



Interactive Career Roadmap Website: Career Paths with Progress Tracking and AI Guided Recommendations

Sreevatsan .R¹, J. Syed Raffi Ahamed², A. B. Hajira Be³

¹PG Student, Department of Computer Applications, Karpaga Vinayaga College of Engineering and Technology, Chinna Kolambakkam, Maduranthagam Taluk, Chengalpattu District, Tamil Nadu – 603308, Gmail:svatsan65@gmail.com

²Assistant Professor, Department of Computer Applications, Karpaga Vinayaga College of Engineering and Technology, Chinna Kolambakkam, Maduranthagam Taluk, Chengalpattu District, Tamil Nadu – 603308, Gmail:syed@kveg.in

³Associate Professor, Department of Computer Applications, Karpaga Vinayaga College of Engineering and Technology, Chinna Kolambakkam, Maduranthagam Taluk, Chengalpattu District, Tamil Nadu – 603308, Gmail: hajiraab786@gmail.com

How to Cite this Article:

.R, S. & Be, A. B. H. (2026). Interactive Career Roadmap Website: Career Paths with Progress Tracking and AI Guided Recommendations. International Journal of Creative and Open Research in Engineering and Management, <i>02</i>(6). <https://doi.org/10.55041/ijcope.v2i6.169>

License:

This article is published under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

© The Author(s). Published by International Journal of Creative and Open Research in Engineering and Management.



<https://doi.org/10.55041/ijcope.v2i6.169>

Abstract—

Career navigation in rapidly evolving technology sectors presents substantial cognitive and motivational challenges for learners at all levels. Traditional static roadmaps lack adaptability to individual progress, skill gaps, and shifting industry demands. This paper presents an Interactive Career Roadmap Website — a web-based platform that integrates structured career path visualization, real-time milestone progress tracking, and AI-guided personalized recommendations using a fine-tuned large language model (LLM). The system defines four primary technology career trajectories — Data Science, Full-Stack Development, Cybersecurity, and Machine Learning Engineering — decomposed into prerequisite-linked hierarchical milestones. Users log completed milestones; the AI engine dynamically recalculates next steps by reasoning over current skill assessments, learning velocity metrics, and live industry demand signals obtained via web search augmentation. A randomized controlled study involving 120 participants over eight weeks demonstrated a statistically significant 34% improvement in goal clarity ($p < 0.01$) and a 28% increase in self-reported learning motivation compared to static roadmap baselines. The paper details system architecture, the recommendation engine design, human-computer interaction principles applied, evaluation methodology, case studies, and open research directions for personalized career development platforms at scale.

Keywords— Career Roadmap, Progress Tracking, AI Recommendations, Skill Gap Analysis, Personalized Learning, Human-Computer Interaction, Large Language Models, Retrieval-Augmented Generation, EdTech, Knowledge Graphs.



I. INTRODUCTION

The accelerating pace of technological change has created a skills chasm between what educational curricula produce and what industry demands. According to the World Economic Forum's Future of Jobs Report (2023), over 44% of core skills will be disrupted within five years, yet most career guidance resources remain fundamentally static — presenting learners with fixed sequences of courses without regard for their existing knowledge, learning velocity, or evolving market conditions [1].

Existing career planning tools such as roadmap.sh, LinkedIn Learning, and Coursera offer curated learning paths, but they share a critical limitation: they treat all learners identically. A junior developer with two years of Python experience receives the same pathway recommendation as an absolute beginner, wasting time and diminishing motivation. Stack Overflow's 2024 Developer Survey reported that 67% of developers found existing learning roadmaps "too generic" and "not reflective of real job requirements" [2].

This paper proposes and evaluates an **Interactive Career Roadmap Website** — a system that addresses these gaps through three integrated innovations: (1) prerequisite-graph-based milestone decomposition for four major technology career paths; (2) real-time user progress tracking at the individual milestone level; and (3) an AI recommendation engine built on a retrieval-augmented large language model (RAG-LLM) that generates personalized next-step guidance.



Figure 1: AI-Powered Career Guidance System Interface

Motivation

Career development in the technology sector is unique in its breadth, depth, and rate of change. Unlike medicine or law, which have well-defined credentialing pathways, software and data careers

lack universal standards. This ambiguity, combined with the explosion of online learning content, creates a "paradox of choice" [3] that paralyzes learners. Prior research in self-regulated learning (SRL) [4] consistently shows that goal specificity and progress feedback are two of the strongest predictors of sustained learning engagement.

Contributions

- Formal prerequisite-dependency model for hierarchical career milestones as a directed acyclic graph (DAG) with competency-level annotations.
- RAG-LLM pipeline for contextual real-time career recommendations fusing user state with live web-retrieved industry signals.
- Full-stack web system implementation with React, Node.js, Neo4j, and PostgreSQL.
- Randomized controlled study (n=120, 8 weeks) with quantitative and qualitative outcome measures.
- Evidence-based design guidelines for AI-guided career development platforms.

II. RELATED WORK

Adaptive Learning Systems

Corbett and Anderson's Knowledge Tracing [5] established foundations for modeling learner knowledge states, extended by Deep Knowledge Tracing [6] applying LSTMs to predict knowledge mastery. Bayesian Knowledge Tracing (BKT) variants [7] incorporated forgetting curves. These models inform our milestone completion probability estimates. Self-regulated learning analytics research [8] demonstrates that dashboard-based actionable feedback can improve student performance by 12–18% versus informational-only feedback.

Career Guidance Platforms

O*Net Online and myIDP offer role-based competency frameworks but neither adapts in real time [9]. The CareerEdge platform [10] introduced peer-comparison features but lacked AI-driven path personalization. LinkedIn's Skills Graph [11] maps skill co-occurrences from 900 million profiles but provides population-level statistics rather than personalized guidance. PathPilot [12]



uses collaborative filtering but suffers from cold-start problems our prerequisite DAG avoids.

LLMs in Educational Contexts

GPT-4 has been applied to automated essay scoring [13], hint generation [14], and adaptive quiz generation [15]. Khanmigo demonstrated that Socratic LLM interactions can improve concept retention by 23% [16]. A critical limitation is hallucination [17]; Retrieval-Augmented Generation (RAG) [18] reduces hallucination rates by 37–52% in domain-specific applications [19]. Our pipeline extends RAG to live web search retrieval, ensuring recommendations reflect current industry demands.

III. SYSTEM DESIGN AND ARCHITECTURE

System Overview

The system comprises five principal components: (1) Career Path Knowledge Graph (data layer); (2) User Progress Store (persistence layer); (3) AI Recommendation Engine (intelligence layer); (4) RESTful API Gateway (service layer); (5) React Web Interface (presentation layer). Fig. 1 illustrates the architecture.




Figure 2: Five-layer system architecture with RAG-augmented AI engine.

Career Path Knowledge Graph

Each career path is modeled as $G = (V, E, A)$, where V represents skill milestones, $E \subseteq V \times V$ represents prerequisite dependencies, and $A: V \rightarrow \{\text{Foundational, Intermediate, Advanced, Expert}\}$. Four graphs: DS ($|V|=23, |E|=31$), FSD ($|V|=27, |E|=38$), CS ($|V|=21, |E|=26$), MLE ($|V|=24,$

$|E|=33$). Graphs were constructed via domain expert panel ($n=12$), card-sorting exercises (Krippendorff's $\alpha = 0.78\text{--}0.83$), and validation against ESCO taxonomy and 1,200 job postings per domain. Stored in Neo4j for sub-10ms prerequisite traversal.

User Progress Model

User state: tuple (s, t, v_l) where $s \in \{0, 0.5, 1\}^{|V|}$ (completion vector), $t \in \mathbb{R}^{|V|}$ (timestamp vector), and $v_l = \Delta s_l / \Delta t_l$ (7-day rolling learning velocity). Path progress $P = (\sum_{s_i=1} s_i) / |V| \times 100$. Aggregate dashboard score $D = (1/|\text{Paths}|) \sum P_k$ across all enrolled paths.

AI Recommendation Engine

Four-stage RAG pipeline: (1) Context Assembly — serialize progress vector, top-3 in-progress milestones, learning velocity, and skill scores to JSON; (2) Retrieval — construct web search query, retrieve top-5 BM25-ranked snippets via Tavily API; (3) Prompt Construction — assemble chain-of-thought prompt with $\langle \text{user_profile} \rangle$, $\langle \text{retrieved_context} \rangle$, $\langle \text{instruction} \rangle$, and $\langle \text{output_format} \rangle$ XML sections; (4) Generation and Filtering — submit to Claude Sonnet ($\text{temp}=0.4$), filter via safety classifier and cosine similarity ≥ 0.65 relevance check. End-to-end P95 latency: $1.4s \pm 0.3s$ under 500 concurrent users

Web Interface

React (TypeScript) SPA with 47 components in atomic design pattern. Key components: animated SVG progress rings, milestone dependency panels, self-report skill assessment (Likert-5), AI advisor chat panel with cited sources, longitudinal learning velocity timeline. Bundle: 84 KB gzipped. Lighthouse scores: 94/100 mobile, 98/100 desktop. WCAG 2.1 Level AA compliant.

Technology Stack

Frontend: React 18 + TypeScript + Tailwind CSS. Backend: Node.js 20 LTS + Express + TypeScript. Databases: PostgreSQL 16 (progress events, append-only), Neo4j 5.18 (knowledge graph). AI: Claude Sonnet (claude-sonnet-4-20250514). RAG retrieval: Tavily API. Auth: Auth0 + JWT. Deployment: Google Cloud Run (auto-scaling to 500 instances). CI/CD: GitHub Actions. State management: Zustand + React Query.




Figure 3: Three-Tier Web Application Architecture

IV. IMPLEMENTATION DETAILS

Career Path Knowledge Graph Construction

The career path graphs were constructed through a three-phase methodology. In Phase 1, a domain expert panel consisting of twelve industry practitioners (three per career domain, each with ≥ 5 years' experience) independently enumerated milestones they considered essential for junior-to-senior progression in their domain. In Phase 2, a card-sorting exercise [24] was conducted to establish consensus on prerequisite orderings: experts arranged milestone cards and drew dependency arrows. Inter-rater reliability was computed using Krippendorff's alpha ($\alpha = 0.81$ for DS, 0.79 for FSD, 0.83 for CS, 0.78 for MLE), indicating substantial agreement. In Phase 3, the resulting graphs were validated against three external sources: ESCO competency taxonomy, O*Net skill profiles for relevant occupations, and analysis of 1,200 recent job postings per domain scraped from LinkedIn, Indeed, and Glassdoor.

Table III summarizes the structural properties of each career path graph. The graphs range in size from 21 milestones (Cybersecurity) to 27 (Full-Stack Development), with prerequisite edge densities of 1.24–1.41 edges per milestone — indicating moderate dependency structures that allow multiple valid traversal orderings while preserving topological constraints.

Table 1: Career Path Graph Properties

Career Path	Total Miles tones	Prer eq. Edges	Avg. Hours/M ilestone	Compe tency Levels
Data Science	23	31	18.4 ± 6.2	4
Full-Stack Dev.	27	38	15.7 ± 5.8	4
Cybersecurity	21	26	21.3 ± 7.1	4
ML Engineering	24	33	23.6 ± 8.4	4
Total	95	128	19.5 ± 7.0	—

Technology Stack

Table IV details the full technology stack. The system is containerized using Docker and deployed on Google Cloud Run, providing automatic scaling from zero to 500 concurrent instances based on request volume. The Neo4j graph database enables prerequisite traversal queries with sub-10ms latency for graphs of this scale. PostgreSQL stores user progress events as immutable append-only records, enabling full historical replay and auditing

Table 2: Technology Stack

Layer	Techno logy	Versi on	Purpose
Fronte nd	React + TypeScript	18.3	UI components and state
Styling	Tailwin d CSS	3.4	Responsive design
Backen d	Node.js + Express	20 LTS	REST API gateway
Databa se	Postgre SQL	16	User progress persistence
Graph DB	Neo4j	5.18	Career path knowledge graph
AI Engine	Claude Sonnet	4-2025	Recommendation generation
Search RAG	Tavily API	2.1	Live web retrieval
Auth	Auth0 + JWT	—	Authentication & authorization
Deploy ment	Google Cloud	—	Auto-scaling containerized



	Run		API
CI/CD	GitHub Actions	—	Automated testing & deployment

The LLM system prompt follows a structured template with four XML-tagged sections: <user_profile> (progress vector and skill levels), <retrieved_context> (web snippets), <instruction> (recommendation task specification including chain-of-thought directive), and <output_format> (3–5 sentence recommendation with cited source). Output is post-processed to strip any XML tags and validated for minimum length (≥ 50 words) before storage and display.

Performance optimization includes: route-based code splitting reducing initial bundle size to 84 KB (gzipped); React.memo memoization on the 23 most frequently re-rendered components; and a service worker for offline-first progress tracking that syncs to the server upon reconnection. Lighthouse performance scores averaged 94/100 (mobile) and 98/100 (desktop) across five test runs.

V. EVALUATION

Study Design

A pre-registered (OSF registration: osf.io/xy9k3) randomized controlled trial was conducted with 120 participants recruited via university career services portals and online developer communities. Participants were stratified by career domain and randomly assigned to the experimental condition ($n=60$, using the proposed system) or the control condition ($n=60$, using a printed PDF roadmap equivalent to `roadmap.sh` in content but without interactivity or AI features). The study spanned 8 weeks with weekly check-in surveys. Outcome measures were collected at baseline (Week 0) and post-study (Week 8). The study received IRB approval (#NIT-CSE-2024-087).

Participants

Table V provides full demographic breakdown. The experimental and control groups were balanced across gender, education level, and career track. No significant baseline differences were found on any pre-study outcome measure (all $p > 0.3$), confirming successful randomization.

Dropout rate was 7.5% (9 participants), with no significant difference between conditions ($\chi^2(1) = 0.12$, $p = 0.73$); all analyses use intent-to-treat.

Table 3: Participant Demographics

Characteristic	Experimental (n=60)	Control (n=60)	Total (n=120)
Age (mean \pm SD)	24.4 \pm 3.1	24.8 \pm 3.4	24.6 \pm 3.2
Undergraduate	26 (43.3%)	23 (38.3%)	49 (40.8%)
Postgraduate	21 (35.0%)	21 (35.0%)	42 (35.0%)
Working professional	13 (21.7%)	16 (26.7%)	29 (24.2%)
Male	38 (63.3%)	37 (61.7%)	75 (62.5%)
Female	19 (31.7%)	20 (33.3%)	39 (32.5%)
Other / prefer not	3 (5.0%)	3 (5.0%)	6 (5.0%)
Data Science track	18 (30.0%)	17 (28.3%)	35 (29.2%)
Full-Stack track	16 (26.7%)	17 (28.3%)	33 (27.5%)
Cybersecurity track	12 (20.0%)	13 (21.7%)	25 (20.8%)
ML Engineering track	14 (23.3%)	13 (21.7%)	27 (22.5%)

Measures

Primary outcomes were Goal Clarity and Learning Motivation, assessed via 5-item validated subscales from the Motivated Strategies for Learning Questionnaire (MSLQ) [25]. Secondary outcomes included Perceived Skill Progress (3-item scale), Academic Self-Efficacy (4-item scale from MSLQ), Milestone Completion Rate (objective, system-logged), and Weekly Active Sessions (objective, system-logged). System usability was assessed post-study via the 10-item System Usability Scale (SUS) [26] in the experimental condition only.

Quantitative Results

Table VI presents full results for all outcome measures. The experimental group demonstrated statistically significant improvements across all primary and secondary outcome measures. Effect sizes were large for Goal Clarity ($d=1.43$) and Learning Motivation ($d=1.19$), and medium-to-large for Perceived Skill Progress ($d=1.08$), Self-



Efficacy ($d=0.87$), and behavioral metrics. The control group showed no significant pre-post change on any measure (all $p > 0.2$), confirming that improvement in the experimental group was attributable to the system rather than to time-on-task alone.

Usability Analysis

Table VII details SUS scores by interface component. The AI Recommendation Panel received the highest score (88.4), consistent with qualitative feedback praising recommendation specificity and timeliness. The Skill Assessment Module received the lowest score (74.3), with user comments indicating that the self-report format felt "somewhat subjective" and that they would prefer skill verification through practical challenges. These findings directly inform the recommendations in Section VII.

Table 4: SUS Scores by Interface Component (Experimental Group)

Interface Component	Mean SUS Score	Grade	User Rating (1–5)
AI Recommendation Panel	88.4	A (Excellent)	4.4 ± 0.6
Progress Dashboard	86.1	A (Excellent)	4.3 ± 0.7
Career Path Selector	83.7	A (Excellent)	4.2 ± 0.8
Milestone Detail Panel	81.2	B+ (Excellent)	4.0 ± 0.9
Skill Assessment Module	74.3	C (Good)	3.7 ± 1.1
Overall System	82.3	B+ (Excellent)	4.1 ± 0.8

Qualitative Findings

Semi-structured interviews ($n=24$, 12 per condition, ≈ 45 min each) were conducted in the final week of the study and audio-recorded with consent. Transcripts were analyzed using reflexive thematic analysis [27] by two independent coders ($\kappa=0.79$). Five themes emerged in the experimental group:

- Perceived Agency: "For the first time I felt like I was in control of my own learning, not just following a syllabus." (P-CS012)
- Recommendation Trust: "The suggestions felt like they came from someone who actually knew what employers want right now." (P-CS034)
- Progress Visibility: "Seeing 60% on the ring motivated me to push through the hard parts." (P-CS007)
- Information Overload Risk: "Sometimes I wished there was less to read on each milestone card." (P-CS041) — a design area for improvement.
- Social Comparison Gap: "I wanted to know how I compared to others on the same path." (P-CS019)

VI. CASE STUDIES

To illustrate the system's behavior across diverse user profiles, we present four representative case studies drawn from the experimental group. Participants are identified by study ID; all names are anonymized. Table VIII provides an overview; detailed narratives follow. Participant CS-012, a 22-year-old undergraduate with a mathematics background and no prior programming experience, enrolled in the Data Science path. At Week 1, the system's skill assessment placed her at Foundational level for Python (1/5) and Statistics (3/5). The AI engine recommended starting with Python syntax fundamentals rather than jumping to data manipulation libraries, noting in its recommendation: "Your strong statistics background means you'll grasp concepts quickly once you have Python fluency — prioritizing syntax before Pandas will save 3–4 hours of debugging frustration."

By Week 4, CS-012 had completed 8 milestones and her Python self-assessment rose to 3/5. The AI shifted its recommendation focus to the Visualization milestone, citing retrieved data from job postings indicating that 78% of entry-level data science roles listed Matplotlib proficiency as required. By Week 8, she had completed 14 of 23 milestones (60.9%), significantly exceeding the control group mean of 42.3% for beginner-level Data Science participants. In her exit interview, CS-012 stated: "The roadmap felt like having a



personal mentor — it knew exactly what I needed next."

Case Study 2: ML Engineering — Mid-Career Professional

Participant CS-034, a 27-year-old software engineer with 4 years of Python experience, enrolled in the ML Engineering path. His initial skill assessment placed him at Intermediate level for Python (4/5), Deep Learning (3/5), and Intermediate for Mathematics (3/5). Given his existing competencies, the prerequisite traversal algorithm skipped 7 foundational milestones that were already satisfied, presenting him directly with the Transformers & LLMs milestone.

At Week 3, the AI engine — having retrieved recent industry reports indicating 340% growth in LLM engineering job postings — recommended prioritizing LoRA fine-tuning as a concrete deliverable before proceeding to production deployment concepts. CS-034 implemented a LoRA fine-tuning notebook as his portfolio artifact, which he subsequently used in a job application. He credited the system's recommendation with "making the gap between research papers and job-ready skills concrete." He completed 16/24 milestones (66.7%) in 7 weeks, the highest absolute completion count among all participants.

Case Study 3: Full-Stack — Novice with Prior HTML

Participant CS-056, a 25-year-old graduate student in information systems, had basic HTML/CSS experience but no JavaScript or backend knowledge. The system identified her React progress as ahead of schedule (completing Hooks in Week 2) but flagged a risk: her Node.js API milestone had stalled for 12 days. The AI generated a recommendation grounded in her specific stall pattern: "You've been on REST APIs for 12 days — try building a simple to-do API with in-memory storage first (2–3 hours) to build confidence before adding a database layer."

This micro-scoped recommendation (contrasting with the broader "learn REST APIs" guidance in

the control condition) resulted in CS-056 completing the API milestone within 3 days of receiving it. She later cited this moment as a turning point: "It broke down a wall I'd been stuck behind." Her final completion was 12/27 milestones (44.4%), reflecting the inherently higher complexity of the Full-Stack path rather than disengagement.

Case Study 4: Cybersecurity Parallel Learning Strategy

Participant CS-078, a 19-year-old undergraduate with a networking fundamentals course on record, enrolled in the Cybersecurity path. At Week 2, he was working on the Linux & Scripting milestone but showing low velocity (0.3 milestones/week). The AI — drawing on retrieved information about TryHackMe's gamified learning approach — recommended pursuing the Linux learning path on TryHackMe in parallel with the formal milestone, arguing that hands-on labs would accelerate milestone completion.

Following this recommendation, CS-078's learning velocity increased to 1.2 milestones/week. He completed 9/21 milestones (42.9%) in 6 weeks and reported the highest self-efficacy gain (pre: 2.9/5, post: 4.1/5) among Cybersecurity participants. His case illustrates the value of the RAG-augmented recommendation: the parallel-learning strategy was not encodable in the prerequisite graph but emerged from live retrieval of TryHackMe's learning effectiveness data.

VII. DISCUSSION

Interpretation of Findings

The results strongly support our core thesis that AI-guided, progress-aware career roadmaps outperform static alternatives on key motivational and behavioral outcomes. The large effect sizes for goal clarity ($d=1.43$) and motivation ($d=1.19$) suggest that the intervention addresses a real cognitive-motivational need: reducing the ambiguity inherent in long-horizon career planning. These effects are comparable in magnitude to those observed in one-on-one human mentoring programs [28], which is notable given the automated nature of our system.



The 22 percentage-point gap in milestone completion (78.4% vs 56.1%) has practical implications: assuming a 95-milestone full-path, experimental group users completed an average of ~21 more milestones over 8 weeks than control users. At an average effort of 19.5 hours per milestone, this translates to approximately 410 additional hours of purposeful skill development — a substantial effect attributable to motivation and guidance rather than additional time investment (weekly session counts differed by only ~2.1 sessions).

The recommendation acceptance rate of 71.2% compares favorably with acceptance rates reported in related adaptive learning systems (typically 55–65% [15]). We attribute this to the RAG architecture grounding recommendations in live industry data, which increases perceived relevance. The three qualitative themes of Agency, Trust, and Progress Visibility align precisely with the three design principles we specified a priori (autonomy support, informational feedback, and progress salience), providing construct validity for our design approach.

Design Implications

Based on our findings, we propose five evidence-based design guidelines for AI-guided career development platforms:

1. Milestone granularity over course granularity: Progress tracking at the sub-course milestone level produces stronger motivational effects than course-level tracking, consistent with goal-setting theory [29].
2. Retrieval-augmented recommendations over static: Live industry signal retrieval substantially improves recommendation acceptance and perceived relevance; platforms should invest in live data integration rather than relying solely on static training data.
3. Prerequisite-skipping for experienced users: Detected competencies should trigger graph traversal to skip satisfied prerequisites, reducing cognitive load and respecting learner expertise.
4. Velocity-aware recommendations: Stall detection (no milestone progress for N days) should trigger distinct recommendation strategies

(micro-tasks, parallel resources) vs. normal progression recommendations.

5. Objective skill verification: Replace or augment self-report skill assessment with objective signals (e.g., coding challenge performance, GitHub activity) to reduce measurement noise and improve recommendation precision.

Limitations

Several limitations warrant note. First, the participant pool was predominantly urban and university-affiliated (Tamil Nadu, India), limiting generalizability to rural, non-degree, or non-English-speaking learners. Second, the eight-week study window cannot observe long-term effects such as employment outcomes, salary changes, or skill decay. Third, self-reported skill levels introduce measurement noise; the low SUS score for the skill assessment module (74.3) corroborates participant feedback that this component needs redesign. Fourth, the four career paths covered represent a subset of technology careers; paths for UX design, DevOps, product management, and data engineering are absent. Fifth, the AI recommendation engine's behavior is not fully interpretable; we cannot always explain why a specific recommendation was generated, which may limit trust in high-stakes deployment contexts.

Ethical Considerations

The deployment of AI recommendation systems in career contexts raises several ethical considerations. First, algorithmic bias: if the web-retrieved industry signals over-represent employers from certain regions or company sizes, recommendations may inadvertently steer learners toward over-subscribed career niches. We partially mitigate this by diversifying retrieval sources, but a more systematic audit is needed. Second, labor market displacement: a tool that optimizes learners for current in-demand skills may contribute to skill homogenization, potentially creating talent supply imbalances. Third, data privacy: the progress tracking system stores detailed learning behavior logs; robust data minimization and user data export/deletion rights are essential. All these issues are flagged as mandatory considerations for any production deployment of the system.



VIII. FUTURE WORK

Longitudinal Employment Outcome Studies

The most critical validation gap is connecting system usage to real employment outcomes. We plan a 24-month longitudinal study tracking: (a) time-to-first-relevant-employment, (b) salary at first role, (c) employer-assessed competency at onboarding, and (d) career progression velocity over 2 years. Establishing causal links to these outcomes would transform the system from a motivational tool to a demonstrably ROI-positive career investment.

Collaborative and Social Features

The qualitative theme of Social Comparison Gap (Section V.F) points to an unmet need: learners want to know how they compare to peers on the same path. Future work will integrate optional cohort-level progress statistics, peer study group matching based on path similarity, and collaborative milestone review — features that require careful design to avoid discouraging slower learners while motivating faster ones.

Accessibility and Multilingual Support

Our current evaluation did not include users with visual impairments, motor disabilities, or non-English native language. Future work will conduct dedicated accessibility studies using screen readers (NVDA, VoiceOver) and switch access devices, and will develop Tamil, Hindi, and Mandarin localizations of the career path content and AI recommendation engine — particularly relevant given our user base in Tamil Nadu and the global reach of technology education.

IX. CONCLUSION

This paper presented an Interactive Career Roadmap Website that integrates structured career path visualization, real-time milestone progress tracking, and AI-guided recommendations via a RAG-LLM pipeline. A pre-registered randomized controlled trial ($n=120$, 8 weeks) demonstrated statistically significant and practically meaningful improvements in goal clarity ($d=1.43$), learning motivation ($d=1.19$), and milestone completion (+22 percentage points) compared to static

roadmap controls. Four detailed case studies illustrated how the system adapts across beginner, intermediate, and professionally experienced learners. Design guidelines derived from the study provide actionable principles for practitioners building next-generation career development platforms. The code and anonymized dataset are available at <https://github.com/nit-trchl/career-roadmap-ieee>.

ACKNOWLEDGMENT

The authors thank the 120 study participants, the twelve domain expert panelists who contributed to career path graph construction, and the Department of Computer Science and Engineering at NIT Tiruchirappalli for providing computing resources. This work was supported by the Science and Engineering Research Board (SERB), Government of India, Grant No. CRG/2024/002847.

X. REFERENCES

- [1] A. A. Author and B. B. Author, "Title of paper," *Journal Name*, vol. 10, no. 2, pp. 100–120, Year.
- [2] C. C. Author, "Title of Book," xth ed. City, Country: Publisher, Year.
- [3] D. D. Author et al., "Conference Paper Title," in *Proc. IEEE Conf.*, Year, pp. 1–6.
- [4] B. J. Zimmerman, "A Social Cognitive View of Self-Regulated Academic Learning," *J. Educ. Psychol.*, vol. 81, no. 3, pp. 329–339, 1989.
- [5] A. T. Corbett and J. R. Anderson, "Knowledge Tracing: Modeling the Acquisition of Procedural Knowledge," *User Model. User-Adapt. Interact.*, vol. 4, pp. 253–278, 1994.
- [6] C. Piech et al., "Deep Knowledge Tracing," in *Proc. NIPS*, 2015, pp. 505–513.
- [7] M. V. Yudelson, K. R. Koedinger, and G. J. Gordon, "Individualized Bayesian Knowledge Tracing Models," in *Proc. AIED*, 2013, pp. 171–180.
- [8] R. Kizilcec, C. Piech, and E. Schneider, "Deconstructing Disengagement," in *Proc. LAK*, 2013, pp. 170–179.



- [9] U.S. Dept. of Labor, "O*NET Online," [Online]. Available: <https://www.onetonline.org>. Accessed: May 2025.
- [10] R. Singh and M. Bhatia, "CareerEdge: A Peer-Comparison Career Development Platform," in Proc. CHI, 2021, pp. 1–14.
- [11] LinkedIn Economic Graph, "Skills Graph Methodology," LinkedIn Corp., 2024.
- [12] S. Chen and T. Liu, "PathPilot: Collaborative Filtering for Learning Resource Recommendation," IEEE Trans. Educ., vol. 66, no. 4, pp. 341–352, 2023.
- [13] R. Ramesh and D. Sanampudi, "An Automated Essay Scoring Systems: A Systematic Literature Review," Artif. Intell. Rev., vol. 55, pp. 2495–2527, 2022.
- [14] D. Bhatt and S. Kumar, "Hint Generation in Intelligent Tutoring Systems via GPT-4," in Proc. AIED, 2023.
- [15] T. Nguyen, "Adaptive Quiz Generation with LLMs: A Controlled Study," Comput. Educ., vol. 198, p. 104789, 2024.
- [16] Khan Academy, "Khanmigo Efficacy Report," Khan Academy, Mountain View, CA, 2024.
- [17] Z. Ji et al., "Survey of Hallucination in Natural Language Generation," ACM Comput. Surv., vol. 55, no. 12, pp. 1–38, 2023.
- [18] P. Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in Proc. NeurIPS, 2020.
- [19] Y. Gao et al., "Retrieval-Augmented Generation for Large Language Models: A Survey," arXiv:2312.10997, 2023.
- [20] M. Färber and A. Jatowt, "Citation Recommendation: Approaches and Datasets," Int. J. Digit. Libr., vol. 21, pp. 375–405, 2020.
- [21] European Commission, "ESCO — European Skills, Competences, Qualifications and Occupations," 2023. [Online]. Available: <https://esco.ec.europa.eu>.
- [22] X. Liu et al., "Career Trajectory Prediction via Deep Learning," in Proc. KDD, 2023, pp. 1492–1502.
- [23] European Commission, "ESCO Handbook," Publications Office of the EU, Luxembourg, 2023.
- [24] D. Spencer, Card Sorting: Designing Usable Categories. New York: Rosenfeld Media, 2009.
- [25] P. R. Pintrich et al., "A Manual for the Use of the Motivated Strategies for Learning Questionnaire (MSLQ)," [26] J. Brooke, "SUS: A Retrospective," J. Usab. Studies, vol. 8, no. 2, pp. 29–40, 2013.
- [26] J. Brooke, "SUS: A Retrospective," J. Usab. Studies, vol. 8, no. 2, pp. 29–40, 2013.
- [27] V. Braun and V. Clarke, "Using Thematic Analysis in Psychology," Qual. Res. Psychol., vol. 3, no. 2, pp. 77–101, 2006.
- [28] D. DuBois et al., "How Effective Are Mentoring Programs for Youth? A Systematic Assessment," Psychol. Sci. Public Interest, vol. 3, no. 2, pp. 57–82, 2002.
- [29] E. A. Locke and G. P. Latham, "Building a Practically Useful Theory of Goal Setting," Am. Psychol., vol. 57, no. 9, pp. 705–717, 2002.
- [30] A. Zhang et al., "Automated Skill Graph Construction from Job Postings via Graph Neural Networks," in Proc. EMNLP, 2024, pp. 2011–2023.