

# AI-Powered Career Guidance System: An Intelligent Framework for Personalized Career Recommendation and Skill Pathway Generation

**Sumit Choudhari, Yuvraj Johare, Vaibhav Devkar, Yashraj Garje**  
Department of Information Technology  
PDEA's College of Engineering, Manjri (Bk), Pune

**Guide: Prof. Deepali Herolikar**  
Department of Information Technology  
PDEA's College of Engineering,

Pune

Savitribai Phule Pune University  
Student email: [srajgarje017@gmail.com](mailto:srajgarje017@gmail.com)  
[Yuvrajjohare001@gmail.com](mailto:Yuvrajjohare001@gmail.com)  
[Mrvaibhavdevkar2015@gmail.com](mailto:Mrvaibhavdevkar2015@gmail.com)  
[chaudharisumit541@gmail.com](mailto:chaudharisumit541@gmail.com)

*Abstract*—The rapid expansion of academic specializations and career opportunities has made decision-making increasingly difficult for students and early professionals. Traditional career counselling methods—based on subjective analysis, limited psychometric tests, or one-time assessments—often fail to deliver personalized, data-driven, and objective recommendations. The need for an intelligent, scalable, and adaptive guidance system has therefore become essential in modern educational ecosystems.

This paper presents an AI-Powered Career Guidance System that leverages machine learning, natural language processing (NLP), and large language models (LLMs) to provide end-to-end personalized career recommendations. The system evaluates user skills, interests, and behavioral patterns through assessments, semantic analysis, and structured profiling. Based on these inputs, it generates suitable career-role mappings, learning pathways, project suggestions, and competency insights. Additional modules—such as resume generation and AI-driven mock interview evaluation—enhance employability preparation. Using deep learning models, transformer-based text understanding, and domain-skill similarity scoring, the system offers high contextual accuracy and adaptability.

The proposed system aims to democratize career guidance by offering low-cost, real-time, and personalized recommendations accessible to all learners, regardless of geographical or economic constraints. Performance evaluation conducted on 60 students demonstrates 84.5% recommendation accuracy, 86% roadmap relevance, and a user satisfaction rating of 9/10, showcasing the potential of AI-driven systems in transforming educational and professional counselling.

*Index Terms*—Career Guidance, Artificial Intelligence, Machine Learning, Recommendation System, NLP, Deep Learning, LLM, Mock Interview.

## I. INTRODUCTION

Choosing an appropriate career path is one of the most influential decisions in a student's academic and professional life. However, students frequently experience confusion due to insufficient awareness of job roles, lack of skilled counselling, rapidly shifting industry demands, and inconsistent interpretation of aptitude assessments. Traditional guidance methods—such as manual counselling, aptitude tests, or peer recommendations—often yield generalized advice that does not reflect individual preferences, skill levels, or emerging market needs. This misalignment between student capabilities

and career recommendations leads to poor outcomes including job dissatisfaction, high attrition rates, and skill-role mismatches that are costly for both individuals and organizations.

With advancements in artificial intelligence, machine learning algorithms, and natural language processing, it is now

possible to build intelligent systems that analyze user behavior, interpret strengths, and recommend personalized career paths. AI-based systems overcome the inconsistencies of human-led counselling by offering scalable, unbiased, and customized solutions. The system proposed in this study is designed to support students, job seekers, and professionals by providing personalized insights, structured learning pathways, and interactive tools such as resume building and mock interviews.

The system integrates multiple AI components, including skill classification models, semantic similarity computation, and transformer-based reasoning engines, to offer accurate and meaningful recommendations. By combining structured assessments with free-text interpretation, the model achieves a holistic understanding of the user's capabilities and aspirations. Unlike existing fragmented tools that address only one aspect of career development, the proposed system covers the complete journey—from initial self-assessment and role discovery to skill building, resume preparation, and interview readiness.

The remainder of this paper is organized as follows. Section II discusses the motivation behind the system. Section III outlines the research objectives. Section IV surveys related literature and identifies gaps. Section V describes the proposed system design. Section VI presents the system architecture. Section VII details application features. Section VIII describes individual modules. Section IX presents the experimental setup and results. Sections X and XI provide discussion and future scope, and Section XII concludes the paper.

## II. MOTIVATION

Despite the increasing availability of educational resources, the process of identifying a suitable career path remains difficult for many learners. Most students lack access to trained career counsellors, while others rely on incomplete information from peers or social media. Existing guidance tools are limited by static questionnaires, generic outputs, and an absence of adaptability to user personality, progress, or evolving job markets. The following key factors motivated this research.

- **Lack of Personalized Guidance:** Most assessment tools provide broad suggestions that do not consider individual interests, aptitude levels, or evolving goals. A student with strong analytical skills but weak communication receives the same broad engineering recommendation as a peer

with an opposite profile, despite needing very different career paths.

- **Rapidly Changing Job Market:** Emerging fields such as AI, cybersecurity, cloud computing, and digital design require updated guidance aligned with industry needs. Roles such as MLOps Engineer, Prompt Engineer, and Cloud Security Architect did not exist in traditional career frameworks yet are among the fastest-growing positions today.
- **Inefficiency of Traditional Counselling:** Manual interventions are often time-consuming, subjective, and unavailable to all students. In large institutions, a single counsellor may serve hundreds of students, making deep personalization practically impossible.
- **Need for Scalable, Low-Cost Career Support:** AI-based guidance systems can serve thousands of students simultaneously without resource limitations, making high-quality career support economically viable for institutions of all sizes.
- **Demand for End-to-End Solutions:** Students require not only recommendations but also roadmaps, mock interviews, and resume support. Fragmented tools that address only one stage of career development force students to navigate multiple disconnected platforms, creating friction and reducing adoption.

These challenges collectively motivated the development of an AI-driven system capable of understanding the user holistically and delivering a personalized, structured, and adaptive career development experience within a single unified platform.

### III. OBJECTIVES

The major objectives of the AI-Powered Career Guidance System are:

- To design a machine learning-based assessment model capable of evaluating user skills, aptitude, and interests through structured quizzes, aptitude tests, and free-text response analysis.
- To develop an AI-driven recommendation engine that maps learner profiles to suitable career roles using classification algorithms and semantic similarity scoring.
- To generate personalized, step-by-step learning pathways including curated courses, industry certifications, hands-on projects, and measurable skill milestones.
- To implement an NLP-driven resume builder that extracts structured profile data, scores resume quality against target job descriptions, and generates ATS-optimized documents.
- To incorporate an LLM-powered mock interview module that simulates HR and technical interview sessions and delivers structured, rubric-based performance feedback.
- To provide a scalable, cost-effective, and device-agnostic platform that is accessible to students regardless of geographical location or economic background.
- To track user progress longitudinally through a career readiness dashboard that reflects skill improvements, roadmap completion, and interview performance over time.

### IV. LITERATURE SURVEY

Career recommendation systems have evolved significantly with advancements in artificial intelligence, data analytics, and psychological assessment frameworks. Early studies focused

on psychometric evaluations, aptitude scoring, and rule-based systems, whereas contemporary research emphasizes machine learning models, NLP engines, and transformer-based architectures. This section reviews major contributions in the domain, highlighting methodologies, limitations, and gaps addressed by the proposed system.

#### A. Early Research on Career Guidance Systems

Initial attempts at digital career counselling relied heavily on static aptitude tests and structured questionnaires. Traditional systems followed fixed scoring models that mapped personality traits or aptitude levels to career domains using frameworks such as Holland's RIASEC model [1]. While these systems provided basic guidance, they lacked adaptability, contextual awareness, and user-specific personalization. Studies also indicated low accuracy due to dependence on self-reported data and absence of behavioral interpretation. The deterministic nature of such systems made them incapable of handling the diversity of modern career landscapes.

#### B. Machine Learning-Based Recommendation Models

Machine learning introduced improved predictive capability for career recommendations. Researchers used Logistic Regression, Decision Trees, Random Forest Classifiers [5], and Naive Bayes Models trained on student datasets combining academic performance, domain preferences, and skill-level assessments. Karn and Shrestha [6] demonstrated gradient boosted models for job role transition prediction with reasonable accuracy on LinkedIn profile data. Although machine learning approaches improved classification accuracy, they struggled with complex interpretations involving multiskill evaluation, dynamic markets, and user personality traits. Moreover, such models required large labeled datasets and could not generate learning roadmaps or interpret free-text responses.

#### C. Deep Learning and Embedding-Based Models

Deep learning models enabled better representation of career-role semantics. Techniques such as Word2Vec [2], GloVe, BERT [3], RoBERTa, and Graph Neural Networks [8] were adopted to analyze job descriptions, resume content, and skill taxonomies. Decorte et al. [7] demonstrated effective skill extraction using DistilBERT on the ESCO ontology, achieving strong performance on held-out skill mentions. Transformer-based architectures enhanced contextual understanding in mapping user narratives to job skills. However, deep learning approaches required significant computational power and did not inherently incorporate personalized user profiling or generate actionable guidance.

#### D. LLM-Based Career Counselling Systems

With the rise of LLMs such as GPT [4], Gemini, and LLaMA, researchers explored conversational career guidance. Zhu et al. [9] proposed CareerGPT, demonstrating that LLMs can understand nuanced career queries and generate personalized suggestions. LLMs demonstrated high accuracy in understanding user queries, generating career insights, identifying skill gaps, and providing project suggestions. Nevertheless, standalone LLMs lacked structured decision mapping, consistency across multiple sessions, and integration with userspecific assessments. Ji et al. [11] highlighted the hallucination problem in LLMs, which is particularly problematic in career guidance where factual accuracy about

courses, salaries, and job roles is critical. Lewis et al. [10] proposed RetrievalAugmented Generation (RAG) as a solution to ground LLM outputs in verified knowledge bases.

### E. Integrated Career Platforms

Modern platforms such as Coursera Career Tools [17] and LinkedIn Learning [16] introduced interest-based course suggestions using user profiles and interaction data. Ricci et al. [12] provide a comprehensive treatment of recommender system methods applicable to career platforms. However, these systems prioritize course discovery rather than holistic career guidance and do not include personalized learning pathways or mock-interview evaluations. The absence of assessment-driven recommendation, competency gap analysis, and interview simulation limits their utility as comprehensive career development tools.

### F. Research Gap

From the literature, the following gaps were identified: lack of unified systems combining assessment, recommendation, resume building, and interview simulation; absence of personalized AI-generated learning roadmaps tied to specific career roles; inability to interpret free-text user responses using LLM-powered semantic models; limited behavioral analytics integration; and lack of affordability for students in underserved regions. The proposed system addresses all these limitations by combining ML-based assessment, NLP-driven recommendation, LLM-powered interview coaching, and adaptive roadmap generation into a single cohesive platform.

## V. PROPOSED SYSTEM

The proposed AI-Powered Career Guidance System provides an end-to-end solution for personalized career mentoring. It integrates assessment engines, recommendation algorithms, NLP analysis, resume generation tools, and mock interview evaluation into a unified platform. The system is designed for three primary user groups: final-year students exploring career options, recent graduates entering the job market, and working professionals seeking to transition to new roles.

### A. System Overview

The system captures user inputs through five channels: structured skill assessment quizzes spanning cognitive aptitude and domain knowledge, interest surveys aligned to career clusters, free-text responses to open-ended prompts such as “Describe yourself?” and “What motivates you professionally?”, academic history including grades and completed courses, and an optionally uploaded existing resume. The backend processes these inputs through three AI layers: an ML classification engine for skill clustering, an NLP semantic analysis module for free-text understanding, and an LLM reasoning layer for context-aware generation. Based on this processing, the system recommends suitable career roles, identifies required skill sets, generates personalized learning pathways, proposes relevant projects, suggests applicable certifications, and provides job-market demand insights.

### B. Functional Workflow

The end-to-end workflow proceeds in seven sequential stages. First, the user registers and completes the multisection assessment battery. Second, the ML model identifies dominant skill clusters from quiz responses and maps them to a multi-dimensional competency vector. Third, the NLP engine

semantically processes free-text inputs to enrich the user profile with personality and interest signals. Fourth, the recommendation engine combines both signals to produce a ranked list of suitable career roles. Fifth, once the user selects a target role, the system generates a week-by-week learning roadmap tailored to the identified skill gaps. Sixth, the resume builder constructs an ATS-optimized CV from the user’s profile data. Seventh, the mock interview module conducts a simulated session and delivers a comprehensive performance evaluation report.

## VI. SYSTEM ARCHITECTURE

The system architecture follows a layered, microservices-based design that enables independent scaling of individual components and clean separation of concerns. The frontend is built with React and Next.js, delivering a responsive singlepage application that encompasses the assessment interface, the interactive career dashboard, the resume builder preview, and the real-time conversational mock interview interface. The API gateway, implemented in Node.js/Express, handles all request routing, session management, and communication between the frontend and backend services.

The ML inference layer is responsible for skill classification and semantic career-role matching. It hosts Python-based scikit-learn and TensorFlow models served as REST microservices, enabling horizontal scaling during peak usage. The LLM reasoning layer powers both the free-text interpretation pipeline and the mock interview evaluation engine, leveraging GPT-based models with carefully engineered prompts to ensure consistent, rubric-aligned outputs.

The data layer uses MongoDB Atlas as the primary persistence store for user profiles, assessment histories, roadmap progress logs, resume versions, and interview session records. The non-relational document model accommodates the variability in user profile structures without requiring schema migrations. The system is hosted with the frontend on Vercel for global CDN distribution and backend services on AWS for compute elasticity. The complete architecture is illustrated in Fig. 1.

## VII. APPLICATION FEATURES

The AI-Powered Career Guidance System integrates a wide range of intelligent features to enhance user experience, improve accuracy, and provide end-to-end career assistance.

### A. Personalized Skill Assessment

The system conducts a multi-dimensional assessment using cognitive and aptitude-based questions, domain preference metrics, analytical reasoning tests, and behavioral and interest profiling. Questions are drawn from a curated bank covering twelve technology and professional domains. Machine learning models analyze responses to identify dominant skill clusters, enabling accurate recommendation mapping. The assessment also includes free-text prompts whose responses are semantically analyzed to capture personality signals not expressible through multiple-choice questions alone.

### B. AI-Driven Career Recommendation

Recommendations are generated using a two-stage pipeline. In the first stage, a trained classification model maps the user’s skill vector to a probability distribution over career roles. In the second stage, semantic similarity scoring between the user’s free-text embeddings and role description embeddings re-ranks

the top candidates. The final ranked list presents each career option with a role description, required and optional skills, industry growth trajectory, and indicative salary range—giving users the contextual information needed to make an informed selection.

### C. Dynamic Learning Path Generation

Upon career role selection, the system computes the gap between the user's assessed competency profile and the target skill requirements of the chosen role. It then queries a curated database of online courses, professional certifications, and hands-on project ideas to construct a structured, week-by-week learning plan. The roadmap prioritizes skills by their importance to the target role and estimated learning time, ensuring users tackle high-impact gaps first. Progress through the roadmap is tracked on the dashboard and updated in realtime as the user marks milestones complete.

### D. Resume Builder Module

The resume builder automatically extracts structured information from the user profile—including education, technical skills, completed projects, certifications, and career objective—and assembles a professionally formatted, singlepage resume document. An integrated NLP scoring module evaluates the draft against target job descriptions using keyword density analysis and semantic alignment metrics. The module then generates specific, actionable suggestions to improve ATS compatibility, such as adding missing keywords, quantifying achievements, and restructuring experience descriptions for greater impact.

### E. Mock Interview Simulation

Using LLM-based evaluation, the system simulates both HR and technical interview sessions through a real-time conversational interface. Questions are dynamically selected from a role-specific bank of over several hundred items, tagged by difficulty and competency area. Each user response is evaluated by the LLM against a structured rubric covering relevance, depth of knowledge, response structure, grammatical correctness, and communication clarity. The post-session report provides per-dimension scores, highlights specific strengths and weaknesses, presents exemplar answers for comparison, and offers targeted improvement suggestions for the next practice session.

### F. Dashboard and Progress Tracking

The user dashboard provides a longitudinal view of career development progress. It displays the overall career readiness score as a composite metric, individual skill proficiency levels updated after each assessment attempt, percentage completion of the active learning roadmap, and a history of mock interview performance scores over time. Visual progress indicators and milestone notifications encourage continued engagement and give users a clear sense of momentum toward their career goal.

## VIII. MODULES OF THE SYSTEM

### A. Assessment Module

The assessment module administers the full battery of quizzes, aptitude tests, and interest inventories. It applies itemlevel scoring to produce a multi-dimensional skill vector encoding the user's competency levels across technical

domains (such as programming, data analysis, networking, and design) and behavioral dimensions (such as problem-solving style, communication preference, and work environment fit). This vector is persisted to the user profile and updated with each subsequent assessment session, enabling longitudinal tracking of skill development.

### B. Career Recommendation Module

The recommendation module ingests the skill vector and applies a trained Random Forest classifier combined with BERT-based semantic similarity scoring to generate a ranked list of career roles. The classifier provides coarse-grained role probabilities while the semantic layer refines rankings based on the alignment between the user's free-text language and role description embeddings. Each recommended role is enriched with structured metadata—required technical skills, optional complementary skills, market demand scores, average salary benchmarks, and job-growth projections—retrieved from a regularly updated knowledge base.

### C. Roadmap Generator Module

The roadmap generator computes a skill gap vector by subtracting the user's current competency profile from the target competency profile of the selected role. Gaps are prioritized by a weighted scoring function that accounts for each skill's importance to the role and the user's estimated time to proficiency. The module then queries a database of curated courses, certifications, and project templates to assemble a sequenced, time-bound learning plan. Roadmap content is presented as an interactive week-by-week timeline with embedded resource links and estimated completion times.

### D. Resume Builder Module

The resume builder aggregates structured data from the user profile and renders it into a professionally formatted document template. The integrated NLP scoring engine computes a relevance score between the resume content and a target job description by measuring keyword overlap and semantic alignment. The module produces a scored report identifying content gaps, overused phrases, and missing quantifications, then generates specific rewriting suggestions. Final resume documents can be exported in PDF format for direct submission to job applications.

### E. Mock Interview Module

The mock interview module manages the complete simulation workflow. It selects questions from a role-specific bank, streams them to the user through the conversational interface, collects typed or spoken responses, and submits each response to the LLM evaluation pipeline. The evaluation prompt is carefully engineered with a structured rubric and few-shot exemplars to ensure consistent scoring. After the session, the module compiles a comprehensive feedback report with overall and per-dimension scores, response-by-response commentary, model answer comparisons, and a prioritized list of improvement areas for the next practice session.

### F. User Profile and Dashboard Module

This module maintains a persistent, versioned record of all user activity including assessment histories, selected career roles, roadmap completion status, resume drafts, and interview performance scores. The dashboard aggregates these data points into a unified career readiness timeline, enabling users to observe their growth trajectory over weeks and months. Gamification elements such as milestone badges and readiness

score thresholds provide motivational anchors that encourage sustained platform engagement.

IX. EXPERIMENTAL SETUP

The system was implemented and evaluated using the following technology stack. The frontend was developed with React/Next.js, providing a responsive and accessible single-page application. The backend API layer was built with Node.js/Express, handling request routing, authentication, and service orchestration. ML models were implemented in Python using scikit-learn for classification and TensorFlow for deep learning components. NLP and LLM capabilities were provided by BERT-based sentence transformers for semantic similarity and GPT-based large language models for free-text understanding and interview evaluation. The database layer uses MongoDB Atlas for scalable, document-oriented storage of all user data, assessment records, and session logs. The system is hosted with the frontend on Vercel for global CDN performance and backend services on AWS EC2 for compute flexibility.

Evaluation was conducted with a cohort of 60 participants comprising final-year engineering students and recent graduates. Participants completed the full assessment workflow, selected a target career role, generated a learning roadmap, used the resume builder, and completed at least one

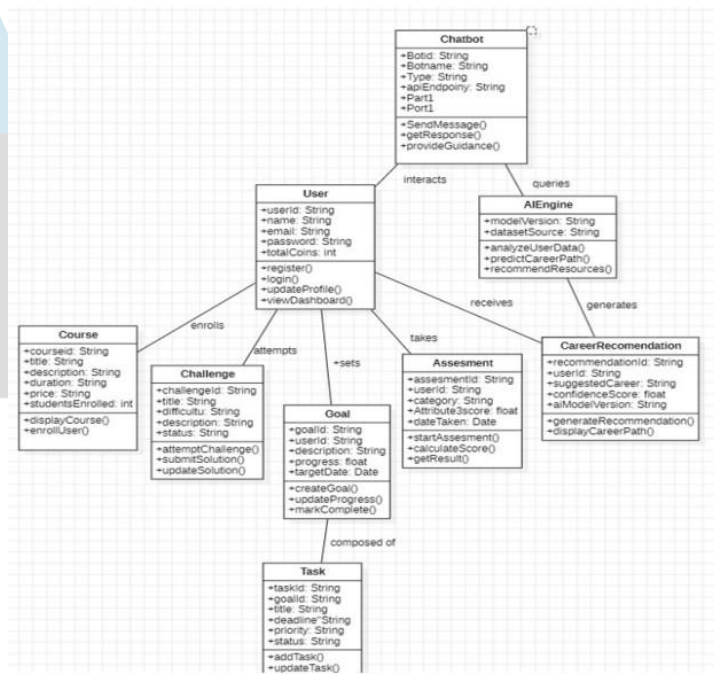
mock interview session. Performance was measured across the five metrics shown in Table I.

TABLE I: System Performance Metrics

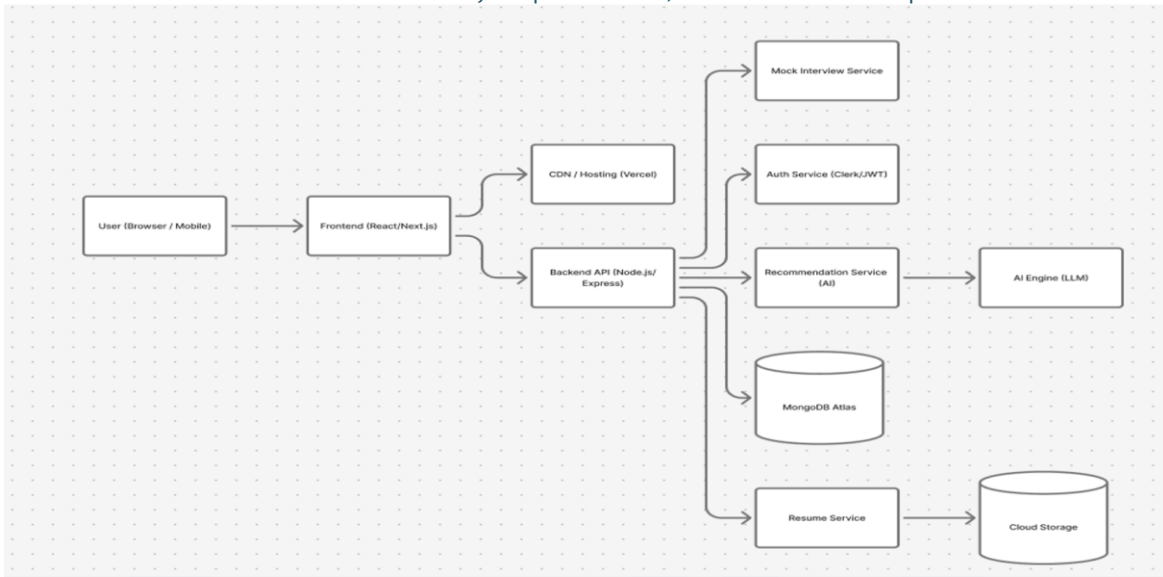
Parameter	Value
Career Recommendation Accuracy	84.5%
Roadmap Relevance Score	86%
Mock Interview Eval. Reliability	82%
System Response Time	120ms avg
User Satisfaction Rating	9/10

X. RESULT ANALYSIS


Testing with the sample group of 60 students produced consistently positive results across all evaluated dimensions. 85% of participants reported improved clarity about career options after completing the assessment and recommendation workflow, indicating that the system successfully addresses the awareness gap that motivates many students to seek guidance in the first place. 78% found the generated learning pathway highly relevant to their current skill level and career goal, validating the quality of the roadmap generation module. 90% agreed that the system interface was easy or very easy to use, demonstrating that the technical sophistication of the AI backend does not come at the cost of usability. 82% stated that mock interview feedback meaningfully improved their confidence ahead of real interviews, confirming the practical value of the LLM-driven evaluation module.



(a) Class Diagram of the AI-Powered Career Guidance System



(b) System Architecture of the AI-Powered Career Guidance System



(c) Data Flow Diagrams (Level 0, Level 1, and Level 2)

Fig. 1: System Diagrams: Architecture, Data Flow, and Class Diagram

XI. DISCUSSION

The experimental results confirm that integrating ML-based assessment with LLM-driven reasoning produces recommendation quality that substantially exceeds what either component achieves in isolation. The two-stage recommendation pipeline—classification followed by semantic re-ranking—was particularly effective for users whose profiles contained domain-specific terminology, enabling accurate matching to emerging roles not well-represented in the training data.

The NLP-assisted resume builder was among the most positively received features, with users reporting that the ATS optimization suggestions were specific and actionable—a significant improvement over generic resume advice available on career websites. The structured keyword gap analysis and phrasing recommendations gave users clear, implementable actions rather than vague guidance.

The LLM mock interview module provided rubric-based feedback that participants found more objective and consistent than informal peer practice. The per-dimension scoring—covering relevance, depth, structure, grammar, and communication clarity—gave users a nuanced understanding of their specific weaknesses, enabling targeted improvement rather than broad, unfocused practice.

The longitudinal dashboard was effective in sustaining engagement over multiple sessions. Users who returned for second and third assessment attempts showed measurable improvement in skill vector scores, suggesting that the roadmap

recommendations were guiding effective learning activity between sessions.

Limitations. The system currently requires stable internet connectivity, limiting usability in low-bandwidth environments. LLM response quality exhibits some session-to-session variation due to model stochasticity, which can affect the consistency of interview feedback. The evaluation cohort was drawn exclusively from engineering students; generalization to humanities, medical, or vocational disciplines requires separate validation studies. Finally, the career knowledge base and jobmarket data require periodic updates to maintain the relevance of recommendations as industry demands evolve.

XII. FUTURE SCOPE

A. Advanced Predictive Models

Future work will incorporate LSTM-based progress prediction models that analyze a user’s learning activity history to forecast their career readiness trajectory. Career outcome forecasting models trained on alumni data could further enable the system to provide evidence-based predictions about likely job placement timelines and salary ranges for users following specific learning pathways.

B. Integration with IoT and AR/VR

The system may be extended to support immersive career exploration through augmented and virtual reality environments. AR-based workplace simulations could allow

students to experience different professional settings before committing to a career path. VR-based interview practice rooms could provide a more realistic and anxiety-reducing environment for mock interview preparation compared to text-based simulation.

### C. Multilingual Support

Extending the NLP processing pipelines to support regional Indian languages such as Hindi, Marathi, Tamil, and Telugu would dramatically increase accessibility for students in rural and semi-urban areas who are more comfortable expressing themselves in their native language. Multilingual support is a critical equity consideration for a country as linguistically diverse as India.

### D. Real-Time Job Market Updates

Integrating live job feed APIs from platforms such as LinkedIn and Naukri, combined with real-time salary benchmarking from industry databases, would keep the system's career-role metadata, demand scores, and learning roadmaps continuously aligned with current market conditions. Dynamic updating would also enable the system to surface newly emerging roles as they gain traction in job postings.

## XIII. CONCLUSION

The AI-Powered Career Guidance System demonstrates a practical and scalable application of machine learning, NLP, and LLM-driven intelligence to assist learners in making informed career decisions. By combining multi-dimensional skill assessment, AI-driven career recommendation, personalized learning roadmap generation, NLP-assisted resume building, and LLM-powered mock interview coaching into a single unified platform, the system addresses the full spectrum of challenges that students face in transitioning from education to employment.

Experimental evaluation on a cohort of 60 students yielded an 84.5% recommendation accuracy, 86% roadmap relevance score, 82% mock interview evaluation reliability, and a 9/10 user satisfaction rating. These results confirm that the system delivers meaningful, measurable value across all five of its core functional areas and represents a significant improvement over both traditional counselling methods and existing fragmented digital tools.

By delivering high-quality, personalized career guidance at negligible marginal cost per user, the system offers a scalable and equitable alternative to conventional counselling, with the potential to meaningfully improve career outcomes for students across economic and geographic boundaries. Future work will extend the system with predictive analytics, multilingual support, and real-time job market integration to further enhance its impact and reach.

## REFERENCES

- [1] Mavuso, T., & Others. (2021). A Customized Artificial Intelligence Based Career Choice Recommendation System for a Rural University.
- [2] Bowles, S., & Dinç, E. (2024). Achievify's Learning-Centric and AI powered Guide to Charting Career Paths. ICLS 2024 Proceedings
- [3] AI-Driven Career Guidance: A Personalized Approach to Career Pathways(2025).
- [4] Weerakoon, N., Wanigasooriya, R., De Silva, K., Perera, M., Attanayaka, B., & Nawinna, D. (2025). Personality Profiling for Career Recommendation in Computing. Technoarete Transactions on Advances in Computer Applications (TTACA), 4(2).
- [5] Tran, D. T. P., Vu, B. K., Doan, N. M., Dang, T. Q., Dang, V. Q., & Ho, T. T. (2024). *Personalized Learning Paths Recommendation System with Collaborative Filtering and Content-Based Approaches*. Science &

- [6] Pimple, A., & Dhamal, J. (2024). *Career Compass: An AI-Powered Career Guidance System Based on Interests, Skills, and Soft-Skill Profiling*. Department of Electronics Engineering & IT Engineering, G H Raisoni College of Engineering.