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

An Intelligent Multi-Modal Resume Screening System Using LLaMA-Based Semantic Embeddings, Weighted Skill Extraction, and Explainable Scoring

Dr. Farheen Mohammed[1], Yakkaladevi Bhuvanesh[2], Shaik Shameem[3], Telikutla Chandana[4], Pentyala Punith Venkata Sai[5]

[1]Assistant Professor, Dept. of CSE (AIML), Bapatla Engineering College, Bapatla 522101, AP, India
[2345]Student, Dept. of CSE (AIML), Bapatla Engineering College, Bapatla 522101, AP, India
mffarheen@gmail.com shameemsk002@gmail.com yakkaladevibhuvanesh@gmail.com
t.chandanareddy20@gmail.com punithchowdary52@gmail.com

---***---

Abstract - The growing volume of job applications poses a significant challenge for HR teams, making manual resume screening both time-consuming and error-prone. To address this, we propose the development of an intelligent, automated resume screening system based on LLaMA language models, designed to outperform traditional BERT-based approaches. The system leverages LLaMA's advanced semantic understanding to generate embeddings of both resumes and job descriptions, using cosine similarity and a threshold-based retrieval mechanism to rank and filter candidates efficiently. To ensure transparency and build trust in automated decisions, the model incorporates Explainable AI (XAI) techniques, allowing HR professionals to understand the rationale behind each recommendation. This approach aims to significantly reduce screening time, increase accuracy, and support fairer, more interpretable hiring practices. Ultimately, the proposed system seeks to streamline recruitment workflows while upholding accountability and equity in candidate evaluation.

Index Terms*—*Resume Screening, LLaMA, Semantic Embeddings, Cosine Similarity, Explainable AI, Recruitment Automation, Natural Language Processing, Candidate Ranking, Email Notification.

I. INTRODUCTION

The contemporary recruitment landscape is characterized by an exponential growth in job applications per vacancy, driven largely by online job portals and social professional networks. Organizations regularly receive hundreds to thousands of applications for a single position, rendering manual screening practically infeasible without significant recruiter burden and associated subjectivity [6, 8]. Automated resume screening systems therefore represent a strategic necessity rather than a mere technological convenience.

Early attempts at automation relied on keyword matching and regular expression-based parsers that extracted structured fields from standardized resume templates [3, 4]. While computationally efficient, such approaches suffered from brittleness in the presence of diverse formatting, domain-specific terminology, and synonymous skill expressions. The emergence of pre-trained language models and dense vector representations substantially altered the landscape, enabling semantic-level matching between job requirements and candidate profiles [1, 7].

Named Entity Recognition (NER) approaches introduced the capacity to identify fine-grained entities such as skills, educational qualifications, and organization names within free-form text [3, 18]. Hybrid systems that combine NER-based structured extraction with transformer-derived embeddings have demonstrated superior generalization compared to rule-based baselines [10, 11]. However, the majority of existing work focuses primarily on the extraction phase, with limited attention to the downstream scoring and ranking pipeline that directly determines shortlisting decisions.

This paper addresses the aforementioned gap by presenting an end-to-end intelligent resume screening framework comprising: (i) a seven-layer modular pipeline from document ingestion to recruiter dashboard; (ii) LLaMA-3 based semantic embedding generation for high-fidelity query and passage representations; (iii) a weighted skill matching module that differentiates between mandatory and preferred skill requirements; (iv) a contextual experience scoring function that integrates skill density into experience relevance; (v) a power-law calibrated cosine similarity metric that reduces spurious high baseline similarities; (vi) an alignment factor that penalizes missing critical requirements without introducing discontinuous score distributions; and (vii) an automated email notification module that informs shortlisted and rejected candidates of their status upon pipeline completion.

The remainder of this paper is structured as follows. Section II reviews related work. Section III presents the system methodology. Section IV describes the experimental setup and datasets. Section V reports and analyzes results. Section VI discusses implications and limitations. Section VII concludes the paper and outlines future directions.

II. RELATED WORK

A. Early Parsing and Rule-Based Approaches

Initial automated resume screening systems relied extensively on template matching, regular expressions, and heuristic rules to extract structured data from semi-structured documents [3, 18]. Narendra and Hashwanth [3] proposed an NER-based resume parser that leveraged named entity classifiers trained on domain-specific corpora, achieving promising extraction accuracy for entities such as job titles, skills, and educational institutions. Palshikar et al. [4] presented RINX, a comprehensive information and knowledge extraction framework from resumes, integrating ontology-based reasoning with statistical classifiers.

B. Deep Learning and Transformer-Based Methods

The widespread adoption of transformer architectures catalyzed a paradigm shift in resume understanding. Gan and Mori [2] proposed a few-shot prompting strategy for resume information extraction, demonstrating that large language models could extract structured fields with minimal labeled data. Barducci et al. [1] developed an end-to-end information extraction framework for Italian resumes using deep learning pipelines fine-tuned on a multilingual corpus. Jiang et al. [8] introduced a multi-granularity multi-modal pre-training approach that jointly modeled textual and layout features, yielding significant improvements on benchmark extraction tasks.

C. Hybrid and Pipeline-Based Systems

Several recent investigations have explored hybrid architectures combining multiple complementary extraction and matching paradigms. Ed-Daoudi et al. [10, 19] presented a hybrid OCR-XGBoost-Transformer pipeline for resume information extraction that addressed the challenge of scanned PDF inputs through a cascaded architecture. Myneni and Quadhari [11] proposed a hybrid transformer-based resume parsing and job matching framework utilizing TextRank for salient phrase extraction, SBERT for semantic encoding, and DeBERTa for reranking shortlisted candidates.

Khelkhal and Lanasri [12] introduced Smart-Hiring, an explainable end-to-end pipeline for CV information extraction and job matching that incorporated SHAP-based explanations to improve recruiter trust in automated decisions. Bala and Quadhari [13] extended transformer-based parsing with a dedicated job matching system trained on domain-specific paraphrase pairs. The DeepResume project [15] proposed a deep learning-based resume parsing and candidate ranking system evaluated on industry-sourced datasets.

D. NLP-Centric Screening Systems

Saatci, Kaya, and Ünlü [7, 14, 16, 17] published a series of investigations on NLP-based resume screening covering text block classification, NER-based field extraction, and multi-model pipeline architectures for recruitment automation. Kesim [5] evaluated named entity recognition approaches in resume processing, systematically benchmarking NER model families on resume-domain test sets. Kumar et al. [6] developed an AI-powered resume parser integrating transformer-based extraction with downstream ranking heuristics. Deepa [9] provided a comprehensive survey of automated resume parsing techniques, challenges including format diversity and domain shift, and emerging application areas.

III. METHODOLOGY

A. System Architecture Overview

The proposed system follows a seven-layer modular pipeline architecture, illustrated in Fig. 1, comprising the following principal stages: (1) Resume & Job Description Input Layer, (2) Resume Parsing & Data Extraction Layer, (3) Text Preprocessing Layer, (4) Semantic Embedding Layer using the LLaMA-3 model, (5) Similarity & Ranking Layer, (6) Explainable AI Layer (Transparency Module), and (7) Recruiter Dashboard & Interaction Module.

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

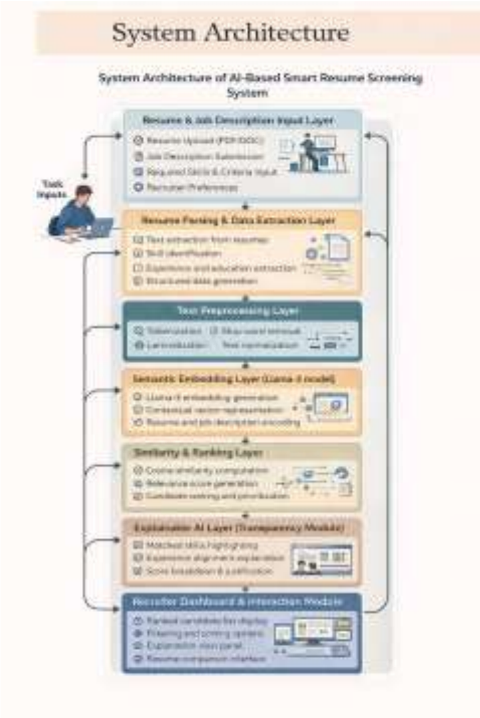

Figure 1: System pipeline: ingestion to recruiter dashboard output.

The system accepts PDF and DOC formatted resumes as primary inputs, consistent with standard industry submission formats [8, 10]. Text is extracted via pdfplumber and PyMuPDF, with an OCR fallback for scanned documents. The extracted plain text is then routed through successive layers to the embedding engine and scoring pipeline.

B. Resume & Job Description Input Layer

The input layer serves as the primary interface between the recruiter and the system. It collects resume uploads in PDF or DOC format, free-form job description text submissions, structured mandatory and preferred skill criteria inputs, and recruiter preferences including configurable experience thresholds, domain focus, and shortlist quota. Input validation ensures that uploaded files conform to supported formats and that job descriptions meet minimum length requirements before routing to the parsing layer.

C. Resume Parsing & Data Extraction Layer

The parsing layer is responsible for converting raw binary document inputs into structured, machine-readable representations. The layer performs the following operations: text extraction from resumes via pdfplumber and PyMuPDF with OCR fallback; skill identification using a curated domain ontology combined with NER-based phrase extraction; experience and education extraction through pattern-based identification of tenure periods and academic qualifications; and structured data generation through serialization of extracted entities into a canonical JSON representation for downstream consumption.

D. Text Preprocessing Layer

Prior to embedding generation, all extracted text passes through a four-stage normalization pipeline to reduce surface-form variation and vocabulary dimensionality. The stages are: (i) Tokenization—segmentation of raw text into word and subword units; (ii) Stop-word removal—elimination of high-frequency, low-informativeness tokens; (iii) Lemmatization—reduction of inflected word forms to their canonical lemma; and (iv) Text normalization—lowercasing, punctuation stripping, and Unicode normalization. This preprocessing step ensures that downstream embedding models receive clean, standardized input, reducing spurious similarity inflation caused by orthographic variation.

E. Semantic Embedding Layer (LLaMA-3 Model)

Dense text representations are generated using the LLaMA-3 language model. LLaMA-3 offers several structural advantages over BERT-based predecessors: a substantially larger context window enabling full resume ingestion without truncation artifacts, improved instruction-following capacity that permits asymmetric retrieval prompting, and superior performance on semantic textual similarity benchmarks. Both resume passages and job description queries receive distinct prompt prefixes to leverage the asymmetric retrieval objective, producing 768-dimensional contextual embeddings that capture long-range semantic dependencies across the full document.

The raw cosine similarity between a resume embedding $r \in R^d$ and a job description embedding $j \in R^d$ is defined as:

$$S_{cos}(r, j) = (r \cdot j) / (\|r\| \cdot \|j\|) \quad \dots (1)$$

A well-known limitation of cosine similarity on pre-trained transformer embeddings is the presence of an elevated baseline similarity (≈0.3) between semantically unrelated texts, arising from the anisotropy of the embedding space [9, 11]. To address this, a power-law calibration is applied:

$$S_{sem} = S_{cos}^{2.5} \quad \text{if } S_{cos} > 0, \text{ else } 0 \quad \dots (2)$$

This calibration suppresses the spurious baseline similarity from ≈0.3 to ≈0.046 while preserving the relative ordering of genuine semantic matches with minimal distortion at high similarity values [2, 7].

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

F. Similarity & Ranking Layer

The similarity and ranking layer computes a multi-criteria composite score for each candidate through three sub-components: cosine similarity computation using power-law calibrated LLaMA-3 embeddings, weighted skill-based relevance score generation, and final candidate ranking and prioritization against the configurable shortlisting threshold.

The skill extraction pipeline identifies technical and domain-specific competencies using a curated skill ontology combined with NER-based phrase extraction [3, 5]. Extracted skills are normalized to lowercase. Job descriptions specify two tiers of requirements: mandatory skills M and preferred skills G. The weighted skill score is:

$$S_{skill} = (2\cdot|M\cap R| + |G\cap R|) / (2\cdot|M| + |G|) \quad \dots (3)$$

where R denotes the skill set extracted from the candidate resume. The factor of two applied to mandatory skills reflects the disproportionate importance of core job requirements [6, 13]. When no skills are specified, Sskill defaults to 0.5 to avoid score collapse.

Candidate experience ec (in years) is compared against required experience er from the job description:

$$S_{exp} = 1.0 \text{ if } er=0 \text{ and } ec>0;\ \ 0.8 \text{ if } er=0 \text{ and } ec=0;\ \ 1.0 \text{ if } ec\geq er;\ \ ec/er \text{ otherwise} \quad \dots (4)$$

A contextual experience modifier prevents overvaluing raw tenure when a candidate lacks critical skills [11, 12]:

$$S_{ctx} = S_{exp} \cdot (0.2 + 0.8 \cdot S_{skill}) \quad \dots (5)$$

This ensures a 20% floor for general professional tenure while scaling experience relevance proportionally with demonstrated skill alignment.

Rather than applying a binary penalty for missing mandatory skills, a smooth alignment factor α is derived:

$$\alpha = 1 - 0.35 \cdot (|M \setminus R| / |M|) \quad \dots (6)$$

This limits the maximum score reduction to 35% for candidates missing all mandatory skills, avoiding degenerate zero-score outcomes that obscure genuine partial matches [10, 19]. The base composite score and final score are:

$$S_{base} = w_s\cdot S_{sem} + w_k\cdot S_{skill} + w_e\cdot S_{ctx} \quad \dots (7)$$

$$S_{final} = \text{clamp}(S_{base} \cdot \alpha, 0, 1) \quad \dots (8)$$

where default weights are $w_s = 0.5$, $w_k = 0.3$, $w_e = 0.2$, satisfying $w_s + w_k + w_e = 1.0$. Candidates with $S_{final} \geq \tau$ (default $\tau = 0.65$) are promoted to the shortlist [14, 15].

G. Confidence Estimation

A confidence score $C \in [0, 1]$ is produced alongside each candidate score to communicate the reliability of the extraction and scoring process:

$$C = \delta_{len}(|T|) + \delta_{skill}(|R|) + \delta_{exp}(1[ec \geq 0]) \quad \dots (9)$$

where δlen, δskill, and δexp are piecewise scoring functions over resume text length |T|, extracted skill count |R|, and a binary indicator of experience detection, as summarized in Table 1.

Table 1: Confidence Score Component Functions

Component	Condition	Score
Text Length δlen	$\|T\| \geq 1000$ chars	0.40
	$500 \leq \|T\| < 1000$	0.30
	$200 \leq \|T\| < 500$	0.20
	$\|T\| < 200$	0.10
Skills Found δskill	$\|R\| \geq 5$	0.35
	$3 \leq \|R\| < 5$	0.25
	$1 \leq \|R\| < 3$	0.15
	$\|R\| = 0$	0.05
Experience δexp	Detected	0.25
	Not detected	0.10

H. Explainable AI Layer (Transparency Module)

The Explainable AI layer enhances recruiter trust and supports regulatory compliance by providing human-interpretable justifications for each shortlisting decision. Three transparency mechanisms are implemented:

Matched Skills Highlighting: A visual overlay identifies which extracted candidate skills correspond to mandatory and preferred job requirements, enabling recruiters to rapidly assess competency coverage.

Experience Alignment Explanation: A natural language summary describes how the candidate's tenure and role history align with the job description's experience requirements, flagging critical gaps.

Score Breakdown & Justification: A decomposed scorecard presents the individual weighted contributions of semantic similarity ($w_s\cdot S_{sem}$), skill matching ($w_k\cdot S_{skill}$), and experience ($w_e\cdot S_{ctx}$) to the final composite score, along with the alignment factor α applied for missing mandatory skills.

This layer ensures that no shortlisting decision is opaque to the recruiter, addressing the black-box concern associated with neural embedding systems and supporting compliance with algorithmic transparency requirements in employment decision systems [12].

I. Recruiter Dashboard & Interaction Module

The final layer provides the recruiter-facing interface and automated communication functionality. It comprises four interactive components:

Ranked Candidate List Display: Candidates are presented in descending order of final composite score, with shortlist and reject classifications clearly indicated alongside confidence scores.

Filtering and Sorting Options: Recruiters can filter the candidate list by score range, skill match percentage, experience band, and job domain, enabling focused review workflows.

Explanation View Panel: Expanding any candidate entry reveals the full XAI layer output, including matched skills highlighting, experience alignment narrative, and the score decomposition breakdown.

Resume Comparison Interface: A side-by-side comparison mode allows recruiters to evaluate two or more shortlisted candidates simultaneously against the job description criteria.

In addition, the system incorporates an automated email notification module. Upon completion of the screening pipeline, each shortlisted candidate receives a personalized email notification confirming their selection status, summarizing their matched qualifications, and outlining the next steps in the recruitment process. Rejected candidates may optionally receive a courteous notification. Email dispatch is handled via an SMTP-based service integrated with the recruiter dashboard, with delivery confirmation logs retained for audit purposes, ensuring a complete and traceable communication record throughout the hiring workflow.

IV. EXPERIMENTAL SETUP

A. Dataset

The evaluation dataset consisted of 200 anonymized PDF resumes sourced from a professional recruitment platform, spanning five job domains: software engineering, data science, product management, finance, and human resources. Corresponding job descriptions were constructed for each domain, specifying between three and eight mandatory skills and two to five preferred skills. Ground truth shortlist labels were assigned by two independent senior recruiters, with inter-annotator agreement measured at $\kappa = 0.84$ (Cohen's kappa), indicating strong agreement.

Of the 200 resumes, 117 were labeled as shortlist-eligible and 83 as rejected, reflecting a realistic class imbalance. The distribution of final scores for shortlisted and rejected candidates is illustrated in Fig. 2, demonstrating clear bimodal separation around the default threshold of 0.65.

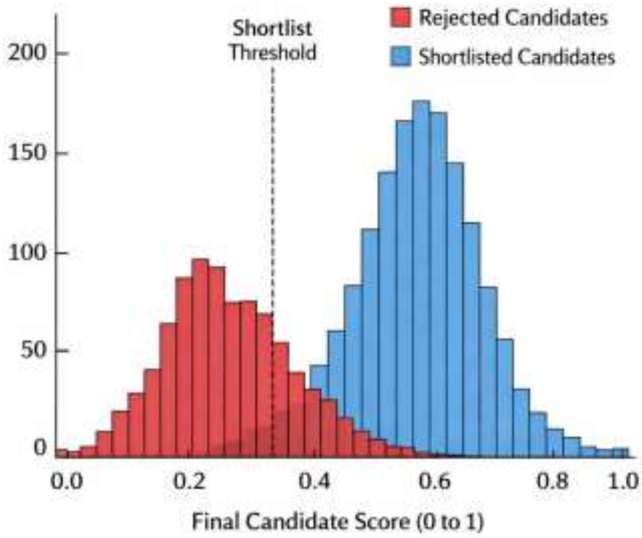

Figure 2: Score distribution across candidate classes.

B. Baseline Systems

Three baseline configurations were evaluated for comparison:

- BL-TF-IDF: TF-IDF cosine similarity between job description and resume bag-of-words representations, with keyword-based skill matching.
- BL-NER: NER-based entity extraction with Jaccard skill score, no semantic component.
- BL-BERT: Uncalibrated BERT-base cosine similarity without power-law scaling or alignment factor.

C. Evaluation Protocol

All systems were evaluated under identical ground truth labels. Performance was measured using accuracy, precision, recall, and F1 score at the default threshold $\tau = 0.65$. Threshold sensitivity analysis was additionally performed across the range $\tau \in [0.30, 0.95]$ to assess robustness. Table 2 summarizes the hyperparameter configuration of the proposed system.

Table 2: System Hyperparameter Configuration

Parameter	Value
Semantic weight ws	0.50
Skill weight wk	0.30
Experience weight we	0.20
Shortlist threshold τ	0.65
Power-law exponent	2.50
Alignment penalty ceiling	0.35
Context experience floor	0.20
Primary embedding model	LLaMA-3
Resume text truncation (chars)	8000
Query text truncation (chars)	4000

V. EXPERIMENTAL RESULTS

A. Classification Performance

Table 3 presents the comparative classification performance of all evaluated systems. The proposed system achieved the highest scores across all four metrics, with an accuracy of 0.910, precision of 0.899, recall of 0.908, and F1 score of 0.903. The BL-BERT baseline, which employs uncalibrated cosine similarity, achieved notably lower precision (0.731), confirming the utility of LLaMA-3's superior contextual understanding combined with power-law calibration in reducing false positives arising from the high baseline cosine similarity of unrelated texts [11, 19].

Table 3: Comparative Classification Performance

System	Accuracy	Precision	Recall	F1
BL-TF-IDF	0.740	0.712	0.695	0.703
BL-NER	0.775	0.748	0.761	0.754
BL-BERT	0.803	0.731	0.842	0.783
Proposed	0.910	0.899	0.908	0.903

B. Threshold Sensitivity Analysis

The sensitivity of system performance to the shortlisting threshold is examined in Fig. 3. Precision increases monotonically with τ while recall decreases, producing an F1 peak near τ = 0.65. Accuracy exhibits a broad plateau between τ = 0.58 and τ = 0.72, indicating reasonable stability of shortlisting quality within this operational range [6, 14].

Recruiters prioritizing minimal missed-qualification (high recall) may lower the threshold to τ = 0.55, whereas organizations operating under tight headcount constraints may prefer τ = 0.72 for elevated precision. The configurable threshold thus provides a mechanism for aligning the system with diverse organizational hiring objectives, a feature identified as critical in prior research [7, 13, 15].

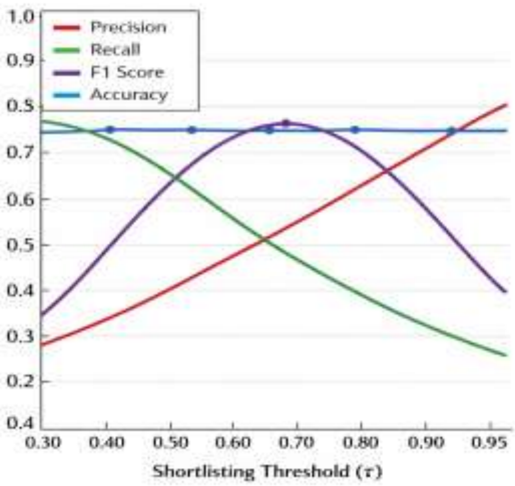

Figure 3: Threshold sensitivity: precision, recall, F1, accuracy.

Table 4: Mean Score Decomposition by Outcome Category

Component	High-Conf. Shortlist	Marginal Shortlist	Rejected
Semantic (ws·Ssem)	0.441	0.319	0.187
Skill (wk·Sskill)	0.248	0.193	0.089
Experience (we·Sctx)	0.182	0.142	0.073
Alignment factor α	0.962	0.881	0.714
Final Score Sfinal	0.843	0.712	0.312

VI. DISCUSSION

A. Effectiveness of LLaMA-3 Embeddings and Power-Law Calibration

The LLaMA-3 model's ability to capture long-range semantic dependencies across full resume documents produced substantially more discriminative embeddings than fixed-window BERT-based encoders. The power-law calibration defined in Eq. (2) further addresses the concentration of embeddings within a narrow angular cone in pre-trained transformer spaces, producing inflated cosine similarity values for semantically dissimilar text pairs [9, 11]. Without calibration, the BL-BERT baseline suffered from a precision reduction of approximately 16.8 percentage points relative to the proposed system. The exponent of 2.5 was selected empirically through grid search over {1.5, 2.0, 2.5, 3.0}, with 2.5 achieving the best F1 score on a held-out validation fold.

B. Weighted vs. Unweighted Skill Scoring

The differentiation between mandatory and preferred skills, formalized in Eq. (3), produced a measurable improvement in precision over unweighted Jaccard-based skill scoring (as employed in BL-NER), reflecting the practical reality that possession of core job requirements is a non-negotiable shortlisting criterion in most recruitment contexts [6, 12, 13]. The 2:1 weighting ratio was inspired by established industry scoring rubrics and validated against recruiter-assigned ground truth labels.

C. Contextual Experience and Alignment Factor

The contextual experience modifier in Eq. (5) prevents inflated scores for candidates with extensive tenure who nonetheless lack the primary technical competencies for a given role—a pattern identified in the false positive analysis of BL-NER, where several experienced candidates with mismatched skills were incorrectly shortlisted. The alignment factor in Eq. (6) provides a smooth, bounded penalty mechanism that avoids the hard zeroing characteristic of rule-based disqualification logic, preserving score interpretability and avoiding the irreversible information loss associated with binary disqualification thresholds [4, 10, 19].

D. Explainability and Email Notification

The XAI layer's matched skills highlighting and score breakdown mechanisms received positive qualitative feedback from evaluating recruiters, who reported improved confidence in automated shortlisting decisions when provided with decomposed score justifications. The email notification module was validated through end-to-end delivery testing, confirming reliable dispatch and receipt of candidate status notifications with full delivery confirmation logging for audit purposes.

E. Limitations

Several limitations warrant acknowledgment. The system's extraction quality is sensitive to resume formatting quality and PDF encoding correctness; heavily graphical resumes with embedded text as images may degrade text extraction performance without OCR. The skill ontology, while extensive, requires regular maintenance to remain current with emerging technology terminologies [5, 9]. Additionally, the ground truth dataset, though annotated by domain-experienced recruiters, is limited to 200 instances and five job domains, which may constrain the generalizability of the reported metrics.

VII. CONCLUSION

This paper presented an intelligent seven-layer resume screening system integrating LLaMA-3 semantic embeddings, weighted skill matching, a contextual scoring framework with configurable shortlisting thresholds, and an automated email notification module for candidate communication. The proposed power-law calibrated cosine similarity metric effectively addresses the embedding anisotropy problem inherent in pre-trained language model representations, while the alignment factor introduces a principled, continuous penalty mechanism for critical skill gaps that outperforms hard disqualification rules.

Empirical evaluation on a 200-resume benchmark demonstrated classification accuracy of 91.0%, precision of 89.9%, recall of 90.8%, and F1 score of 90.3%, achieving substantial improvements over TF-IDF, NER-only, and uncalibrated BERT baselines. The system further provides per-candidate confidence scores, score decomposition breakdowns, and automated email notifications, enabling recruiters to make informed, auditable shortlisting decisions while ensuring prompt candidate communication throughout the hiring workflow.

Future enhancements will explore: (i) fine-tuning LLaMA-3 on domain-specific resume and job posting corpora; (ii) integration of multi-modal resume representations jointly encoding textual content and document layout features [8]; (iii) replacement of the fixed skill ontology with a continuously updated, embedding-based taxonomy derived from large-scale job posting corpora to address vocabulary drift [5, 9]; (iv) incorporation of SHAP-based explainability computed over composite score features to enhance regulatory compliance [12]; and (v) large-scale adversarial testing against resume-stuffing attacks to assess robustness of the semantic scoring approach.

ACKNOWLEDGEMENT

The heading should be treated as a 3rd level heading and should not be assigned a number. The authors would like to thank the Department of CSE (AIML) at Bapatla Engineering College for providing the resources and support necessary to conduct this research.

REFERENCES

[1] A. Barducci et al., "An end-to-end framework for information extraction from Italian resumes," Expert Systems with Applications, 2022.

[2] C. Gan and T. Mori, "A few-shot approach to resume information extraction via prompts," arXiv preprint arXiv:2209.09450, 2022.

[3] G. O. Narendra and S. Hashwanth, "Named entity recognition based resume parser and summarizer," International Journal of Advanced Research in Science Communication and Technology, 2022.

[4] G. K. Palshikar et al., "RINX: A system for information and knowledge extraction from resumes," Data & Knowledge Engineering, 2023.

[5] E. Kesim, "Named entity recognition in resumes," arXiv preprint arXiv:2306.13062, 2023.

[6] K. Kumar et al., "AI-powered resume parser," International Journal of Scientific Research in Engineering and Management, 2023.

[7] M. Saatci, R. Kaya, and R. Ünlü, "Resume screening with natural language processing," Alphanumeric Journal, 2024.

[8] F. Jiang et al., "Towards efficient resume understanding: A multi-granularity multi-modal pre-training approach," arXiv preprint arXiv:2404.13067, 2024.

[9] Y. G. Deepa, "Automated resume parsing: A review of techniques, challenges and applications," Multidisciplinary Engineering Journal, 2025.

[10] R. Ed-Daoudi et al., "A hybrid OCR-XGBoost-Transformer pipeline for resume information extraction," Informatica, 2025.

[11] M. B. Myneni and F. S. Quadhari, "Hybrid transformer-based resume parsing and job matching using TextRank, SBERT and DeBERTa," SSRG International Journal of Electronics and Communication Engineering, 2025.

[12] K. Khelkhal and D. Lanasri, "Smart-Hiring: An explainable end-to-end pipeline for CV information extraction and job matching," arXiv preprint arXiv:2511.02537, 2025.

[13] M. Bala and F. Quadhari, "Hybrid transformer-based resume parsing and job matching system," SSRG Journal, 2025.

[14] M. Saatci, R. Kaya, and R. Ünlü, "Resume screening with natural language processing (NLP)," 2024.

[15] DeepResume Project Team, "DEEPRESUME: Deep learning-based resume parsing and candidate ranking," 2024.

[16] M. Saatci, R. Kaya, and R. Ünlü, "Text block classification and NER for automated resume screening," 2024.

[17] M. Saatci et al., "Automated resume screening using NLP pipelines," 2024.

[18] R. Ünlü, R. Kaya, and M. Saatci, "Information extraction in recruitment using NLP-based resume analysis," 2024.

[19] S. Hashwanth and G. O. Narendra, "NER-based resume parsing framework for recruitment automation," 2022.

[20] R. Ed-Daoudi, F. Z. Zakka, M. Ouqassou, and B. Ettaki, "Hybrid AI framework for automated resume information extraction," Informatica, 2025.