



A.C.E: Where Ambitions Meet AI-Driven Career Intelligence

Liby P Job¹, Tania Khalid², Sreehari CS³, Steffin Sibi John⁴, Vyshnav PK⁵, Rena Manar Majeed⁶

¹Assistant Professor, Department of Linguistic Studies, Yenepoya Deemed to be University, Bengaluru, Karnataka, India.

^{2,3,4,5,6}UG - Bachelors in Computer Application, Yenepoya Deemed to be University, Bengaluru, Karnataka, India.

Email ID: liby.p.blr@yenepoya.edu.in¹, 230858@yenepoya.edu.in², 330547@yenepoya.edu.in³, 431698@yenepoya.edu.in⁴, 530353@yenepoya.edu.in⁵, 630053@yenepoya.edu.in⁶

Abstract

With global job markets evolving faster than ever and data-driven industries becoming the norm, there is a growing need for career guidance tools that are smart, flexible, and built to scale. This paper introduces A.C.E. (Aspire, Connect, Evolve), an AI-powered web platform designed to deliver personalised, context-aware career advice using machine learning and natural language processing (NLP). The system works with both structured data — such as academic records and skill scores — and unstructured text input, including user interests, feedback, and career-related information. Through NLP techniques like text preprocessing, semantic analysis, feature extraction, and context interpretation, raw text is transformed into meaningful insights that support better decision-making. These insights feed into a recommendation engine that ranks career paths and generates tailored skill development plans based on current labour market trends. Built with modern web technologies and Python-based AI toolkits, the platform is designed to be modular, scalable, and capable of processing information in real time. The results demonstrate that incorporating NLP-based semantic analysis meaningfully improves the accuracy, relevance, and adaptability of career recommendations. Ultimately, A.C.E. offers an intelligent and accessible platform that helps close the gap between education skill development, and the real demands of today's job market.

Keywords: Natural language processing, semantic analysis, artificial intelligence, career recommendation system, machine learning, text mining, feature extraction, context-aware recommendation, education data mining, intelligent decision support system, career prediction, and intelligent web platform.

1. Introduction

The global workforce is undergoing rapid transformation due to the Fourth Industrial Revolution (4IR). The Fourth Industrial Revolution has been marked by the advancement of artificial intelligence, automation, and big data analytics. The recent projections indicated that 40 to 50% of job skills are likely to be transformed in the next ten years. This indicates the need for a smart system for career guidance. The traditional system for career guidance has been based on psychological tests and a set of defined relationships between academic and professional careers. These systems provide a general idea about the potential career choices for a person

but fail to take into account the recent trends in the job market. Students are opting for academic careers without considering the recent trends in the job market. The traditional system for career guidance lacks personalization and the integration of recent trends in the job market. To overcome these challenges in the traditional system for career guidance, this research proposes a smart system for career guidance using A.C.E (Aspire, Connect, Evolve), an intelligent system for career guidance using various artificial intelligence concepts. The contributions of this research are as follows:

- Proposing a career recommendation system



using multiple factors with the help of artificial intelligence concepts.

- Proposing a natural language processing system for processing unstructured interest descriptions and computing semantic similarity. Integrating geographic labor-demand indicators for context-aware career rankings.
- Implementing a cloud-compatible three-tier architecture for scalable deployment.
- Conducting empirical evaluations through structured experiments to assess precision, stability, and response performance.

2. Related Work

In the past two decades, recommender systems, educational data mining, and natural language processing have made significant strides. However, their use in career guidance systems has not been as widespread.

2.1.Recommender Systems

Recommendation engines are popular in e-commerce, media platforms, and social networks [1], [2]. These systems usually rely on:

- Content-based filtering, which links user preferences with item attributes.
- Collaborative filtering, which identifies trends among individuals who exhibit comparable actions.
- Hybrid techniques, which incorporate both tactics to improve the accuracy of recommendations.

The majority of conventional recommendation algorithms are not particularly aware of employment market developments, notwithstanding their popularity in commercial settings [7], [8], [9]. There is minimal personalization offered by career guidance platforms that rely on static mappings between academic subjects and professions.

2.2.Educational Data Mining

Using techniques including classification, clustering, and regression analysis, educational data mining (EDM) looks for patterns in student data [6]. These methods aid in predicting student performance and identifying learning trends. While adaptive learning systems and student performance have been successfully predicted by EDM, little study has been

done to establish a connection academic insights with career path suggestions driven by job market demand.

2.3.Natural Language Processing in Career Systems

Text input is converted into structured representations by natural language processing [3], [4], [5]. Typical pipelines comprise of:

- Tokenization.
- Stop-word removal.
- TF-IDF feature extraction.
- Semantic similarity calculation using cosine similarity.
- Clustering or classification.

While NLP techniques can deepen understanding of user interests, many career platforms still mainly rely on keyword matching instead of interpreting semantics. Furthermore, it is scarce to combine real-time labor market data with NLP outputs.

2.4.Research Gap

Academic records, labor market data, and semantic interest modeling are rarely integrated [7], [8].

3. System Architecture

The display, application, and data layers comprise the three-tier architecture of the A.C.E. platform. Scalability, maintainability, and cloud deployment are made possible by this modular configuration.

3.1.Presentation Layer

React.js is used in the frontend's construction to provide an interactive area where users may input academic information, skill sets, and descriptions of their interests. Important characteristics consist of:

- Role-based dashboards for students, teachers, parents, and administrators.
- A display of ranked career recommendations.
- Skill gap visualization.
- Career readiness score.
- Real-time recommendation updates.

Token-based authentication is used in secure RESTful APIs to communicate with backend services.

3.2.Application Layer

The essential system logic and AI processing modules are located on the application layer. Backend API (Node.js + Express)

The backend service oversees:

- User authentication and authorization.

- Data validation and preprocessing.
- API routing.
- Communication with the AI micro service.

3.2.1. Authentication Module

Secure access control is achieved using JSON Web Tokens (JWT) along with role-based authorization.

3.2.2. AI Micro service (FastAPI – Python)

The AI engine operates as a separate micro service responsible for:

- Data normalization
- NLP processing
- Feature vector creation
- Weighted scoring computation
- Career ranking generation

AI tasks can be independently scaled thanks to this micro service architecture without interfering with other services.[8]

3.3. Data Layer

MongoDB is used in the database layer to provide flexible storage for semi-structured and structured data. Among the primary collections are: As shown in Table 1.

Table 1 Core Database Collections

Collection Name	Description
Users	Stores authentication and role data Student Profiles
Academic records and personal details	Skill scores and competency mapping
Skills	Interests
NLP-processed interest vectors	Career domain metadata
Careers	Market Trends
Industry growth and demand indicators	System interaction tracking
Activity Logs	

3.4. Architecture Overview

As Shown in Figure 1.

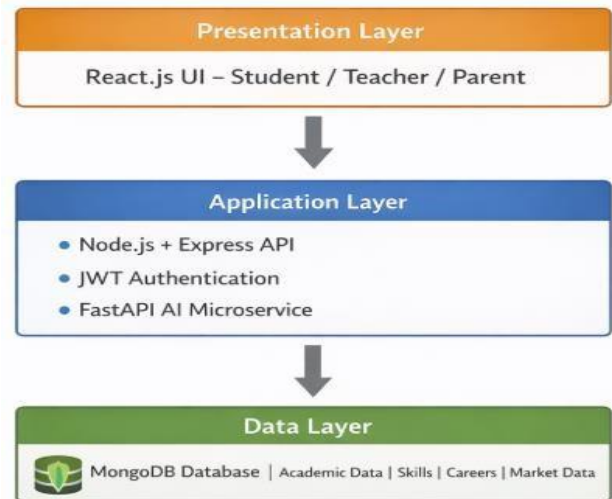


Figure 1 System Architecture of A.C.E Platform

4. Methodology

4.1. Data Collection

The system collects data gathered from a number of sources, including but not limited to:

- Academic grades
- Technical and soft skill ratings
- Certifications
- Text descriptions of interests
- Labor market demand indicators

4.2. Weighted Multi-Factor Scoring Model

Career recommendation scores are calculated using a weighted linear model:

$$\text{FinalScore}_i = w_1 A_i + w_2 S_i + w_3 C_i$$

Where:

A_i = Academic score

S_i = Skill score

C_i = Context factor (labor demand)

Weights used in the experiments are:

$$w_1 = 0.4, w_2 = 0.4, w_3 = 0.2$$

Academic Score Calculation

Normalized subject grades and particular weightings for each area are used to calculate academic scores.

$$X_{\text{norm}} = (X - X_{\text{min}}) / (X_{\text{max}} - X_{\text{min}}).$$

4.3. NLP Processing Pipeline

This pipeline is applied to process user interest descriptions:

- Tokenization.
- Stop-word removal.

- TF-IDF feature extraction.
- Cosine similarity calculation.
- Career cluster mapping.
- Text normalization.
- Stemming or lemmatization.
- Extracting keywords to identify skills.

4.4. Algorithm Flow

As Shown in Figure 2

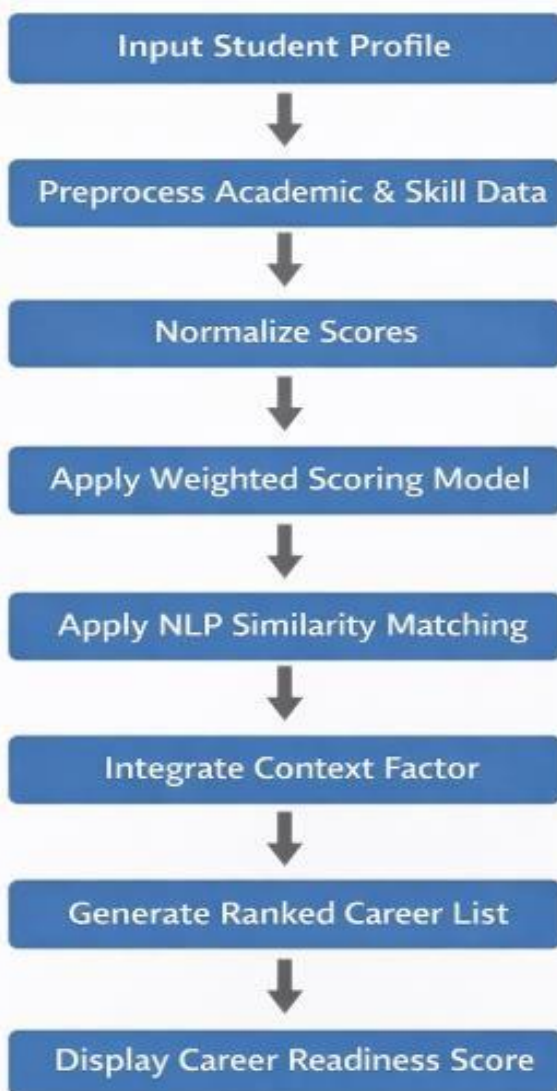


Figure 2 Recommendation Algorithm Flow

5. Experimental Evaluation

5.1. Dataset Construction

For assessment, a synthetic dataset was made. A list that includes 300 student profiles. The student profiles show how the students are divided into academic

streams, which is as follows:

- Science: 120 students
- Commerce: 100 students
- Humanities: 80 students

Each profile contains:

- Subject grades (5 to 8 subjects).
- Technical skill ratings.
- Soft skill scores.
- Certification indicators.
- Textual interest descriptions.

The system evaluated 50 predefined career domains, each with skill requirements, subject weights, and demand indices. The subsequent was applied to model labor demand:

Demand Index $\in [0.7, 1.3]$

5.2. Evaluation Metrics

Performance was assessed using a number of indicators.

Recommendation Precision

$P = \text{Correct Top Recommendations} / \text{Total Profiles}$

The precision observed is 89%.

Ranking Stability Index (RSI)

$RSI = 1 - (\text{Observed Ranking Variance} / \text{Maximum Possible Variance})$

Observed RSI is 92%.

Mean API Response Time

The average latency for system inference is 1.3 seconds.

Context Sensitivity Score (CSS)

This gauges how sensitive rankings are to shifts in the job market.

The observed CSS is 0.84.

5.3. Comparative Benchmark

A conventional rule-based career guidance system was contrasted with the A.C.E. model.

- As a result, the rule-based system depended on pre-established rules, it was less adaptable to a variety of user inputs.
- A conventional rule-based career guidance system was contrasted with the A.C.E. model.
- Because the rule-based system depended on pre-established rules, it was less adaptable to a variety of user inputs. As Shown in Table 2.

Table 2 Comparative Performance

Feature	Baseline Mode	A.C.E Model
Personalization	Low	High
Market Adaptability	Absent	Integrated
Semantic Interest Processing	Keyword-based	NLP-based
Skill Gap Identification	Limited	Advanced
Scalability	Limited	High
Recommendation Precision	76–78%	89%

5.4.NLP Ablation Study

For the purpose of gauging the NLP module's contribution, an ablation experiment was carried out.

- Configuration 1: Structured data only.
- Configuration 2: Structured + NLP semantic similarity. As Shown in Table 3.

Table 3 NLP Impact Evaluation

Configuration	Precision
Without NLP	78%
With NLP	89%

6. Performance Analysis

The computational complexity of the recommendation engine is dependent on three primary operations:

- Structured computation of scores.
- TF-IDF feature extraction.
- Similarity matching.

Overall complexity is:

$$O(nm + nk)$$

Where:

- n = number of career domains
- m = TF-IDF feature size

- k = academic/skill parameters

Real-time inference remains effective when n is 50 and m is less than 1000.

- Load testing was used to confirm stable performance during concurrent queries: As Shown in Table 4.

Table 4 System Response Time for Different Numbers of Users

User	Response Time
50	1.3 s
100	1.45 s
200	1.7 s

7. Discussion

The A.C.E framework is seen to enhance the accuracy, personalization, and scalability of career suggestion in comparison to traditional models, based on experimental data. The alignment of user interests and professional domains is enhanced by semantic interpretation via NLP, while labor demand is considered to make career suggestions relevant to changing workforce needs. The weighted multi-factor scoring system ensures clear and interpretable decision logic, which is crucial for institutional acceptability in contrast to black box deep learning models. There are also certain limitations, such as domain limitations in career domains and simulated labor demands, while transformer models and labor market APIs could be future directions for improvement. [9]

Conclusion

A system of career recommendations based on artificial intelligence, A.C.E., which connects scholastic achievements, personal inclinations, and changing labor market trends, is proposed in this paper. The proposed system consists of:

- Multi-factor scoring.
- NLP-based interest analysis.
- Labor demand indicators with context-aware capabilities.

The experimental results demonstrated 89% accuracy in recommendations and high ranking stability,



outperforming traditional rule-based systems. The proposed system is applicable to career counselling platforms and prominent education institutions owing to its micro service-based AI engine and three-tier system design. Taking all of this into consideration, the A.C.E. framework is a useful tool for wise career counselling in a changing labor market.

References

- [1]. P. Resnick and H. R. Varian, "Recommender systems," *Communications of the ACM*, vol. 40, no. 3, pp. 56–58, 1997.
- [2]. R. Burke, "Hybrid recommender systems: Survey and experiments," *User Modeling and User-Adapted Interaction*, vol. 12, no. 4, pp. 331–370, 2002.
- [3]. G. Salton and C. Buckley, "Term-weighting approaches in automatic text retrieval," *Information Processing & Management*, vol. 24, no. 5, pp. 513–523, 1988.
- [4]. C. D. Manning, P. Raghavan, and H. Schütze, *Introduction to Information Retrieval*. Cambridge, U.K.: Cambridge University Press, 2008.
- [5]. D. Jurafsky and J. H. Martin, *Speech and Language Processing*, 3rd ed., Pearson, 2024.
- [6]. C. Romero and S. Ventura, "Educational data mining: A review of the state of the art," *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, vol. 40, no. 6, pp. 601–618, 2010.
- [7]. F. Ricci, L. Rokach, and B. Shapira, *Recommender Systems Handbook*, 2nd ed. New York, NY, USA: Springer, 2015.
- [8]. M. J. Pazzani and D. Billsus, "Content-based recommendation systems," in *The Adaptive Web*. Berlin, Germany: Springer, 2007, pp. 325–341.
- [9]. J. L. Herlocker, J. A. Konstan, and J. Riedl, "An empirical analysis of design choices in neighborhood-based collaborative filtering algorithms," *Information Retrieval*, vol. 5, no. 4, pp. 287–310, 2002.