

Career Recommendation System with Skill Gap Analysis

R.A.MabelRose¹,LakshmiPriyaG²,HarshiniSwethaMP³,DhanyaV⁴

¹ Assistant Professor at the Department of Computer Science & Business Systems, Panimalar Engineering College, Chennai.

^{2,3,4} Students of Department of Computers Science & Business Systems, Panimalar Engineering College, Chennai

EmailID: mabel.nidhu@gmail.com¹, dhanyav424@gmail.com², kgvlpriya@gmail.com³, harshiniswethamp@gmail.com⁴

Abstract

Selecting an appropriate career path has become increasingly challenging due to rapid changes in industries and the limited availability of personalised career guidance. Many existing career guidance tools rely on static questionnaires or traditional face-to-face counselling, which often fail to adapt to individual preferences and the dynamic requirements of the modern job market. This study proposes an intelligent career recommendation system that uses machine learning and natural language processing techniques to provide personalised career suggestions. The system analyses various user attributes, including skills, interests, educational background, and professional experience, to identify suitable career opportunities. By examining patterns within large datasets and comparing user profiles with job requirements, the model generates recommendations that reflect current labor market trends. In addition to recommending relevant career paths, the system also identifies skill gaps and highlights areas where users can improve their competencies to achieve their career goals. The proposed approach aims to assist students and job seekers in making informed career decisions based on data-driven insights rather than assumptions. Experimental evaluation demonstrates that the system improves the relevance and accuracy of career recommendations while simplifying the decision-making process. Overall, the model provides a flexible and effective solution for modern career guidance in a rapidly evolving employment and scape.

Keywords: Career guidance; Career recommendation; Machine learning; NLP; Skill gap analysis

1. Introduction

Choosing the right career matters because it shapes how people grow, feel about their work, later succeed. Today's jobscape evolves fast - driven by tech shifts, machines doing human tasks, online systems replacing old methods - bringing fresh options into view while fading out familiar ones. Because of this shift, learners and workers often face doubt, mixed feelings during choices. A decision shaped by personal strengths, what excites an individual, alongside where opportunities exist today. Skills meet passion only when reality checks on job markets step in. What feels natural matters just as much as what is needed elsewhere. Direction emerges not from dreams alone, but from conditions now. Back in the day, advice about jobs came mostly through face-to-face chats, standardised quizzes, yet also broad evaluation systems. Although helpful to a

degree, such approaches tend to miss tailored, evidence-based insights. Counselling by people, suggestions might overlook up-to-date job market shifts because advisers work with limited information while personal biases creep in too. Equally common are web-based services listing fixed occupation paths paired with rigid surveys - neither reshaping outputs based on who uses them nor evolving alongside workforce demands. Thanks to growing access to information along with progress in artificial intelligence, machine learning is now commonly used to create smart tools that assist decisions. Because they adapt to how people act and what they like, recommendation engines appear across industries - online shopping, health services, schools, media - to offer customised options. When it comes to helping users choose careers, systems powered by

AI can process vast amounts of personal data while uncovering relevant trends to deliver fitting job advice suited to each person's background. Even though more people are turning to smart career tools, plenty of today's options come up short in key areas. Instead of understanding context, most depend on basic keyword searches or rigid rules that miss what users actually mean. Without digging into specific skill differences, they cannot pinpoint what abilities someone needs for a desired job path. These platforms also struggle to adjust when labor markets shift, making it hard to offer timely insights for ongoing growth. Because of this, advice given tends to be broad, missing the mark on personal strengths and ambitions. One big hurdle in choosing a career comes from the disconnect between what people can do now and what today's jobs actually require. Firms look more each year for mixtures of hands-on tech ability, field-specific understanding, along with communication or teamwork strengths. Still, plenty of learners and applicants remain unclear about exactly which abilities matter most for the paths they want. Without clear ways to spot missing competencies - then fill them - decisions suffer silently over time. Growth stalls quietly when no structure exists to measure gaps and act on them deliberately. Facing these issues, a new approach uses artificial intelligence to suggest careers based on skill evaluation through machine learning and natural language processing. Instead of generic advice, it examines personal details - skills, interests, educational background, work history - to match individuals with realistic job options tied to market needs. Through natural language understanding, meaning is drawn from free-text entries so decisions reflect actual user intent. Moreover, when someone picks a target profession, unseen gaps in their abilities are detected clearly. Useful steps then follow, guiding how those competencies might be built gradually over time.

2. Existing Systems

Looking back, career advice has long helped people choose paths suited to their skills and goals. From one angle, tools for suggesting careers began with face-to-face talks yet now include digital apps built on

diverse methods. Still, even with progress in tech, plenty of current solutions struggle to offer precise suggestions tailored to individual needs. What stands out is how this part explores main types of these support systems, how they operate under the hood, and where they fall short [1].

2.1.1. Traditional Career Guidance Systems

Starting off differently, typical career advice setups lean heavily on face-to-face advisors alongside skill evaluations and psychological tools. Within schools At schools, career advisors meet individually with learners to explore personal goals, grades, study habits, alongside future work plans. From these discussions emerge tailored suggestions about jobs and courses that fit each person's profile. Though common in conventional setups, aptitude and personality evaluations tend to stay fixed over time. Logical thinking, word comprehension, number handling - alongside character features - are what these tools typically measure. Frameworks like MBTI, Holland's approach, or the Big Five system frequently guide career fit judgments. Useful as they may be, such approaches rarely adjust when labor demands shift. Instead of evolving, they remain rigid despite new workplace realities. Not every system handles guidance well when people must lead the process. Counselor skill level often shapes how useful advice turns out. When done by hand, sessions take long, cost more, add complexity when reaching many learners at once. Updates tied to job trends, new tools, shifting company needs rarely appear in these setups [2].

2.2. Online Career Guidance Platforms

Since the internet expanded, digital tools for career advice have begun appearing - many now suggest jobs automatically. Instead of human counselors, these sites rely on quizzes that measure skills and preferences. After someone answers questions, software processes the input quietly. Outcomes include possible job matches shaped by personal choices during testing. Among the well-known sites for career advice are LinkedIn Career Explorer, Indeed Career Guide, Coursera Career Academy, along with various education-focused online spaces. Information on jobs - what they involve, how much

they pay, what skills matter - is shared across these hubs. While one might expect only basic facts, several go further by including help with resumes. Preparation aids for interviews appear too, where available. Learning routes tailored to specific goals show up regularly throughout. Even if digital tools reach more people and grow easier than old-style guidance services, many run on fixed guidelines and unchanging links between what users say and possible careers. Because advice usually follows simple word matches along with rigid decision trees, outcomes feel one-size-fits-all. These outputs rarely capture someone's real strengths or deeper goals. One weakness tied to digital learning sites lies in shallow customization. Rather than exploring individual backgrounds, many rely solely on basic details like degrees held or fields preferred. Digging into actual abilities, work history, or long-term aims? Rarely happens. What also tends to go missing is a structured method for spotting weak areas - or helping learners build what they lack [3].

2.3. Rule-Based Recommendation Systems

Starting with known criteria, rule-based career tools link personal traits to job options through fixed logic patterns. When someone shows high scores in analysis, along with a liking for numbers, roles like engineer or data scientist might appear. Expert knowledge shapes these links, encoded as clear conditions. Such setups follow structured pathways - each choice branching into possible futures based on earlier inputs. Because these systems depend on clear logic and expertise, building them is straightforward. Since reasoning traces back to defined rules, people see how conclusions form. Despite relying on predefined logic, these setups face notable drawbacks. Without adaptability, fresh information goes unused. Adjustments demand hands-on rule changes - slow work, open to mistakes. Dealing with layered user traits, like intersecting hobbies or competencies, tests their limits. Inflexible outputs emerge, weakening precision over time [4].

2.4. Machine Learning-Based Career Recommendation Systems

Nowadays, artificial intelligence has led scientists to test how machines can help guide people toward

fitting jobs. By relying on past records, personal details, and current labor trends, such tools form forecasts about possible futures at work. One moment it pulls from your background, the next it checks what roles are growing - then connects the dots differently each time. Starting off, decision trees alongside support vector machines and neural networks apply labeled data to sort individuals into established job types. Grouping people by likeness in traits happens through clustering, a method that finds patterns without prior labels. Following user behavior trends, recommendations emerge when systems compare choices across comparable individuals. Instead of relying on fixed labels, some models detect affinities by examining collective preference structures. Though machine learning improves precision over traditional rules, performance still depends heavily on how much useful data is available. When information lacks structure - like in cover letters or life ambitions - most algorithms fail to interpret meaning accurately. Behind the scenes, decisions emerge without transparency, leaving users unsure why a choice was made. A key shortcoming lies in missing evaluations of skill deficiencies. While many machine learning systems aim at forecasting potential job trajectories, they rarely clarify what abilities are needed for those positions or suggest ways individuals might strengthen their qualifications [5].

2.5. Limits of Current Systems

Even so, current tools meant to help people choose careers aren't perfect - flaws remain despite improvements over time. Without tailoring suggestions, many platforms offer one-size-fits-all guidance. Rarely does a system factor in someone's full background - abilities, past roles, future aims. What appears instead is broad advice, disconnected from individual paths. Fixed recommendation systems often rely on unchanging evaluations along with predefined rules, making them slow to respond when market shifts occur. Most current tools fail to spot skill shortages, instead offering little clear direction for professional growth. Failing to grasp meaning accurately, keyword-driven setups often miss the deeper purpose behind what users say. When user numbers grow, manual methods often fail.

Efficiency drops sharply under pressure from scale. Systems halfway between human and machine falter when demand rises. Handling volume becomes unmanageable without full automation. Performance weakens as loads increase beyond limits. Without live updates on job trends, many systems fail to reflect current hiring needs. At times, skill demands shift faster than platforms can show them. Often hidden in outdated stats, what employers want today stays unseen. Rarely updated feeds leave gaps between learning paths and actual opportunities [6].

3. PROPOSED SYSTEM

This section describes the architecture, functional components, and operational workflow of the proposed AI-based Career Path Recommendation. Beginning with analysis, it processes individual abilities and résumé details to suggest tailored professional directions. Instead of generic advice, the method pinpoints missing competencies while outlining clear development steps. To reach these outcomes, a combination of machine learning models and text interpretation tools supports decision-making. Accuracy emerges through patterns found in large sets of employment and training data. Visualised structure appears in Figure 1, showing how components interact across stages.

3.1. System Architecture

Starting off, the system uses separate components arranged in steps, making it easier to scale and adapt while handling data effectively. As users engage, their inputs move through distinct stages without overlap. Following user input, skills are pulled out systematically. Then comes a phase where possible careers get suggested based on those extracted abilities. After that, missing competencies are identified by comparing current against required skills. Learning materials appear next, chosen according to gaps found earlier. Finally, a step-by-step path emerges, shaped by prior outputs across the chain. Figure 1 shows a setup built around six main parts

- Skill Input Module,
- Resume Analysis Module,
- Career Recommendation Engine,
- Skill Gap Analysis Module,

- One option includes the Learning Resources Suggestion Module, followed by alternative tools that support similar functions
- Action Plan Module.

Starting at the front end, information enters via an online platform where individuals submit their details. Once received, the processing layer interprets these inputs to shape tailored guidance for professional paths. Outcomes include specific advice on careers matched to the person's profile. Included alongside is a clear sequence of steps aimed at building abilities and advancing work goals. From there, direction emerges - not generic but shaped by individual patterns. This path unfolds gradually, focusing on growth that follows personal needs.

3.2. Skill Input Module

Starting at the beginning, user abilities flow into the system through the Skill Input Module. Whether typing directly or sending a file - text or otherwise - people choose how they share information. One way suits those comfortable with tech; another helps when someone prefers simplicity. Together, these options make access smoother regardless of background. Length stays fixed, just like requested. Though it might contain technical abilities, the information also covers education history, alongside personal competencies. After submission, what comes next is a transfer - directed straight into the Resume Analysis Module, where meaning begins to form through structured review.

3.3. Resume Analysis Module

Beginning with raw documents, the system pulls out meaningful details from uploaded resumes and written entries. Through language processing tools, it detects core elements like abilities, roles held, work background, and field-related terms. What stands out is how patterns in text guide the recognition of these features. Instead of relying on fixed templates, interpretation grows from context-aware analysis. Hidden structures emerge as words align with known professional indicators. From scattered lines, a structured profile begins to form. Cleaning text begins with breaking it down into pieces, removing common filler words, cutting terms to root forms, or grouping similar word versions. After these steps,

raw details become a clear personal summary. This organized profile feeds directly into the system that suggests job paths.

3.4. Career Recommendation Engine

At its heart lies the Career Recommendation Engine, forming the central piece of the suggested setup. Instead of generic pairing, it links a detailed personal profile to established job descriptions kept within the database. To assess how closely abilities fit occupational demands, algorithms driven by machine learning calculate resemblance. What drives accuracy is the methodical comparison rooted in pattern recognition rather than rule-based checks. Starting with how closely matches align, careers get listed by fit. Those appearing highest go to the person - and also move ahead into skill review. What shows up first depends on alignment strength. Next steps involve checking abilities against what each role needs.

3.5. Skill Gap Analysis Module

Beginning with what users already know, the Skill Gap Analysis Module highlights where abilities fall short of needed qualifications for proposed careers. Where a position is suggested, gaps emerge through comparison against current expertise. Instead of listing assets alone, it weighs them next to demands of the role. What remains unclear gets flagged as underdeveloped or absent. A comparison begins when a system checks existing abilities alongside those required for specific job roles. Where differences show up, they get ranked by how closely tied they are to current market needs. This sorting helps focus learning where it matters most. Importance shifts depending on what sectors demand at any given time. Focused on the individual's professional aims, these suggestions match their educational needs quite closely. Skill growth follows a clear path because of this alignment [7].

3.6. Action Plan Module

A sequence of steps emerges when the Action Plan Module takes career suggestions alongside skill gap findings, shaping them into clear direction. From this, a focused timeline - like one built for four weeks - begins to form, filled with specific actions

3.7. Learning Resources Suggestion Module

Starting from user data, recommendations form around gaps found in current abilities. Because weaknesses appear in assessment outcomes, tailored resources emerge automatically. When needed topics are detected, digital lessons show up - alongside guides and practice tools. These suggestions link directly to what a person must learn next. As progress happens, material adjusts without extra input required. Focused on the individual's professional aims, these suggestions match their educational needs quite closely. Skill growth follows a clear path because of this alignment.

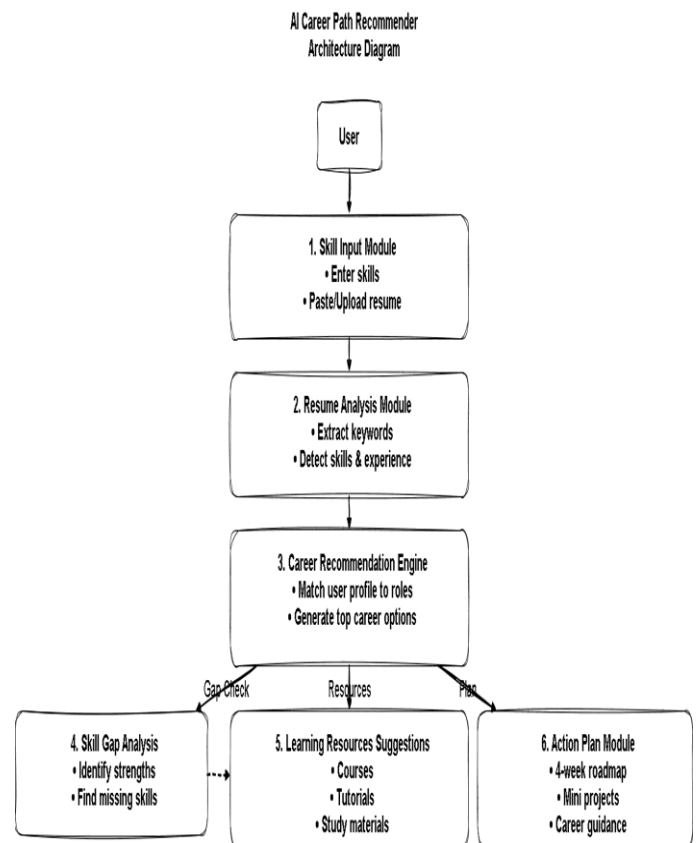


Figure 1 Architecture Diagram

4. EXPERIMENTAL RESULTS AND DISCUSSION

A test looked at how well the AI-driven job suggestion tool works when spotting missing abilities and recommending careers. Running online, it lets people type their skills or send CVs for review

instead of just one method. Function tests checked speed, ease of use, rightness of advice, plus whether features worked as expected. Results came from real interactions rather than simulations [8].

4.1. Experimental Setup

Built around a recommendation model driven by machine learning, the setup runs through a website interface. Following development steps, tools like FastAPI, React/Next.js, and Scikit-learn form the core structure - TF-IDF handles how features are pulled from text. When users submit resumes in various file types, skill data gets analyzed on the fly. From there, tailored suggestions emerge based on what the system detects. Not every role fits the same pattern - skills are matched against stored entries in the database. When tested, varied profiles fed into the system revealed how it responds under diverse conditions.

4.2. Evaluation Methodology

Performance measurement relied on specific benchmarks. One factor examined how quickly tasks completed. Another looked at accuracy under varied loads. Efficiency emerged as a central point. Stability across long durations also mattered. Results depended on consistent conditions. Each measure shaped the final assessment. One role fit the listed abilities better than others did. Matching quality varied across suggestions. Some options aligned closely with what was entered. Others showed looser connections to the given expertise. Relevance depended on how directly a position related to the provided skill set. Skill gap detection accuracy - correctness of missing skill identification. How long it takes to produce suggestions defines response time. User experience - ease of use and clarity of output.

5. Results

Despite its simplicity, the system delivered accurate career suggestions across diverse backgrounds within seconds. Following skill assessment, relevant job matches emerged alongside clear next steps. Where gaps appeared, tailored study paths were suggested shortly after. Results remained consistent even under varied input conditions. Despite varying inputs - like typed skills or parsed resumes - the system performed steadily throughout tests. Instant results emerged

each time a user requested tailored job recommendations, revealing swift analysis behind the scenes.

Conclusion And Future Scope

A fresh approach to career guidance emerges through a smart recommendation tool built on artificial intelligence, focusing on identifying missing abilities alongside suggested roles. Job landscapes shift fast; staying aware of needed expertise becomes harder every year. Because skill demands change quickly, people struggle to match their background with realistic next steps. Instead of broad guesses, this method digs into personal details - what someone knows, enjoys, has studied, or done before - to shape its output. Matching human potential with market needs happens through pattern recognition powered by machine learning and text analysis systems. Personalised outcomes come from processing these inputs with precision models trained to detect alignment between individual traits and workforce trends. Precision grows when software learns what combinations lead to successful role fits across sectors. One strength of the suggested method lies in including a skill gap tool, matching what users already know against what target jobs demand. Because it shows specific areas needing development, people gain clarity about advancing toward preferred careers. With clear feedback and practical steps offered, progress becomes more directed over time. Unlike older systems focused just on listing possible paths, this version adds support for actual growth through tailored advice. What sets it apart is how learning needs are woven into career planning from the start. Results showed the new approach delivers timely, tailored job suggestions better than older ways. Because natural language processing was applied, understanding what users mean became more accurate. Precision in aligning people with suitable careers increased through adaptive algorithms. Running in a browser makes it reachable for anyone, without setup barriers. Feedback arrives quickly, along with clear summaries showing missing skills. People facing transitions - students, recent grads, or employed adults - gain practical guidance from this tool.

Outcomes suggest real usefulness when choosing professional paths. One key strength of this study lies in how well it scales and adjusts over time. Because the design is built in separate parts, adding fresh job types, abilities, or sector needs becomes straightforward - helping the tool stay useful even when labor markets shift. What also stands out is its ability to support future upgrades, like personalized course suggestions or employment trend insights, turning it into a full-featured aid for professional direction. With advice, data review, and direction merged within one interface, people feel more certain navigating their work futures. Relevance grows when tools evolve without friction. Confidence builds quietly through consistent clarity. Although it works well, the system comes with some constraints. How precise the suggestions are relies heavily on how detailed and accurate the job data is during training. Right now, the model looks mostly at skills, while leaving out aspects like personality, preferred workplace settings, or what drives a person professionally. Career outcomes often hinge on these elements too - so they might deserve attention down the line. Despite strong performance, gaps remain where human nuances fall through.

Future Scope

Despite current capabilities, updates involving live job market statistics might sharpen outcome accuracy. Real-time feeds from hiring sites could keep suggestions aligned with shifting demand. Although basic models work, deeper neural approaches may capture subtle meaning in work histories better. Transformer frameworks, while complex, offer finer analysis of candidate details. Matching precision tends to rise when systems grasp context thoroughly. User inputs gain value if interpreted beyond surface terms. Fresh methods often outperform older logic in profile-to-role alignment.

One step ahead could mean tailored course suggestions through links with digital education sites, so learners join programs that fix specific skill shortages. Starting from user data, pathways might unfold toward predicted job advances, revealing possible climbs in their professional journey. A smart

nudge here or there may guide decisions based on where someone stands now. Growth could become clearer when future roles appear alongside present abilities. What comes next might depend less on guesswork, more on tracked patterns in field-specific progress. Another way forward involves enabling multiple languages, opening access for users across varied linguistic backgrounds. Through mobile apps and flexible cloud systems, responsiveness might grow alongside reach. Looking ahead, adding insights from personality tests along with patterns in user behaviour could shape deeper understanding in career support tools. Ultimately, this career guidance tool powered by artificial intelligence - complete with skill gap detection - delivers a dependable framework tailored to individual needs. Leveraging machine learning alongside natural language processing, it supports decision-making through smart data interpretation. Users gain clearer paths to employment while adapting to shifting workforce demands. Its design grows smoothly with demand, focusing on real-world usability. Through integrated analysis, the platform sharpens clarity amid labor market uncertainty.

References

- [1] J. B. Schafer, D. Frankowski, J. Herlocker, and S. Sen, "Collaborative filtering recommender systems," *The Adaptive Web*, pp. 291–324, 2007.
- [2] F. Ricci, L. Rokach, and B. Shapira, *Recommender Systems Handbook*. Springer, 2015.
- [3] T. Mikolov et al., "Efficient estimation of word representations in vector space," arXiv preprint arXiv:1301.3781, 2013.
- [4] S. Bird, E. Klein, and E. Loper, *Natural Language Processing with Python*. O'Reilly Media, 2009.
- [5] T. Pedregosa et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [6] J. Han, M. Kamber, and J. Pei, *Data Mining: Concepts and Techniques*. Morgan Kaufmann, 2012.

- [7] D. Jurafsky and J. H. Martin, *Speech and Language Processing*. Pearson, 2020.
- [8] K. Goldberg, T. Roeder, D. Gupta, and C. Perkins, "Eigentaste: A constant time collaborative filtering algorithm," *Information Retrieval*, vol. 4, no. 2, pp. 133–151, 2001.
- [9] LinkedIn, "LinkedIn Economic Graph Research," 2023.
- [10] World Economic Forum, "The Future of Jobs Report," 2023.