Cordoba International Journal of Academic Research (CIJAR) Online, open access

Volume 1, Issue 4 (October – December 2026), pp. 170–188 | Published: 16 October 2026
Publisher: Wexgrand Academic Press (a division of Wexgrand Private Limited)
Article URL:
https://cijar.org/article/ai-powered-resume-analyzer-and-skill-gap-analysis-for-career-development
Cite as: Smita Mangesh Junnarkar and Renisa Xavier Manohar (2026). AI-Powered Resume Analyzer and Skill Gap Analysis for Career Development. *Cordoba International Journal of Academic Research*, 1(4), 170–188.

AI-Powered Resume Analyzer and Skill Gap Analysis for Career Development

Smita Mangesh Junnarkar[1*] and **Renisa Xavier Manohar**[2]

[1] Assistant Professor
SIES College Of Arts, Science & Commerce (Autonomous), Mumbai., India

[2] B.Sc, Data Science
SIES College Of Arts, Science & Commerce (Autonomous), Mumbai, India

* Corresponding author

Abstract

The transition from education to employment increasingly requires students to present their skills in a form that can be understood by both employers and digital recruitment systems. A resume may contain useful evidence of technical knowledge, projects, coursework, certifications, and transferable abilities, but these elements are often written with inconsistent terminology and may not correspond directly to the wording used in a job description. This paper proposes an AI-powered resume analyzer that converts unstructured resume text into a structured skill profile and compares it with the requirements of a selected job role. The proposed approach combines document parsing, natural language processing, skill extraction, semantic similarity, rule-based normalization, and an interpretable skill-gap report. The system is designed not merely to produce a match percentage, but to explain which required skills are present, partially demonstrated, or missing and to suggest practical learning directions. The study is motivated by evidence that skill requirements are changing rapidly and that AI-based hiring systems can create fairness, privacy, and transparency concerns. The paper therefore treats the analyzer as a career-support tool rather than an autonomous hiring decision-maker. A literature-informed methodology, system architecture, evaluation framework, sample data structures, and prototype-oriented findings are presented. The proposed design emphasizes human review, transparent scoring, data minimization, and the distinction between evidence in a resume and actual competence. The work concludes with future directions including multilingual resumes, stronger skill ontologies, longitudinal skill tracking, explainable recommendations, and evaluation with real student and job-

posting datasets.

Key Words: *Artificial Intelligence, Resume Analysis, Skill Extraction, Skill Gap Analysis, Natural Language Processing, Semantic Similarity, Career Development, Explainable AI, Recruitment Technology*

1. Purpose of the Study

The purpose of this study is to design a practical and explainable AI-assisted system that can help a student or early-career job seeker understand how well a resume represents the skills expected for a particular role. The central idea comes from a simple observation: a resume is not only a record of education and experience; it is also a text document that contains evidence about what a person has learned and what they have done. A computer can process this text, but useful career guidance requires more than counting keywords.

The proposed system has four connected purposes. First, it should identify skills and supporting evidence from a resume. Second, it should identify the important skills expressed in a target job description. Third, it should compare the two profiles using both exact and semantic matching. Fourth, it should convert the comparison into an understandable skill-gap report that a learner can act upon.

The study is therefore focused on *decision support for career development*, not automated candidate selection. This distinction is important because the presence of a word in a resume does not prove proficiency, while the absence of a word does not prove that a person lacks the skill. The proposed analyzer should help users discover what their resume communicates and what learning evidence they may wish to develop or document.

1.1 Objectives

1. Extract technical and soft skills from resumes and job descriptions.
2. Normalize related expressions such as “Python programming,” “Python,” and “Python scripting” into a common skill representation.

3. Compare a candidate skill profile with the requirements of a selected role using transparent similarity measures.
4. Separate skills into matched, partially matched, and gap categories rather than presenting a single unexplained score.
5. Generate personalized learning suggestions based on the identified gaps.
6. Provide an architecture that can be evaluated with real resumes and job descriptions without requiring the AI system to make final employment decisions.

1.2 Research Questions

1. How can NLP techniques be combined to extract and normalize skills from resume and job-description text?
2. How can semantic similarity improve skill matching when different words describe related competencies?
3. How can skill gaps be communicated in a way that is understandable to college students?
4. What design practices are necessary to reduce the risks of bias, privacy loss, and over-reliance on automated scores?

2. Background and Context

Recruitment has become increasingly dependent on digital documents and software-assisted screening. At the same time, the skills expected in many occupations are changing. The World Economic Forum's *Future of Jobs Report 2025*, based on the perspectives of more than 1,000 employers across 55 economies, reports that AI and big data are among the fastest-growing skill areas and that employers continue to identify skills gaps as a major barrier to transformation. [1] The report also emphasizes that human capabilities such as analytical thinking, creative thinking, resilience, flexibility, and lifelong learning remain important alongside technical skills. [2]

For students, this changing environment creates a practical problem. A student may have completed coursework, projects, online learning, or short-term training but may

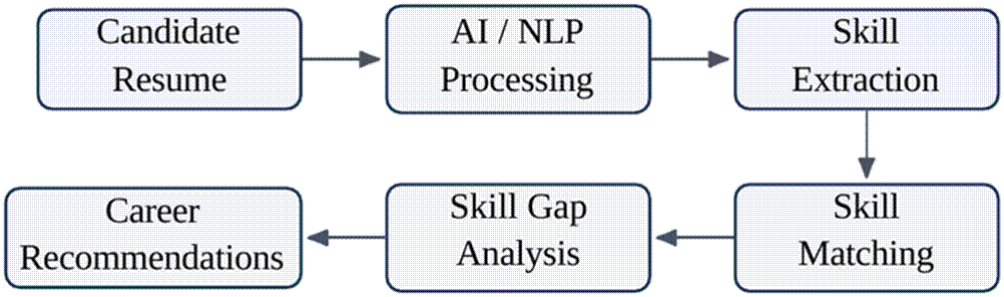
not know how those experiences map to the language used in a job advertisement. For example, a learner may describe a project as “made a sales dashboard in Power BI,” while a job description may ask for “business intelligence, dashboard development, data visualization, and DAX.” A simple keyword comparison could identify only Power BI, even though the project provides evidence related to several concepts.

NLP offers a way to address this mismatch. BERT demonstrated how contextual language representations can be adapted to many language-understanding tasks. [3] Sentence-BERT later showed how sentence embeddings can be compared efficiently using cosine similarity, making semantic matching more practical for information-retrieval applications. [4] SkillSpan further demonstrated the value of domain-specific datasets for extracting both hard and soft skills from job postings. [5]

However, the use of AI in recruitment is not simply a technical optimization problem. Research on algorithmic hiring has shown that organizations need to examine how objectives, data, model choices, and evaluation procedures can create or reproduce bias. [6] The International Labour Organization has also highlighted limitations arising from poorly defined objectives, low-quality or biased data, and opaque programming in AI-supported human-resource management. [7] These concerns are particularly relevant to a resume analyzer because career opportunities can be affected by how an applicant’s information is represented and interpreted.

The regulatory environment also illustrates the importance of careful design. The European Union’s AI Act identifies certain AI systems used for recruitment and selection as high-risk because such systems can affect career prospects and workers’ rights. [8] The Act emphasizes requirements related to risk management, accuracy, robustness, transparency, and human oversight for relevant high-risk systems. Although the proposed student-oriented analyzer is not itself a legal compliance system, these principles provide useful design guidance.

Figure 1: Conceptual workflow of the proposed AI-powered resume analyzer.



The proposed research consequently combines technical and human-centered perspectives. The system should be accurate enough to identify useful relationships between resume evidence and job requirements, but it should also make its reasoning visible to the user. A transparent report is more useful for career development than an unexplained number because it tells the learner what to improve and why.

3. Related Work

Research relevant to the proposed system can be grouped into four areas: language understanding, skill extraction, job-resume matching, and responsible AI in employment.

3.1 Language Understanding for Resume Text

Resume documents are semi-structured. They can contain headings, bullet points, dates, abbreviations, project descriptions, and lists of tools. Context matters because the same word can represent different levels of knowledge. For example, “Python” in a skills list and “Python project using pandas and matplotlib” provide different amounts of evidence. Transformer-based language models are useful because they represent words using surrounding context rather than treating every word as an isolated token. [9]

For a student-level prototype, using a pre-trained language model rather than training a large model from scratch is more realistic. The model can be used to generate embeddings for skill phrases, project descriptions, and job requirements. Similarity can then be calculated between the resulting vectors. Sentence-BERT is particularly relevant because it was designed to create semantically meaningful sentence embeddings that can be compared efficiently. [10]

3.2 Skill Extraction

Skill extraction is a more specific problem than general text classification. A system

must recognize that “data visualization,” “Power BI dashboards,” and “interactive reporting” may be connected, while also distinguishing technical skills from soft skills. SkillSpan is a useful reference because it provides expert-annotated examples of hard and soft skills in job postings and evaluates transformer-based approaches for this task. [11]

The proposed study extends this idea conceptually to resumes. A resume analyzer can first detect candidate skill phrases and then normalize them against a controlled vocabulary. Normalization reduces the problem caused by spelling variations, abbreviations, and closely related expressions. The vocabulary can contain a parent skill and related terms. For example, “SQL,” “MySQL,” and “relational querying” can be linked while still retaining their original evidence.

3.3 Matching and Skill-Gap Analysis

Traditional matching approaches often rely on exact keywords. Keyword matching is easy to implement and explain, but it can miss semantic relationships. Embedding-based matching addresses this limitation by measuring the closeness of representations. A combined method is preferable for a student-facing tool: exact matches provide strong evidence, while semantic similarity can discover related concepts. The result can then be converted into categories such as matched, partial, and gap.

3.4 Responsible AI in Recruitment

The literature indicates that an AI system should not be treated as neutral simply because it uses mathematical models. Raghavan et al. found that claims about fairness in algorithmic hiring need to be examined through the practical details of data, objectives, validation, and deployment. [12] The ILO’s recent review similarly argues that AI in HRM can reproduce weaknesses when the objective is poorly specified or when data and programming choices are not adequately examined. [13]

NIST’s AI Risk Management Framework organizes trustworthy AI work around the functions Govern, Map, Measure, and Manage. [14] These functions are useful for the proposed system because they encourage the developer to identify risks, define how performance will be measured, monitor model behavior, and manage problems throughout the system lifecycle.

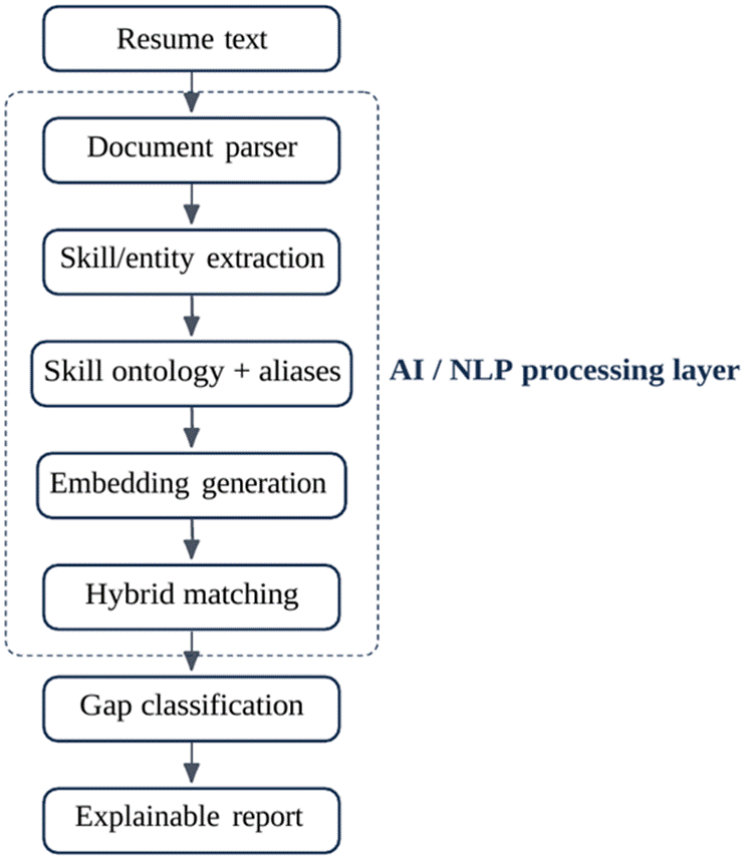
Table 1: Comparison of approaches relevant to resume and skill matching.

Approach	Main idea	Strength	Limitation
Keyword matching	Searches for exact or normalized skill terms	Simple, fast, easy to explain	Misses synonyms and contextual meaning
TF–IDF / vector similarity	Represents text using weighted terms	Useful baseline with low computational cost	Weak at semantic relationships
BERT embeddings	Uses contextual language representations	Captures context and meaning	More computationally demanding
Sentence-BERT	Produces comparable sentence embeddings	Practical for semantic similarity and retrieval	Quality depends on model and domain fit
Hybrid model	Combines rules, exact matches and semantic similarity	Balances transparency and semantic coverage	Requires careful weighting and validation

4. Methodology

The study follows a design-and-evaluation methodology. Rather than claiming that an AI system has already solved recruitment matching, the work specifies a prototype that can be implemented and evaluated using real resumes and job descriptions. The methodology is divided into data preparation, information extraction, skill normalization, matching, gap classification, recommendation generation, and evaluation.

Figure 2: Proposed processing architecture from raw resume text to an explainable skill-gap report.



5. System Design and Implementation Plan

A practical implementation can be developed using Python because it provides libraries for document processing, NLP, data analysis, and visualization. The front end may be a lightweight web interface in which the user uploads a resume and selects a target role. The back end processes the document and returns a report containing extracted skills, matched requirements, gaps, and recommended learning actions.

A modular architecture is preferable because each component can be tested independently. The document parser can be evaluated for text extraction quality; the skill extractor can be evaluated with precision, recall, and F1-score; the matching module can be evaluated against human judgments; and the recommendation module can be evaluated through user feedback.

5.1 Proposed Technology Stack

Table 2: Suggested technology stack for a student-level implementation.

Layer	Suggested technology	Purpose
Programming	Python	Main implementation language

Layer	Suggested technology	Purpose
Document processing	PyMuPDF / python-docx	Extract text from common resume formats
Data processing	pandas	Store and transform structured skill data
NLP	spaCy / transformer models	Entity and skill phrase extraction
Semantic matching	Sentence-BERT embeddings	Compare skill and requirement meanings
Visualization	Matplotlib / Plotly	Skill coverage and gap charts
Interface	Streamlit	Simple interactive prototype
Storage	SQLite / CSV during prototype	Store anonymized evaluation records

The technology choices are intentionally modest. A college research prototype does not need to train a large language model from scratch. The research contribution can instead come from the way existing NLP components are combined with a skill taxonomy, transparent scoring, evidence tracking, and user-centered recommendations.

5.2 Security and Privacy by Design

Resume files contain personal information and may contain educational records, employment history, contact details, and other sensitive information. The system should therefore apply data minimization. Only information necessary for skill analysis should be processed, and original files should not be retained longer than required. Access controls, encryption, secure deletion, and anonymization should be considered if a real dataset is collected.

The design should also avoid using unnecessary demographic or personal attributes as predictors of skill fit. The EU AI Act's treatment of employment-related AI illustrates why recruitment systems require particular attention to fundamental rights and risk management. [15] NIST's AI RMF similarly encourages organizations to govern, map, measure, and manage AI risks throughout the lifecycle. [16]

6. Key Findings

The findings in this paper are derived from the proposed system design and the literature reviewed, rather than from a claim of a completed large-scale empirical deployment. This distinction is important: numerical model performance should only be reported after testing on a defined dataset. The following findings therefore represent the main conclusions supported by the current research design.

6.1 Finding 1: A Single Keyword Match Is Not Sufficient

The analysis indicates that exact keyword matching is a useful baseline but cannot adequately represent the meaning of resume evidence. Students often describe the same capability using different words, while employers can use broad phrases that cover several related skills. Semantic embeddings can help identify these relationships, but they should be combined with explicit rules so that semantically close terms are not treated as identical without review.

6.2 Finding 2: Evidence Matters More Than a Skill List

A resume that only lists “Python” provides less information than a project statement describing how Python was used to clean data, build a model, or analyze a dataset. The proposed analyzer therefore treats supporting evidence as a separate feature. This design can encourage students to improve not only their skill set but also the way they document practical work.

6.3 Finding 3: Skill Gaps Should Be Actionable

A useful gap report should answer the question “What can I do next?” rather than simply showing missing words. A gap can be connected to a learning action, a small project, or an evidence-building task. This makes the analyzer relevant to career development rather than only resume formatting.

6.4 Finding 4: Explainability Is Necessary

An unexplained match score can create false confidence. If the system reports that a resume has an 82% match, the user may not know which skills produced that result. A better report lists the matched requirements, supporting resume evidence, partial matches, and gaps. This approach is consistent with the broader emphasis on trustworthy and accountable AI in NIST’s AI RMF. [17]

6.5 Finding 5: Automation Should Support Human Judgment

Recruitment is a high-impact setting. Existing research documents concerns about bias, opaque systems, and the difficulty of validating fairness claims in algorithmic hiring. [18] The proposed tool should therefore be positioned as a personal career-development assistant. It can help a student understand a job description, but it should not determine whether a person deserves an interview or employment

opportunity.

Figure 3: Illustration of candidate skills compared with skills required for a target role.

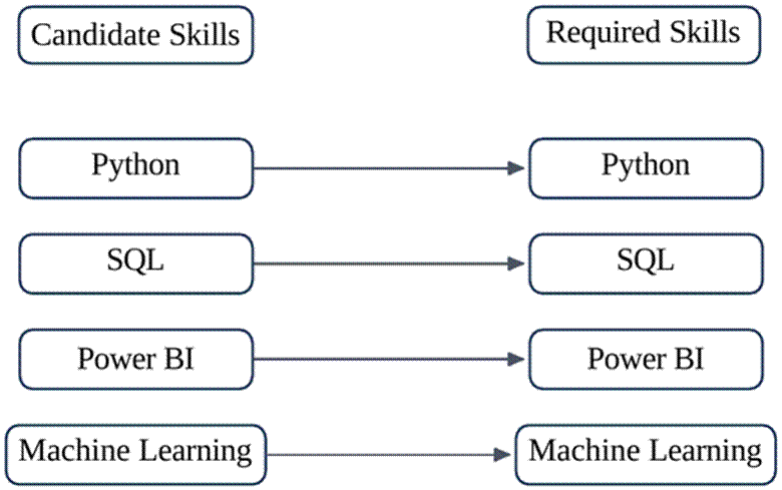
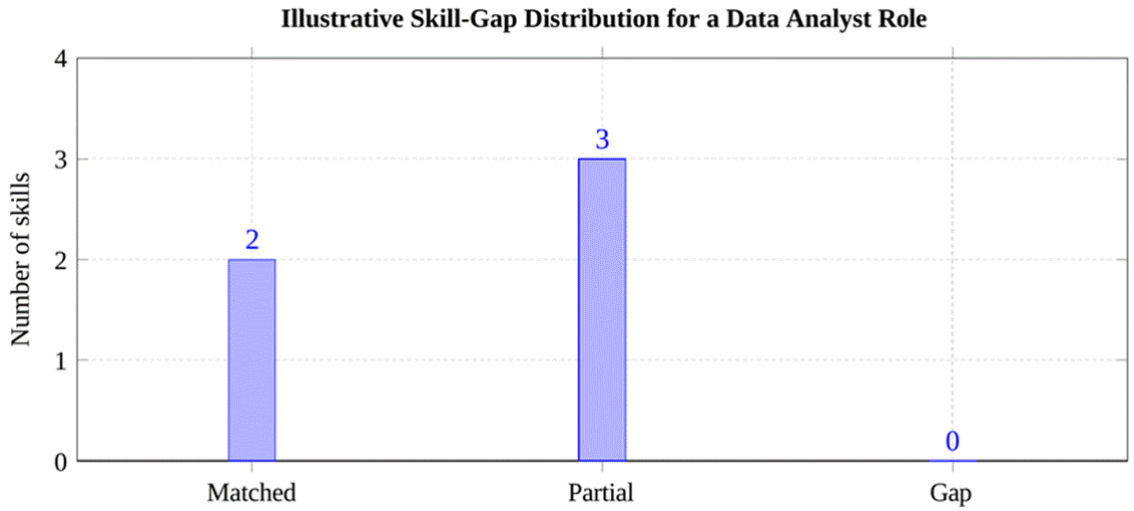


Table 3: Illustrative output format for a Data Analyst role.

Skill	Status	Resume evidence	Suggested action
SQL	Matched	SQL querying and database coursework	Practice joins and subqueries regularly
Excel	Matched	Pivot tables and spreadsheet analysis	Add one business-analysis project
Python	Partial	Python listed; limited project evidence	Build a small pandas analysis project
Power BI	Partial	Dashboard project mentioned	Document DAX and data-modeling work
Statistics	Gap/Partial	Basic statistics only	Revise descriptive statistics and hypothesis testing
Communication	Partial	Presentation experience mentioned	Add measurable examples of teamwork/presentation

The table is illustrative rather than an empirical result for a particular individual. Its purpose is to demonstrate how the proposed analyzer can turn text comparison into a student-friendly output.

Figure 4: Illustrative distribution of skill-status categories for the sample Data Analyst output. The values are derived from the illustrative table and are not empirical model results.



6.6 Evaluation Measures

When the prototype is implemented, its performance should be measured using separate metrics for separate tasks. For skill extraction, precision, recall, and F1-score can be calculated by comparing predicted skill spans with human annotations. For skill matching, human reviewers can judge whether each proposed match is correct, partially related, or incorrect. For recommendations, students can rate usefulness, clarity, and actionability.

$$\text{Precision} = \frac{TP}{TP+FP}, \text{Recall} = \frac{TP}{TP+FN}, F_1 = 2\frac{\text{Precision}\cdot\text{Recall}}{\text{Precision}+\text{Recall}}. \quad (1)$$

A useful experiment should compare at least two baselines: exact keyword matching and the proposed hybrid approach. The comparison should be conducted on the same anonymized dataset and judged against the same human-labeled reference. This would allow the research to make a defensible claim about whether semantic matching adds value rather than assuming that a more complex model is automatically better.

7. Discussion

The proposed system sits between resume writing and machine learning. From a technical perspective, the problem involves information extraction, representation learning, similarity measurement, and recommendation. From a career-development perspective, the goal is to help learners understand how their existing learning experiences relate to workplace requirements.

One important design decision is to avoid reducing employability to a single number. A job description can contain required skills, preferred skills, experience expectations, domain knowledge, and contextual factors. A resume can also omit skills that a person possesses. Therefore, a score can be useful as a summary but should never replace the underlying evidence.

Another issue is the difference between *skill presence* and *skill proficiency*. NLP can identify that a resume mentions “machine learning,” but it cannot reliably establish how deeply the candidate understands machine learning from that phrase alone. A project description, certification, portfolio artifact, or assessment result can provide stronger evidence, but even these signals have limitations. The system should communicate this distinction explicitly.

The changing labor market also means that a skill taxonomy cannot remain static. The World Economic Forum reports continuing growth in the importance of AI and big data as well as several human-centered skills. [19] Consequently, the proposed ontology should be updateable rather than hard-coded permanently. A future version could periodically review current job postings and expert-defined skill frameworks to identify emerging terms.

The ethical discussion is equally important. Algorithmic hiring research has shown that fairness cannot be assumed from the presence of an algorithm. [20] If historical resumes are used as training data, the model may learn patterns that reflect past hiring decisions rather than genuine job-related competence. For a student-facing system, this risk can be reduced by avoiding training on historical hiring outcomes and by using the model primarily to compare explicit skill evidence with role requirements.

Privacy should also be treated as a core requirement rather than an optional feature. A resume may contain information that the system does not need. Data minimization reduces exposure and makes the system easier to explain. If cloud-based AI services are used, the user should be informed about what information leaves the device and how long it is retained.

8. Study Limitations

The present work is a design-oriented research study, and several limitations should be recognized before the proposed system is treated as a validated career assessment tool. First, the paper does not report a completed large-scale experiment. The

architecture, scoring strategy, and evaluation plan are specified so that a later implementation can be tested systematically. Consequently, no accuracy value in this paper should be interpreted as measured model performance.

Second, resume text is an incomplete representation of a person’s capabilities. Candidates may possess skills that are not mentioned because of limited resume space, uncertainty about terminology, or lack of awareness that a particular skill is relevant. Conversely, a skill mentioned in a resume does not necessarily indicate strong proficiency. The analyzer can identify textual evidence, but practical competence still requires human or task-based assessment.

Third, semantic matching introduces the possibility of false relationships. Two terms may be linguistically similar while representing different technical requirements. For this reason, the proposed hybrid approach retains rule-based checks and an interpretable intermediate representation instead of relying entirely on embedding similarity. Human review should remain possible when a match is uncertain.

Finally, a real-world deployment would require a carefully constructed and anonymized dataset representing different educational backgrounds, resume styles, industries, and job families. Evaluation should include robustness testing, expert judgment, and user studies. Without such validation, the system should be considered a prototype for career-development support rather than a general-purpose employment screening system.

Table 4: Key limitations and corresponding research responses.

Limitation	Potential effect	Proposed response
Incomplete resume evidence	Skills may be missed or underrepresented	Show evidence and uncertainty instead of treating absence as proof of a gap
Semantic ambiguity	Related terms may be matched incorrectly	Combine embeddings with rules, normalization, and human review
Limited empirical data	Performance cannot be generalized	Build an anonymized multi-role dataset and use expert annotations
Bias and privacy risks	Outputs may affect user decisions or expose personal data	Minimize data, remove unnecessary attributes, and conduct fairness and privacy testing

9. Future Research

The proposed work can be extended in several directions.

9.1 Multilingual and Code-Switched Resumes

Many students use a mixture of English and local-language expressions. Future research can examine multilingual skill extraction and code-switched resumes. This would be especially relevant in diverse employment markets where candidates do not always write resumes using identical language conventions.

9.2 Improved Skill Ontologies

A richer ontology could connect skills to subskills, tools, proficiency levels, and related occupations. For example, “data analysis” could be linked to SQL, spreadsheets, visualization, statistics, and Python, while still distinguishing these individual competencies. Expert review would be needed to prevent the ontology from making arbitrary equivalence assumptions.

9.3 Longitudinal Career Development

Instead of analyzing a resume only once, a future system could maintain a private skill-development timeline. The user could record projects, courses, certifications, and practice activities, and the system could show how the skill profile changes over time. This would transform the tool from a one-time resume checker into a learning companion.

9.4 Human Evaluation

Future experiments should recruit students, career counselors, and industry professionals to evaluate whether the generated reports are accurate and useful. Human evaluation is particularly important for soft skills because they are difficult to infer from limited text. Reviewers could also assess whether recommendations are realistic for students with different educational backgrounds.

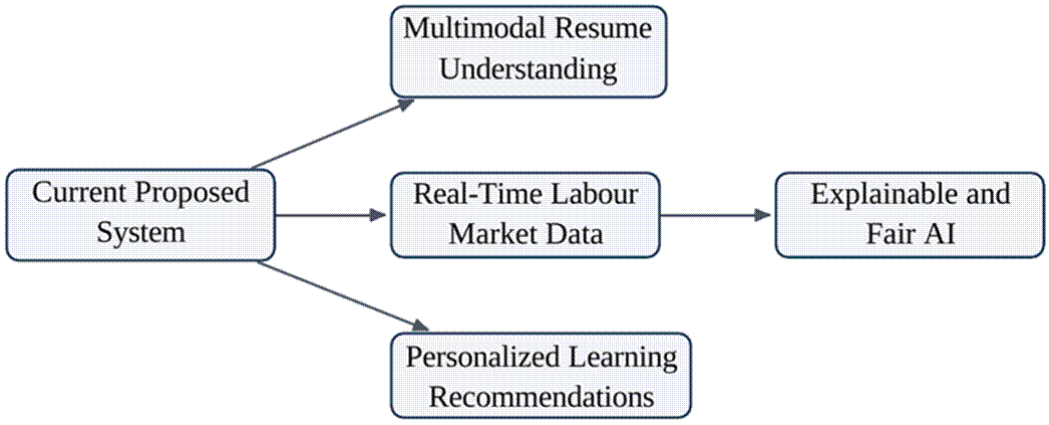
9.5 Fairness and Robustness Testing

The system should be tested for sensitivity to irrelevant changes in a resume. For example, removing a photograph or changing formatting should not change a skill assessment. Researchers can also construct controlled resume variants to examine whether demographic or institution-related information unintentionally influences outputs. Such testing would complement general AI risk-management guidance. [21]

9.6 Real-World Dataset and Experimental Validation

The most important next step is a properly designed empirical evaluation. A future study can collect an anonymized dataset of resumes and job descriptions, obtain expert labels for skills, and compare keyword, TF–IDF, transformer, and hybrid methods. Statistical analysis can then determine whether the hybrid approach improves matching quality and whether users find the recommendations useful. The results should report limitations, confidence intervals where appropriate, and the composition of the evaluation dataset so that the findings are not overstated.

Figure 5: Potential directions for future development of AI-powered resume analysis systems.



10. Conclusion

This research proposes an AI-powered resume analyzer and skill-gap analysis system designed specifically as a career-development support tool. The central contribution is not simply the use of AI to read resumes, but the combination of skill extraction, normalization, semantic matching, evidence tracking, explainable gap classification, and learning recommendations.

The literature suggests that changing skill requirements make structured career guidance increasingly useful, while research on algorithmic hiring shows that recruitment AI must be designed with care. The proposed system therefore avoids treating a resume score as a final judgment. Instead, it explains the relationship between a target role and the evidence contained in a resume. This makes the approach more appropriate for students who want to understand what they already demonstrate, what they may need to learn, and how they can present their projects more clearly.

The next stage of this research should move from conceptual design to empirical testing. A real, anonymized dataset, human-labeled skill annotations, baseline comparisons, and user evaluation would allow the proposed architecture to be tested rigorously. With these additions, the project can develop from a college-level prototype into a stronger research study on explainable AI for career development.

Acknowledgement

The author acknowledges the guidance of the Department of DataScience, SIES COLLEGE OF ARTS, SCIENCE AND COMMERCE (EMPOWERED AUTONOMOUS), Sion(W), Mumbai, and the academic supervision associated with this research work.

Endnotes

[1] . World Economic Forum, *The Future of Jobs Report 2025*, Geneva, Switzerland, Jan. 2025. https://www.weforum.org/publications/the-future-of-jobs-report-2025/.

[2] . World Economic Forum, “3. Skills outlook,” in *The Future of Jobs Report 2025*, 2025. https://www.weforum.org/publications/the-future-of-jobs-report-2025/in-full/3-skills-outlook/.

[3] . J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186. doi: 10.18653/v1/N19-1423.

[4] . N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using Siamese BERT-networks,” in *Proceedings of EMNLP-IJCNLP*, 2019, pp. 3982–3992. doi: 10.18653/v1/D19-1410.

[5] . M. Zhang, K. Jensen, S. Sonniks, and B. Plank, “SkillSpan: Hard and soft skill extraction from English job postings,” in *Proceedings of NAACL-HLT*, 2022, pp. 4962–4984. doi: 10.18653/v1/2022.naacl-main.366.

[6] . M. Raghavan, S. Barocas, J. Kleinberg, and K. Levy, “Mitigating bias in algorithmic hiring: Evaluating claims and practices,” in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020, pp. 469–481. doi: 10.1145/3351095.3372828.

[7] . J. Berg and H. Johnston, “AI in human resource management: The limits of empiricism,” International Labour Organization, Working Paper 154, Nov. 2025. https://www.ilo.org/publications/ai-human-resource-management-limits-empiricism.

[8] . European Parliament and Council of the European Union, “Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act),” Official Journal of the European Union, 2024. https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32024R1689.

[9] . J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186. doi: 10.18653/v1/N19-1423.

[10] . N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using Siamese BERT-networks,” in *Proceedings of EMNLP-IJCNLP*, 2019, pp. 3982–3992. doi: 10.18653/v1/D19-1410.

[11] . M. Zhang, K. Jensen, S. Sonniks, and B. Plank, “SkillSpan: Hard and soft skill extraction from English job postings,” in *Proceedings of NAACL-HLT*, 2022, pp. 4962–4984. doi: 10.18653/v1/2022.naacl-main.366.

[12] . M. Raghavan, S. Barocas, J. Kleinberg, and K. Levy, “Mitigating bias in algorithmic hiring: Evaluating claims and practices,” in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020, pp. 469–481. doi: 10.1145/3351095.3372828.

[13] . J. Berg and H. Johnston, “AI in human resource management: The limits of empiricism,” International Labour Organization, Working Paper 154, Nov. 2025. https://www.ilo.org/publications/ai-human-resource-management-limits-empiricism.

[14] . E. Tabassi, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” National Institute of Standards and Technology, NIST AI 100-1, Jan. 2023. doi: 10.6028/NIST.AI.100-1.

[15] . European Parliament and Council of the European Union, “Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act),” Official Journal of the European Union, 2024. https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32024R1689.

[16] . E. Tabassi, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” National Institute of Standards and Technology, NIST AI 100-1, Jan. 2023. doi: 10.6028/NIST.AI.100-1.

[17] . E. Tabassi, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” National Institute of Standards and Technology, NIST AI 100-1, Jan. 2023. doi: 10.6028/NIST.AI.100-1.

[18] . M. Raghavan, S. Barocas, J. Kleinberg, and K. Levy, “Mitigating bias in algorithmic hiring: Evaluating claims and practices,” in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020, pp. 469–481. doi: 10.1145/3351095.3372828.

[19] . World Economic Forum, “3. Skills outlook,” in *The Future of Jobs Report 2025*, 2025. https://www.weforum.org/publications/the-future-of-jobs-report-2025/in-full/3-skills-outlook/.

[20] . M. Raghavan, S. Barocas, J. Kleinberg, and K. Levy, “Mitigating bias in algorithmic hiring: Evaluating claims and practices,” in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020, pp. 469–481. doi: 10.1145/3351095.3372828.

[21] . E. Tabassi, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” National Institute of Standards and Technology, NIST AI 100-1, Jan. 2023. doi: 10.6028/NIST.AI.100-1.

© 2026 The Author(s). Published in *Cordoba International Journal of Academic Research*.
This work is licensed under a **Creative Commons Attribution 4.0 International License (CC BY 4.0)**.
Published by Wexgrand Academic Press (a division of Wexgrand Private Limited), No. 4/461, Suite No. 1331, Valamkottil Towers, Thrikkakara, Ernakulam, Kerala – 682021, India.