

# IntelliCode: A Multi-Agent LLM Tutoring System with Centralized Learner Modeling

Jones David<sup>1</sup>, Shreya Ghosh<sup>2</sup>

<sup>1</sup>School of Computer Science and Engineering, VIT-AP University, India

<sup>2</sup>School of Electrical and Computer Sciences, IIT Bhubaneswar, India

Correspondence: jones.22bce8135@vitapstudent.ac.in, shreya@iitbbs.ac.in

## Abstract

LLM-based tutors are typically single-turn assistants that lack persistent representations of learner knowledge, making it difficult to provide principled, transparent, and long-term pedagogical support. We introduce **IntelliCode**, a multi-agent LLM tutoring system built around a centralized, versioned learner state that integrates mastery estimates, misconceptions, review schedules, and engagement signals. A StateGraph Orchestrator coordinates six specialized agents: skill assessment, learner profiling, graduated hinting, curriculum selection, spaced repetition, and engagement monitoring, each operating as a pure transformation over the shared state under a single-writer policy. This architecture enables auditable mastery updates, proficiency-aware hints, dependency-aware curriculum adaptation, and safety-aligned prompting. Our demo showcases an end-to-end tutoring workflow: a learner attempts a DSA problem, receives a conceptual hint when stuck, submits a corrected solution, and immediately sees mastery updates and a personalized review interval. We report validation results with simulated learners, showing stable state updates, improved task success with graduated hints, and diverse curriculum coverage.<sup>12</sup>

## 1 Introduction

Large Language Models (LLMs) have rapidly expanded the possibilities of automated tutoring, yet most existing systems remain fundamentally reactive: each query is treated in isolation, with little continuity or awareness of a learner’s evolving knowledge (Wu et al., 2025). Human tutors, in contrast, maintain rich, persistent models of student understanding that support targeted feedback, curriculum scaffolding, and long-term learning trajectories (VanLEHN, 2011). This discrepancy limits

<sup>1</sup>Live System: <https://intellicode.redomic.in>,

<sup>2</sup>Video Demo: <https://youtu.be/o08bZfele0U>

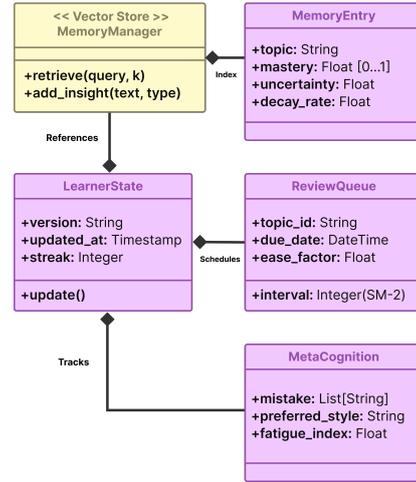


Figure 1: Unified Learner State Schema.

the pedagogical reliability of current LLM tutors, which often provide inconsistent hints, overlook dependencies between concepts, and fail to account for systematic misconceptions (Chu et al., 2025; Mukherjee and Ghosh, 2025). Recent frameworks such as GenMentor (Wang et al., 2025) and SocraticLM (Liu et al., 2024) demonstrate the promise of multi-agent orchestration and structured dialogue in stabilizing tutoring behavior. However, these systems typically rely on ephemeral or implicit memory, lacking an explicit, auditable learner model shared across agents. At the same time, decades of work on learner modeling from Bayesian Knowledge Tracing (BKT) (Corbett and Anderson, 1995) and Deep Knowledge Tracing (DKT) (Piech et al., 2015) to memory networks (Zhang et al., 2017) and PFA (Pavlik Jr et al., 2009), underscore the importance of accurate mastery estimation for personalization. Yet, few LLM-based tutors integrate such formal models with generative reasoning in a persistent, updateable state.

**Our goal is to bridge this gap.** We introduce **IntelliCode**, a multi-agent LLM tutoring system built around a centralized, versioned learner state that

serves as the single source of truth for all pedagogical decisions. Unlike prior systems, IntelliCode enforces a single-writer policy through a *StateGraph Orchestrator*, ensuring that every mastery update, hint intervention, or curriculum choice results from a coherent, formally validated transformation of the learner model. This design mitigates drift, prevents conflicting agent outputs, and enables transparent, multi-turn personalization.

IntelliCode integrates well-established instructional principles: mastery estimation, graduated hinting, dependency-aware curriculum planning, and spaced repetition; with modern LLM capabilities. The learner state encodes mastery vectors with uncertainty, misconceptions, review schedules, and behavioral signals. Six specialized agents (Skill Assessment, Learner Profiler, Pedagogical Feedback, Content Curator, Progress Synthesizer, and Engagement Orchestrator) each operate as pure functions over this shared state. For example, if a learner repeatedly omits a base case in recursion problems, the Profiler records a misconception, the Feedback agent adapts its hint level, and the Curator adjusts upcoming tasks accordingly.

**Contributions.** This work makes the following contributions: (1) a centralized, versioned learner state and single-writer orchestration mechanism enabling consistent, auditable multi-turn tutoring; (2) a suite of pedagogical agents that implement mastery estimation, graduated hinting, curriculum planning, and spaced repetition as pure transformations over the shared state; and (3) a fully functional, end-to-end LLM tutoring system demonstrating stable interaction, interpretable decision-making, and robust content coverage through simulated learner studies.

Figure 1 highlights the unified learner state that drives our adaptation policies. By grounding agent behavior in this structured representation while leveraging LLM reasoning for high-variance tasks such as hinting and code analysis, IntelliCode offers a transparent and pedagogically consistent alternative to memory-less conversational tutoring.

## 2 System Architecture

We frame adaptive personalized education as a Partially Observable Markov Decision Process (POMDP). At each timestep, the learner state  $\mathbf{S}_t$  maintains mastery vectors, SM-2–based review schedules, engagement metrics, and metacognitive memory. Observations  $\mathcal{O}_t$  reflect noisy behavioral

signals such as submissions, errors, and hint requests, while actions  $\mathcal{A}_t$  correspond to pedagogically meaningful interventions including content recommendation, graduated hinting, and schedule adjustments. The reward function  $R_t$  trades off mastery gains with penalties for excessive hint usage and inefficient solve times (details in Appendix B). This formulation enables principled adaptation while supporting the modular multi-agent design showcased in the demo.

### 2.1 Orchestrator Overview

At the core of IntelliCode is the **StateGraph Orchestrator**, the only component permitted to write to the persistent learner record. It maintains a synchronized in-memory copy of the learner state and coordinates all interactions among the six pedagogical agents. When an event occurs, the orchestrator routes it to the appropriate agents, aggregates their outputs, validates the proposed state changes, and then commits them as an atomic update. When multiple agents produce overlapping deltas, the orchestrator applies a fixed priority ordering—Skill Assessment, then Learner Profiler, then remaining agents—and validates each proposed update against the learner state schema before committing. Malformed or conflicting deltas are rejected and logged for audit. Similar to the coordination strategy in GenMentor (Wang et al., 2025), this mechanism prevents conflicting writes, enforces safety and schema constraints, and ensures that the system behaves predictably over multi-turn, long-term tutoring sessions. The orchestrator thus provides the reliability and auditability necessary for principled learner modeling.

### 2.2 Trigger Types and Routing

The orchestrator reacts to pedagogically meaningful events and dispatches them to the relevant agents. These triggers operationalize the full workflow demonstrated in the system:

- **on\_submission:** A code or answer submission triggers Skill Assessment, followed by the Learner Profiler and Pedagogical Feedback.
- **on\_hint\_request:** A learner request for help triggers the Pedagogical Feedback agent, informed by current proficiency and hint history.
- **on\_session\_check:** A daily check-in triggers

Table 1: Roles and responsibilities of the six pedagogical agents.

Agent	Responsibility
Pedagogical Feedback	Provides proficiency-aware, five-level graduated hinting without solution disclosure.
Content Curator	Selects personalized problems based on mastery, dependencies, and the 40/50/10 curriculum policy.
Engagement Orchestrator	Monitors motivation, pacing, and disengagement signals to issue supportive nudges.
Skill Assessment	Performs hybrid evaluation using test-case execution and semantic code review.
Learner Profiler	Estimates mastery deltas, identifies misconceptions, and infers behavioral trends.
Progress Synthesizer	Schedules reviews using an enhanced SM-2 mechanism with context-aware adjustments.

the Content Curator and Engagement Orchestrator.

- **on\_daily\_generation:** The system generates the day’s personalized problem set via the Content Curator.
- **on\_review\_due:** When an SM-2 review is due, the Progress Synthesizer and Content Curator are invoked.

These triggers allow IntelliCode to respond adaptively to the learner’s evolving behavior, maintaining pedagogical continuity across sessions.

### 2.3 Overview of System Agents

Each agent operates as a pure transformation over the shared learner state, producing structured outputs that the orchestrator validates and integrates. Together, the six agents support assessment, personalization, pacing, hinting, and review scheduling. Table 1 summarizes their responsibilities, and Figure 2 shows how the orchestrator mediates their communication. This design ensures that all instructional decisions are grounded in a consistent, auditable learner state and remain traceable throughout the tutoring trajectory.

The overall data flow is illustrated in Figure 2. The StateGraph Orchestrator mediates all communication between agents and the persistent learner state, ensuring that instructional decisions remain traceable and consistent across sessions.

## 3 Learner State and Agent Adaptation

IntelliCode maintains a centralized, versioned learner state that governs all pedagogical decisions. The state is initialized from historical activity and updated using a Bayesian Knowledge Tracing (BKT)–inspired mechanism that incorporates difficulty, recency, hint usage, and solve-time effects. Spaced repetition is managed by an enhanced SM-2 scheduler that computes personalized review intervals based on recall quality and

interaction history. This shared representation enables consistent, long-term adaptation rather than isolated single-turn responses.

To illustrate the update process, consider a learner who solves a recursion problem correctly but requests several hints and exceeds the expected solve time. The BKT-inspired update assigns a modest mastery gain, attenuated by the heavier reliance on hints, while the Progress Synthesizer schedules an earlier review to reinforce retention. The Content Curator then interprets recursion as lying in the learner’s “growth” region and adjusts future problem selection accordingly.

### 3.1 Agent Behaviors and Pedagogical Logic

All six agents operate as pure transformations over the learner state, producing structured outputs validated by the orchestrator before being committed as atomic updates. While the architecture supports fully generative agents, deterministic logic is used for the Learner Profiler and Content Curator for reproducibility, whereas higher-variance components (e.g., hinting, code analysis) leverage LLMs.

**Learner Profiler** The Profiler acts as the diagnostic backbone of the system, identifying mastery deltas, misconceptions, and behavioral trends such as fatigue or decreasing velocity. It consumes correctness, topic tags, error patterns, time-on-task, hint usage, and the current mastery map. For example, if a learner repeatedly omits base cases in recursion problems, the Profiler records a misconception related to termination conditions, which later guides both hinting and content selection. For new learners, the Profiler initializes mastery with uninformative Beta priors ( $\alpha = 1, \beta = 1$ ) and conservative estimates, allowing the system to warm up within 2–3 submissions as observed outcomes rapidly concentrate the posterior.

**Skill Assessment** This agent performs hybrid evaluation by executing test cases and conducting semantic code review. When tests fail, errors alone

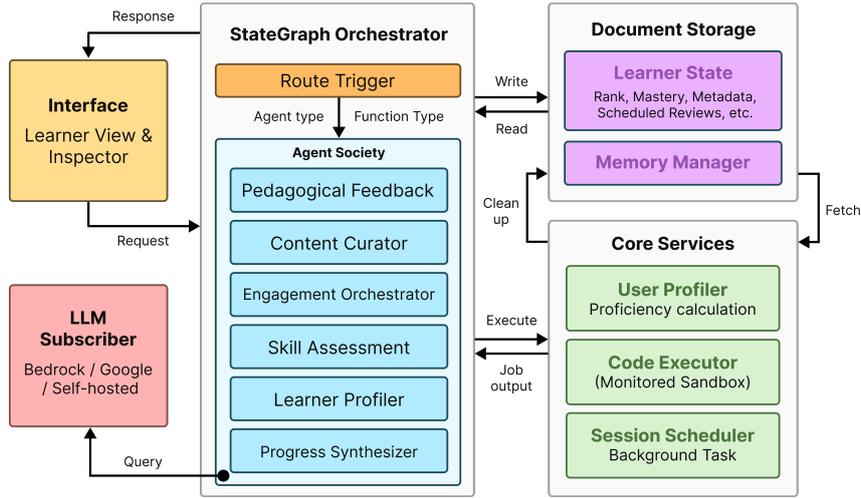


Figure 2: System Architecture: The StateGraph Orchestrator manages the flow between six specialized agents and the persistent learner state. Arrows indicate data flow and trigger events.

are surfaced; when they pass, the agent provides improvement suggestions across categories such as time complexity, space usage, readability, and edge-case coverage. For instance, after a successful merge sort implementation, the agent may recommend reducing auxiliary memory to improve space efficiency.

**Pedagogical Feedback** Informed by Socratic Playground (Zhang and Others, 2024), the Pedagogical Feedback agent employs a five-level graduated hinting protocol:

- **Metacognitive:** prompt the learner to reflect (“What did you try, and what happened?”).
- **Conceptual:** surface the key idea (“This problem relies on identifying a recurrence relation.”).
- **Strategic:** suggest an approach (“Consider breaking the input and solving the halves recursively.”).
- **Structural:** highlight missing logic (“Your solution lacks a base case for empty input.”).
- **Targeted:** point to a region of interest (“Inspect the condition near line 14; termination may not be guaranteed.”).

## System Demonstration

The specificity of hints scales with the learner’s proficiency estimate  $\hat{p}$ . Beginners receive simple analogies and single-step cues, intermediate learners

receive pattern-oriented guidance, and advanced learners receive concise nudges with edge-case emphasis. For the same recursion bug, a beginner might be told to “think of recursion like climbing down a ladder,” while an advanced learner might be prompted to “check whether the termination condition is reachable.”

**Content Curator** The Curator operationalizes the learner state into task selection using a dependency-aware 40/50/10 policy:

$$\text{selection} = 0.4 \times \text{due\_reviews} + 0.5 \times \text{growth\_zone} + 0.1 \times \text{challenge}. \quad (1)$$

Growth-zone items correspond to mastery levels between 0.3 and 0.7, while challenge items target skills below 0.3. The Curator enforces prerequisite dependencies, avoids repetition within a  $k$ -day window, and ensures topic diversity. For example, a learner showing intermediate recursion mastery but weak dynamic programming mastery may receive (i) a recursion review task, (ii) a medium-difficulty DP subproblem, and (iii) a lightweight DP challenge problem.

**Progress Synthesizer** The Progress Synthesizer governs spaced repetition using SM-2 (Wozniak, 1990) and forgetting-curve theory (Ebbinghaus, 1913), augmented with contextual features (Settles and Meeder, 2016; Reddy et al., 2016). Review intervals shrink when hints are heavily used, expand when solutions are fast and confident, and tighten when predicted recall drops near the due date. If the recall probability for a graph traversal concept falls below threshold, the agent prepones

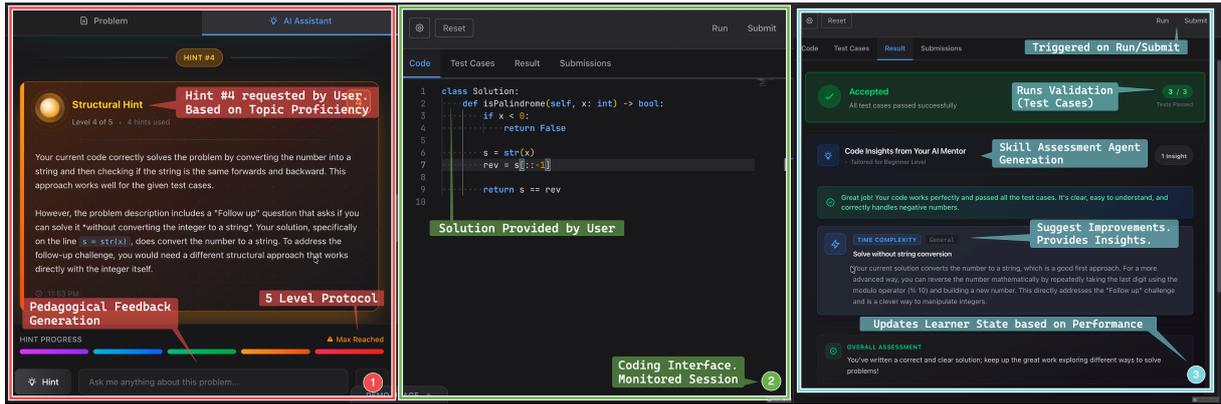


Figure 3: Learner interface featuring the graduated hinting mechanism (Left), the Code Editor (Middle), and proficiency-aware code analysis (Right).

the review—even if the learner has not recently interacted.

**Engagement Orchestrator** Finally, the Engagement Orchestrator monitors motivational signals. It issues supportive prompts after broken streaks, encourages re-engagement after periods of inactivity, and suggests simpler variants when failure streaks accumulate. All interventions are rate-limited and phrased non-judgmentally. For instance, after multiple failed attempts on tree problems, the system may suggest: “Would you like to revisit the easier ‘binary tree basics’ exercise before trying again?”

Collectively, these agents form IntelliCode’s adaptation engine, grounding every hint, problem selection, and review decision in a unified, auditable learner state.

The IntelliCode platform is implemented using FastAPI for the backend, React for the frontend interface, and LangGraph for multi-agent orchestration. A persistent, graph-structured learner model is maintained in ArangoDB, enabling the system to track mastery, misconceptions, and review schedules across sessions. All LLM-powered agents (Pedagogical Feedback, Skill Assessment, Engagement Orchestrator) use Google Gemini 2.5 Flash via API without fine-tuning. Median end-to-end latency for a single agent call is approximately 2–3 seconds, dominated by LLM inference. The backend serves concurrent users via asynchronous request handling in FastAPI; API throughput is rate-limited to 5 requests per second per user to comply with provider quotas.

The demo showcases the full adaptive tutoring loop. A learner begins at a curriculum roadmap and is assigned a data structures and algorithms (DSA) problem selected by the Content Curator using the

40/50/10 policy. If the learner struggles, the Pedagogical Feedback agent produces a proficiency-aligned hint—for example, a Level 2 conceptual cue such as “This problem requires identifying the recurrence pattern.” After incorporating the hint, the learner submits a correct solution. This submission triggers a sequence of coordinated agent behaviors. The Skill Assessment agent validates correctness and offers semantic feedback; the Learner Profiler updates the learner’s mastery estimate for recursion; and the Progress Synthesizer schedules a spaced-repetition review two days later, reflecting the hint usage and solve-time profile. The system interface allows the learner to view mastery trajectories, upcoming reviews, and past interactions, making the adaptation process transparent. This end-to-end interaction exemplifies how IntelliCode integrates real-time assessment, graduated hinting, curriculum adaptation, and spaced repetition into a coherent, state-driven teaching cycle.

## 4 Evaluation Protocols

We evaluate IntelliCode along offline, online, and fairness dimensions to demonstrate the reliability of its learner modeling, content adaptation, and multi-agent interactions.

**Offline Metrics** Our offline analysis focuses on validating the fidelity of the learner model. We measure **Mastery Calibration** using Expected Calibration Error (ECE), which bins predicted mastery values and compares them against observed success rates to assess whether the model’s confidence is well-aligned with actual outcomes. We omit per-problem discriminative metrics such as Brier score and AUROC, as individual problem success in our simulation is influenced by stochastic code gener-

ation and error injection; aggregate mastery is not designed to predict single-problem outcomes, but rather to track long-term learning trajectories. The **Content Policy** is assessed by tracking topic coverage and diversity to ensure balanced curriculum exposure. To isolate the contribution of the multi-agent architecture, we compare against a **stateless baseline** in which learners receive random problem selection, a single attempt per problem, and no agent support (no hints, orchestrator, or code analysis).

**Online Metrics** We also track performance in live interactions. **Learning Gains** are estimated through pre/post mastery changes on held-out assessments. **Engagement** is monitored through streak retention, inactivity gaps, and voluntary practice rates. System **Efficiency** is assessed via median end-to-end latency, with a target of under 500 ms. **Safety** is evaluated through hint acceptance rates (aiming for  $\geq 70\%$ ) and verification that the system never discloses full solutions.

**Fairness Analysis** To ensure equitable support across learner profiles, we compare learning gains, hint levels, and pacing behaviors across proficiency deciles, targeting an interquartile range within 15% of the median. Component-level ablation studies isolating the individual contributions of the Content Curator, Pedagogical Feedback agent, and SM-2 scheduler are planned for future work.

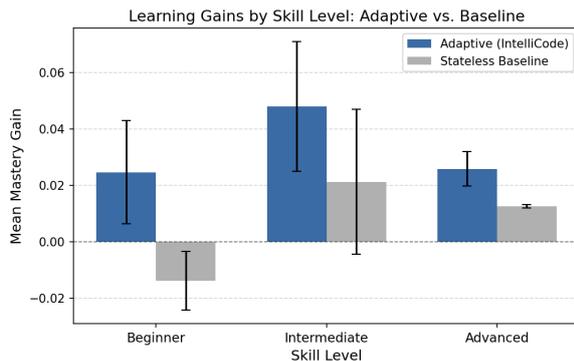


Figure 4: Adaptive IntelliCode consistently outperforms the stateless baseline across all skill levels. Baseline beginners show negative mastery growth in the absence of scaffolding.

### Validation with Simulated Learners

To assess the responsiveness and stability of the multi-agent architecture, we conducted controlled simulations using agent-based learner personas (Wu et al., 2025), drawing on methodolo-

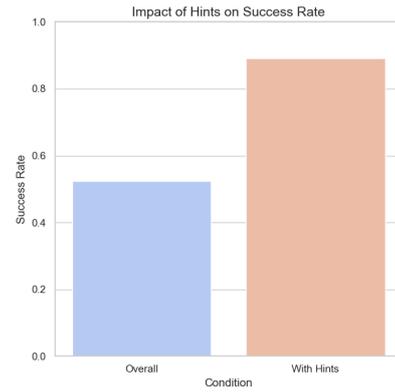


Figure 5: Success rates with and without hint utilization.

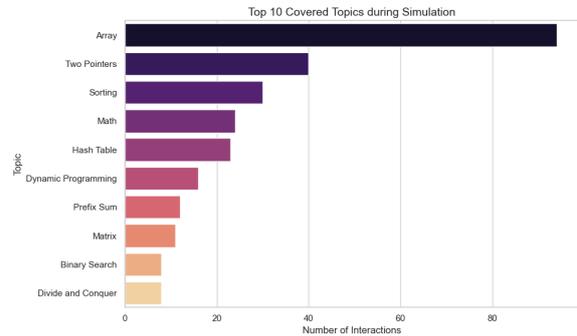


Figure 6: Top 10 topics covered during the simulation.

gies from generative agent societies (Park et al., 2023). While these simulations cannot substitute for human studies in evaluating educational efficacy, they provide a controlled comparison of the system’s adaptation capabilities across diverse cognitive profiles. Ten synthetic learners (4 beginner, 4 intermediate, 2 advanced) each completed a 14-day learning trajectory under two conditions: (1) the full adaptive IntelliCode system with mastery-based problem selection, graduated hinting, and up to three attempts per problem; and (2) a stateless baseline with random problem selection, a single attempt, and no agent interactions.

### 4.1 Comparative Results

Table 2 summarizes the results. The adaptive system achieved a mean mastery gain of 3.42%, compared to 0.54% for the stateless baseline—a  $6.3\times$  improvement (Figure 4). Learning velocity was similarly differentiated: adaptive learners gained 0.0029 mastery points per active day, versus 0.0004 for the baseline. Notably, baseline beginners experienced *negative* mastery growth ( $-1.39\%$ ), suggesting that without adaptive scaffolding, weaker learners regress due to poorly matched problem

Table 2: Adaptive vs. stateless baseline across 10 simulated learner trajectories over a 14-day horizon.

Metric	Adaptive	Baseline
Mastery Gain (Mean)	0.0342	0.0054
ECE ( $\downarrow$ )	0.043	0.286
Success Rate	89.1%	52.4%
Learning Velocity (gain/day)	0.0029	0.0004
Topics Covered	38	36

difficulty and the absence of corrective feedback. In contrast, adaptive beginners achieved positive gains (+2.46%).

Mastery calibration, measured by ECE, showed a large separation: 0.043 (adaptive) versus 0.286 (baseline). This indicates that the adaptive system’s mastery estimates are well-aligned with observed success rates, whereas the baseline’s predictions are poorly calibrated. The graduated hinting mechanism contributed meaningfully: tasks in which simulated learners requested hints exhibited a success rate of 89.1%, compared with 52.4% without hints (Figure 5), confirming that the Pedagogical Feedback agent provides effective conceptual guidance without disclosing solutions.

## 4.2 Content Coverage

The Content Curator maintained strong diversity across topics (Figure 6), demonstrating the ability of the 40/50/10 policy to avoid topic starvation while respecting prerequisite relationships. Even in extended sessions, the system preserved balanced coverage across skill areas, validating the orchestrator’s ability to coordinate long-term learning arcs and manage curriculum progression.

## 5 Conclusion

In this paper, we presented IntelliCode, a principled multi-agent LLM framework for adaptive education built around a persistent, auditable learner state. By integrating formal mastery update rules, proficiency-aware graduated hinting, dependency- and fairness-aware curriculum adaptation, and safety-aligned prompting, the system offers transparent and consistent multi-turn tutoring capabilities. Controlled comparison against a stateless baseline demonstrates that the adaptive multi-agent architecture produces substantially higher learning gains and better-calibrated mastery estimates, while maintaining robust content coverage and effective hinting interventions. Future work will extend IntelliCode through larger-scale human studies, offline policy evaluation, and federated learn-

ing integrations to strengthen privacy guarantees. We envision IntelliCode as a foundation for next-generation educational systems that blend modern LLM reasoning with established principles from learner modeling, instructional design, and cognitive science.

## Limitations

While IntelliCode demonstrates promising architectural and pedagogical capabilities, there are a few limitations. First, the mastery estimates rely on BKT/DKT-inspired proxies that require sufficient interaction scale and careful calibration; cold-start learners, in particular, necessitate conservative priors and may experience reduced personalization in early sessions. To mitigate this, the system initializes mastery with uninformative priors and applies conservative problem selection during the first few sessions, progressively refining estimates as interaction data accumulates. Second, LLM-driven components introduce variability due to model drift, occasional refusals, and cost constraints. Our guardrails and validation schemas mitigate these issues but cannot fully eliminate them. Finally, rigorous fairness evaluation requires diverse and representative datasets, and remains vulnerable to selection bias, behavior-signal noise, LLM drift, data leakage, and survivorship bias. These considerations underscore the need for large-scale, longitudinal human studies in future work.

## Ethical Considerations

The deployment of LLM-based agents in educational settings necessitates careful attention to accuracy, dependency, and privacy. While IntelliCode integrates verifiers like the *Skill Assessment* agent to validate code logic, generative components such as the *Pedagogical Feedback* agent remain susceptible to hallucinations or plausible but incorrect explanations. Consequently, the system is designed to function as a supplemental tutor rather than a replacement for formal instruction, and we recommend its use under the guidance of human educators who can monitor for potential deviations.

To mitigate the risk of learner over-dependence on AI assistance, we implemented strict graduated hinting protocols. However, we acknowledge that prolonged reliance on automated scaffolding may impact unassisted problem-solving capabilities. Our design prioritizes metacognitive prompting over direct solution disclosure to foster genuine

skill acquisition.

Regarding data privacy, all learner interactions are processed with strict minimization principles. Personally Identifiable Information (PII) is redacted prior to agent ingestion, and the *Learner Profiler* operates on anonymized mastery integers rather than raw user profiles. Finally, while our curriculum policy incorporates fairness constraints to ensure equitable topic coverage, the underlying datasets used for cold-start calibration may inherently reflect historical biases, requiring ongoing monitoring of learning outcomes across diverse demographic groups.

## Acknowledgments

SPARC (Scheme for Promotion of Academic and Research Collaboration) Phase-III (Project ID: 3385) and TIH IIT Tirupati (IITNiF/TPD/2024-25/P16) partially supported this research work.

## References

- Zhendong Chu, Shen Wang, Jian Xie, Tinghui Zhu, Yibo Yan, Jingheng Ye, Aoxiao Zhong, Xuming Hu, Jing Liang, Philip S. Yu, and Qingsong Wen. 2025. [LLM agents for education: Advances and applications](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 13782–13810, Suzhou, China. Association for Computational Linguistics.
- Albert T Corbett and John R Anderson. 1995. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction*, 4(4):253–278.
- Hermann Ebbinghaus. 1913. *Memory: A contribution to experimental psychology*. Teachers College, Columbia University.
- Jiayu Liu, Zhenya Huang, Tong Xiao, Jing Sha, Jinze Wu, Qi Liu, Shijin Wang, and Enhong Chen. 2024. [SocraticLM: Exploring socratic personalized teaching with large language models](#). In *Proceedings of the Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Arka Mukherjee and Shreya Ghosh. 2025. [mmJEE-eval: A bilingual multimodal benchmark for evaluating scientific reasoning in vision-language models](#). In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 2268–2290, Mumbai, India. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics.
- Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#). *Preprint*, arXiv:2304.03442.
- Philip I Pavlik Jr, Hao Cen, and Kenneth R Koedinger. 2009. Performance factors analysis—a new alternative to knowledge tracing. In *Artificial Intelligence in Education*, pages 531–538. IOS Press.
- Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J Guibas, and Jascha Sohl-Dickstein. 2015. Deep knowledge tracing. In *Advances in neural information processing systems*, pages 505–513.
- Sashank Reddy, Igor Labutov, Siddhartha Banerjee, and Thorsten Joachims. 2016. Unbounded human learning: Optimal scheduling for spaced repetition. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1815–1824.
- Burr Settles and Brendan Meeder. 2016. A trainable spaced repetition model for language learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1248–1258.
- KURT VanLEHN. 2011. [The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems](#). *Educational Psychologist*, 46(4):197–221.
- Xinyi Wang, Jiacheng Li, Yu Zhang, and Others. 2025. Llm-powered multi-agent framework for goal-oriented learning. In *Proceedings of the AAI Conference on Artificial Intelligence*.
- Piotr A Wozniak. 1990. Optimization of learning. *Master’s thesis, University of Technology in Poznan*.
- Tao Wu, Jingyuan Chen, Wang Lin, Mengze Li, Yumeng Zhu, Ang Li, Kun Kuang, and Fei Wu. 2025. [Embracing imperfection: Simulating students with diverse cognitive levels using llm-based agents](#). *Preprint*, arXiv:2505.19997.
- Jiani Zhang, Xingjian Shi, Irwin King, and Dit-Yan Yeung. 2017. Dynamic key-value memory networks for knowledge tracing. In *Proceedings of the 26th international conference on World Wide Web*, pages 765–774.
- Yifan Zhang and Others. 2024. SPL: Socratic playground for learning with large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*.

## A Resources and Availability

To support reproducibility and further research, we provide open access to our platform components and simulation data under the **MIT License**:

- **Live Demo:** <https://intellicode.redomic.in>

- **Frontend:** <https://github.com/Redomic/Intellicode-frontend>
- **Backend:** <https://github.com/Redomic/Intellicode-backend>
- **Simulation Framework:** [https://github.com/Redomic/intellicode\\_student\\_sim](https://github.com/Redomic/intellicode_student_sim)

## B Mathematical Formulation Details

We formulate adaptive personalized education as a Partially Observable Markov Decision Process (POMDP):

$$\text{POMDP} = (\mathcal{S}, \mathcal{A}, \mathcal{O}, T, R, \gamma, b_0) \quad (2)$$

### B.1 State Space

The learner state at time  $t$  is:

$$\mathbf{S}_t = \{m_t, r_t, e_t, p_t, M_t, v_t\} \quad (3)$$

where:

- $m_t$ : mastery vector,  $m_{t,i} \in [0, 1]$  for topic  $i \in T$
- $r_t$ : review schedule, items with  $(q_{id}, \text{topics}, d_{\text{due}}, \text{interval}, \text{EF}, n_{\text{reviews}})$
- $e_t$ : engagement state, streak, last-seen timestamp, recent activity windows
- $p_t$ : preferences, skill level, modality, time budget, opt-outs
- $M_t$ : long-term memory, structured text sections on trends, misconceptions, insights
- $v_t$ : version, timestamp for auditing

We also maintain uncertainty  $u_{t,i}$  per topic, encoded as Beta parameters  $(\alpha_{t,i}, \beta_{t,i})$ .

### B.2 Observation Space

Observations are partial, noisy signals of state:

$$\mathcal{O}_t \in \{ \text{submission}, \text{hint\_request}, \text{session\_start}, \text{due\_review} \} \quad (4)$$

For each submission, we observe:

$$o_t = (q_{id}, y, \tau, h_{\text{cnt}}, \text{errors}, t_{\text{solve}}) \quad (5)$$

where  $y \in \{0, 1\}$  (pass/fail),  $\tau$  (timestamp),  $h_{\text{cnt}}$  (hints used), errors (semantic signals),  $t_{\text{solve}}$  (time on task).

### B.3 Action Space

The orchestrator (via agents) selects actions:

$$\mathcal{A}_t \in \{ \text{recommend\_item}, \text{hint}(l), \text{adjust\_schedule}, \text{intervene}, \text{feedback}(d) \} \quad (6)$$

where  $l \in \{1, 2, 3, 4, 5\}$  is hint level,  $d$  is feedback detail level.

### B.4 Reward Function

The reward proxies learning progress while penalizing inefficiency:

$$R_t = \underbrace{w_m \Delta m_t}_{\text{mastery gain}} + \underbrace{w_r \mathcal{K}[\text{review\_success}]}_{\text{retention}} - \underbrace{w_h h_{\text{cnt}}}_{\text{hint penalty}} - \underbrace{w_t \max(0, t_{\text{solve}} - \mu_t)}_{\text{time penalty}} \quad (7)$$

with regularizers for fairness (topic coverage) and engagement.

## C Learner State Updates

### C.1 State Initialization

From historical submissions, we compute initial mastery using a recency-weighted exponential moving average:

$$m_{t,i}^{(0)} = 0.6 \text{ success\_rate}_i + 0.4 \text{ recent\_success\_rate}_i + \mathcal{N}(0, \sigma_0^2) \quad (8)$$

We initialize Beta parameters as  $(\alpha_0, \beta_0) = (1, 1)$  (uninformative prior), and review queue empty with ease factor  $\text{EF}_0 = 2.5$ .

### C.2 Mastery Update Rule

Upon outcome  $y \in \{0, 1\}$  on a question tagged with topics  $Q$ :

$$m_{t,i} \leftarrow \begin{cases} \min(1, m_{t,i} + \alpha w_d w_r (1 - m_{t,i})) & (y=1), \\ \max(0, m_{t,i} - \beta w_d^{-1} w_r m_{t,i}) & (y=0) \end{cases} \quad (9)$$

where:

- $w_d \in \{0.8, 1.0, 1.2\}$  maps difficulty  $\in \{\text{Easy}, \text{Medium}, \text{Hard}\}$  (inverted for failure penalties)
- $w_r = \exp(-\Delta t / \tau_{\text{upd}})$  decays with recency

- Hint/time penalties:  $m_{t,i} \leftarrow m_{t,i} - \eta_h h_{\text{used}} - \eta_t \max(0, t_{\text{solve}} - \mu_i)$
- Momentum smoothing:  $m_{t,i} \leftarrow (1 - \lambda)m_{t,i} + \lambda m_{t,i}^{\text{new}}$  reduces jitter

### C.3 Proficiency Composite

Overall proficiency is a weighted composite:

$$\hat{p} = \sum_k w_k s_k \quad (10)$$

where  $s_k$  includes: topic mastery average (0.40), expertise rank (0.25), self-reported skill (0.20), recent success rate (0.10), streak normalization (0.05).

### C.4 Spaced Repetition Updates (SM-2)

For a review item with ease factor EF, we derive quality score  $q \in \{0, 1, 2, 3, 4, 5\}$  from:

$$q = \begin{cases} 5 & \text{fast, no hints} \\ 4 & \text{solved, minor delay} \\ 3 & \text{solved with hints} \\ \leq 2 & \text{failed/forgot} \end{cases} \quad (11)$$

Update:

$$EF' = \max(1.3, EF - 0.8 + 0.28q - 0.02q^2) \quad (12)$$

Intervals:  $I_1 = 1, I_2 = 6, I_n = \text{round}(I_{n-1} \cdot EF')$  days. Predicted recall:

$$R(\Delta t) = \exp(-\Delta t / \tau) \quad (13)$$

where  $\tau \propto EF'$ .

## D Reproducibility Details

1. Seeded random splits; fixed topic-graph snapshot.
2. Frozen prompt versions and role texts.
3. Logged hyperparameters:  $\alpha, \beta, w_d, \tau_{\text{upd}}, \lambda$ , proficiency weights.
4. Validation schemas for agent outputs (JSON specs).
5. Behavior signal preprocessing with masking rules.
6. Ablation code and offline evaluation scripts.