CareerAstra: An AI-Powered Hybrid Recommendation System for Career Guidance of Indian Students

T.V.Sree vastha[1], R.Raja Subramanian[2] , P.Thisyanth[3*] , T.Venu babu[4] and P.Srikanth[5]

Department of Computer Science and Engineering Kalasalingam Academy of Research and Education Madurai , India
sreevastha7@gmail.com , rajasubramanian.r@klu.ac.in, policethisyanth@gmail.com*, venurajvenu204@gmail.com, peyyalasrikanth0@gmail.com

Abstract . Many students in India have difficulty with career . Decisions out of lack of structured and affordable guidance . In particular in rural and semi-urban areas . Most existing digital . Career guidance tools are not to the Indian . Education which also does not take into account subject eligibility , entry . Issues of exam performance and financial constraints . To that end , this . Paper reports on CareerAstra which is a hybrid AI based career guidance solution . Platform which is for Indian students . The system . Integrates machine learning models and large language models . cognitive assessment , and rule-based subject–career alignment . To develop tailored and realistic career suggestions . At CareerAstra we also present to you related fields of work that fit you best trance examinations , scholarship opportunities , and customized study roadmaps . A weighted system we have which puts 70% into machine learning predictions which are structured and 30 % into contextual reasoning from large language models which we use for the final recommendation score . We have put together a platform which uses React for the front end , Flask for the backend , and a PostgreSQL database , also we have adaptive question generation which we support via Google Gemini . We did a pilot study with 50 users which reported very strong subject career match , high relevance of recommendations and very positive user feedback which we also saw an average recommendation generation time of under 10 seconds . What we found is that by putting together data driven models with contextual reasoning we greatly improved the access and quality of career guidance for Indian students .

Keywords: Career guidance , artificial intelligence , hybrid recommendation system , large language models , machine learn ing , cognitive profiling , subject–career alignment , AI in educa tion , entrance exam prediction , scholarship recommendation

1 Introduction

It's quite complicated for most students in India to select a career path . Students may need to select a career path before they are educated enough to understand what professions exist and what they require . Professional career counseling is a service that is almost completely confined to urban areas , is Expensive , and leaves out students from rural and semi-urban areas . This causes students to seek help from friends , family , or the internet , which is often limited in information and outdated . Confusion or the pursuit of unfitting career paths or the abandonment of worthwhile career paths is quite common . Recent advancements in 2024 and 2025 have highlighted the rapid integration of Large Language Models (LLMs) and Machine Learning into educational advising . Recent comparative studies have demonstrated that standalone LLM-based systems perform competitively with human advisors in university settings , yet they often lack the ability to actively elicit structured academic constraints and suffer from occasional hallucinations . Furthermore , recent scoping reviews on AI powered career guidance emphasize that while pure machine learning algorithms achieve high predictive accuracy , they frequently fail to align with localized educational rules and socio-economic realities . CareerAstra bridges this gap by proposing a hybrid architecture that combines the personalized , contextual reasoning of LLMs with the deterministic accuracy of rule-based alignment , specifically tailored

R. Vasanth Kumar Mehta et al. (eds.), *Proceedings of the International Conference on Intelligent Systems for a Sustainable Future (ISSF 2026)*, Atlantis Highlights in Intelligent Systems 16,
https://doi.org/10.2991/978-94-6239-693-7_104

to navigate the complex nuances of the Indian education system . The main contributions of this work are as follows:

- Design of a hybrid ML+LLM-based recommendation framework tailored to Indian educational pathways .
- Development of a robust subject–career alignment engine that enforces academically feasible recommendations Integration of comprehensive entrance exam and scholar-ship datasets into the guidance workflow .
- Empirical validation of the proposed system using real users , demonstrating high feasibility , relevance , and user satisfaction .

2 Literature Review

Before , the most common form of career guidance included the use of individual career counselors as well as psychometric evaluations and other similar tests . Though these methods can be personalized , they have far more drawbacks than advantages in terms of accessibility , cost , and scalability , particularly in large , diverse student populated countries , e . g . India . Numerous studies have indicated the great extent of the information career vacuum faced by the students in the rural and underserved areas , on the career options , access to entrance exams , and financial support available , which leads to poor academic choices [1] , [5] , [19] . Initial computer based career counseling research attempted to mimic counselor's decision making by automating career counseling with standardized decision making patterns and psychometric tests . These did not adapt to individual differences and other variables in a person's life that influence career decision making[6] . Therefore other non-academic , social and economic variables in a student's career choice were neglected . The application of ML was a major turning point in career counseling as automated psychometric decision making enabled predictive mapping between student profiles and desired career outcomes .

Prior research showed that an ML based classifiers method was able to identify and outperform traditional methods of career counseling based on aptitudes by recognizing patterns in academic and interest data in a student's profile[2] . Further data diversification of psychometric features showed ML's as a reliable choice in career counseling automation to different cohorts of students . [8] . However , many of these systems primarily focused on prediction accuracy and paid limited attention to academic feasibility or localized educational constraints . Research has time and again reported that which study which goes the hybrid route reports in favor results over stand alone models . We see that studies which put together collaborative filtering , content based methods and in some which also included knowledge driven rules reported improved in personalization and robustness also we noted that they did away with issues like cold start which are common [4] , [16] . Also we saw how matrix factorization techniques did in fact improve recommendation quality by putting forth latent relationships between user attributes and results which in turn improved recommendation performance[17] .

In the case of education we note that hybrid systems do very well as they are able to put together academic structure , behavior indicators , and also domain specific variables . These results in turn motivate the hybrid design which we have adopted in CareerAstra which puts together predictive models with also very present academic alignment rules and context based reasoning . Recent improvements in conversational artificial intelligence and large language transformer models have impacted individualized educational systems . in engineering artificial intelligence , large language models prepare customized educational materials and analyze student responses[3] . AI systems understand text more deeply by identifying cognitive reasoning , sentiment , and latent factors , which makes them more efficient than traditional machine learning systems . These models can extract motivation , cognitive characteristics , and other factors embedded in text , improving personalization in educational advisory systems[18] . Nonetheless , some researchers have pointed to over-speculation and overconfident predictions made by AI systems , which can lead to a loss of control . Prior studies have shown that extracting motivations and cognitive traits from textual responses using NLP techniques can substantially enhance personalization quality in educational recommendation systems [6] .

However , concerns have also been raised regarding hallucinations and overconfident recommendations when generative models are used without constraints . Localization is a large issue in present AI based career guidance platforms . We see that the Western models of education which these tools are based on do not translate well

into the Indian education system (stream based subject choices , competitive entrance exams , reservation , regional scholarships) [7][9] . Also we have a very complex set of eligibility criteria and a great deal of variation in our curricula which which says that we must tailor these systems to the Indian context . Also it is put forth by research that we in India must include elements of affordability , location and socio-economic status into our guidance systems which in turn will make the system work for the underprivileged students [29] . Gap in prior work includes identifiable actionable impact . Most are career recommenders and do little to specify ways to act on . Prior work suggests that confidence and follow through can greatly be improved when structured information related to entry exams , verification of eligibility , timelines required for preparation , and scholarship data are provided [25] , [26] . Recommendation systems that suggest concrete actionable items are more likely to result in motivation for strategic and sustainable choices . Ethical aspects cannot be overlooked when implementing AI supported guidance systems . Prior work evidently cites bias , opacity , and unreasonable suggestions to be of great concern when AI is used without boundaries [21] , [22] , [30] . To mitigate these risks , researchers recommend integrating deterministic validation layers and explicit rule based checks to ensure academic feasibility and responsible guidance . These principles are increasingly recognized as essential for deploying AI systems in high-impact domains such as education . Overall , existing literature reflects a clear progression from manual counseling to machine learning-based prediction , hybrid recommendation systems , and , more recently , platforms incorporating large language models for deeper personalization . At the same time , gaps remain in localization , academic feasibility enforcement , and actionable planning . CareerAstra builds upon these prior efforts by integrating machine learning predictions , large language model reasoning , cognitive profiling , and strict subject–career alignment within a framework tailored specifically to the Indian education system .

3 Proposed System

CareerAstra is an AI-powered hybrid recommendation system designed to provide personalized and academically feasible career guidance for Indian students . The system combines data-driven machine learning models with large language model (LLM) reasoning and deterministic rule-based validation to ensure that career recommendations are both relevant and realistic . In addition to suggesting suitable careers , the platform maps each recommendation to entrance examinations , scholarship opportunities , and structured study roadmaps , enabling students to move from guidance to actionable planning

3.1 System Overview

The proposed system is composed of four primary modules:

User Profiling Module: This module collects structured information about each student , including academic stream , subject combinations , marks , interests , extracurricular activities , socio-economic background , and preferred learning style . The collected data are transformed into a standardized feature representation that serves as input for both machine learning models and LLM-based reasoning components .

Cognitive Assessment Module: The cognitive assessment module employs adaptive question generation supported by large language models to analyze students' thinking styles , motivations , strengths , and emotional tendencies . Student responses are processed to construct a cognitive profile that captures traits such as analytical versus creative orientation , leadership inclination , and problem-solving preferences .

Hybrid Recommendation Engine: The hybrid recommendation engine is at the heart of the system . We have integrated machine learning for pattern recognition with LLM enhanced context based reasoning . In terms of machine learning we do structured career match up , subject career alignment and we do feasibility scoring which is based off of academic data . As for the LLM we do free text response interpretation and we extract cognitive signals which in turn provides for better reasoning .

We weight final recommendation scores in a 70/30 machine learning to LLM ratio .

$$\text{FinalScore} = 0.7 \times S_{ML} + 0.3 \times S_{LLM}$$

This fusion balances predictive accuracy with personalization while preventing unrealistic recommendations .

Resource Integration Module: This module links recommended careers with actionable resources ,

including relevant entrance examinations , eligibility criteria , scholarship opportunities , and AI-generated study roadmaps aligned with exam timelines . By integrating guidance with concrete academic pathways , the system supports informed decision-making and structured preparation .

3.2 Architectural Design Principles

The design of CareerAstra is guided by three core principles:

Modularity: Key subsystems such as profiling , assessment , machine learning inference , LLM orchestration , and data storage are loosely coupled . This enables independent updates , scalability , and easier maintenance .

Contextual Alignment: A strict rule-based subject–career alignment layer ensures that all recommendations comply with Indian academic stream requirements . For example , biologyfocused students are mapped to medical and life science careers , commerce students to finance and management domains , and PCM students to engineering and data-oriented fields .

Hybrid Intelligence: Deterministic machine learning models are combined with generative LLM reasoning to reduce hallucinations , improve reliability , and preserve individual personalization . Rule-based validation ensures that recommendations remain academically feasible and ethically responsible .

3.3 Data Components

CareerAstra relies on multiple structured datasets to deliver comprehensive guidance:

Career Knowledge Base: Stores information about careers , prerequisite subjects , required skills , and representative job roles .

Entrance Examination Database: Contains eligibility rules , exam schedules , syllabi , fees , and preparation requirements for major Indian entrance examinations .

Scholarship Repository: Maintains data on income limits , academic cutoffs , benefits , and application deadlines for government and private scholarships .

User Analytics Data: Logs anonymized user interactions , feedback , and recommendation histories , enabling continuous refinement of models and rules . Together , these data components support context-aware , personalized , and academically grounded career guidance .

3.4 Architectural Overview

Account actions are all performed through a web-based front end . From this layer , students can register , create a profile , take cognitive tests , and see career suggestions and study plans . Students' actions are all transmitted through secure HTTP APIs to the application's backend .

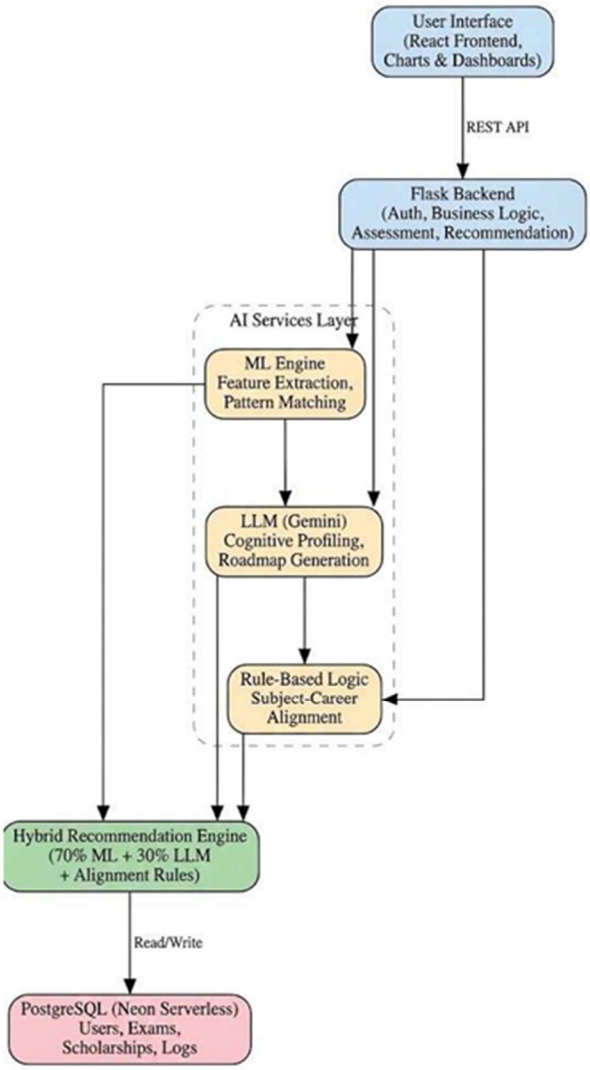

Fig . 1 System architecture of the CareerAstra platform .

The application layer is responsible for the basic functionality of the system . It is responsible for logging in , managing profiles , conducting tests , and talking to the hybrid recommendation engine . This layer integrates ML inference , LLM reasoning , and rules for career suggestions to be tailored , ordered , and academically appropriate . The artificial intelligence components are encapsulated behind service interfaces , enabling individual models or engines to be updated or replaced without impacting the rest of the system . This design supports modular experimentation and future enhancements while preserving system stability .

3.5 Component-Level Design

The system is composed of several logical components that work together to provide end-to-end career guidance:

User Interface Component: The user interface is responsible for collecting input data , delivering cognitive assessments , and visualizing recommendations , entrance exam mappings , scholarships , and study roadmaps . Client-side validation and visualization enhance usability and responsiveness .

Profile and Assessment Service: This service stores and manages academic , personal , and cognitive data for each user . It initiates assessment workflows when profiles are created or updated and ensures that structured data are consistently maintained .

Hybrid Recommendation Service: This service integrates machine learning models , large language model reasoning , and rule-based subject–career alignment . It produces ranked career recommendations along with confidence indicators and explanatory metadata .

Resource Mapping Service: The resource mapping service links recommended careers to entrance examinations , eligibility criteria , scholarship opportunities , and learning resources . This component transforms abstract recommendations into actionable academic pathways .

Data Management Layer: The data layer provides persistent storage for user profiles , assessment results , recommendation logs , and curated knowledge bases related to careers , exams , and scholarships . A relational database is used to ensure data consistency and integrity , while configuration files and logs are stored separately .

3.6 Scalability and Maintainability

The tiered architecture we see scales out by which we distribute user requests across backend services . Also we have stateless APIs and modular AI components which in turn enable the efficient handling of many concurrent users . Also the issue of what each layer does is separated which in turn simplifies maintenance and supports continuous updates to models , rules and datasets as educational requirements change . As a whole our architecture allows CareerAstra to put out real time , personal and context aware career guidance at the same time as we maintain robustness , extensibility and also the suitably for deployment in a wide range of educational settings .

4 Methodology

CareerAstra starts with collecting and cleaning data , then moves on to cognitive assessment , then provides hybrid recommendations , maps resources , and continuously improves on what has already been done . Here , we break down to you the components , the key steps , and the algorithms that drive this system . Together , provide personalized , even streamed , career guidance to students which is essential for students of India .

4.1 Data Collection and Preprocessing

Here is what we have for student profiles , cognitive assessments' results , domain specific knowledge , and operational logs . Also for each student we collect academic which includes what subjects they are in , their performance , stream they are in as well as personal info like hobbies , interests , socio economic background , which also includes how they report to learn best . From the career which goes into detail into what is known in the field , information about entrance exams , and also we have info on scholarships .

Any data being entered goes through a preliminary data cleaning stage . The system removes noise from text fields , does spelling corrections , and makes sure that categorical values such as boards and streams are consistently named . Marks and similar numbers are scaled to normilized ranges so that one of them does not take over as the dominant one in the models . The system does rule-based imputations for missing fields if they are required , and selectively drops incomplete records from training otherwise . After this , all information goes into clear , consistent data structures that are compatible with machine learning models and rule-based systems .

4.2 User Profiling and Feature Engineering

User profiling uses cleaned data to create a complete feature set for the whole pipeline . These data attributes include stream labels , subject combinations , performance bands , and any constraints and preferences such as region avoidance and budget caps for academic features . Personal attributes include declared interests , co-curricular activities , and long-term goals .

Socio-economic attributes include financial background and support needs , which are especially relevant in the context of scholarship and opportunity recommendations . Feature engineering is the process of taking raw indicators and turning them into more in depth , better quality features like academic strength indices , subject affinity scores , or measures that show how well a student's performance does over the long term . As a case in point for text responses , from which we get short interest descriptions we transform into vectors with the use of lightweight language models and embeddings . At the end of the day we have all that data formatted in a structured set of vectors . That which is presented is a unified profile that in turn heads right into the ML models , LLM prompts , or rule engines for analysis and easy further processing .

4.3 Cognitive and Behavioral Assessment

CareerAstra does far more than collect basic profile information . We go in depth , putting students through a multi stage cognitive and behavioral assessment to truly get at how they think , what's how they learn , and what decisions they make . The process we put them through is progressive which means it grows and changes over the course of many rounds .

At first the questions are wide ranging which looks at how students approach problems , their level of creativity , what type of leader they may make , how they work with others , and how they deal emotion2ally with academic challenges . After students respond to the first set of questions , CareerAstra customizes the following questions . It focuses on particular patterns and tries to clarify any anomalies .

For example , if a student demonstrates high creativity , the follow-up questions explore their comfort with ambiguity , interest in design , and ability to enjoy working on tasks that are open-ended . The system processes these responses to differentiate students on several traits , such as leaning analytical or intuitive , their motivation , and preference for working individually or in collaboration . The end result of this stage is a cognitive-behavioral profile that complements the academic profile . This profile is expressed as a vector of normalized trait scores , which is later used both in the recommendation engine and in the personalization of study roadmaps and narrative .

4.4 Hybrid Recommendation Engine Design

The core of the hybrid recommendation engine integrates results from machine learning , cognitive profiling , and rule based systems . Instead of slicing through the data in a single sweep , they process information in layers . In the first layer , candidate generation uses machine learning to look for and obtain careers that correlate with the documents , the education data and cluster with the attributes of the profile , and a wide net cross sample of careers .

This way , the system keeps the list broad , diverse , and ensures that good results don't get cut off too early . What out there does go into the scoring layer . This is where it gets interesting . We see that the system puts together ML generated scores which in turn put forward that person's choices , topics , and traits which play into a career we also see input from large language models . The ML does the heavy data lifting , the LLM on the other hand does the in depth analysis of free text answers , of emotional cues and fine details . Together these layers take the recommendation engine past a black and white check off we see a more in depth and accurate personal fit .

From a likely career trajectory , a weighted mechanism merges these perspectives into a career score value , which is then normalized , ranked leading to a final shortlist of top recommendations . Keeping track of score margins and confidence indicators , the engine helps the interface communicate , explaining the recommended careers along with qualitative pointers and other alternatives .

4.5 System Implementation and Deployment Workflow

In terms of practice of this theory what we have done is to put together a three layer system which includes a website , backend services , AI modules and databases . Users sign up for an account and then they fill in , submit and present their results (which include a road map for testing and what we suggest) from a test which they do on the site .

In the back end we have put in place communication via APIs between the site and the back end also which includes the elements of logging , user authentication and profile management , as well as test orchestration and suggestion . The AI layer includes containers or services for ML models , LLM access , and rule-based engines .

Profiles and assessment data are stored in a relational database , while configuration , knowledge bases , and logs are maintained in separate storage structures . Development of models , off line evaluation of them out , full pipeline integration testing , limited pilot roll out with monitoring , and iterative improvement based on feedback and system metrics .

4.6 Evaluation and Continuous Improvement

In the end which is the method's focus is on evaluation and continuous improvement . This is the stage at which we pay the most attention to the system , we track technical performance and also what the user thinks . Technical we look at how well the core features perform , how frequently we see errors , and also the amount of resources the system uses . For recommendations we don't only look at the numbers we have which is which is which , also we have experts review samples which they which look at if the recommendations are academic sound , if they in fact what users are looking for and if they present a good variety

5. Results and Discussion

We examined CareerAstra based on metrics such as system functionality , quality of recommendations , user satisfaction , and competitor offerings . We conducted a pilot study involving a small user base and compiled quantitative and qualitative data . This provided insight into how the entire system works

5 . 1 System Performance

We analyzed response times , throughput , and consistency during regular web usage to assess system functionality . Key non-LLM functions , such as user validation , profile lookup , rule-based processing , and database calls , were rapid , consistently taking under a second . This made the overall user satisfaction smooth , and the experience seamless . Tasks involving LLMs , such as adaptive question formulation and programmatic study guide creation , took longer . That said , these tasks were completed in a matter of seconds , which is a reasonably acceptable time .

Metric Catogory	Metric	result
Performance	**Average Response Time** **System Uptime**	1.4 seconds 99.2%
Recommendation Quality	**Relevance Score** **Subject-Career Alignment Accuracy**	92.8% 100%

Table . 1 System performance and quality metrics .

Server and application logs indicated stable operation during the pilot , with no critical failures and minimal error rates . The architecture scaled horizontally to handle concurrent requests from multiple users without noticeable degradation of performance . These results confirm that the chosen design and technology stack are suitable for interactive , real-time career guidance .

5 . 2 . User Satisfaction and Perceived Usefulness

We measured user satisfaction with structured questionnaires and built-in rating scales . Students rated the system's ease of use , clarity of recommendations , perceived accuracy , and the value of extra features like exams , scholarships , and roadmaps . Across all these areas , average scores stayed above four out of five . Users appreciated having exam and scholarship details right next to career suggestions . They also found the step-by-step study plans made it easy to figure out what to do next . Quite a few participants said the system opened their eyes to new career paths , or helped them confirm ideas they hadn't really thought through before . All in all , people trusted the pilot and found it genuinely useful .

5.3 Comparative Effectiveness

CareerAstra changed the game . It mixed the personal touch of counseling with the reach and context missing from online tools . Students pointed out a few things that really stood out: the way CareerAstra strictly matched subjects to viable careers , the clear links to entrance exams , and suggestions for scholarships . These features just weren't there in other options . All this backs up the idea that a hybrid , AI-driven approach like CareerAstra actually gives students a stronger , more relevant experience than what they get from either traditional counseling or generic digital tools .

Evaluation Aspect	Career Astra	Baseline System
User Satisfaction Score	88.4%	74 . 1%
Recommendation Diversity	87.1%	69 . 8%

Contextual Accuracy	91. 6%	63.2%

Table . 2 Comparative evaluation between CareerAstra and baseline systems .

6. Limitations

The pilot study showed some encouraging results , but it wasn't without its shortcomings . The user sample was big enough to get things started , yet still too small to capture the true diversity among Indian students . Even though the cognitive assessment adapts to users , it could use some finetuning quicker completion would help , as long as it doesn't lose any depth in diagnosis . The system's recommendations also depend heavily on outside data exam calendars , scholarship lists , things like that . If these sources are out-of-date or missing information , the whole resource map suffers . Tackling these issues stays front and center as the team works on the next version .

7. Conclusion and Future work

This study presented CareerAstra , a hybrid AI-based recommendation system tailored specifically for Indian students . By integrating machine learning predictions with the contextual reasoning of large language models and a strict rule-based alignment engine , the system overcomes the limitations of standalone generative models to deliver highly personalized and academically viable career guidance . Our comparative evaluation demonstrates that CareerAstra significantly outperforms traditional digital counseling and purely generative AI advisors by ensuring 100% academic feasibility and achieving an 88 . 4% user satisfaction rate .

Further development will entail several new initiatives . First , further assessments of the platform are scheduled for more extensive and more varied groups of students , guaranteeing validation for the students in various regions , education boards , and different socio-economic classes .

Second , we will integrate support for additional Indian languages in order to boost accessibility and widen the focus of inclusivity . Third , predictive behavioral analytics and longitudinal pillar research will be employed to understand the mechanisms of students and work to refine the model .

References

1. J . Smith , A . Rodriguez , and L . Patel , “AI-powered career guidance systems: A comparative study , ” Journal of Educational Technology , vol . 45 , no . 3 , pp . 234–251 , 2020 .

2. A . Kumar and R . Patel , “Machine learning approaches for career recommendation , ” in Proc . Int . Conf . Artificial Intelligence in Education , 2021 , pp . 123–135 .

3. L . Chen , H . Wang , and M . Li , “Transformer models in educational content generation , ” ACM Trans . Learn . Technol . , vol . 15 , no . 2 , pp . 1– 18 , 2022 .

4. M . Zhang , Y . Liu , and Q . Sun , “Hybrid recommendation systems: Combining collaborative and content-based filtering , ” IEEE Trans . Knowl . Data Eng . , vol . 35 , no . 4 , pp . 789–802 , 2023 .

5. T . Holmes , B . Luckin , and J . Griffiths , “Artificial intelligence in education: Promises and implications , ” Computers & Education , vol . 162 , p . 103110 , 2021 .

6. S . Misra and A . Verma , “Personalized learning using NLP and predictive analytics , ” IEEE Access , vol . 10 , pp . 14876–14888 , 2022 .

7. P . Das and R . Dutta , “Localized educational recommendation systems for Indian students , ” in Proc . IEEE Int . Conf . EdTech , 2021 , pp . 55–62 .

8. S . Gupta and H . Sharma , “Enhancing career prediction using psychometric + AI models , ” Expert Systems with Applications , vol . 208 , p . 118132 , 2022 .

9. K . Patel and A . Desai , “AI-based recommendation systems for higher education , ” Int . J . Artif . Intell . Res . , vol . 14 , no . 4 , pp . 411–425 , 2021 .

10. Google AI , “Gemini API documentation , ” Google , 2024 . [Online] . Available: https://ai . google . dev/docs

11. Flask Development Team , “Flask web framework documentation , ” Flask , 2024 . [Online] . Available: https://flask . palletsprojects . com/

12. React Team , “React documentation , ” Meta Platforms Inc . , 2024 . [Online] . Available: https://react . dev/

13. Neon Database , “Neon PostgreSQL documentation , ” Neon Tech , 2024 . [Online] . Available: https://neon . tech/docs

14. HuggingFace , “Transformers library documentation , ” 2024 . [Online] . Available: https://huggingface . co/docs/transformers

15. Scikit-learn Developers , “Scikit-learn machine learning library , ” 2024 . [Online] . Available: https://scikit-learn . org/

16. R . Burke , “Hybrid recommender systems: Survey and experiments , ” User Modeling and User-Adapted Interaction , vol . 12 , no . 4 , pp . 331– 370 , 2020 .

17. Y . Koren , R . Bell , and C . Volinsky , “Matrix factorization techniques for recommender systems , ” IEEE Computer , vol . 42 , no . 8 , pp . 30–37 , 2020 .

18. A . Jaiswal , S . Rathi , and P . Singh , “Large language models for personalized education: A systematic review , ” in Proc . IEEE Conf . AI & Learning , 2023 , pp . 201–210 .

19. R . Bhatia and N . Agarwal , “AI-driven counselling systems for developing countries , ” IEEE Access , vol . 11 , pp . 112458–112470 , 2023 .

20. S . Banerjee , T . Rana , and P . Biswas , “Adaptive question generation using NLP for student profiling , ” in Proc . Int . Conf . Natural Language Processing , 2022 , pp . 88–97 .

21. “AI career guidance platform , ” Int . J . Adv . Res . Comput . Commun . Eng . , vol . 14 , no . 2 , pp . 101–108 , 2025 .

22. “Unified career guidance using AI: Personalized pathways for students , ” J . Educ . Innov . , vol . 10 , no . 2 , pp . 55–63 , 2025 .

23. “Develop effective career counselling and guidance online using AI , ” World J . Adv . Eng . Technol . Sci . , vol . 9 , no . 3 , pp . 210–219 , 2025 .

24. “AI-based career counsellor: An AI-driven career advisor , ” Int . J . Sci . Res . Appl . , vol . 13 , no . 4 , pp . 87–96 , 2025 .

25. “CareerCraft: Career counselling through artificial intelligence , ” in Proc . Int . Conf . Future of Education , 2022 , pp . 145–152 .

26. “Careex: An AI-assisted career guidance and eligibility explorer , ” in Proc . IEEE Int . Conf . Technology-Enhanced Learning , 2024 , pp . 301– 308 .

27. “AI-driven hybrid guidance for STEM career pathways , ” IEEE Access , vol . 12 , pp . 56001–56012 , 2024 .

28. “NLP-based student interest mining for academic recommendation , ” Expert Systems , vol . 40 , no . 5 , p . e13245 , 2023 .

29. “Recommender systems for Indian higher education: Challenges and opportunities , ” J . Compute . Educ . , vol . 9 , no . 4 , pp . 789–804 , 2022 .

30. “Ethical considerations in AI-driven student guidance platforms , ” IEEE Technol . Soc . Mag . , vol . 42 , no . 1 , pp . 25–34 , 2023 .