

AspireNextGen: An Explainable AI-Powered Personalized Career and Skill Recommendation System

GARV DANWANI¹, NIYAZ AHMED²

^{1,2} DDCS, Greater Noida Institute of Technology, Greater Noida, India.

Abstract- Career decision-making has become increasingly complex due to rapid changes in job roles, evolving skill demands, and the abundance of unstructured information available to learners. Students and early professionals often lack personalized, data-driven guidance that clearly connects their current skills to suitable career paths and actionable learning directions. This paper presents AspireNextGen, an AI-powered personalized career and skills advisor designed to recommend career paths, identify role-specific skill gaps, and generate structured guidance for learners. The system employs a hybrid recommendation approach that combines structured role-skill mapping with semantic similarity analysis of user intent. A controlled natural language generation module translates analytical results into clear, step-wise career guidance. Unlike purely conceptual frameworks, AspireNextGen is implemented as a working prototype with a modular backend, normalized database design, and interactive frontend. The system demonstrates strong potential for explainable, scalable, and deployable career guidance in educational and early professional contexts.

Keywords: Career Recommendation System, Skill Gap Analysis, Explainable AI, Educational Data Mining, Personalized Learning

I. INTRODUCTION

The modern job market is characterized by rapid technological advancement, frequent emergence of new roles, and continuous evolution of required skills. As a result, students and early professionals often struggle to identify career paths that align with their current capabilities and long-term aspirations. Traditional career guidance platforms typically rely on static questionnaires or generic roadmaps, which fail to adapt to individual profiles and changing industry demands.

Recent AI-driven guidance systems and conversational assistants have attempted to address this gap. However, many such systems lack structured grounding in skills data, provide limited explainability, or generate advice that is difficult to

validate and act upon. There is a clear need for career guidance systems that are both personalized and transparent.

This research introduces AspireNextGen, an explainable AI-powered career advisor that bridges structured skill intelligence with semantic understanding of user intent. The system focuses on providing clear career recommendations, prioritized skill gaps, and human-readable guidance while maintaining reproducibility and practical feasibility.

II. LITERATURE REVIEW

Career guidance and recommendation systems have evolved significantly over the past several decades, driven by changes in labor markets, educational structures, and advances in artificial intelligence. This section reviews the historical development of career guidance systems, examines contemporary AI-driven approaches, and identifies the research gaps that motivate the proposed AspireNextGen framework.

2.1 Early Career Guidance and Rule-Based Expert Systems

The earliest computerized career guidance systems were predominantly rule-based expert systems, developed to replicate the decision-making process of human counselors. These systems relied on predefined rules derived from psychometric tests, aptitude assessments, and interest inventories. Decisions were typically made through deterministic if-else logic and decision trees.

While these systems offered high transparency and interpretability, they were inherently limited by their static nature. Rule maintenance required continuous expert intervention, and the systems struggled to adapt to emerging job roles or evolving skill requirements. Furthermore, such systems often produced overly rigid recommendations, failing to account for nuanced user intent or interdisciplinary career paths.

2.2 Data-Driven and Machine Learning-Based Career Recommendation Systems

With the rise of data mining and recommender system research, career guidance systems began incorporating machine learning techniques such as classification, clustering, and collaborative filtering. These approaches utilized historical user data, academic records, and occupational datasets to infer career suitability.

Although machine learning-based systems improved adaptability and scalability, many operated as black-box models, offering limited explanation for their recommendations. Users were often presented with predicted career paths without insight into why a particular role was suggested or what skills needed improvement. This lack of explainability reduced trust, particularly in educational contexts where guidance decisions have long-term consequences.

2.3 Educational Data Mining and Skill-Centric Approaches

Educational Data Mining (EDM) introduced a shift toward learner profiling, skill inference, and personalized learning paths. Research in this area emphasized analyzing academic performance, competency development, and learning behavior to guide educational decisions.

However, a significant portion of EDM-based systems focused primarily on course or curriculum recommendation, rather than holistic career alignment. Skill gaps were often identified implicitly, without structured prioritization or clear mapping to real-world job roles. As a result, learners received fragmented guidance that did not clearly connect learning activities to career outcomes.

2.4 Natural Language Processing and Conversational Career Advisors

Recent advancements in Natural Language Processing (NLP) and large language models have enabled conversational career advisors capable of interacting with users through free-text dialogue. These systems excel at understanding user intent expressed in natural language and generating fluent, human-like responses. Despite their strengths, purely NLP-driven systems often lack structured grounding in verified skill or role data. Recommendations may vary across interactions

and risk being generic or overly confident without sufficient factual basis. Moreover, the absence of deterministic evaluation mechanisms makes systematic assessment and academic validation challenging.

2.5 Explainability and Human-Centered Career Guidance Systems

Explainability has emerged as a critical requirement in AI systems applied to education and decision support. Human-centered AI research emphasizes transparency, user trust, and actionable insights over raw predictive performance.

In the context of career guidance, explainability involves not only identifying suitable roles but also explicitly highlighting skill gaps, reasoning pathways, and next steps. Systems that fail to provide interpretable explanations often struggle to gain adoption among learners, educators, and institutions.

2.6 Comparative Analysis of Existing Approaches and Identified Gaps

Table 1: Comparison of Existing Career Guidance Approaches and Research Gaps

Approach Category	Key Characteristics	Limitations Identified	AspireNext Gen Contribution
Rule-Based Expert Systems	Transparent decision rules	Rigid, difficult to scale, outdated roles	Hybrid logic with dynamic skill updates
ML-Based Recommenders	Data-driven predictions	Limited explainability	Explicit skill-gap reasoning
EDM-Based Systems (Educational Data Mining)	Learner profiling, course-focused	Weak career-role alignment	Direct role-skill mapping
NLP / Conversational	Natural-language	Generic or	Structured grounding +

onal Advisors	interactio n	unground ed advice	controlled generation
Skill Assessme nt Tools	Identify competen cies	Lack holistic career context	Integrated recommen dation pipeline

2.7 Research Gap and Motivation

The literature indicates a clear gap between personalization, explainability, and practical feasibility in existing career guidance systems. Rule-based systems offer transparency but lack adaptability, while modern AI-driven advisors provide flexibility at the cost of interpretability and reliability.

AspireNextGen addresses this gap by adopting a hybrid career recommendation framework that integrates structured role-skill intelligence with semantic analysis of user intent and controlled natural language guidance. By grounding recommendations in a normalized skill database and limiting AI generation to explanation rather than decision-making, the proposed system achieves a balance between accuracy, trust, and usability.

III. PROPOSED TECHNOLOGY

AspireNextGen is designed as a modular, API-driven system that integrates user profiling, recommendation logic, and explanation generation. The architecture separates concerns across frontend, backend, data, and AI layers, ensuring scalability, maintainability, and clarity.

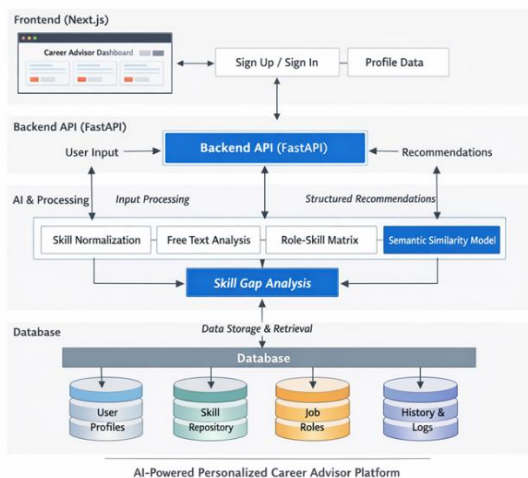


Fig. 1 Architecture

3.1 System Architecture Overview

The system comprises four primary layers:

1. User Interaction Layer: A web-based interface that collects user skills, interests, and intent, and presents recommendations in a structured dashboard.
2. Backend API Layer: A FastAPI-based service responsible for validation, orchestration, and recommendation logic.
3. Data Layer: A normalized relational database storing users, skills, domains, job roles, and role-skill mappings.
4. AI Processing Layer: Modules for semantic similarity computation, skill gap analysis, and guidance generation.

3.2 System Methodology and End-To-End Workflow

The operational workflow of AspireNextGen begins with user onboarding and profile completion. Users provide explicit skills, select interest domains, and optionally describe their career intent in free text. This multi-modal input design reflects realistic user behavior. User-entered skills undergo normalization to remove duplicates, resolve formatting inconsistencies, and map inputs to canonical database identifiers. In parallel, implicit skills are extracted from free-text descriptions by matching known skill vocabularies.

For recommendation, a Role \times Skill matrix is dynamically constructed for the selected domain. Semantic representations of job roles and user intent are generated using TF-IDF vectorization followed by dimensionality reduction. Cosine similarity is used to rank roles based on relevance. Skill gaps are identified by comparing required role skills against the user's current skill set.

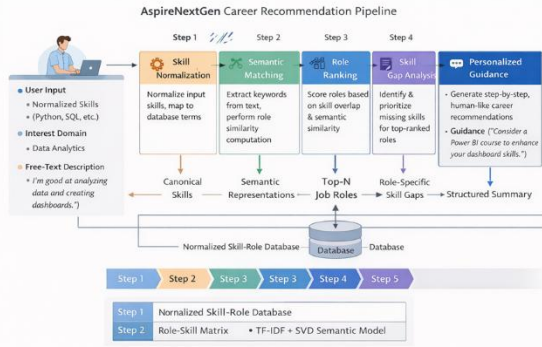


Fig. 2: Pipeline

3.3 AI-BASED GUIDANCE GENERATION

To improve usability, AspireNextGen employs a controlled natural language generation module to convert structured recommendation outputs into step-wise, human-readable guidance. The language model operates strictly on verified system outputs and is used to explain results rather than generate independent recommendations.

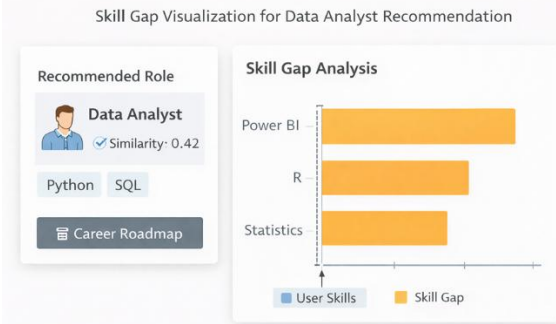
4. RESULTS AND DISCUSSION

The system was evaluated through functional testing and case-based analysis using realistic user profiles. AspireNextGen consistently produced relevant career recommendations, clearly identified skill gaps, and generated structured guidance aligned with user intent. The hybrid recommendation approach reduced ambiguity compared to static systems and improved interpretability relative to purely conversational AI advisors.

This table presents representative user profiles used to evaluate the AspireNextGen system. Each profile reflects a realistic combination of skills, interests, and self-described intent commonly observed among students and early professionals. The corresponding recommended roles demonstrate how the system aligns user capabilities with suitable career paths while highlighting prioritized skill gaps.

User Profile	Declared Skills	Interest Domain	User Intent (Free Text)	Top Recommended Role(s)	Identified Skill Gaps
Profile A	Python, SQL, Excel	Data Analytics	"I enjoy finding patterns in business data and creating dashboards."	Data Analyst	Power BI, Statistics, Data Visualization
Profile B	Java, OOP, Data Structures	Software Development	"I like building applications and solving logical problems."	Backend Developer	System Design, REST APIs, Databases
Profile C	HTML, CSS, JavaScript	Web Development	"I want to create interactive and responsive websites."	Frontend Developer	React, State Management, Performance Optimization
Profile D	Python, NumPy, Pandas	AI / Machine Learning	"I'm interested in working with data and building intelligent systems."	Machine Learning Engineer	Linear Algebra, Model Deployment, MLOps
Profile E	Excel, Communication	Business & Management	"I like working with people and understanding business processes."	Business Analyst	SQL, Data Modeling, Process Analysis

Table 1: Sample User Profiles and Recommended Career Roles



Comparative Analysis of Career Guidance Approaches To contextualize AspireNextGen, a comparative analysis was conducted across rule-based systems, LLM-only advisors, and the proposed hybrid approach.

Table 2: Comparison of Career Guidance System Approaches

Criterion	Rule-Based Systems	LLM-Only Advisors	Hybrid Approach (AspireNextGen)
Decision Logic	Expert-defined rules and decision trees	Generative reasoning from prompts	Structured skill intelligence + semantic matching
Explainability	High but rigid	Limited and non-deterministic	High and interpretable
Personalization	Limited to predefined rules	Conversational but often generic	Skill- and intent-driven personalization
Consistency	Deterministic but brittle	Variable across runs	Deterministic core with controlled generation
Hallucination Risk	None	High due to ungrounded generation	Low (LLM used only for explanation)
Maintenance Effort	High (manual rule updates)	Moderate (prompt tuning)	Moderate (data updates, model tuning)
Scalability	Poor as rules grow	High but difficult to audit	High and maintainable
Practical Deployability	Easy but outdated	Costly and reliability-sensitive	Deployable and resource-efficient
Suitability for Education	Limited adaptability	Requires safeguards	Well-suited for explainable guidance
Future Extension	Low	High with strong constraints	High with feedback loops and embeddings

This table compares three widely observed paradigms in career guidance systems: traditional rule-based expert systems, large language model (LLM)-only advisory systems, and the hybrid approach adopted in AspireNextGen. The comparison highlights differences in explainability, personalization, feasibility, and long-term scalability.

V. CONCLUSION AND FUTURE WORK

5.1 CONCLUSION

This research presented AspireNextGen, an explainable AI-powered career and skills advisor that combines structured skill intelligence with semantic analysis and controlled guidance generation. The system demonstrates how hybrid approaches can deliver personalized, transparent, and actionable career recommendations while remaining feasible for real-world deployment.

5.2 FUTURE WORK

Future enhancements include integrating embedding-based semantic models, incorporating user feedback loops for adaptive personalization, adding quantitative evaluation metrics, and linking recommendations with live learning and labor market data sources.

REFERENCES

- [1] F. Ricci, L. Rokach, and B. Shapira, *Recommender Systems Handbook*, 2nd ed. New York, NY, USA: Springer, 2015.
- [2] C. Romero and S. Ventura, *Educational Data Mining: An Introduction*. Boca Raton, FL, USA: CRC Press, 2010.
- [3] C. C. Aggarwal, *Recommender Systems: The Textbook*. Cham, Switzerland: Springer, 2016.
- [4] Organisation for Economic Co-operation and Development (OECD), *The Future of Education and Skills: Education 2030*. Paris, France: OECD Publishing, 2018.
- [5] R. Guidotti et al., “A Survey of Methods for Explaining Black Box Models,” *ACM Comput. Surv.*, vol. 51, no. 5, pp. 1–42, 2019.
- [6] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman, “Indexing by Latent Semantic Analysis,” *J. Amer. Soc. Inf. Sci.*, vol. 41, no. 6, pp. 391–407, 1990.
- [7] D. J. Power, *Decision Support Systems: Concepts and Resources for Managers*. Westport, CT, USA: Greenwood Publishing Group, 2002.
- [8] P. Brusilovsky and E. Millán, “User Models for Adaptive Hypermedia and Adaptive Educational Systems,” in *The Adaptive Web*, Berlin, Germany: Springer, 2007, pp. 3–53.
- [9] M. L. Rodriguez, R. Patel, and S. Bright, “Skill Gap Analysis Using Data-Driven Educational Models,” *Int. J. Educ. Technol. High. Educ.*, vol. 17, no. 1, pp. 1–15, 2020.
- [10] A. K. Singh and P. K. Gupta, “AI-Based Career Guidance Systems: A Review,” *Int. J. Adv. Res. Comput. Sci.*, vol. 10, no. 4, pp. 45–52, 2019.