International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 12 December 2025
Page No: 4232-4239
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0620

An AI-Driven Placement Ecosystem for Automated Skill Matching

Linett Sophia D[1], Mohan C[2], Nagalingam L [3], Praveen C[4], Rajakumar R[5]

[1]Department of Artificial Intelligence and Data Science / Assistant Professor Erode Sengunthar Engineering College, Erode, India.

[2,3,4,5]Department of Artificial Intelligence and Data science, Erode Sengunthar Engineering College, Erode, India.

Emails: *linisofia89@gmail.com[1], mohanmss7707@gmail.com[2], nagalingamlakshmanan2004@gmail.com[3], ccpraveen789@gmail.com[4], raja.kumar02104@gmail.com[5]*

Abstract

One of the biggest challenges in today's recruit- ment systems is making sure that candidate resumes match job descriptions. This matching directly affects the quality of hiring decisions and the productivity of the organization. To accomplish this, we propose a supervised fine-tuning method for semantic resume–job matching that leverages Sentence-BERT (SBERT) embeddings to match candidates to job descriptions with high accuracy. Our approach represents both resumes and job descriptions in a shared embedding space. This allows the method to use high-quality computation of similarity for the retrieval of top-k job match rankings. The model was fine- tuned and trained on a labeled dataset of resume–job pairs, and evaluated using Spearman and Pearson correlation coefficients to assess agreement with ground truth relevance, with additional metrics of top-k retrieval, namely Precision, Recall and Normalized Discounted Cumulative Gain (NDCG). Experimental results show that the fine-tuned method outperformed the pretrained baseline, achieving high correlations, precision, and accuracy in identifying relevant candidates. This work demonstrates the use of embedding, along with supervised fine- tuning, can improve accuracy and applicability of resume-job matching approaches. The experimental analysis shows that the fine-tuned model consistently gets higher performance scores than the pretrained baseline.

Keywords: *Sentence-BERT, Supervised Fine-Tuning, Se- mantic Embeddings, Spearman Correlation, Pearson Correlation, Similarity-Based Ranking.*

1. Introduction

Recruitment today has moved from manually checking resumes to using automated systems that can handle large numbers of applications. Despite this progress, the main issue remains how to efficiently and accurately match resumes to job descriptions. Traditional methods that rely on keyword searches or rule-based filtering often miss important meanings and context. This shortcoming leads to poor matches between candidates and jobs, higher hiring costs, and longer recruitment processes. Recent advancements in natural language processing (NLP) and deep learning allow the use of embedding-based models, which map text into a common semantic space. These mod- els enable more accurate similarity measurements between resumes and job postings. Techniques like Sentence-BERT (SBERT) perform well in semantic similarity tasks by learning sentence-level embeddings that maintain contextual meaning [1]. Fine-tuned models applied to resume-job datasets improve the ability to find suitable candidates by recognizing specific nuances in the field, outperforming traditional methods like TF-IDF or basic BERT embeddings [6]. This paper introduces an AI-driven placement system that uses supervised fine-tuning of SBERT embeddings to boost resume-job matching accuracy. By training on labeled resume- job pairs and evaluating with correlation and retrieval-based metrics, the system offers clear similarity scores and ranked recommendations. This approach not only shows better performance than basic pre-trained models but also provides real-world benefits for recruitment platforms by cutting down manual screening efforts and ensuring better matches between candidates and jobs.

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 12 December 2025
Page No: 4232-4239
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0620

2. Literature Review

The automation of recruitment processes has received a lot of attention in recent years because of the need for efficient and accurate candidate selection. Traditional methods, which often depend on keyword-based searches and manual filtering, struggle to capture the deeper meanings between resumes and job descriptions. To address these issues, embedding-based methods like Sentence-BERT (SBERT) have been developed. SBERT uses a Siamese architecture to create meaningful embeddings, allowing for high-quality similarity calculations in natural language tasks[1]. Several specific adaptations of SBERT have appeared in recruitment applications. For example, conSultantBERT was created to fine-tune SBERT for matching job postings with candidate resumes. This model showed that training for spe- cific tasks leads to better connections between resumes and job requirements compared to general embeddings[2]. Similarly, CareerBERT built on this research by creating a shared embedding space for resumes and ESCO job categories, which improves understanding and relevance within the field[3]. Embedding-based methods have also been used in applicant tracking systems. Ravi Varma Kumar Bevara and Nishith Reddy Mannuru introduced Resume2Vec, which turned traditional tracking systems into smart platforms that capture semantic similarities between candidates and roles. Their work highlighted that embeddings significantly improve candidate ranking by representing both resumes and jobs in a unified space[4]. In a study, Asmita Deshmukh compared SBERT with BERT for resume screening tasks and found that SBERT achieved better results due to its optimized design for semantic similarity[5]. Other researchers have looked into how Natural Language Processing (NLP) can be integrated into recruitment systems. Mehtap Saatci, Ramazan Ünlü and Rukiye Kaya showed that NLP techniques could automate the early stages of resume screening by extracting structured data from unstructured resumes, thus reducing manual work[6]. Mohammed Shinan K S and Goutham Praveen P P further developed this idea by applying machine learning to categorize and screen resumes, demonstrating how AI-based classification models can make recruitment processes smoother[7]. In addition, Charan S N, Devananda S N and Suhas G K suggested hybrid job recommendation systems that combine content-based and col- laborative filtering, showcasing how embedding models can provide personalized career guidance [8]. Beyond recruitment, studies have examined broader job- matching algorithms and their real-world effects. Prabakaran Paranthaman. pointed out the efficiency gained from automat- ing resume screening with machine learning, especially in large-scale recruitment situations where manual screening is impractical [9]. Bryar A. Hassan. looked at different machine learning job matching methods, noting that embedding-driven systems outperform traditional heuristic methods in terms of scalability and flexibility [10]. These studies provide a solid groundwork for further research, driving the supervised fine- tuning of SBERT for matching resumes and job descriptions recommendations. From reviewing past studies, it's clear that most existing systems concentrate on general semantic similarity, not fully optimizing embeddings for recruitment-specific settings. Al- though several works have used pretrained models like BERT and SBERT for text similarity, they often miss the subtle connections between resumes and job descriptions. In contrast, this research presents a fine-tuned matching model based on SBERT, specifically designed for job-resume and resume-job retrieval tasks. This model aims to improve semantic under- standing and relevance to the domain. To ensure a strong evaluation, this work uses a dual-directional analysis to assess both job-to-resume and resume-to-job matching. This method more accurately reflects real-world recruitment situations than one-way matching systems. Additionally, a thorough performance analysis has been performed, comparing the fine-tuned model against pretrained and traditional baseline methods using several metrics, including Spearman, Pearson, Precision, Recall, and NDCG. The results clearly show improved matching accuracy, indicating that the fine-tuned embeddings exceed the performance of pretrained models in all key metrics. This improvement highlights the model's ability to offer more reliable and context-aware candidate-job recommendations, supporting the development of smarter recruitment systems.

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 12 December 2025
Page No: 4232-4239
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0620

3. Methodology

3.1. Methodology Overview

The proposed system follows a structured pipeline that enables accurate and efficient resume and job matching. At a high level, the method combines natural language processing (NLP), pre-trained transformer models, and supervised fine-tuning to capture the relationships between candidate resumes and job descriptions. The process starts with data preparation, where raw text from resumes and job postings is cleaned and standardized. Both document types are transformed into dense vector representations using the Sentence-BERT (SBERT) all-MiniLM model, selected for its balance between computational efficiency and semantic accuracy. These embeddings are then projected into a shared vector space, allowing for direct similarity comparisons.

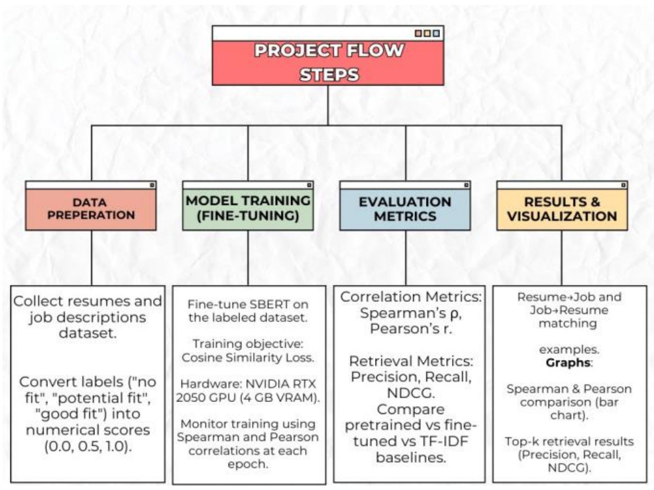

Figure 1 Project Flow Diagram

To improve performance, the model undergoes fine-tuning on a labeled dataset of resume and job pairs, ensuring the learned representations match recruitment-specific meanings. The training process uses GPU resources (NVIDIA RTX 2050), with optimization based on similarity loss functions. Once trained, the model produces interpretable similarity scores and ranked recommendations, which can suggest the best candidates for a job or retrieve the most relevant job opportunities for a specific resume. The entire workflow can be shown as a flow diagram, illustrating the steps from raw data to final ranked recommendations. This diagram highlights the key stages, including data preprocessing, embedding generation, supervised fine-tuning, similarity computation, and evaluation. The flow of the system is shown in Fig. 1 (Project Flow Diagram). It highlights the main steps from data preprocessing to deployment.

3.2. Dataset Preparation

The dataset used in this study came from publicly available, unlicensed sources. This ensured there were no restrictions on its use for research and experimentation. Unlike traditional natural language processing pipelines that often need a lot of text cleaning and normalization, our approach minimized preprocessing. We kept both resumes and job descriptions in their original form to maintain their meaning and context. This decision was based on the design of transformer- based models like all-MiniLM-L6-v2. These models are built to handle raw text inputs well. They use subworld tokenization and contextual embeddings, which reduces the need for manual steps such as removing stop words, lemmatizing, or stemming. As a result, the pipeline stays simple, making it easier to integrate into real-world recruitment platforms without extra burden. By skipping unnecessary preprocessing, we allow the model to work with real resume and job description text. This improves its usefulness in practical deployment situations.

3.3. Model Architecture

The foundation of the proposed system is Sentence-BERT (SBERT), specifically the lightweight variant all-MiniLM- L6-v2, which works well for semantic similarity tasks with limited computational resources. Unlike traditional bag-of- words or TF-IDF methods, SBERT generates dense vector embeddings that capture the semantic meaning of sentences rather than just focusing on surface-level word overlap. This makes it especially effective for matching resumes and job descriptions, where similar concepts may use different words. The SBERT architecture is built on the Transformer frame- work. It features several layers of self-attention mechanisms that help the model learn the context of each token in a sequence. In the MiniLM variant, the architecture is simpli- fied to 6 Transformer layers with 384 hidden dimensions. This allows for faster processing and lower memory usage without sacrificing much accuracy. After generating token embeddings, they are combined using a mean-pooling strat- egy, which creates a fixed-length sentence embedding suitable for similarity calculations. In our system, both resumes and job descriptions are separately processed through the same SBERT encoder. This shared-weights design ensures that both inputs are mapped into a common embedding space, where semantically related resume-job pairs are placed closer together. During training, the model is fine-tuned using supervised similarity objectives so that relevant pairs score higher in similarity than irrelevant ones.

Figure 2 shows the overall architecture of SBERT and its workflow within our system. The diagram illustrates how tokenized inputs are sent into the Transformer encoder, com- bined into sentence embeddings, and then compared in the embedding space to evaluate semantic similarity.

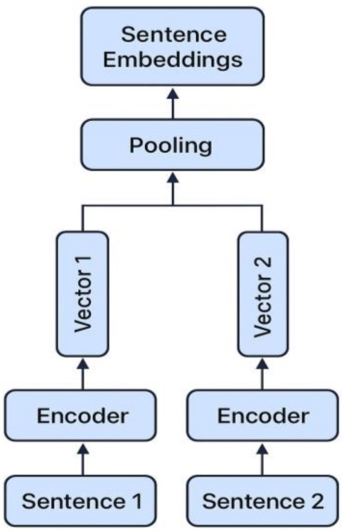

Figure 2 S-BERT Architecture

3.4. Training

The training phase of the proposed system used the all- MiniLM-L6-v2 variant of Sentence-BERT. This variant was fine-tuned to better capture the connection between resumes and job descriptions. To ensure a thorough evaluation, two different datasets were used: one for training and another for testing. This separation prevents data leakage and provides an unbiased assessment of the model's ability to generalize. The training dataset included labeled resume-job pairs. Each pair was marked to show how relevant the candidate profile was to the job description. These labels guided the fine- tuning process. The testing dataset, however, was completely new during training, which ensured that the performance metrics genuinely reflected real-world performance. The fine-tuning took place on an NVIDIA GeForce RTX 2050 GPU with 4 GB VRAM. This GPU was selected for its good balance of computing power and memory efficiency. The model was trained using a batch size of 16 and a learning rate of 2e-5, over a period of 15 epochs. The training objective was Cosine Similarity Loss, which helps make the embeddings of relevant resume-job pairs closer in the vector space compared to non-relevant pairs. The training progress was recorded, including the loss reduction across epochs. Figure 3 shows a snapshot of the training phase, illustrating how the model gradually improved as it learned to map resumes and job descriptions into a meaningful embedding space.

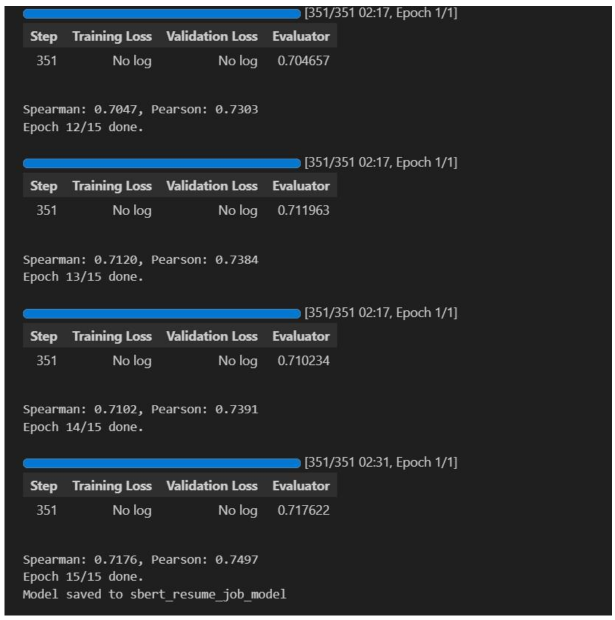

Figure 3 Fine-Tuning Image

3.5. Deployment

The trained model can be used on a server or in a local environment to make real-time predictions. The model processes input data, like student profiles and job descriptions, to produce similarity scores for matching skills to jobs. The deployment is set up to

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 12 December 2025
Page No: 4232-4239
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0620

be modular. This allows for future integration of the model into applications or platforms for automated recommendations without changing the main model. The system can use GPU acceleration to speed up processing, ensuring it can efficiently handle many inputs at once.

4. Experiments and Results

4.1. Dataset Description

The experiments used an unlicensed dataset with two parts: a training set having over 6,200 entries and a testing set with more than 1,700 entries. Each entry includes three columns: resume_text, job_description_text, and a label that shows how relevant a candidate profile is to a job description. The model used the textual data directly without any preprocessing be- cause it relies on sentence embeddings to understand meaning. Separating the training and testing datasets makes sure that evaluation happens on new data, leading to an unbiased assessment of how well the model works. This dataset was only for research and training in this study, and there was no redistribution or commercial use.

4.2. Experimental Setup

The experiments used an unlicensed dataset with two parts: a training set having over 6,200 entries and a testing set with more than 1,700 entries. Each entry includes three columns: resume_text, job_description_text, and a label that shows how relevant a candidate profile is to a job description. The model used the textual data directly without any preprocessing be- cause it relies on sentence embeddings to understand meaning. Separating the training and testing datasets makes sure that evaluation happens on new data, leading to an unbiased assessment of how well the model works. This dataset was only for research and training in this study, and there was no redistribution or commercial use.

4.3. Bidirectional Matching

To evaluate semantic matching, we conducted two types of comparisons between candidate resumes and job descriptions. The job-to-resume direction checks how well candidate resumes rank for specific jobs. The resume-to-job direction looks at how well jobs rank for particular candidates. Figure 4 shows matching examples for both models. The fine-tuned SBERT model gives higher similarity scores to relevant pairs and ranks them more accurately, showing a better understanding of the specific field. The pre-trained model has lower scores and sometimes misranks relevant matches, which highlights the advantages of fine-tuning.

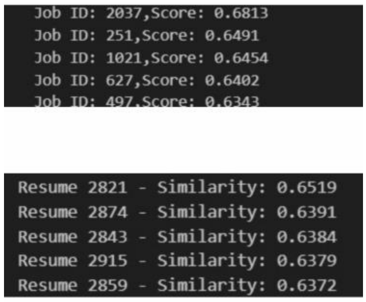

Figure 4 Pre-trained SBERT Model Bidirectional Resume–Job Matching

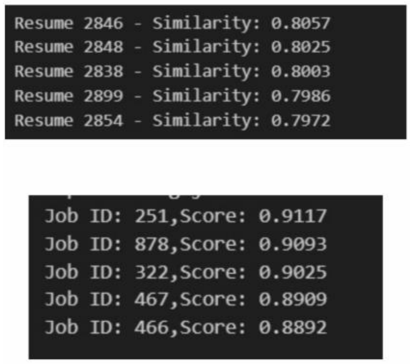

Figure 5 Fine-tuned SBERT Model Bidirectional Resume–Job Matching

4.4. Evaluation Metrics

The proposed resume-job matching system's effectiveness was measured using correlation measures and ranking-based retrieval metrics. These metrics were chosen to evaluate the system thoroughly, balancing statistical agreement with ground-truth labels and the performance of ranked candidate retrieval in real-world situations. First, correlation-based metrics like Spearman's rank correlation and Pearson's correlation coefficient were used. Spearman's correlation evaluated whether the predicted similarity scores maintained a consistent ranking order with the ground-truth labels. This ensured that more relevant matches received higher

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 12 December 2025
Page No: 4232-4239
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0620

scores. Pearson's correlation focused on the linear relationship between predictions and actual relevance values. It provided insight into how closely the model's outputs matched the annotated dataset. Figure 5 illustrates the comparison of these two metrics for both the pretrained and fine-tuned models, showing that the fine-tuned model consistently achieved higher correlation scores. This confirms the advantages of supervised training.

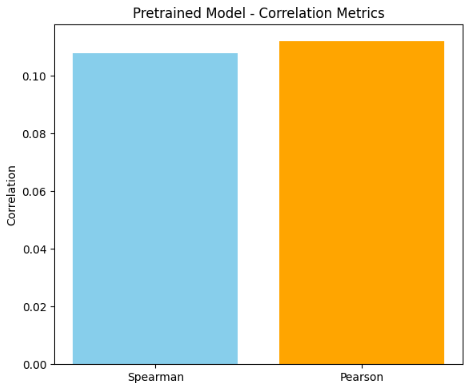

Figure 6 Pre-trained SBERT Model

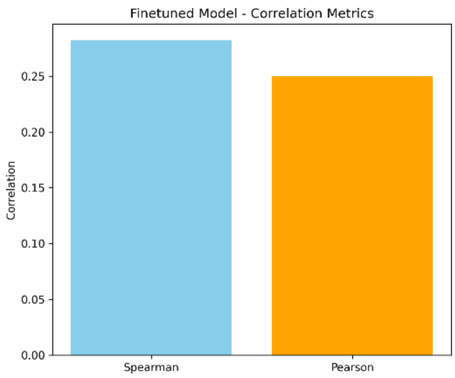

Figure 7 Fine-tuned SBERT Model

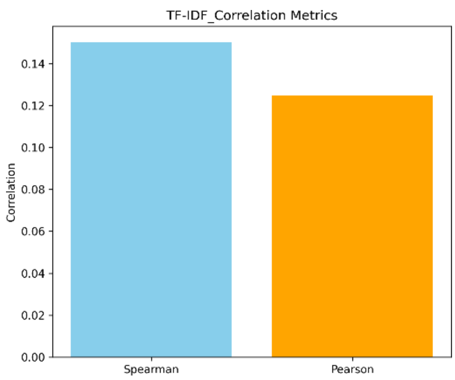

Figure 8 TF-IDF Correlation Metrics

In addition to correlation, ranking-based metrics were used to assess retrieval quality. Precision measured the proportion of relevant resumes retrieved within the top-k recommendations. This reflects the system's ability to return accurate matches quickly. Recall captured how many relevant resumes were retrieved overall. This emphasizes the system's ability to reduce missed candidates. Finally, Normalized Discounted Cumulative Gain (NDCG) assigned greater importance to relevant items that appear higher in the ranking. This is crucial in recruitment systems where the order of retrieved candidates significantly influences decision-making. Figure 6 shows the performance of the pretrained and fine-tuned models on these ranking-based metrics, highlighting the improvement in top- k retrieval after fine-tuning.

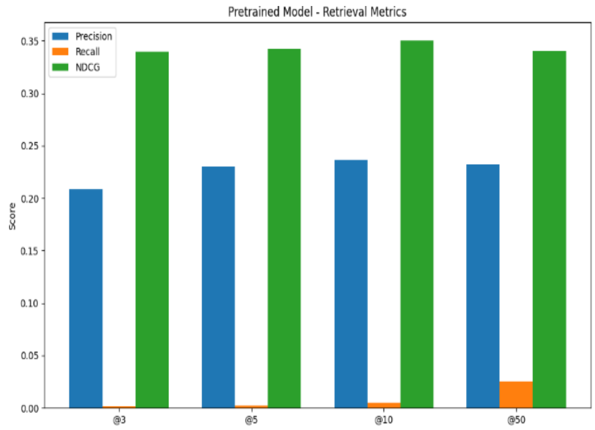

Figure 9 Pre-trained SBERT Model

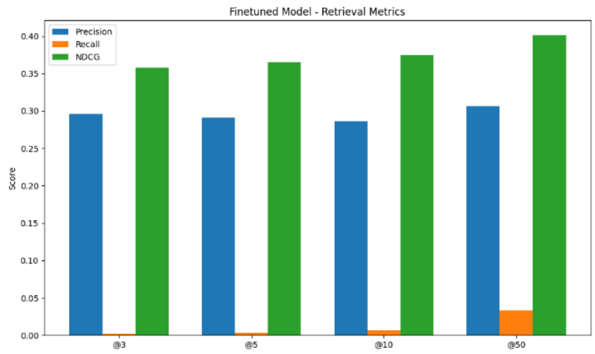

Figure 10 Fine-tuned SBERT model Fig. 6. Retrieval Metrics

By combining correlation analysis with ranking evaluations, the system's performance was validated from two different perspectives. The correlation metrics confirmed a statistical link between predicted

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 12 December 2025
Page No: 4232-4239
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0620

scores and labeled relevance. Mean- while, the ranking metrics demonstrated practical effective-ness in real-world recruitment tasks where ranked candidate lists are essential. Together, these evaluations show that the fine-tuned model is a more accurate and helpful approach to semantic resume-job matching, shown in Figure 7 to 10.

4.5. Comparison With Baseline

To test how effective the fine-tuned SBERT model is, we compared it with two baseline methods: the pretrained SBERT model and the classical TF-IDF method. The pretrained SBERT baseline shows semantic similarity without training for any specific domain. TF-IDF serves as a basic measure based on word frequency overlap. The earlier results (see Fig. 5 and Fig. 6) clearly show that the fine-tuned SBERT outperformed both baselines. While TF-IDF could identify surface-level textual similarity, it missed deeper semantic connections, resulting in lower retrieval accuracy. The pretrained SBERT performed better than TF-IDF but still didn't have task-specific optimization. In contrast, the fine-tuned model showed significant improvements in all evaluation metrics, especially in ranking relevant resumes and job descriptions higher in the retrieval list. This comparison shows that adding fine-tuning not only improves semantic understanding but also directly affects retrieval quality, making the system more useful for real-world recruitment [11-16].

Conclusion and Future Work

In this work, we developed a supervised fine-tuning approach for semantic resume and job matching using Sentence-BERT embeddings. By training on a dedicated dataset of resume and job pairs, the system learned a shared embedding space that improved the accuracy of similarity scoring and retrieval. The fine-tuned model performed better than the pretrained baseline, with noticeable improvements in correlation metrics and top-k retrieval scores. Compared to traditional methods like TF-IDF, which depend on lexical overlap, our approach captures semantic meaning. This allows for more precise and context-aware matching. These results show that fine- tuned embedding-based models offer a more effective solution for real-world recruitment systems by providing interpretable similarity scores and better retrieval quality. In summary, the proposed framework lays a strong foundation for semantic matching in recruitment systems. With ongoing improvements in data, model architectures, and inter-pretability, it offers potential to advance the automation and efficiency of hiring processes.

References

[1]. N. Reimers and I. Gurevych, "Sentence-BERT: Sen- tence Embeddings using Siamese BERT-Networks," 2019.
[2]. Dor Lavi, Volodymyr Medentsiy, and David Graus, "conSultantBERT: Fine-tuned Siamese Sentence- BERT for Matching Jobs and Job Seekers," 2021.
[3]. Julian Rosenberger, Lukas Wolfrum,Sven Weinzierl, and Patrick Zschech, "CareerBERT: Matching Re- sumes to ESCO Jobs in a Shared Embedding Space," 2023.
[4]. T. Sharma and P. Singh, "Resume2Vec: Transform- ing Applicant Tracking Systems with Embedding- Based Matching," 2025.
[5]. Asmita Deshmukh and Anjali Raut Dahake , "Com- paring bidirectional encoder representations from transformers and sentence-BERT for automated re- sume screening ," 2023.
[6]. Mehtap Saatci, Rukiye Kaya and Ramazan Ünlü, "Resume Screening with Natural Language Process- ing (NLP)," 2023.
[7]. Mohammed Shinan K S, Goutham Praveen P P and Ms. Varsha C R , "AI-Powered Resume Screening and Categorization Using Machine Learning," 2023.
[8]. Charan S N, Suhas G K,Yathisha L, and Devananda S N , "Job Recommendation System: Content-Based and Collaborative Filtering for Predictive Job Rec- ommendation Systems," 2023.
[9]. Dr.Prabakaran Paranthaman, B R Celia, Dr V Vijay- alakshmi and Dr. V. Senthil Kumaran, "Resume Screening Automation: Enhancing Recruitment Effi- ciency with Machine Learning Algorithms," 2022.
[10]. A. K. Jirjees et al., "Machine Learning for Recruit- ment: Analyzing Job Matching Algorithms," 2023.
[11]. Junshu Jiang, Songyun Ye, Wei Wang, Jingran Xu and Xiaosheng Luo, " Learning

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 12 December 2025
Page No: 4232-4239
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0620

Effective Represen- tations for Person-Job Fit by Feature Fusion," 2020.
[12]. [Chen Yang, Yupeng Hou, Yang Song, Tao Zhang, Ji- Rong Wen, Wayne Xin Zhao, "Modeling Two-Way Selection Preference for Person-Job Fit," 2022.
[13]. Chen Zhu et al, "Person-Job Fit: Adapting the Right Talent for the Right Job with Joint Representation Learning," 2018.
[14]. Hao Chen et al, ".Professional Network Matters: Connections Empower Person-Job Fit," 2023.
[15]. Xia Xue, Jingwen Wang, Bo Ma, Jing Ren, Yue Chang, Hongyu Wang, "Fine-Grained Semantics- Enhanced Graph Neural Network Model for Person- Job Fit," 2025.
[16]. Xia Xue et al, "Graph-based adaptive feature fusion neural network model for person-job fit," 2025.