

# PAL: Personal Adaptive Learner

Megha Chakraborty, Darssan L. Eswaramoorthi, Madhur Thareja, Het Riteshkumar Shah, Finlay Palmer, Aryaman Bahl, Michelle A Ihetu, Amit Sheth

<sup>1</sup> Artificial Intelligence Institute, University of South Carolina

## Abstract

AI-driven education platforms have made some progress in personalisation, yet most remain constrained to static adaptation—predefined quizzes, uniform pacing, or generic feedback—limiting their ability to respond to learners’ evolving understanding. This shortfall highlights the need for systems that are both context-aware and adaptive in real time. We introduce PAL (Personal Adaptive Learner), an AI-powered platform that transforms lecture videos into interactive learning experiences. PAL continuously analyzes multimodal lecture content and dynamically engages learners through questions of varying difficulty, adjusting to their responses as the lesson unfolds. At the end of a session, PAL generates a personalized summary that reinforces key concepts while tailoring examples to the learner’s interests. By uniting multimodal content analysis with adaptive decision-making, PAL contributes a novel framework for responsive digital learning. Our work demonstrates how AI can move beyond static personalization toward real-time, individualized support, addressing a core challenge in AI-enabled education.

**Code** — <https://tinyurl.com/3c3vx2zn>

**Video** — <https://tinyurl.com/yc42pj55>

## 1 Introduction

Artificial intelligence is increasingly used in digital education to improve personalization and engagement. However, most lecture-based platforms remain static: learners watch videos passively, receive uniform quizzes, receive limited feedback. Such systems fail to adapt to individual learners in real time, often leaving students either under-challenged or overwhelmed, both of which reduce motivation and retention. Addressing this gap is key for scalable, equitable, and engaging learning environments.

Existing adaptive learning approaches, such as pretests or rigid rule-based adjustments, capture only coarse measures of ability (Claned 2025; Hyperspace 2025). They struggle to keep learners consistently within their optimal “learning zone” where the material is neither too easy nor too difficult.

We present the **Personal Adaptive Learner (PAL)**, a system that transforms lecture videos into interactive learning sessions. PAL integrates multimodal video analysis with an

adaptive question engine that periodically inserts multiple-choice questions during playback. A Hybrid Reinforcement Learning (RL) algorithm dynamically adjusts the difficulty of the question based on the accuracy, streaks, and timing of the responses, ensuring that learners remain engaged while progressively challenged. At the end of each lesson, PAL produces a personalized summary that reinforces key concepts and connects them to the specific interests of the learner.

The contributions of this demo are threefold:

- **PAL system:** A platform that unites lecture content with adaptive, real-time questioning.
- **Hybrid RL algorithm:** An adaptive decision making engine that balances stability with exploration for personalized learning.
- **Live demonstration:** A web-based demo showcasing how PAL sustains engagement and provides tailored post-learn feedback.

PAL highlights how AI can move beyond static personalization toward individualized real-time learning support.

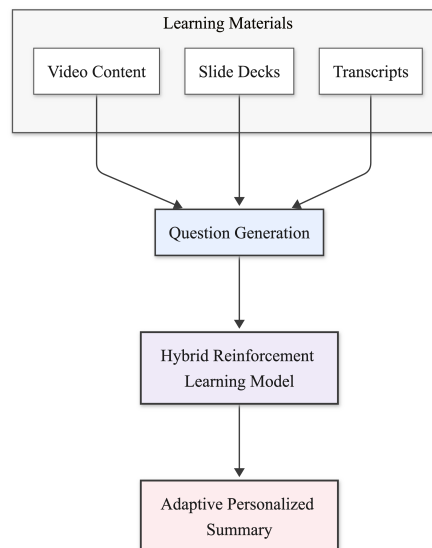


Figure 1: Architecture Diagram of PAL

## 2 Video-to-Question Dataset Generation

Given a lecture video  $V$  and transcript  $T = \{(t_i, u_i)\}_{i=1}^n$ , PAL uses a four-stage pipeline to generate timestamped, difficulty-rated questions (Koppolu et al. 2025; Singh et al. 2025).

**Transcript Analyzer.** Scans  $T$  for candidate points via linguistic cues (e.g., “is defined as”) or every  $N$  sentences (Garrison et al. 2006).

**Context Validator.** At each point, extracts a frame, applies OCR for slide text, and uses LLaVA-mini for visual description. These combine with transcript text into a context bundle (Abinaya et al. 2024).

**Question Generator.** Produces Q-A pairs from context. Difficulty is heuristic-based: factual (Easy), conceptual (Medium), applied (Hard). Supports MCQs and lightweight LLM generation (Du and Cardie 2018; Liu et al. 2025a).

**Difficulty Rater.** Labels difficulty via rules (“what is”  $\rightarrow$  Easy, “why/how”  $\rightarrow$  Medium, “apply/predict”  $\rightarrow$  Hard), with LLM fallback.

**Master Pipeline.** The integrated `tri_plus_one_pipeline` outputs structured JSON ( $q, a, d, t, c$ ) for adaptive learning, quizzes, and study aids.

This modular design is lightweight, extensible, and ensures contextual accuracy by combining speech, vision, and text.

## 3 Adaptive Difficulty via Hybrid RL

**Problem.** At interaction  $t$ , the learner state  $\mathbf{x}_t \in R^6$  encodes (skill, recent accuracy, normalized response time, streak momentum, learning velocity, confidence) (Zheng 2024; Liu et al. 2025b). We choose difficulty  $a_t \in \{\text{Easy, Medium, Hard}\}$  to maximize long-term reward:

$$\pi^* = \arg \max_{\pi} E \left[ \sum_{k=0}^{\infty} \gamma^k R_{t+k+1} \mid a_t \sim \pi(\mathbf{x}_t) \right], \gamma \in [0, 1).$$

**Reward (composite).** We use a compact additive shaping:

$$r_t = r_{\text{acc}} + r_{\text{time}} + r_{\text{prog}} + r_{\text{mom}},$$

where  $r_{\text{acc}}$  is 1 for a correct response and  $-0.5$  for an incorrect one, while  $r_{\text{time}} \in [0, 0.3]$ ,  $r_{\text{prog}} \in [0, 0.2]$ , and  $r_{\text{mom}} \in [0, 0.1]$  encourage timely answers, sensible difficulty progression, and streak momentum respectively.

**Statistical prior (IRT with stability buffers).** We impose a smooth prior over difficulties using a two-parameter logistic (2PL) model:

$$p_{\text{stat}}(d \mid \mathbf{x}_t) \propto \sigma(a_d(\theta_t - b_d)),$$

where  $a_d$  and  $b_d$  are the item discrimination and difficulty parameters, and  $\theta_t$  is the learner skill estimate. To stabilize adaptation, we apply asymmetric thresholds (promote if recent accuracy  $\geq 0.75$ , demote only if  $\leq 0.35$ ) together with short cooldown and hold periods, which prevent rapid back-and-forth switching between levels (Matteucci, Mignani, and Veldkamp 2012; Lee, Peña-Mora, and Park 2006).

**RL head (bandit + Q-learning).** With  $\varepsilon$ -greedy selection (Vivanti 2021; Cannelli et al. 2023) and three actions, we update:

$$Q_{t+1}(a_t) \leftarrow Q_t(a_t) + \alpha \left[ R_t + \gamma \max_{a'} Q_t(a') - Q_t(a_t) \right].$$

**Hybrid stochastic blending.** We form a confidence-scaled mixture:

$$\pi_t(d \mid \mathbf{x}_t) = (1 - w_t) p_{\text{stat}}(d \mid \mathbf{x}_t) + w_t p_{\text{RL}}(d \mid \mathbf{x}_t), \quad (1)$$

$$w_t = \min(w_{\text{max}}, w_0 + \kappa c_t \text{progress}_t). \quad (2)$$

so the probability of choosing the RL suggestion rises with evidence ( $c_t$ ) and session progress, capped by  $w_{\text{max}}$  (e.g., 0.8). Both heads are updated every step; the statistical prior stabilizes early phases while RL learns a personalized preference.

## 4 Personalized Post-Lesson Summary

At the end of a lecture, PAL generates a personalized summary based on the full learner profile (Essa, Celik, and Human-Hendricks 2023). The summary highlights **Territory Mastered** and **Discovery Zone**.

**Stage 1 - Smart extraction with semantic search.** First, the backend preprocesses the whole lecture: it breaks the transcript into sentences and turns each sentence into a numeric vector with a SentenceTransformer (Ruas and Grosky 2017). This creates a “semantic map” of the lecture. When a student asks about a concept (for example, “Correlation vs Causation”), the system converts that concept into a vector too and does a fast similarity search on the semantic map. That way it finds the most relevant sentences even if they use different words or are spread out. The process ignores conversational noise and collects a compact, context-rich block of raw text to use in the next step.

**Stage 2 - Context-aware synthesis with Llama 3.2.** The collected text (which may be fragmented) is sent to our Llama 3.2 1B Instruct model. Instead of a simple “summarize” command, the model receives a carefully designed prompt that makes it act like an expert teaching assistant. The prompt tells the model to use the lecture text as the main source but also allows it to fill gaps and clarify awkward phrasing using its own knowledge. The prompt can also adapt based on the student’s needs (e.g., request more detail for topics the student struggled with). The result is a clear, well-structured explanation that stays faithful to the lecture while improving readability and completeness.

## 5 Discussion

PAL (Personal Adaptive Learner) adapts instruction by generating timestamped, difficulty-rated questions from lecture videos through a multimodal, agentic pipeline. Unlike static AI education tools, PAL integrates transcript analysis, visual context validation, question generation, and difficulty rating to align content with learner performance in real time. The evaluation of PAL will be made available online through an interactive interface, allowing readers to explore example transcripts, generated questions, and difficulty ratings. Future work will expand feedback loops, integrate collaborative learning, and evaluate classroom impact.

## References

- Abinaya, G.; Rao, G. D.; Gopal, P. S. R.; Kaushik, M.; and Vignesh, B. S. V. 2024. Automated Document Processing: Combining OCR and Generative AI for Efficient Text Extraction and Summarization. In *2024 International Conference on Innovative Computing, Intelligent Communication and Smart Electrical Systems (ICSES)*, 1–5. IEEE.
- Cannelli, L.; Nuti, G.; Sala, M.; and Szehr, O. 2023. Hedging using reinforcement learning: Contextual k-armed bandit versus Q-learning. *The Journal of Finance and Data Science*, 9: 100101.
- Claned. 2025. The Role of AI in Personalized Learning. Accessed: 2025-09-14.
- Du, X.; and Cardie, C. 2018. Harvesting paragraph-level question-answer pairs from wikipedia. *arXiv preprint arXiv:1805.05942*.
- Essa, S. G.; Celik, T.; and Human-Hendricks, N. E. 2023. Personalized adaptive learning technologies based on machine learning techniques to identify learning styles: A systematic literature review. *IEEE Access*, 11: 48392–48409.
- Garrison, D. R.; Cleveland-Innes, M.; Koole, M.; and Kappelman, J. 2006. Revisiting methodological issues in transcript analysis: Negotiated coding and reliability. *The internet and higher education*, 9(1): 1–8.
- Hyperspace. 2025. The Future of Personalized Education. Accessed: 2025-09-14.
- Koppolu, H. K. R.; Gadi, A. L.; Motamary, S.; Dodda, A.; and Suura, S. R. 2025. Dynamic Orchestration of Data Pipelines via Agentic AI: Adaptive Resource Allocation and Workflow Optimization in Cloud-Native Analytics Platforms. *Metallurgical and Materials Engineering*, 31(4): 625–637.
- Lee, S.; Peña-Mora, F.; and Park, M. 2006. Reliability and stability buffering approach: Focusing on the issues of errors and changes in concurrent design and construction projects. *Journal of construction engineering and management*, 132(5): 452–464.
- Liu, X.; Li, R.; Huang, M.; Liu, Z.; Song, Y.; Guo, Q.; He, S.; Wang, Q.; Li, L.; Liu, Q.; et al. 2025a. Thus spake long-context large language model. *arXiv preprint arXiv:2502.17129*.
- Liu, Z.; Gong, C.; Fu, X.; Liu, Y.; Chen, R.; Hu, S.; Zhang, S.; Liu, R.; Zhang, Q.; and Tu, D. 2025b. Ghpo: Adaptive guidance for stable and efficient llm reinforcement learning. *arXiv preprint arXiv:2507.10628*.
- Matteucci, M.; Mignani, S.; and Veldkamp, B. P. 2012. Prior distributions for item parameters in IRT models. *Communications in Statistics-Theory and Methods*, 41(16-17): 2944–2958.
- Ruas, T.; and Grosky, W. 2017. Keyword extraction through contextual semantic analysis of documents. In *Proceedings of the 9th International Conference on Management of Digital EcoSystems*, 150–156.
- Singh, A.; Ehtesham, A.; Kumar, S.; and Khoei, T. T. 2025. Agentic retrieval-augmented generation: A survey on agentic rag. *arXiv preprint arXiv:2501.09136*.
- Vivanti, R. 2021. Can Q-learning solve Multi Armed Bandits? *arXiv preprint arXiv:2110.10934*.
- Zheng, T. 2024. Dynamic difficulty adjustment using deep reinforcement learning: A review. *Applied and Computational Engineering*, 71: 157–162.