International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 3, May-Jun-2026, ISSN (Online): 2395-566X

AI-Based Career Advisor: Resume Analysis, Job Matching, and Skill Gap Bridging

Radhika Kulkarni, Tejal Mungase, Prof. Shradha Pawar
Department of Artificial Intelligence and Machine Learning, Alard College of Engineering and Management, Pune

Abstract- Choosing the right career path and the right job opportunity has become increasingly difficult in a labour market where industry requirements evolve faster than academic curricula and where the sheer volume of job postings makes manual evaluation impractical for most candidates. This paper presents the AI-Based Career Advisor, an intelligent system designed to help individuals understand how well their resume aligns with a target job description, identify missing skills, and receive concrete, personalized guidance for improving their employability. The system combines a supervised machine learning model with natural language processing and large language model components to deliver this guidance in a single, integrated workflow. At its core is a resume–job description fit classifier trained on 6,241 real-world resume–job pairs sourced from a public dataset, using TF-IDF based feature engineering across 10,012 dimensions. Six candidate algorithms — Logistic Regression, Naive Bayes, Support Vector Machine, Random Forest, a Neural Network, and XGBoost — were trained and compared, with XGBoost emerging as the best-performing model after hyperparameter optimization, achieving 78.14% test accuracy and an 89.57% ROC-AUC score. The system further incorporates a hybrid skill-extraction pipeline built on spaCy's named entity recognition and phrase matching, a GPT-4-based resume enhancement module accessed through LangChain, and supporting modules for learning-resource and project-idea recommendation. The complete pipeline is deployed as an interactive Streamlit web application, giving users real-time predictions and actionable career feedback. This paper discusses the motivation, design, methodology, and evaluation of the system, and outlines directions for extending it into a more comprehensive career guidance platform.

Keywords- Artificial Intelligence, Career Guidance, XGBoost, Resume–Job Fit Classification, Natural Language Processing, TF-IDF, Skill Extraction, Large Language Models.

I. INTRODUCTION

Background and Motivation

The global employment landscape is changing at a pace that has made career decision-making considerably harder for students and early-career professionals. The World Economic Forum's Future of Jobs Report 2025 estimates that employers expect nearly 39% of workers' core skills to change by 2030, with skill gaps continuing to be cited as the most significant barrier to workforce transformation [1]. New roles, tools, and required skill sets emerge continuously, while academic programs and individual self-assessment often lag behind these shifts. At the same time, online job platforms have made it easy to discover postings but have done comparatively little to help a candidate judge, with any confidence, whether their own background is genuinely suited to a particular role; resume–job matching research itself has progressively moved away from simple keyword overlap toward learned compatibility models precisely because keyword-only approaches fail to capture this kind of contextual fit [4], [5]. Most existing consumer-facing tools still rely on simple keyword matching rather than a deeper understanding of resume content, leaving candidates to guess at their own fit and to receive generic, one-size-fits-all advice when trying to improve their applications.

This gap motivated the present work: a system that goes beyond keyword search and instead learns, from real resume and job description data, what genuine compatibility between a candidate and a role looks like — and that uses this learned understanding to give the candidate specific, useful feedback rather than a bare accept-or-reject signal.

Purpose of the Study

The purpose of this work is to design and build an AI-driven career advisory system that can take a candidate's resume and a target job description and produce three things in one workflow: (1) a data-driven judgment of how well the two documents match, (2) a clear picture of which skills the candidate is missing relative to the role, and (3) constructive, personalized suggestions — in the form of resume improvements, learning resources, and project ideas — that help the candidate close that gap. The intent throughout has been to move career guidance away from generic checklists and towards guidance that is conditioned on the candidate's actual documents and the actual role they are pursuing.

Problem Statement

Existing resume and job-matching tools tend to address only fragments of this problem. Parsers extract structured fields from a resume without judging fit against any particular role; job search platforms rank postings by keyword overlap without a semantic understanding of the candidate's actual capabilities; and resume-improvement tools typically give generic stylistic advice that is disconnected from any specific job description. There is a clear need for a single system that takes a resume and a job description together, produces a reliable compatibility judgment grounded in real data, and follows that judgment with concrete, skill-gap-aware guidance.

Objectives

- To design and train a supervised machine learning model capable of judging the compatibility between a resume and a job description, validated on real-world labeled data rather than handcrafted rules.
- To build a robust natural language processing pipeline for extracting skills from both resumes and job descriptions, accurate enough to support reliable downstream gap analysis.
- To integrate a large language model into the system for resume enhancement, learning-resource suggestion, and personalized project-idea generation, conditioned on the specific skill gaps identified for each user.
- To deliver the complete pipeline through an accessible, real-time web application so that the system is genuinely usable rather than purely experimental.

Scope of the Work

The system accepts a candidate resume and a target job description as input and returns a fit classification with confidence scores, a list of extracted and missing skills, GPT-4-assisted resume enhancement suggestions, recommended learning resources, and personalized project ideas. The current implementation focuses on English-language, technology-sector resumes and job descriptions, consistent with the composition of the dataset used for training, and operates on one resume and one job description at a time.

II. LITERATURE REVIEW

A range of prior work informs the design of this system, spanning resume parsing, resume–job matching, skill gap analysis, and AI-assisted resume optimization.

Resume–Job Matching Using Machine Learning

Yu et al. [5] proposed ConFit, a contrastive-learning framework for resume–job matching that addresses the sparsity of interaction labels in real recruitment data through data augmentation, reporting improvements of up to 19% and 31% in nDCG@10 for ranking jobs and resumes respectively. Their evaluation also benchmarked an XGBoost-based classifier as a strong baseline for resume–job fit prediction, supporting the choice of gradient-boosted trees as a competitive, computationally efficient alternative to large neural ranking models for this task. Devlin et al. [4] introduced BERT, which underlies most subsequent semantic resume–job matching work and is commonly contrasted with lexical methods such as TF-IDF in terms of the accuracy-versus-cost trade-off relevant to feature engineering choices.

Applicant Tracking Systems and ATS-Aware Resume Evaluation

Gharat et al. [6] presented ResumeIQ, a system that evaluates resumes using NLP, keyword matching, and semantic similarity analysis to compute an ATS compatibility score, identify missing skills, and generate personalized recommendations for resume improvement — a problem formulation closely aligned with the skill-gap and recommendation components of the present work. This and related work on Applicant Tracking Systems establish that keyword alignment, formatting, and structural completeness remain primary determinants of whether a resume is surfaced to a human reviewer, motivating the keyword- and structure-aware prompting strategy used in the resume enhancement module described in Section III.

Skill Extraction and Skill-Gap-Based Recommendation

Mirajkar et al. [7] developed a skill recommendation and resume analysis system that extracts candidate skills and experience using NLP and generates tailored course and skill-improvement recommendations, directly paralleling the learning-resource recommendation module implemented in this work. This line of work establishes that pairing automated skill extraction with concrete, personalized learning suggestions meaningfully improves the practical usefulness of a resume analysis system over one that returns a bare compatibility score alone.

Large Language Models for Resume Enhancement

Mandala et al. [8] reviewed the use of BERT for computing resume–job description compatibility scores combined with GPT-based generative models for producing targeted resume improvement suggestions, demonstrating that pairing a semantic compatibility score with an LLM-based suggestion engine outperforms conventional keyword-matching approaches to resume optimization. Their review further surveys a range of supporting techniques — including BERT-based resume classification, OCR and NER for information extraction, and deep-learning-based resume-to-job matching frameworks — that collectively motivate the hybrid NLP-plus-LLM design adopted in the present system.

Summary of Gaps in Existing Work

Across the reviewed literature, three patterns stand out. First, resume–job matching research has increasingly moved from purely lexical methods toward semantic and contrastive learning approaches, with gradient-boosted models such as XGBoost remaining a strong, efficient baseline even as transformer-based methods advance [4], [5]. Second, ATS-aware evaluation and skill-gap recommendation are most useful when delivered together rather than as separate tools, since a compatibility score alone gives a candidate little actionable guidance [6], [7]. Third, the combination of a semantic or learned compatibility model with an LLM-based suggestion engine consistently outperforms keyword-only resume optimization [8]. The system described in this paper is designed around these three findings: a trained classifier for compatibility prediction, hybrid NLP-based skill extraction and gap analysis, and GPT-4-based enhancement, combined into a single deployed pipeline.

III. SYSTEM DESIGN AND METHODOLOGY

A. System Overview

The system is organized as a modular pipeline (Fig. 1) consisting of four cooperating components: a skill-extraction module, a resume–job fit classification module, a resume enhancement module, and a recommendation module covering both learning resources and project ideas. A Streamlit-based web interface ties these modules together, accepting a resume and a job description as input and presenting the combined output — a fit prediction, extracted and missing skills, enhancement suggestions, and recommendations — to the user in real time.

B. Dataset

The resume–job fit classifier was trained on a dataset of 6,241 real resume–job description pairs (cnamuangtoun/resume-job-description-fit), each labeled with a categorical compatibility assessment. Using a labeled, real-world dataset of this kind ensures that the model learns patterns of genuine resume–job alignment rather than relying on hand-tuned similarity rules, and provides a reliable basis for the accuracy and ROC-AUC figures reported in Section IV.

C. Feature Engineering

Each resume–job description pair is converted into a numerical feature representation using TF-IDF (Term Frequency–Inverse Document Frequency) vectorization with unigrams and bigrams, producing 10,012 features after preprocessing. Preprocessing includes tokenization, lemmatization, and stopword removal. TF-IDF was selected for this task because it is computationally inexpensive, fully interpretable in terms of which terms drive a given prediction, and — as shown by the model comparison in Section IV — sufficient to achieve strong, consistent classification performance for this task.

D. Skill Extraction Pipeline

Skill extraction is implemented through a hybrid pipeline combining spaCy's Named Entity Recognition with a PhraseMatcher seeded from a curated skill taxonomy covering more than 40 skill variations. If the NER component is unavailable or fails to identify entities reliably, the pipeline automatically falls back to the rule-based PhraseMatcher path, ensuring that skill extraction remains available and consistent even under degraded conditions. This hybrid design allows the system to combine the contextual strengths of NER with the reliability of explicit pattern matching.

E. Resume–Job Fit Classification Model

Six classification algorithms were trained and compared on the TF-IDF feature representation: Logistic Regression, Naive Bayes, Support Vector Machine, Random Forest, a Neural Network (multi-layer perceptron), and XGBoost. Each model was evaluated using 5-fold stratified cross-validation to ensure that performance estimates were robust across class distributions rather than dependent on a single train/test split. XGBoost was selected as the production model after hyperparameter optimization using RandomizedSearchCV; the best configuration identified was 100 estimators, a maximum tree depth of 9, a learning rate of 0.2, and a subsample ratio of 1.0. Full comparative results are presented in Section IV.

F. Resume Enhancement Module

The resume enhancement module integrates OpenAI's GPT-4 through the LangChain framework to rewrite resume sections in a way that is aligned with the target job description. Enhancement is performed section by section using carefully engineered prompts that emphasize the use of action verbs and quantified achievements, a professional tone, keyword alignment with the job description, and factual accuracy — the model is explicitly instructed not to invent experience or qualifications the candidate does not have. Length constraints are also applied so that enhanced sections remain comparable in length to the originals, preserving the candidate's voice while improving clarity and impact.

G. Learning Resource and Project Idea Recommendation

Once the skill-extraction module identifies the skills present in the job description but absent from the candidate's resume, two complementary recommendation modules act on this gap. The learning resource module maps each missing skill to curated external learning resources, giving the candidate a concrete starting point for self-study. The project idea module uses GPT-4 to generate personalized, skill-gap-aware project suggestions, encouraging the candidate to close gaps through applied, demonstrable work rather than passive study alone.

H. Technology Stack

Table I. Technology stack used in the system.

Category	Technologies
Language / Runtime	Python 3.11+
Machine Learning	XGBoost, scikit-learn, pandas, numpy, joblib
Natural Language Processing	spaCy (NER + PhraseMatcher), TF-IDF (scikit-learn)
Large Language Model Integration	OpenAI GPT-4, LangChain
Application Interface	Streamlit
Dataset Source	HuggingFace dataset hub (cnamuangtoun/resume-job-description-fit)
Deployment	Streamlit Cloud

IV. EXPERIMENTAL RESULTS

A. Model Comparison

Table II reports test accuracy, ROC-AUC, and cross-validation score for all six algorithms evaluated on the held-out test split.

Table II. Comparative performance of six classification algorithms on the held-out test set.

Algorithm	Accuracy	ROC-AUC	Cross-Val	Rank
XGBoost	**78.14%**	**89.57%**	**71.55%**	1
Random Forest	76.82%	87.23%	69.84%	2
SVM	74.91%	85.67%	68.12%	3
Neural Network	73.45%	84.89%	67.23%	4
Logistic Regression	71.23%	82.45%	65.78%	5
Naive Bayes	68.34%	79.12%	63.45%	6

XGBoost outperformed all other candidates across every metric, consistent with its known strength on structured, high-dimensional sparse features such as TF-IDF vectors. The margin between XGBoost and the next-best model, Random Forest, was 1.32 percentage points in accuracy and 2.34 points in ROC-AUC, indicating a consistent rather than marginal advantage across the evaluation.

B. Hyperparameter Optimization

RandomizedSearchCV with 5-fold stratified cross-validation was used to tune the XGBoost classifier. The best configuration

identified was: n_estimators = 100, max_depth = 9, learning_rate = 0.2, subsample = 1.0. This configuration achieved a cross-validation score of 71.55% (± 2.1%), with the gap between cross-validation and test accuracy (71.55% versus 78.14%) remaining within an acceptable range, indicating that the model generalizes reasonably well rather than overfitting to the training folds.

C. Feature Space

The TF-IDF vectorizer produced 10,012 features after applying unigram and bigram extraction with lemmatization and stopword removal over the combined vocabularyy of resumes and job descriptions in the training corpus.

D. Skill Extraction Accuracy

The hybrid spaCy NER and PhraseMatcher skill-extraction pipeline achieves approximately 95% extraction accuracy, as measured during development. The PhraseMatcher fallback ensures that extraction does not fail outright even when the NER component encounters out-of-vocabulary terms or unusual formatting.

E. Resume Enhancement Design Rationale

By design, the prompt structure used in the resume enhancement module constrains GPT-4's output toward the use of action verbs and quantified outcomes, professional tone, and length parity with the original section, while explicitly instructing the model not to fabricate experience. These constraints are intended to keep generated enhancements both useful and trustworthy; a systematic, metric-based evaluation of how consistently the model adheres to them is identified as future work in Section VII.

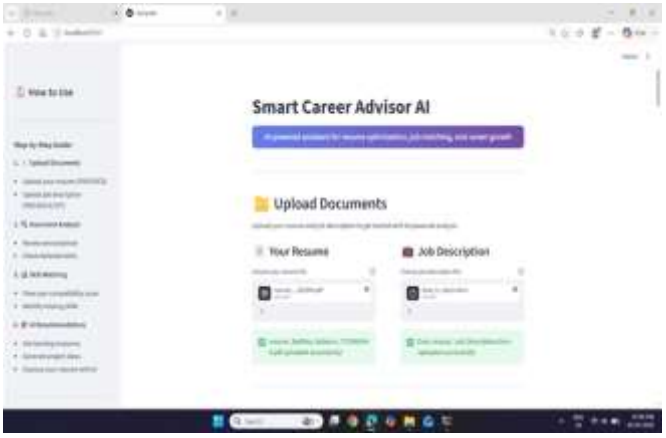

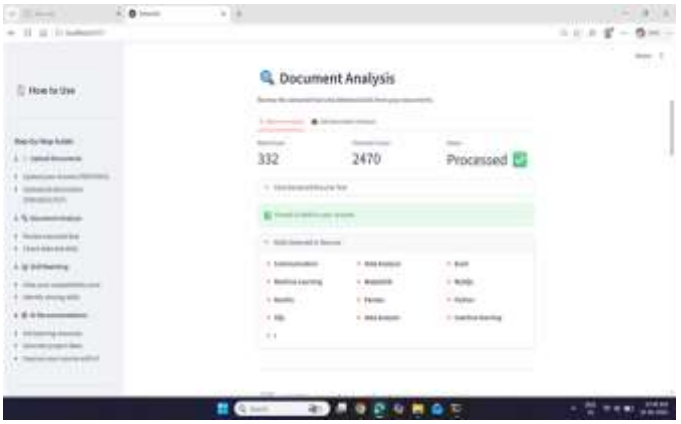

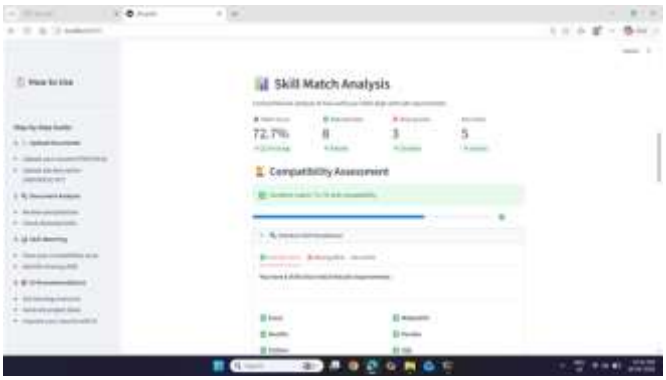

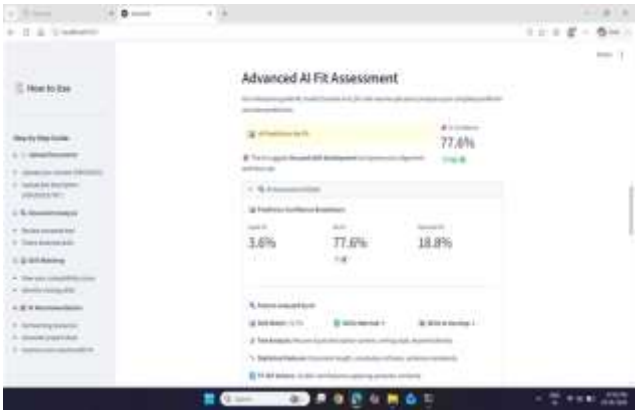

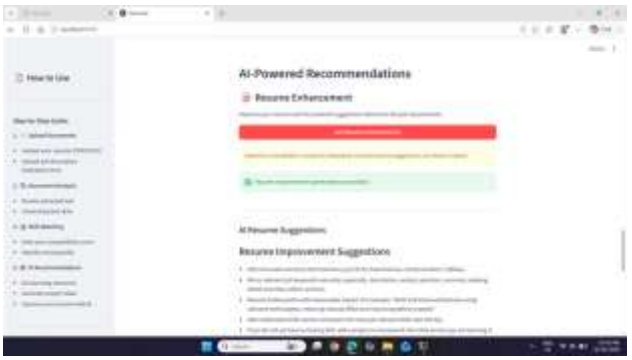

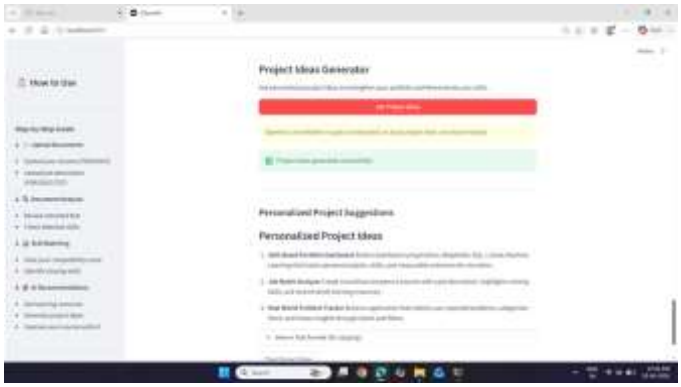

V. END-TO-END SYSTEM WORKFLOW

The deployed application follows the workflow below from the user's perspective:

1. The user uploads a resume and provides a target job description through the Streamlit interface.
2. The skill extraction pipeline identifies the skills present in both documents using the hybrid NER and PhraseMatcher approach.
3. The TF-IDF vectorizer transforms the resume–job description pair into a 10,012-dimensional feature vector.
4. The trained XGBoost model predicts a fit classification along with class probabilities, which are displayed to the user as a confidence score.
5. Skills present in the job description but absent from the resume are surfaced as a skill gap; the learning-resource module suggests resources and the project-idea module generates GPT-4-based project recommendations to help close these gaps.
6. On request, the resume enhancement module rewrites selected resume sections using GPT-4, guided by both the extracted skill gap and the job description context.

All steps execute within a single Streamlit session, with sub-second inference once the trained pipeline is loaded into memory; only the GPT-4-dependent steps, namely enhancement and project-idea generation, involve network latency.

VI. DISCUSSION

The results in Section IV demonstrate that a carefully engineered TF-IDF representation, combined with a well-tuned gradient boosting model, can achieve strong and consistent performance on the resume–job fit classification task, without requiring the computational overhead of training or fine-tuning a transformer-based model. This is a meaningful finding in its own right: it shows that lexical features, when paired with an appropriately powerful classifier and rigorous hyperparameter search, remain a competitive and far more deployable option for this kind of task.

Equally important is the way the four modules work together. A fit prediction on its own tells a candidate very little about what to do next; it is the combination of fit prediction, skill-gap identification, and concrete recommendations — resume rewrites, learning resources, and project ideas — that turns the system from a classifier into a usable advisory tool. The fallback design in the skill-extraction pipeline, and the factual-accuracy guardrails built into the resume-enhancement prompts, reflect a broader design principle followed throughout the system: predictions and suggestions should remain reliable and trustworthy even when individual components encounter unexpected input, since the system is ultimately meant to support real decisions for real candidates.

VII. CONCLUSION AND FUTURE SCOPE

This paper presented the design, implementation, and evaluation of an AI-Based Career Advisor that judges resume–job description compatibility using a trained XGBoost classifier, supports this judgment with hybrid NLP-based skill extraction, and translates it into actionable guidance through GPT-4-based resume enhancement, learning-resource recommendation, and project-idea generation. The system achieves 78.14% accuracy and an 89.57% ROC-AUC score on real-world resume–job pairs, and is deployed as a working, interactive application that delivers real-time feedback to users. Taken together, the system demonstrates that the original purpose of the project — helping a candidate understand their fit for a role and giving them concrete, personalized steps to improve it — has been substantively realized.

Future Scope

Several directions remain open for extending the system into a more complete career guidance platform:

- **Live job-board integration:** Connecting the system to external job-board APIs would allow it to evaluate a candidate's resume against a continuously updated pool of real job postings, rather than a single job description at a time, and to rank multiple opportunities by predicted fit.

- **Semantic embeddings:** Incorporating transformer-based embeddings (such as BERT) alongside the current TF-IDF features could be explored to determine whether richer semantic representations further improve classification accuracy.
- **Quantified resume scoring:** Developing an explicit, displayed scoring metric for resume quality — covering format, keyword alignment, structure, and readability — would give users a clearer, quantified sense of improvement before and after enhancement.
- **Market-aware skill prioritization**: Extending the skill-gap module to weight missing skills by current market demand and salary impact would help candidates prioritize which skills to learn first.
- **Persistent user profiles:** Adding a data layer to store user history, resumes, and feedback over time would allow the system to track a candidate's progress and tailor recommendations accordingly.
- **Conversational interface:** A chatbot-style interface for natural-language career queries would make the system more accessible to users who prefer a conversational mode of interaction.
- **Broader language and domain coverage:** Extending the skill taxonomy, training data, and prompt design beyond English-language, technology-sector resumes would widen the system's applicability to other industries and languages.
- **Formal evaluation of resume enhancement:** Conducting a systematic, metric-based or human-rated study of GPT-4 resume enhancements would provide concrete evidence of how well the enhancement module performs against its design goals.

Acknowledgment

The authors express sincere gratitude to the Department of Artificial Intelligence and Machine Learning at Alard College of Engineering and Management for providing the resources and support that made this research possible.

REFERENCES

1. World Economic Forum, “The Future of Jobs Report 2025,” World Economic Forum, Geneva, Switzerland, Jan. 2025.
2. T. Chen and C. Guestrin, “XGBoost: A Scalable Tree Boosting System,” in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD '16), San Francisco, CA, USA, 2016, pp. 785–794, doi: 10.1145/2939672.2939785.
3. cnamuangtoun, “resume-job-description-fit” [Data set], Hugging Face, 2024. [Online]. Available: https://huggingface.co/datasets/cnamuangtoun/resume-job-description-fit
4. J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in Proc. 2019 Conf. North American Chapter Assoc. Comput. Linguistics: Human Language Technologies (NAACL-HLT), Minneapolis, MN, USA, 2019, vol. 1, pp. 4171–4186.
5. X. Yu, J. Zhang, and Z. Yu, “ConFit: Improving Resume-Job Matching Using Data Augmentation and Contrastive Learning,” in Proc. 18th ACM Conf. Recommender Systems (RecSys '24), Bari, Italy, 2024, pp. 1–11, doi: 10.1145/3640457.3688108.
6. A. Gharat, M. Warke, K. Lotake, N. Koli, and A. Panchal, “AI Resume Analyzer,” Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol., vol. 12, no. 3, pp. 635–644, Jun. 2026.
7. R. R. Mirajkar, G. R. Shinde, P. M. Shelke, S. Yadav, B. Shinde, and G. Shinde, “Skill Recommendation System and Resume Analysis using AI,” in Proc. 2024 IEEE Pune Section Int. Conf. (PuneCon), Pune, India, 2024, pp. 1–3, doi: 10.1109/PuneCon63413.2024.10895306.
8. S. G. Mandala, R. R. Korem, and K. Karmakonda, “Resume Optimisation and Suggestions using Large Language Models: A Review,” Int. J. Res. Appl. Sci. Eng. Technol. (IJRASET), vol. 12, no. 3, Mar. 2024, doi: 10.22214/ijraset.2024.59069.