An Explainable AI Approach to Resume Ranking Using Semantic Similarity and Skill-Based Matching

Harshal Chinchole
Department of Information Technology
Shri Sant Gajanan Maharaj College of Engineering Shegaon, Maharashtra, India
workspaceofharsh@gmail.com

Ranjit Dhakare
Department of Information Technology
Shri Sant Gajanan Maharaj College of Engineering Shegaon, Maharashtra, India
ranjitdhakare0505@gmail.com

Tulsi Totala
Department of Computer Science and Engineering
Shri Sant Gajanan Maharaj College of Engineering Shegaon, Maharashtra, India
tulsitotala@gmail.com

Abstract—As the number of online job applications continues to grow, artificial intelligence (AI) has become essential for screening resumes. However, most existing AI models operate as black boxes: they produce a ranking without any justification, which limits the confidence recruiters place in them. This paper proposes an explainable and transparent AI (XAI) framework for resume ranking. The system combines contextual language understanding, obtained through Sentence-BERT embeddings, with a domain-specific skill matching engine. Experiments show that a purely semantic model can misrank candidates by over-weighting broad textual similarity, a weakness that the proposed hybrid approach corrects. By combining a wide semantic view with strict skill-based scoring, the system ranks candidates correctly and provides a clear explanation for every decision. The framework therefore achieves a balance between accuracy and transparency, which is what real-world human resource systems require in order to establish trust and fairness.

Keywords—explainable AI (XAI), resume parsing, AI hiring, Sentence-BERT, skill-based matching, transparent recruitment

I. INTRODUCTION

The digitization of job applications has significantly changed the recruitment process. Organizations receive a very large number of applications for every job advertisement, and screening them manually is time-consuming and inefficient. As a result, automated resume analysis systems have become popular for screening and ranking applicants against job requirements.

Conventional resume screening tools rely on keyword matching and rule-based systems, which are weak at capturing contextual meaning. With advances in machine learning and Natural Language Processing (NLP), more sophisticated methods such as TF-IDF, word embeddings, and transformer-based models have been proposed to improve matching accuracy. In particular, Bidirectional Encoder Representations from Transformers (BERT) and Sentence-BERT (SBERT) have proved effective for semantic similarity tasks.

Despite these advances, the major drawback of contemporary AI-based recruiting systems is their lack of interpretability. Deep learning models are often black boxes: they produce a final score but no explanation of how it was calculated. This is unacceptable in a real-life setting, where recruiters need justifiable and explainable decisions to guarantee fairness, accountability, and trust.

Explainable Artificial Intelligence (XAI) has emerged as a potential solution to these problems. Although methods such as LIME and SHAP can provide post-hoc explanations, they can be computationally expensive and do not necessarily reflect the actual reasoning of complex models. Furthermore, most existing resume ranking systems treat explainability as an add-on to the decision-making process rather than an integral part of the system.

To resolve these drawbacks, this paper presents an explainable AI framework for resume ranking that combines semantic matching with interpretable scoring. The proposed method integrates SBERT-based embeddings with a domain-specific skill matching mechanism to provide both accuracy and transparency. Semantic similarity is converted into domain-specific relevance through a weighted scoring model so that candidates can be ranked more reliably.

The main contributions of this work are as follows:

- A hybrid resume ranking model that combines semantic similarity with skill-based scoring.
- A built-in explainability component that offers feature-level interpretation of results.
- An analysis that demonstrates the weakness of purely semantic models in recruitment settings.
- A lightweight and easy-to-use system that can be deployed in real-world settings.

II. RELATED WORK

Automated resume processing has received considerable attention because organizations receive a very large number of applications. This section reviews four key areas: resume parsing, resume–job matching, deep learning-based matching, and explainability in NLP.

A. Resume Parsing

Early resume parsers were rule-based and template-driven; they were fast but failed on irregular layouts. Statistical approaches based on Support Vector Machines (SVM) and Conditional Random Fields (CRF) followed, and more recently deep learning and NLP techniques have been used to extract skills, experience, and education from unstructured resumes with high accuracy [1]. The limitation of these systems is that they perform extraction only; they neither rank candidates nor explain their output.

B. Resume–Job Matching

TF-IDF keyword matching is simple but lacks contextual sensitivity. NLP-based systems [2] perform better through the addition of engineered features, but they also require large datasets. The gap in these approaches is that they capture lexical overlap rather than actual knowledge of the domain.

C. Deep Learning-Based Matching

Dense vector representations combined with cosine similarity, particularly BERT and SBERT models [4], are more successful than earlier methods. Their disadvantages are black-box decision-making and extensive data labelling requirements.

D. Explainable AI in NLP

LIME and SHAP explain model decisions by highlighting the most influential words [5], [6], [7]. Their limitations are that the approximations are unstable, they incur additional computation, and they do not truly demystify deep models [3].

E. Research Gap

No existing system combines accurate matching with reliable explanation. Current tools can either match, extract, or explain, but none provides all three at the feature level (skills, experience, and education) simultaneously.

F. Our Approach

This work builds a single system that incorporates explainability directly into the resume ranking process rather than adding it as an afterthought. Recruiters can immediately understand why Candidate A is ranked above Candidate B on every metric.

III. METHODOLOGY

A. System Overview

The proposed system accepts resumes and a job description, matches them using semantic similarity, and provides justification for the resulting ranking. As shown in Fig. 1, it is divided into four key stages: (1) data parsing, (2) SBERT embedding generation, (3) multi-factor scoring, and (4) explanation generation.

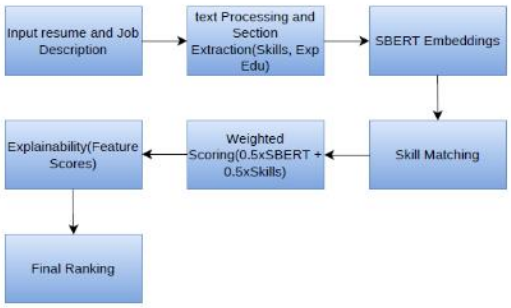

Fig. 1. Proposed explainable resume ranking system architecture.

B. Data Preprocessing

Resumes in PDF or DOCX format are first parsed and divided into the conventional sections: skills, experience, and education. Each section is then tokenized, stop-words are removed, and the text is normalized independently.

C. Semantic Embedding Using Sentence-BERT

Sentence-BERT (SBERT) is used as the embedding model because it is a transformer designed to capture context in sentence-level similarity tasks.

Let R represent the resume text and J represent the job description. SBERT encodes these into dense vector representations:

$$E_R = \text{SBERT}(R), \quad E_J = \text{SBERT}(J) \tag{1}$$

The semantic similarity between the resume and the job description is then computed as the cosine similarity of the two embeddings:

$$Sim(R, J) = (E_R \cdot E_J) / (\|E_R\| \, \|E_J\|) \tag{2}$$

D. Multi-Factor Scoring Mechanism

To make the result interpretable, the overall similarity is decomposed into section-level components corresponding to the resume sections. Let:

- S_{skill} = similarity between the skills section of the resume and the job description,
- S_{exp} = similarity between the experience sections,
- S_{edu} = similarity between the education sections.

Each component is computed from SBERT embeddings, and the components are combined using a weighted sum:

$$Score = w_1 S_{skill} + w_2 S_{exp} + w_3 S_{edu} \quad (3)$$

where $w_1 + w_2 + w_3 = 1$. In this work the weights are set to $w_1 = 0.5$ (skills), $w_2 = 0.3$ (experience), and $w_3 = 0.2$ (education). These values reflect the relative importance that recruiters typically assign to these sections when hiring.

E. Explainability

The explainability component presents the following information to the recruiter for every candidate:

1) Section-level scores (for example, Skills = 85%, Experience = 60%).
2) The key matching terms, highlighted in bold or colour.
3) A warning flag for every required skill that is missing from the resume.

F. Workflow

The complete workflow of the proposed system, from input to ranked output, is summarized in Fig. 2.

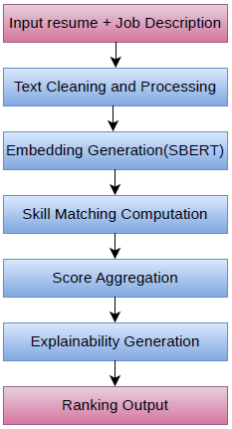

Fig. 2. Workflow of the proposed resume ranking framework.

G. Advantages

The proposed methodology offers several advantages over existing systems:

- Provides semantic matching using transformer-based embeddings.
- Ensures interpretability through multi-factor scoring.
- Eliminates reliance on purely black-box models.
- Enables actionable insights for both recruiters and candidates.

IV. EVALUATION

A. Dataset and Experimental Setup

To test the effectiveness of the proposed explainable resume ranking framework, a domain-specific dataset was created consisting of job descriptions and candidate resumes related to cybersecurity. Since no standardized public resume–job matching dataset was available, a controlled experimental dataset was hand-curated. The dataset includes:

- One cybersecurity job description that defines the required skills, such as vulnerability analysis, penetration testing, SIEM tools, Python automation, and network security.
- Four candidate resumes of different relevance levels (high, medium, and low).

To obtain a ground truth for evaluation, each resume was graded manually with a relative relevance level based on domain knowledge. All textual data were preprocessed using standard NLP normalization and tokenization. Semantic embeddings were generated using the Sentence-BERT model all-MiniLM-L6-v2, and similarity was calculated using cosine similarity.

B. Baseline Methods

Two baseline approaches were used to evaluate the effectiveness of the proposed method.

1) *TF-IDF Based Similarity:* Resumes and job descriptions were converted into TF-IDF vectors, and matching scores were calculated using cosine similarity. Although this approach is interpretable, it does not provide contextual understanding.
2) *Semantic Matching Using SBERT:* Dense vector embeddings of the resume and job description were generated using Sentence-BERT, and semantic similarity scores were computed using cosine similarity. This approach captures contextual relationships, but it is a black-box model and is not interpretable.

C. Proposed Method

The proposed system combines semantic similarity obtained from SBERT with a domain-based skill matching score. The final score is calculated as:

$$Score = 0.5 \times S_{semantic} + 0.5 \times S_{skill} \quad (4)$$

where $S_{semantic}$ represents the SBERT-based cosine similarity and S_{skill} represents the keyword-based skill matching score. This hybrid formulation ensures that semantic meaning and domain coverage are both captured, in addition to interpretability.

D. Experimental Results

The effectiveness of the baseline approaches and the proposed technique was measured by analysing the ranking of candidates against the job description. Table I summarizes the comparison.

Table I. Comparative Analysis Of Methods

Method	Observation
TF-IDF	Poor performance due to lack of semantic understanding
SBERT	Better semantic similarity but produced incorrect ranking in some cases
Proposed Method	Achieved the correct ranking with improved interpretability

The SBERT model achieved better matching than TF-IDF; however, it ranked some candidates too high because it relied on general semantic similarity rather than domain-specific expertise. The proposed approach resolved this problem through skill-based scoring, which made the rankings more accurate and trustworthy.

E. Ranking Results

The final ranking obtained using the proposed system is presented in Table II.

Table II. Final Ranking Results

Rank	Candidate	Final Score
1	Rahul Sharma	0.625
2	Priya Verma	0.402
3	Sneha Kulkarni	0.389
4	Aman Gupta	0.343

The results indicate that the system ranked the most deserving candidate first and less relevant candidates lower. The small differences among the mid-ranked candidates are explained by overlapping skills and semantic similarities, which is consistent with real recruitment situations.

F. Explainability Analysis

One of the most significant contributions of the proposed framework is its ability to provide explanations for the ranking decisions. The system output contains feature-level contributions, as illustrated in Fig. 3, which make the decision process transparent. Table III shows the explainability output for each candidate.

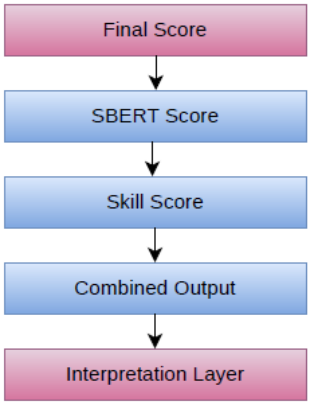

Fig. 3. Explainability module showing feature-level contribution.

TABLE III. Explainability Output

Candidate	SBERT Score	Skill Score	Final Score	Interpretation
Rahul Sharma	0.645	0.605	0.625	Strong alignment across skills and experience
Priya Verma	0.540	0.263	0.402	Moderate match; insufficient experience
Sneha Kulkarni	0.488	0.289	0.389	Good domain expertise but limited coverage
Aman Gupta	0.581	0.105	0.343	High semantic similarity but lacks required skills

The explainability module reports the contribution of each component to the final score, which allows recruiters to understand how candidates were ranked.

G. Discussion

The experimental findings lead to a number of important observations:

SBERT-based semantic models are much more effective at capturing contextual similarity than traditional TF-IDF approaches. However, semantic methods can produce incorrect rankings because they include no domain weighting.

Skill-based scoring improves the accuracy of the ranking because it focuses on the expertise that is actually required.

The proposed framework is more transparent because it produces interpretable outputs, which black-box models do not provide.

These results show that combining semantic understanding with domain-specific features produces more credible and understandable recruitment systems.

V. CONCLUSION

This paper has presented an explainable AI-based resume analysis and ranking framework that combines semantic similarity with domain-aware skill

matching. The proposed approach overcomes the drawbacks of traditional and deep learning-based systems by improving both ranking accuracy and interpretability. Experimental findings show that although semantic models such as SBERT have strong contextual understanding, they are vulnerable to producing incorrect rankings without domain-specific weighting. The proposed system is more reliable and transparent because it includes a skill-based scoring component and a weighted aggregation mechanism.

The framework also provides feature-level explanations, which support better decision-making and greater confidence in automated recruitment systems. Overall, the proposed solution is a feasible and scalable approach for contemporary job recruitment, with scope for further improvement in areas such as bias mitigation and real-time deployment.

REFERENCES

[1] Y. G. Deepa, A. Sindhu, A. Shruthi, and B. Neha, "Automated resume parsing: A review of techniques, challenges and future directions," *International Journal of Multidisciplinary Research and Growth Evaluation*, vol. 6, no. 2, pp. 1065–1069, 2025.

[2] S. Bhor, H. Shinde, V. Gupta, V. Nair, and M. Kulkarni, "Resume parser using natural language processing techniques," Dept. of Computer Engineering, Pillai College of Engineering, Navi Mumbai, India, 2022.

[3] B. Gupta, "Explainable AI for automated student feedback: A comparative study of LIME and SHAP on text classification," Heritage Xperiential Learning School, India, 2025.

[4] D. Lavi, V. Medentsiy, and D. Graus, "conSultantBERT: Fine-tuned Siamese Sentence-BERT for matching jobs and job seekers," in *Proc. RecSys in HR Workshop*, 2021.

[5] M. Danilevsky et al., "A survey of the state of explainable AI for natural language processing," in *Proc. AACL-IJCNLP*, pp. 447–459, 2020.

[6] V. Bidve et al., "Use of explainable AI to interpret the results of NLP models for sentimental analysis," *Indonesian Journal of Electrical Engineering and Computer Science*, vol. 35, no. 1, pp. 511–519, 2024.

[7] M. A. Shukla, "Interpreting BERT using LIME and SHAP," 2023.