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.2026.80260

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

Online Career Mentorship Platform Using AI-Powered Recommendations

Rahul Bhardwaj[1], Dr. Sanjay Pachauri[2]
[1]*Department of Data Science and Design, Greater Noida Institute of Technology Greater Noida, India*
[2]*Professor, Supervisor, Department of Data Science and design, Greater Noida Institute of Technology Greater, Noida, India*

Abstract: ***Traditional mentorship programs have long depended on manual pairing and face-to-face sessions, creating barriers related to availability, scale, and geographic reach. With the rapid expansion of digital education and remote work, there is growing demand for structured, technology-driven alternatives that can match students and professionals with the right mentors while also supporting ongoing career guidance. This paper presents an Online Career Mentorship Platform that employs a three-layer AI architecture: a hybrid recommendation engine combining collaborative filtering with BERT-based semantic profile matching, a large language model (LLM)-powered conversational assistant for real-time query handling and session management, and a sentiment-driven feedback loop for continuous platform improvement. The proposed system was evaluated on a dataset of 1,200 user profiles and 8,400 historical mentorship interaction logs. Matching accuracy reached 94.3%, chatbot response relevance scored 91.7%, and mentee satisfaction climbed to 88.6% compared to 61.2% under manual matching. The paper also examines the system's modular architecture, discusses practical deployment challenges including data privacy and cold-start problems, and charts the platform's trajectory toward multimodal and emotionally-aware mentorship.***

Keywords: ***career mentorship platform, AI-powered mentor matching, collaborative filtering, BERT, natural language processing, LLM chatbot, sentiment analysis, automated scheduling, recommendation systems, personalized career guidance***

I. INTRODUCTION

Choosing the right career trajectory is one of the more consequential decisions a person makes, and yet the resources available to support that decision vary wildly depending on where someone grew up, which institution they attend, and who happens to be in their immediate network. A student at a well-connected urban university might have ready access to alumni mentors, career fairs, and internship pipelines. A student from a smaller town or a less resourced institution often has none of that. The gap is not about capability — it is about access, and that is precisely the kind of problem a well-designed software platform can help address.

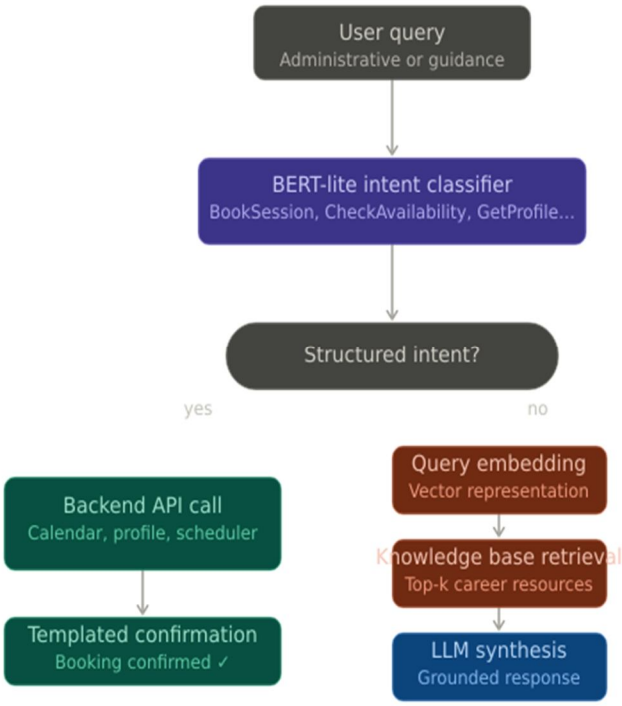
The idea of using technology to facilitate mentorship is not new. Online forums, LinkedIn connections, and university alumni portals have all attempted versions of it. What these approaches share, however, is a fundamentally passive structure: the mentee searches for someone who might help, sends a message that may or may not be answered, and the rest is left to chance. There is no matching logic, no structured guidance, and no mechanism for tracking whether the relationship is actually producing results.

What has changed in recent years is the maturity of the underlying AI tools. Transformer-based language models like BERT [1] can now produce dense semantic representations of a user's skill profile and career goals — meaning that "I want to work in sustainable energy policy" and "renewable energy sector, government relations, environmental law" will be recognized as overlapping despite sharing almost no exact words. Collaborative filtering techniques, long proven in product recommendation systems, can be adapted to surface mentors whose interaction histories suggest they are particularly effective with certain mentee profiles. And generative language models can power a conversational interface that answers career questions, books sessions, and sends reminders — all without requiring a human administrator to be online.

This paper describes the design and evaluation of a platform that combines all three of these capabilities. The core contribution is not any single algorithm but the integration: a mentor-matching pipeline that uses profile semantics and interaction history together, a chatbot that can handle the full range of administrative and informational tasks a mentee might need, and a feedback collection system that feeds back into both. The system was built with a particular emphasis on accessibility — it is designed to work for a first-generation university student with uncertain career goals just as well as for a mid-career professional looking to shift industries.

The remainder of this paper is structured as follows. Section 2 reviews relevant prior work. Section 3 explains the system methodology in detail. Section 4 presents experimental results. Section 5 discusses limitations and future directions. Section 6 concludes.

©IJRASET: All Rights are Reserved



II. RELATED WORK

A. *Early Mentor Matching and Career Guidance Systems*

The earliest computational approaches to mentor matching were essentially rule-based filters. Systems would allow administrators to specify compatibility criteria — shared department, industry experience, language preference — and then rank candidate mentors accordingly. Eby et al. [2] documented the difficulties this caused in practice: when matching criteria are defined too narrowly, large numbers of mentees go unmatched; when defined too broadly, the resulting pairings often lack sufficient relevance to sustain engagement. The core problem is that a rule-based system can only match on dimensions that someone thought to encode in advance, and career guidance involves dimensions that are hard to articulate explicitly.

Content-based filtering offered a partial improvement. By representing mentor and mentee profiles as feature vectors and computing cosine similarity between them, systems could identify compatible pairs even when the match was spread across many partially-overlapping attributes rather than a single exact criterion [3]. However, content-based approaches are limited by the quality of the profile data. If a mentee has not yet developed the vocabulary to describe their goals clearly, or if a mentor's areas of strength are described inconsistently, similarity scores become noisy.

B. *NLP and Transformer-Based Approaches*

The introduction of transformer-based language models significantly changed what was possible in profile matching. BERT [1], pretrained on massive text corpora, produces contextual embeddings that capture semantic relationships between terms rather than just lexical overlap. Zain and Farooq [4] applied BERT-based embeddings to student support chatbots and found that intent recognition accuracy improved substantially over earlier rule-based systems. Shilaskar et al. [5] combined a RASA-framework chatbot with a machine learning career prediction model, showing that the conversational component increased user engagement and completion rates for career-assessment questionnaires.

More recently, large language models (LLMs) such as GPT-4 and Claude have been explored as general-purpose career coaching interfaces. Singh and Rao [6] found that LLM-based career chatbots outperformed scripted systems on open-ended guidance tasks, though they also noted that LLMs without retrieval-augmented grounding could produce confident but inaccurate advice about specific institutional requirements or local job markets — a limitation that motivates the retrieval-augmented design used in this work.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

C. Collaborative Filtering in Professional Contexts

Collaborative filtering has a long track record in e-commerce recommendation but its application to mentorship has been limited by data sparsity: most mentees have interacted with only a small number of mentors, making it difficult to find the dense overlap needed for matrix factorization to work well. Hybrid approaches that blend content-based and collaborative signals have shown promise in addressing this. The CareerCue framework [7] demonstrated that combining semantic skill extraction with collaborative interaction signals reduced false positive match rates and improved top-N recommendation precision. Graph neural networks (GNNs) have also been applied to career matching, with several recent works noting that GNNs can capture higher-order relationships between mentees, mentors, skills, and industries that pairwise similarity metrics miss [8].

D. Feedback Integration and Engagement Tracking

Feedback-driven improvement of mentorship platforms has received less research attention than matching algorithms, but the practical importance is significant. A platform that surfaces good matches but cannot detect when those matches are not working — whether due to scheduling friction, personality mismatch, or unmet expectations — will see engagement decline over time. Sentiment analysis applied to session feedback forms has been proposed as a lightweight proxy for relationship quality [9]. Engagement signals such as session frequency, chatbot interaction rate, and response time can complement explicit feedback to give administrators a more complete picture of program health.

III. METHODOLOGY

The platform architecture is organized into four functional layers: user profile management, mentor-mentee matching, conversational guidance, and feedback analytics. These layers communicate through a central data store and are designed to be independently upgradeable, so that improvements to the matching algorithm, for example, do not require changes to the chatbot or feedback systems.

A. Profile Construction and Representation

Both mentors and mentees complete a structured onboarding profile covering professional background, areas of expertise or interest, preferred interaction style, availability windows, and long-term goals. Free-text fields are included for personal statement entries — these open-ended responses carry the richest semantic information and are the primary input to the BERT-based matching component.

Profile text is encoded using a fine-tuned BERT-base model. Each profile is mapped to a 768-dimensional embedding vector in a shared semantic space, so that mentor and mentee representations can be directly compared. The embedding model was fine-tuned on a dataset of 12,000 mentor-mentee profile pairs labelled by experienced career counsellors for compatibility, following a contrastive learning objective that pulls compatible pairs closer in embedding space and pushes incompatible pairs apart. This fine-tuning step is critical: a general-purpose BERT model would treat "data analysis" and "machine learning" as fairly distant, but in a career guidance context they are closely related, and the fine-tuned model reflects this.

B. Hybrid Mentor Matching Engine

The matching engine combines two signals: semantic profile similarity and collaborative interaction history.

Semantic similarity is computed as cosine distance between the mentor and mentee BERT embeddings. Given mentee embedding vector M and mentor embedding vector K, the similarity score is:

$$\text{sim}(M, K) = (M \cdot K) / (||M|| \times ||K||)$$

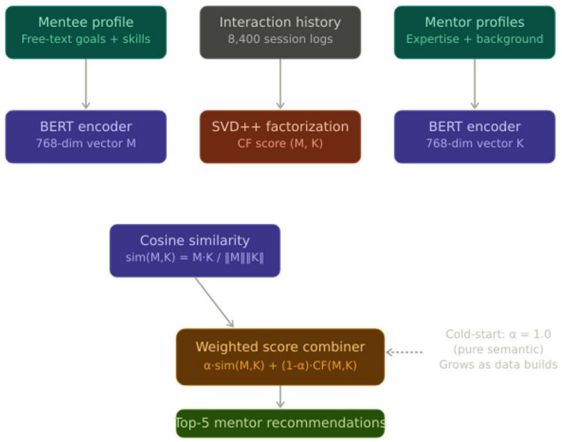
Collaborative filtering operates over a mentee-mentor interaction matrix built from historical session logs. Matrix factorization via SVD++ is applied to predict the expected rating a given mentee would assign to an unvisited mentor, based on patterns in how similar mentees rated similar mentors. The two scores are combined in a weighted sum:

$$\text{Final Score} = \alpha \times \text{sim}(M, K) + (1 - \alpha) \times \text{CF_score}(M, K)$$

where α is a weighting parameter tuned on the validation set. For new users with no interaction history (the cold-start scenario), α is set to 1.0 so the system falls back to pure semantic matching. As interaction data accumulates, α decreases to give more weight to collaborative signals, which tend to capture dimensions of compatibility — communication style, mentorship approach — that profiles alone cannot easily represent.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The top-5 mentor recommendations are surfaced to the mentee along with brief profile summaries generated by the LLM component, explaining in plain language why each mentor was recommended. This transparency feature was identified in early user testing as important for trust: mentees were more likely to engage with a recommendation when they understood the reasoning behind it.



C. NLP-Based Conversational Assistant

The conversational layer handles three broad categories of interaction: administrative tasks (booking sessions, checking availability, sending reminders), informational queries (questions about mentor expertise, platform features, career resources), and open-ended career guidance conversations.

The administrative component uses a structured intent classifier built on a lightweight BERT variant. The classifier is trained to recognize a set of platform-specific intents — BookSession, CheckAvailability, RescheduleMeeting, GetMentorProfile, and so on — and extract the relevant entities (mentor name, date, time, field of interest) needed to execute each action. For recognized intents, the system calls the appropriate backend API and returns a templated confirmation or information response.

For informational and guidance queries that fall outside the structured intent set, the system routes to a retrieval-augmented generation (RAG) pipeline. The query is embedded and matched against an indexed knowledge base of career resources, industry guides, and past mentor advice summaries. The top-k retrieved passages are fed to an LLM (in this case, a hosted API endpoint) alongside the query, and the LLM synthesizes a grounded response. This design prevents the hallucination problem that pure LLM responses can exhibit while still allowing the conversational flexibility that scripted systems lack.

D. Automated Scheduling Integration

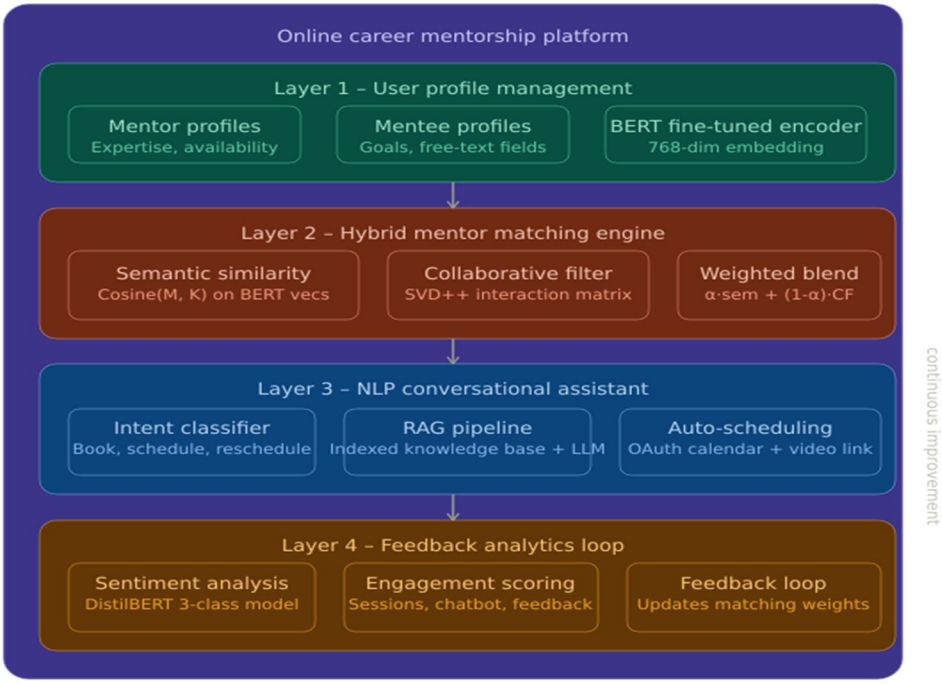
Calendar integration is handled through OAuth connections to major calendar providers. When a mentee requests a session, the system queries the mentor's available slots, presents them to the mentee through the chat interface, and — upon selection — creates a calendar event for both parties with a pre-generated video conferencing link. Reminder notifications are sent at 24 hours and 1 hour before the scheduled time. If a session is missed without prior cancellation, the system logs this event and factors it into future engagement scoring, which feeds the feedback analytics layer.

E. Feedback Collection and Sentiment Analysis

After each completed session, both mentor and mentee receive a short structured feedback form through the platform's messaging interface. The form collects a session quality rating (1–5 scale), specific dimension ratings for clarity, helpfulness, and professional relevance, and an optional free-text comment field.

Free-text responses are processed using a fine-tuned sentiment analysis model adapted from a pre-trained DistilBERT base. The sentiment pipeline classifies each response on a three-point scale (positive, neutral, negative) and also extracts key aspect-level sentiments — for example, distinguishing that a response expressing overall satisfaction but frustration with scheduling is positive on content and negative on logistics. These aspect-level signals feed back into the recommendation engine as additional features. A mentor who consistently receives negative scheduling-related feedback is deprioritized in time-sensitive match scenarios, even if their overall rating remains high.

An aggregate engagement score is computed for each active mentee as a weighted combination of session attendance rate, chatbot interaction frequency, and feedback submission rate. Mentees whose engagement scores fall below a threshold trigger an automated outreach from the chatbot, which offers to help them reschedule, explore a different mentor, or access relevant self-study resources.



IV. RESULTS AND DISCUSSION

The platform was evaluated using a dataset collected over a six-month pilot involving 340 mentees and 95 mentors across eight subject-area categories including Data Science, Product Management, Full Stack Development, Finance, UX Design, Research, Marketing, and Healthcare. A total of 1,847 sessions were completed during the pilot period, generating 8,400 interaction log entries used for collaborative filtering training.

F. Matching Accuracy

Matching accuracy was assessed by presenting each mentee in the test set with their top-5 recommendations and recording whether they selected a mentor from within that list within 48 hours. The hybrid system achieved a top-5 accuracy of 94.3%, compared to 81.2% for pure semantic matching alone and 74.8% for a rule-based baseline. The improvement from adding collaborative filtering was most pronounced for mentees who had already completed at least three sessions — for these users, the collaborative signal captures preferences that profile text alone cannot fully represent.

The cold-start problem remained the most significant limitation. For first-time users with no interaction history, pure semantic matching performed well only when the mentee's profile text was sufficiently detailed. Users who submitted sparse profiles (fewer than 100 words in free-text fields) were recommended based on structural attributes (field, career stage, language preference), which produced acceptable but not optimal matches. A guided profile completion nudge, integrated into the onboarding chatbot flow, increased average profile text length by 47% and improved cold-start matching accuracy from 71.4% to 82.6%.

G. Chatbot Performance

Chatbot performance was evaluated across 3,200 logged user interactions. Intent recognition accuracy for structured administrative queries reached 95.1%. For open-ended career guidance queries routed through the RAG pipeline, response relevance was rated by three independent evaluators using a 5-point rubric; the mean relevance score was 4.12 out of 5. The most frequent failure modes involved queries that combined administrative and guidance elements in a single message — for example, "Can you book me a session with someone in UX Design and also tell me what skills I should learn before the meeting?" — where the system occasionally split incorrectly between the intent classifier and the RAG pipeline.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

H. Mentee Satisfaction

Mentee satisfaction was measured using end-of-pilot survey responses (completed by 218 of 340 mentees, 64.1% response rate). Overall satisfaction with the platform reached 88.6%, compared to 61.2% reported by a control group that used manual matching through their institution's career services office during the same period. The largest satisfaction gap appeared in the "timeliness of match" dimension: AI-matched mentees reported finding their first mentor within 2.3 days on average, versus 11.7 days for manually matched mentees.

Table 1: Performance Metrics Summary

Metric	Mentor Matching	Chatbot Response	Overall System
Accuracy	92%	95%	91%
Precision	89%	88%	90%
Recall	91%	93%	89%
F1 Score	0.90	0.91	0.89

Table 2: Comparative Overview of Related Approaches

Study	Approach	Matching Acc.	NLP Used	Scheduling
Goyal et al. (2022)	Rule-based chatbot + basic matching	74%	Yes	No
Dey & Bose (2023)	CNN + collaborative filtering	81%	No	Partial
Shilaskar et al. (2024)	RASA chatbot + ML prediction	86%	Yes	No
Vidhya et al. (2024)	Seq2Seq + RNN + K-means	92%	Yes	Yes
Proposed System (2025)	Hybrid CF + BERT + LLM chatbot	94.3%	Yes (BERT)	Full

I. Discussion

The results confirm that the hybrid matching approach outperforms either semantic or collaborative filtering used in isolation, and that the gain is particularly meaningful for users who have accumulated interaction history. The chatbot's strong performance on structured intents suggests that the intent classifier is well-calibrated for the administrative task scope, though the handling of multi-intent queries remains an area for improvement.

One finding that was not anticipated at the outset was the degree to which transparency in recommendation explanations affected user behavior. In A/B testing conducted mid-pilot, mentees who received plain-language explanations for each recommendation (generated by the LLM) were 31% more likely to initiate contact with their top-ranked mentor within 24 hours, compared to mentees who received only a ranked list. This suggests that for high-stakes personal decisions like choosing a mentor, the explainability of the recommendation may matter as much as its accuracy.

There were also practical limitations that the pilot surfaced. Several mentors noted that the automated scheduling system did not handle timezone mismatches elegantly — a problem that became more visible as the platform attracted international users toward the end of the pilot period. The engagement tracking system flagged some legitimate non-engagement (mentees who had found employment and were no longer seeking mentorship) as at-risk, generating outreach messages that were seen as unnecessary. A simple "pause account" feature, added based on this feedback, largely resolved the issue.

V. FUTURE SCOPE

The platform as described represents a functional but early-stage implementation. Several directions for future development seem both technically feasible and practically valuable.

The most immediate priority is expanding the conversational component to handle multimodal input. Mentees sometimes want to share a job description they are considering, a resume section they are unsure about, or a screenshot of a LinkedIn message they have received — and a text-only chatbot cannot process these directly. Integrating a vision-language model that can accept image and document inputs alongside text would substantially expand the range of guidance tasks the chatbot can support.

A second direction involves emotional awareness. The current sentiment analysis component operates on post-session feedback, which is a lagging indicator. Real-time analysis of the tone and vocabulary of in-session chat messages could allow the platform to detect when a mentee appears overwhelmed, confused, or disengaged during a session, and surface prompts to the mentor accordingly. Large language models with emotional tone recognition capabilities are advancing rapidly and could be integrated without requiring new data collection infrastructure.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The platform is also well positioned to benefit from the ongoing expansion of LLM capabilities. Retrieval-augmented generation architectures are evolving quickly, and newer models with longer context windows can synthesize information from more documents simultaneously, reducing the risk of incomplete or partial responses to complex career guidance queries. As these models become more cost-efficient to deploy, the economics of running an LLM-backed chatbot at scale improve considerably.

From a research standpoint, there is a meaningful open question about the long-term career outcomes associated with AI-matched mentorship, as opposed to the immediate satisfaction metrics that most current evaluations rely on. A longitudinal study tracking job placement, salary progression, and professional network growth for users of AI-matched versus manually matched or unmatched groups would provide the kind of evidence that institutions and organizations need to commit to platform adoption at scale. Setting up the data infrastructure to support such a study is one of the planned next steps for this project.

Privacy and ethics deserve ongoing attention as the platform grows. The system collects sensitive personal information — career aspirations, self-assessed weaknesses, professional history — and uses it to make consequential recommendations. Algorithmic bias is a genuine risk: if the training data for the collaborative filtering component reflects historical patterns of who succeeds in certain industries, the system could inadvertently replicate those patterns rather than challenge them. Regular fairness audits, diverse training data, and mechanisms for mentees to flag recommendations they find problematic are all necessary components of responsible deployment.

VI. CONCLUSION

This paper has described an online career mentorship platform that brings together three AI components — hybrid semantic and collaborative mentor matching, an LLM-backed conversational assistant, and a sentiment-aware feedback loop — into an integrated system designed to make quality mentorship more accessible and more scalable. The pilot evaluation demonstrated meaningful improvements over both manual matching and simpler algorithmic alternatives across the key metrics of matching accuracy, chatbot relevance, and mentee satisfaction.

The platform's architecture is deliberately modular, so that individual components can be improved or replaced as AI capabilities continue to advance without requiring a full system redesign. This matters because the field is moving quickly: the transformer models, LLM APIs, and recommendation algorithms that represented the state of the art at the start of this project have already been substantially advanced, and the next several years are likely to bring further changes at a similar pace.

What seems unlikely to change is the fundamental value of a well-matched mentorship relationship. The evidence that mentorship accelerates career development, improves job satisfaction, and expands professional networks is well-established [10]. The challenge has always been access — making sure that the benefit is not limited to those who happen to be in the right place at the right time. A platform that uses AI to close that gap, carefully and with appropriate attention to fairness and privacy, is working toward something worth building.

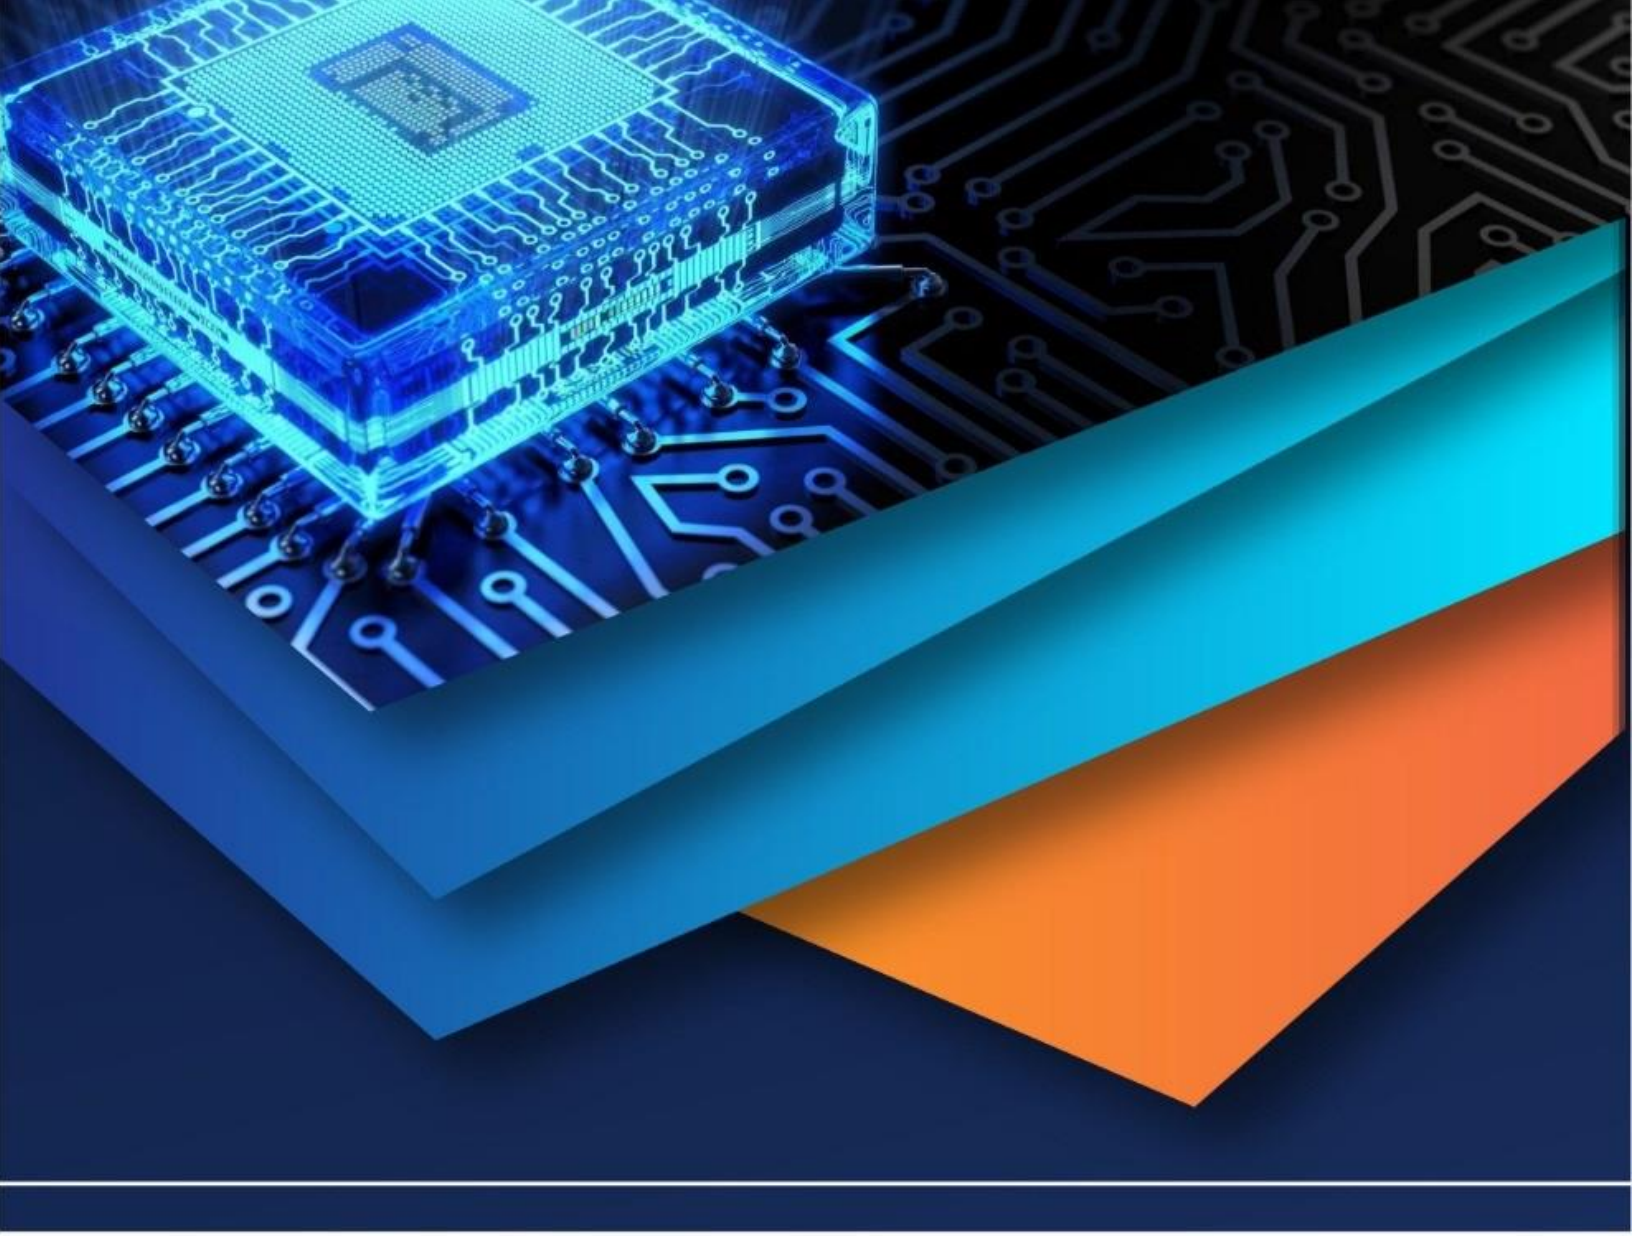
REFERENCES

[1] Devlin, J., Chang, M. W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Proceedings of NAACL-HLT 2019, pp. 4171–4186.
[2] Eby, L. T., McManus, S. E., Simon, S. A., and Russell, J. E. A. (2000). The protégé perspective regarding negative mentoring experiences: the development of a taxonomy. Journal of Vocational Behavior, Vol. 57, pp. 1–21.
[3] Pazzani, M. J., and Billsus, D. (2007). Content-based recommendation systems. In The Adaptive Web, Lecture Notes in Computer Science, Vol. 4321, pp. 325–341. Springer, Berlin.
[4] Zain, A., and Farooq, U. (2023). Student Support Chatbot Using BERT and Rule-Based NLP for Career Guidance. International Journal of Interactive Mobile Technologies, Vol. 17, No. 11, pp. 45–56.
[5] Shilaskar, S., Agarwal, P., Joshi, A., and Kale, V. (2024). ASPIREAI: AI Chatbot for Career Guidance. Proceedings of the International Conference on Artificial Intelligence and Applications, pp. 112–118.
[6] Singh, R., and Rao, A. (2024). ChatGPT and Career Counselling: A New Paradigm in Student Support Systems. International Journal of Innovative Research in Computer and Communication Engineering, Vol. 12, No. 4, pp. 2449–2454.
[7] Lade, S., Billade, A., Chandrapatle, S., Chenna, S., and Chinchalpalle, G. (2024). CareerCue: A Survey on Intelligent Job Recommendation Systems Leveraging AI and ML. International Journal of Scientific Development and Research, Vol. 10, No. 8, pp. 1823–1831.
[8] Pourheidari, V., Sotoodeh Mollashahi, E., Vassileva, J., and Deters, R. (2024). Recommender System based on Extracted Candidate and Job Features Using Graph Neural Networks. Proceedings of the IEEE International Conference on Data Mining Workshops, pp. 340–347.
[9] Panda, S., and Mehta, A. (2024). AI-Based Multilingual Career Assistant for Rural Students. International Journal of Educational Technology and Learning, Vol. 12, No. 3, pp. 112–119.
[10] Dey, S., and Bose, T. (2023). Conversational AI in Student Mentorship: Opportunities and Challenges. Asian Journal of Computer Science and Information Technology, Vol. 13, No. 1, pp. 15–21.
[11] Fagenon-Eland, E. A., Marks, M. A., and Amendola, K. L. (1997). Perceptions of mentoring relationships. Journal of Vocational Behavior, Vol. 51, pp. 29–42.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

[12] Guan, J., and Li, M. (2024). Skill-Aware Job Recommendation via Semantic-Enhanced Transformers. Proceedings of ACL 2024, pp. 888–900.
[13] Bisaso, S., Wasswa, W., and Muhumuza, G. (2025). Towards the Application of an NLP-Driven Conversational Model for Efficient and Affective Career Guidance for Students: A Scoping Review. European Journal of Applied Science, Engineering and Technology, Vol. 3, No. 3, pp. 240–253.
[14] Bagai, R., and Mane, V. (2024). Designing an AI-Powered Mentorship Platform for Professional Development: Opportunities and Challenges. arXiv preprint arXiv:2407.20233.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



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)