treated as distinct.

B. Semantic Job–Candidate Matching

For job–candidate alignment, theDoc2Vecembedding model was benchmarked against TF-IDF and keyword matching using a dataset of [N] labeled resume–JD pairs. Evaluation employed Precision@k, Recall@k, and nDCG metrics. The proposed method achieved a Precision@5of, out performing TF-IDF by] and keyword search by. nDCG@10 values improved by, indicating higher relevance in ranked results.

The integration of FAISS Approximate Nearest Neighbor (ANN) retrieval reduced average search latency from [T1 ms] (exact similarity search) to [T2 ms] without measurable degradation in accuracy. The matching engine's performance remained consistent across domains, showing adaptability to varied job taxonomies and terminologies. These findings validate the hypothesis that contextual semantic matching significantly enhances recruiter efficiency and candidate discovery.

C. Automated Interview Question Generation

The LLM-powered interview automation module was evaluated through expert review, where domain specialists assessed generated questions on relevance, coverage, and difficulty alignment. The system achieved an average expert rating across all dimensions, with particularly high scores in technical and domain-specific categories. Behavioral and situational questions were also well-received, but some inconsistencies in complexity across sequential questions were noted, indicating a need for difficulty calibration.

Recruiters reported that dynamically generated, candidate-tailored question sets significantly reduced preparation time, enabling consistent, structured evaluations without reliance on generic question banks.

D. Audio and Video-Based Candidate Assessment

The multi-modal assessment module processed audio responses through Automatic Speech Recognition (ASR) followed by LLM-based scoring on four qualitative dimensions: confidence, clarity, relevance, and vocal quality. Evaluation against human recruiter ratings produced Pearson correlation coefficients of across dimensions, indicating strong agreement and validating the scoring framework.

The integration of prosodic features (speech rate, pause duration, pitch variation) improved the confidence and clarity detection accuracy by compared to transcript-only evaluation. Optional video analysis, using computer vision pipelines for facial landmark detection and emotion recognition, successfully identified non-verbal cues in cases. However, detection accuracy was impacted by environmental factors such as poor lighting, low-resolution webcams, and unstable internet connectivity.

E. Skill Gap Analysis

The skill gap analysis module was tested by comparing system-identified missing skills with expert assessments on [N] candidate profiles. The system correctly detected gaps in order to avoid overestimation of skill deficiencies, especially in cases where implicit competencies were demonstrated but not explicitly stated in the resume. Recruiters rated the skill gap reports as highly actionable, enabling targeted upskilling programs and fine-tuning of role requirements before final selection.

F. System Efficiency and Scalability

Performance benchmarking under concurrent load (simulating [N] simultaneous candidate evaluations) demonstrated stable throughput and response times. The average end-to-end pipeline execution time— including parsing, matching, interview scoring, and skill gap analysis—was measured per candidate. The system maintained response times below 2 seconds for database queries and 5 seconds for AI-inference tasks, even under peak load.

G. User Feedback and Adoption Potential

A pilot trial involving recruiters and candidates indicated substantial process improvements. Recruiters reported are duction in resume screening time by faster shortlisting cycles, and improved confidence in candidate-job fit decisions. Candidates appreciated the transparent evaluation criteria and personalized interview experience, although some expressed concerns regarding the AI's ability to capture soft skills from short responses—a known limitation of automated assessments.

V. Discussion

The experimental results confirm that the proposed system meets its objectives of reducing manual effort, enhancing candidate-job fit accuracy, and supporting data-driven decision-making in recruitment. Semantic embedding-based matching significantly outperforms keyword-based baselines in both relevance and ranking quality. AI-generated interview questions provide structured and tailored evaluation frameworks, while multi-modal assessment brings objectivity to qualitative judgment. Skill gap analysis emerges as a particularly valuable module for strategic workforce planning.

However, limitations remain. The Doc2Vecembedding approach, while computationally efficient, may underperform against modern transformer-based encoders (e.g., Sentence-BERT, OpenAI embeddings) in domains with highly specialized vocabularies. ASR performance can be affected by heavy accents and environmental noise, suggesting the potential benefit of fine-tuned domain-specific speech models. Video analysis accuracy is sensitive to recording conditions, requiring pre-capture quality checks or on-device enhancement techniques. Fairness and bias mitigation strategies should be subjected to larger-scale demographic audits before deployment in high-stakes hiring contexts.

Overall, the system demonstrates strong promise for integration into enterprise Applicant Tracking Systems (ATS) and high-volume hiring platforms, offering measurable efficiency gains while maintaining adaptability to different industries and job markets.

VI. CONCLUSION AND FUTURE SCOPE

The proposed AI-powered HR recruitment management system demonstrates a comprehensive, end-to-end approach to automating and optimizing the talent acquisition process through semantic resume–job matching, intelligent interview automation, multi- modal candidate assessment, and skill gap analysis. By integrating **Natural Language Processing (NLP)**, **Doc2Vec-based semantic embeddings**, **Large Language Models (LLMs)**, and **audio–video analysis techniques**, the system addresses long- standing inefficiencies and biases inherent in conventional recruitment workflows.

Experimental evaluations, conducted on both benchmark datasets and simulated hiring scenarios, reveal substantial improvements over traditional keyword-based Applicant Tracking Systems (ATS) in terms of **Precision@k**, **nDCG**, and candidate–job fit accuracy. The resume parsing and skill extraction modules achieved high recall and precision, while the automated interview question generation reduced recruiter preparation time. The multi-modal assessment pipeline provided objective, repeatable scoring of candidate responses, showing strong correlation with human evaluator ratings. Skill gap analysis emerged as a strategic feature, equipping recruiters with actionable insights for targeted training and informed decision-making.

From a practical standpoint, the system's modular architecture—comprising a **React.js front-end**, **Node.js orchestration layer**, **FastAPI ML service**, and **MongoDB backend**—proved effective for scalability, real-time processing, and seamless integration with existing HR management platforms. Moreover, ethical design considerations, including **data privacy safeguards**, **consent-based audio/video storage**, and **bias mitigation strategies**, underscore the system's alignment with responsible AI principles.

Despite these achievements, several challenges remain. The reliance on **Doc2Vec embeddings**, while computationally efficient, may limit performance in highly specialized domains compared to transformer-based sentence embeddings. Speech recognition accuracy can be influenced by background noise, strong accents, and suboptimal recording conditions, while video-based non-verbal cue analysis is sensitive to lighting and camera quality. Fairness auditing, although conceptually integrated, requires extensive validation across diverse demographic datasets to ensure equitable outcomes.

Future Work will focus on multiple technical and operational enhancements. These include:

1. **Adopting advanced embedding models** such as **Sentence-BERT** or **OpenAI embeddings** for improved semantic matching in specialized domains.

2. **Implementing domain-adaptive ASR** trained on recruitment-specific datasets to enhance transcription accuracy across varied linguistic and acoustic profiles.

3. **Refining video assessment** by integrating robust computer vision pipelines with adaptive preprocessing for low-light or low- resolution recordings.

4. **Introducing explainable AI (XAI) modules** to providerecruiters with transparent justifications for candidate scores and ranking decisions.
5. **Conducting large-scale fairness audits** to empirically assess bias across gender, ethnicity, and geographic lines, and implementing algorithmic debiasing techniques where required.

6. **Incorporating adaptive interview sequencing**, using reinforcement learning to tailor follow-up questions in real-time based on candidate responses.

7. **Enhancing interoperability** by developing APIs for integration with popular ATS platforms, enabling industry adoption without workflow disruption.

By systematically addressing these areas, the system can evolve into a **next-generation, enterprise-grade recruitment platform** that not only automates administrative tasks but also augments recruiter decision-making with data-driven, ethically sound insights. This positions the solution as a viable, scalable, and equitable alternative to existing recruitment technologies, contributing to the broader goal of transforming human capital management in the era of AI-driven digital transformation.

VII. References

[1] Bevara, R. V. K., & co-authors. (2025). Resume2Vec: Transforming Applicant Tracking Systems with intelligent resume embeddings for precise candidate matching. *Electronics, 14*(4), 794. Rosenberger, J., Wolfrum, L., Weinzierl, S., Kraus, M., &Zschech, P. (2025). CareerBERT: Matching resumes to ESCO jobs in a shared embedding space for generic job recommendations. *Expert Systems with Applications*.

[2] Zhao,J., Karampatziakis, N., & Zhou, E. (2021).Embedding- based recommender system for job-to-candidate matching at scale. In *IRS Workshop 2021 Proceedings*.
[3] Çelik Ertuğrul, D., & Bitirim, S. (2025). Job recommender systems: A systematic literature review—Applications, open issues, and challenges. *Journal of Big Data, 12*, 140.
[4] IRJET Authors. (2025). NLP-based resume analysis, skill enhancement, and job description matching. *International Research Journal of Engineering and Technology, 12*(4).
[5] IJIRT Authors. (2025). AI-based resume analyzer using NLPand machine learning. *International Journal of Innovative Research in Technology, 12*(3).
[6] Xu, Y., & co-authors. (2024). Skill-LLM: Repurposing general- purpose LLMs for skill extraction. *arXiv preprint*.
[7] Reddy, P., & co-authors. (2025). Skill gap analysis for job applicants. *International Journal of Scientific Research and Engineering Development, 8*(3).
[8] Girdzijauskas,S.,& co-authors.(2023).Enhancing skills-demand understanding through job-ad analytics.*AppliedSciences, 13*(10), 6119.

[9] Agarwal,R., & co-authors. (2023). Distilling large languagemodels using skill–occupation graphs for HR tasks. *arXivpreprint*.
[10] Koutsoumpis,A., & co-authors. (2024). Beyond traditional interviews: Psychometric analysis of automated video interviews. *Computers in Human Behavior, 149*.
[11] Nikolaou, I., & co-authors. (2024). Are automated video interviews smart enough? Reliability, validity, and bias of ML-based cognitive assessments. *Preprint / conference working paper*.
[12] Raghavan,M., Barocas, S., Kleinberg, J., & Levy, K. (2024).Fairness and bias in algorithmic hiring: A multidisciplinary survey. *ACM Computing Surveys*.
[13] Trautwein,Y., & co-authors. (2025). Opening the "blackbox" of HRM algorithmic biases. *Journal of Business Research*.
[14] Sharma, P., & co-authors. (2025). Comparing BERT and Sentence-BERT for automated resume screening.*Preprint*.