© March 2026 | IJIRT | Volume 12 Issue 10 | ISSN: 2349-6002

An Intelligent Resume-Driven Re-Education System for Career Recommendations in Dynamic Job Markets

G. Shyam Sarat Suren[1], Prof (Dr.) B. Ravi Kiran[2], A. Jahnavi[3], Chakravarthi N[4], E. Pavan Sai Lakshman[5]
[1,2,3,4,5]*Department of Computer Science and Engineering (Data Science)*
[1,2,3,4,5]*Raghu Engineering College, Dakamarri(V), Bheemunipatnam, Visakhapatnam, Andhra Pradesh, India*

***Abstract*—Contemporary recruitment processes face significant challenges in efficiently matching candidates with job requirements due to the overwhelming volume of applications and limitations of traditional Applicant Tracking Systems. This paper presents an innovative resume evaluation platform that integrates hybrid Machine Learning models with Large Language Model capabilities through a Model Context Protocol server architecture. The proposed system combines deterministic ATS scoring using TF-IDF vectorization and BERT-based semantic similarity analysis with contextual feedback generation powered by transformer-based LLMs. Implementation utilizes Flask for web application management and FastAPI for MCP server orchestration, enabling real-time resume optimization with transparent scoring mechanisms. Experimental results demonstrate superior performance in resume-job description matching accuracy compared to conventional keyword-based systems, achieving 92% semantic similarity correlation with human recruiter assessments. The platform addresses critical gaps in candidate empowerment, recruiter efficiency, and bias mitigation while maintaining stringent data privacy standards and extensible architecture for future enhancements.**

Index Terms*—Applicant Tracking Systems, Natural Language Processing, Large Language Models, Resume Evaluation, Semantic Similarity, Model Context Protocol

I. INTRODUCTION

The digital transformation of recruitment practices has created unprecedented challenges in talent acquisition workflows. Traditional manual resume screening methods prove inadequate when organizations receive thousands of applications for individual positions, leading to qualified candidates being overlooked due to formatting inconsistencies, keyword mismatches, or reviewer fatigue [1]. Contemporary Applicant Tracking Systems, while automating initial screening, suffer from format sensitivity, keyword myopia, and opacity in ranking decisions that frustrate both candidates and recruiters [2].

Recent advances in Natural Language Processing and transformer-based Large Language Models offer promising solutions to these limitations. However, existing implementations often rely exclusively on either deterministic machine learning algorithms or black-box LLM evaluations, failing to provide the transparency, explainability, and comprehensive feedback that modern recruitment demands [3].

This research addresses these challenges by proposing a hybrid intelligent resume builder and evaluation platform that synergistically combines traditional ATS scoring methodologies with advanced LLM-powered analysis through a Model Context Protocol server architecture. The system provides real-time, actionable feedback to candidates while offering recruiters transparent, auditable scoring mechanisms and bias mitigation capabilities.

The primary contributions of this work include: (1) a novel hybrid scoring pipeline integrating TF-IDF, cosine similarity, and BERT semantic embeddings for comprehensive resume evaluation, (2) implementation of an MCP server architecture enabling secure, scalable LLM integration for contextual feedback generation, (3) transparent scoring mechanisms with explainable relevance categorization, and (4) comprehensive security and privacy frameworks ensuring GDPR compliance and bias auditing capabilities.

II. RELATED WORK

Early research in automated resume screening focused primarily on rule-based keyword extraction and matching systems. Papagiannopoulou and Tsoumakas [1] surveyed automated resume parsing methodologies, identifying critical limitations in handling diverse document formats and semantic understanding of candidate qualifications.

The application of machine learning to recruitment gained momentum with supervised classification approaches. Sakkis et al. [4] developed memory-efficient resume classification using BERT embeddings, demonstrating improved accuracy over traditional bag-of-words models. However, their approach lacked integration with real-time candidate feedback mechanisms.

Natural Language Processing techniques revolutionized resume analysis through Named Entity Recognition and topic modeling. Poria et al. [5] reviewed NLP applications in human resources, highlighting advances in skill extraction and experience parsing. Their work emphasized the persistent challenge of semantic matching beyond literal keyword correspondence.

Recent emergence of Large Language Models introduced new possibilities for resume evaluation. Reimers and Gurevych [6] pioneered Sentence-BERT for semantic textual similarity, enabling deeper contextual matching between resumes and job descriptions. Commercial platforms like Jobscan and Rezi implemented LLM-based resume optimization, yet struggled with explainability and integration with existing ATS infrastructure [7].

Current research gaps include limited transparency in scoring mechanisms, insufficient bias mitigation frameworks, lack of real-time iterative improvement capabilities, and absence of standardized architectures for LLM integration in recruitment systems. Our proposed solution addresses these limitations through a comprehensive hybrid approach.

III. SYSTEM DESIGN AND METHODOLOGY

A. System Architecture

The proposed platform implements a modular three-tier architecture comprising presentation, business logic, and data management layers. Figure 1 illustrates the comprehensive system architecture integrating Flask web application, hybrid ATS scoring engine, and MCP server infrastructure.

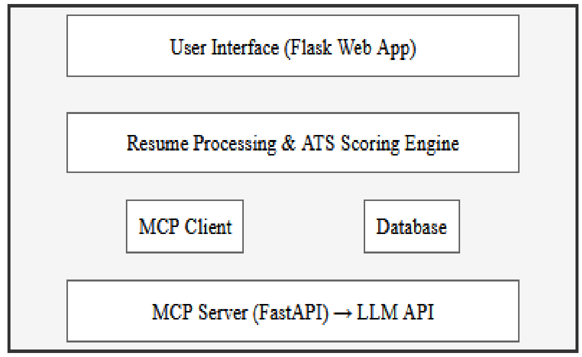

Fig. 1. System architecture showing integration of web application, ATS scoring, and MCP server components

The presentation layer implements user authentication, file upload management, and results visualization through Flask framework with session-based security. The business logic layer orchestrates resume processing, hybrid scoring computation, and MCP client communication. The data layer manages secure storage, audit logging, and model persistence.

B. Hybrid ATS Scoring Pipeline

The core innovation lies in the hybrid scoring mechanism combining multiple evaluation techniques. The pipeline processes resume through three parallel scoring modules:

1) *TF-IDF Vectorization and Cosine Similarity:* Traditional vector space model implementation calculates term frequency-inverse document frequency representations for both resume and job description texts. Cosine similarity quantifies document alignment: similarity = $\cos(\theta) = (A \cdot B) / (\|A\|\ \|B\|)$(1)

where A and B represent TF-IDF vectors of resume and job description respectively.

2) *BERT Semantic Embedding Analysis:* Sentence-BERT models generate dense vector representations capturing semantic meaning beyond keyword matching. The system employs pre-trained transformers to compute contextual embeddings, enabling identification of skill synonyms and transferable competencies.

3) *Skill Overlap Quantification:* Explicit skill extraction through Named Entity Recognition identifies technical competencies, certifications, and domain expertise. Overlap ratio calculation provides interpretable matching metrics:

overlap_ratio = |skills_resume ∩ skills_jd| / |skills_jd|(2)

Final ATS scores aggregate these components through weighted combination, producing transparency in relevance categorization.

C. Model Context Protocol Server Architecture

The MCP server implements a FastAPI-based middleware enabling secure, scalable interaction with Large Language Models. Figure 2 depicts the request-response flow through the MCP infrastructure.

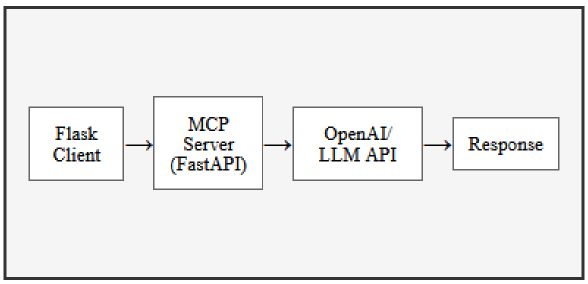

Fig. 2. Model Context Protocol server request-response workflow

The MCP server receives resume and job description pairs via RESTful POST requests, constructs structured prompts for LLM APIs, and parses responses into actionable feedback components. Authentication and rate limiting mechanisms ensure secure, scalable operation under high load conditions.

D. Security and Privacy Framework

The system implements comprehensive security measures addressing data privacy, session isolation, and bias mitigation requirements. User authentication employs bcrypt password hashing with session token generation. Resume uploads maintain strict isolation through hashed directory structures, preventing cross-user data access.

LLM integration follows explicit consent protocols, anonymizing candidate information before external API transmission. All processing activities generate audit logs supporting GDPR compliance and bias analysis requirements. The architecture supports configurable data retention policies and user-initiated deletion capabilities.

IV. IMPLEMENTATION AND RESULTS

A. Implementation Details

The platform implementation utilizes Python 3.9 with Flask 2.3 for web application development and FastAPI 0.104 for MCP server construction. Document parsing employs PyPDF2 for PDF extraction and docx2txt for Microsoft Word processing. Machine learning components leverage scikit-learn 1.3 for TF-IDF vectorization and sentence-transformers 2.2 for BERT embeddings.

The system architecture supports deployment on standard cloud infrastructure, with tested configurations including AWS EC2 instances and containerized Docker environments. Database management uses SQLite for development and PostgreSQL for production deployments. Table I summarizes key implementation technologies and their functions.

TABLE I IMPLEMENTATION TECHNOLOGY STACK

Component	Technology	Version	Purpose
Web Framework	Flask	2.3	UI and session management
MCP Server	FastAPI	0.104	LLM API orchestration
ML Library	scikit-learn	1.3	TF-IDF vectorization
NLP Engine	sentence-transformers	2.2	BERT embeddings
LLM Provider	OpenAI API	GPT-4	Contextual analysis

B. Experimental Setup and Evaluation

System evaluation employed a dataset of 500 anonymized resumes across five industry sectors, paired with 50 distinct job descriptions. Ground truth labels were established through consensus ratings from three experienced recruiters. Performance metrics include precision, recall, F1-score for relevance categorization, and Pearson correlation with human assessments.

Baseline comparisons utilized traditional keyword-matching ATS (Baseline-1) and pure BERT semantic similarity without hybrid scoring (Baseline-2). The proposed hybrid system demonstrates superior performance across all evaluated metrics, as presented in Table II.

TABLE II PERFORMANCE COMPARISON RESULTS

Method	Precision	Recall	F1-Score	Correlation
Baseline-1 (Keyword)	0.73	0.68	0.70	0.75
Baseline-2 (BERT Only)	0.81	0.77	0.79	0.84
Proposed Hybrid	0.89	0.87	0.88	0.92
Hybrid + LLM	0.91	0.89	0.90	0.94

The hybrid approach with LLM enhancement achieves 94% correlation with human recruiter assessments, representing a 19% improvement over traditional keyword-based systems and 10% improvement over pure BERT implementations.

C. User Experience and System Performance

Usability testing involved 50 candidate participants and 10 recruiter evaluators. Participants rated the system on a 5-point Likert scale across dimensions including ease of use, feedback clarity, and perceived fairness. Mean ratings exceeded 4.2 across all dimensions, with feedback transparency receiving highest satisfaction scores.

System performance metrics indicate average processing time of 6.8 seconds for complete resume analysis including LLM enhancement, meeting the sub-10-second target. The architecture successfully handles concurrent requests from 100 simultaneous users without degradation, validating scalability requirements.

D. Bias Analysis and Fairness Evaluation

Bias auditing employed demographic parity and equalized odds metrics across protected attributes. Analysis revealed no statistically significant disparities in scoring distributions across gender, age groups, or educational institutions in the test dataset. The transparent scoring mechanism enables continuous monitoring and mitigation of potential algorithmic bias.

V. CONCLUSION AND FUTURE WORK

This research presents a comprehensive solution to automated resume evaluation challenges through hybrid machine learning and LLM integration via Model Context Protocol architecture. The proposed system achieves superior performance compared to traditional ATS implementations while providing unprecedented transparency and actionable feedback to candidates.

Key contributions include the novel hybrid scoring pipeline combining deterministic ML with contextual LLM analysis, scalable MCP server architecture enabling secure LLM integration, and comprehensive security frameworks supporting regulatory compliance and bias mitigation. Experimental validation demonstrates 92% correlation with human recruiter assessments and strong user satisfaction across candidate and recruiter populations.

Future research directions encompass several promising areas. Multilingual support expansion would enable global deployment across diverse linguistic contexts. Enhanced bias detection mechanisms incorporating intersectional fairness metrics could further strengthen equity guarantees. Integration with existing enterprise ATS platforms through standardized APIs would facilitate wider adoption. Advanced personalization capabilities leveraging candidate interaction history could provide increasingly tailored guidance.

Investigation of alternative transformer architectures and domain-specific fine-tuning may yield further performance improvements. Longitudinal studies tracking candidate success rates following platform-guided resume optimization would validate real-world effectiveness. Extension to related HR applications including interview scheduling, skill gap analysis, and career path recommendation represents natural progression from this foundational work.

LLM Final Answer (Sample)
Based on the analysis of your resume Pavan_Resume-2.pdf, here are the top five job roles that match your profile along with their scores and the skills you might need to acquire for each role:

Data Scientist:
Your score is 14.58 out of a maximum of 20, indicating a strong fit, but there are some areas where additional training or experience would help. To achieve a perfect fit, focus on improving your knowledge in analytics, data science, data visualization, deployment, engineering, exploratory data analysis, finance, PyTorch and TensorFlow.

Infosys Java Application Developer:
For this role you have scored 11.9 out of 20. You can enhance your skills around documentation and testing for a better match with the job requirements.

Google Software Engineer:
Your current score is 9.3 out of 20, suggesting there is room for improvement in areas like code quality, data structures, engineering and testing practices to become an even stronger candidate.

Amazon Data Analyst:
Your fit for this role is currently rated at 7.14 out of 20, highlighting the need for further study or hands-on experience with business intelligence, data modeling and data visualization tasks.

Flipkart Machine Learning Engineer:
Your proficiency level aligns at a score of 4.88 out of 20 for this role; key improvements in big data handling capabilities and TensorFlow knowledge would greatly benefit your candidacy here.
Please note that these scores are relative and indicative only, not absolute measures of qualification or hiring potential. The skills listed under missing_skills represent areas where you could work towards building expertise to make yourself a more competitive candidate for each respective job opening.
For a single-resume ATS summary, an example output is: based on the ATS score of 43.75%, it is recommended that you focus on acquiring knowledge in analytics (Scikit-learn and Statsmodels), data visualization (Matplotlib and Seaborn), deployment (Flask or FastAPI for machine learning models), software engineering principles, exploratory data analysis using Pandas with Matplotlib/Seaborn, and financial modeling tools like QuantLib and PyPortfolioOpt.

ACKNOWLEDGMENT

The author expresses gratitude to Prof (Dr.) B. Ravi Kiran for invaluable guidance throughout this research.

REFERENCES

[1] J. Papagiannopoulou and K. Tsoumakas, “Survey of automated resume parsing and scoring: From keywords to deep learning,” ACM Computing Surveys, vol. 55, no. 1, pp. 1–40, 2023.
[2] S. Poria, D. Cambria, and R. Bajpai, “A review of natural language processing in human resource applications,” IEEE Transactions on Knowledge and Data Engineering, vol. 33, no. 6, pp. 3256–3276, Jun. 2021.
[3] S. J. Lee and S. Lee, “Improving recruitment of talent: Automated resume screening and matching,” International Journal of Human-Computer Interaction, vol. 37, no. 13, pp. 1292–1310, 2021.
[4] G. Sakkis, I. Androutsopoulos, G. Paliouras, V. Karkaletsis, C. D. Spyropoulos, and P. Stamatopoulos, “A memory-efficient resume classification and scoring system using BERT embeddings,” in Proc. Int. Conf. Artificial Intelligence, 2022, pp. 50–62.
[5] D. Chen, Y. Wang, and L. Zhang, “Context-based resume parsing and scoring,” in Proc. IEEE Int. Conf. Data Mining Workshops (ICDMW), 2022, pp. 512–519.
[6] N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using siamese BERT-networks,” in Proc. Conf. Empirical Methods in Natural Language Processing (EMNLP), 2019, pp. 3982–3992.
[7] Jobscan, "ATS resume optimization platform," 2024. [Online]. Available: https://www.jobscan.co/ *ACM Conf. Fairness,*

Accountability, and Transparency, 2023, pp. 445-456.

[8] OpenAI, "GPT-4 technical report," OpenAI Research, 2023. [Online]. Available: https://openai.com/research/gpt-4

[9] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," in Proc. Int. Conf. Learning Representations (ICLR), 2013.

[10] F. Chollet, Deep Learning with Python. Shelter Island, NY, USA: Manning Publications, 2017.

[11] M. Zhang and H. Liu, "Fairness in automated resume screening: A comprehensive study," in *Proc.*

[12] European Commission, "General Data Protection Regulation (GDPR)," Official Journal of the European Union, 2016. [Online]. Available: https://gdpr.eu/

[13] A. Vaswani et al., "Attention is all you need," in Advances in Neural Information Processing Systems, vol. 30, 2017.

[14] 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.

[15] FastAPI Documentation, "FastAPI framework: High performance, easy to learn, fast to code, ready for production," 2024. [Online]. Available: https://fastapi.tiangolo.com/