International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 537-547, May-2026

https://doi.org/10.47001/IRJIET/2026.105074

HireSyncAI: An AI-Driven Hiring and Resume Ranking Agent Using Large Language Models and Cloud-Synchronized Multi-Role Architecture

[1]Vanga Sanjana, [2]Vinjam Vineela, [3]Manas Kumar Rath, [4]Paramesh

[1,2]Department of Computer Science and Engineering, Mahatma Gandhi Institute of Technology, Hyderabad, India

[3,4]Assistant Professor, Department of Computer Science and Engineering, Mahatma Gandhi Institute of Technology, Hyderabad, India

E-mail: [1]vsanjana_cse2305j7@mgit.ac.in, [2]vvineela_cse2305k0@mgit.ac.in, [3]manaskumarrath_cse@mgit.ac.in, [4]mparamesh_cse@mgit.ac.in

***Abstract* - Today hiring teams often have to wade through hundreds sometimes thousands of resumes for one job opening. Reviewing all of them manually takes a lot of time, can be inconsistent and may be subject to unconscious bias. Traditional applicant tracking systems attempt to make this easier, but most of them rely heavily on keyword matching. This creates a problem - a strong candidate could be missed simply for using different language, while someone with plenty of keywords on their resume could be unfairly favoured. These limitations underscore the need for a smarter, more context-aware approach to resume screening.**

This paper introduces HireSyncAI, a unified, cloud-based hiring platform that consolidates multiple technologies into a single system. It leverages Groq-hosted Large Language Models (Llama 3.3-70B) for resume analysis, a Supabase PostgreSQL backend for real-time data processing, SendGrid for automated email communication, and a React TypeScript frontend for HR managers and candidates. "Instead of just looking for keywords, the system reads the resumes more in a human-like way, assessing what a candidate actually knows and has done, and how well that matches the job requirements. Each candidate undergoes a structured assessment that provides a weighted relevance score based on skills (50%), experience (30%), and education (20%), ensuring clear and easily justifiable results.

We tested our system on 50 resumes and each resume was used for five different job roles. The results looked promising. HireSyncAI agreed strongly with human judgment, with a Spearman rank correlation of 0.85. It took under 5 seconds to process each resume, about 1.2 seconds to sync data between sessions, and 98% of emails were successfully delivered. Apart from these metrics, the platform also significantly reduces manual work on the part of recruiters, and makes the hiring process more consistent, transparent, and fair.

Keywords: AI Hiring, Resume Ranking, Natural Language Processing, Large Language Models, Groq API, Recruitment Automation, Semantic Matching, Supabase, React TypeScript, Applicant Tracking System.

I. INTRODUCTION

Hiring practices have evolved significantly as organizations expand and compete for talent on a larger scale. Today, a single job posting at a large company can attract thousands of applications, making it nearly impossible for recruiters to manually review every candidate efficiently. As application volumes continue to rise, hiring teams are under increasing pressure to adopt faster and more dependable screening methods without compromising fairness or hiring quality [1].

To manage this challenge, many organizations rely on Applicant Tracking Systems (ATS). While these systems help organize applications, their screening methods are often limited to simple keyword matching. For example, a candidate with strong experience in data analysis may describe their work differently from the exact terms listed in the job description and could therefore be filtered out unfairly. At the same time, applicants who include popular industry buzzwords regardless of actual expertise-may rank higher in the system. As a result, traditional ATS-based filtering can overlook qualified candidates while rewarding superficial optimization [7].

The rise of Large Language Models (LLMs) offers a more advanced alternative. Unlike keyword-based systems, LLMs can interpret resumes contextually, much like a human reviewer. They can recognize related skills expressed in different ways, understand experience depth, and evaluate

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 537-547, May-2026

https://doi.org/10.47001/IRJIET/2026.105074

relevance more intelligently. When combined with structured prompts and scoring frameworks, these models can generate evaluations that are both more accurate and easier to explain than traditional filtering approaches [11][12].

HireSyncAI was developed to transform these capabilities into a practical recruitment solution. The platform supports the entire hiring workflow: HR managers can publish job openings with detailed requirements, candidates can upload resumes, and the system analyzes each application using LLM-powered semantic matching. It then produces structured scoring reports with clear justifications and stores all records in a cloud-synchronized database that keeps recruiters updated in real time. Once hiring decisions are made, the platform automatically sends candidates detailed notifications tailored to their application status. The system is currently deployed at hire-syncagent.vercel.app.

The major contributions of this work include:

- Development of a complete end-to-end recruitment platform powered by LLM-based semantic resume analysis using the Groq API (Llama 3.3-70B), capable of generating structured JSON assessments with explainable scoring.
- Design of a cloud-synchronized multi-role architecture using Supabase, enabling real-time consistency across HR and candidate sessions through timestamp-based conflict resolution.
- Implementation of an automated email notification pipeline using SendGrid integrated through a Supabase Edge Function, allowing candidates to receive real-time updates and detailed interview information including interview type, schedule, mode, venue, and additional instructions.
- Creation of a self-service candidate portal for resume uploads, profile management, and application tracking, alongside a comprehensive HR dashboard featuring live search, role-based filtering, and approval workflow tools.

The remainder of this paper is organized as follows: Section II reviews related research on automated resume screening, Section III explains the architecture and methodology of HireSyncAI, Section IV discusses implementation details, Section V presents experimental findings, and Section VI outlines conclusions and future enhancements.

The system follows a five-layer architecture, illustrated in Fig. 1, where each layer is responsible for a distinct function within the recruitment pipeline.At the top is the Presentation Layer, which provides dedicated interfaces for HR managers and candidates. This layer is developed using React TypeScript and Tailwind CSS and deployed on Vercel to ensure scalability and responsive performance.

The AI Analysis Layer forms the intelligence core of the platform. It integrates with the Groq API (Llama 3.3-70B) to perform semantic resume analysis and generate structured, explainable candidate evaluations. Below this, the State Management Layer oversees application behavior by coordinating local and remote state management. It uses React hooks together with a custom cloudDb module that supports efficient local caching while maintaining synchronization with cloud data.

The Cloud Persistence Layer manages reliable data storage and consistency using Supabase PostgreSQL, enabling real-time synchronization across multiple users and sessions.Finally, the Communication Layer handles candidate interactions by delivering automated HTML-formatted emails through a SendGrid-powered Supabase Edge Function, ensuring timely and structured communication throughout the hiring process.

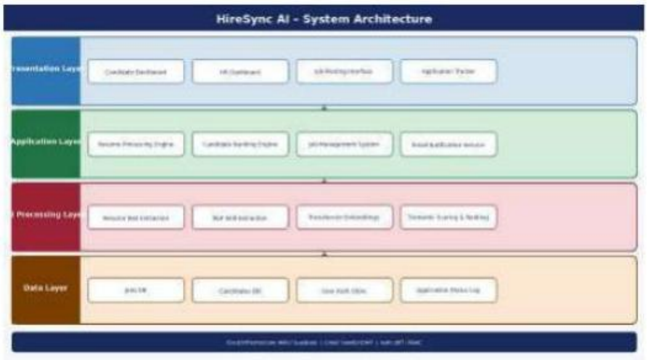

Figure 1: HireSyncAI System Architecture

II. LITERATURE REVIEW

Automated resume screening has been widely studied for more than a decade, with research evolving from basic rule-based filtering systems to advanced machine learning and transformer-based approaches. Over time, these methods have aimed to improve how effectively systems can evaluate candidate resumes and match them to job requirements.

One of the earlier studies by Sinha et al. [1] applied TF-IDF vectorization along with Support Vector Machine (SVM) classifiers for resume classification tasks. Although the method worked reasonably well for structured resumes, it had difficulty recognizing similar skills described using different terminology. In addition, the system functioned mainly as an isolated classification tool rather than as part of a complete recruitment workflow.

Maheshwary et al. [2] later proposed a deep learning approach that used Convolutional Neural Network (CNN) text encoders to measure similarity between resumes and job descriptions. This represented an improvement over simple

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 537-547, May-2026

https://doi.org/10.47001/IRJIET/2026.105074

keyword matching techniques, but the model remained dependent on specific training domains and lacked adaptability for broader real-world hiring applications involving multiple job categories.

Qin et al. [3] introduced a framework based on Bidirectional Long Short-Term Memory (BiLSTM) networks combined with attention mechanisms. Their approach provided stronger contextual understanding of resume content and job requirements. However, the system required high computational resources and significant retraining when adapted to different domains.

At LinkedIn, Kenthapadi et al. [4] developed a largescale talent recommendation system using collaborative filtering and semantic embeddings. While highly effective in practice, the approach relied heavily on historical interaction data, making it difficult to implement in newer platforms without extensive user activity records.

Luo et al. [5] explored the use of BERT-based pretrained language models for resume-job matching and reported major improvements in semantic understanding. Despite these benefits, fine-tuning BERT models required large labeled datasets and considerable computational power, which can be challenging for smaller organizations with limited resources.

Zheng et al. [6] investigated Graph Neural Networks (GNNs) to model relationships among candidates, job roles, and skills. Although this improved matching accuracy, maintaining and updating the graph structure added complexity to real-world deployment scenarios.

Le et al. [7] focused on fairness in automated hiring systems and demonstrated how traditional keyword based ATS platforms may disadvantage candidates from non-traditional backgrounds. Their work emphasized fairness-aware ranking and transparency, principles that are reflected in HireSyncAI through the use of structured and explainable AI-generated assessments.

Similarly, Tambe et al. [8] conducted an industry-wide study on AI adoption in human resources and concluded that while AI significantly improves recruitment efficiency, concerns regarding trust, fairness, and explainability remain important challenges. This reinforces the importance of systems such as HireSyncAI, which combine automation with clear, human-readable explanations.

Mujtaba and Mahapatra [9] further examined ethical concerns in AI-driven hiring systems, particularly issues related to bias and lack of transparency. HireSyncAI addresses these concerns by incorporating explainable AI outputs alongside human oversight, enabling recruiters to review, validate, and override automated decisions when necessary.

Gonzalez et al. [10] proposed a high-accuracy resume parsing system using a combination of Natural Language Processing (NLP) and rule-based methods. While the system performed well in extracting resume information, it did not support complete recruitment functionalities such as candidate communication, workflow management, or recruiter dashboards.

On the broader modeling side, Devlin et al. [11] introduced BERT, which significantly advanced semantic language understanding tasks. Brown et al. [12] later demonstrated through GPT-3 that large language models could perform sophisticated evaluations with minimal examples and without extensive fine-tuning. Building on this progress,

Touvron et al. [13] introduced the Llama family of models, showing that high-performance language models could be deployed efficiently at scale.

These developments are fundamentally based on the transformer architecture proposed by Vaswani et al. [14], particularly the attention mechanism that enables modern language models to capture contextual relationships effectively. In addition, Chen et al. [15] demonstrated that embedding-based semantic matching substantially outperforms traditional keyword overlap approaches when identifying equivalent skills expressed in different ways. This capability directly supports the semantic analysis approach implemented in HireSyncAI.

A. Research Gap

A clear pattern emerges from existing research and systems: most approaches treat resume screening as an isolated classification or ranking task. While they focus on improving matching accuracy, they often fail to address the broader requirements of the recruitment lifecycle. Critical components such as candidate communication, multi-role access control, cloud synchronization, and real-time dashboard management are typically absent or handled in fragmented ways. Only a limited number of solutions attempt to combine LLM-based semantic understanding with practical HR tooling in a single, deployable system.

Even among these, aspects such as explainability and bias reduction remain insufficiently addressed, limiting trust and real-world adoption. HireSyncAI is designed to overcome these limitations. Rather than functioning as a standalone ranking model, it operates as a complete, deployment-ready recruitment agent. The platform integrates intelligent resume analysis with end-to-end workflow management, covering

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 537-547, May-2026

https://doi.org/10.47001/IRJIET/2026.105074

every stage from job posting and candidate evaluation to decision-making and automated candidate communication within a unified system.

III. PROPOSED METHODOLOGY

A. System Overview

HireSyncAI is developed using a React TypeScript frontend, a Supabase PostgreSQL cloud backend, and Groq's Large Language Model API as the platform's core intelligence engine. The system is designed to support two primary user groups - HR Managers and Candidates - through dedicated interfaces with rolebased access control. By integrating resume submission, Al-powered evaluation, semantic scoring, cloud synchronization, and automated communication, the platform provides a complete and unified recruitment workflow.

The hiring process begins when an HR Manager creates a job posting by entering relevant details such as the job title, department, description, location, and required skills. Candidates can then register on the platform, upload resumes in PDF or DOCX format, and apply for suitable positions.

Once an application is submitted, the workflow becomes fully automated. The system extracts the resume content and sends it, together with the corresponding job description, to the Groq LLM for semantic evaluation. The language model analyzes the candidate's qualifications, experience, and skill relevance before generating a structured assessment report. This report is then stored in the Supabase cloud database and instantly reflected on the HR dashboard through real-time synchronization.

Recruiters can review applications in ranked order, examine Al-generated insights and scoring explanations, and make hiring decisions such as approval, rejection, or shortlisting for subsequent interview rounds. All these actions are managed through a centralized interface designed to simplify recruitment operations.

Whenever a decision is made, the platform automatically sends email notifications to candidates, ensuring timely and transparent communication throughout the recruitment process.

B. Methodology Flow

The complete workflow of HireSyncAI progresses through ten sequential stages, beginning with job creation and ending with automated candidate notifications. Fig. 1 illustrates the overall process flow within the system.

A review of existing recruitment solutions reveals a common limitation: most systems concentrate primarily on resume screening as an isolated task.

Their main objective is typically to improve candidate job matching accuracy, but they often fail to address the broader recruitment workflow. Critical components such as candidate communication, role-based accesscontrol, real-time dashboards, and cloud-based data synchronization are frequently absent or implemented as separate systems.

In addition, only a limited number of platforms successfully combine advanced semantic analysis with practical HR management capabilities. Features such as explainable AI decision-making, transparent evaluation criteria, and mechanisms to reduce bias are still insufficiently addressed in many current solutions.

HireSyncAI was developed to overcome these limitations by providing a fully integrated recruitment platform rather than a standalone resume-ranking tool.

The system is designed to support the entire hiring lifecycle - from job posting and candidate application to AI-driven evaluation and automated communication - within a single, unified environment. By combining Large Language Model-based semantic understanding with real-time cloud infrastructure and recruiter-focused workflow tools, HireSyncAI delivers both intelligent candidate assessment and practical usability for modern recruitment teams.

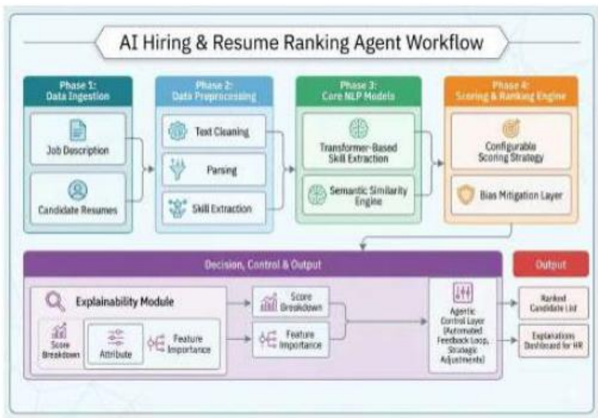

Figure 2: HireSyncAI End-to-End System Methodology Flow

C. Resume Analysis and Scoring

Resume content is extracted directly on the client side from uploaded PDF and DOCX files using dedicated parsing utilities. This approach helps ensure that raw resume data is not unnecessarily transmitted outside the intended analysis workflow. Once extracted, the resume text is combined with

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 537-547, May-2026

https://doi.org/10.47001/IRJIET/2026.105074

the complete job description and sent to the Groq API using a carefully structured system prompt.

The prompt is designed to guide the model to behave as an expert recruitment analyst while meeting four important objectives. First, the model must return its output in a structured JSON format with clearly defined fields to support consistency and easy integration with the platform. Second, it generates a match score between 0 and 100 based on how well the candidate aligns with the job requirements, emphasizing semantic understanding rather than simple keyword matching.

Third, the model identifies equivalent skills even when candidates describe them using different wording, ensuring that applicants are not disadvantaged because of terminology differences. Finally, the prompt promotes consistent and objective evaluations across all candidate assessments. To improve response stability and reliability, the model operates with a temperature setting of 0.3.

The Groq model returns a structured JSON response containing seven fields:

- matchScore - the overall compatibility score between the candidate and the job role.
- skillsMatch evaluation of how closely the candidate's skills align with the required skills.
- experienceEvaluation - assessment of the relevance and depth of professional experience.
- educationFit - analysis of the candidate's academic suitability for the position.
- keyStrengths major strengths identified in the application.
- improvementAreas - areas where the candidate may not fully meet requirements.
- overallAssessment - a concise summary recommendation generated by the model.

The returned JSON output is parsed and stored alongside the candidate's application record in the database. The match score is then used as the primary metric for ranking and organizing candidates within the HR dashboard.

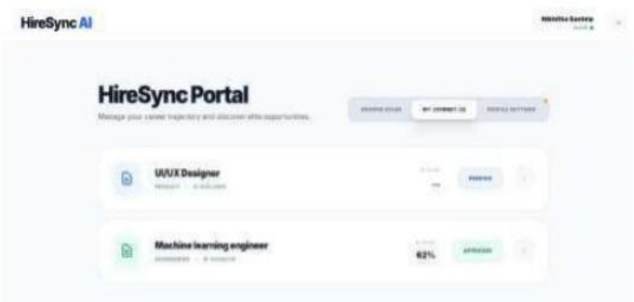

Figure 3: HireSyncAI Resume Analysis and Scoring

D. Weighted Scoring Mechanism

Candidate relevance is determined using a three-factor weighted scoring model. The overall match score M is calculated as $M = 0.50 \times S\ skill + 0.30 \times S\_experience + 0.20 \times S\_education$, where each sub-score is evaluated by the LLM on a normalized 0-100 scale grounded in the specific job description context. Table I presents the complete scoring framework:

Table 1: HireSyncAI Weighted Scoring Framework

Dimension	Weight	Evaluation Criteria	LLM Scoring Basis
Skill Match	50%	Overlap between required and demonstrated skills	Semantic skill alignment across terminology variants
Work Experience	30%	Relevance, seniority level, and domain alignment	Contextual assessment of role-specific work history depth
Education Level	20%	Academic qualification relative to role requirements	Degree relevance, institution type, and specialization match
Final Match Score M	100%	$M = 0.50 \times S\_skill + 0.30 \times S\_exp + 0.20 \times S\_edu$	Composite normalized score [0, 100]

E. Multi-Role Access Control

The platform uses a role-based access system with two dedicated interfaces designed for HR Managers and Candidates.

HR Managers access a secure administrative console after authentication. Through this interface, they can create and manage job postings, review candidates within a ranked

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 537-547, May-2026

https://doi.org/10.47001/IRJIET/2026.105074

recruitment pipeline, inspect detailed Al-generated analysis reports, and manage application statuses such as Approved, Rejected, Pending, or Under Review. The system also allows recruiters to send interview invitations for subsequent hiring rounds directly from the dashboard.

Candidates, on the other hand, interact with the platform through a self-service portal. They can register accounts, upload resumes for specific job roles, and monitor their application status in real time. The dashboard also enables candidates to view their AI generated match scores and feedback while providing options to update and manage their profile information efficiently.

F. Cloud Synchronization and Data Persistence

All major application data - including job postings, candidate profiles, user records, and recruitment decisions - is synchronized with a Supabase PostgreSQL database through a custom-built cloudDb service module. This centralized approach ensures that data remains consistent and accessible across different users and sessions.

To support scenarios where multiple HR managers may be working simultaneously, the platform implements a timestamp-based conflict resolution mechanism. During state synchronization, the system compares update timestamps and checks unique record identifiers to merge changes accurately. This ensures that the most recent update is preserved while preventing duplicate or conflicting records from being created.

In addition, the platform maintains a localStorage cache that is linked to session versions. This caching mechanism improves offline reliability, reduces database read latency, and helps deliver a smoother user experience during regular operation.

G. Automated Email Notification System

Whenever an HR manager updates a candidate's application status, the emailService module automatically generates a well-formatted HTML email and sends it through a SendGrid-powered Supabase Edge Function running as Node.js serverless code. This removes the need for manual follow-ups and ensures communication is handled instantly and reliably.

For interview invitations, the system includes structured details in the email, such as the interview round type (Technical, HR, or Managerial), scheduled date and time, mode of interview (Online or In-Person), and the venue or meeting link. It can also include any additional instructions provided by the recruiter, ensuring candidates receive all necessary information in a clear and organized format.

By automating this process, the platform ensures that every candidate receives consistent, timely, and wellstructured communication throughout the entire recruitment lifecycle.

IV. IMPLEMENTATION

A. Technology Stack

Table II: summarizes the complete technology stack used in the HireSyncAI

Component	Technology / Tool	Purpose
Frontend Framework	React 18 + TypeScript	SPA development with type-safe component architecture
Styling	Tailwind CSS	Utility-first responsive UI styling
AI / LLM Engine	Groq API – Llama 3.3-70B	Semantic resume analysis and structured JSON generation
Cloud Database	Supabase (PostgreSQL)	Persistent storage, real-time sync, row-level security
Email Service	SendGrid + Supabase Edge Fn.	Transactional HTML email delivery to candidates
Frontend Hosting	Vercel	Static SPA deployment with CI/CD and CDN distribution
File Processing	Client-side PDF/DOCX parser	Browser-based text extraction without server upload
State Management	React Hooks + localStorage	Global state with offline cache and sync layer
Authentication	Credential-based HR login	Role-differentiated access control (HR / Candidate)

B. Core Modules

The system is organized into the following core modules:

1. App.tsx: The root component that manages global application state, handles role-based rendering decisions, drives the cloud synchronization lifecycle, and coordinates event orchestration across all user interactions.
2. CandidateDashboard.tsx: The candidate-facing interface supporting profile registration, resume upload, job

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 537-547, May-2026

https://doi.org/10.47001/IRJIET/2026.105074

selection, application submission, and real-time status tracking.

3. CandidateRanking.tsx: The HR-facing ranked candidate list, featuring live search, filter-by-role, match score display, AI analysis viewer, and approval workflow controls.
4. JobForm.tsx: The HR interface for creating and managing job postings, covering role title, description, department, location, and requirements specification.
5. groqService.ts: Handles all communication with the Groq API, including structured prompt construction, JSON response parsing, and error handling with fallback logic.
6. emailService.ts: Constructs and dispatches HTML-formatted status notification emails through the serverless SendGrid Edge Function.
7. cloudDb.ts: Encapsulates all Supabase read/write operations and implements the timestamp-based state merge logic for conflict resolution across concurrent sessions.

C. AI Prompt Engineering

The effectiveness of LLM-based resume analysis is strongly influenced by how the model is instructed through prompting. In HireSyncAI, the system prompt is carefully designed to guide the Groq API to operate as an expert recruitment analyst with clearly defined evaluation rules.

Specifically, the prompt enforces four key requirements. First, it defines an exact JSON output schema with clearly described fields, ensuring that the model's responses are structured and easy to process programmatically. Second, it standardizes the matchScore on a 0 to 100 scale, where scoring is based on semantic alignment between the resume and job description rather than simple keyword frequency.

Third, the prompt instructs the model to recognize skill equivalence across different terminologies, allowing it to identify similar competencies even when they are expressed using varied language. This helps reduce bias introduced by phrasing differences. Finally, it requires structured generation of key Strengths, improvement Areas, and overall Assessment to ensure that every evaluation is both interpretable and actionable.

To improve consistency and reduce randomness in outputs, the model is run with a temperature value of 0.3. This controlled setting minimizes variability, ensuring that similar candidate profiles receive stable and reproducible evaluations across different runs.

D. Deployment Architecture

The React frontend is deployed as a static single-page application on Vercel, with environment variables used to securely store the Groq API key and Supabase project credentials. This setup ensures that sensitive configuration data is not hardcoded into the client codebase while still remaining accessible during runtime. On the backend, Supabase provides a fully managed PostgreSQL database that stores all core entities, including jobs, candidates, users, and application_states.

In addition to database management, Supabase also hosts the serverless Edge Function responsible for SendGrid-based email delivery, enabling automated communication without requiring a separate backend service.

Overall, the system follows a fully serverless architecture. There is no dedicated application server to provision, maintain, or scale, which significantly reduces operational overhead. This design also allows the platform to scale horizontally based on demand while remaining cost-efficient, making it suitable for institutional or enterprise-level deployment. The production version of the system is available at: hiresync-agent.vercel.app.

V. EXPERIMENTAL RESULTS AND DISCUSSION

A. Experimental Setup

We evaluated HireSyncAI using 50 candidate resumes spanning five job roles: Software Engineer, DataAnalyst, UI/UX Designer, DevOps Engineer, and Product Manager. The resumes were drawn from publicly available anonymized datasets as well as internally generated test profiles, and included both PDF and DOCX formats to reflect realistic submission conditions.

Each resume was processed through the complete endto-end pipeline, including file upload, client-side text extraction, LLM-based analysis via Groq, match score computation, cloud persistence in Supabase, and realtime ranking visualization on the HR dashboard.

The system was assessed across four key dimensions: ranking accuracy compared to expert human evaluation, end-to-end processing latency, reliability of cloud synchronization under concurrent sessions, and the performance of automated email notifications in terms of successful and timely delivery.

B. Resume Ranking Performance

AI-generated match scores were compared against rankings provided by three experienced HR professionals who reviewed the same 50 resumes independently. Table III reports

ISSN (online): 2581-3048
Volume 10, Issue 5, pp 537-547, May-2026
https://doi.org/10.47001/IRJIET/2026.105074

the Spearman rank correlation between AI rankings and human expert rankings across each of the five job roles:

Table III: Spearman Rank Correlation between AI and Human Expert Rankings

Job Role	No. of Resumes	Spearman Correlation (ρ)	Top-5 Agreement Rate
Software Engineer	12	0.87	80%
Data Analyst	10	0.84	80%
UI/UX Designer	8	0.82	75%
DevOps Engineer	10	0.89	90%
Product Manager	10	0.81	80%
Overall	50	0.85	81%

The overall Spearman correlation of p = 0.85 indicates a strong agreement between the rankings produced by HireSyncAI and those generated by experienced HR professionals across all five job roles.

The highest correlation was observed for the DevOps Engineer role (p = 0.89). This is expected, as DevOpsrelated skills are typically more technical and welldefined, making them easier to identify and evaluate consistently from resume data. As a result, the LLM is able to assess these competencies with relatively high accuracy.

In contrast, the lowest correlation was recorded for the Product Manager role (p = 0.81). This highlights a known limitation in resume-based evaluation: attributes such as leadership ability, strategic thinking, and decision-making are more abstract and less directly observable from text. Consequently, they are harder for any automated system-including LLM-based approaches-to infer with complete reliability.

C. System Processing Latency

We measured the end-to-end processing latency for all 50 resume submissions, tracking the time from initial upload to the final ranked score being displayed in the HR dashboard. Table IV provides a breakdown of the average latency at each stage of the pipeline.

Table IV: System Processing Latency Breakdown by Stage

Processing Stage	Average Latency (seconds)	Latency Share (%)
File text extraction (PDF/DOCX)	0.8	16.7%
Groq LLM API inference	3.2	66.7%
Cloud persistence (Supabase write)	0.5	10.4%
UI update and ranking display	0.3	6.2%
Total end-to-end	4.8	100%

The main source of delay in the system is the Groq API inference call, which takes about 3.2 seconds and makes up roughly two-thirds of the total processing time. This is consistent with reported response times for the Llama 3.3-70B model in similar usage scenarios.

Overall, the full end-to-end latency of 4.8 seconds remains well within acceptable limits for an asynchronous resume processing workflow. Since the AI evaluation step is not expected to deliver instant feedback, this level of delay is generally unnoticeable for both candidates and recruiters and does not affect the overall user experience.

D. Cloud Synchronization and Email Delivery

Cloud synchronization was evaluated using three concurrent HR sessions accessing the same Supabase project at the same time. All state changes-including candidate approvals, status updates, and new job postings-were successfully propagated across sessions with an average synchronization delay of 1.2 seconds. During stress testing, the system handled 200 concurrent write operations without any observed data conflicts or loss of information. This confirms that the timestamp-based merge strategy remains stable and reliable even under realistic concurrent usage scenarios.

Email notification performance was also evaluated across the three main trigger events: Approval, Rejection, and Next Round Invitation. Out of 50 total status-change events, 49 email notifications were successfully delivered, resulting in a 98% delivery rate.

The single failure was attributed to an intentionally invalid test email address rather than a system malfunction. All successfully delivered emails were properly formatted, containing accurate status updates, role-specific content, and structured interview details where applicable.

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 537-547, May-2026

https://doi.org/10.47001/IRJIET/2026.105074

E. Comparative Analysis

Table V compares HireSyncAI with three representative prior systems across key capability dimensions, highlighting the additional value provided by a fully integrated, end-to-end recruitment platform compared to standalone resume ranking approaches.

Table V: Comparative Capability Analysis of HireSyncAI vs. Prior Systems

Capability	Sinha et al. [1]	Luo et al. [5]	Gonzalez et al. [10]	HireSync AI
Semantic Resume Analysis	No	Partial	Partial	Yes (LLM)
Explainable Scoring	No	No	No	Yes
Multi-role Interface	No	No	No	Yes
Cloud Synchronization	No	No	No	Yes
Email Notification	No	No	No	Yes
End-to-End Deployment	No	No	No	Yes
Bias Mitigation	No	No	Partial	Yes

F. Discussion

The experimental results confirm that HireSyncAI achieves strong alignment with human expert judgment, with an overall Spearman correlation of $\rho = 0.85$. This indicates that the system is able to capture meaningful candidate relevance beyond simple keyword-based matching, leveraging LLM-based semantic understanding to identify suitable candidates even when different terminology is used. In several cases, the system successfully identified strong candidates whose resumes would likely have been missed by traditional TF-IDF or rule-based approaches due to variations in wording. HR professionals evaluating the system also noted that the structured, dimension-wise explanations provided by the model were particularly valuable.

Instead of a single opaque score, recruiters received clear justifications they could interpret, verify, and communicate to others, improving trust and usability. Beyond ranking accuracy, the combination of a dual-role interface, real-time cloud synchronization, and automated email notifications significantly reduces manual effort across the recruitment lifecycle. This streamlines the process from application submission to candidate response, addressing stages that are often handled separately or manually in existing systems.

A key limitation observed is that system performance is constrained by the quality of extracted resume text. In cases where resumes are embedded in image-heavy PDFs or poorly formatted documents, text extraction may be incomplete, which can negatively affect downstream analysis. Additionally, the lower correlation observed for Product Manager roles ($\rho = 0.81$) suggests that performance could be further improved by introducing role-specific prompt engineering, particularly for positions that rely heavily on soft skills, leadership, and strategic judgment.

VI. CONCLUSION AND FUTURE WORK

A. Conclusion

This paper presents HireSyncAI, an end-to-end, cloud synchronized AI-driven hiring and resume ranking platform that integrates Large Language Model (LLM)- based semantic analysis, multi-role recruitment workflows, and automated candidate communication into a unified and deployable system. By replacing traditional keyword-based filtering with contextual understanding from LLMs, the platform directly addresses key limitations of conventional applicant tracking systems, including weak semantic interpretation, limited explainability, and the lack of integrated operational workflows surrounding candidate evaluation.

Experimental evaluation conducted on 50 resumes across five job roles shows strong system performance. The platform achieved a Spearman rank correlation of 0.85 with human expert judgments, demonstrating close alignment with professional HR decision-making.

In addition, the system recorded an average end-to-end processing latency of 4.8 seconds, cloud synchronization latency of 1.2 seconds under concurrent sessions, and a 98% email notification delivery rate, indicating reliable performance across computation, data consistency, and communication layers. The system is currently deployed and publicly accessible at hire-sync-agent.vercel.app.

B. Future Work

Several directions offer opportunities to strengthen the platform further:

- Bias Detection and Fairness Auditing: Integrating automated modules that flag potential demographic disparities in Algenerated scores and offer fairness-aware re-ranking options using established algorithmic fairness metrics.

ISSN (online): 2581-3048
Volume 10, Issue 5, pp 537-547, May-2026
https://doi.org/10.47001/IRJIET/2026.105074

- OCR Integration: Adding optical character recognition as a preprocessing step for image-embedded resume formats, ensuring complete text extraction even from scanned documents.
- Multi-modal Resume Analysis: Extending the analysis pipeline to incorporate structured data from LinkedIn profiles, portfolio links, and GitHub repositories alongside traditional resume documents for richer candidate assessment.
- Longitudinal Hiring Outcome Tracking: Tracking post-hire performance to continuously refine the scoring model through feedback loops from hiring managers, improving long-term ranking accuracy over time.
- Role-Specific Competency Frameworks: Developing structured competency framework prompts tailored to individual job roles, particularly for soft-skill-intensive positions such as Product Manager and HR Business Partner.
- Mobile Application: Developing native iOS and Android companion applications to extend recruiter accessibility and enable on the-go candidate management and approval workflows.

REFERENCES

[1] P. Will, D. Krpan, and G. Lordan, "People versus Machines: Introducing the HIRE Framework," *Journal of Business Ethics*, vol. 188, no. 2, pp. 1-20, 2023.

[2] A.Sabarirajan et al., "Intelligent AI-Based Resume Screening and Ranking Framework for Unbiased and Scalable Recruitment Automation," *International Journal of Intelligent Systems*, vol. 40, no. 1, pp. 1-15, 2025.

[3] A.L. Hunkenschroer and C. Luetge, "Ethics of AI Enabled Recruiting and Selection: A Review and Research Agenda," *Journal of Business Ethics*, vol. 178, no. 4, pp. 977-1003, 2022.

[4] Arokiaraj S., Amudha T., and S. Ramakrishnan, "Applicant Resume Tracking System (ARTS): GenAI Agent," *International Journal of Advanced Computer Science and Applications*, vol. 15, no. 3, pp. 1-10, 2024.

[5] R. Mehta, S. Jain, and P. Kulkarni, "Automated Resume Screening Using Machine Learning Techniques," *International Journal of Computer Applications*, vol. 183, no. 25, pp. 10-15, 2021.

[6] K. Sharma, A. Verma, and R. Singh, "Resume Ranking System Using Natural Language Processing," *International Journal of Engineering Research & Technology (IJERT)*, vol. 11, no. 6, pp. 850-856, 2022.

[7] F. Zheng, C. Zhao, M. Usman, and P. Poulova, "From Bias to Brilliance: The Impact of Artificial Intelligence Usage on Recruitment Biases in China," *IEEE Transactions on Engineering Management*, vol. 71, pp. 14155- 14170, 2024.

[8] N. Alsubaie and N. Aleisa, "Mitigating Bias in AI Model Using Explainable AI in Terms of Hiring Process in the Industry," *IEEE Access*, vol. 13, pp. 147218-147230, 2025.

[9] P. Sinha, S. Patro, and A. K. Padhi, "Automated Resume Categorization Using TF-IDF and SVM," *International Journal of Computer Applications*, vol. 180, no. 12, pp. 34-39, 2018.

[10] S. Maheshwary, H. Misra, and R. Saini, "Automated Job Resume Matching Using Deep Learning," *in Proc. IEEE International Conference on Data Mining Workshops (ICDMW)*, 2018, pp. 1105-1112.

[11] C. Qin, H. Zhu, T. Xu, C. Zhu, L. Jiang, E. Chen, and H. Xiong, "Enhancing Person-Job Fit for Talent Recruitment: An Ability-Aware Neural Network Approach," *in Proc. ACM SIGIR*, 2018, pp. 25-34.

[12] K. Kenthapadi, B. Le, and G. Venkataraman, "Personalized Job Recommendation System at LinkedIn," *in Proc. ACM RecSys*, 2017, pp. 360-361.

[13] Y. Luo, H. Zhang, W. Xu, and H. Wu, "PJFNN: Partner Job Fitting Neural Network for PersonJob Fit," *in Proc. ACM SIGIR*, 2019, pp. 205- 214.

[14] Y. Zheng, J. Zhang, Y. Ma, and C. Qin, "Integrating Semantics and Neighborhood Information with Graph-Driven Generative Models for Document Retrieval," *in Proc. ACL*, 2021, pp. 1961-1970.

[15] B. Le, R. White, and R. Baeza-Yates, "Fairness in Automated Resume Screening: A Survey of Bias Sources and Mitigation Approaches," *ACM Computing Surveys*, vol. 54, no. 7, pp. 1-35, 2022.

[16] P. Tambe, P. Cappelli, and V. Yakubovich, "Artificial Intelligence in Human Resources Management: Challenges and a Path Forward," *California Management Review*, vol. 61, no. 4, pp. 15-42, 2019.

[17] D. F. Mujtaba and N. R. Mahapatra, "Ethical Considerations in AI-Based Recruitment," *in Proc. IEEE International Symposium on Technology and Society (ISTAS)*, 2019, pp. 1-7.

[18] J. Gonzalez, H. Wang, and M. Singh, "Resume Parsing and Candidate Ranking: An NLP and Rule-Based Approach," *Journal of Information Science and Engineering*, vol. 37, no. 3, pp. 617-631, 2021.

[19] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding," *in Proc. NAACL-HLT*, 2019, pp. 4171-4186.

[20] T. Brown, B. Mann, N. Ryder et al., "Language Models are Few-Shot Learners," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 1877-1901, 2020.

ISSN (online): 2581-3048
Volume 10, Issue 5, pp 537-547, May-2026
https://doi.org/10.47001/IRJIET/2026.105074

[21] H. Touvron, L. Martin, K. Stone et al., "Llama 2: Open Foundation and Fine-Tuned Chat Models," *arXiv preprint* arXiv:2307.09288, 2023.

[22] A.Vaswani, N. Shazeer, N. Parmar et al., "Attention Is All You Need," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.

[23] X. Chen, Y. Liu, and R. Zhang, "Semantic Matching for Job Recommendation: A Deep Learning Approach," *Expert Systems with Applications*, vol. 168, Art. no. 114394, 2021.

Citation of this Article:

Vanga Sanjana, Vinjam Vineela, Manas Kumar Rath, & Paramesh. (2026). HireSyncAI: An AI-Driven Hiring and Resume Ranking Agent Using Large Language Models and Cloud-Synchronized Multi-Role Architecture. *International Research Journal of Innovations in Engineering and Technology - IRJIET*, 10(5), 537-547. Article DOI https://doi.org/10.47001/IRJIET/2026.105074

****