Journal for Research and Practice in College Teaching
http://journals.uc.edu
2025, Volume 10, Issue 2

Integrating Contemporary Topics and Methodologies into Experiential Learning for Business Analytics Students

Ganesh Mani Prasad Chalasani
Carnegie Mellon University

This reflection explores the integration of contemporary topics and methodologies into experiential learning for business analytics students at Carnegie Mellon University. It highlights curriculum innovations, including the incorporation of digital twins, collective intelligence, and large language models (LLMs), and describes a novel AI-augmented system for capstone project team formation. The authors also allude to insights from a pilot study on collaborative deliberation methods using swarm intelligence platforms. Key takeaways emphasize the pedagogical value of real-world, data-driven projects and the importance of scaffolding support to address domain complexity and diverse student backgrounds. The work potentially demonstrates how LLMs can enhance both instructional design and educational operations research.

Ganesh Mani (GM):

Over the past three years, I have taught Experiential Learning courses at the Tepper School of Business at Carnegie Mellon University, a top-ranked institution known for its rigorous AI and analytics programs. In this reflection, I share insights from redesigning the capstone preparation course and curriculum for Master of Science in Business Analytics students—both full-time and part-time—with a focus on aligning instructional content with emerging trends in data science and artificial intelligence.

In the first year of teaching the course, I recognized a need to modernize the analytics curriculum, which had traditionally emphasized optimization techniques. To better equip students for the evolving demands of real-world, data-driven analytics projects, I introduced three contemporary themes: digital twins, collective intelligence, and large language models (LLMs). This shift was not only content-driven but also pedagogically motivated—to foster critical (computational) thinking, adaptability, and interdisciplinary awareness.

Because of travel commitments to CMU Africa and the IEEE Africon Conference during the (Fall 2023) semester, I invited my colleague and collective intelligence expert, Professor Anita Woolley, to deliver a guest lecture. Her session provided students with a research-grounded perspective on group decision-making and emergent intelligence, complementing the technical focus of the course.

Although students initially found these topics somewhat disconnected, their engagement increased significantly during the LLM module, particularly during a PaLM 2 hackathon we facilitated in collaboration with Google. This hands-on experience allowed students to explore the capabilities of generative AI in a structured, collaborative setting.

These approaches proved valuable as students tackled a range of capstone projects—some with abundant data, others with very little. For example, digital twin simulation methodologies were especially useful in generating synthetic data for projects lacking real-world datasets. Digital twins are virtual, abstract representations of entities or systems (IBM, 2021). Just as flight simulators allow rookie pilots to train for rare scenarios, a more everyday example might involve simulating service at a coffee shop to estimate wait times and make staffing decisions.

My collaborator, Prasad Chalasani (PC), played a key role in supporting the capstone project assignments during the past academic year. The primary objective was to assign student teams—each consisting of 4 to 5 members—to suitable capstone projects. Building on a legacy optimization script originally developed by a former faculty member and partially inspired by Magnanti & Natarajan (2018), we enhanced the system by integrating contemporary Generative AI capabilities.

Specifically, we employed large language models (LLMs) to extract structured metadata from project descriptions, identifying the key skills required for each project. These were then heuristically matched with students' self-reported skills and project preferences. This AI-augmented approach streamlined the team formation process and improved alignment between student capabilities and project demands.

Optimization is a foundational tool in business analytics but setting up such models—especially for multi-objective problems—can be daunting. Our approach demonstrates how LLMs can now assist not just in coding but in problem formulation itself, enabling what might be called "vibe optimization." Below, PC outlines the technical details of the process.

PC:

LLM-Enhanced Optimization for Educational Resource Allocation: We developed an intelligent optimization system for matching students to capstone projects by leveraging OpenAI's *o1* model for mathematical problem formulation combined with computational solving via Python's

PuLP library (Dunning et al., 2011). Our approach demonstrates how large language models can enhance educational operations research beyond traditional coding assistance.

LLM-Guided Problem Architecture: Rather than immediately implementing a solution, we first consulted the *o1* model to properly formulate this complex multi-objective optimization challenge. The problem required balancing student preferences, project skill requirements, and pedagogical constraints around team composition.

Through iterative collaboration with the LLM, we evolved from basic preference-matching to a sophisticated synergy-coverage model treating student teams as collaborative units. This key insight—facilitated by the LLM—shifted our focus from matching individual students to requirements toward optimizing how teams' composite skills meet project needs, better reflecting the collaborative nature of capstone work.

Mathematical Model Development: Our collaboration with *o1* produced a Mixed Integer Linear Program featuring binary assignment variables, continuous skill coverage variables, and an objective function that maximizes student preference satisfaction while minimizing skill gaps through penalized shortfall variables. The LLM guided us toward an elegant formulation where teams collectively contribute skills to projects, with gaps captured as penalties in the optimization. Roughly speaking, this means that the optimizer is discouraged from finding solutions where students' self-assessed skill-levels fall short of the levels required for the project.

Implementation and System Design: We implemented dual optimization approaches—individual skill-matching and team synergy methods with flexible parameterization allowing administrators to balance student preferences against skill coverage requirements. Our system generates comprehensive analytics for both student assignments and project skill coverage.

The complete pipeline extracts project requirements from task descriptions, processes student survey responses capturing preferences and self-assessed skills, handles incomplete survey responses through synthetic student generation, and produces detailed reports for administrators and students.

Educational Innovation Impact: Our work demonstrates using LLMs for fundamental problem architecture rather than mere implementation. The *o1* model helped us recognize that effective capstone allocation isn't just about individual preference satisfaction, but creating teams where collective capabilities align with learning objectives—a more pedagogically sound approach.

This human-LLM collaboration resulted in a production-ready system balancing student satisfaction with educational outcomes, showcasing how LLMs can enhance educational operations research from problem conception through implementation. Our framework provides a replicable model for institutions facing similar resource allocation challenges, demonstrating the potential for AI-assisted educational administration that maintains pedagogical integrity while optimizing student experiences.

GM:

In the current academic year (2024-25), I also served as the inaugural Director of Collaborative AI for our school and was able to further amplify the human-machine teaming aspects. For instance, I extended the pedagogical explorations through a pilot study conducted in collaboration with the Eberly Center for Teaching Excellence. The study investigated the impact of different collaborative deliberation methods on student engagement and decision-making in an experiential learning environment. Using a new platform—*Thinkscape* from Unanimous.ai (Rosenberg et al., 2025)—that integrates LLMs with collective intelligence, we compared traditional group formation and whiteboard use with real-time virtual deliberation supported by an Infobot. A parallel comparison was conducted with Zoom breakout rooms among the geographically distributed part-time student cohort. The task involved a simplified financial asset allocation scenario, allowing us to focus on process rather than domain complexity.

Key metrics included consensus efficiency, decision quality, and inclusivity. The results revealed meaningful differences in how each method supported engagement and decision-making, with notable insights into the scalability and potential of Conversational Swarm AI. This research contributes to the growing body of innovative educational practices aimed at enhancing student participation, improving decision quality, and fostering rapid consensus-building—skills increasingly vital in both academic and organizational contexts.

One key takeaway from the analytics capstone experience is the value of immersing students in real-world, industry-relevant challenges. These projects provide a unique opportunity to engage with current issues and apply analytical skills in a practical setting. However, they also present several pedagogical challenges: ensuring access to mature and usable data—essential for meaningful analysis—while accommodating the diverse backgrounds and skill levels of students, all within the tight timeframe of a single semester.

Some projects, such as one focused on nuclear energy strategy, underscore the added complexity of domain-specific knowledge. In these cases, students often require time to build foundational understanding through industry and regulatory whitepapers or through multiple knowledge transfer sessions with capstone sponsors, who are typically subject matter experts. This highlights the importance of scaffolding support and resources—such as faculty advisors, alumni experts, and LLMs—to help students ramp up quickly and deliver actionable insights within a limited timeframe.

Journal for Research and Practice in College Teaching 2025, Volume 10, Issue 2
http://journals.uc.edu

Acknowledgements
GM gratefully acknowledges the support of the Eberly Center at Carnegie Mellon University, including the Generative AI in Teaching as Research (GAITAR) fellowship. Special thanks to H. Elisabeth Ellington for her leadership in the writing workshops and for encouraging submission to this journal.

References

IBM. (2021, August 5). What is a digital twin? IBM. https://www.ibm.com/think/topics/what-is-a-digital-twin

Magnanti, T. L., & Natarajan, K. (2018). Allocating students to multidisciplinary capstone projects using discrete optimization. Interfaces, 48(3), 204–216. https://doi.org/10.1287/inte.2017.0940

Dunning, I., Mitchell, S., & O'Sullivan, M. (2011). PuLP: A linear programming toolkit for Python. Optimization Online. https://optimization-online.org/?p=11731

Rosenberg, L., Schumann, H., Dishop, C., Willcox, G., Woolley, A., & Mani, G. (2025). Large-scale group brainstorming and deliberation using swarm intelligence and generative AI. In Proceedings of the 27th International Conference on Enterprise Information Systems (ICEIS 2025), Volume 1 (pp. 867–872). SCITEPRESS. https://doi.org/10.5220/0013379800003929