Grenze International Journal of Engineering and Technology, June Issue

PathFinder AI: An Agentic AI Career Guide and Technical Interview Assistant

Shreya Singh[1], Ami Khatri[2], Bhavy Shah[3], Akshad Jain[4], Himanshu Gohil[5] and Nikita Raichanda[6]
[1-5]Student, Artificial Intelligence and Data Science, Dwarkadas J. Sanghvi College of Engineering, Mumbai, India
[6]Assistant Professor, Artificial Intelligence and Data Science, Dwarkadas J. Sanghvi College of Engineering, Mumbai, India
Email: shreyasingh16april@gmail.com, amikhatri100704@gmail.com, bhavysshah40@gmail.com, akshadjain25@gmail.com, himanshugohil405@gmail.com,Nikita.raichada@djsce.ac.in

***Abstract*— The current job market environment is one that lacks cohesion among career resources, generalist tips, and lack of personalized, efficient career preparation solutions. Traditional job websites and inflexible resume screeners neglect to consider time as a crucial variable, leaving individuals with disorganized learning material and without any clear roadmap to prepare for interviews. In this paper, we present PathFinder AI – an innovative career platform based on Agentic AI technology and intended to provide affordable, personal guidance. PathFinder AI consists of a three-layer architecture system, using the Node.js and Express frameworks combined with large language models through APIs offered by OpenAI and Gemini. An important contribution of this project is the Deadline Aware Scheduling Algorithm, which reverse-engineers a preparation curriculum anchored to the user's specific interview date. Through semantic skill-gap analysis — comparing vectorized user resumes against target job descriptions — the system pinpoints high-priority competencies and assigns curated learning resources in real time. The platform further incorporates five core agentic modules: a voice-driven AI Mock Interviewer built on the Whisper Medium speech recognition model, a Semantic CV Evaluator, an Automated Cold Outreach tool, and a Real-Time Job Search engine. Prompt engineering with bias awareness mandates an evaluation framework based on skill assessment alone without consideration of any demographic data points during AI evaluations. Evaluation outcomes demonstrate that the PathFinder AI turns the usual disorganized, stressful process of preparation into a well-organized process that enhances candidate preparedness.**

***Index Terms*—Agentic AI, Career Guidance System, Large Language Models (LLMs), Whisper AI, Resume Parsing, Skill Gap Analysis, Generative AI, Natural Language Processing (NLP), Time-Bound Scheduling.**

I. INTRODUCTION

A. Background

The past decade has seen a tremendous change in the area of recruitment and career growth. In the past, job preparation was a very simple process because a job seeker would acquire a degree, prepare a common resume and place applications to the job advertisements posted in the paper or permanent job boards. However, the new digital economy has developed a complex and fiercely competitive workplace where skills are formed at a very high rate.

Grenze ID: 01.GIJET.12.2.61_24

Assessments in the present industry show that the average length of applicability of a technical skill has dropped to under five years, necessitating consistent education and skill development. There is still a wide gap between finding a job opportunity and getting ready to get it, despite the presence of such sites as LinkedIn, Indeed, and Glassdoor with millions of job posts available, not to mention Massive Open Online Courses (MOOCs), which provide inexhaustible educational opportunities. Job seekers, particularly students and new graduates, often feel inundated by information and cannot find a way to make all the information available to them into a single preparation strategy. Concurrently, the Artificial Intelligence (AI) world has evolved to Generative and Agentic AI. The large language models (LLMs) have demonstrated the ability to not only be able to understand the text but also to reason, plan, and execute complex instructions. This breakthrough in technology provides us with a golden chance to overcome the career preparedness bottleneck. The PathFinder AI is a result of EdTech and Generative AI collaborating together to use the powers of contemporary LLMs to change the messy interview preparation process into an organized, evidence-based process.

B. Problem Statement

Despite the abundance of job-searching sites present online, there are many obstacles faced by people seeking jobs that are not covered by those sites. Some of the problems may be classified as follows:

Information Overload and Analysis Paralysis: Job aspirants who apply for IT jobs like Software Development Engineer (SDE) or a Data Scientist faces an overwhelming number of study subjects – from Data Structures, System Design, Frameworks, Databases to Behavior rounds. Without a prioritized roadmap, this abundance backfires: candidates invest more cognitive energy in deciding what to study than in the act of studying itself, a phenomenon commonly referred to as analysis paralysis.

One-Size-Fits-All Guidance: The vast majority of career preparation solutions currently rely on preset learning programs which presuppose a common starting point of each student. Such solutions do not take into account the differences in individual knowledge levels, such as someone being an expert in Python and knowing nothing about statistics. As a result, a person invests time into mastering those things that he or she already knows, instead of dealing with more impactful knowledge gaps. Modern ATS does not address the problem and focuses merely on keywords.

Absence of Time-Constrained Planning: One of the major drawbacks of most learning programs is that they have no consideration for deadline management. The typical Full Stack Development curriculum roadmap will be developed without a specific time limit and thus cannot facilitate any adjustment within the content in case of a few days left until an interview with the candidate.

Tool Fragmentation Across the Hiring Pipeline: In today's day and age, there is an eclectic range of tools available to one seeking employment; however, these tools are compartmentalized in such a way that information exchange cannot occur between the various phases of recruitment: For instance, no information gathered during resume gap analysis is ever passed to the study planner.

C. Motivation

The driving force behind PathFinder AI is an inequity in career coaching that stems from a disparity in fairness. Whereas there exists a lot of merit in human career coaches, the service that they provide can unfortunately be afforded by only a small portion of the student body and recent graduates. Through the creation of an Agentic AI powered decision engine that incorporates all of the wisdom of a career advisor into its programming, we seek to democratize career coaching. As equally crucial is the automation of an intellectually strenuous task of reconciling syllabi with resumes.

D. Research Contributions

The article contributes the following aspects to the field of educational technology driven by AI:

- Agentic Coordination Model: We propose a modular agentic system where autonomous agents formulate, implement, and refine career guidance strategies instead of merely responding to user requests.
- Time-Bound Curriculum Rescheduling: We propose a groundbreaking rescheduling technique that periodically adjusts the learning schedule according to the proximity of the user's interview, transitioning from extensive learning to crash courses as the deadline draws near.
- Secure Resume Management: Resume data will be stored in memory via a RAG model where no sort of storage mechanism is used. This ensures complete data security for the users.
- Comprehensive Career Preparation Platform: Our platform integrates resume assessment, roadmap formulation, AI-assisted mock interviews, outreach, and job searching into a seamless process.

II. LITERATURE REVIEW

The development of automated recruitment has changed the focus from less sophisticated keyword matching to sophisticated semantic understanding.

In [1], Jagwani et al. presented a resume-job matching system that uses the Latent Dirichlet Allocation (LDA) and Natural Language Processing (NLP). Their method identified thematic clusters such as skills and education background reaching an accuracy of nearly 82 percent in predicting a fit between candidates and jobs. However, it was highly sensitive to the changes made in the hyperparameters and struggled with the disorganized, real-life resumes with complex formatting. Li et al. [2] examined more risks associated with automated hiring by analyzing the potential of word embeddings to entangle and support national origin-related biases. The study revealed using the Word Embedding Association Test (WEAT) that pre-trained embeddings are always associated with certain cultural groups with negative traits, and that resumes with particular names or institutional affiliations are rated worse than others. The significant limitation of their results is that the bias is embedded at the level of vectors before the start of classification.

Harris [3] conducted research on age-related discrimination in AI-facilitated recruitment processes, where the use of the candidate's year of graduation and total experience as numerical indicators discriminates against older candidates. Kumar and Grosz [4] took the discussion further and spoke on larger recruitment recommendation systems rather than the individual resume assessment. Their paper noted that current disparities are exacerbated by exposure bias in which historically underrepresented groups are introduced to fewer valuable opportunities. Although metrics of technical fairness do exist, they rarely align well with legal nondiscrimination conceptions, and pose compliance burdens to any organization that applies such metrics.

Armstrong et al. [5] conducted assessments of GPT-family models, and they found that even state-of-the-art LLMs are characterized by significant racial and gender biases in providing qualitative feedback.

Wilson and Caliskan [6] extended the analysis to intersectional demographic bias and showed that retrieval-based screening systems have adverse effects on candidates occupying the intersection of race and gender with the main negative influence being on Black women. Their findings reveal that the analysis of demographic features alone does not reflect the cumulative disadvantage of people with compounded marginalized identities.

Everitt et al. [7] proposed using the fine-tuned Large Language Models like FLAN-T5 to produce the exclusionary language in job advertisements, which demonstrates that AI can minimize bias. Samadhiya et al. [8] studied conventional machine learning pipelines and found that choices during feature engineering such as including employment gaps may inadvertently propagate occupational segregation. The study has shown that information included in the engineering of features such as the inclusion of employment gaps or the use of gendered pronouns could bring in occupational segregation trends unwittingly.

In the study by Deshpande et al. [9], Fair-TF-IDF was suggested as an effective algorithm designed to reduce lexical demographic bias at the preprocessing stage of the process. An et al. [10] proposed a model of systematic assessment to identify racial and gender biasness in their hiring mechanisms through the use of LLMs. They could evaluate the effect of modifications in name or pronoun on ranking results in a systematic way by writing synthetic resumes with only deviation in demographic factors. According to Albaroudi [11], AI-based recruitment is an area where past research studies have been consolidated, and he found out that addressing issues related to algorithmic bias requires a multi-faceted approach.

A Neural Fair Collaborative Filtering framework of job recommendation systems suggested by Islam et al. [12] incorporates fairness-related loss functions into the training of the model. The approach worked well in increasing the equity of exposure and maintaining relevance in recommendations.

Devaraju [13] provided an overview of NLP in recruitment noting that despite the fact that semantic alignment is improved by transformer models, these models increase the level of black-box obscurity. Finally, Zhang [14] studied the issue of algorithmic bias in job recommendation systems through user audits and found out that similar candidate profiles are provided with different job recommendations simply because of gender indicators.

It is clear from these studies that the problem of algorithmic bias does not just occur during the selection process, but is present even earlier, at the stage of exposing applicants to potential employment opportunities.

III. SYSTEM ANALYSIS

The system analysis phase outlines the functional and non-functional requirements that dictate the PathFinder AI structure. This phase was important in determining the scope of the study.

A. Functional Requirements

The system has been developed using specific operational modules as detailed in Table I.

TABLE I. FUNCTIONAL REQUIREMENTS SPECIFICATION

Req ID	Module	Description
FR-01	User Interface	Resume Processing: The system needs to accept PDF documents and guarantee uniform parsing of document formats.
FR-02	User Interface	Objective Definition: Users need to specify their desired Job Role (e.g., SDE) and the deadline for interviews.
FR-04	Backend	Data Sanitization: A preprocessing stage needs to eliminate non-standard symbols and extraneous whitespace to avoid AI hallucinations.
FR-05	Core AI	Identifying Skill Gaps: The system should evaluate the resume vector in relation to job criteria to identify particular lacking skills.
FR-06	Core AI	Scheduling That Considers Deadlines: The algorithm needs to effectively allocate learning modules, emphasizing high-value subjects when time is limited.
FR-07	Content Gen	Resource Collection: The system should obtain curated educational resources (documents, tutorials) for every recognized gap.
FR-09	Chat Module	Contextual Uncertainty Resolution: A chatbot should maintain the context of the created roadmap for subsequent inquiries.

B. Non-Functional Requirements

- Latency (NFR-01): The entire creation of a roadmap spanning 7 days should ideally be accomplished in 15-30 seconds with the use of Gemini 2.5 Flash.
- Scalability (NFR-02): Non-blocking IO should be used on the Express.js server to prevent overloading of the application due to simultaneous requests.
- Privacy (NFR-03): Design for privacy: Resumes belonging to users will only exist in working memory and not in the database in permanent form. Also, vector embeddings will be anonymized.
- Ethical AI (NFR-04): Prompt engineering should clearly inform the model that personally identifiable information, such as names, gender, and addresses, should be ignored.

C. Feasibility Analysis

Feasibility studies were performed before development:

- Technical Feasibility: The integration between OpenAI and Pinecone APIs in combination with Node.js has been extensively studied. It is due to the availability of libraries that can parse PDFs (‘pdf-parse’) and hence extract text data from various formats.
- Operational Feasibility: No specific requirements regarding hardware exist, as a web browser is all the user needs. The user interface design is highly intuitive and easy to use with minimal training for users.
- Financial Feasibility: The cost incurred by making API calls (Gemini/OpenAI) is negligible as the project is merely a prototype. The "serverless" architecture ensures lower maintenance costs.

IV. ARCHITECTURE AND DESIGN OF THE SYSTEM

PathFinder AI has a well-built architecture that allows managing a significant number of requests while also maintaining data privacy.

A. Architecture of Data Flow

Data stream organization to achieve efficiency and security:

- Input: A user uploads a PDF resume and fills in certain parameters: Job Position, Due Date via Frontend.
- Processing: The file stream gets uploaded on the Node.js backend, and pdf-parse parses the document, while RegEx filters do the normalization.
- Vectorization (Pinecone): This is a breakthrough feature. The cleaned text gets segmented into semantic parts: Skills, Experience, Education. Those are further processed using an embedding model to get high-dimensional vectors, stored in Pinecone Vector Database, thus making it possible for the AI to perform semantic searches and maintain "long-term memory" of the user's profile without processing the PDF file each time the user makes a request.
- Reasoning (Gemini 2.5 Flash): The backend forms an organized prompt, consisting of the vectorized context of the user and a job description, which then gets passed to the Google Gemini 2.5 Flash API.
- Output: The AI generates a JSON file with skill gaps, roadmap, and resources recommendations, displayed in the Frontend.

B. Algorithm Design: Scheduling with Deadline Awareness

The core rationale of creating roadmaps is guided by an algorithmic approach to adjust the course intensity level. The algorithm is also provided below in pseudo-code form. Such reasoning ensures that a person who has only two days remaining does not feel burdened with complex multi-hour courses but rather receives concise cheat sheet documents and quick interview questions.

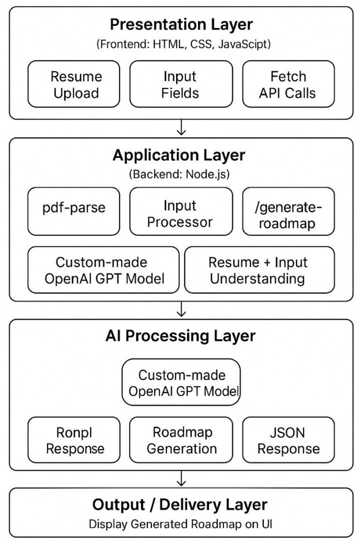

Fig. 1. PathFinder AI System Architecture.

Listing 1. Deadline-Aware Scheduling Logic

```
function generateRoadmap(gaps, days):
    curriculum = []
    if days < 5:
        priority = "HIGH_YIELD_ONLY"
        for skill in gaps:
            topic = fetchInterviewQuestions(skill)
            curriculum.append(topic)
    else:
        priority = "DEEP_DIVE"
        for skill in gaps:
            topic = fetchFullCourse(skill)
            curriculum.append(topic)
    schedule = distribute(curriculum, days)
    return schedule
```

C. Scheduling Strategy Justification

The suggested scheduling approach could be mathematically formulated as a problem of allocating resources based on priorities. We assume that $G = \{g_1, g_2, ..., g_n\}$ are the skill gaps recognized and the learning utilities U_i and time taken to learn c_i. The objective is to optimize the total of learning utilities for the mentioned skills subject to the preparation duration T_r^{bm}:

$$\text{maximize } \Sigma\, w_i \cdot U_i \cdot x_i \quad \text{subject to } \Sigma\, c_i \cdot x_i \leq T\_rem,\; x_i \in \{0,1\} \quad (4)$$

where xi is the decision variable and wi is the time-dependent decaying weight factor used in prioritization.In case of $T_r^{bm} < 5$ days, the system utilizes a greedy approach, choosing topics based on their utility in descending order. In other words, it solves an optimization problem, where the constraint operates like the maximum weight constraint in the 0/1 Knapsack problem. If $T_r^{bm} \geq 5$ days, the constraint no longer plays its role, and the system enters intensive mode with the aim of maximizing concept depth. This approach is grounded by the diminishing returns effect in spaced repetition studies, according to which, in case of extremely constrained time resources, broad but shallow knowledge coverage is better.

D. Component Level Design

- Resume Parsing Engine: The module is composed of the ETL (Extract-Transform-Load) pipeline. The module parses PDF binary streams, extracts text, and performs cleaning actions to remove any artifacts that can confuse the LLM.
- AI Orchestration Module: This module manages interactions within the Intelligence layer. The Gemini instance, API Keys, and "System Prompt" ("You are a specialized technical interviewer...") are managed by this module. The CoT prompting method is used here to ensure that the resume gets evaluated before roadmap creation.
- Pinecone Integration: Our approach differs from stateless architecture in the following: thanks to Pinecone integration, our system has the ability to "remember" information regarding the user. For example, when the user asks a follow-up question via the chatbot ("Explain how my project relates to this position"), the system searches for relevant parts of the resume in Pinecone

V. Core Product Features

PathFinder AI is a system that integrates five distinct agentic modules to embrace the overall recruitment process

A. AI Interview Simulator with "Hire/No Hire" Decision

This is a simulation of a high-pressure technical interview conducted via the voice-based application.

- Operation: The application generates audio questions pertaining to the resume history of the user. The verbal response generated by the user is then transcribed to text and delivered to the server.
- AI Role: The Medium model of STT has been implemented for this process because it is better equipped to deal with other accents and noise as compared to other smaller models.
- TTS: The TTS feature adds more functionality to the dialogue loop in the form of enunciation of questions posed by the AI.
- Conclusion: This process produces a unique output – either "a hire" or "no hire" along with a score (e.g., 7/10).

B. Agentic Pathway Creator with Alerts

Mitigation strategy to address the issue of poor time management during independent learning:

- Deadline-Conscious Algorithm: This algorithm is the opposite of the curriculum design. If T_deadline < 5 days, then this mode will trigger the Crash Course Mode, where the focus shifts to high-frequency interview questions rather than theory.
- Notification System: For compliance, it uses daily reminders such as "Day 2: Complete Linked List Module," using browser notifications.

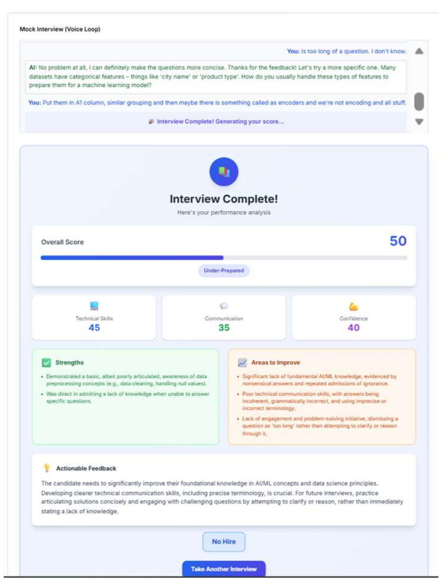

Fig. 2. AI Simulation Interviewer evaluating a candidate

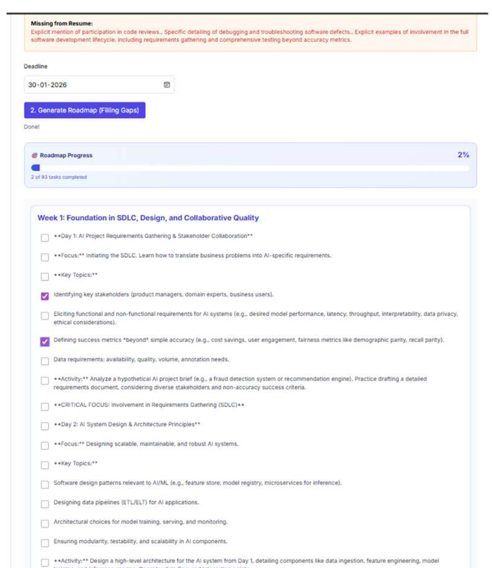

Fig. 3. Deadline-Conscious Route Planner

C. Semantic CV Evaluator

- Vector Similarity: Rather than a keyword-based approach, this model works based on the cosine similarity of resume and job description vectors.
- Gap Analysis: Semantic gap analysis highlights implied concepts within the job description that the candidate's resume misses (e.g., Job Description asks for Scalable Systems, the user has System Design experience, but no scalable solutions).
- Optimization: Gives specific recommendations to rewrite the bullet points to optimize ATS score.

D. Automated Communication (Cold Email/Referral)

Networking is now done in an organized manner to minimize any form of resistance:

- Template Creation: In accordance with the project requirements of the user and the industry of thetarget firm, the AI analyses the target market.
- Customization: It develops Cold Emails for the recruiters and LinkedIn Referred Messages that integrate hooks seamlessly in order to integrate the user experience and the business needs.

E. Instant Job Hunt and Career Advisor

- Live Aggregation: It integrates the features offered by HerKey, Adzuna, and Gemini Search in order to fetch real-time job postings from vector profiles of the users.
- Career Coach Chatbot: The RAG chatbot is used here for users to ask open-ended questions like

"How reasonable is the pay?" and "Explain to me about this roadmap topic." Pinecone is integrated to provide customized answers.

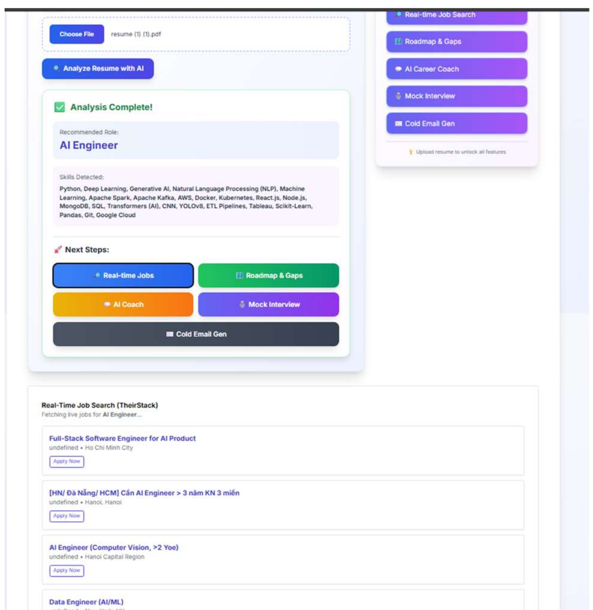

Fig. 4. Real-Time Job Search Results.

VI. Implementation Details

The design follows a modular approach across four levels.

A. Interface Level

HTML, CSS, and JavaScript were used to build the frontend, which supports file uploads using 'FormData' objects. An interactive dashboard displays the server's response in the form of JSON. The central point is the dashboard of the user, showing the results of the automatic analysis with skills discovered and immediate access to tools.

B. Application Layer (Backend)

Node.js comprises the backend that acts as the controller.

- Input Processing: Prior to file processing, the backend verifies the file type (PDF only) and size (under 5 MB).
- Text Standarization: A regex-based cleansing procedure removes any noise (\n+ replaced by spaces) for more efficient token utilization.
- Prompt Engineering: Using the job profile and timeline information provided by the user, the backend creates dynamic prompts.

C. AI Processing Layer

This layer communicates with the OpenAI/Gemini APIs.

- Voice Processing: The audio blobs are transmitted to the endpoint of the Mock Interviewer at the Medium. The content is then transcribed and finally inputted into the LLM where it is semantically analyzed.
- Context Management: It optimizes the 'Context Window' by removing unnecessary resume fields (such as Hobbies) before the data was submitted to the model.

D. Safety and Enhancement

Safety was a crucial issue throughout the implementation.

- Rate Limiting: We implemented strict rate limiting on the API endpoints to prevent abuse and manage costs associated with the Gemini API.
- Data Cleaning: This is responsible for cleaning all user input to avoid attacks such as Prompt Injection attacks, where an attacker may try to change the system instructions.
- Temporary Storage: In the course of our privacy policy, user resumes are not saved to disk but in memory only to the extent that the operation is performed during the session.

VII. Experimental Setup and Results

In order to assess Pathfinder AI accurately, an evaluation framework was established emphasizing Semantic Accuracy, Latency, and User Satisfaction.

A. Dataset Description

The analysis was carried out based on a specially created dataset including N=50 anonymous users' profiles received after the process of informed consent had been executed among university students and young professionals. The profiles consisted of three categories: Student (n = 22), Career Pivot (n = 18), and Experienced Professional (n = 10). A set of profiles was formed by combining a PDF version of the resume with a job posting description of the target position available for employment on LinkedIn and Adzuna. The target skill deficiencies for each profile were marked independently by two experts from the relevant field (κ inter-rater reliability = 0.84), and in the event of disagreement, the consensus was taken. The preparation time for all profiles ranged from 2 to 30 days (mean = 11.4 days, standard deviation = 6.8 days).

B. Quantitative Metrics

- Skill Gap Detection Accuracy Pgap: True skill gaps are denoted as Gtrue = J/C. Predicted skill gaps are given by Gpred.

$$P\_gap = |G\_pred \cap G\_true| / |G\_pred|, \;\; R\_gap = |G\_pred \cap G\_true| / |G\_true| \quad (1)$$

The F1-Score is calculated as:

$$F1\_gap = 2 \cdot P\_gap \cdot R\_gap / (P\_gap + R\_gap) \qquad (2)$$

On average across all 50 cases, Student, Senior, and Pivot, the proposed system attained
an F1-Score of 0.92 (σ = 0.04). This indicates that the model is highly precise in detecting skill gaps relative to human labeling. Precision and recall measures were recorded as 0.91 (σ = 0.05) and 0.93 (σ = 0.03), respectively. The precision and recall metrics were 0.91 (σ = 0.05) and 0.93 (σ = 0.03), respectively.

- Curriculum Density Coefficient ρ: This assesses the effectiveness of the generated study schedule with respect to time constraints.

$$\rho(T\_rem) = \Sigma(w\_i \cdot U\_i) / T\_rem \qquad (3)$$

Where N is the number of gaps, U_i is the basic learning unit, w_i is the complexity factor, and T_rem is the time left. The algorithm successfully calculated ρ in order to switch on "Crash Course" mode when T_rem < 5 days.

C. System Performance

- Inference Time: With the aid of Gemini 2.5 Flash, the average inference time for roadmap creation was 15.2 s ($\sigma = 1.8$ s), which is a 40% reduction from the initial inference time of GPT-4 models (25.4 s).
- Search Speed: Average vector search using Pinecone was 87 ms ($\sigma = 12$ ms).

D. User Satisfaction

The qualitative analysis using the sample of $N = 25$ subjects (selected from the described database) used a 5-point Likert scale. The mean satisfaction level was found at the mark of 4.2/5 ($\sigma = 0.6$). In particular, 84% selected “Deadline Awareness” as the most influential criterion, while the “Hire / No Hire” decision module gained praise due to the readiness signal it provided before real interviews were conducted. The Net Promoter Score (NPS) amounted to +72 (“Excellent”).

E. Baseline Comparison

For the thorough evaluation of the system performance, PathFinder AI is benchmarked against two other conventional methods regularly used in career counseling systems:

- Deterministic Rule-Based Scheduler: A scheduler that divides the available time for preparation equally among all identified skill gaps, without deadline-aware scheduling and semantic consideration.
- Pure LLM Career Counseling System: A method that asks the Large Language Model (LLM – Gemini 2.5 Flash) to generate a preparation plan based on the resume and job description of the user, skipping the scheduling step and using no semantic vector representation of skill gaps.

The performance is measured by three metrics: F1-Score for semantic gap detection, average system response time, and the user satisfaction rate (qualitative study, $N = 25$).

TABLE II: BASELINE COMPARISON OF CAREER GUIDANCE APPROACHES

Method	F1-Score ($\mu \pm \sigma$)	Avg. Latency (s)	User Satisfaction (%)
Rule-Based Scheduling System	0.78 ± 0.06	8.4 ± 1.1	61
LLM-Only Guidance System	0.85 ± 0.05	18.7 ± 2.3	73
PathFinder AI (Proposed)	0.92 ± 0.04	15.2 ± 1.8	84

The proposed system is superior to the two baseline systems on all three metrics. The rule-based system attains the shortest latency thanks to the determinism of the process, but its allocation algorithm leads to the poorest results in terms of the F1-Score and user satisfaction, which indicates that a uniform scheduling algorithm is not enough for customized preparation.

The second algorithm exhibits reasonable coverage of semantic terms but does not have any prioritization strategy, causing poor results in user satisfaction and latency.

F. Ablation Study

For isolating the role played by each module in architecture, an ablation study has been performed by disabling the critical components. Three architectures have been compared to the full architecture, as described in Table III.

TABLE III: ABLATION STUDY RESULTS

Configuration	F1-Score ($\mu \pm \sigma$)	User Satisfaction (%)
Full System (PathFinder AI)	0.92 ± 0.04	84
w/o Scheduling Module	0.86 ± 0.05	70
w/o Semantic Skill-Gap Analysis	0.83 ± 0.06	68
w/o Both (LLM-Only Equivalent)	0.85 ± 0.05	73

Without the schedule module, F1-Score decreases by 6.5 percentage points while user satisfaction goes down by 14 percentage points, since, due to equal time distribution, deadlines do not play any role, resulting in poor study schedules for users.

Without the semantic skill gap analysis module, there is the greatest loss of F1-Score (-9 pp), as, without taking advantage of semantically inferred skills, the system defaults to plain keyword-based competency identification. In comparison, the equivalent setup without either module achieves medium levels of satisfaction (73%), since it leverages wide-ranging information provided by the LLM, albeit without a well-defined structure and prioritization. All of these findings demonstrate that the suggested design is more than a sum of its parts.

VIII. CONCLUSION

The introduction of PathFinder AI is a noteworthy architectural development in educational technology since it shifts the paradigm from a passive form of knowledge retrieval to a proactive interactive system that aims to influence users' performance. The integration of the inferential abilities of LLMs and a bespoke deadline-oriented schedule allows the system to address the problem of choice overload faced by candidates when looking for employment. The findings confirm a more general hypothesis — the ability of Generative AI to successfully conduct structured planning and demonstrate it by creating personalized deadlines and priorities for each candidate based on an intelligent learning curriculum. Finally, the three-layer decoupled architecture confirms that high-volume and real-time AI inference can be performed efficiently without sacrificing system latency. The inclusion of bias-proof prompts and a focus on the privacy of users' data provides a better ethical standard for AI-assisted hiring solutions, ensuring they remain beneficial for their intended audience. Overall, PathFinder AI is not just another collection of curated resources but a personal and evidence-based journey towards becoming employable.

FUTURE SCOPE

Despite the success that has been achieved in developing this version of PathFinder AI, the future development trends in AI-coaching solutions seem to be heading in the direction of multimodal, more immersive experiences. For instance, a proposed upgrade would involve the use of computer vision and sophisticated audio analysis to create a fully functional mock interview setup. Through real-time monitoring of non-verbal cues, such as eye contact and facial expressions as well as pitch changes and other vocal nuances, the system would provide personalized feedback on soft skills – areas in which a candidate usually succeeds or fails in an interview but which are not addressed in the text-based solutions. In addition to interview training features, the platform has plans to link to the user's productivity ecosystem, such as calendar APIs like Google Calendar or Microsoft Outlook, so that the system could generate personalized schedules for the user. Furthermore, hyper-personalization via the creation of company profiles will be an important area of development, allowing users to prepare using interview experiences based on the evaluation criteria of particular companies.

Finally, for ensuring sustained involvement, the intention is to integrate gamification that takes behavioral psychology into consideration. With the help of longitudinal profiles that measure skill advancement on an ongoing basis, the platform would be able to track the path of competence and employ feedback loops to motivate career enhancement.

ACKNOWLEDGMENT

The authors wish to express sincere gratitude to the Department of Artificial Intelligence and Data Science at Dwarkadas J. Sanghvi College of Engineering for their continuous support and guidance throughout this research.

REFERENCES

[1] V. Jagwani et al., “Resume Evaluation through Latent Dirichlet Allocation and NLP for Effective Candidate Selection,” 2023.
[2] S. Li et al., “National Origin Discrimination in Deep-learning-powered Automated Resume Screening,” 2023.
[3] C. Harris, “Mitigating Age Biases in Resume Screening AI Models,” 2023.
[4] S. Kumar et al., “Fairness of Recommender Systems in the Recruitment Domain,” 2023.
[5] M. Armstrong et al., “The Silicon Ceiling: Auditing GPT’s Race and Gender Biases in Hiring,” 2024.
[6] K. Wilson et al., “Gender, Race, and Intersectional Bias in Resume Screening via Language Model Retrieval,” 2024.
[7] J. Everitt et al., “Large Language Models for Detecting Bias in Job Descriptions,” 2024.
[8] I. Samadhiya, “Machine Learning Bias in Resume Screening (Practical Approaches),” 2024.
[9] A. Deshpande et al., “Mitigating Demographic Bias in AI-based Resume Filtering,” 2020.
[10] X. An et al., “Measuring Gender and Racial Biases in Large Language Models,” 2025.
[11] I. Albaroudi, “A Comprehensive Review of AI Techniques for Addressing Algorithmic Bias in Hiring,” 2024.
[12] M. Islam et al., “Debiasing Career Recommendations with Neural Fair Collaborative Filtering,” 2021.
[13] S. Devaraju, “NLP in AI-Driven Recruitment Systems,” 2022.
[14] L. Zhang, “Understanding Algorithmic Bias in Job Recommender Systems,” 2022.