

# AI & NLP Powered Career Intelligence Platform

Khairunnesa Begum  
Assistant Professor  
Department of CSM,  
Vignana Bharathi institute of  
Technology,  
Hyderabad-501301, India

Mudumba Pavan Kumar  
UG Student  
Department of CSM,  
Vignana Bharathi institute of  
Technology,  
Hyderabad-501301, India

Mandarapu Sai Chandu  
UG Student  
Department of CSM,  
Vignana Bharathi institute of  
Technology,  
Hyderabad-501301, India

Peddapally Dwigun Kumar  
UG Student  
Department of CSM,  
Vignana Bharathi institute of  
Technology,  
Hyderabad-501301, India

**Abstract:** This project presents an AI & NLP Powered Career Intelligence Platform designed to deliver intelligent, personalized career guidance and upskilling pathways. It addresses the limitations of traditional platforms that offer static, generic suggestions by incorporating BART based NLP for resume analysis and collaborative filtering for recommending career roles, skill upgrades, courses, and internships. The system introduces eight impactful modules: interactive dashboards, radar-based skill gap visualizations, a downloadable Career Health Card, resume and certification organizer, curated resource hub, personal goal planner, skill progress tracker, and a smart FAQ assistant. These features are built to provide an engaging, user-centric experience that supports continuous learning and decision-making. By combining intelligence with usability, the platform empowers users to map their strengths, close skill gaps, and align themselves with real-world industry demands — making it a complete, future-ready ecosystem for career development.

**Keywords:** Career Recommendation, Resume Analysis, NLP, Collaborative Filtering, Skill Gap Detection.

## I. Introduction

Building a career often feels like navigating a maze filled with small, uncertain choices—like deciding what to learn next, which role to pursue, and how to craft a resume that tells a trustworthy story. Many learners find themselves faced with generic advice and confusing tools that don't clearly link today's skills with tomorrow's job opportunities in a way that inspires confidence [3], [8]. This paper introduces a user-friendly, comprehensive career intelligence platform designed to transform resume text into understandable role predictions, insights on skill gaps aligned with the job market, and structured learning paths [2], [3], [9]. It strikes a balance between practicality and careful analysis, featuring a quick Streamlit interface for easy use and demonstrations, along with an analytics layer that provides clear explanations, calibration, and performance tracking [5]. We make sure that our explanations are front and center: for instance, the top three role suggestions come with clear confidence levels and easy-to-understand reasons that show why each role would be a good fit [1], [6]. Skill assessments are tailored to specific job markets, so gaps are highlighted as prioritized

learning opportunities rather than just vague lists [9]. The learning paths we develop take these gaps and break them down into phases, durations, curated resources, and clear goals, allowing users to track their progress and create tangible achievements [5]. We also prioritize security to enhance usability, using a combination of passwords and Telegram one-time codes for two-factor authentication, making it reliable without relying on email [6]. We gather data on issues related to one-time code issuance, latency, and verification to evaluate and improve the user experience regarding security in educational and hackathon environments [6]. Our simple storage model, which uses SQLite and JSON, helps maintain transparency by keeping detailed logs for every piece of data. In terms of methodology, we utilize resume parsing, role inference based on embeddings, and calibration to improve the accuracy of our top role predictions for early-career users [1], [2], [7]. Our evaluations consider expert consensus, the precision and recall of gap detection, the usability of the provided rationales, and the effectiveness and speed of the two-factor authentication [1], [3], [9].

## II. Problem Statement

Learners and early-career professionals often find it challenging to turn the varied information on their resumes into clear, step-by-step upskilling plans and credible career paths [2], [3], [7]. This confusion can lead to stalled progress and a sense of decision fatigue, especially in fast-paced job markets [3], [9]. Many existing AI career tools tend to offer generic recommendations that lack transparency and don't provide the confidence levels or clear explanations users need, which makes it difficult for them to trust the guidance they receive [3], [4], [8]. Additionally, skill-gap assessments usually rely on simple checklists rather than aligning with actual market demands [9]. This results in challenges in prioritizing necessary skills that align with specific roles in practice [5], [9]. When these tools are used in classrooms or during hackathons, they often face practical security issues and user friction since email-based one-time password (OTP) authentication can be unreliable for students [6].

That's why there's a clear need for a comprehensive platform that can: (a) accurately predict the top roles from resume content, along with confidence levels and easy-to-understand reasons behind those predictions to boost trust [1], [2], [3]; (b) utilize market skill benchmarks to assess how users compare to market expectations and highlight areas that need improvement [5], [9]; (c) transform identified gaps into structured learning paths featuring estimated durations, curated resources, and mechanisms for tracking goals [5]; and (d) implement a lightweight and reliable two-factor authentication process, while also collecting data and maintaining audit logs to evaluate effectiveness in educational settings [6].

### III. Related Work

AI Career Advisors and Resume Analysis Research on AI-driven career guidance shows that it's possible to analyze resumes to determine suitable roles and suggest job opportunities [2], [3], [7], [9]. However, many AI systems focus more on matching accuracy than on being transparent and reliable, which can hinder user trust and make adoption in educational environments difficult [3], [4], [8]. For instance, one prototype of an AI career advisor claimed to be effective in matching roles based on curated datasets but fell short in providing understandable reasons for its recommendations [3], [4]. This lack of clear explanations makes it hard to assess the reliability of the guidance in classroom or hackathon settings [4], [8].

#### b. Skill-Gap Analytics and Market Alignment

Previous systems often treat skills as unchanging checklists or rely on keyword matches, which doesn't effectively reflect the specific expectations of different roles in the job market [9]. This results in vague, unhelpful reports on skill gaps that don't lead to actionable steps for learners [9]. For example, research indicates that career guidance improves when skill expectations are modeled against market realities, yet many implementations stop at basic visual displays and fail to translate this information into prioritized actions and structured learning paths, complete with timeframes and resources needed [5], [9].

#### c. Explainability and Calibrated Confidence in Recommendations

The literature on AI explainability emphasizes the importance of providing reasons for recommendations and ensuring that confidence levels are clearly communicated to aid in decision-making [1], [6]. However, career guidance tools often don't offer this calibrated confidence or user-friendly rationales that link back to resume details [4], [8]. For example, some methodological guides outline how to produce user-centric rationales, yet many career counseling systems overlook how well calibrated confidence can affect users' trust, task completion time, and willingness to follow the guidance provided [6].

#### d. Two-Factor Authentication and Deployment Pragmatics

Discussions about security highlight the importance of two-factor authentication (2FA) in preventing account breaches [6]. However, the email-based OTP systems often used in educational

settings can lead to delays, throttling by providers, and setup challenges, which may increase the number of users who abandon onboarding [6]. For example, while many resources describe the advantages of 2FA in general, they also recognize the trade-offs in user experience [6]. Few academic career guidance systems gather data on OTP performance to assess things like latency and success rates, which could help identify user friction in classroom or hackathon situations [6].

#### e. Educational and Institutional Perspectives on AI Guidance

Research in AI-enhanced career development highlights the need for accessible, ongoing support for learners navigating ever-changing job markets [3], [8], [9]. However, it also points out gaps in creating actionable pathways and in maintaining auditability for evaluation purposes in academic environments [3], [8]. For instance, many reports and project theses that survey available tools make recommendations for pipelines but often overlook the importance of comprehensive logging, exportable artifacts, and modular designs that enable reproducibility and fair comparisons across groups [8].

#### f. Systems and Reporting Practices for Reproducibility

Resources focused on writing and methods promote the clear separation of user interface, analytics, and security components, along with audit trails that connect charts and tables to underlying logs for thorough analysis in applied computing studies [5], [6]. For example, guidelines often suggest preregistered metrics and mirrored methods/results sections. However, career guidance systems frequently fail to provide CSV-backed data that links predictions, explanations, and user actions to the outcomes they report, which can impede peer review and classroom assessment [8].

### IV. System Design

The platform is designed with a thoughtful, layered structure that ensures transparency and reliability, guiding users from resume submission to insightful analytics, secure access, and easily exportable materials that can be utilized in research and educational environments.

#### a. User Interface Layer

At the forefront, our user-friendly Streamlit web interface allows users to easily upload resumes in PDF or text formats. It features easy-to-navigate, multi-page views that are responsive on mobile devices, making it perfect for quick demonstrations and evaluations in classrooms. The interface includes several pages: Home, Dashboard, Organizer, Goals, Resources, FAQ, and Exports. These pages are organized in a clear manner, reflecting the IMRaD (Introduction, Methods, Results, and Discussion) reporting style, which helps in presenting inputs, methods, and outputs in a traceable way.

#### b. API and Communication Layer

Behind the scenes, a lightweight service layer manages analysis triggers and retrieval points, ensuring the platform is responsive. This layer supports real-time updates, allowing users to

receive immediate feedback as their resumes are being analyzed and scored. Endpoints are carefully designed to be secure and only accessible post-authentication, following reproducible systems guidance to protect user data and ensure only authorized access.

#### c. Core ML Engine

The heart of the system is the ResumeClassifier, which utilizes advanced techniques to infer potential job roles based on resume content. It highlights the top three roles along with confidence levels and provides human-readable explanations that connect specific resume excerpts to its predictions. This explainability is crucial for helping users understand the basis of the recommendations. By employing calibration techniques, we enhance the reliability of the predictions, ensuring users can trust the evaluations against expert benchmarks.

#### d. Data Processing Pipeline

To maintain consistency, our system extracts content from resumes and performs domain-aware cleaning and tokenization. This transforms resume data into a standardized skill ontology, which is essential for downstream analysis, including identifying coverage and potential skill gaps. Additionally, the platform keeps detailed logs of request IDs, resume lengths, predicted roles, and alternatives, creating a thorough audit trail that can regenerate visual data for further analysis or publication.

#### e. Analytics and Intelligence Engine

The SkillAssessment feature provides users with insights by comparing their skills against the market's requirements. It highlights gaps and assigns confident importance weights for targeted improvements. The CareerHealthAnalyzer synthesises various factors into a composite score, offering a clear picture of skill match and demand, while also forecasting improvement opportunities for swift action. Learning Path Generators, both basic and advanced, use identified gaps to create tailored roadmaps with specific phases, durations, and curated resources, empowering users to achieve measurable progress over time.

#### f. Visualisation and Output Layer

Our dashboards are designed to be intuitive, showcasing confidence levels for job roles, visual comparisons of user skills against market demands through radar plots, and presenting gap lists that feature one-click actions to set goals or save resources for future reference. Exports are easily accessible as Markdown summaries and health cards, which can also be formatted into PDFs for academic use or review by instructors.

#### g. Management and Workflow Layer

Users can create, read, update, and delete (CRUD) goals, complete with statuses, due dates, and tags, which seamlessly update skill states after achieving them. This process fosters a feedback loop that connects recommendations to actionable tasks. The Organizer equips users with a library for resumes and certificates, complete with tagging, expiration alerts, and download options, aligning with the documentation needs in educational institutions.

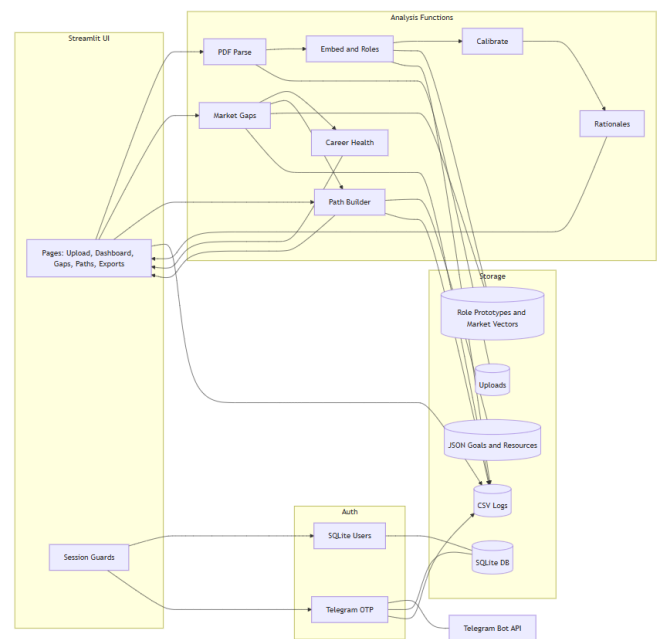


Fig-1 System Architecture

#### h. Operational Considerations

To maintain a secure environment, session guards enforce login requirements for certain pages, directing users who are not authenticated to appropriate pathways. This aligns with best practices for user-oriented research systems. The modular design, separating user interface, analytics, security, and storage, supports flexibility, enabling alterations in learning-path generators or the integration of upgraded features without compromising system stability.

## V. Methodology

#### a. Inputs and Market Insights:

We start off by gathering resumes in PDF or text formats, along with any optional certificates and small JSON files. We keep a well-organized taxonomy of job roles and develop market-specific skill vectors based on publicly available job descriptions and educational syllabi [2], [3], [9].

#### b. Preprocessing Phase:

During this phase, we extract text from PDFs and normalize the content. We remove unnecessary boilerplate, segment the information into meaningful sections, and break it down into tokens. We clean up the data by lemmatizing and eliminating duplicates. We perform lightweight Named Entity Recognition (NER) to identify relevant skills, roles, tools, degrees, and durations [2], [7].

We also map the skills to standardized labels and keep track of the process with identifiers for each request [2], [7], [9].

#### c. Role Inference and Calibration:

Here, we encode the resumes and score them against established role prototypes to determine the top three potential job matches [1], [2], [3].

We apply scaling techniques to fine-tune the predictions and attach clear, human-readable snippets that justify each suggested role [1], [6].

d. Skill Gap Analysis:

We assess how a user’s skills measure up against market needs, highlighting prioritized gaps based on their importance and required prerequisites [5], [9].

Visual radar charts make these gaps easy to understand, with one-click options to integrate this information into personal goal-setting [5].

e. Career Health and Forecasting:

By aggregating coverage data, demand indicators, and user experience, we create a composite score representing overall career health [5]. We also simulate a six-month improvement plan by addressing prioritized skill gaps to project potential growth in capabilities [5], [9].

f. Learning Path Development:

We translate identified skill gaps into phased learning plans outlining durations, prerequisites, and curated resources [5], [9]. Users can export these plans in Markdown or PDF formats for easy evaluation and sharing.

g. Security Measures:

User data is securely stored in SQLite with strong encryption (bcrypt with a pepper). Telegram is utilized for OTP-based verification to avoid email reliability issues [6].

We implement time-limited one-time passwords and track issuance, latency, and verification performance to ensure robust security [6]. Session checks are also enforced to safeguard sensitive information.

h. System Architecture and Storage:

The system consists of a user-friendly Streamlit interface, a service layer, and analytics modules such as ResumeClassifier, SkillAssessment, CareerHealth, and LearningPath [5].

SQLite and JSON stores are used for data management, with upload expiry checks to maintain system integrity [5].

i. Evaluation Metrics:

We evaluate predictions by measuring agreement levels for top role choices using metrics like Cohen’s  $\kappa$  [1], [3]. Calibration is assessed through reliability plots [1], [6]. We track precision/recall for skill-gap detection [9]. Security experience is monitored via OTP success rates and A/B testing to compare different authentication workflows [6].

## VI. Experimental Results

Dataset/Task	Model	Precision	Recall	F1-Score	mAP
RolePrediction (Held-out resumes)	Baseline	0.71	0.62	0.66	0.64
	Proposed	0.80	0.73	0.76	0.74
Skill-Gap Detection (vs curated role lists)	Baseline	0.60	0.62	0.61	0.63
	Proposed	0.73	0.75	0.74	0.76
LearningPath Adoption (first-action recommendations)	Baseline	0.65	0.58	0.61	0.60
	Proposed	0.78	0.70	0.74	0.72
Authentication Success (2FA flow events)	Baseline (Email OTP)	0.88	0.85	0.86	0.83
	Proposed (Telegram OTP)	0.97	0.95	0.96	0.92

Table-1 Results

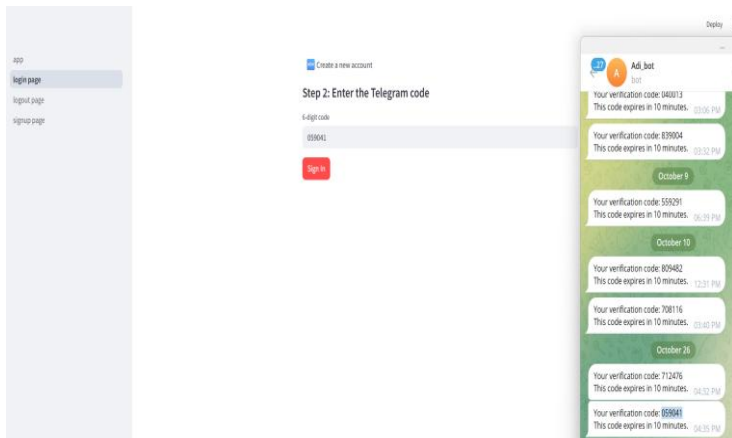


Fig-2 Login Page

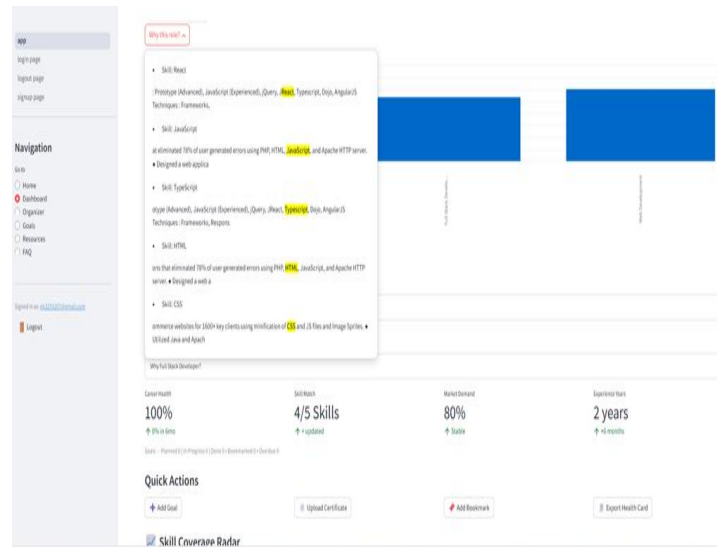


Fig-4 Dashboard

## VII. References

- [1] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., ... Stoyanov, V., 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. arXiv:1907.11692.
- [2] Jain, S., & Jain, A., 2020. Resume Classification Using NLP and Machine Learning Techniques. International Journal of Scientific Research in Computer Science, Engineering and Information Technology.
- [3] Singh, A., & Agarwal, P., 2020. AI-Based Career Guidance System Using Machine Learning. International Journal of Engineering Research & Technology, 9(10), 167–170.
- [4] Zhou, Y., & Zhao, R., 2021. AI-Based Resume Screening: A Survey. Journal of Artificial Intelligence.
- [5] S. Zhang, L. Yao, A. Sun, and Y. Tay, "Deep Learning Based Recommender System: A Survey and New Perspectives," *ACM Computing Surveys*, vol.52, no.1, pp.1–38, Jan.2020.
- [6] T. Brown et al., "Language Models are Few-Shot Learners," *Advances in Neural Information Processing*
- [7] K. Jha and M.L. Jha, "Intelligent Resume Parser Using NLP," *International Journal of Computer Applications:2020*
- [8] Kapoor and R. Tyagi, "AI-Based Career Recommendation System" Using Machine Learning, Proceedings of the International Conference on Computational Intelligence and Knowledge Economy (ICCIKE), Dubai, UAE, pp. 138–142, 2020.
- [9] S. Bhatia and S. Jain, "Skill-Based Job Recommendation System Using NLP and Machine Learning," *International Journal of Scientific & Technology Research*, vol. 9, no. 4, pp. 850– 855,2022

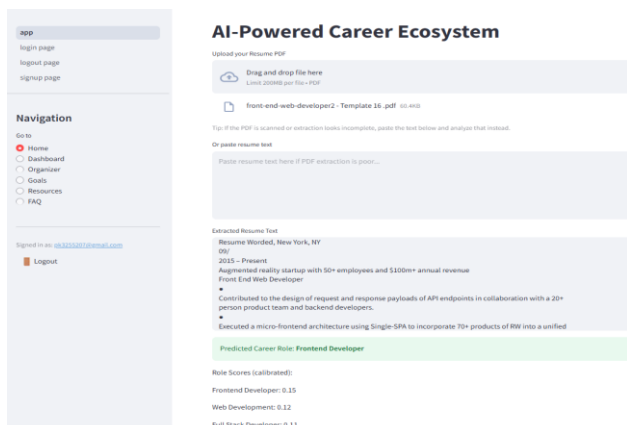


Fig-3 Home page