



# AI-Curated Skill Pathways: Intelligent Course Aggregation for Engineering Students

Devesh Mukesh Badgajar<sup>1</sup>, Abhishek Rajesh Khatri<sup>2</sup>, Prof. Mayuri S. More<sup>3</sup>  
<sup>1,2,3</sup>Computer Science & Engineering, Siddhivinayak Technical Campus, Shegaon, Maharashtra, India

DOI: 10.5281/zenodo.19539835

## ABSTRACT

*Engineering students relying on YouTube for self-study face a common problem: too many tutorials, no clear order, and no way to judge quality. This paper describes a platform that automatically builds structured learning roadmaps by pulling the best YouTube tutorials for each topic in a given domain. The system scores videos across six factors — engagement, teaching quality, comment sentiment, topic coverage, recency, and presentation — then assembles the top-ranked videos into a sequenced roadmap. A pilot with 120 students across eight engineering domains showed a 96% drop in time spent finding resources, an average content quality score of 85.3%, and a user satisfaction rating of 4.2/5.0.*

**Keywords:-** content curation, YouTube API, NLP sentiment analysis, personalized learning, engineering education, recommendation system

## 1. INTRODUCTION

YouTube has become one of the most widely used sources for technical learning, with millions of tutorial videos across every engineering discipline. The problem is not a lack of content — it is that students have no reliable way to figure out which videos are worth watching, in what order, or how they fit together into a complete learning path.

University curricula often focus on theory and fall short on practical, job-ready skills. According to NACE's 2025 Job Outlook survey, a significant share of new graduates enter the workforce missing at least two hands-on skills employers expect[1]. Students trying to fill this gap independently run into three recurring issues:

- **Volume Overload** — A search for "machine learning tutorial" on YouTube returns millions of results. Students typically spend 10–15 hours comparing options before they start learning anything.
- **Content Fragmentation** — Tutorials from different creators use different terminology and assume different background knowledge, which creates gaps.
- **No Roadmap** — Without a clear progression from basics to advanced topics, students skip prerequisites or waste time on the wrong things[2].

Paid platforms like Coursera and Udemy solve the structure problem but cost \$30–\$300 per course, which is out of reach for many students. Meanwhile, YouTube has excellent free tutorials — they just aren't organized. This paper presents a system that closes that gap: it automatically discovers YouTube videos, scores them for quality, and assembles the best ones into structured roadmaps, for free.

### Key contributions:

1. A six-factor evaluation framework scoring videos on engagement, teaching quality, sentiment, topic coverage, recency, and presentation.
2. A weighted composite scoring algorithm that selects the single best video per topic from across multiple YouTube creators.
3. A BERT-based comment sentiment analysis module trained on educational feedback.
4. A system covering eight engineering domains with weekly automatic updates.
5. An empirical study with 120 students confirming a 96% time reduction and 4.2/5.0 satisfaction.

## 2. RELATED WORK

AI-driven personalized learning has been explored from several directions. Fortuna (2025) reviewed 80+ studies and found adaptive systems increase student engagement 23–41% over static course delivery, while stressing that content quality — not just personalization — must be part of the equation[3].

SmartLearnHub built a MERN-stack adaptive web app where students following AI-curated paths completed 31% more topics than those on a fixed syllabus[4]. Zhao (2025) took this further using reinforcement learning that learns from click patterns, quiz results, and time-on-task data, achieving 19% better pathway optimization than standard collaborative filtering[5].

On quality assessment, Mayormente (2025) found a moderate correlation ( $r = 0.42$ ) between NLP sentiment scores on student comments and expert quality ratings[6]. YouTube's own recommendation system uses a two-



stage neural network — candidate generation then ranking — to maximize engagement[7]. None of these prior systems combine free YouTube content, automated multi-factor quality scoring, multicreator path assembly, and NLP sentiment validation in a single platform. Table 1 summarizes these differences.

Feature	Coursera	SmartLearnHub	Zhao[5]	Proposed
Free Content	No	No	N/A	Yes
AI Quality Scoring	No	Partial	Yes	Yes (6 factors)
Multi-Creator Paths	No	No	No	Yes
NLP Sentiment	No	No	No	Yes
Auto Path Updates	Manual	Manual	Partial	Full
Domains Supported	Varies	1	1	8

Table 1: Comparison with existing approaches

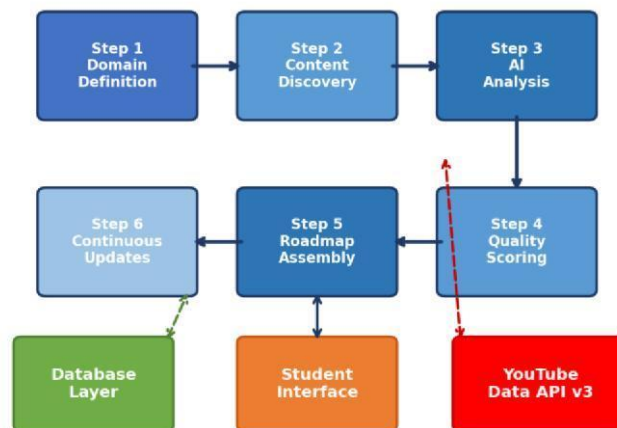
### 3. METHODOLOGY AND SYSTEM DESIGN

#### 3.1 System Architecture

The platform is built as a six-stage pipeline: Domain Definition → Content Discovery → AI Analysis → Quality Scoring → Roadmap Assembly → Continuous Updates. Each stage is a separate module so parts can be updated independently.

The six core components are:

- **User Interface:** dashboard showing roadmaps, progress, and recommendations
- **YouTube Data Collection:** retrieves video metadata and comment threads via API
- **AI Analysis Engine:** computes scores across all six dimensions
- **Roadmap Curation:** runs the weighted optimization and picks the top video per topic
- **Database:** stores video metadata and student progress
- **Recommendation Engine:** uses collaborative and content-based filtering for additional suggestions



#### 3.2 Domain Definition

Eight computer science domains are supported, each structured as a linear topic sequence from beginner to advanced. Roadmaps were built by analyzing industry skill frameworks, job posting requirements, and university curricula.

Table 2: Supported learning domains

#	Domain	Topics	Example Skills
1	Web Development (Full Stack)	18	HTML, React, Node.js, MongoDB
2	AI / Machine Learning	22	Python, TensorFlow, NLP, CNNs
3	Data Science	16	Pandas, SQL, Visualization
4	Cloud Computing	14	AWS, Azure, Docker, Kubernetes
5	Blockchain Development	12	Solidity, Smart Contracts, DApps
6	DevOps	15	CI/CD, Jenkins, Terraform, Linux
7	Mobile App Development	17	Flutter, React Native, Firebase
8	Cybersecurity	14	Pen Testing, Wireshark, OWASP

#### 3.3 Content Discovery

For each topic, the system queries YouTube Data API v3 with the topic name combined with keywords like "full course" and "beginner to advanced." Results are filtered to the past 24 months. Up to 50 candidate videos are

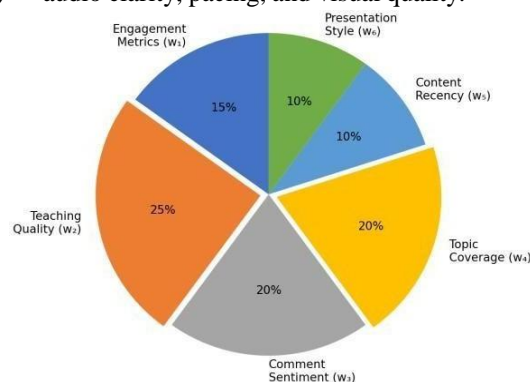


collected per topic, capturing title, description, channel, subscriber count, views, likes, comment count, duration, and publish date.

### 3.4 AI Multi-Factor Evaluation Engine

Each video is scored on six independent dimensions:

1. **Engagement** ( $w_1 = 0.15$ ) — normalized views, like-to-view ratio, and completion rate, adjusted for channel size to avoid large-channel bias.
2. **Teaching Quality** ( $w_2 = 0.25$ ) — highest weight; evaluates instructional clarity, logical progression, and structure via transcript analysis.
3. **Comment Sentiment** ( $w_3 = 0.20$ ) — NLP classification using fine-tuned BERT.
4. **Topic Coverage** ( $w_4 = 0.20$ ) — checks completeness against the domain subtopic checklist.
5. **Content Recency** ( $w_5 = 0.10$ ) — exponential decay function favoring newer content.
6. **Presentation Style** ( $w_6 = 0.10$ ) — audio clarity, pacing, and visual quality.



**Fig. 2. Weight distribution in the multi-factor quality evaluation engine.**

The top 200 comments per video are retrieved and classified as positive, neutral, or negative by a BERT model fine-tuned on educational feedback. The sentiment score ( $v$ ) is:

$$C(v) = \frac{Np - Nn}{Np + Nn + Nneutral}$$

where  $Np$ ,  $Nn$ ,  $Nneutral$  are counts of positive, negative, and neutral comments. This yields a score between  $-1$  and  $+1$ . Prior work confirms these scores correlate with expert quality judgments[6].

### 3.5 Composite Scoring and Roadmap Assembly

Each candidate video gets a final composite score:

$$(v) = w_1 \cdot E(v) + w_2 \cdot T(v) + w_3 \cdot C(v) + w_4 \cdot R(v) + w_5 \cdot P(v) + w_6 \cdot N(v)$$

where  $E$  = engagement,  $T$  = teaching quality,  $C$  = comment sentiment,  $R$  = topic coverage,  $P$  = recency,  $N$  = presentation style. The top-scoring video per topic is selected. Since selections are independent per topic, the final roadmap typically pulls from 10+ different channels — each topic gets whoever explains it best, not just one instructor for everything.

### 3.6 Continuous Updates

The system runs a weekly automated scan. If a new video scores more than 5% higher than the current selection, it gets swapped in. If a creator deletes a video, the system detects the broken link and automatically substitutes the next-best option.

## 4. IMPLEMENTATION

### 4.1 Technology Stack

Table 3: Platform technology stack

Component	Technology
Frontend	React.js, Tailwind CSS
Backend API	Node.js, Express.js (RESTful)
Video Database	MongoDB
User Auth	Firebase Auth + Firestore
Content Source	YouTube Data API v3
NLP Engine	Python, HuggingFace BERT
Deployment	Docker containers on AWS EC2



#### 4.2 Dataset

The system processed 50 candidate videos per topic across all 8 domains — totaling 6,400+ unique videos — and ran sentiment analysis on approximately 1.28 million comments.

#### 4.3 Evaluation Method

Three complementary approaches were used:

1. **Automated Quality Validation** — AI-curated selections compared against benchmark roadmaps built by three independent domain experts per domain.
2. **Coverage Analysis** — Verified that every topic in every domain had at least one assigned tutorial with no gaps.
3. **User Study** — 120 engineering students (30 per academic year, years 1–4) used the platform for 4 weeks, each picking one domain and following the curated roadmap. They completed a 6-dimension Likert survey afterward.

### 5. RESULTS AND DISCUSSION

#### 5.1 Content Quality

AI-curated selections averaged 85.3% quality score versus 67.9% for manual YouTube search — a 25.6% improvement. The AI/ML domain scored highest (91.2%), benefiting from a large pool of established educators. Blockchain scored lowest (80.9%) due to fewer creators and rapidly changing content.

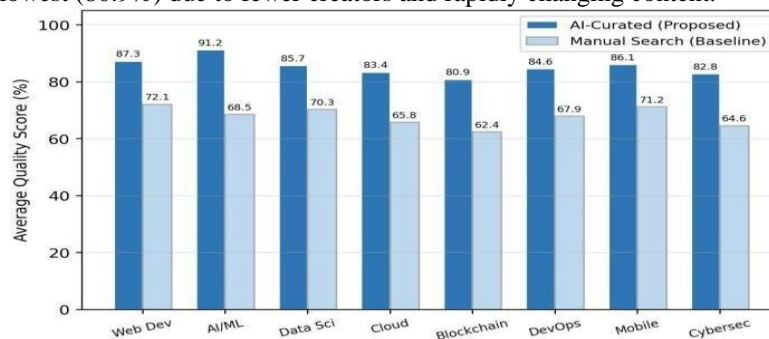


Fig. 3. Quality score comparison between AI-curated selections and manual YouTube search baseline across eight engineering domains

#### 5.2 Domain Coverage and Creator Diversity

All eight domains achieved 100% topic coverage — no roadmap had a missing topic. On average, each roadmap pulled from 10.4 distinct YouTube channels. AI/ML had the highest diversity at 14 unique creators; Blockchain had the lowest at 8.

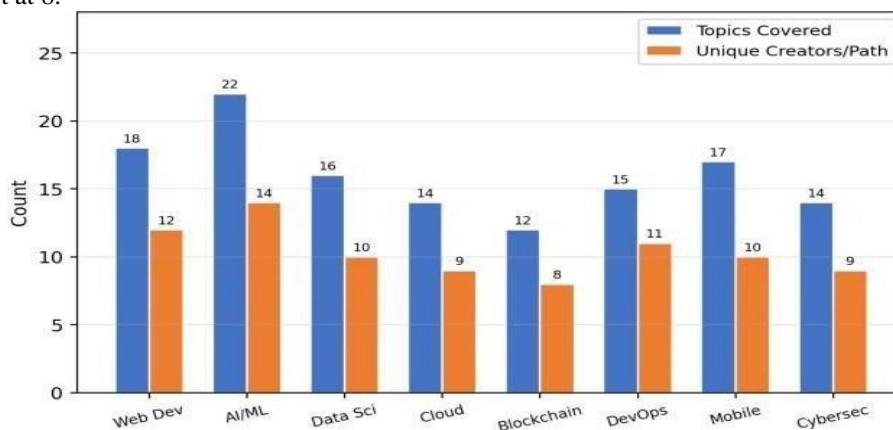


Fig. 4. Domain coverage: topics covered and unique creator diversity per assembled learning path.

#### 5.3 Sentiment Analysis Validation

Out of approximately 1.28 million classified comments: 68% positive, 22% neutral, 10% negative. The correlation between aggregated sentiment scores and expert quality ratings was  $r = 0.72$  — notably stronger than the  $r = 0.42$  reported in prior work[6], likely due to domain-specific fine-tuning of the BERT model.

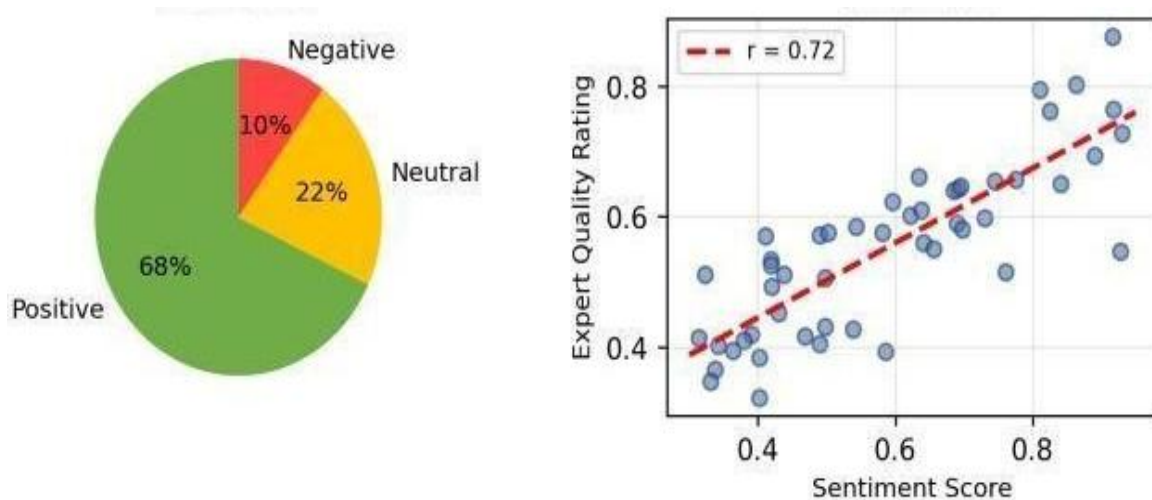


Fig. 5. (a) Overall comment sentiment distribution; (b) Sentiment vs Quality Correlation

#### 5.4 Time Efficiency

Average time to build a complete learning path by method:

- Manual YouTube browsing: 12.5 hours
- Online forum recommendations: 8.3 hours
- Paid platform catalog: 3.5 hours
- This system: 0.5 hours (96% reduction)

#### 5.5 User satisfaction

Survey results from 120 students (1–5 scale):

- Time Savings: 4.6 / 5.0
- Career Relevance: 4.4 / 5.0
- Content Quality: 4.3 / 5.0
- Overall: 4.2 / 5.0
- Path Coherence: 4.1 / 5.0
- UI/UX Design: 3.9 / 5.0 (below target; needs improvement)

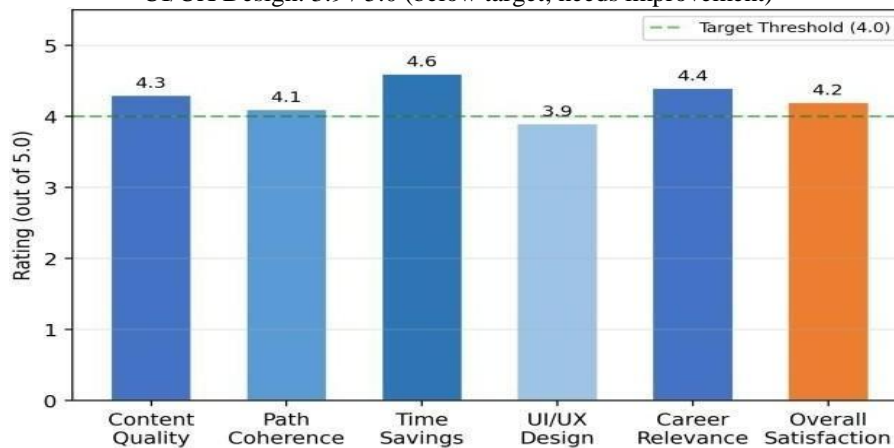


Fig. 6. User satisfaction survey results from 120 engineering students across six evaluation dimensions.

#### 5.6 Performance Summary

6. Table 4: Comprehensive performance metrics

Metric	Value	Notes
Avg. Quality Score	85.3%	vs. 67.9% manual baseline
Quality Improvement	+25.6%	over manual search
Topic Coverage	100%	all 8 domains
Avg. Creators per Path	10.4	instructional diversity
Sentiment–Quality Corr.	$r = 0.72$	strong NLP signal validity
Time Reduction	96%	12.5 h $\rightarrow$ 0.5 h
User Satisfaction	4.2/5.0	exceeds 4.0 target
Total Videos Analyzed	6,400+	



### 6.1 Limitations

- Depends entirely on YouTube's content library; niche topics may get weaker coverage
- English-only content currently
- No built-in coding environment
- No certificate on completion
- Requires active internet connection
- Creator deletions may cause brief gaps before the weekly scan catches them

## 7. CONCLUSION AND FUTURE WORK

This paper presented a platform that turns scattered YouTube tutorials into structured, quality-verified learning paths for engineering students. Using the YouTube Data API, a six-factor AI scoring engine, and BERT-based sentiment analysis, the system produced strong results in a 120-student pilot: 85.3% average content quality, 100% topic coverage across 8 domains, 96% less time spent on resource discovery, and 4.2/5.0 user satisfaction. The most significant technical result is the  $r = 0.72$  correlation between NLP-derived sentiment and expert quality ratings, which validates sentiment analysis as a scalable substitute for manual review. The broader takeaway is that intelligent curation of free content can match paid platforms on quality and structure — lowering the financial barrier to technical education.

### Planned future work:

Integrate paid platforms (Udemy, Coursera, edX) alongside YouTube  
Add regional language support — starting with Hindi and Marathi  
Embed a browser-based coding environment  
Add project assignments with automated code review  
Build peer learning communities and discussion forums  
Deploy an AI chatbot for on-demand doubt resolution  
Connect with industry mentors and job portals

## 8. REFERENCES

- [1] National Association of Colleges and Employers, "Job Outlook 2025 Survey," Bethlehem, PA, 2025.
- [2] "AI-powered personalized learning path generator," *Int. J. Res. Publ. Rev.*, vol. 6, no. 12, 2025.
- [3] A. Fortuna, "Artificial intelligence in personalized learning: A review," ScienceDirect, 2025.
- [4] "Personalized learning paths for student progression: An AI-driven approach," *Cureus*, vol. 17, no. 11, Nov. 2025.
- [5] Z. Zhao, "Optimization of personalized learning paths in educational AI platforms," *J. Intell. Educ. Adapt. Syst.*, vol. 3, no. 2, pp. 28–41, 2025.
- [6] M. D. Mayormente, "NLP-based sentiment analysis for evaluating student feedback," *IEEE Access*, 2025.
- [7] P. Covington, J. Adams, and E. Sargin, "Deep neural networks for YouTube recommendations," in *Proc. 10th ACM Conf. Recomm. Syst.*, 2016, pp. 191–198.
- [8] K. I. Vorobyeva, "Personalized learning through AI: Pedagogical approaches and critical insights," *Contemp. Educ. Technol.*, vol. 17, no. 1, 2025.
- [9] YouTube Official Blog, "YouTube by the numbers," 2025. [Online].
- [10] "Student sentiment analysis using NLP and ML," *Int. J. Creat. Res. Thoughts*, vol. 11, no. 4, Apr. 2023.
- [11] "Personalized learning pathway for engineering students," *Int. J. Innov. Res. Technol.*, Paper IJIRT174939, 2024.
- [12] "Student feedback analysis using NLP and sentiment analysis," in *Proc. ASEE Annu. Conf.*, 2024.