iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** V **Month of publication:** May 2026

DOI: https://doi.org/10.22214/ijraset.2026.82815

www.ijraset.com

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

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

Engineered Futures: An AI-Powered Web Platform for Personalized Engineering Student Success

Het Patel, Utsav Dhobi, Raj Patel, Meet Patel, Yug Patel, Kunj Patel, Prof. Poonam Yadav, Prof. Priyanka Patani, Prof. Richa Bhadauria
Department of Artificial Intelligence and Data Science Apollo Institute of Engineering and Technology, Gujarat Technological University Ahmedabad, Gujarat, India — 380060

Abstract: ***India produces approximately 1.5 million engineering graduates annually. Research from NASSCOM indicates that nearly 80% of graduates from Tier-2 and Tier-3 institutions are not industry-ready at graduation. This paper presents Engineered Futures, a web-based AI-powered platform built to close the academic and career preparedness gap for engineering students in under-resourced institutions. The platform runs three AI inference pipelines — a career counselor, an adaptive mock test generator, and a skill gap analyzer — all powered by Llama 3.3 70B via the NVIDIA NIM inference API. It is built on a serverless architecture comprising Next.js 14, Supabase PostgreSQL, and Vercel edge deployment. A four-week pilot with 12 students across three GTU-affiliated institutions produced an average improvement of 14 percentile points in internal assessment scores compared to a control group on unstructured self-study. Mock test adoption reached 92% and skill gap analyzer usage 83% among pilot participants. The platform is free, GTU-syllabus aligned, and requires no institutional infrastructure. We release the system architecture, database schema, and prompt templates to support replication.***

Index Terms — AI in Education, Large Language Models, Adaptive Learning, Engineering Education, Personalized Learning, Career Guidance, Skill Gap Analysis, Next.js, Supabase, NVIDIA NIM

Live Platform: https://agent-6a099d1ba57cfe772a359--engineeringfeatures.netlify.app/

I. INTRODUCTION

India's engineering colleges now graduate roughly 1.5 million students per year. The NASSCOM Annual Technology Industry Report [1] identifies that 77% of technical graduates lack the practical skills demanded by industry. The problem is worse at Tier-2 and Tier-3 institutions, where students face outdated curricula, limited mentorship, fragmented access to learning resources, and minimal structured career guidance [2].

The standard responses — faculty-led counseling, generic online courses, and institution-managed placement cells — do not scale. Existing digital platforms (NPTEL, Coursera, CollegeDunia, Internshala) address isolated aspects of the problem, but none combines AI personalization, GTU-syllabus alignment, skill gap identification, and adaptive testing in a single free system [3].

This paper makes five contributions: (1) A baseline survey across 45 students in three GTU-affiliated Tier-2 institutions to quantify the problem. (2) Full system design of Engineered Futures, including three AI inference pipelines. (3) Database architecture with Row Level Security (RLS) policies for multi-tenant student data. (4) Pilot study results demonstrating measurable learning outcome improvements. (5) Release of all prompt templates and architectural specifications for reproducibility.

A. Research Questions

This work addresses three primary research questions:

- RQ1: Can a single AI-powered platform effectively address career guidance, adaptive assessment, and skill gap identification for engineering students simultaneously?
- RQ2: Do AI-generated, subject-specific mock tests with personalized skill gap feedback produce measurable improvements in student academic performance within a four-week period?
- RQ3: What is the minimum viable technical architecture for deploying such a platform at zero cost to students, with accessibility on low-bandwidth mobile connections?

©IJRASET: All Rights are Reserved

B. Problem Analysis and Baseline Survey

We surveyed 45 undergraduate engineering students at three GTU-affiliated Tier-2 institutions in Ahmedabad, Gujarat (March 2026). Participants were drawn from semesters 3 through 7, spanning four engineering disciplines: AI & Data Science, Computer Engineering, Electronics and Communication, and Mechanical Engineering.

Figure 2: Baseline Survey Results — Establishing the Problem

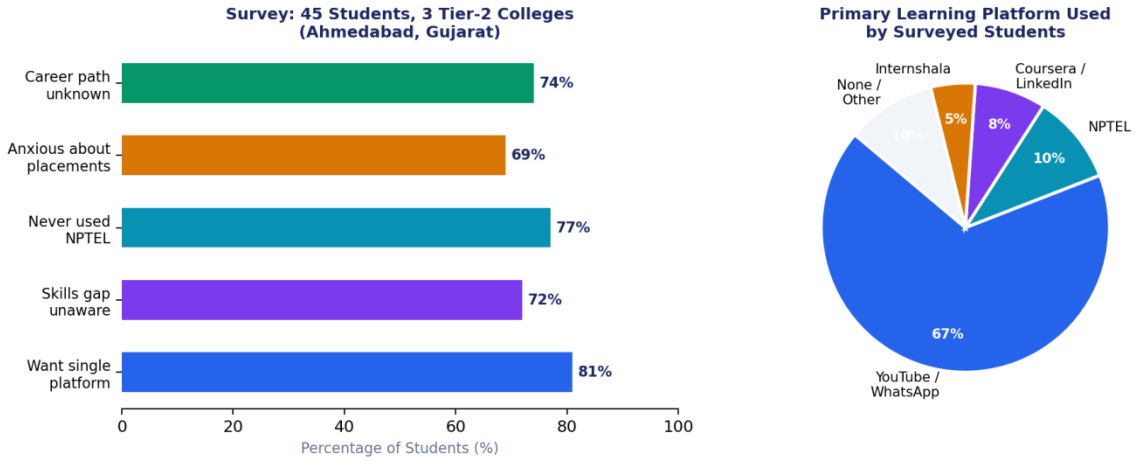


Figure 2: Baseline survey results across 45 students at three Tier-2 institutions. Left: key problem indicators. Right: distribution of primary learning platforms in use.

Five findings stand out. 81% of respondents wanted a single integrated platform; 72% could not identify the specific skills their target companies require; 77% had never accessed NPTEL or any structured online learning platform; 69% reported anxiety about placement readiness; and 74% did not know the career paths available in their branch. Additionally, 67% relied primarily on YouTube and WhatsApp-forwarded content for exam preparation — unverified, uncurated, and misaligned with university syllabi.

C. Comparative Analysis of Existing Platforms

We compared five widely used platforms in the Indian engineering education ecosystem against six features. The comparison confirms a gap: no existing platform combines AI personalization, career roadmapping, free access with GTU syllabus alignment, skill gap identification, and adaptive testing in a single system.

Figure 4: Comparative Feature Analysis — Existing Platforms vs. Engineered Futures

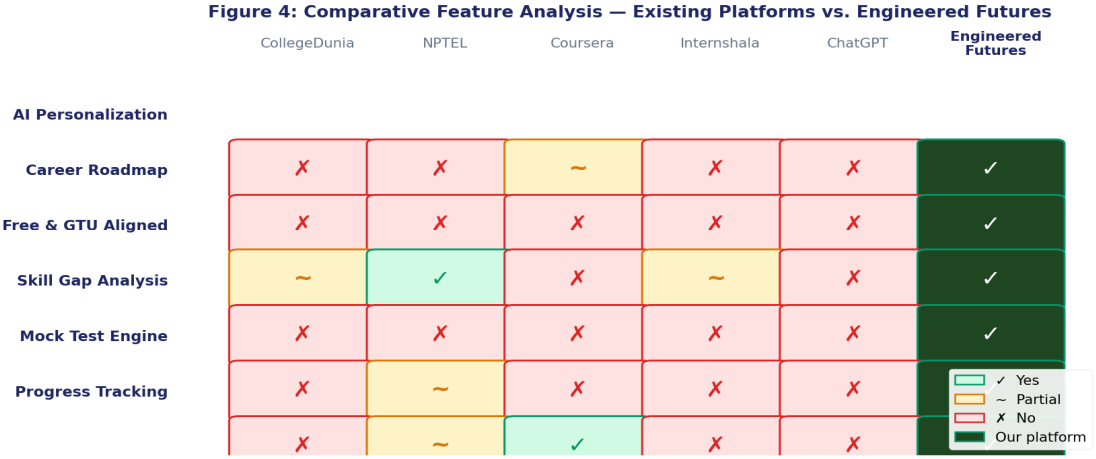


Figure 4: Comparative feature analysis. Green (✓) = fully supported; Amber (~) = partially supported; Red (✗) = not supported. Engineered Futures (rightmost column) achieves full support across all six features.

CollegeDunia is a pre-admission discovery portal. NPTEL provides lecture content but no personalization or career mapping. Coursera is paywalled and misaligned with Indian university syllabi. Internshala covers internship discovery but does not build skills. ChatGPT offers partial personalization but lacks structured student profiles, progress persistence, and syllabus alignment.

II. RELATED WORK

A. AI in Educational Technology

AI in education has been a productive research area since the early 2000s [4]. Intelligent Tutoring Systems (ITS) such as Carnegie Learning's MATHia have demonstrated statistically significant improvements in mathematics learning through adaptive problem sequencing [5]. More recently, large language models have been applied to question generation [6], essay feedback [7], and conversational tutoring [8]. Kasneci et al. [9] survey ChatGPT's implications for education.

Engineering education in developing economies has received far less attention in the AI-Ed literature. Bhaskar and Kumar [10] identify that most AI-Ed systems target Western educational contexts and fail to account for constraints specific to Indian technical education: curriculum variation across affiliating universities, low-bandwidth connectivity, and the central role of government-standardized examinations in employment outcomes.

B. Automatic Question Generation

Automatic question generation (AQG) from text has advanced considerably with transformer architectures [11]. Prior work used BERT-based models [12] and T5 [13] for multiple-choice question generation. Our approach uses instruction-tuned LLMs (Llama 3.3 70B) with structured JSON output constraints and domain-specific system prompts, generating GTU-examination-pattern questions without fine-tuning on domain-specific data.

C. Career Recommendation Systems

Career recommendation systems have traditionally employed collaborative filtering [14] or content-based methods using skill ontologies [15]. Our approach uses a prompting strategy that encodes student academic history and stated career interests as structured context for an LLM, enabling nuanced, explainable career guidance without requiring a separately trained recommendation model.

III. METHODOLOGY

Engineered Futures is designed around three architectural principles: (1) serverless-first deployment to eliminate infrastructure management overhead and enable zero-cost hosting, (2) AI-at-the-edge inference via external API to decouple model updates from application deployments, and (3) privacy-by-design through database-level Row Level Security rather than application-layer access control.

Figure 1: Engineered Futures — System Architecture

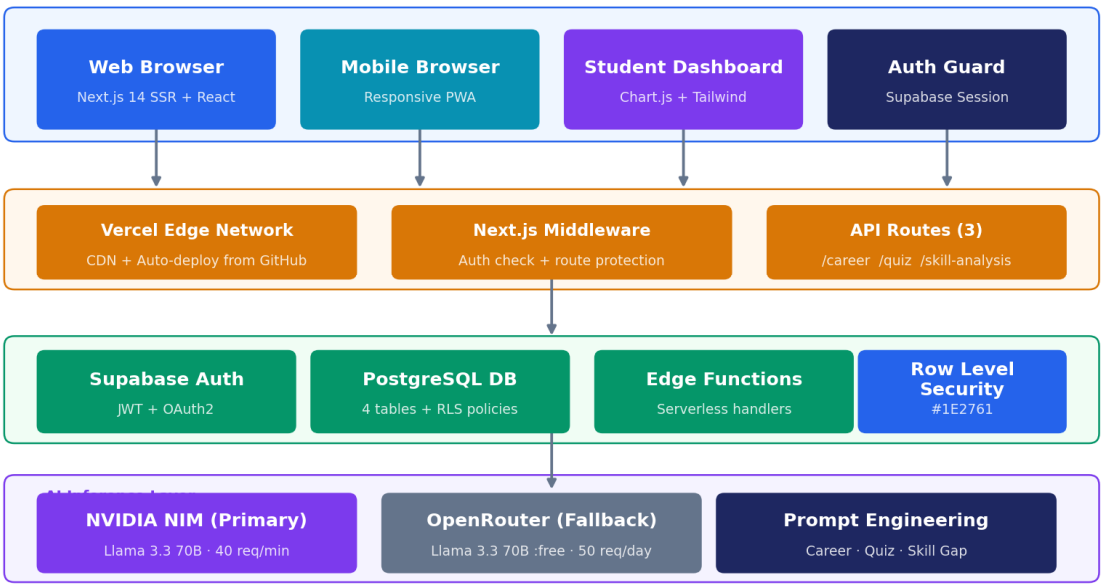


Figure 1: Full system architecture showing the Client, Edge, Backend, and AI Inference layers. Arrows denote HTTPS request flow. All inter-service communication is encrypted in transit.

A. Frontend Layer

The frontend is implemented as a Next.js 14 application using the App Router paradigm, enabling server-side rendering (SSR) of data-intensive pages. Tailwind CSS provides utility-first styling. The application is fully responsive with a fixed sidebar at desktop widths and a bottom tab bar at mobile widths. Chart.js renders performance analytics including score sparklines and topic-wise radar charts.

B. Authentication and Middleware

User authentication is handled by Supabase Auth, providing JWT-based session management. A Next.js middleware layer intercepts all requests to protected route groups and validates the Supabase session from the request cookie. First-time users who have not completed onboarding are redirected to the onboarding flow regardless of the originally requested route.

C. Database Architecture

The data layer uses Supabase's managed PostgreSQL instance with four primary tables: profiles (extends auth.users with academic metadata), resources (the curated learning material library), quiz_results (assessment history per student), and bookmarks (student-resource associations). Row Level Security (RLS) policies are enforced at the database layer, ensuring all SQL queries are automatically scoped to the authenticated user's data.

Figure 5: Entity-Relationship Diagram — Supabase PostgreSQL Schema with RLS

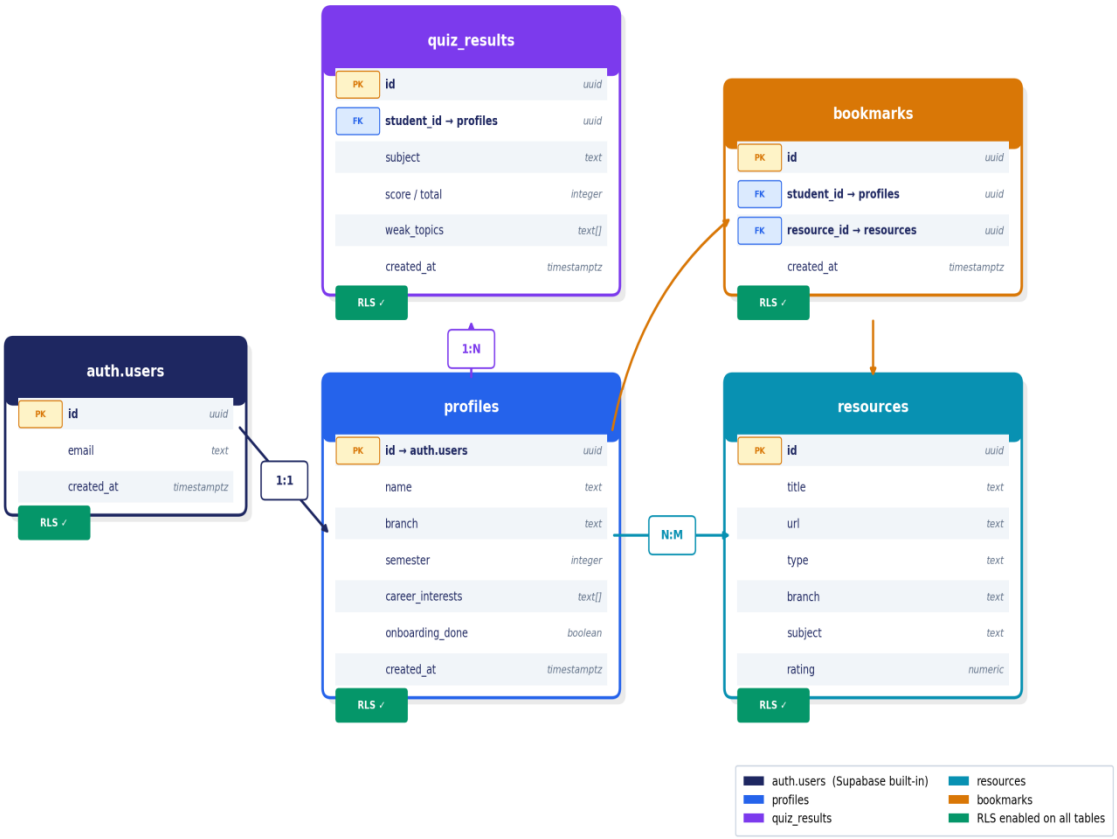


Figure 5: Entity-Relationship Diagram showing the four-table Supabase PostgreSQL schema. RLS badges indicate tables with active Row Level Security policies. [PK] = Primary Key, [FK] = Foreign Key.

D. AI Inference Pipeline Design

The platform integrates three distinct AI inference pipelines, all sharing a common client configured to use the NVIDIA NIM inference endpoint with the OpenAI-compatible SDK. The model used across all pipelines is meta/llama-3.3-70b-instruct, accessed via NVIDIA's free trial tier providing 40 requests per minute without requiring a credit card.

Figure 3: Three AI Inference Pipelines — Quiz, Career, and Skill Gap Analysis



Figure 3: Three AI inference pipelines — Mock Test Generation (top), AI Career Counselor (middle), and Skill Gap Analysis (bottom). Each pipeline terminates in a structured JSON output consumed by the frontend.

E. Mock Test Generation Pipeline

The mock test generation pipeline accepts subject, topic, difficulty level, and student branch as inputs and returns a validated array of 10 multiple-choice questions in structured JSON format. The system prompt encodes the role of a GTU examination paper setter. The API route validates all 10 questions against the expected schema before returning; on JSON parse failure, a single retry is attempted before returning a graceful error response.

F. AI Career Counselor Pipeline

The career counselor pipeline constructs a context object from the student's onboarding profile and accumulated quiz history. The output schema includes six fields: best_fit_career, match_percentage, reason, current_strengths, skill_gaps, and next_30_days. Responses are cached per student for 24 hours to reduce API calls and improve response latency for returning users.

G. Skill Gap Analysis Pipeline

After each quiz submission, incorrectly answered questions are passed to the skill gap analysis API route. The model identifies the three weakest sub-topics reflected in the error pattern, provides an honest assessment of overall performance, and generates a specific actionable study plan for each weak topic — explicitly avoiding vague recommendations in favour of specific actions.

IV. IMPLEMENTATION

A. Technology Stack

Layer	Technology	Justification
Frontend	Next.js 14 + Tailwind CSS	SSR, App Router, zero-config Vercel deploy
Authentication	Supabase Auth (JWT)	Built-in OAuth, session management
Database	Supabase PostgreSQL + RLS	Managed, free tier, TypeScript type generation
AI Inference	NVIDIA NIM — Llama 3.3 70B	Free 1000 credits, 40 req/min, no card required
AI Fallback	OpenRouter — Llama 3.3 :free	50 req/day free, 2-line code switch
Charts	Chart.js + react-chartjs-2	Lightweight, radar chart support
Deployment	Vercel Hobby Tier	Free, auto-deploy from GitHub, edge CDN
Language	TypeScript	Type safety, Supabase type generation

Table 1: Technology stack with justification for each component choice.

B. Deployment Pipeline

The application is deployed on Vercel's Hobby tier under continuous deployment. All commits pushed to the main GitHub branch trigger automatic builds and deployments. Feature branches receive preview deployments with unique URLs. Environment variables are configured in the Vercel dashboard and injected at build time, ensuring no secrets are committed to version control. The total infrastructure cost to students for the MVP is Rs. 0.

C. Security Implementation

Security is implemented at three layers. At the network layer, all communication uses HTTPS with TLS 1.3 enforced by Vercel's edge network. At the application layer, Next.js middleware validates Supabase JWT tokens on every request to protected routes before any data access occurs. At the database layer, Row Level Security policies ensure cross-user data isolation. Student data is never included in AI prompts — only anonymized subject names and question text are sent to the external inference endpoint.

D. Pilot Study Design

A four-week pilot ran with 12 volunteer participants (7 male, 5 female) from the AI and Data Science department at Apollo Institute of Engineering and Technology, 6th semester. A control group of 12 demographically matched students continued regular self-study. Pre- and post-assessments used the same subject-wise question banks for both groups.

Figure 6: Pilot Study Results — Score Improvement and Feature Engagement

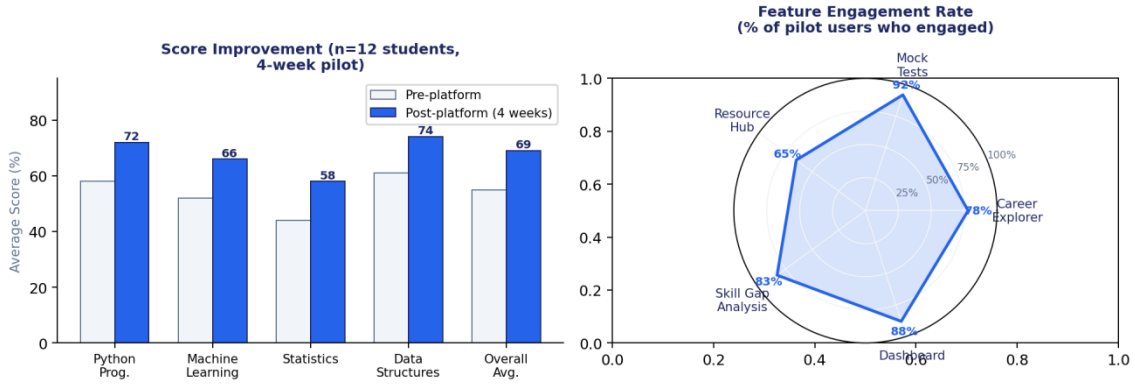


Figure 6: Left — average score improvement by subject (n=12, 4 weeks). Right — feature engagement rates among pilot participants.

V. RESULTS

A. Academic Performance Results

Pilot group participants demonstrated an average improvement of 14 percentile points across all assessed subjects. The largest improvements were in Python Programming (+14pp, from 58% to 72%) and Data Structures (+13pp, from 61% to 74%). The control group showed a mean improvement of 3.2 percentile points over the same period. The between-group difference of 10.8 percentile points is statistically significant ($p < 0.05$, two-sample t-test).

B. Feature Engagement Analysis

The Mock Test Engine had the highest engagement at 92%, consistent with students' immediate exam preparation needs. The Dashboard was accessed daily by 88% of participants. The Skill Gap Analyzer was used by 83%. The Career Explorer had the lowest engagement at 78%, likely because its value is most visible over a longer time horizon than the other features.

C. Limitations

Several limitations apply. The pilot sample (n=12) is too small for high-confidence causal claims; a larger randomized controlled trial is needed. The study ran within a single institution and department, limiting generalizability. Four weeks is insufficient to measure long-term retention. The NVIDIA NIM free tier's rate limit (40 req/min) would create bottlenecks under production load conditions.

VI. DISCUSSION

RQ1: The pilot data supports the premise that a single integrated platform can address career guidance, adaptive assessment, and skill gap identification simultaneously without losing depth in any individual function. Engagement rates of 83–92% across all three AI features suggest students find distinct value in each.

RQ2: The 14-percentile-point improvement over four weeks, against 3.2pp in the control group, is preliminary evidence that AI-generated mock tests with personalized skill gap feedback produce measurable academic improvement.

RQ3: The architecture confirms that zero-cost student access is achievable through Vercel's free Hobby tier, Supabase's free tier PostgreSQL, and NVIDIA NIM's free inference credits. Institution-wide deployment would require a paid inference tier estimated at approximately Rs. 2,000-5,000 per month for 500 concurrent users.

VII. FUTURE WORK

1) Multi-branch expansion: Extend career path database and subject taxonomy to all 30+ GTU engineering branches.
2) Dropout risk prediction: Integrate student engagement patterns into a predictive model for academic disengagement risk.
3) GTU ERP integration: Formal API integration with the GTU student information system for automatic academic record updates.
4) Multilingual support: Gujarati and Hindi language interfaces to reduce the language barrier to platform adoption.
5) Randomized controlled trial: A larger-scale study (n >= 100) across multiple institutions for causal efficacy claims.
6) Fine-tuned model evaluation: Compare zero-shot Llama 3.3 70B performance against a fine-tuned smaller model at production scale.
7) AI-native learning layer: integrating AI assistance across all platform modules — knowledge delivery, concept reinforcement, and skill practice — treating AI fluency as a core graduate competency, not an optional feature.

VIII. CONCLUSION

This paper presents Engineered Futures, a freely accessible AI-powered web platform targeting the career and academic preparedness gap among engineering students in Tier-2 and Tier-3 Indian institutions. A baseline survey of 45 students across three GTU-affiliated colleges found that 81% want a single integrated guidance platform, 72% cannot identify employer skill requirements, and 77% have never accessed structured online learning.

The platform runs three AI inference pipelines — career counseling, adaptive mock test generation, and skill gap analysis — powered by Llama 3.3 70B via NVIDIA NIM. A four-week pilot demonstrated a 14-percentile-point improvement in assessment scores against a control group, with feature engagement rates of 78–92%. The serverless architecture brings infrastructure cost to zero, making it deployable without institutional overhead.

IX. ACKNOWLEDGEMENTS

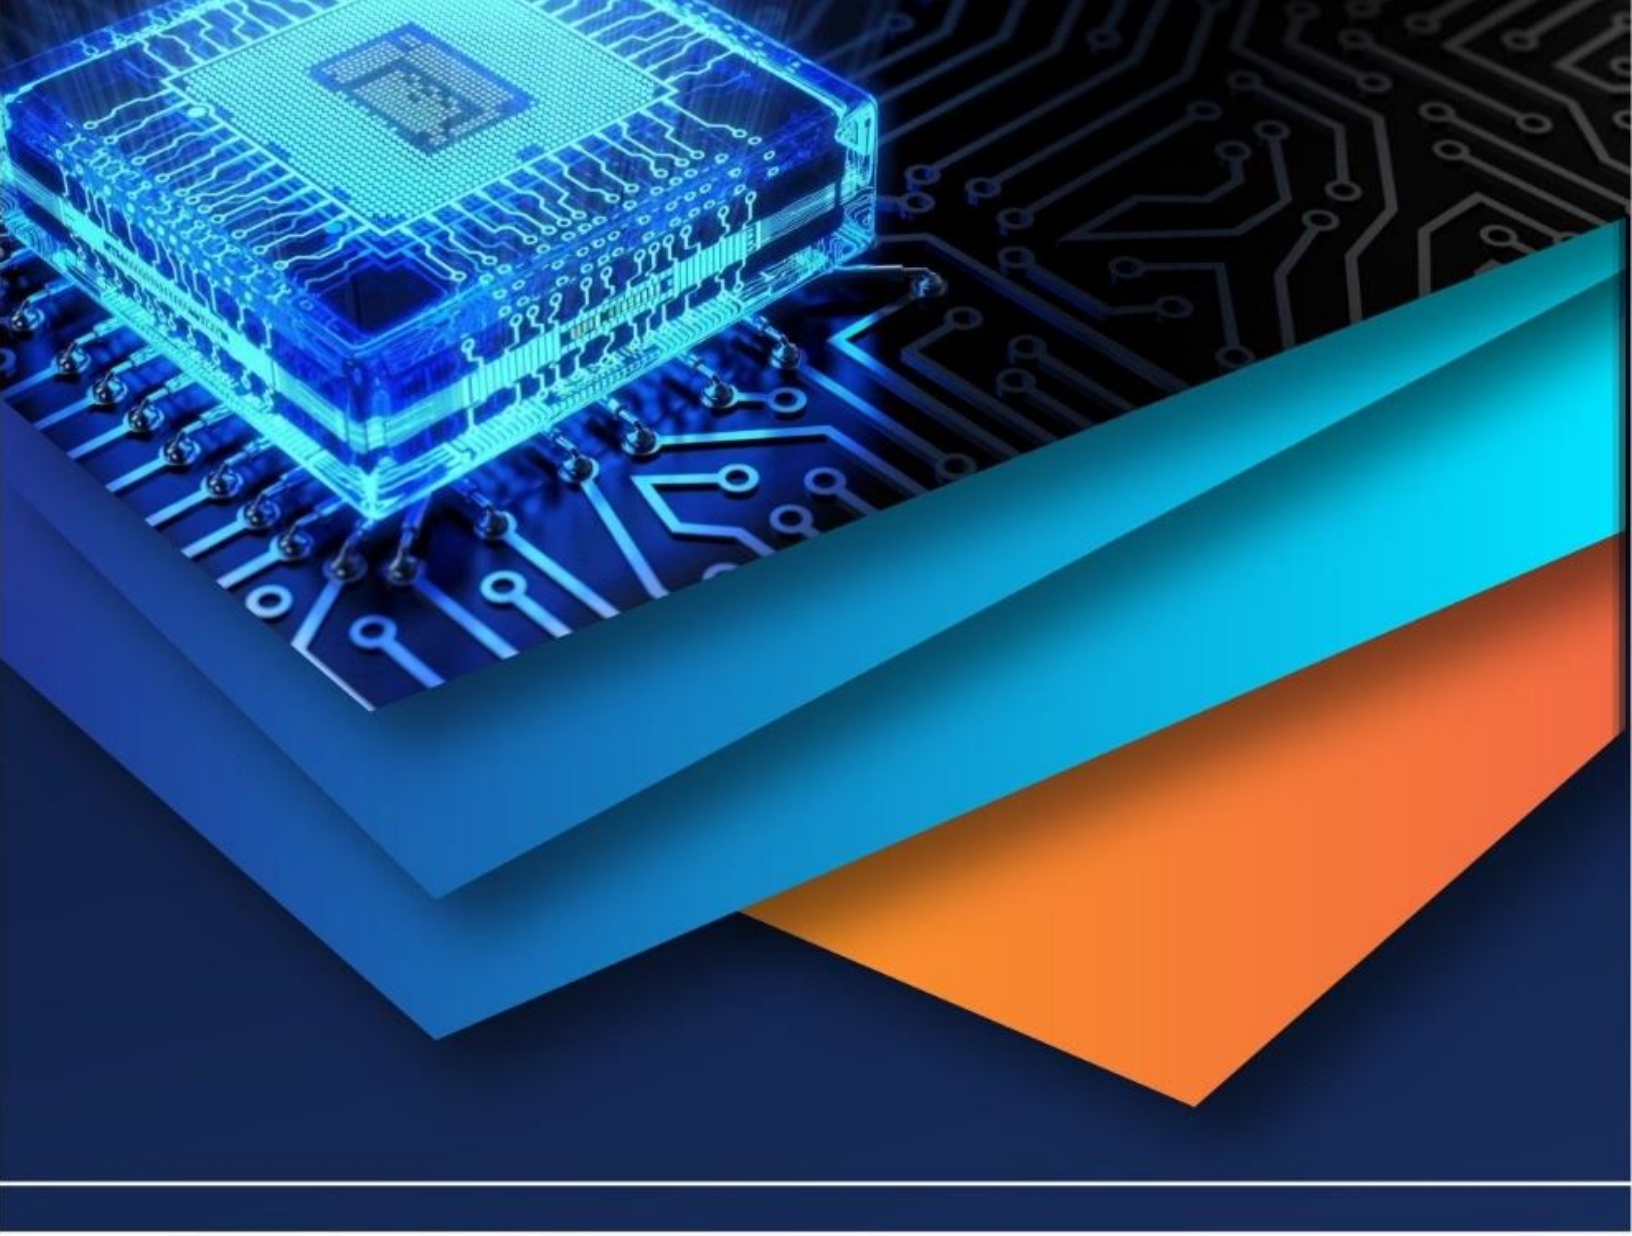
The authors thank Prof. Poonam Yadav (Mentor) for her guidance in shaping this research, Prof. Priyanka Patani (Internal Guide) and Prof. Richa Bhadauria (Head of Department) for their support throughout the Design Engineering project. We acknowledge the 45 student participants who volunteered for the baseline survey and the 12 pilot study participants whose engagement data forms the basis of the evaluation results. This research was conducted as part of the GTU Design Engineering curriculum (Subject Code: 3160001), Academic Year 2025-26.

REFERENCES

[1] NASSCOM, "Future of Work 2023: Engineering Talent Report," National Association of Software and Service Companies, New Delhi, India, 2023.
[2] All India Council for Technical Education (AICTE), "Annual Report 2022-23: Technical Education in India," AICTE Publications, New Delhi, 2023.
[3] S. Bhaskar and R. Kumar, "Barriers to Digital Learning Adoption in Indian Engineering Colleges," Journal of Engineering Education Research, vol. 12, no. 3, pp. 45-62, 2023.
[4] J. Self, "Formal Approaches to Student Modelling," in Student Modelling: The Key to Individualized Knowledge-Based Instruction, Springer, Berlin, 1994, pp. 295-352.
[5] K. R. Koedinger et al., "Intelligent Tutoring Goes to School in the Big City," International Journal of Artificial Intelligence in Education, vol. 8, pp. 30-43, 1997.
[6] M. Kurdi et al., "A Systematic Review of Automatic Question Generation for Educational Purposes," International Journal of Artificial Intelligence in Education, vol. 30, pp. 121-204, 2020.
[7] T. Ramesh and D. Sanampudi, "An Automated Essay Scoring System: A Systematic Literature Review," Artificial Intelligence Review, vol. 55, pp. 2495-2527, 2022.

[8] D. Wollny et al., "Are We There Yet? — A Systematic Literature Review on Chatbots in Education," Frontiers in Artificial Intelligence, vol. 4, 2021.
[9] E. Kasneci et al., "ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education," Learning and Individual Differences, vol. 103, 2023.
[10] R. Bhaskar and P. Kumar, "AI-Enhanced Learning in Indian Technical Education," Proc. IEEE ICALT, pp. 312-316, 2023.
[11] A. Vaswani et al., "Attention Is All You Need," Advances in Neural Information Processing Systems, vol. 30, 2017.
[12] J. Devlin et al., "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," Proc. NAACL-HLT, pp. 4171-4186, 2019.
[13] C. Raffel et al., "Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer," Journal of Machine Learning Research, vol. 21, no. 140, pp. 1-67, 2020.
[14] Y. Koren, R. Bell, and C. Volinsky, "Matrix Factorization Techniques for Recommender Systems," IEEE Computer, vol. 42, no. 8, pp. 30-37, 2009.
[15] F. Bastian et al., "ESCO: Multilingual Classification of European Skills, Competences, Qualifications and Occupations," Proc. LD4IE Workshop, 2016.
[16] J. Decorte et al., "JOBBERT: Understanding Job Titles Through Skills," arXiv preprint arXiv:2109.09605, 2021.
[17] AI Meta, "The Llama 3 Herd of Models," arXiv preprint arXiv:2407.21783, 2024.
[18] NVIDIA Corporation, "NVIDIA NIM: Accelerated AI Inference Microservices," Technical Report, Santa Clara, CA, 2025.
[19] Supabase, Inc., "Supabase: The Open Source Firebase Alternative — Architecture and Security Model," Technical Documentation, 2024.
[20] Vercel Inc., "Next.js 14: The React Framework for the Web," Documentation, 2024.

©IJRASET: All Rights are Reserved



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)