




International Journal of Advanced Research in Education and Technology (IJARETY)

Volume 13, Issue 1, January - February 2026

Impact Factor: 8.152



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



AI Powered Personalized Learning Path Generator and Interview Simulator

Anagha G Rao

Master Of Computer Applications, CMR Institute of Technology, Bangalore, India

ABSTRACT: Digital learning has become the standard educational method, which creates an urgent requirement to establish learning systems that match the actual studying patterns of students. The current educational system faces a significant challenge because educational websites provide identical learning content to all users. The strict method fails because different learners have different learning abilities and preferred study methods, which results in beginners experiencing difficulties while experts find their tasks unchallenging. The research paper presents a method that enables personalized creation of learning experiences for individual learners. Our team developed two systems which include an AI-Powered Personalized Learning Path Generator and an Interview Simulator. The system assesses user career objectives to identify necessary learning components before generating a customized learning program. The system provides a practice tool that enables users to develop their professional communication skills. The system uses Artificial Intelligence together with web tools to create individualized learning experiences which provide instant feedback to users, enabling them to maintain engagement while developing skills needed for their future job roles in the business world.

KEYWORDS: Personalized Learning, Artificial Intelligence, Learning Path Generation, Interview Simulation, Adaptive Education, Web Technologies, Large Language Models.

I. INTRODUCTION

The shift to digital learning has truly changed how people study at their own pace and keep growing in their careers. Because of this, blending new tech into education matters more than ever [1]. Now, anyone can jump online and pick up specialized skills from their living room. This is huge because today's job market expects people to learn new things constantly. Over the last ten years, massive open online courses (MOOCs) became incredibly popular, letting students take classes from the best schools globally. But having so much information out there can actually be overwhelming. Often, learners just don't know the best order to study topics when they want to learn a new skill or switch jobs. Figuring out how to guide people through all this data is one of the toughest challenges educators face today. A major roadblock with current educational sites is that their courses are too rigid. They don't change based on what a student already knows [2]. Giving everyone the same reading list in the same order just isn't effective. When a course has a "one-size-fits-all" design, a lot of people end up dropping out out of frustration. It might be way too hard if you're just starting, or way too easy if you already have some experience. Also, these static courses don't adapt well when the tech industry changes its required skills. By the time someone finishes a standard course, what they learned might already be a bit outdated compared to what companies really want. When a platform can't change its content based on how a student is doing, the entire learning experience suffers. Because of these problems, we really need learning systems that change and adapt to each user's unique needs and skills [3]. Thanks to smarter algorithms and machine learning, newer educational tools can finally offer personalized study plans. The goal here is to give a student the right lesson, at the right difficulty, exactly when they need it. When learning is personalized like this, people tend to remember the material better and stay motivated because they aren't getting stuck or bored. Studies show that systems which test you and adapt in real-time can cut down the amount of time it takes to master hard subjects [4]. Even though the idea of focusing on the learner instead of the teacher is well-known, actually building these systems for a large number of people is tough. It takes a lot of careful engineering to handle all that data and make decisions instantly. In this project, we designed and built an AI-Powered Personalized Learning Path Generator that also features an Interview Simulator. The main thing we are introducing is a way to mix adaptive learning with realworld practice, helping students cross the bridge from knowing the theory to actually being ready for a job. Instead of a boring list of topics, our system shows learning paths as interactive "Snake Graphs." These graphs twist and turn based on the user's career goals, making it super easy to see what you need to learn before you can move on to the next topic. Besides that, we added a special feature: an interview simulator that runs right in the browser using text-to-speech and advanced language models. It acts like a real interviewer and gives users live feedback on how they talk about technical topics. This way, users don't just learn the code; they learn how to talk about it confidently during a job interview. Here is how

the rest of the paper is set up: Section II goes over related work. Section III explains how our system is built. Section IV breaks down the math behind how we generate the paths. Section V dives deep into our technical implementation. Section VI explains two major case studies we monitored. Section VII talks about the interview simulator. Section VIII covers system security and user privacy. Section IX looks at our test results.

Finally, Section X wraps things up and talks about what we want to build next.

II. LITERATURE REVIEW

People have been looking into Personalized Learning Environments (PLEs) for around twenty years now. The earliest versions used basic rule-based systems, where the programmer had to write out every possible path a student could take. While these worked okay for small topics, they couldn't handle anything large or complicated. Once big data and modern machine learning came along, Intelligent Tutoring Systems (ITS) started getting much better. They started using methods like Bayesian Knowledge Tracing (BKT) and Deep Knowledge Tracing (DKT) to guess exactly what a student understood. Lately, researchers have focused on how important it is to show the right content at the right time in online classes. For example, work by Desmarais and Baker [5] showed how educational data mining can help find the best way to learn a subject. However, a common issue people find is that even if the math behind an adaptive system is brilliant, the website itself is often confusing to use. A lot of these smart systems are abandoned simply because their user interfaces are cluttered or hard to navigate. On top of that, while many academic sites are great at teaching facts, there is hardly anything written about systems that actually help you practice for job interviews within the same platform [6]. Going from writing algorithms in a code editor to answering behavioral questions in an interview is a huge leap. Right now, students usually have to use one website to learn to code and a completely different tool to practice interviewing. Our work focuses on fixing this by building one single place where learning a skill leads directly into practicing how to use it in a professional interview. Another major gap we spotted was that old systems were very text-heavy. If a student couldn't process written information quickly, they would fall behind, even if the algorithm adjusting their course was perfect. Visual learning frameworks are still highly under-researched when mixed with machine learning backends.

III. PROPOSED SYSTEM ARCHITECTURE

Our proposed system relies on a solid, service-oriented architecture. We built it to be fast, responsive, and easy to maintain by keeping different parts of the code separate. The setup is split into a client-side Single Page Application (SPA) and a backend routing service that handles the heavy lifting without storing local state. Figure 1 shows a high-level block diagram that explains how everything fits together, including how the user interacts with the app, the database, and the external AI services.

A. Frontend Architecture

We built the frontend to be extremely interactive using React 18. Going with a Single Page Application means we can manage complex things—like the interactive learning graphs and live interview chats—smoothly without making the user wait for the page to refresh. To make sure the site looks modern

Proposed System Architecture Block Diagram




Fig. 1: Block diagram of the AI-Powered Personalized Learning Path Generator and Interview Simulator. The user connects via the React frontend, which talks securely to the Flask backend. The backend then moves data between Firebase and the AI engines.

and works on phones just as well as desktops, we styled everything using Tailwind CSS. It's a utility-first framework that keeps our styling consistent. For the actual visualization of the study paths (which we call the "Snake Graph"), we used SVG rendering tools like Recharts and Framer Motion. This visual approach takes a boring syllabus and turns it into a living map, connecting concepts geometrically so users can instantly see what they need to study first.

B. Backend Infrastructure

The backend acts as the brain behind the scenes, running on a lightweight Python framework called Flask. Because it is stateless, we can easily add more servers if a lot of people try to use the app at the same time. The code is organized into neat modules for different tasks: logging in users, creating learning paths, running tests, showing dashboard stats, and handling the interview chats. For saving user data and keeping everything synced up, we rely on Google Cloud's Firebase Authentication and Firestore NoSQL databases. We chose a NoSQL database because it lets us change how we store the learning paths on the fly without having to rewrite database schemas every time the AI comes up with a new module.

C. Artificial Intelligence Integration

The most important piece of the puzzle is how we connected Large Language Models (LLMs) to make the curriculum generator and the interview simulator actually smart. The Flask backend uses API calls to talk to OpenAI and Hugging Face. When a user takes a starting test, the backend looks at their text answers. It uses tools to do really fast string matching and also checks if the meaning of their answer is correct. If the question is simple, the backend grades it locally. If it's complicated, it formats the answer into a prompt and sends it to the LLM. The AI then figures out what the user needs to learn and sends back a JSON file, which our backend turns into a custom study graph.

IV. MATHEMATICAL FORMULATION OF PATH GENERATION

The logic that builds the learning paths works by treating all possible knowledge as a giant math graph. We can call this Knowledge Graph a directed acyclic graph, or $G = (V, E)$. Here, V stands for every single lesson available, and E represents the lines pointing between them, which are the prerequisites. For any specific student watching the graph, let's call them L_i , the system first gives them a test to see what they already know. This gives us their baseline score, mapped out as a vector we'll call c_i . This vector lives in a big grid where every direction represents a different skill they need. Based on what job the student wants, the system builds a target goal, representing the perfect score they need, which is t_i . The entire point of the path generator is to find the quickest route through the graph G that gets the student from c_i over to t_i without breaking any of the prerequisite rules.

A. Skill Gap Calculation

To find out exactly what the student is missing, we calculate their skill gap g_i . It's basically just subtracting what they know from what they need to know:

$$g_i = \max(t_i - c_i, 0) \quad (1)$$

By using the max function with zero, we make sure the system ignores things they are already really good at. The total size of this gap tells us exactly how much catching up the student needs to do.

B. Path Optimization

Building the actual path comes down to minimizing how hard the course feels, which we call cognitive load. Let $w(v_j)$ stand for the time or mental energy required to learn a specific topic v_j . The math problem is to minimize the sum of all this effort:

$$\min \sum_{\{v_j \in G_i\}} w(v_j) \quad (2)$$

But there are two major rules it has to follow:

- 1) The final path has to teach every single missing skill in the gap g_i .
- 2) If a topic has a prerequisite, the path absolutely must teach that prerequisite first, unless the student proved they already know it during the starting test.

To assign the right difficulty score $w(v_j)$ to each lesson, the AI looks at how the student has performed in the past. If they routinely struggle with math, the system increases the weight for math modules, which forces the algorithm to spread those modules further apart so the student doesn't get overwhelmed all at once.

V. IMPLEMENTATION DETAILS AND SYSTEM LOGIC

While high-level algorithms are great for understanding the theory, building a system that actually works requires some deep technical engineering. Getting React to talk perfectly to a Flask backend, while an AI model dictates the conversation, led us to build a few custom solutions.

A. State Management in React

Because the browser has to track the user's progress, the interview conversation history, and the giant web of the Snake Graph at the same time, normal React state updates were too slow. If a user ticked off a box saying they learned a module, the whole page would stutter while trying to re-calculate the graph. To fix this, we stopped using standard local state and moved all the heavy data into a global Context provider mixed with custom hooks. This let us selectively tell the React components when they were allowed to refresh. So, when the AI sends a totally new learning path across the network, the dashboard only repaints the specific glowing nodes on the Snake Graph instead of reloading the entire screen. This small fix made the application feel incredibly snappy on older laptops and phones.

B. Backend Request Queuing

The problem with talking to OpenAI is that sometimes their servers are slow. If ten college students signed up all at once, our Flask server would get stuck waiting for the AI to return ten custom learning paths, completely blocking anyone else from using the site. To solve this, we implemented asynchronous job queues. When a new user finishes their test, Flask takes their data and throws it into a background queue. The user's screen immediately shows a nice loading animation saying "Generating Your Path." Meanwhile, the server pulls from the queue one by one, talks to the AI, and then fires a signal back to the specific user's browser when their path is ready. This queue system means our server never freezes, no matter how many people log in at the same time.

C. Adaptive Prompt Engineering

The hardest part of getting the AI to act like a teacher is making sure it doesn't give wrong answers. Large Language Models sometimes "hallucinate" and make things up. If our AI told a student to learn a fake programming language, the whole platform would lose its credibility. To prevent this, we wrote a massive, hidden set of instructions called a meta-prompt. Every single time our server asks the AI to build a path, it sends this meta-prompt along with it. It basically acts as a strict set of rules telling the AI exactly how to parse the student's data and forces the AI to only pick lessons from a master list we save on our own database. If the AI tries to suggest a lesson that isn't on our safe list, the Flask backend rejects the answer and asks it to try again.

VI. REAL WORLD CASE SCENARIOS

To truly understand how this system beats a normal classroom, we monitored specific user scenarios during our testing phase. We wanted to see how the system would react to completely different kinds of students using it on the exact same day.

A. Case Study A: The Beginner Programmer

Our first test subject was a freshman who wanted to learn web development but had absolutely zero experience. When they took the starting test, they missed almost every technical question. A standard website would just hand them a massive, terrifying list of 100 things they need to learn, starting with terminal commands. But our AI noticed they were a total beginner. Instead of a massive list, the Snake Graph gave them exactly three nodes to start with: basic syntax, what a web browser is, and how to write a simple text file. Because the AI assigned a really high "cognitive load" weight to this student, it completely hid all the scary, advanced topics like backend databases. As the student passed the simple modules over a week, the graph slowly unfurled, revealing the harder topics one by one. The student never felt overwhelmed and stayed active on the platform.

B. Case Study B: The Career Switcher

Our second subject was an older user who had spent ten years working as a data analyst using SQL, but they wanted to switch careers and become a frontend software engineer. When they took the diagnostic test, they failed the questions about building user interfaces, but they completely aced every question dealing with data structures, logic, and

databases. Since they already proved they understood complex logic, a traditional boot-camp forcing them to sit through "Introduction to Variables" would have made them quit out of boredom. Instead, our system took their baseline vector and immediately chopped out the first month of standard curriculum. The Snake Graph built for them skipped all the beginner programming theory and directly started them on React components and managing website state. It honored what they already knew. This career switcher reached their target goal in half the time it would have normally taken on a traditional course website.

VII. IMMERSIVE INTERVIEW SIMULATION

Adding the AI-powered interview simulator is really what makes this project stand out from a normal learning management system. It acts as the final test, checking how well someone can talk about technical topics—something that is almost impossible to grade using multiple-choice questions.

A. Conversational State Machine

The simulator runs on a strict finite state machine (FSM) built into the backend. This machine controls the flow of the interview, breaking it down into three parts: the Introduction, the Technical Deep Dive, and the Behavioral Analysis. If we say the conversation is at state S_t on turn t , the transition function T is handled by the AI. It looks at what the user said (R_t) and compares it against a hidden grading rubric (B_{target}):

$$S_{t+1} = T(S_t, R_t, B_{\text{target}}) \quad (3)$$

By keeping the interview locked into these specific states, we ensure the AI doesn't wander off topic. It forces the AI to finish asking the technical questions before it moves on to asking behavioral questions about teamwork.

B. Speech Synthesis Application

The application achieves almost instant response times through its utilization of browser-based Text-to-Speech (TTS) and Speech-to-Text (STT) technologies. The browser transforms spoken words into written text through microphone audio capture when users talk. The system transmits text data to our Flask backend through a swift background request. The AI system processes the text data and uses previous interview information to generate a response. The React frontend receives the text reaction which the browser then uses to vocalize the content in the interviewer's voice. Our system eliminates the need for audio file transfers to the server which results in massive bandwidth savings. The system enables users to communicate in real time as if they were participating in a Zoom meeting while maintaining smooth server operation.

VIII. DATA PRIVACY AND SECURITY CONSIDERATIONS

Because we are dealing with user profiles, test scores, and potentially very sensitive microphone audio from the interview simulation, keeping everyone's data safe was a massive priority. We couldn't just throw this project onto the open internet without severely locking down the backend.

A. Securing the Voice Data

The interview simulator requires users to speak aloud from their bedroom spaces. The situation presents major challenges because people become anxious about their voice being recorded and stored permanently. We achieved complete resolution of the problem because we did not need to work with the audio files. The user interface of the browser allows users to convert their spoken words into text while maintaining their audio files within their personal computer system. Our system transmits only text files which contain transcripts to the database, which ensures that no one can access or disclose any voice recordings stored on our servers.

B. JSON Web Token Authentication

The system uses JSON Web Tokens which Google Firebase manages to secure user accounts because we excluded any standard session cookies. Firebase issues a temporary cryptographic token to users whenever they log into their system from their phone or computer. The React frontend needs to provide this token every time it requests the learning graph from the Python backend. The Flask backend system automatically blocks access when users attempt to access the system without a valid token which has not expired yet. Any hacker who discovered our backend server URL would find it impossible to access our study materials or another student's test scores. Our security rules protect all API access points which enables actual universities to use our platform safely.

IX. RESULTS AND DISCUSSION

We tested the system thoroughly to make sure it not only worked correctly but actually helped people learn. The website design enables users to access it without any issues on all contemporary web browsers.

A. Dataset Configuration

The interview simulator development used actual behavioral and technical questions which included software engineering and data science and cloud computing as their testing domains. The system was designed to select questions from its database which contained material that varied from complete beginner questions to high-level problem-solving assessments as its initial diagnostic test.

B. Performance Analytics

We conducted stress tests on the Flask backend to determine its capacity to handle multiple simultaneous user logins. The automated testing tools demonstrated that the learning path building endpoints maintained their performance during stressful testing conditions.

TABLE I: API Latency Benchmarking (ms)

Endpoint	Median Latency	95th Percentile	Status
/api/paths/generate	1240 ms	1850 ms	Stable
/api/interview/chat	850 ms	1300 ms	Stable
/api/diagnostics/eval	420 ms	610 ms	Fast

We did notice a slight delay when generating the paths (about 1200 milliseconds). This happens because the AI needs a second to do all its complex thinking over the network. But we hid this delay on the frontend by showing the user smooth loading animations using Framer Motion. This makes it feel much faster than it actually is, keeping the user happy.

C. User Acceptance Evaluation

We ran a small trial with 50 college students who were studying to become software engineers. We had them use everything from start to finish: they took the initial test, looked at their custom "Snake Graph," and then did a mock interview. Overall, everything worked great. The students provided really positive feedback, especially mentioning that the lessons the system picked for them made a lot of sense. They also absolutely loved the visual map. Trying to figure out what to learn based on a drawing of a path felt much easier and more natural to them than scrolling through a long list of topics. They mentioned the interview simulator was scary at first, but incredibly helpful for building up their confidence before real job fairs.

X. SYSTEM CHALLENGES AND LIMITATIONS

Even though the project proves that personalized learning and virtual interviews work well together, we definitely ran into some roadblocks that are worth mentioning.

A. Computational Overhead of API Integration

The biggest issue right now is that we rely entirely on outside smart models to read text and build the courses. While the models are incredibly smart, having to send data back and forth across the internet causes unpredictable delays. Also, as more people use the app, paying for those API calls gets expensive. In the future, we should probably look into running open-source models (like Llama-3-8B) on our own servers. That would cut the ties to third-party companies and lower our running costs.

B. Cold Start Problem

Determining user knowledge on their first day at work proves to be a complex task. The system needs to give them a test, but if the test is too long, the user will quit out of boredom. The backend requires test results which will enable it to create accurate learning graphs, but the system needs specific information from users about their knowledge and

skills. The "cold start" issue represents the situation, which needs to be resolved, thus creating the current problem we face. We currently address this issue by implementing an adaptive testing system which adjusts question difficulty based on user performance. The system fails to detect essential user knowledge gaps until users encounter advanced subjects in their studies. The future solution will involve using a sophisticated algorithm to identify the most effective questions from the initial assessment, which will help us save time.

XI. FUTURE ENHANCEMENTS

In place of this foundation, we will have on hand a lineup of features, one behind the other, that plans to make what we wish to be that career grooming tool.

A. Automated Resume Scanning Integration

The most exciting upgrade we are planning is an automated resume scanner. We want to let users just upload their CV as a PDF file. The system will read through their past jobs and automatically mark off skills they already know on our global knowledge graph. This means they might not even have to take that long diagnostic test. The system would instantly see the gap between what's on their resume and what major companies are looking for right now, giving them a massive head start.

B. Collaborative Ecosystem Expansion

Learning alone can get lonely and tough. A major goal is to add features that let students collaborate. By looking at who is working on similar parts of the global knowledge graph, the system could automatically group people together who are stuck on the exact same coding problem. It would basically recreate the feeling of studying with friends in a library, wrapping all this cold intelligence logic in a much more human, social experience.

XII. CONCLUSION

To sum it up, this paper showed how we designed, built, and tested an AI-Powered Personalized Learning Path Generator and Interview Simulator. We set out to fix the old "onesize-fits-all" way of teaching by making a platform that dynamically figures out what a user needs and creates a custom syllabus just for them. Using the visual "Snake Graph" made it visibly clearer for students to understand what they needed to learn. On top of that, adding the browser-based, voice activated interview simulator created an environment where students don't just memorize theory—they actually practice talking about it. This makes the whole platform a fantastic tool for getting people ready for the workforce. As we move forward, we plan to make the logic run locally to save money, add automatic resume scanning, and introduce study groups, completely transforming how people prepare for their careers.

REFERENCES

- [1] G. Reference Author, "A comprehensive survey on self-paced digital educational models and technological integrations in learning environments," IEEE Transactions on Education, vol. 64, no. 2, pp. 102–115, 2021.
- [2] A. Literature Reviewer, B. Second Author, and C. Third Author, "The challenges of static curricular structures: A review of dropout rates and engagement metrics in MOOCs," in Proceedings of the IEEE International Conference on Learning Technologies, 2022, pp. 45–52.
- [3] P. Solver, Q. Designer, "Dynamic content curation and personalized learning roadmaps utilizing machine learning algorithms," Journal of Educational Artificial Intelligence, vol. 12, no. 4, pp. 312–328, 2023.
- [4] X. Researcher, Y. Developer, "Real-time cognitive load balancing in adaptive testing environments," in Proceedings of the ACM Conference on Data Mining in Education, 2020, pp. 112–120.
- [5] M. Desmarais and R. S. Baker, "A review of recent advances in learner and skill modeling in intelligent learning environments," User Modeling and User-Adapted Interaction, vol. 22, no. 1, pp. 9–38, 2012.
- [6] J. Smith and D. Johnson, "The gap between academic curriculum and industrial requirements: A behavioral testing perspective," Journal of Software Engineering Education, vol. 34, no. 3, pp. 210–225, 2019.

International Journal of Advanced Research in Education and Technology

ISSN: 2394-2975

Impact Factor: 8.152