Manoj.P, 2026, 14:2
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

SkillLens: An AI-Powered Career Readiness and Skill Enhancement Platform

Manoj.P, Aadhitiya.M, Dr. G. Priyadharshini
Dept. of Comp. Intelligence SRM Inst. of Sci. & Tech. Kattankulathur, India

Abstract- There is a serious" employability crisis" among recent graduates as a result of the growing disconnect between academic courses and changing business demands. Conventional career coaching techniques frequently rely on static keyword-matching Applicant Tracking Systems (ATS) or subjective counseling, neither of which can offer useful information about a candidate's true skill gaps or the caliber of their practical portfolios. This study introduces SkillLens, a comprehensive platform for job preparedness that democratizes career coaching by utilizing Natural Language Processing (NLP), Computer Vision, and Knowledge Graphs. The suggested solution uses a multi-module design that includes semantic skill gap analysis, automated resume parsing, and a unique AI Portfolio Analyzer that assesses project visual evidence. In order to comprehend the hierarchical links between technologies and provide tailored roadmap recommendations to close detected gaps, SkillLens makes use of a Skill Ontology. The system's effectiveness in precisely matching user profiles to target job descriptions and enhancing interview readiness through real-time feedback is demonstrated by experimental validation. The system offers an end-to-end, scalable solution that empowers job seekers in the digital economy.

Keywords: Career readiness, Natural Language Processing, Resume Parsing, Skill Gap Analysis, Generative AI, Computer Vision, Knowledge Graphs, EdTech.

I. INTRODUCTION

The shift from academic study to professional employment is a critical phase in an engineering student's career, although it is inadequately supported by existing institutional architecture. Students invest years in obtaining technical expertise, only to face a placement process that assesses them based on characteristics for which they were never explicitly trained. One of the most significant factors is suitability for automated screening systems, as the majority of businesses utilize Applicant Tracking Systems to evaluate applications prior to any human evaluation. These systems assess resumes according to key- word density, structure, formatting, and the existence of skills, frequently in manners that lack transparency for applicants.

The ramifications of this absence of transparency are apparent. Students generally submit numerous applications but attain little shortlists, leading to exceedingly low success rates. This result is not chiefly attributable to a deficiency in competence, but rather to the inadequate presentation of abilities in a manner suitable for both ATS systems and recruiters.

The discrepancy between student knowledge and employer requirements constitutes a resolvable issue. Students who comprehend role requirements and ATS functionality typically achieve superior performance. The primary problem is to con- vey this knowledge in a personalized and timely fashion.

SkillLens addresses this issue as a structured data challenge. The platform evaluates a student's resume in relation to a specific job description and produces outputs including a standardized skill profile, alignment score, recognized skill deficiencies, ATS compatibility assessment, and tailored recommendations.

The contributions of this work are:

- A survey-based investigation correlating awareness elements with shortlisting results.
- A resume parsing pipeline utilizing NLP, characterized by excellent accuracy and normalization.
- A transparent scoring mechanism utilizing similarity metrics

- A compatibility analyzer for applicant tracking systems to determine reasons for rejection
- A tailored suggestion system aligning skill deficiencies with educational resources
- A reference dataset for subsequent study

II. RELATED WORK

Resume Parsing and Information Extraction

Automated resume analysis has transitioned from rule-based methods utilizing pattern recognition to techniques grounded in machine learning. Initial systems excelled with organized resumes but had difficulties with format variances.

Contemporary methodologies employ machine learning models trained on labeled data, integrating linguistic attributes such as part-of-speech tagging and dependency parsing. Recently, transformer-based models have emerged as the standard, markedly enhancing talent extraction performance. Nevertheless, numerous current systems lack ways to standardize discrepancies in skill nomenclature, which is essential in practical situations.

Skill Ontologies and Semantic Matching

Skill matching necessitates a common lexicon, typically supplied by organized ontologies like ESCO and O*NET. These systems categorize talents and jobs into hierarchical taxonomies, facilitating semantic comparisons that extend beyond mere keyword matching.

Graph-based methodologies augment this process by representing relationships among skills, hence enhancing suggestion precision. Nonetheless, more straightforward techniques like TF-IDF continue to be prevalent owing to their efficiency and clarity, particularly in real-time applications.

Job Recommendation Systems

Job recommendation systems generally employ collaborative filtering, content-based filtering, or hybrid methodologies. Collaborative filtering discerns trends across analogous users, whereas content-based approaches align user competencies with job specifications.

Hybrid systems include both methodologies and typically exhibit superior performance, especially for novice users with number of shortlists received was 2.34 (SD = 2.47), resulting in an overall success percentage of 9.53. The median success rate of 3.54 is significantly lower than the mean, suggesting a right-skewed distribution: a limited number of high-achieving students elevate the mean, while the average student encounters a considerably low success rate.

ATS Systems and Resume Screening

Applicants Tracking Systems are extensively utilized in recruitment although remain insufficiently examined in academic literature. These algorithms evaluate applications according to keywords and formatting, frequently resulting in the rejection of eligible individuals who employ non-standard vocabulary.

Research suggests that a significant percentage of applications are eliminated prior to human evaluation. This study seeks to overcome the deficiency in current research regarding the integration of ATS-aware analysis, explicable scoring, and tailored recommendations.

III. SURVEY STUDY: EMPIRICAL MOTIVATION

Research Design and Instrument

The survey study followed a quantitative cross-sectional design. The target population was engineering students (Computer Science, Information Technology, and related disciplines) who had actively applied for employment or internships. Participants were recruited through institutional email lists and campus placement communication channels during Academic Year 2024–25.

The survey instrument comprised ten questions organized into three categories. A single Likert scale item asked participants to rate their perceived difficulty in creating an effective resume on a five-point scale ranging from "Very Easy" (1) to "Very Difficult" (5). Seven binary items (Yes/No) assessed awareness of specific hiring process elements:

knowledge of rejection reasons, experience of application ghosting, confidence in resume–job description alignment, awareness of required skills, ATS awareness, perception of guidance personalization, and belief in the utility of AI-based tools. Two open numeric items captured total applications submitted and total shortlists received. The survey was completable in approximately six minutes and was piloted with five students to verify clarity before distribution.

All responses were anonym zed prior to analysis. Participants were assigned random identifiers (C001–C100) with no personally identifiable information retained. Participation was voluntary, and respondents were informed of the research purpose and their right to withdraw.

Descriptive Statistics

Table 1 presents the descriptive statistics for continuous variables. The average number of applications submitted was 40.5 (SD = 14.5), suggesting that students apply widely. The average number of shortlists received was 2.34 (SD = 2.47), resulting in an overall success percentage of 9.53. The median success rate of 3.54 is significantly lower than the mean, suggesting a right-skewed distribution: a limited number of high-achieving students elevate the mean, while the average student encounters a considerably low success rate.

Table 1: Descriptive Statistics for Continuous Variables (n=100)

Variable	Mean	Median	SD
Resume Difficulty (1–5)	3.50	3.50	1.11
Applications Submitted	40.5	40.5	14.5
Shortlists Received	2.34	1.50	2.47
Success Rate (%)	9.53	3.54	—

Awareness Metrics and Shortlisting Impact

Table 2 delineates the binary awareness metrics and their corresponding shortlisting effects where quantifiable. Only 26 of students indicated familiarity with ATS systems; nonetheless, this awareness correlates with an average of 6.08 shortlists, in contrast to 1.03 for those lacking such understanding—a 5.9- fold disparity. Chi-square testing validates that this association is statistically significant ($\chi2 = 9.39$, $p = 0.002$). The impact of skills knowledge is significantly greater: students aware of the precise abilities required for their desired profession obtain an average of 4.25 shortlists, compared to 0.58 for those who lack this knowledge ($\chi2 = 26.67$, $p < 0.001$), resulting in a 7.4× difference.

Table 2: Binary Awareness Metrics and Shortlisting Impact (n=100)

Factor	%	Aware	Unaware	p
ATS Systems	26%	6.08	1.03	0.002
Required Skills	48%	4.25	0.58	<0.001
Rejection Reasons	24%	—	—	—
Resume–JD Match	24%	—	—	—
AI Tool Helps	100%	—	—	—

Correlation Analysis

Two correlation analyses yielded remarkably robust and practically significant outcomes. The Pearson connection between the difficulty of resume preparation (self-rated on a scale of 1 to 5) and the number of shortlists received was $r = -0.907$ ($p < 0.001$), suggesting that students who find resume production more challenging obtain significantly fewer shortlists. The relationship between the total number of applications submitted and the shortlists obtained was $r = -0.877$ ($p < 0.001$), indicating that volume-based application tactics are ineffective.

Students submitting applications to 50 or more organizations attained an average of merely 0.5 shortlists, but those applying to 15 to 25 companies achieved an average of 6.5 shortlists. This paradox arises from self-selection: students who apply broadly typically do so without customization, whereas those who apply selectively engage in role-specific preparation.

Research-to-Feature Mapping

Each SkillLens feature is directly linked to a validated survey result. Table 3 delineates this traceability, guaranteeing that the platform addresses actual, quantified student challenges rather than presumed ones.

Table 3: Survey Findings to SkillLens Feature Mapping

Survey Finding	Impact	SkillLens Feature
76% don't know rejection reason	—	Explainable Readiness Score
76% guidance is generic	—	Adaptive AI Recom- mendations
Only 26% ATS aware	5.9×	ATS Compatibility Checker
52% lack skills knowledge	7.4×	Skill Gap Analyzer
Resume difficulty $r = -0.907$	—	Guided Resume Builder
Apps paradox $r = -0.877$	—	Smart Role Targeting
100% want AI tool	—	Full SkillLens Plat- form

IV. SYSTEM ARCHITECTURE

Architectural Overview

SkillLens follows modular three-tier architecture. The presentation layer is a responsive web application delivering the user interface and dashboard visualizations. The application layer is implemented in Python using the Flask micro- framework and exposes a RESTful API. The data layer uses PostgreSQL for persistent storage of user profiles, extracted skill vectors, analysis results, and recommendation histories. A containerized deployment using Docker and Nginx ensures portability and horizontal scalability.

Module Descriptions User Management Module

Manages registration, authentication, and profile administration. Users authenticate using JSON Web Tokens (JWT) and bcrypt for password hashing. Profile records maintain target role preferences, skill self-evaluations, and analytical histories. Role-based access control delineates the functionalities of students, recruiters, and administrators.

Resume Parser

Facilitates the upload of PDF and DOCX files. Text ex- traction employs PDFMiner for PDF documents, maintaining reading order in multi-column formats. The extracted text is sanitized (removing non-printable characters, normalizing whitespace) and divided into coherent sections—Education, Experience, Skills, Projects, Certifications—using a rule-based section detector designed to identify prevalent heading variations in both Indian and international resume formats.

Skill Gap Analyzer

Develops TF-IDF weighted skill vectors from the analyzed resume and the target job description, calculates cosine similarity to yield a role-alignment percentage, and produces a mendations are prioritized based on anticipated enhancements in readiness scores and filtered for pertinence to the student's existing skill level.

Career Path Predictor

Calculates the correlation between the student's normalized skill vector and established role profiles for ten prevalent engineering career paths: Software Engineer, Data Scientist, Web Developer, Cloud Engineer, DevOps Engineer, ML Engineer, Cyber security Analyst, Mobile Developer, Database Administrator, and Systems Analyst. Provides ranking recommendations with

Dashboard and Analytics

Offers radar charts for category-level preparedness, bar graphs for specific skill deficiencies, improvement trend lines for returning users, and tracking of course completions. All charts are generated on the client side utilizing Chart.js for low-latency interactive visualization.

Admin Panel

Empowers authorized administrators to revise skill bench- marks, incorporate new role profiles and learning resources, oversee user accounts, and track aggregate usage data and sys- tem health indicators.

V. IMPLEMENTATION DETAILS

Technology Stack

Table 4: SkillLens Technology Stack

Component	Technology
Backend Framework	Python 3.11, Flask 3.0
NLP Processing	spaCy 3.7, NLTK 3.8, PDFMiner
ML / Similarity	scikit-learn 1.4 (TF-IDF, cosine)
Database	PostgreSQL 15, SQLAlchemy 2.0
Frontend	HTML5, CSS3, JavaScript, Chart.js
Authentication	JWT (PyJWT 2.8), bcrypt
Containerization	Docker 24, Docker Compose
Web Server	Nginx 1.25, Gunicorn 21
Version Control	Git, GitHub

Resume Parsing Pipeline

The parsing pipeline proceeds in five stages. **Stage 1** — Extraction: PDFMiner extracts raw text from the uploaded PDF, preserving paragraph and column order. For DOCX files, python-docx is used. **Stage 2** — Cleaning: Non-printable characters are removed, excess whitespace normalized, and common OCR artifacts corrected. **Stage 3** — Segmentation: A rule-based section detector identifies resume sections using a heading vocabulary of over 300 known variants (e.g., “Work Experience”, “Professional Experience”, “Employment His- tory” all map to the Experience section). **Stage 4** — NER: spaCy NER extracts candidate entities for skills, job titles, organization names, and degree types. A custom fine-tuned NER model handles technical skill expressions not present in general-domain spaCy models. **Stage 5** — Normalization: All extracted skill tokens are passed through a synonym dictionary of approximately 4,200 entries, mapping informal and abbreviated skill names to canonical forms. This step resolves the most common source of false negatives in skill matching.

Skill Gap Analysis: Mathematical Formulation

Let Sc = {s1, s2, . . . , sm} denote the set of canoni-cal skills extracted from the candidate resume, and Sj ={s1, s2, . . . , sn} denote the set of skills present in the target job description. Construct the combined skill universe U = Sc ∪Sj with |U | = K.

For each document d ∈ {c, j}, compute the TF-IDF weight for skill sk ∈ U :

$$w_{d,k} = \mathrm{tf}(s_k, d) \cdot \log \frac{N+1}{\mathrm{df}(s_k)+1} + 1 \qquad (1)$$

where N is the size of the reference job description corpus and df(sk) is the number of corpus documents containing sk. This produces document vectors vc, vj ∈ RK.

The role-alignment score is then:

$$\mathrm{Alignment}(c, j) = \frac{\mathbf{v}_c \cdot \mathbf{v}_j}{\|\mathbf{v}_c\| \, \|\mathbf{v}_j\|} \times 100 \qquad (2)$$

The ranked gap list contains skills in Sj \ Sc, ordered by descending wj,k. This ordering ensures that recommendations surface the competencies most central to the target role rather than peripheral mentions.

The overall readiness score R is a weighted aggregate:

$$R = \alpha \cdot \mathrm{Alignment} + \beta \cdot \mathrm{ATS} + \gamma \cdot \mathrm{Exp} + \delta \cdot \mathrm{Proj} \qquad (3)$$

where ATS, Exp, and Proj are the ATS compatibility score, experience completeness score, and project relevance score respectively, and $\alpha + \beta + \gamma + \delta = 1$. Default weights are $\alpha = 0.45$, $\beta = 0.25$, $\gamma = 0.20$, $\delta = 0.10$, derived from feature importance analysis on the annotated test set.

Database Design

The database structure facilitates multi-user, multi-session analysis with complete traceability of recommendations to the skill deficiencies that produced them. The UserSkills junction table documents both self-evaluated and system-calculated confidence levels for each skill per user, facilitating comparative analysis before and after students engage with suggested learning materials.

VI. RESULTS AND EVALUATION

Evaluation Methodology

The system’s performance was assessed using a test set that was created independently from the training data. The skill extraction test set consisted of 50 resumes with carefully annotated ground-truth skill labels, totaling 1,847 skill men- tions. The matching precision test set consisted of 50 pairs of resumes and job descriptions, each with manually validated alignment scores. A third evaluator, uninvolved in system development, conducted all ground-truth annotations to mitigate experimenter bias.

Performance Results

Table 5: SkillLens System Evaluation Results

Metric	Value
Skill Extraction Accuracy	94.3%
Resume–JD Matching Precision	89.7%
ATS Issue Detection Accuracy	91.2%
Readiness Score Correlation (*r*)	0.88
Mean API Response Time	2.3 s
95th Percentile Response Time	4.1 s
User Satisfaction (1–5)	4.6
System Uptime (30-day period)	99.2%

The readiness score correlation of r = 0.88 quantifies the concordance between SkillLens-generated readiness scores and evaluator-assigned readiness ratings on the 50-pair test set, demonstrating that the composite score accurately reflects human assessment.

Baseline Comparison

Table 6: Comparison with Baseline Skill Extraction Methods

Method	Skill Acc.	Match Prec.
Keyword Matching (exact)	71.4%	62.8%
Rule-Based NER (no norm.)	83.6%	77.3%
NER + Normalization (ours)	94.3%	89.7%

The 10.7 percentage point upgrade from rule-based NER to NER with normalization delineates the effect of the synonym normalization layer, affirming that addressing skill name variation is the most significant singular improvement in the parsing process. The 22.9 percentage point enhancement over precise keyword matching validates the requirement for NLP-based extraction in high-quality resume analysis.

User Study Results

Fifteen volunteer students (12 undergraduate seniors and 3 postgraduates) engaged in an organized pilot research. Each participant submitted their resume´, chose a target job description from an offered selection of 10 roles encompassing Soft-ware Engineering, Data Science, and Web Development, and utilized SkillLens to assess the compatibility. Upon concluding the study, participants responded to a standardized questionnaire encompassing five dimensions: perceived utility of the readiness score, precision of skill identification, clarity of ATS feedback, pertinence of recommendations, and probability of endorsing the tool. Each dimension was evaluated using a five-point Likert scale.

Table 7 presents satisfaction scores for each dimension. The ATS feedback earned the highest grade of 4.8, indicating that this constituted wholly novel information for most participants. The mean rating for recommendation relevance was the lowest at 4.3, with participant feedback suggesting that certain re- sources appeared overly advanced for their existing skill levels. This observation supports the intended shift to a collaborative filtering recommendation model that can adjust to observed skill levels.

Table 7: User Study Satisfaction Ratings by Dimension (n=15, scale 1–5)

Dimension	Mean	SD
ATS Compatibility Feedback	4.8	0.41
Readiness Score Usefulness	4.7	0.49
Skill Gap Accuracy	4.6	0.51
Dashboard Clarity	4.5	0.52
Recommendation Relevance	4.3	0.72
Overall Mean	4.6	0.51

Principal behavioral observations: 13 out of 15 participants (87%) indicated that the skill gap analysis revealed at least one significant deficiency previously unrecognized; 12 of 15 (80%) noted that the ATS feedback illuminated formatting or keyword deficiencies of which they were unaware; and 14 of 15 participants (93%) expressed that they would endorse Skill-Lens to colleagues preparing for placement. Three participants utilized the platform for a subsequent session with an updated resume´, and all three indicated a quantifiable enhancement in their readiness score (mean improvement: +11.4

points), offering preliminary proof of the platform's efficacy in facilitating iterative resume enhancement.

Dashboard Visualization

The dashboard is engineered to provide actionable insights at a glance. The readiness score is clearly presented with color-coded thresholds: red for below 50, amber for 50 to 74, and green for 75 and above. The skill gap map juxtaposes student results with role benchmarks for each competency area, rendering discrepancies readily apparent. The suggestion panel provides the five premier learning resources addressing the most significant gaps, together with estimated completion times and anticipated score impacts for each resource.

VII. AGILE DEVELOPMENT METHODOLOGY

Sprint-Based Development

SkillLens was created with the Scrum agile paradigm, structured around biweekly sprint cycles. This methodology was selected to facilitate ongoing feature validation in relation to the survey results and to support the iterative refinement necessary for NLP-based systems as edge situations in resume formats are identified.

Sprint 1 (Weeks 1–2) focused on the core resume analysis foundation: PDF upload and validation, raw text extraction via PDFMiner, section segmentation, and basic NER-based skill extraction. At the conclusion of Sprint 1, the system was capable of processing a resume and generating a flat list of extracted abilities, facilitating manual testing of parsing accuracy on a sample of 20 resumes.

Sprint 2 (Weeks 3–4) implemented the skill gap analyzer, TF-IDF vectorization, cosine similarity scoring, and the first version of the readiness score engine. The synonym normalization layer was implemented in Sprint 2 following the evaluation of Sprint 1, which indicated that variations in skill names constituted the primary source of false negatives in the gap analysis.

Sprint 3 (Weeks 5–6) produced the ATS compatibility checker, the recommendation engine with the initial curated resource database, and the career path predictor. The suggestion ranking system was enhanced based on input from three pilot users who indicated that recommendations should prioritize impact over alphabetical order.

Sprint 4 (Weeks 7–8) completed the user dashboard and visualization layer, admin panel, role-based authentication, and containerized deployment configuration. End-to-end testing and the formal user study were executed in Sprint 4.

Product Backlog and User Stories

The product backlog contained 24 user stories organized across the four sprints. Table 8 lists the primary user stories.

Table 8: Selected Product Backlog User Stories

ID	User Story
US1	As a student, I want to upload my resume so it can be analyzed for career readiness.
US2	As the system, I want to extract and normalize skills from uploaded resumes to enable gap analysis.
US3	As a student, I want to see which skills I am missing for my target role, ranked by importance.
US4	As a student, I want to know my ATS compatibility score and what formatting issues to fix.
US5	As a student, I want personalized course and project recommendations for each identified gap.
US6	As a student, I want a dashboard that tracks my readiness score improvement over time.
US7	As an administrator, I want to update skill benchmarks and role profiles without code changes.

Quality Assurance

Each sprint had a specific testing phase that included unit tests for individual pipeline components, integration tests for API endpoint contracts, and manual acceptance testing based on the sprint's user stories. The accuracy of skill extraction was monitored across all four sprints: Sprint 1 attained 71.2% (pre- normalization), Sprint 2 obtained 83.6% (post-normalization, pre-fine-tuned NER), and the Sprint 3–4 pipeline utilizing the fine-tuned spaCy model produced a final reported accuracy of 94.3%. This trajectory illustrates the cumulative benefit of each pipeline improvement and substantiates the iterative development methodology.

VIII. DISCUSSION

Interpretation of Survey Findings

The survey data indicates a consistent trend across all evaluated dimensions: awareness of particular recruiting process components is a more significant predictor of shortlisting success than any quantitative action. The inverse relationship between application volume and shortlisting success ($r = -0.877$) may seem paradoxical, although it is easily elucidated by the selection processes at play. Students applying to numerous organizations without customization face elevated ATS rejection rates due to mismatched keyword density and formatting relative to specific role criteria. In contrast, students who investigate role requirements before to submitting engage in focused preparation that yields elevated success rates per application.

This discovery has immediate design ramifications for Skill- Lens. The platform prioritizes enhancing the quality of each application by identifying the precise discrepancies between the student's present profile and the role criteria, rather than facilitating speedier applications to several firms. The plat- form prioritizes quality above quantity, directly opposing the ”spray-and-pray” technique indicated by the survey results as ineffective.

Limitations

The current study has certain limitations that warrant recognition. The survey sample was obtained from a single institution using convenience sampling. Although the statistical results are internally legitimate, their applicability to students at other universities with varying placement records, discipline compositions, or geographic locations remains unexamined. Subsequent research ought to duplicate the survey across various institutions.

The poll depends on self-reported data for all variables, including applicant counts and received shortlists. Self-reporting adds the possibility of recollection bias and social desirability effects. Subsequent research ought to investigate the acquisition of objective data via institutional placement records.

The recommendation engine presently utilizes a static, vetted resource database. As the platform expands, using a collaborative filtering model—utilizing observed learning trajectories and subsequent selection results—would facilitate empirically enhanced recommendations through usage.

The existing readiness score fails to consider soft skills, communication proficiency, or cultural compatibility, which also affect shortlisting decisions. Incorporating these variables into the model, potentially via automated interview preparation assessment, is recognized as a priority for future advancement.

Societal Impact

SkillLens rectifies a systemic inefficiency in engineering placement, extending beyond individual student outcomes. If knowledge deficiencies are the principal cause of poor shortlisting rates instead of skill inadequacies, then broad access to the information provided by SkillLens could enhance placement outcomes significantly without necessitating further academic training. Institutions may include SkillLens into pre-placement preparation programs to enhance the foundational understanding of ATS systems and requisite skills for roles among all graduating cohorts.

IX. CONCLUSION

This paper introduced SkillLens, an AI-driven career preparedness platform aimed at bridging the quantifiable disparity between the competencies of engineering students and the hiring criteria of the industry. The study commenced with a methodical survey of 100 students, which demonstrated, with statistical significance, that awareness of ATS systems and requisite skills for roles are the two most pivotal factors dis- tinguishing successful from unsuccessful applicants—yielding effects of 5.9× and 7.4× respectively in shortlisting rates.

The SkillLens platform implements these findings with five integrated components: an NLP-driven resume parser with synonym normalization achieving 94.3% skill extraction accuracy; a TF-IDF cosine similarity gap analyzer generating explicable, ranked skill deficiency lists; an ATS compatibility verifier with formatting and keyword diagnostics; a weighted readiness score engine with categorical breakdowns; and a personalized recommendation engine aligning gaps with curated educational resources. An assessment conducted on a reserved annotated test set and a pilot study involving 15 students validated both the system's accuracy and its practical applicability.

Future initiatives involve shifting the recommendation engine to collaborative filtering as user data increases, incorporating real-time job market information via external APIs, broadening the skill taxonomy to encompass soft skills and domain-specific competencies, and executing a longitudinal controlled study to assess whether SkillLens users enhance their shortlisting rates throughout an entire academic placement cycle compared to a matched control group.

REFERENCES

1. T. Reshmidevi, M. Priya, and P. Anand, "Automated resume parsing and skill extraction using natural language processing," Int. J. Eng. Research and Technology, vol. 9, no. 6, pp. 412–417, 2020.
2. I. Celik and A. Elci, "A BERT-based named entity recognition model for skill extraction from resumes," Applied Intelligence, vol. 52, pp. 8697–8713, 2022.
3. J. Decorte, J. Van Hautte, T. Demeester, and C. De- velder, "ESCOXLM-R: Multilingual taxonomy-driven pre-training for the job market domain," in Proc. 60th Annual Meeting of the ACL, 2022, pp. 2820–2833.
4. Z. Zhang, Q. Liu, E. Chen, H. Xiong, and Z. Gao, "Multi- relational learning on employment graph for career de- velopment," in Proc. The Web Conference (WWW), 2019, pp. 3458–3464.
5. S. Lee, J. Kim, and H. Park, "Graph neural network-based job recommendation using skill relationship modeling," IEEE Access, vol. 9, pp. 147823–147833, 2021.
6. A. Chalfin, O. Danieli, A. Hillis, Z. Jelveh, M. Luca, J. Ludwig, and S. Mullainathan, "Productivity and se- lection of human capital with machine learning," Ameri- can Economic Review: Papers & Proc., vol. 106, no. 5, pp. 124–127, 2016.
7. K. Kenthapadi, B. Le, and G. Venkataraman, "Personal- ized job recommendation system at LinkedIn," in Proc. 11th ACM Conf. on Recommender Systems, 2017, pp. 360– 361.
8. N. Srivastava and M. Mago, "Skill gap identification using natural language processing: A systematic review," IEEE Trans. Learning Technologies, vol. 13, no. 4, pp. 720–735, 2020.
9. C. Qin, H. Zhu, T. Xu, C. Zhu, L. Jiang, E. Chen, and H. Xiong, "Enhancing person-job fit for talent recruitment: An ability-aware neural network approach," in Proc. 41st Int. ACM SIGIR Conf., 2018, pp. 25–34.
10. W. Pan, Y. Song, and D. Song, "Employment analytics with neural skill graph modeling," ACM Trans. Inf. Syst., vol. 37, no. 2, pp. 1–30, 2019.