Agentic Orchestration for Adaptive Educational Recommendations: A Multi-Agent LLM Framework for Personalized Learning Pathways

Naina Chaturvedi
nc832@cs.rutgers.edu
Rutgers University
USA

Dr. Ananda Gunawardena
andy.guna@cs.rutgers.edu
Rutgers University
USA

Abstract

Educational personalization represents a unique challenge for recommender systems: learners require not just content recommendations, but dynamic curriculum adaptation, real-time feedback, and proactive intervention strategies that evolve over extended timescales. We present a novel multi-agent architecture that treats educational personalization as an emergent property of specialized agent collaboration rather than a monolithic recommendation model. Our framework deploys 18+ coordinated agents organized in a four-tier hierarchy spanning perception, domain expertise, coordination, and strategic planning. Through deployment on a learning platform serving 6,000+ active users, we demonstrate that hierarchical agent orchestration enables recommendation capabilities unachievable by single-model approaches: parallel domain-specific analysis, temporal stratification from millisecond feedback to multi-month roadmap generation, and graceful degradation under partial failures. We present the architectural principles, coordination protocols, and preliminary evidence that agentic systems offer a promising paradigm for next-generation personalized learning systems. Our work contributes both a concrete implementation blueprint and theoretical foundations for applying multi-agent LLM orchestration to complex recommendation domains beyond education.

CCS Concepts

• **Information systems → Recommender systems**; • **Computing methodologies → Artificial intelligence**; **Multi-agent systems**.

Keywords

Multi-Agent Systems, LLM Orchestration, Educational Recommendations, Personalized Learning, Adaptive Systems, Agentic AI

1 Introduction

Recommender systems have traditionally focused on matching users to items—products, movies, articles—based on historical preferences and collaborative signals. Educational recommendation presents fundamentally different challenges: learners don't merely consume content but actively transform through interaction with it. A student struggling with recursion today may excel at it tomorrow, invalidating yesterday's difficulty assessments. Unlike static item catalogs, educational content must adapt to each learner's evolving knowledge state, learning velocity, and long-term career objectives.

Recent advances in Large Language Models (LLMs) have enabled conversational tutoring systems [1, 2, 7, 9], yet these predominantly employ monolithic architectures where a single model attempts to simultaneously maintain pedagogical strategies, assess learner progress, generate personalized content, and predict future needs. This architectural choice creates three critical limitations: (1) cognitive overload as models struggle to balance competing objectives, (2) absence of specialization preventing deep expertise in any single recommendation facet, and (3) brittle failures where errors cascade across the entire system.

We propose treating educational personalization as a **multi-agent coordination problem** where specialized LLM-powered agents collaboratively generate adaptive learning experiences. Drawing inspiration from multi-agent systems research [10, 11], we architect a four-tier hierarchy: perception agents continuously observe learner behavior, domain agents provide specialized expertise in programming/system design/research, coordination agents orchestrate responses across domains, and strategic agents generate long-term personalized roadmaps. Unlike traditional recommender architectures, agents operate at different temporal scales—from real-time syntax feedback (milliseconds) to career pathway recommendations (months)—enabling true adaptive personalization.

Our contributions are threefold:

(1) **Novel Architecture**: A hierarchical multi-agent framework for educational recommendation that decomposes the personalization problem into 18+ specialized agents with explicit coordination protocols.
(2) **Implementation Evidence**: Deployment results from a production platform serving 6,000+ active learners across 5 career tracks, demonstrating emergent behaviors unachievable by single-model systems.
(3) **Design Principles**: Generalizable architectural patterns for applying agent orchestration to complex recommendation domains requiring expertise specialization, temporal stratification, and graceful degradation.

This work addresses the workshop's call for agentic systems in recommendation use-cases while providing concrete deployment insights for LLM-driven personalization at scale.

WSDM Companion '26, Boise, ID, USA

ACM ISBN 979-8-4007-2358-2/2026/02
https://doi.org/10.1145/3779211.3795741

2 Related Work

2.1 LLM-Driven Personalization

Recent work has explored LLMs for personalized systems. Chudziak and Kostka [1] demonstrated adaptive math tutoring using multi-agent architectures with graph-based knowledge retrieval, while Zhang et al. [2] simulated classroom interactions with LLM-empowered agents that model teacher-student dynamics. Chu et al. [4] provided a comprehensive survey of LLM agents in education, identifying key capabilities including memory, tool use, and planning. Maity and Kumar [5] explored generative AI's impact on intelligent tutoring systems. Recent work on LLM-based agents [19–22] and multi-agent frameworks [23–25] demonstrates sophisticated reasoning capabilities. However, these systems primarily employ monolithic models or loosely coordinated agents without explicit orchestration hierarchies for educational recommendation.

2.2 Multi-Agent Recommender Systems

Classical multi-agent approaches to recommendation have explored collaborative filtering with competing agents [11] and negotiation-based preference learning [14, 15]. Recent work on neuro-symbolic frameworks [3] has begun integrating LLMs into agent architectures for education, while general surveys on LLM-based agents [27, 28] highlight their potential across domains. However, comprehensive end-to-end hierarchical systems for educational personalization deployed at scale remain unexplored. We contribute the first fully-realized multi-tier agent architecture specifically designed for adaptive educational recommendation.

2.3 Adaptive Learning Systems

Intelligent Tutoring Systems (ITS) have long pursued personalization through knowledge tracing and adaptive instruction [30, 31]. Modern approaches increasingly leverage neural networks for skill assessment and content sequencing [32], with meta-analyses demonstrating their effectiveness [33]. Recent work emphasizes the transformative potential of AI-powered personalized learning [7] and the importance of learner choice in adaptive systems [6]. Student simulation for conversational tutors [8] and broader discussions on AI's evolution in education [9, 34, 35] highlight growing interest in agent-based approaches. Our contribution lies in demonstrating that treating personalization as emergent agent coordination—rather than centralized planning—enables capabilities beyond traditional ITS architectures.

3 Problem Formulation

Educational recommendation differs from traditional item recommendation in three fundamental ways:

Non-stationary User Models: A learner's knowledge state K_t evolves continuously through interaction: $K_{t+1} = f(K_t, I_t, R_t)$ where I_t represents learning interactions and R_t represents system responses. Traditional collaborative filtering assumes stable preferences; educational systems must model transformation.

Multi-Objective Optimization: Educational recommendations balance competing objectives: immediate understanding vs. long-term retention, challenge vs. frustration prevention, breadth vs. depth. Single reward formulations inadequately capture this complexity.

Temporal Stratification: Recommendations span multiple timescales simultaneously—real-time feedback on code syntax (ms), lesson sequencing (hours), skill progression (days), and career pathway planning (months). Monolithic models struggle to reason coherently across these scales.

We formulate educational personalization as a multi-agent sequential decision problem where specialized agents $\{A_1, ..., A_n\}$ collaborate to generate adaptive experiences. Each agent A_i maintains domain-specific expertise and observes different aspects of the learner state. The coordination mechanism determines when and how agents contribute to recommendations.

4 Multi-Agent Architecture

Our framework organizes agents into a four-tier hierarchy. We describe each tier and its coordination mechanisms.

4.1 Tier 1: Perception Agents

Perception agents continuously monitor learner activity and generate behavioral insights:

- **Behavioral Analytics Agent**: Tracks session patterns, engagement signals (time-on-task, revisit frequency), and learning velocity across topics.
- **Error Pattern Miner**: Analyzes mistake sequences to identify persistent misconceptions. For example, detecting repeated confusion between deep copy and shallow copy in data structures.
- **Affective Computing Agent**: Infers emotional states from interaction patterns—rapid topic switching may indicate frustration, sustained focus suggests flow states.

These agents operate asynchronously, generating insights that propagate to higher tiers without blocking the request-response cycle. This design enables proactive recommendations before learners explicitly request help.

4.2 Tier 2: Primary Domain Agents

Domain agents provide specialized expertise for different educational domains:

Code Analysis Agent: Evaluates programming submissions across multiple dimensions. Delegates to sub-agents:

- *Syntax Checker*: Identifies language-specific errors
- *Structure Analyzer*: Assesses code organization, modularity
- *Performance Analyzer*: Evaluates algorithmic complexity

System Design Agent: Assesses architectural decisions for distributed systems. Sub-agents include:

- *Scalability Evaluator*: Analyzes bottlenecks, load patterns
- *Tradeoff Analyst*: Examines CAP theorem considerations
- *Metrics Agent*: Evaluates monitoring strategies

Research Agent: Supports academic writing and critical thinking with sub-agents for clarity checking and argumentation evaluation.

The hierarchical structure (primary agent → sub-agents) enables parallel execution while maintaining coherent domain-level synthesis.

4.3 Tier 3: Coordination Agents

The **Central Coordinator** serves as the orchestration hub:

(1) **Intent Analysis**: Classifies incoming learner queries to determine required domains (e.g., "Why is my distributed cache slow?" triggers both Code Analysis and System Design agents).
(2) **Contextual Routing**: Routes requests to appropriate primary agents with learner history context.
(3) **Response Aggregation**: Synthesizes multi-agent outputs into coherent feedback, resolving conflicts when agents disagree.
(4) **Adaptive Difficulty**: Adjusts response complexity based on accumulated learner model (novice vs. advanced).

The **Quality Assurance Meta-Agent** reviews all outputs before delivery, checking for:

- Pedagogical soundness (no overwhelming detail for beginners)
- Technical accuracy across domain boundaries
- Consistency with prior feedback

This meta-review layer enables graceful degradation: if a domain agent produces flawed output, QA can request re-generation or provide corrective context.

4.4 Tier 4: Strategic Agents

Strategic agents operate on longer timescales:

Roadmap Generator: Creates personalized multi-month learning paths by:

- Analyzing career objectives (e.g., "Machine Learning Engineer")
- Assessing current skill levels across 50+ technical competencies
- Generating prerequisite-respecting skill graphs
- Recommending project sequences aligned with industry standards

Prediction Agent: Forecasts learner trajectories:

- Dropout risk using engagement time-series analysis (72% accuracy)
- Time-to-mastery estimates for upcoming skills
- Struggle area identification before failure occurs

Intervention Recommender: Generates proactive engagement strategies:

- Personalized micro-lesson suggestions when progress stalls
- Difficulty adjustments to maintain flow state
- Gamification trigger recommendations based on motivation profiles

Strategic agents don't participate in real-time request processing but periodically update the shared learner model that other tiers consume.

4.5 Agent Communication Protocol

Agents communicate through structured JSON payloads enabling explicit coordination. The protocol supports:

- **Parallel Execution**: Primary agents receive requests simultaneously, with sub-agents executing in parallel within each domain.
- **Context Preservation**: Learner state propagates through all agent interactions.
- **Failure Isolation**: Sub-agent failures don't cascade; primary agents aggregate available results.

5 Mining Agents and Insight Generation

Seven specialized mining agents collectively generate 50+ insight types that feed the recommendation pipeline:

5.1 Pattern Discovery

Identifies recurring behavioral sequences across the learner population and within individual learning trajectories. These patterns inform recommendation strategies without requiring explicit rules.

5.2 Trend Analysis

Tracks temporal evolution at multiple granularities—platform-wide trends and individual learner trajectories.

5.3 Predictive Analytics

Generates forward-looking recommendations including dropout risk scores, completion probability estimates, and struggle area predictions. The key innovation is closed-loop insight generation: mining agents observe system behavior, generate insights, those insights trigger recommendations, and the recommendation outcomes feed back to mining agents.

6 Implementation and Deployment

We deployed this architecture on a production learning platform.

6.1 Technical Stack

- **LLM Backend**: GPT-4 powers primary and coordination agents; GPT-3.5-turbo for sub-agents where lower latency is critical
- **Vector Store**: Pinecone for semantic search across learning materials and past interactions
- **Orchestration**: Custom Python framework managing agent lifecycle and communication
- **Learner Modeling**: PostgreSQL storing skill graphs, engagement timeseries, and error histories

6.2 Request Flow Example

Consider a learner query: "My distributed cache keeps returning stale data." The coordinator analyzes intent and routes to Code Analysis and System Design agents. Each agent evaluates the query from their domain perspective. The coordinator aggregates responses, and the QA Meta-Agent verifies explanation complexity matches learner level. This entire flow completes in 2-4 seconds with parallel agent execution.

7 Preliminary Results

While comprehensive controlled evaluation is ongoing, we present preliminary indicators from production deployment:

Table 1: System Scale Metrics

Metric	Value
Active learners	6,000+
Waitlist	2,000+
Agent count	18+
Primary domain agents	4
Specialized sub-agents	11
Mining agents	7
Meta-agents	1
Insight types generated	50+
Career tracks personalized	5
Real-world projects	8
Parallel sub-agent calls/query	2-5

7.1 Emergent Coordination Behaviors

Without explicit programming, we observe cross-domain collaboration, adaptive scaffolding, and proactive intervention. Mining agents successfully trigger re-engagement strategies 48 hours before predicted dropout events.

7.2 User Engagement Metrics

Comparing pre- and post-agent-deployment cohorts:

Table 2: Engagement Impact

Metric	Before	After
Avg. session duration	32 min	47 min
Questions/session	2.1	4.7
Return rate (week 2)	67%	81%
Project completion	43%	58%

These improvements suggest that specialized agent expertise and proactive recommendations meaningfully enhance learning experiences.

7.3 System Reliability

The multi-agent architecture demonstrates graceful degradation with fallback mechanisms and quality assurance checks catching 8% of domain agent outputs needing refinement before delivery.

8 Discussion

8.1 Advantages of Agentic Orchestration

Our implementation reveals three key benefits: expertise specialization, temporal stratification, and failure isolation. Unlike monolithic models, specialized agents develop deep domain expertise and operate at appropriate timescales while containing failures.

8.2 Challenges and Limitations

Agent communication adds latency (2-4 seconds vs. 1 second for monolithic models). Multiple agents must maintain coherent learner models. Traditional recommendation metrics inadequately capture educational outcomes. Coordinating 18+ agents requires carefully designed prompts and communication protocols.

8.3 Generalization Beyond Education

The architectural principles extend to healthcare (specialist medical agents), financial planning (risk/tax/portfolio agents), career coaching (skill-gap/culture/salary agents), and any domain requiring expertise specialization, temporal stratification, and graceful degradation.

9 Future Work

We identify five promising research directions: formal coordination protocols, multi-objective optimization, explainability and trust, lifelong learning agent architectures, and benchmark development for measuring emergent coordination capabilities.

10 Conclusion

We presented a multi-agent architecture for educational personalization that treats adaptive learning as an emergent property of specialized LLM agent collaboration. Our four-tier hierarchy demonstrates that agentic orchestration enables recommendation capabilities unachievable by monolithic models: parallel domain analysis, temporal stratification, and graceful degradation.

Deployment on a platform serving 6,000+ learners provides preliminary evidence that specialized agents coordinated through explicit protocols can meaningfully enhance learning experiences. Beyond education, the architectural principles generalize to complex recommendation domains requiring coordinated decision-making across multiple timescales and objectives. The future of recommendation lies not in building larger monolithic models, but in orchestrating specialized agents that collectively achieve what no single model can.

11 Ethical Considerations

Deploying AI agents in educational contexts raises important concerns around algorithmic bias, privacy, over-reliance, and transparency. We continuously audit for disparate impact, encrypt all learner data with explicit consent, design scaffolding that fades as mastery grows, and work toward agent-coordination visualizations that make decision-making processes interpretable.

References

[1] J. A. Chudziak and A. Kostka, "Ai-powered math tutoring: Platform for personalized and adaptive education," *arXiv preprint arXiv:2507.12484*, 2024.

[2] Z. Zhang, D. Zhang-Li, J. Yu, L. Gong, J. Zhou, Z. Hao, J. Jiang, J. Cao, H. Liu, Z. Liu, L. Hou, and J. Li, "Simclass: Simulating classroom education with llm-empowered agents," *arXiv preprint arXiv:2406.19226*, 2024.

[3] R. Hare and Y. Tang, "Toward generalized autonomous agents: A neuro-symbolic ai framework for integrating social and technical support in education," *arXiv preprint arXiv:2508.18406*, 2025.

[4] Z. Chu, S. Wang, J. Xie, T. Zhu, Y. Yan, J. Ye, A. Zhong, X. Hu, J. Liang, P. S. Yu, and Q. Wen, "Llm agents for education: Advances and applications," *arXiv preprint arXiv:2503.11733*, 2025.

[5] S. Maity and A. Kumar, "Generative ai and its impact on personalized intelligent tutoring systems," *arXiv preprint arXiv:2410.10650*, 2024.

[6] B. Clément, H. Sauzéon, D. Roy, and P.-Y. Oudeyer, "Improved performances and motivation in intelligent tutoring systems: Combining machine learning and learner choice," *arXiv preprint arXiv:2402.01669*, 2024.

[7] F. St-Hilaire et al., "A new era: Intelligent tutoring systems will transform online learning for millions," *arXiv preprint arXiv:2203.03724*, 2022.

[8] Z. Liu, S. X. Yin, G. Lin, and N. F. Chen, "Personality-aware student simulation for conversational intelligent tutoring systems," *arXiv preprint arXiv:2404.06762*, 2024.

[9] J. Wang et al., "Evolution of ai in education: Agentic workflows," *arXiv preprint arXiv:2504.20082*, 2025.

[10] M. Wooldridge, *An Introduction to MultiAgent Systems*, John Wiley & Sons, 2nd edition, 2009.

[11] J. Ferber, *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*, Addison-Wesley, 1999.

[12] D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang, D. Song, and J. Steinhardt, "Measuring mathematical problem solving with the math dataset," in *Proceedings of NeurIPS*, 2021.

[13] Korn Ferry, "The global talent crunch," Technical report, Korn Ferry Institute, 2023.

[14] T. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohmé, "Coalition structure generation with worst case guarantees," *Artificial Intelligence*, vol. 111, no. 1-2, pp. 209–238, 1999.

[15] S. Fatima, M. Wooldridge, and N. R. Jennings, "Multi-issue negotiation with deadlines," *Journal of Artificial Intelligence Research*, vol. 27, pp. 381–417, 2004.

[16] H. V. D. Parunak, "A survey of environments and mechanisms for human-human stigmergy," in *Environments for Multi-Agent Systems II*, pages 163–186, Springer, 2006.

[17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proceedings of NeurIPS*, pages 5998–6008, 2017.

[18] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.

[19] T. Brown et al., "Language models are few-shot learners," in *Proceedings of NeurIPS*, volume 33, pages 1877–1901, 2020.

[20] L. Ouyang et al., "Training language models to follow instructions with human feedback," in *Proceedings of NeurIPS*, volume 35, pages 27730–27744, 2022.

[21] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. Chi, Q. Le, and D. Zhou, "Chain-of-thought prompting elicits reasoning in large language models," in *Proceedings of NeurIPS*, volume 35, pages 24824–24837, 2022.

[22] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, and Y. Cao, "React: Synergizing reasoning and acting in language models," in *Proceedings of ICLR*, 2023.

[23] J. S. Park, J. C. O'Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein, "Generative agents: Interactive simulacra of human behavior," in *Proceedings of UIST*, pages 1–22, 2023.

[24] Q. Wu, G. Bansal, J. Zhang, Y. Wu, S. Zhang, E. Zhu, B. Li, L. Jiang, X. Zhang, and C. Wang, "Autogen: Enabling next-gen llm applications via multi-agent conversation," *arXiv preprint arXiv:2308.08155*, 2023.

[25] S. Hong et al., "Metagpt: Meta programming for a multi-agent collaborative framework," *arXiv preprint arXiv:2308.00352*, 2023.

[26] C. Qian, X. Cong, C. Yang, W. Chen, Y. Su, J. Xu, Z. Liu, and M. Sun, "Communicative agents for software development," *arXiv preprint arXiv:2307.07924*, 2023.

[27] Z. Xi et al., "The rise and potential of large language model based agents: A survey," *arXiv preprint arXiv:2309.07864*, 2023.

[28] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, W. X. Zhao, Z. Wei, and J.-R. Wen, "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, pp. 1–26, 2024.

[29] M. O. Riedl, "Human-centered artificial intelligence and machine learning," *Human Behavior and Emerging Technologies*, vol. 1, no. 1, pp. 33–36, 2019.

[30] K. R. Koedinger, J. C. Stamper, E. L. McLaughlin, and T. V. Nixon, "Using data-driven discovery of better student models to improve student learning," in *Proceedings of International Conference on Artificial Intelligence in Education*, pages 421–430, Springer, 2012.

[31] K. VanLehn, "The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems," *Educational Psychologist*, vol. 46, no. 4, pp. 197–221, 2011.

[32] V. Aleven, B. M. McLaren, J. Sewall, M. van Velsen, O. Popescu, S. Demi, M. Ringenberg, and K. R. Koedinger, "Example-tracing tutors: Intelligent tutor development for non-programmers," *International Journal of Artificial Intelligence in Education*, vol. 26, no. 1, pp. 224–269, 2016.

[33] J. A. Kulik and J. D. Fletcher, "Effectiveness of intelligent tutoring systems: A meta-analytic review," *Review of Educational Research*, vol. 86, no. 1, pp. 42–78, 2016.

[34] N. Pinkwart, "Another 25 years of aied? challenges and opportunities for intelligent educational technologies of the future," *International Journal of Artificial Intelligence in Education*, vol. 26, no. 2, pp. 771–783, 2016.

[35] I. Roll and R. Wylie, "Evolution and revolution in artificial intelligence in education," *International Journal of Artificial Intelligence in Education*, vol. 26, no. 2, pp. 582–599, 2016.