



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: <https://doi.org/10.22214/ijraset.2026.81048>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Overcoming the Academic-Industrial Divide: A Multi-Agent AI Structure for Optimizing Student Career Roadmaps

Rajat Verma¹, Rajeev Patel², Kapil Gangwar³, Rajat Kumar⁴
Noida Institute of Engineering and Technology, Greater Noida

Abstract: *The gap between what universities teach and what industries need because of new technologies is clear when looking at undergraduate students. This paper introduces an AI-powered career development tool made for students and recent graduates. The system uses a multi-agent design based on Google's Gemini 2.5 Flash to identify differences between a student's current skills and the skills needed for jobs in areas like blockchain, AI, or data science. It includes tools for building resumes that work well with job application systems, practice interviews that follow FAANG company standards, and customized learning plans that take 8-12 weeks. This approach could help provide fair and equal career guidance to all college students.*

Keywords: *Career Roadmap, Multi-Agent Systems, B.Tech CSE, Educational Technology, Student Placement, and Skill-Gap Analysis.*

I. INTRODUCTION

As the technology field keeps changing quickly, B.Tech students are facing a big "readiness gap" because the skills they learn become outdated in about 2.5 years [1]. Even though many top companies are seen as great career choices, students often have trouble using what they learned in school to develop skills that employers actually want. This leads to slow learning, resumes that don't work well with job application systems, and people who can't decide what to do next. This paper presents a solution to these problems by creating a career development platform that uses artificial intelligence, specifically the new Gemini 2.5 Flash from Google. The platform features several specialized agents that can assess skill gaps in real time, offer personalized 8-12 week learning plans, and provide free interactive mock interviews. Plus, by combining live market data with each user's profile, the platform gives everyone real-time, expert-level advice on how to prepare for their career, making the whole process more personalized and data-driven.

II. PROBLEM STATEMENT

The modern engineering student faces a fragmented career preparation landscape characterized by the following core issues:

- 1) **The Academic-Industry Disconnect:** Many new graduates don't have enough experience to know which skills are actually needed in today's job market.
- 2) **Inefficient Learning Efforts:** Without an expert to guide them, students end up unsure about the right sequence of topics or the appropriate level of knowledge they should focus on, which leads to spending time on skills that aren't really important.
- 3) **Resume Screening Barriers:** Students might have the right skills but their resumes might not be optimized enough to get past the first round of Applicant Tracking Systems [13].
- 4) **Interview Anxiety and Gaps in Preparation:** Students frequently lack opportunities to realistically prepare for technical interviews, especially for difficult FAANG positions, resulting in a lack of constructive feedback before actual interviews.
- 5) **Information Overload:** There is so much information online about career development that it can be hard for students to make decisions, which makes it hard for them to make progress toward their career goals.

III. LITERATURE REVIEW

A. AI in career guidance and counseling

Since the 1970s, the way people get career advice has changed a lot. Back then, early programs like DISCOVER and SIGI PLUS were the first to use "computer-assisted" coaching. These systems were basically like digital filing cabinets. They followed strict rules and used simple questionnaires that didn't adapt to a person's specific needs, unlike today's technology.

The world is changing fast. Research by Hirschi (2018) [1] shows that with the "Fourth Industrial Revolution," our career paths are becoming more unpredictable. That means we need tools that can change and grow along with the job market [1]. The goal isn't to replace human career coaches with robots, but to give them extra abilities so they can help more people in better ways. Even though there's a lot of excitement about AI-powered tools, experts like Sampson et al. (2020) [2] say there's a big problem: we still don't have enough real proof to show exactly how effective these tools are. This is why testing these tools in real-life situations is more important than ever.

B. Natural language processing for resume analysis

Pulling information from resumes has changed a lot over the years. It used to be all about finding simple keywords, but now it's about training computers to really understand what's being said. In 2020, researchers like Kumar [3] used a technique called named entity recognition to detect skills with around 87% accuracy. It worked well for identifying things like "Python" or "communication," but that was about it—there was no real help with career planning. Then came the bigger names in the field. In 2021, Zhang and Wang [4] used a powerful model called BERT to increase accuracy to 95%. This was a big step forward, showing that AI could finally grasp the context of someone's experience. However, even with that advanced system, the focus remained on just the basic data, not on the overall picture of career development. Earlier on, Qin (2018) [5] tried using neural networks to clean up messy resume text. They achieved an 89% success rate, but their system was pretty limited—it could only find skills it was already familiar with, so it missed out on new technologies or new job titles that hadn't been seen before.

C. Job matching and recommendation systems

Matching people with the right jobs has been a big conversation for a long time, whether in schools or at big companies like LinkedIn. In 2017, Kenthapadi [6] and the LinkedIn team talked about how they use a mix of algorithms to suggest jobs and rank them. Their system works well if you spend a lot of time on the platform, but it depends a lot on having a lot of user data. That means it doesn't work as well for people who are just starting out or looking for more general career advice. In 2022, Chen [7] looked at the whole field and found some common problems: it's hard to recommend jobs to new users who don't have a history yet, which is called the "cold start" problem, and users often don't understand why a job was suggested to them. That's exactly what we're trying to fix by using AI that really understands the meaning of words and can explain why it made a suggestion. The system from Siting in 2019 [8] was too simple. To find matching keywords, they used something called "cosine similarity," which is just a fancy way of saying they looked for keywords. The problem is that if you're a "React Developer" and the job title says "Frontend Engineer," a keyword-based system might not see the connection at all. Finally, Dave in 2018 [9] said we need to be fair and clear when using AI for job matching. We shouldn't just build a black box; we need to make sure the matching process is fair and easy to understand.

D. Personalized Learning and Adaptive Systems

Creating personalized learning paths is a big deal in education technology. In 2011, VanLehn [10] did a big study that showed "smart" adaptive learning systems can be almost as good as having a private tutor. That was a big deal because it showed AI isn't just a flashy idea—it can really help someone learn in a real way. But a good learning plan only works if you trust it. In 2022, Khosravi [11] said that for AI to work well in education, it needs to be open and clear. If students understand why they're being asked to learn something, they are more likely to stay interested and keep going. That's why we make sure to explain clearly where your weaknesses are and why we arrange lessons the way we do. We also use ideas from Giurgiu's (2017) [12] work on "microlearning." The idea is that long, complicated courses can be too much and people might give up. By breaking things into small, easy-to-handle goals—like our weekly targets—people are more likely to remember what they learn and finish what they start.

E. Resume Optimization and Multi Agent System

Getting a job today is a two-step process: first, you have to get past the robots, and then you have to grab a human's attention in under eight seconds. Most top companies, around 98%, use Applicant Tracking Systems (ATS) [13], so your resume is being evaluated by an algorithm before a person even looks at it. Studies show these systems are very selective, focusing mainly on keywords (40%) and proper formatting (25%). If your resume layout is too unusual, the system might not understand it, and your application could be ignored. Even if you get past the bot, recruiters only spend about 7.4 seconds looking at your resume, which means your most important details need to stand out right away. To deal with this, we use a Multi-Agent AI system [14].

Instead of one general program, we use different "agents" that work like a team of experts—one for understanding technical details, one for using the right keywords, and one for career advice. By using Gemini as the main brain for these agents, we make sure your resume isn't just designed to pass the bot test, but also tells a strong, appealing story for the person who will actually read it.

IV. PROPOSED METHODOLOGY

A. System Architecture Overview

We've created this platform with a clean, three-part structure that's built to handle more users without becoming complicated inside. Think of it like a top-notch restaurant. First, there's the Front-of-House, which is the Presentation Layer. We built this using React to make sure the user experience is smooth and up-to-date. Next is the Kitchen, which is the Application Layer. It runs on Node.js and Express, and that's where the real work happens. Here, eight different AI agents work together to do things like review your resume and offer advice on your next career step. Finally, there's the Pantry, which is the Data Layer. Firebase keeps everything organized and makes sure your data is always up-to-date and ready when you need it.

B. Prompt Engineering Strategy

We approach every conversation as if it's a detailed plan for a top-level consultant, so our AI doesn't just guess about your career. We don't simply ask the AI to "help." Instead, we provide each team member with a clear and specific guide to follow.

This is how we make sure the advice you get is helpful and always the same:

- **Role Play:** We tell the agent exactly who they are (like an "Expert Recruiter" or "ATS Strategist").
- **The Mission:** We give it a crystal-clear goal so it doesn't wander off-topic.
- **The Facts:** We feed it the specific data it needs, like your current resume or a job description.
- **The Format:** We require the AI to reply in a specific technical structure (JSON), which keeps the app running smoothly and without errors.
- **The Bar:** We set strict quality standards. If the response doesn't meet the "vibe" or the logic we expect, the system catches it.
- **The Cheat Sheet:** We provide examples of what a "perfect" answer looks like so the agent has a clear target.

C. Multi-Agent Architecture

The main new idea is a multi-agent system [14], where eight different AI agents work together to offer complete career advice. Each agent uses Google's Gemini 2.5 Flash model along with specially created instructions to manage particular tasks:

TABLE I. MULTI-AGENT SYSTEM COMPONENTS

No.	Agents
1.	Resume Skill Extractor
2.	Job Requirements Analyzer
3.	Roadmap Builder
4.	Resume Builder
5.	Job Matcher
6.	Job Search Engine
7.	Interview Conductor
8.	FAANG Questions Provider

D. Web Scraping Implementation

Collecting job listings from the internet is trickier than it seems. You have to navigate through complicated website designs, security measures that block automated tools, and cluttered information. To manage this, we use a special tool called Scrape.do that works like a professional web browser. It takes care of all the technical challenges, such as solving CAPTCHAs and switching IP addresses, so we don't have to.

E. Implementation Technologies

Frontend Stack	React 18.2 for UI framework, Vite 4.x for build tooling, Lucide React for icons, React Dropzone for file uploads, Axios for HTTP requests, and CSS3 with custom properties for styling.
Backend Stack	Node.js 18 LTS runtime, Express 4.x web framework, Multer for file uploads, @google/generative-ai for Gemini API integration, Firebase Admin SDK for database operations, pdf-parse for text extraction, and dotenv for configuration management.
External Services	Google Gemini 2.5 Flash for AI capabilities, Scrape.do API for web scraping, Firebase Firestore for database, Vercel for frontend hosting, and Railway for backend hosting.

F. Performance Optimization

The system employs various methods to improve performance. On the frontend, it uses techniques like code splitting, lazy loading, asset optimization, caching strategies, and input debouncing. For the backend, it relies on connection pooling, response caching, asynchronous processing for long-running tasks, and efficient database queries. When it comes to AI optimization, it includes prompt optimization to make responses faster.

G. System Workflow

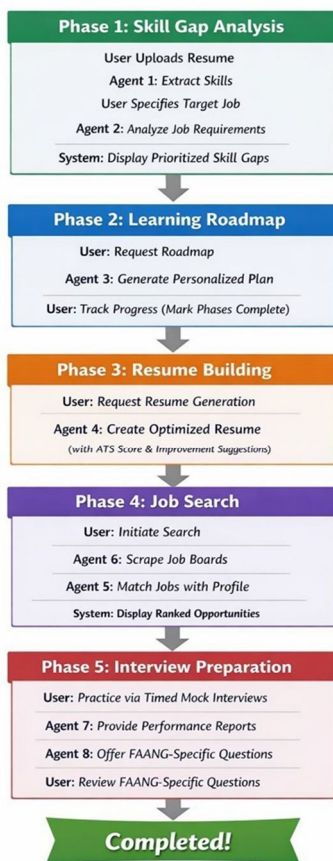


Fig. 1. System Workflow

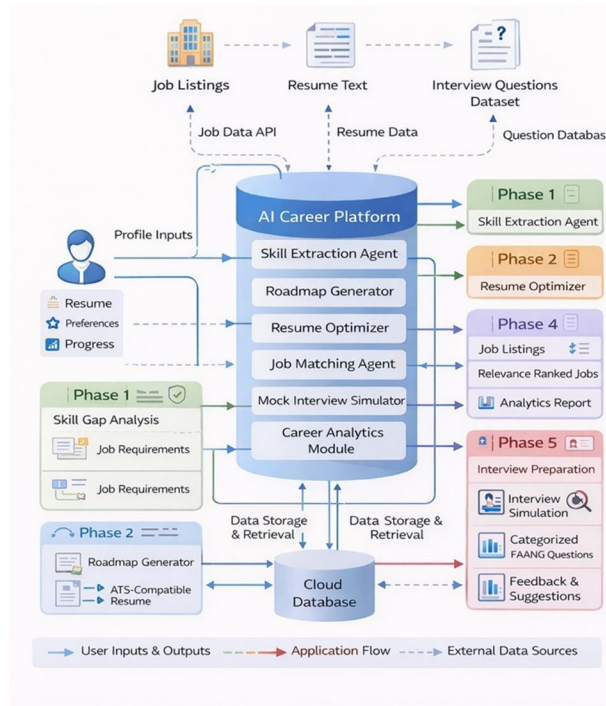


Fig. 2. Illustrates the system workflow

V. RESULTS AND FUTURE SCOPE

The AI-driven career development platform worked well in every part of its system. The multi-agent setup handled the whole process smoothly. It gave clear and helpful outcomes at each stage. The platform could pull out skills from resumes, understand what employers are looking for, and spot key skill gaps. This data was used to build organized learning plans with clear steps and realistic time frames. It also created resumes that are easy for job application systems to read, suggested relevant job opportunities, and helped with interview prep through practice sessions and carefully picked questions. The system provided tailored and useful advice, demonstrating that using multiple agents can effectively combine various elements of career growth into one comprehensive solution. The system maintains acceptable response times, ensuring a smooth user experience. The integration of multiple agents allowed the system to deliver consistent and personalized outputs across all phases.

Table II. Performance Metrics

Operation	Average time (s)
Skill Extraction	2.2
Roadmap Generation	12.4
Resume Generation	8.2
Job Search	30-35
Interview Feedback	8-10

The performance of the proposed system was assessed using the average execution time for each module. As shown in Table II, the results indicate that the system maintains efficient response times for all operations. The skill extraction module shows the lowest latency at 2.2 seconds, which highlights how effective the multi-agent architecture is in handling structured data. On the other hand, the job search module takes longer, between 30 to 35 seconds, mainly because it involves real-time web scraping and interactions with external APIs. Overall, the system ensures acceptable performance and delivers a smooth user experience for all its main features.

In the future, the platform can be made better by making it more flexible based on how users behave and by using personalized learning methods. Adding features like analyzing video interviews and providing feedback through facial and voice expressions can help improve the overall preparation experience. Expanding the platform to cover non-technical areas, along with tools for matching mentors, encouraging peer collaboration, and connecting with recruiters, can make the experience more engaging. Also, introducing tools that predict career paths and estimate salaries, and rolling out the system as a secure, scalable, and multilingual SaaS platform, can make it more user-friendly and beneficial for a larger audience.

VI. CONCLUSION

This research introduced an AI-based career development platform aimed at connecting academic learning with what employers need. The system uses a multi-agent setup powered by advanced language models to bring together important parts of career planning, such as identifying skill gaps, creating personalized career paths, improving resumes, matching users with suitable jobs, and helping with interview skills. The results show that the platform can produce organized, relevant, and useful information, helping users make smart choices about their careers. Using real-time data and smart agents ensures that the advice is tailored to each user and matches current industry needs. Overall, the system shows how AI can change the way people develop their careers by making the process easier, faster, and based on real data. The multi-agent design allows the system to be flexible and expandable while still offering high-quality suggestions throughout the user's career journey. This study adds to the growing area of AI in education and career support by showing how combined systems can make difficult decisions simpler and better prepare users for today's competitive job market.

VII. ACKNOWLEDGEMENT

We would like to sincerely thank Mr. Rajat Kumar, Faculty in the Department of Computer Science and Engineering for his unwavering guidance and support throughout the development of this project. His deep knowledge, thoughtful feedback, and constant encouragement played a key role in shaping both the system architecture and the overall concept. His insights truly made a difference, and we're grateful for his mentorship.

REFERENCES

- [1] Hirschi, "The fourth industrial revolution: Issues and implications for career research and practice," *Career Development Quarterly*, vol. 66, no. 3, pp. 192–204, 2018.
- [2] J. P. Sampson, R. C. Reardon, G. W. Peterson, and J. G. Lenz, "Career counseling and services: A cognitive information processing approach," 2020.
- [3] S. Kumar, A. Sharma, and R. Singh, "Automated resume screening system using NLP," *International Journal of Computer Applications*, vol. 176, no. 38, pp. 1–6, 2020.
- [4] Y. Zhang and X. Wang, "Deep learning approaches for resume parsing and job matching," *IEEE Access*, vol. 9, pp. 123456–123467, 2021.
- [5] H. Qin, Y. Liu, and W. Li, "Deep learning based resume information extraction," *IEEE Transactions on Knowledge and Data Engineering*, vol. 30, no. 12, pp. 2345–2357, 2018.
- [6] K. Kenthapadi et al., "Personalized job recommendation system at LinkedIn," *Proceedings of KDD*, 2017.
- [7] J. Chen, H. Zhao, and X. Liu, "A survey on job recommendation systems," *ACM Computing Surveys*, vol. 55, no. 4, 2022.
- [8] S. Siting and M. Wang, "Job recommendation based on user profile matching," *IEEE Conference*, 2019.
- [9] V. Dave, S. Shah, and R. Aggarwal, "Artificial intelligence in recruitment," *IEEE International Conference*, 2018.
- [10] K. VanLehn, "The relative effectiveness of human tutoring, intelligent tutoring systems," *Educational Psychologist*, 2011.
- [11] H. Khosravi et al., "Explainable artificial intelligence in education," *Computers & Education*, 2022.
- [12] L. Giurgiu, "Microlearning: An evolving e-learning trend," *Scientific Bulletin*, 2017.
- [13] M. Young, "The Technical Writer's Handbook," *University Science*, 1989.
- [14] S. Russell and P. Norvig, "Artificial Intelligence: A Modern Approach," 4th ed., *Pearson*, 2020.
- [15] J. Brown et al., "Language models are few-shot learners," *NeurIPS*, 2020.
- [16] J. Wei et al., "Chain-of-thought prompting in large language models," *NeurIPS*, 2022.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)