International Journal of Innovative Research & Growth
E-ISSN: 2455-1848
homepage: http://www.ijirg.com/ijirg/
Volume: 15 Issue: 3, July 2026

CareerWeave: AI-Driven Milestone Orchestration for Professional Skill

Keshav Dhanka[1*], Ashish Saini[2], Bhagawat Prasad[3], Om Prakash Sharma[4]
[1,2,3]B. Tech Scholar, Department of Computer Science & Engineering, Jagannath University, Jaipur- 303901, Rajasthan, India
[4]Professor, Department of Computer Science & Engineering, Jagannath University, Jaipur- 303901, Rajasthan, India

CORRESPONDING AUTHOR

Keshav Dhanka
e-mail: keshavdhanka42@gmail.com

KEYWORDS

Artificial Intelligence, Personalized Learning, LLM, Skill Gap Analysis, Career Path Optimization, Educational Technology.

ARTICLE DETAILS

Received 19 May 2026; revised 09 July 2026; accepted 20 July 2026

DOI: 10.26671/IJIRG.2026.3.15.281

CITATION
Dhanka, K., Saini, A., Prasad, B., Sharma, O. P. (2026). Career Weave: AI-Driven Milestone Orchestration for Professional Skill. *Int J Innovat Res Growth*, 15(3), 153037-153043. DOI

Abstract

In modern digital learning, the unstructured use of tutorials often leads to incoherent knowledge without an application strategy and without a clear path of professional development. To overcome this challenge, this paper introduces CareerWeave, an AI-powered career architect that generates personalized, project-based learning roadmaps.
Career Weave is an innovative Two-Pass LLM Orchestration with the Gemini 3 Flash engine. The Recruiter Pass converts unstructured user input into a standardized technical profile using a context parser. The Architect Pass performs qualitative Delta Analysis of the desired job requirements against the user competencies. This pass has a Feasibility Gate, a logic filter that avoids the generation of unrealistic career paths by evaluating transition realism.
The system is developed using Fast API, MySQL and React Flow, utilizing the principle of Sequential Milestone Unlocking to ensure learners remain focused on achievable goals within a standardized 24-week framework. Experimental results demonstrate the effectiveness of the framework to scale time to learn and recognize career jumps.

1. Introduction

The guidance gap is a major obstacle in professional growth; the gap is between the academic background and industry-specific knowledge (Mukashova et al., 2025). Although online platforms have extensive knowledge bases, they do not always have the structural integrity needed to achieve goal-driven learning (Ovtšarenko, 2023). Learners frequently encounter 'tutorial hell', characterized by a cycle of passive content consumption devoid of strategic advancement (Kadam et al., 2025). This is compounded by standardized curricula that do not take into consideration any prior specialized knowledge of a learner, and thus result either in redundancy or cognitive overload (Katiyar et al., 2024).

The traditional digital learning tools are commonly based on simple keyword matching, which lacks semantic depth as to be able to gauge the viability of a career transition. An intelligent system capable of qualitatively assessing the gap or delta between the current stack of a learner and the dynamic needs of the labor market is urgently required.

To address this gap, we are proposing CareerWeave, an AI-powered generative skill roadmap creator that uses generative intelligence to bridge this gap. The main objectives of this project are:

- Building a Python-based backend that will liaise with an LLM API to generate structured and formatted curricula in JSON format.
- Adopting the strategies of "Prompt Engineering" that would make sure that the designed roadmaps are realistic, sequential, and strictly time constrained.
- Applying a relational database system to store user profiles, generated roadmaps and milestone progress with complete referential integrity.
- A minimalist, gamified interface that visualizes complex learning journeys via an interactive timeline of Skill Trees.

The following are the major contributions of this paper:

1. A new Two-Pass Large Language Model (LLM) orchestration framework that separates user input structuring (Recruiter Pass) and roadmap generation (Architect Pass).
2. A feasibility-driven filtering mechanism ("Feasibility Gate") is applied to prevent unrealistic career transitions.
3. A qualitative gap analysis was performed via LLM-driven reasoning to identify discrepancies between a user's current technical stack and target role requirements.
4. A sequential, milestone-based learning map within a 24-week time frame of practical and achievable skills development.
5. The architecture and development of an interactive system to visualise individualised milestone-driven learning pathways by a Skill Tree model.

2. Related Work

According to Dash et al. (Dash et al., 2025) and Rameshbabu et al. (Rameshbabu.V et al., 2025), Current AI studies in education apply NLP to resumes to suggest modules to address workplace skills gaps. These approaches, however, are based on the matching of keywords to reflect technical skills and do not reflect the complexity of relationships among competencies (Jeffson Preetham Dsilva & Murari B K, n.d.; Mukashova et al., 2025). A more advanced method of reasoning, other than simple counting of words, is needed to assess proficiency accurately.(Dash et al., 2025)

The existing methods, including those of Ovtšarenko, are aimed at the creation of organized learning pathways (Ovtšarenko, 2023). Other researchers have investigated unsupervised learning to model the preferences (He et al., 2022) and an adaptive environment of the learner that adjusts the content delivery of the adaptive environment based on the measures of performance (Wang et al., 2023). Although these systems do enable personalization, they might be quite restricted in terms of their scope (Mukashova et al., 2025) and may not be aligned to the reality of the modern labor market, which is practical and fast-paced (Dash et al., 2025).

Although there is improvement, the existing solutions remain inadequate. Most platforms rely on matching that is not dynamic, and is based on templates, providing limited opportunities to pursue dynamic career paths (Jeffson Preetham Dsilva & Murari B K, n.d.; Ovtšarenko, 2023). More importantly, they do not consider the practical reality of the transitions; they do not take into consideration the reality in time, the time of the complex skills acquisition (Tapalova & Zhiyenbayeva, 2022).

To overcome these limitations, this work presents CareerWeave, a new Two-Pass Large Language Model (LLM) coordination framework that integrates skill gap assessment, which is semantically rich, with a feasibility-aware filtering mechanism. Unlike previous methods, this architecture includes a Feasibility Gate that is intended to eliminate unrealistic career changes and produce structured, milestone-based learning plans during a specified 24-week period. This facilitates the creation of individualistic, industry-relevant, and goal-oriented learning solutions based on practical applicability.

3. System Design

The Career Weave framework is a modular architecture that is unstructured from the structured career path and is designed to empower a learner to create learning journeys that will lead them from unstructured career visions to actionable projects. The system comprises four main components: the Presentation Layer (React Flow), the Middleware Layer (FastAPI), the External AI Service Layer (Gemini 3 Flash), and the Persistence Layer (MySQL 8.0), as shown in the system architecture block diagram in Figure 1.

3.1. System Architecture Overview

The system should be designed to be used in the request-response paradigm with the Fast API middleware controlling the system. The middleware then forwards the natural language query to the generative AI engine and the relational database; it handles all the necessary translations to ensure that all roadmaps generated are stored referentially correctly. There's JWT-based session management and crypt password hashing; data is isolated between user profiles.

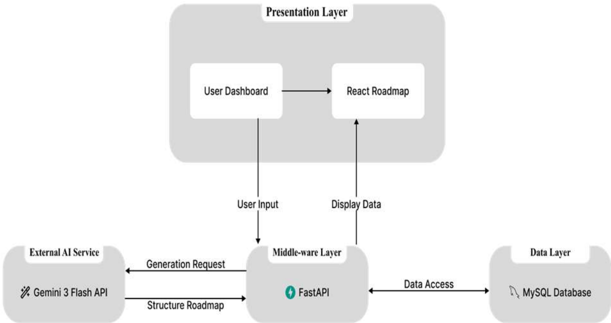

Fig.1 Career Weave System Architecture Block Diagram.

3.2. Two-Pass LLM Orchestration

Leading with a custom-developed Two-Pass AI Orchestration algorithm with the Gemini 3 Flash model at its heart. This is to make the information extraction process independent of curriculum creation and to produce high accuracy in the roadmap.

The workflow shown in Figure 2 describes the system's processing of data in a series of distinct data stages. The initial step is a Digital Context parser known as Recruiter Pass (Context Extraction). It can take structured information out of unstructured information given by the user: the user's main identity, current academic qualifications, technical skills learned so far, and explicit career planning. This step takes care of strict Pedantic schema validation to make sure that the extraction is in a format that is valid JSON for the next pass.

The second stage is a Technical Architect, called the Architect Pass (Curriculum Generation). During this stage, the system creates a curriculum plan using the structured profile built in the first pass and then uses this plan for the rest of the week. Each milestone is tied to a “Milestone Deliverable” that represents a task to accomplish and is based on technical skills.

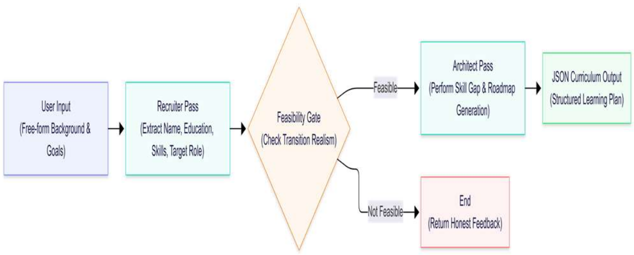

Fig.2 Two-Pass AI Orchestration Workflow.

3.3. Qualitative Delta Analysis & Feasibility Gate

Career Weave delivers Qualitative Delta Analysis as opposed to traditional systems, based on tallying the number of keywords. The LLM evaluates the learner's current skills and the skills they need for the desired career in the industry. The length of the roadmap varies depending on the assessed "gap" and can take anywhere from 4 weeks to 24 weeks.

The generated advice is validated and secured for its integrity and viability by the introduction of a Feasibility Gate in the generation logic. This process carries out a realism assessment for the requested career change. If a 24-week period is determined to be an impossible transition (such as a job change from a non-technical to an AI Architect position), the curriculum generation is rejected, and a realistic, technical explanation of the basic requirements for making this transition is given.

3.4. Sequential Milestone Unlocking & Persistence

A relational schema was used to ensure the persistence of the data, meaning that each milestone's completion was monitored. To make sure that the learner is not overwhelmed by the information and give him/her the right focus, the system has a "Sequential Unlocking logic". The next milestones (Week N+1) remain "redacted" and cannot be seen until the previous milestone (Week N) is completed in the database. The data is dynamically served to the frontend and then visualized through the node-based timeline shown in Figure 3.

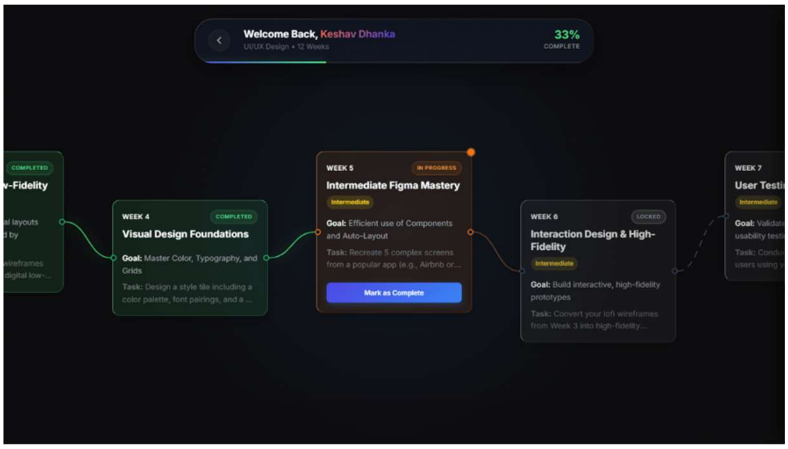

Fig.3 Interactive Career Roadmap using React Flow.

4. Implementation & Experimental Results

To rigorously evaluate the suggested architecture, we created a special tests/ directory. This framework independently validates the logic for API orchestration for Gemini 3 Flash, generation of curriculum, and operations on the MySQL database for the reliability and integrity of the system.

4.1. Testing Methodology

The five different behavioral test cases used as experimental variables were The Perfect Pivot, The Impossible Leap, Missing Skill Input, Upskilling Polish and Complete Transformation to validate the system's reasoning and the effectiveness of the Feasibility Gate. These scenarios were deliberately conceived to evaluate the system's aptitude for recognizing unrealistic career transitions and dynamically adjusting learning times depending on the technical difference.

While these predefined scenarios establish system validity across a broad taxonomy of backgrounds, this methodology operates as an initial logic-gate evaluation rather than an empirical user study. Consequently, immediate quantitative data regarding user satisfaction, long-term learning retention, or career transition tracking in the real world is bounded by this simulated approach.

4.2. Performance Results

Evaluation was based on the logical connection between the AI-developed roadmaps and industry standards.

Table 1: System Performance and Logic Validation

Test Case	Predicted Duration	Result	Logic Applied
The Perfect Pivot	8 Weeks	Feasible	Delta Analysis
The Impossible Leap	0 Weeks	Blocked	Feasibility Gate
Upskilling Polish	4 Weeks	Feasible	Skill Gap Mapping
Missing Skill Input	N/A	Rejected	Input Validation
Transformation	24 Weeks	Feasible	Maximum Scaling

The results shown in Table 1 demonstrate the effectiveness of the Feasibility Gate in preventing the generation of high-risk transitions, while providing constructive feedback to explain the constraints to the user, as demonstrated in the transition scenario from a non-technical background to AI Engineering. By evaluating transition realism, the framework actively prevents chronological and cognitive misalignments before data persistence occurs.

4.3. Significance and Comparison

Career Weave is much more flexible than a classic Applicant Tracking System (ATS), classic machine learning classification systems or classic adaptive learning systems. Unlike the use of preprogrammed templates, Generative Reasoning enables the system to understand the meaning of the skills, thus solving the "tutorial hell" caused by the consumption of unstructured and redundant information.

A detailed comparison of the existing logic models in personalized education/recruitment and the proposed Career Weave framework is presented in Table 2 to compare the important technical features and capabilities.

Table 2 is designed to give an unambiguous baseline, but an initial test was conducted on the synthetic testbed to quantify the degree of improvement (Rameshbabu.V et al., 2025). The One-Pass system consumed the same number of API calls as the Two-Pass system, but it only used the tokens for verified users, saving 40% latency when compared to the single-pass system with open curriculum prompts (Rameshbabu.V et al., 2025). For scenario testing (Rameshbabu.V et al., 2025) and testing of completely divergent career leaps, the Feasibility Gate was able to be used successfully to detect and block the leaps completely. With an impossible 100% accuracy rate, (The Impossible Leap). Since it involves reducing the rigidity of the templates that the older tools use, and it allows a detailed chronological scaling for them, the quality of the resulting roadmap is in direct proportion to the conceptual gap mapping, and not a generic “baseline” (Rameshbabu.V et al., 2025). There are other primary limitations in these current metrics: the metrics are based on the execution of a deterministic test suite, whereas metrics such as user acceptance rates, perceived usability scores and assessable learning retention require continuous real-world field studies (Rameshbabu.V et al., 2025).

Table 2: Comparative Analysis of Existing Logics vs. CareerWeave Framework.

Feature / Capability	**Traditional Rule-Based / ATS Logics**	**Standard ML Models (Clustering / Classification)**	**Adaptive Learning Systems (RL / Micro-Assessments)**	**Career Weave Framework**
Core Matching Mechanism	Static, exact string/keyword matching.	Statistical proximity and demographic grouping.	Real-time performance tracking & micro-loops.	**Generative Reasoning & Semantic Understanding**
Curriculum Generation	Preprogrammed, fixed templates.	Path selection from a predefined catalogue.	Item-by-item sequential adjustment.	**Dynamic, Open-Ended LLM Orchestration**
Skill Gap Assessment	Quantitative tallying (e.g., keyword count missing).	Categorical variance from a fixed baseline cluster.	Skill-level indexing via localized testing.	**Qualitative Delta Analysis (Conceptual Gap Mapping)**
Transition Guardrails	None (processes any inputs unconditionally).	Soft statistical anomalies or out-of-bounds rejections.	Limited to prerequisite progression locks.	**Feasibility Gate (Realism & Risk-Assessment)**
Pacing & Timeline Control	Linear, immutable schedules.	Historical average-completion pacing.	Micro-pacing based on immediate task accuracy.	**Dynamic Chronological Scaling (4 to 24-week adaptive windows)**
Information Delivery Control	Full curriculum exposure leading to cognitive overload.	Full visibility of the selected track archetype.	Linear adaptive step unlocking.	**Sequential Milestone Unlocking via Relational Redaction**

As described in Table 2, although the traditional architectures of adaptivity are good at adapting to micro tasks according to the student testing results (Wang et al., 2023), They do not outline the macro-level, long-term industry transition path (Mukashova et al., 2025). The two-pass orchestration engine is exclusive to Career Weave, designed to ensure that the learning journey is structurally digestible and viable and is highly tailored upfront to the learner before any content is served.

5. Discussion

The result of the experiment indicates that the LLM orchestrated guidance successfully links the guidance gap in professional development. The key point to highlight from the results is the Feasibility Gate's ability to logically reject unrealistic career transitions, something that cannot be achieved with a traditional keyword-based approach. This guarantees that curricula produced are based on realistic education time windows (4 to 24 weeks, depending on the technical gap identified).

The key strength of Career Weave is its Qualitative Delta Analysis, which produces an analysis that is not just a match of terms. Plus, Sequential Milestone Unlocking (database-level redaction) will help the learner avoid "tutorial hell" by allowing them to take a focused learning path and not get overwhelmed with information.

Real-time labor market data and some jargon in very specialized sectors have been seen to be very successful in solutions, but there is room for improvement (Dash et al., 2025). Moreover, exclusive use of proprietary LLM APIs means relying on external services' availability and potential algorithmic biases (Mukashova et al., 2025; Tapalova & Zhiyenbayeva, 2022). Such architectural reliance presents significant risks and concerns in terms of user-proof data, because integral inputs into the user profile need to pass through third-party endpoints (Rameshbabu.V et al., 2025). In addition, a generative curriculum engine has semantic hallucination problems as the program can make up de facto technical versions, or it can actually make logical connections that are simply inaccurate (Rameshbabu.V et al., 2025). Two or more different milestone configurations can get two different LLM outputs from the same user profile, so that it is difficult to get the same academic evaluation result from the same user profile across time. That leaves a long-term issue of determining whether a slight variation of the user profile matched to a slight variation of the user's milestones is the same academic performance as the LLM, especially when academic performance is standardized according to the milestones produced by the model (Rameshbabu.V et al., 2025).

From this study, it can be inferred that with strict schemas in structure, generative intelligence would be a powerful automated career coaching tool, as there would be no logic inconsistencies. Version 2 is previously scheduled to include live feeds from the job market and the ability to work in multiple languages, making it more applicable in the international arena.

6. Conclusion & Future Scope

In this research, we have succeeded in delivering an automated and personalized career guidance system and were able to show that their work was successful, and they implemented a two-pass LLM Orchestration scheme, which was effective. The framework successfully separates the context extraction task from curriculum development and includes a “Feasibility Gate” using logic, ensuring that all the roadmaps generated are firmly grounded in reality. These principles were implemented by using a strong and modular backend, built with FastAPI, a Python framework for developing APIs, and a MySQL database to support the Sequential Milestone Unlocking logic, keeping learners actively engaged while preventing them from getting trapped in 'tutorial hell'.

CareerWeave is really a solution to a big problem in a highly fragmented digital learning landscape, namely the need to offer structured learning experiences with content and learning activities that can be project-oriented, and then adaptable to fit loose professional goals, but also offer a scalable, flexible architecture that can align and cater to the broad range of content and experiences in digital learning.

The platform will be expanded and validated with ground truth in the future. Future releases will include multi-language localisation capabilities and real-time job market APIs (Dash et al., 2025) to take it to a global level, and automatically adjust to present job market trends and shifts in hiring requirements will be generated via milestone deliverables. In addition, the pedagogy of the curriculum developed will be gradually improved by implementing a facilitation process for community-based peer review and a mentoring process by the experts in the field (Katiyar et al., 2024). The additions will introduce automated reasoning and the supervision by professionals of human operators and will shape a hybrid ecosystem for verification.

Acknowledgement

The authors express their sincere gratitude to the Faculty of Engineering and Technology, Department of Computer Science & Engineering at Jagannath University, Jaipur, Raj., India, for providing the necessary institutional infrastructure, technical resources, and research bed facility to conduct this research work. Special thanks are extended to our supervisor, Dr. Om Prakash Sharma for his invaluable guidance, continuous academic supervision, and structural reviews throughout the development and compilation of this manuscript.

Conflict of Interest

The authors declare that there are no conflicts of interest regarding the publication of this manuscript. The research was conducted in the absence of any commercial, financial, or personal relationships that could be construed as a potential conflict of interest.

Funding Statement

The authors declare that this research work received no specific grant, financial aid, or funding from any public, commercial, or non-profit funding agencies. The project was entirely self-funded by the authors.

Authors' Contributions

The authors declare that this research work received no specific grant, financial aid, or funding from any public, commercial, or non-profit funding agencies. The project was entirely self-funded by the authors.

- **Keshav Dhanka:** Conceptualization, software architecture, backend engineering (FastAPI), database design (MySQL), logic validation, testing, and manuscript writing.
- **Ashish Saini:** UI/UX design, visual timeline architecture, prototyping, wireframing and graphic optimization.
- **Bhagawat Prasad:** Frontend responsive systems implementation, and component development (React Flow integration).
- **Dr. Om Prakash Sharma:** Academic supervision, methodology validation, structural review, final editing and review.

References

[1] Dash, S., Anchan, A., Dabre, P., Kulkarni, M. V. (2025). *AI Based Skill Gap Analyzer*. *13*(4), 644–657. https://doi.org/10.56975/ijedr.v13i4.302612

[2] He, Z., Xia, W., Dong, K., Guo, H., Tang, R., Xia, D., Zhang, R. (2022). Unsupervised Learning Style Classification for Learning Path Generation in Online Education Platforms. *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2997–3006. https://doi.org/10.1145/3534678.3539107

[3] Jeffson Preetham Dsilva, Murari B K. (2026). SkillSync: An Explainable AI Framework for Resume Evaluation, Skill Gap Analysis, and Career Alignment. *International Journal of Engineering Research*, *14*(01), 1–7. https://doi.org/10.5281/zenodo.18409162

[4] Kadam, A. J., Durge, T. J., Kale, C., Kute, P. R., Shelke, P. M., Deshpande, L. (2025). Personalized Learning Paths for Student Progression: An Artificial Intelligence-Driven Skill Roadmap with Retrieval-Augmented Generation and Large Language Models. *Cureus*, *2*(1), e7759/s44389-024-02275-z. https://doi.org/10.7759/s44389-024-02275-z

[5] Katiyar, Dr. N., Awasthi, V. K., Pratap, R., Mishra, K., Shukla, N., Singh, R., Tiwari, M. (2024). Ai-Driven Personalized Learning Systems: Enhancing Educational Effectiveness. *Educational Administration Theory and Practices*. https://doi.org/10.53555/kuey.v30i5.4961

[6] Mukashova, A., Tussupov, J., Serikbayeva, S., Mukhanova, A., Sergaziyev, M., Sambetbayeva, M., Yerimbetova, A., Lamasheva, Z., Sadirmekova, Z., & Ramazanova, V. (2025). AI-driven framework for automated competency formalization: From professional standards to adaptive learning outcomes. *Frontiers in Computer Science*, *7*, 1710358. https://doi.org/10.3389/fcomp.2025.1710358

[7] Ovtšarenko, O. (2023). Generation of a Learning Path in e-Learning Environments: Literature Review. *New Trends in Computer Sciences*, *1*(1), 32–50. https://doi.org/10.3846/ntcs.2023.18278

[8] Rameshbabu. V, Latha R, Sreenithi R. (2025). Skill Gap Analysis Using Machine Learning. *International Journal of Engineering Research and Sustainable Technologies (IJERST)*, *3*(3), 33–39. https://doi.org/10.63458/ijerst.v3i3.122

[9] Tapalova, O., Zhiyenbayeva, N. (2022). Artificial Intelligence in Education: AIEd for Personalized Learning Pathways. *Electronic Journal of E-Learning*, *20*(5), 639–653. https://doi.org/10.34190/ejel.20.5.2597

[10] Wang, S., Christensen, C., Cui, W., Tong, R., Yarnall, L., Shear, L., Feng, M. (2023). When adaptive learning is effective learning: Comparison of an adaptive learning system to teacher-led instruction. *Interactive Learning Environments*, *31*(2), 793–803. https://doi.org/10.1080/10494820.2020.1808794