Jiang, B. et al. (Eds.) (2025). Proceedings of the 33rd International Conference on Computers in Education. Asia-Pacific Society for Computers in Education

LLM Agent Collaboration for Educational Strategy Design

Bjanka VRLJIĆ[a*], Ivica BOTIČKI[a] & Kristijan POJE[a]
[a]*Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia*
*bjanka.vrljic@fer.hr

Abstract: Large language models (LLMs) can be organized into multi-agent systems (MAS) that simulate expert panels, but coordination remains a challenge. In education, effective support often requires integrating insights from data analytics, curriculum design, and pedagogy. We present a moderated, embedding-driven MAS framework for educational decision-making. Role-specialized LLM agents generate structured recommendations, which are embedded into a shared semantic space. A central moderator measures alignment and agreement, computes a consensus embedding, and translates it back into natural language. This consensus is returned to the agents in subsequent rounds, guiding convergence while preserving distinct perspectives. A simulated case study illustrates the process in the context of secondary education. The results highlight the potential of embedding-based moderation to produce coherent, interpretable recommendations.

Keywords: Natural language processing, large language models, multi-agent systems, personalized learning, educational decision-making

1. Introduction

LLMs are increasingly used within MAS to simulate expert collaboration and tackle complex tasks (Park et al., 2023; Yang et al., 2024). By assigning complementary roles and encouraging iterative exchange, these systems can approximate the dynamics of human expert panels, where knowledge from multiple domains must be integrated to resolve complex challenges. Education represents a domain where such coordinated reasoning could provide substantial value. Supporting learners often requires expertise spanning data analysis, curriculum planning, and pedagogy. However, existing approaches rarely succeed in combining these perspectives into a unified and interpretable process.

This paper explores how a moderated, embedding-driven MAS framework can be adapted to educational settings. The approach builds on recent advances in consensus generation (Amirkhani & Barshooi, 2022; Mason & Roberts, 2023) and semantic similarity analysis (Chandrasekaran & Mago, 2021; Reimers & Gurevych, 2019). Role-specialized LLM agents interact under the guidance of a central moderator that aggregates their outputs into an embedding-based consensus, then translates it back into natural language. This approach shows promise in blending NLP-driven insights with educational expertise to support transparent, data-informed, and collaborative decision-making in learning environments.

2. System Overview

The proposed system employs a moderated multi-agent workflow in which role-specialized LLMs collaborate under the coordination of a central moderator. The process begins with a high-level problem statement that defines the context, constraints, and desired goals. The moderator decomposes this problem into role-specific subgoals and distributes them to the agents, each instantiated with a prompt defining its expertise and reasoning style.

Agents generate structured outputs containing both a decision and an accompanying explanation. These outputs are transformed into vector representations using embeddings, which provide a shared semantic space for comparison. The moderator evaluates alignment to subgoals and agreement between agents through cosine similarity. It then computes an embedding-based consensus by averaging the embeddings of the agents' outputs.

The system makes the consensus understandable by applying Vec2Text (Morris et al., 2023), a method that converts embeddings back into natural language. This consensus statement is passed to the agents in the following round, alongside the original task description. Agents adapt their reasoning, taking this shared consensus into account while still maintaining their role-specific perspectives. Through repeated cycles, the system gradually moves toward stronger alignment and a coherent final recommendation.

The iterative process concludes once convergence is achieved, further refinements no longer yield substantive changes, or a predefined iteration limit is reached. The resulting output consists of a decision recommendation and an interpretable explanation that reflects the collaborative reasoning of all agents.

3. Application Example

3.1 Scenario Description

To illustrate the operation of the proposed framework, a hypothetical case study was constructed in the context of secondary education. The problem statement is as follows:

A 15-year-old student is underperforming in Algebra, Geometry, and Physics and lacks consistent study habits. The school wants an actionable, sustainable plan that improves grades, builds study discipline, and supports college readiness. Given available tools (Khan Academy, Quizlet, PhET), a peer-mentoring program, and 6–8 hours/week capacity, what should we implement this year to produce the best outcomes?

This scenario provides a structured setting for evaluating how the system manages competing goals, integrates multiple forms of expertise, and produces an actionable educational plan.

3.2 Expert Agents

Three role-specialized agents are instantiated to reflect complementary perspectives in educational planning. The *Education Data Analyst* (A_1) interprets student performance data and identifies predictors of success. The *Curriculum Designer* (A_2) connects learning gaps to curriculum standards, selects suitable resources, and structures a feasible learning plan. The *Pedagogy and Assessment Specialist* (A_3) designs study routines that sustain engagement and integrates assessments to monitor progress. The agents are instructed to recommend one of three intervention strategies:

G_1: Personalized learning path with adaptive practice
G_2: Peer-mentoring and habit-building with technological support
G_3: Mastery-based pacing with embedded formative assessment

3.3 Iterative Refinement

Agents generate structured outputs consisting of a decision and an explanation. These are converted into embeddings and assessed for alignment with subgoals and inter-agent agreement. The moderator computes an embedding-based consensus, which is then inverted into natural language and redistributed to the agents alongside the original problem description. Through this iterative loop, agents adapt their reasoning in response to the shared consensus while retaining their role-specific orientations.

In the initial round, agents diverge in their recommendations, each emphasizing priorities tied to their domain expertise. Subsequent rounds demonstrate increasing convergence, as exposure to the consensus text gradually promotes alignment without enforcing uniformity. By the final round, the agents converge on the recommendation of a personalized learning path with adaptive practice (G_1). Figure 1 illustrates this progression, with principal component analysis (PCA) projections of agent outputs that are initially scattered but gradually move closer together, forming a compact cluster by the final iteration.

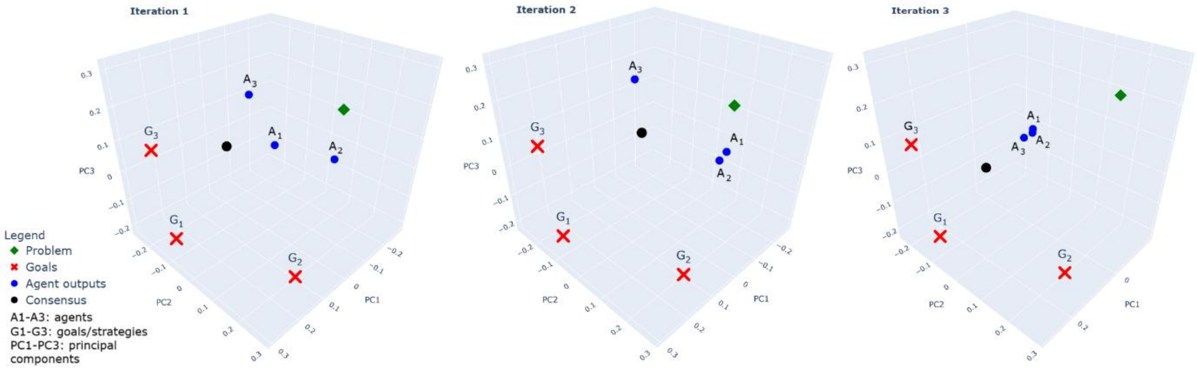

Figure 1. PCA visualization of agent outputs across three iterations.

4. Conclusion

This paper presented a multi-agent framework for educational decision-making, where role-specialized LLM agents iteratively collaborate under the guidance of a central moderator. By combining semantic similarity analysis with embedding-to-text inversion, the system generates an interpretable consensus. Iterative feedback cycles enable agents to refine their reasoning while preserving diverse perspectives. A simulated case study shows clear gains in semantic alignment across iterations, suggesting potential for supporting decision-making in learning environments where multiple expert viewpoints must be integrated. In practice, it could serve as a decision-support tool for teachers, administrators, or policymakers, helping integrate perspectives from data analysis, curriculum planning, and pedagogy into actionable strategies. The system could tailor interventions for individual students or guide resource allocation. However, the framework depends on prompt and embedding quality, which may introduce bias or instability. Future work will validate the framework in authentic educational contexts to ensure its practical value and integration with existing tools and workflows.

References

Amirkhani, A., & Barshooi, A. H. (2022). Consensus in multi-agent systems: a review. *Artificial Intelligence Review, 55*(5), 3897–3935. https://doi.org/10.1007/s10462-021-10097-x

Chandrasekaran, D., & Mago, V. (2021). Evolution of Semantic Similarity – A Survey. *ACM Computing Surveys, 54*(2), 1–37. https://doi.org/10.1145/3440755

Mason, J., & Roberts, L. D. (2023). Consensus moderation: the voices of expert academics. *Assessment & Evaluation in Higher Education, 48*(7), 926-937. https://doi.org/10.1080/02602938.2022.2161999

Morris, J., Kuleshov, V., Shmatikov, V., & Rush, A. (2023). Text Embeddings Reveal (Almost) As Much As Text. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing* (pp. 12448-12460). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.emnlp-main.765

Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (pp. 3982-3992).

Park, J.S., O'Brien, J.C., Cai, C.J., Morris, M.R., Liang, P., & Bernstein, M.S. (2023). Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (pp. 1-22).

Yang, Y., Peng, Q., Wang, J., Wen, Y., & Zhang, W. (2024). LLM-based Multi-Agent Systems: Techniques and Business Perspectives. *arXiv preprint arXiv:2411.14033*.