

Special Issue: Artificial Intelligence and Pedagogical Innovation

- 4. Introduction to the Special Issue on Artificial Intelligence and Pedagogical Innovation**
Eric Breimer, Special Issue Guest Editor, Siena University
Kevin Mentzer, Nichols College
Ira Goldstein, Siena University
David Yates, Bentley University
Michelle Louch, University of Pittsburgh Greensburg
Mark Pisano, Southern Connecticut University
- 7. Comparing Augmented Generative AI Platforms and Programming-Based Tools for Teaching Data Analytics to Undergraduate Business Students**
Wenqi Zhou, Duquesne University
Pinar Ozturk, Duquesne University
Kathleen S. Hartzel, Duquesne University
Carlos Salazar-Betancourt, University of Rhode Island
- 21. Integrating Artificial Intelligence into Undergraduate Business Programs: An Analysis of AACSB-Accredited Schools**
Kiku Jones, Quinnipiac University
Lori Leonard, University of Tulsa
Wendy Ceccucci, Quinnipiac University
Guido Lang, Quinnipiac University
- 31. Teaching Case**
Developing Critical Thinking in Data Analytics Education: A Teaching Case Evaluating ChatGPT Responses to a Visualization
Benjamin Larson, Auburn University of Montgomery
Jeffrey Bohler, Auburn University of Montgomery
Nandini Bolekar, Auburn University of Montgomery
- 44. Teaching Activity**
Introducing Data Analytics and AI Collaboration to Novice Students: An Assignment Using Real Business Data
Thilini Ariyachandra, Xavier University
- 63. Teaching with Generative AI: Ethics and Analytical Confidence in Undergraduate IS Education**
Asish Satpathy, Arizona State University
Jayaram Balachander, Arizona State University
- 75. Invited Article**
Designing for Tension: Balancing the Pace of Trust and the Pace of Innovation in AI-Integrated Higher Education
David Firth, University of Montana
Jason Triche, University of Montana
Isho Tama-Sweet, University of Montana

The **Information Systems Education Journal** (ISEDJ) is a double-blind peer-reviewed academic journal published by **ISCAP** (Information Systems and Computing Academic Professionals). Publishing frequency is five times per year. The first year of publication was 2003.

ISEDJ is published online (<https://isedj.org>). Our sister publication, the Proceedings of the ISCAP Conference (<https://iscap.us/proceedings>) features all papers, abstracts, panels, workshops, and presentations from the conference.

The journal acceptance review process involves a minimum of three double-blind peer reviews, where both the reviewer is not aware of the identities of the authors and the authors are not aware of the identities of the reviewers. The initial reviews happen before the ISCAP conference. All papers, whether award-winners or not, are invited to resubmit for journal consideration after applying feedback from the Conference presentation. Award winning papers are assured of a publication slot; however, all re-submitted papers including award winners are subjected to a second round of three blind peer reviews to improve quality and make final accept/reject decisions. Those papers that are deemed of sufficient quality are accepted for publication in the ISEDJ journal. Currently the target acceptance rate for the journal is under 35%.

Information Systems Education Journal is pleased to be listed in the Cabell's Directory of Publishing Opportunities in Educational Technology and Library Science, in both the electronic and printed editions. Questions should be addressed to the editor at editor@isedj.org or the publisher at publisher@isedj.org. Special thanks to volunteer members of ISCAP who perform the editorial and review processes for ISEDJ.

2026 ISCAP Board of Directors

Amy Connolly
James Madison University
President

Michael Smith
Georgia Institute of Technology
Vice President

Jeff Cummings
Univ of NC Wilmington
Past President

David Firth
University of Montana
Director

Mark Frydenberg
Bentley University
Director/Secretary

Leigh Mutchler
James Madison University
Director

RJ Podeschi
Millikin University
Director/Treasurer

Bryan Reinicke
Rochester Institute of
Technology / Director

Jeffrey Babb
West Texas A&M University
Director/Curricular Matters

Eric Breimer
Siena University
Director/2026 Conf Chair

Thomas Janicki
Univ of NC Wilmington
Director/Meeting Planner

Xihui "Paul" Zhang
University of North Alabama
Director/JISE Editor

Copyright © 2026 by Information Systems and Computing Academic Professionals (ISCAP). Permission to make digital or hard copies of all or part of this journal for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial use. All copies must bear this notice and full citation. Permission from the Editor is required to post to servers, redistribute to lists, or utilize in a for-profit or commercial use. Permission requests should be sent to Kevin Mentzer, Editor, editor@isedj.org.

INFORMATION SYSTEMS EDUCATION JOURNAL

Editors

Kevin Mentzer
Editor
Nichols College

Ira Goldman
Associate Editor
Siena University

David Yates
Associate Editor
Bentley University

Michelle Louch
Teaching Cases & Exercises
Editor
University of Pittsburgh - Greensburg

Mark Pisano
Teaching Cases & Exercises
Associate Editor
Southern Connecticut
State University

Thomas Janicki
Publisher
Univ of NC Wilmington

David Woods
Assistant Publisher
Miami University
Regionals

Paul Witman
Emeritus Editor
(2021-2026)
California Lutheran
University

Jeffry Babb
Emeritus Editor
(2016-2021)
West Texas A&M
University

Donald Colton
Emeritus Editor
(2003-2010)
Brigham Young University
Hawaii

Introduction to the Special Issue on Artificial Intelligence and Pedagogical Innovation

Eric Breimer
Siena University
Special Issue Guest Editor

ISEDJ Editorial Team

Kevin Mentzer
Nichols College

Ira Goldstein
Siena University

David Yates
Bentley University

Michelle Louch
University of Pittsburgh Greensburg

Mark Pisano
Southern Connecticut University

Recommended Citation: Breimer, E., Mentzer, K., Goldstein, I., Yates, D., Louch, M., Pisano, M., (2026). Introduction to the Special Issue on Artificial Intelligence and Pedagogical Innovation. *Information Systems Education Journal*, v24(n6) pp 4-6. DOI# <https://doi.org/10.62273/ZYIS6668>

OVERVIEW

Artificial intelligence, and generative artificial intelligence in particular, has rapidly moved from an emerging technology to a central concern for information systems educators. Its influence is visible across curriculum design, classroom practice, student learning, assessment, academic integrity, analytics education, business-school strategy, and institutional governance. This special issue brings together five papers and one invited editorial that examine how educators,

programs, and universities are responding to this moment.

The contributions in this issue address the integration of GenAI into undergraduate analytics courses, the development of AI collaboration skills among novice learners, the trade-offs between augmented GenAI platforms and traditional programming tools, the state of undergraduate AI programs at AACSB-accredited business schools, the use of AI to develop critical thinking in data analytics education, and the

organizational tensions created by AI adoption in higher education. Collectively, these papers demonstrate that AI in information systems education is not merely a question of tool adoption. Rather, it requires careful attention to pedagogy, ethics, assessment, critical thinking, workforce readiness, and institutional trust.

For information systems educators, the current AI moment is especially important. The IS discipline has long occupied a space between technical possibility and organizational value. We teach students not only how technologies work, but how they shape decisions, processes, organizations, and society. Generative AI makes this bridging role even more important. AI can lower barriers to complex technical work, support personalized feedback, and help students engage with data and ideas in new ways. At the same time, it can produce inaccurate outputs, obscure assumptions, encourage overreliance, complicate authorship, and deepen inequities among students with different levels of access and prior experience.

Several papers in this issue focus on classroom-level implementation, offering practical models for integrating AI into analytics, visualization, and data mining instruction. Others examine AI at the program and institutional levels, asking how business schools and universities should respond structurally to the demands of an AI-augmented world. Across the issue, a shared theme emerges: the goal is not simply to add AI tools to existing courses. Rather, the challenge is to design learning experiences that help students use AI thoughtfully, evaluate AI critically, and understand when human judgment must remain central.

We begin with the invited editorial by Firth, Triche, and Tama-Sweet, *Designing for Tension: Balancing the Pace of Trust and the Pace of Innovation in AI-Integrated Higher Education*. This piece provides a conceptual frame for many of the concerns that are present throughout the special issue. The authors argue that higher education faces a structural tension between the pace of technological innovation and the pace required to build institutional trust. They distinguish between AI literacies, which emphasize critical understanding, ethical judgment, and academic rigor, and AI competencies, which emphasize applied capability, workforce relevance, and experimentation. Rather than treating this tension as a problem to be resolved, the authors encourage institutions to design for it. Their central contribution is a reminder that AI

integration is not only a pedagogical challenge but also an organizational one. Sustainable adoption requires shared values, transparent governance, faculty support, student partnership, and institutional structures that allow innovation to proceed without sacrificing trust.

Satpathy and Balachander, in *Teaching with Generative AI: Ethics and Analytical Confidence in Undergraduate IS Education*, provide a scalable model for integrating GenAI and ethical reasoning into a required undergraduate analytics course. Their intervention embeds tools such as ChatGPT and Google Notebook into existing modules on regression, visualization, and interpretation, while also introducing a "Mindful AI" component focused on fairness, accountability, bias, and societal implications. A particular strength of this paper is its emphasis on preserving foundational analytical skills. Students first conduct analyses using standard methods and then use GenAI to generate explanations, compare interpretations, and evaluate the reliability of AI-generated output. The paper contributes a practical model for faculty who want to modernize analytics instruction without abandoning core learning objectives.

Ariyachandra's teaching activity, *Introducing Data Analytics and AI Collaboration to Novice Students: An Assignment Using Real Business Data*, addresses a related challenge at the introductory level. Many students in introductory IS courses have limited experience with both data analytics and AI tools. Rather than teaching these skills separately, the assignment asks students to work through a real-world business scenario using Starbucks beverage nutritional data. Students move through data exploration, dashboard development, AI-supported refinement, and reflection. The contribution of this paper is its focus on AI collaboration as a learnable skill. Students are not asked to use AI as a shortcut, but as a source of feedback, alternative perspectives, and revision support.

Zhou, Ozturk, Hartzel, and Salazar-Betancourth, in *Comparing Augmented Generative AI Platforms and Programming-Based Tools for Teaching Data Analytics to Undergraduate Business Students*, examine a central instructional trade-off: how should educators balance the accessibility of AI-augmented analytics tools with the technical depth of traditional programming environments? Their study compares SAP Analytics Cloud, including GenAI-supported features, with R GUI in an undergraduate business data mining course. The paper highlights that augmented GenAI tools can reduce technical friction and help students

focus more quickly on interpretation and application. At the same time, traditional programming tools provide flexibility, procedural understanding, and technical rigor. The authors do not present the choice as either/or. Instead, they argue that tool selection should be aligned with pedagogical priorities, student background, faculty readiness, and desired learning outcomes.

Jones, Leonard, Ceccucci, and Lang broaden the conversation from individual courses to undergraduate business curricula in *Integrating Artificial Intelligence into Undergraduate Business Programs: An Analysis of AACSB-Accredited Schools*. Their study reviews AI programs at U.S. universities with AACSB-accredited business schools, examining where undergraduate AI programs are housed and what themes they emphasize. The results show that undergraduate AI programs are far more common in computing and engineering than in business schools. While computing and engineering programs tend to emphasize technical themes, business programs are more likely to focus on AI for value creation, managerial decision-making, and business applications. This paper makes an important contribution by documenting a curricular gap at the program level.

Larson, Bohler, and Bolekar, in *Developing Critical Thinking in Data Analytics Education: A Teaching Case Evaluating ChatGPT Responses to a Visualization*, focus on one of the most important competencies for AI-augmented work: the ability to question AI output. Their teaching case asks students to evaluate ChatGPT's interpretation of a wage-by-tenure visualization that subtly violates linear regression assumptions. By prompting AI, comparing responses, refining prompts, and discussing whether the AI correctly recognizes the analytical issue, students learn that AI can provide useful scaffolding but cannot replace statistical reasoning, domain knowledge, or human judgment.

Taken together, the papers in this issue show that AI pedagogical innovation is occurring at multiple levels. At the assignment level, educators are designing activities that help students collaborate with AI, revise their work, and evaluate outputs. At the course level, faculty are integrating GenAI into analytics instruction while maintaining foundational skills and ethical reflection. At the tool level, instructors are reconsidering the balance between technical rigor and accessibility. At the program level, business schools are evaluating how AI should be represented in

undergraduate curricula. At the institutional level, universities are wrestling with governance, trust, experimentation, and shared responsibility.

Ultimately, this special issue suggests that the future of AI in IS education will depend less on any single tool and more on the quality of our pedagogical choices. Generative AI will continue to change, and today's tools will be replaced or transformed. What should endure are the habits we help students develop: asking better questions, evaluating evidence, recognizing limitations, communicating clearly, considering stakeholders, reflecting ethically, and using technology to augment rather than replace human judgment. The papers in this issue offer a timely and practical contribution to that work. They demonstrate that information systems educators are well positioned not only to respond to AI, but to lead in shaping how AI is understood, taught, governed, and used in higher education.

References

- Ariyachandra, T. (2026). Introducing Data Analytics and AI Collaboration to Novice Students: An Assignment Using Real Business Data. *Information Systems Education Journal*.
- Firth, D., Triche, J., & Tama-Sweet, I. (2026). Designing for Tension: Balancing the Pace of Trust and the Pace of Innovation in AI-Integrated Higher Education. *Information Systems Education Journal*.
- Jones, K., Leonard, L., Ceccucci, W., & Lang, G. (2026). Integrating Artificial Intelligence into Undergraduate Business Programs: An Analysis of AACSB-Accredited Schools. *Information Systems Education Journal*.
- Larson, B. E., Bohler, J. A., & Bolekar, N. (2026). Developing Critical Thinking in Data Analytics Education: A Teaching Case Evaluating ChatGPT Responses to a Visualization. *Information Systems Education Journal*.
- Satpathy, A., & Balachander, J. (2026). Teaching with Generative AI: Ethics and Analytical Confidence in Undergraduate IS Education. *Information Systems Education Journal*.
- Zhou, W., Ozturk, P., Hartzel, K. S., & Salazar-Betancourth, C. (2026). Comparing Augmented Generative AI Platforms and Programming-Based Tools for Teaching Data Analytics to Undergraduate Business Students. *Information Systems Education Journal*.

Comparing Augmented Generative AI Platforms and Programming-Based Tools for Teaching Data Analytics to Undergraduate Business Students

Wenqi Zhou
zhouw@duq.edu

Pinar Ozturk
ozturkp@duq.edu

Kathleen S. Hartzel
hartzel@duq.edu

Duquesne University
Pittsburgh, PA

Carlos Salazar-Betancourth
carlos.salazar@uri.edu
University of Rhode Island
Kinston, RI

Abstract

As data mining and predictive modeling skills become increasingly crucial, educators face challenges in selecting teaching tools that effectively balance technical depth, ease of use, and real-world applicability. This paper examines the use of augmented Generative Artificial Intelligence (GenAI) data mining tools, such as SAP Analytics Cloud, and compares their educational impact to traditional programming tools such as R. Through surveys, performance analysis over two terms, and instructor reflections, we explored how these tools affect learning outcomes and teaching practices differently in an undergraduate business school data mining course. We present a brief overview of GenAI and traditional programming tools, considering their respective strengths and limitations in educational contexts. We particularly examine how each type of tool influences the learning process, technical skill development, and students' ability to apply data mining concepts. This study offers insights into the effectiveness of augmented GenAI and traditional teaching tools and presents pedagogical implications for educators seeking to optimize data analytical skills education in business schools. Unlike prior studies that examine programming tools or investigate GenAI in educational contexts separately, our work presents a direct contrast between the two approaches, allowing for a systematic comparison of their impacts on teaching and learning.

Keywords: Business Analytics Education, Data Literacy, Generative AI in Education, Student Learning Outcomes, Traditional Programming Tools, Tool Integration in Curriculum.

Recommended Citation: Zhou, W., Ozturk, P., Hartzel, K.S., Salazar, C., (2026). Comparing Augmented Generative AI Platforms and Programming-Based Tools for Teaching Data Analytics to Undergraduate Business Students. *Information Systems Education Journal*, v24(n6) pp 7-20. DOI# <https://doi.org/10.62273/LAZQ5486>

Comparing Augmented Generative AI Platforms and Programming-Based Tools for Teaching Data Analytics to Undergraduate Business Students

Wenqi Zhou, Pinar Ozturk, Kathleen S. Hartzel and Carlos Salazar-Betancourth

1. INTRODUCTION

As firms rely on data for competitive advantage, efficiency, and strategy, demand for skilled data analysts is surging. Predictive modeling has become a core tool, used to forecast market trends, improve customer retention, and streamline operations (Layne, 2023).

Despite these advancements, significant gaps in data literacy remain. A recent report highlights that only 21% of the global workforce feels fully confident in their data literacy skills, and companies lose an average of five working days per employee annually due to data-related stress and inefficiencies (Simonet, 2020). A shortage of data literacy is a major barrier to realizing analytics' value, highlighting a gap between tech capabilities and workforce skills. Business education must treat data literacy as core and shift from theory-heavy curricula to hands-on training with practical tools and methods.

Business schools must balance technical depth with accessibility, integrating analytics with core business concepts. Because curricula blend analytical skills with practical application, tools that combine ease of use with rigor are best for diverse learners. As tools become more accessible, teaching must also emphasize interpreting and acting on insights. This paper examines how traditional programming and augmented Generative Artificial Intelligence GenAI tools shape business analytics teaching and learning.

Traditional Programming Tools and Augmented GenAI.

Traditional programming languages, particularly Python and R, remain foundational in both academia and industry for data analysis and predictive modeling (Johnson et al., 2020). Python's extensive libraries and readable syntax make it a popular teaching tool, while R's statistical foundations and packages support rigorous methods and visualization (Tucker et al., 2023). Interfaces such as RStudio, R GUI, and Jupyter Notebooks reduce the steep learning curve, making coding more accessible for students without technical backgrounds

(Çetinkaya & Rundel, 2018; Software Sustainability Institute, 2021). Still, many business students struggle with syntax and debugging before they can meaningfully engage in analysis, creating a trade-off between technical training and business-oriented analytical thinking (Harnowo, 2022; Wymbbs, 2016).

Recent advances in GenAI, mainly those foundational large language models, introduce new possibilities for lowering these barriers. Augmented GenAI tools leverage large language models to generate insights, automate analysis, and support natural language interaction (Bommasani et al., 2021; Feuerriegel et al., 2024). Platforms such as SAP Analytics Cloud (with Smart Discovery) and Power BI Copilot exemplify this shift, offering automated visualizations, pattern detection, and conversational interfaces (Microsoft, 2025; SAP PRESS, 2021). For example, once specifying the data and key measures vs. master data, Smart Discovery can automatically generate dashboards, visualize charts, and provide insights in text. By automating routine tasks, augmented GenAI tools may reduce cognitive load and free students to focus on higher-order skills like critical evaluation and problem-solving, fostering "hybrid intelligence as the integration of humans and AI, leveraging the unique Strengths of both". (Feuerriegel et al., 2024)

Augmented GenAI tools lower entry barriers, letting students move quickly to interpretation and application, but overreliance can promote superficial learning and limit technical depth. Traditional programming in R/Python is harder (syntax, debugging) yet builds analytical rigor and procedural problem-solving. Each cultivates different aspects of critical thinking. Augmented GenAI emphasizes evaluative/strategic reasoning, while programming emphasizes technical/procedural skills. Therefore, educators should balance these approaches to develop both conceptual insight and technical fluency.

Research Questions and Objectives

Augmented GenAI and traditional programming tools present trade-offs: Augmented GenAI offers accessibility and quick answers to business

questions, while programming provides deeper technical understanding and flexibility. Our study asks how these approaches differently affect student learning outcomes, instructional efficiency, and perceptions in undergraduate business data mining, and what mix best integrates them in analytics education.

Our investigation focuses on comparing SAP Analytics Cloud (with GenAI features) and R GUI (traditional programming tool) in an undergraduate business data mining course over two semesters. Through this comparison, we aim to:

1. Compare the effectiveness of SAP Analytics Cloud and R GUI in supporting student mastery of regression-based data mining concepts, using lab assignments and exam performance as benchmarks.
2. Assess the impact of tool selection on student learning outcomes and instructional efficiency, drawing on survey responses, performance scores, and instructor reflections across the two semesters.
3. Evaluate the pedagogical implications of exposing students to both augmented GenAI and traditional programming tools, based on student perceptions from surveys and comparative analysis of instructional time and learning results.

Through surveys, performance assessments, and instructor observations, we sought to understand not just which tool students preferred but how different tools shaped their learning experience and analytical capabilities. This paper presents our findings and their implications for business analytics education.

The remainder of this paper is organized as follows. Section 2 reviews prior research on programming tools in business education and emerging studies on GenAI in learning contexts, identifying the methodological gap our work addresses. Section 3 outlines the course design, instructional interventions, and data collection methods used to compare SAP Analytics Cloud and R GUI across two semesters. Section 4 presents survey results, performance analyses, and instructor reflections to assess student learning outcomes and tool effectiveness. Section 5 discusses the pedagogical implications of our findings, highlighting trade-offs between accessibility and technical depth. Finally, Section 6 concludes with key contributions, limitations, and directions for future research.

2. LITERATURE REVIEW

State of the art

While prior studies have examined programming tool use in business education (e.g., Python, R, or GUI-based environments), (Doyle et al., 2022; Porubän et al., 2024) and emerging research has explored GenAI in general learning contexts, (Alparslan, 2025; Chan & Hu, 2023; Fulara, 2024; Joshi, 2025a, 2025b; Kok Cha & Daud, 2025; Lin & Chen, 2024; Tu et al., 2024; Vieriu & Petrea, 2025) we found no studies directly comparing augmented GenAI analytics platforms with traditional programming tools in undergraduate business data mining courses. This methodological gap is important because it leaves unanswered questions about how different tools influence both short-term learning outcomes and long-term skill development. By comparing SAP Analytics Cloud (with GenAI features) and R GUI in parallel classroom settings, our study provides one of the first empirical tests of this contrast.

The Rise of Data Analytics in Business and Education

Teaching data analysis for decision-making has become a curricular imperative for business schools. It has become increasingly established that business decisions should be grounded in data-driven insights rather than intuition. A growing body of empirical research demonstrates that firms that adopt data-driven decision-making (DDD) practices consistently outperform their peers in terms of productivity, asset utilization, and overall profitability. Brynjolfsson et al. (2011) in a seminal large-scale study of 179 publicly traded firms, found that DDD adoption was associated with 5–6% higher output and productivity, independent of IT spending and organizational structure. Comparable results in emerging economies, such as Gul et al. (2023) findings in Pakistan's banking sector underscore that the strategic benefits of DDD span diverse contexts.

These findings have prompted reforms across business curricula. Organizations expect business graduates to enter the workforce with fluency in data interpretation and modeling. Yet, achieving this in business programs, where students often lack a computer science background, requires pedagogical strategies that reach a balance between analytical rigor and accessibility. The question is no longer whether to teach analytics, but how best to teach it, especially to non-technical learners.

Teaching Data Mining with Traditional Programming Tools in Business Education

Business undergraduates typically enter programming-heavy courses with a limited technical background, which can hinder learning when cognitive overload eclipses conceptual understanding. Doyle et al. (2022) details Temple University's experience introducing JavaScript to over 1,400 business students annually, noting that a significant portion of instructional time is spent introducing students to basic programming syntax and logic. This experience aligns with Luo & Adelopo (2024), who advocate for a student-focused pedagogy that brings together problem-solving and collaborative activities to make programming easier to learn. Still, challenges such as data wrangling and debugging persist.

Porubän et al. (2024) propose task-driven case studies as a solution, enabling students to engage with real-world problems while avoiding poor programming practices. Despite these innovations, programming environments often impose a steep learning curve that risks alienating non-technical students or reducing the time available to explore high-level analytics concepts like clustering or regression. GUI-based environments such as RStudio have helped mitigate this issue, but the broader question remains: How can instructors ensure that students meaningfully learn to analyze data without becoming overwhelmed by code?

Augmented GenAI Tools for Teaching Data Mining in Business Programs

Unlike traditional programming environments, augmented GenAI tools reduce technical friction, allowing students to interact with data using natural language and generate visualizations or models with minimal manual coding. Tu et al. (2024) argues that LLMs are redefining the role of data scientists by shifting focus from syntax to strategic oversight. This evolution aligns with the shift in educational needs, where students must learn to manage AI-augmented tools rather than master every logical and technical detail.

Studies by Chan & Hu, (2023), Grájeda et al., (2023), Kovari, (2025), Lin & Chen, (2024) Ali et al. (2024) demonstrate how AI tools can enhance student engagement and comprehension. Effective results hinge on infrastructure, teacher training, and pedagogy, not just the tool. Ali's review underscores customization and ethics as core to responsible AI use. By generating insights from prompts and automating modeling, GenAI lowers technical barriers while helping students build the interpretive and critical thinking needed for decision-making. Córdova et al. (2024) and

Lin & Chen (2024) offer direct evidence from higher education settings, GenAI tools like ChatGPT can enhance creativity, reduce stress, and support autonomy, especially among learners with lower technical confidence. However, they also caution that overreliance on automation may limit deep understanding and critical thinking, concerns that should be included in curriculum design. These findings are complemented by Vieriu & Petrea (2025), who highlight both improved academic outcomes and ethical concerns, such as privacy and academic integrity. Together, these studies justify the use of GenAI as a complement, not a substitute, for analytical reasoning in business programs.

Tool Selection and Integration Strategies in Analytics Education

The decision to use augmented GenAI or traditional tools in analytics education involves more than technical considerations; it reflects deeper pedagogical priorities. Research on tool integration suggests that blended models can yield the best results, especially when instructors intentionally build pedagogical structures that facilitate learning experiences. Kovari (2025) highlights the value of AI-driven personalization and collaborative learning in higher education, noting that multimodal learning environments supported by AI can enhance engagement and motivation.

From a curriculum design standpoint, frameworks such as Cognitive Load Theory and TPACK (Technological Pedagogical Content Knowledge) offer useful lenses. Gkintoni et al. (2025) synthesizes neuroscience and AI literature to argue that AI-powered systems, especially those that adapt based on cognitive state, can reduce mental load and optimize instructional pacing.

While business education has not yet widely adopted neuroadaptive systems, the principle remains relevant: tools that reduce unnecessary complexity allow students to focus on high-value thinking. According to Bloom's Taxonomy, higher-order cognitive skills involve analysis, evaluation, and creation. Augmented GenAI, when integrated thoughtfully, may support these skills by enabling students to focus less on routine tasks and more on critical assessment and innovative problem-solving. However, overreliance on automation risks limiting opportunities for students to practice these higher-order skills.

Aspect	GUI-based Tools (e.g., Tableau, SAP SAC)	Code-based Tools (e.g., R, Python)
Ease of Use	User-friendly, drag-and-drop interfaces	Steeper learning curve, requires coding knowledge
Flexibility	Limited to built-in features and options	Highly flexible, customizable analyses and models
Learning Focus	Focus on interpretation and visualization	Focus on programming, logic, and statistical rigor
Reproducibility	More difficult to fully reproduce workflows	Easily reproducible with scripts and version control
Scalability	Best for small-to-medium datasets	Handles large datasets, advanced models, automation

Table 1. Comparison of Augmented GenAI and Traditional Tools.

Finally, insights from early education offer a broader pedagogical justification. Kok Cha & Daud (2025) find that AI tools boost early-childhood learning by reducing cognitive load and personalizing support. The pedagogical lesson generalizes adaptive, scaffolded environments help learners of any age when first tackling abstract or technical content.

Existing studies have not directly compared augmented GenAI tools like ChatGPT with R in business schools largely because R is a long-established programming language in analytics curricula, while GenAI has only recently emerged, and scholarship has not yet caught up. Research in business education has historically emphasized traditional tools such as R, Python, and Excel as core competencies (Aasheim et al., 2015; Johnson et al., 2020; Tucker et al., 2023) , whereas GenAI studies have mainly focused on student perceptions, creativity, and emotions in broader higher-education contexts (Chan & Hu, 2023; Córdova et al., 2024; Lin & Chen, 2024) . Methodological hurdles — such as the need to redesign curricula, address academic integrity concerns, and overcome disciplinary silos — further complicate systematic comparisons (Doyle et al., 2022; Luo & Adelo, 2024). In addition, GenAI has often been framed as an

assistive rather than foundational skill (Feuerriegel et al., 2024; Tu et al., 2024), making head-to-head contrasts with R less natural in prior research designs. Combined with the lag of peer-review cycles, these factors explain why direct R vs. augmented GenAI comparisons in higher education have not yet appeared, leaving this gap open for timely investigation.

Business education should blend augmented GenAI tools (for accessibility and engagement) with traditional programming in R/Python (for technical depth). Adoption must align with teaching goals, faculty readiness, and ethical standards, and future research should evaluate effects on learning outcomes, confidence, engagement, and long-term skill development. Table 1 compares the tools examined in this study.

3.METHODOLOGY

Course Structure and Design

Business Intelligence is an undergraduate elective course that covers both data warehousing and mining components offered at a private university in the northeast of U.S. The course is organized into sequential modules, introducing students to increasingly sophisticated

Component	Fall 2022	Spring 2024
Course Content	Regression model, analysis through code-based analysis with GUI	Regression model, analysis through code-based analysis with GUI and augmented GenAI tools
Lecture Materials	Regression slides	Same regression slides
Lab Activities	Wine price analysis using R GUI	Wine price analysis using both SAP SAC and R GUI
Lab Support Materials	<ul style="list-style-type: none"> Background slides R tutorial slides Step-by-step R GUI instructions 	<ul style="list-style-type: none"> Same background slides SAC tutorial materials Step-by-step SAC instructions R tutorial slides Step-by-step R GUI instructions
Assessment by lab assignment and final exam	Questions focusing on analysis and interpretation	Same questions focusing on analysis and interpretation

Table 2. Implementation in Fall 2022 and Spring 2024.

data analysis techniques. This study focuses specifically on the 3-week data mining module, which introduces students to predictive modeling through regression analysis. The module was designed to progressively build students' understanding from theoretical foundations to practical implementation, with learning objectives focused on understanding fundamental regression concepts, mastering predictive model construction, developing model evaluation skills, and applying these techniques to business problems. Its detailed learning objectives are available upon request. The module follows a structured three-phase approach. The first phase establishes conceptual foundations through lectures covering regression concepts, metrics, and business applications. The second phase provides guided practice through a hands-on lab, allowing students to apply theoretical concepts in a controlled environment. The final phase consists of an assignment, where students demonstrate their acquired skills in a new context.

In Fall 2022, the module exclusively used R GUI in the lab activities. The lab structure included background slides introducing the lab, tutorial slides with R-specific instructions, and step-by-step guided practice using the R.GUI. After the lab demonstration, students completed their lab assignments with questions designed to test both technical implementation and conceptual understanding.

The Spring 2024 implementation integrated both SAP Analytics Cloud (SAC) and R GUI into the curriculum, which was systematically developed. The instructor completed formal training in SAP and obtained a badge in analytics and revised the curriculum in response to feedback from the institution's Center for Teaching Excellence.

The SAP badge assures us that the instructor has an in-depth understanding of SAC. The lab retained the same analytical content but was enhanced with tutorials and step-by-step guides for both tools. New SAC-specific materials supported its integration. The assessment questions remain consistent to ensure comparability. A set of the same data mining questions was also administered in the final exam in both semesters. The Spring 2024 dual-tool approach aimed at exposing students to both traditional programming and augmented GenAI tools, enabling a direct comparison of learning outcomes and expanding students' analytical skills. Table 2 details the differences between semesters, outlining lecture content, lab structure, and exam requirements, forming the

basis for evaluating the impact of tool choice on student learning.

Student populations of Fall 2022 and Spring 2024 are comparable. In Fall 2022, the class had 24 students. 11.7% are SCM, 11.7% are marketing, 11.7% are finance, and the rest are IST students. 41.2% have GPA between 3 and 3.5, and 53.3% have GPA above 3.5. In Spring 2024, the class has 24 students. 20% are SCM, 26.7% are marketing, 6.7% are management, 6.7% are finance, and the rest are IST students. 40% have GPA between 3 and 3.5, and 40% have GPA above 3.5.

Data Collection

Data was collected through surveys and performance scores on lab activities and final exam questions.

Surveys

Two structured surveys were administered as required course activities through Canvas to assess students' experiences with the analytical tools in Spring 2024. Both surveys required course activities with non-anonymous responses to ensure completion and enable connection with performance data. The first survey (GenAI Survey, 0.2% of total course score) was conducted immediately after students completed their first SAC module, while the second survey (R and SAC Comparison Survey, 0.15% of total course score) was administered after the data analysis module. Each survey received 100% response rate.

The GenAI Survey focused on students' experience with SAC and its GenAI features. The survey included ten questions designed to assess three key areas: prior experience with analytics tools, user experience with SAC's Smart Discovery feature, and attitudes toward augmented GenAI tools in data mining. Questions explored students' confidence in creating business intelligence reports, perceived difficulties in implementation, and their evolving attitudes toward augmented GenAI tools in data analysis.

The R and SAC comparison survey, administered after students had experience with both tools, contained eight questions examining tool preferences and learning outcomes. The survey investigated students' perceptions of how each tool contributed to their understanding of predictive modeling, their preferences between the tools, and specific likes and dislikes about each platform. Table 3 summarizes the key areas

Survey	Timing	Focus Areas	Key Questions
GenAI Survey	After completing SAC introduction and lab	<ul style="list-style-type: none"> • Prior tool experience • Smart Discovery implementation • GenAI attitudes 	<ul style="list-style-type: none"> • Tool familiarity • Implementation difficulty • Confidence in BI reporting • Future GenAI adoption
R and SAC Comparison	End of Data Mining module	<ul style="list-style-type: none"> • Tool comparison • Learning outcomes • Background information 	<ul style="list-style-type: none"> • Tool preferences • Learning effectiveness • Tool-specific feedback • Programming experience

Table 3. Survey Details.

of investigation in each survey. Details of the survey questions can be found in the appendix.

Lab Activities

The hands-on lab reinforced theory through practical predictive modeling using a real-world dataset with variables like temperature, rainfall, wine age, and population, aiming to predict wine prices. Students began with data exploration and single-variable regressions, then built multiple regression models to improve accuracy. They evaluated models using metrics like R-squared and RMSE, alongside GenAI insights, and made final predictions. In Fall 2022, students used R GUI exclusively, supported by background materials, tutorials, and guided practice in interpreting results and making data-driven recommendations.

The R and SAC comparison survey, administered after students had experience with both tools, contained eight questions examining tool preferences and learning outcomes. The survey investigated students' perceptions of how each tool contributed to their understanding of predictive modeling, their preferences between the tools, and specific likes and dislikes about each platform. Table 3 summarizes the key areas of investigation in each survey. Details of the survey questions can be found in the appendix.

Lab Activities

The hands-on lab reinforced theory through practical predictive modeling using a real-world dataset with variables like temperature, rainfall, wine age, and population, aiming to predict wine prices. Students began with data exploration and single-variable regressions, then built multiple regression models to improve accuracy. They evaluated models using metrics like R-squared and RMSE, alongside GenAI insights, and made final predictions. In Fall 2022, students used R GUI exclusively, supported by background materials, tutorials, and guided practice in interpreting results and making data-driven recommendations.

In Spring 2024, the lab used both SAC and R GUI, allowing students to compare traditional and augmented GenAI methods through the same wine price analysis. SAC tutorials emphasized Smart Discovery and automated features. Students were assessed on implementing regression models, interpreting results, making evidence-based predictions, and comparing approaches. This setup enabled direct comparison of learning outcomes while keeping core objectives consistent. The detailed activity description is available upon request.

4.RESULTS

Survey Results

The GenAI survey revealed evolving attitudes toward AI-augmented analytics tools. Students reported significantly increased confidence in their ability to generate comprehensive business intelligence reports when using GenAI features. They noted reduced time requirements for comprehensive analyses and improved their understanding of the business implications of their findings. Students struggled with GenAI's prompt design and validating automated insights, highlighting the need to balance automation, with critical thinking. Even so, surveys showed GenAI lowered technical barriers and was viewed positively, with students valuing a blend of GenAI and traditional methods and seeing these combined skills as important for their careers

Using both tools led to deeper learning: SAC was preferred for its intuitive, visualization-rich, business-oriented interface, while R GUI offered greater control and analytical depth but came with setup hassles, package issues, and a steep learning curve. Many students had prior Python experience but little exposure to SAC/BI tools. Overall, students favored augmented GenAI tools for ease and speed yet recognized programming's value for building strong analytical skills, suggesting a blended approach best supports understanding and practical application in data mining.

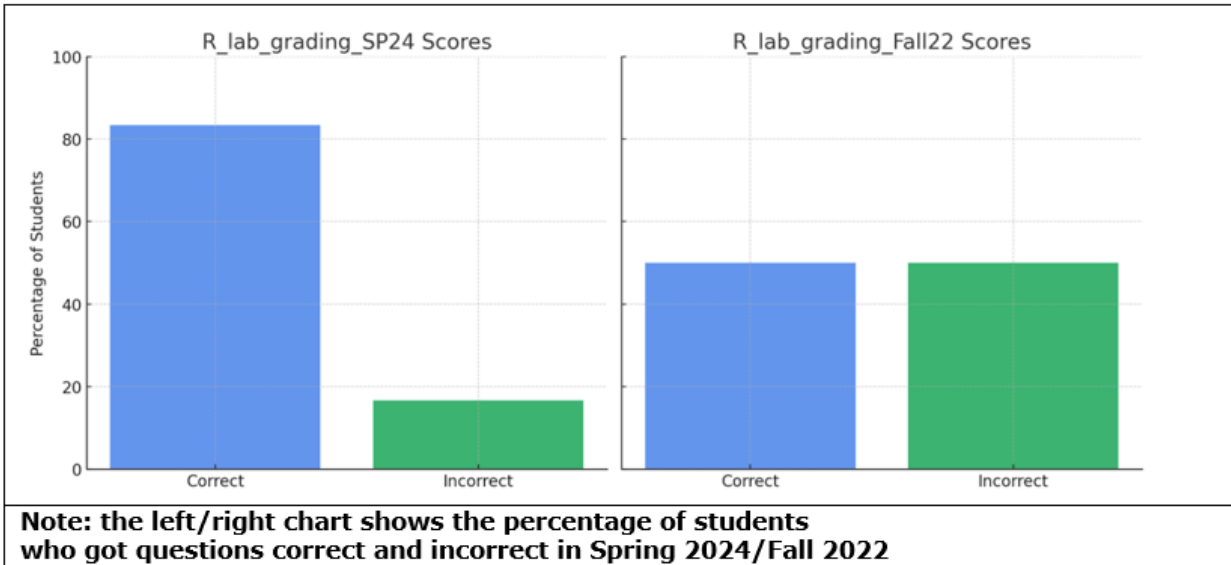


Figure 1. Comparison of student performance in R lab activities.

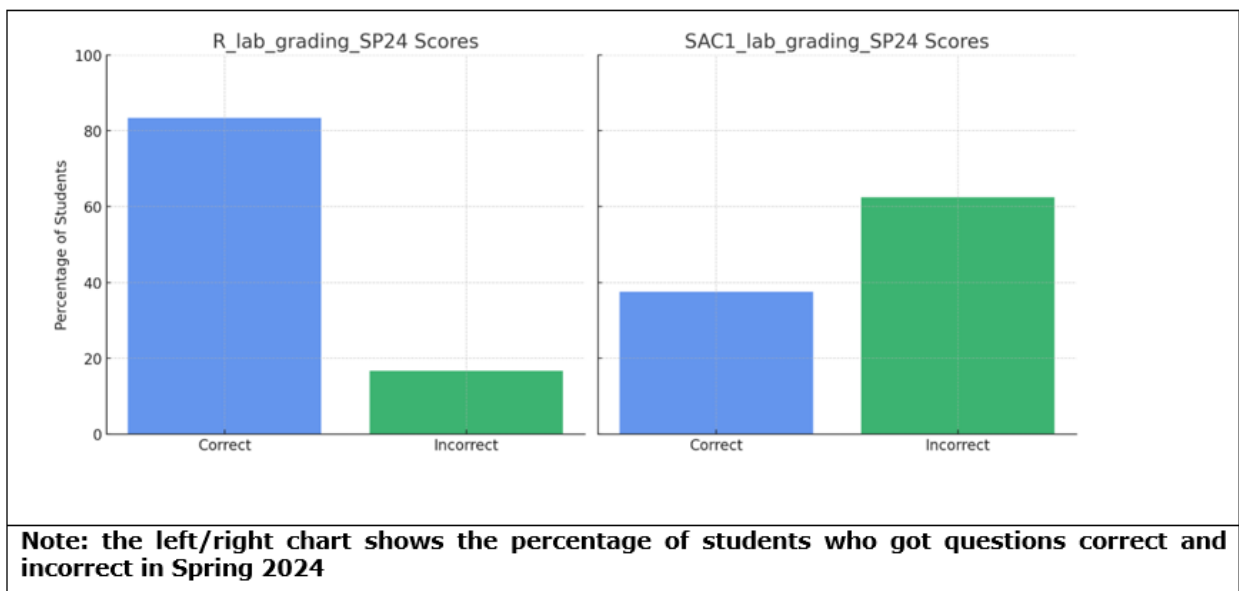


Figure 2. Comparison of student performance in R and SAC lab activities.

Performance Analysis

Performance analysis was conducted using two assessments. The first compared student lab activity outcomes across semesters (Spring 2024 post-implementation vs. Fall 2022 pre-implementation) and within Spring 2024 between R and SAC labs. The second compared student performances on an identical set of final exam questions. To ensure objective and consistent grading, one of the authors, who did not teach this course yet has expertise in data mining and GenAI, also evaluated all the lab activities and

final exam answers in both semesters. Discussions were held until a consensus was reached.

Specifically, in terms of R lab activities, students in Spring 2024 demonstrated significantly higher average performance as compared to Fall 2022 ($t=2.56$, $p=0.014$), with less variation in scores (1.9 in Spring 2024 compared to 2.6 in Fall 2022, on a 0-5 scale) despite receiving the same amount of lecture time and lab activity duration. (See figure 1). We note that the student population of these two cohorts is comparable, as

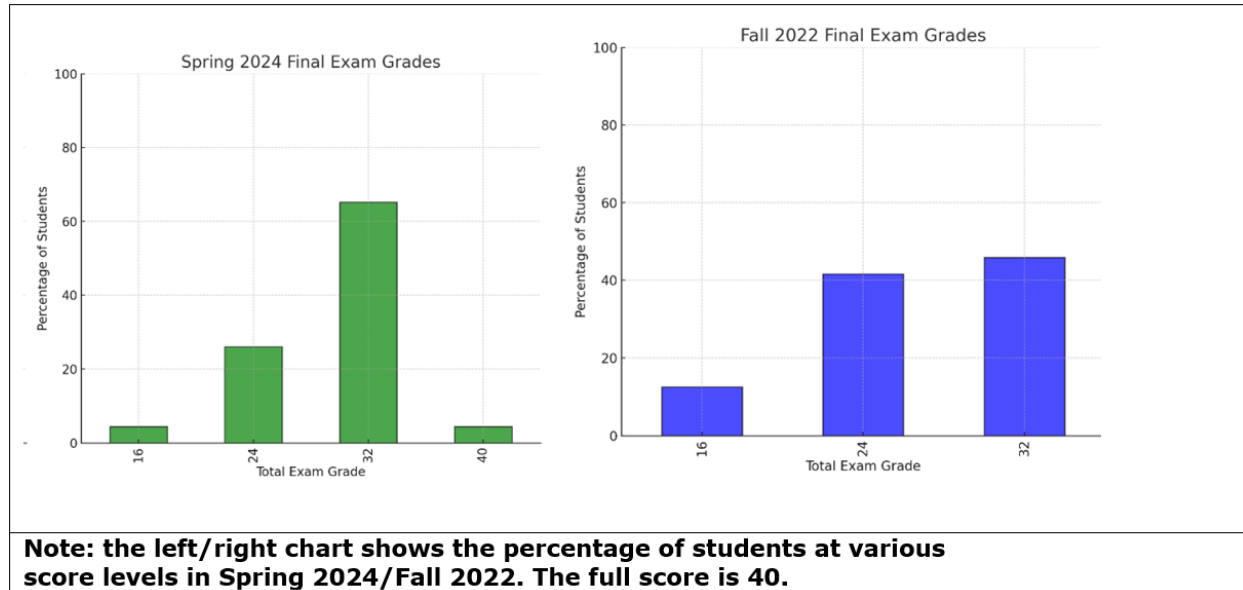


Figure 3. Comparison of data mining questions in final exams.

reported beforehand, and the assessment questions of all lab activities across semesters are the same.

The observed improvement in student learning is likely attributed to the integration of SAC in Spring 2024, which facilitated the instruction and delivery of data mining concepts and skills.

We further analyzed student performance between the student R lab and the SAC lab in Spring 2024. Not surprisingly, given that SAC was presented first, students performed better in the R lab than SAC ($t=4.41$, $p=0.0002$) as reported in Figure 2. This is consistent with the conventional wisdom that repetitive exposure can reinforce learning.

The current research found that reducing technical barriers through augmented GenAI tools can enhance student learning outcomes, particularly among non-technical business students. The observed improvement in Spring 2024 performance on R lab activities mirrors the literature's assertion that GenAI platforms can reduce cognitive load and increase accessibility (Lin & Chen, 2024; Tu et al., 2024). This alignment supports the claim by Grájeda et al. (2023) and Vieriu & Petrea (2025) that augmented GenAI tools promote autonomy and reduce stress, leading to greater productivity, consistency, and performance in academic tasks. Moreover, these findings emphasize the importance of balancing technical skill acquisition with conceptual understanding. Thoughtful SAC integration in Spring 2024 let students focus on

analysis rather than coding, yielding performance gains. These results empirically support the literature's call to strategically use augmented GenAI tools to improve learning and serve diverse learners in business analytics programs.

The results from our second assessment echo the above findings. Analysis of the final exam questions (See Figure 3) indicates that Spring 2024 students achieved a higher average score than their Fall 2022 counterparts ($p=0.07$). This marginal statistical significance indicates meaningful improvements in learning. Furthermore, the standard deviation of student performance was lower in Spring 2024 than in Fall 2022. This exciting finding is likely due to the introduction of SAC, an augmented GenAI data mining tool.

Instructor Reflections

The integration of SAC and R GUI into the curriculum provided a unique opportunity to observe how these tools shaped both the student learning experience and instructional practices. Each tool presented distinct strengths and challenges, which highlighted the trade-offs educators face when choosing tools.

SAC's no-code, Power BI/Tableau-like interface (e.g., Smart Discovery) made labs smoother and lower-stress, letting students focus on concepts and strategic analysis. By contrast, R GUI required coding steps (scripts, data wrangling, package issues), which proved challenging, especially for non-programmers, and demanded

extra troubleshooting even from experienced students.

Student needs for support varied: SAC's user-friendly design reduced one-on-one coaching and encouraged peer help, whereas R's technical complexity (coding, package installs) led to frequent instructor assistance, making R labs far more support-intensive in and out of class.

5. DISCUSSION

This pedagogical study deals with the issues of tool selection, technical depth, and real-world application. This includes questions of both tool depth versus breadth and the conceptual understanding of the business domain versus hands-on execution using tools. Decisions an instructor must make should be based on pedagogical priorities, which require balancing competing learning objectives. These types of choices are not limited to the context of data mining education but are widely generalizable across learning environments.

Using multiple tools creates a depth-versus-breadth trade-off, but the variety can better support diverse learners. Beyond tool choice, cultivating human-AI collaborative problem solving is becoming increasingly essential.

To be employable, students must develop the skills needed to be fluent in using GenAI. This fluency includes the ability to write meaningful GenAI prompts. Prompt writing should be emphasized throughout university curricula. Furthermore, the effectiveness, intelligence, and ease of use of AI augmentation are expected to increase at an exponential rate. Thus, the use of AI augmented tools, like SAC, is not just a means to learn new concepts, like data analysis, it is a skill students must master to participate in the emerging Workforce 5.0.

6. CONCLUSIONS, LIMITATIONS, AND FUTURE RESEARCH

This study provides valuable insights into the use of augmented GenAI tools and traditional programming platforms in business analytics education. It used a survey, performance assessments, and instructor reflection to specifically compare the effectiveness of using traditional programming tools (R-GUI) to augmented GenAI tools (SAC). It also raises the issue of the effectiveness of using multiple tools, in a serialized fashion, to reinforce concepts. Our results showed that the use of a traditional programming tool, R-GUI, offers greater control

and flexibility to students. However, R-GUI has a steep learning curve, and it is difficult to set up the programming environment. Those with a coding background were more capable of tackling R-GUI without issue. Thus, student backgrounds should be a key factor in tool choice.

The augmented GenAI tool, SAC, enhances the students' understanding of the analyses, increases their confidence in performing the analyses, and allows them to become productive more quickly. However, designing good prompts in SAC proves to be challenging. Overall, there is a clear student preference for using SAC. We found that the assessment of student performance when using R-GUI, in a lab assignment (Fall 2022), improves when it is preceded by a SAC lab assignment (Spring 2024). We also observe that students' overall understanding of data mining and analysis concepts increases when both tools are used.

However, several limitations should be considered when interpreting the findings. First, the results are generated in the context of an undergraduate business school and need to be generalized to other disciplines or educational levels with caution. While our sample size of 24 students per semester provides valuable initial insights, the analysis has limited statistical power. Additionally, the focus on SAC and R GUI represents a subset of the available tools for teaching data analytics. Other tools with distinct features, such as Python with Jupyter Notebooks, could yield different insights and learning experiences. Second, this study compares the learning outcomes of two semesters, which may not fully capture the long-term impact of tool use. Additionally, the asymmetric measurement structure of our survey instruments, where one survey focused exclusively on SAC while the other compared SAC and R, constrains the validity of a direct, one-to-one comparison between the tools. The reliance on self-reported survey data also introduces potential biases, as students may overestimate their confidence or minimize challenges due to social desirability bias. Furthermore, the performance metrics used in this study might not comprehensively measure students' conceptual understanding or practical proficiency and may be influenced by potential grader bias. Future studies should also investigate the relative importance of augmented GenAI tools and repeated practice.

Future work should use longitudinal studies to assess GenAI tools' long-term effects on learning and career outcomes, and broader comparisons to evaluate more analytics platforms. It should

also reflect the 2025 shift to augmented AI collaborative, dialogic systems that work alongside humans, ask clarifying questions, deepen human reasoning, and improve via continuous feedback.

7. REFERENCES

- Aasheim, C. L., Williams, S., Rutner, P., & Gardiner, A. (2015). Data Analytics vs. Data Science: A Study of Similarities and Differences in Undergraduate Programs Based on Course Descriptions. *Journal of Information Systems Education*, 26(2), 103–116.
<https://jise.org/Volume26/n2/JISEv26n2p103.html>
- Ali, O., Murray, P. A., Momin, M., Dwivedi, Y. K., & Malik, T. (2024). The effects of artificial intelligence applications in educational settings: Challenges and strategies. *Technological Forecasting and Social Change*, 199.
<https://doi.org/10.1016/j.techfore.2023.123076>
- Alparslan, A. (2025). The Role of Accuracy and Validation Effectiveness in Conversational Business Analytics. *IEEE Access*.
<https://doi.org/10.1109/ACCESS.2025.3540975>
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., & Liang, P. (2021). On the Opportunities and Risks of Foundation Models. *ArXiv Preprint ArXiv:2108.07258*.
<https://doi.org/10.48550/arXiv.2108.07258>
- Brynjolfsson, E., Hitt, L., & Kim, H. H. (2011). Strength in Numbers: How does data-driven decision-making affect firm performance?
<https://doi.org/10.2139/ssrn.1819486>
- Çetinkaya Rundel, M., & Rundel, C. W. (2018). Infrastructure and Tools for Teaching Computing Throughout the Statistical Curriculum. *The American Statistician*, 72(1), 58–65.
<https://doi.org/10.1080/00031305.2017.1397549>
- Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20(1).
<https://doi.org/10.1186/s41239-023-00411-8>
- Córdova, P., Grájeda, A., Córdova, J. P., Vargas-Sánchez, A., Burgos, J., & Sanjinés, A. (2024). Leveraging AI tools in finance education: exploring student perceptions, emotional reactions and educator experiences. *Cogent Education*, 11(1).
<https://doi.org/10.1080/2331186X.2024.2431885>
- Doyle, M., Lavin, A., & Sclarow, S. (2022). Association for Information Systems Association for Information Systems Teaching Programming to 1,400 Business Students per year Teaching Programming to 1,400 Business Students per year.
https://aisel.aisnet.org/treos_amcis2022/17
- Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66(1), 111–126. <https://doi.org/10.1007/s12599-023-00834-7>
- Fulara, A. S. (2024). Comparative Analysis of Artificial Intelligence (GenAI) in Business Intelligence Platforms. *International Journal of Computer Trends and Technology*, 72(4), 95–101.
<https://doi.org/10.14445/22312803/ijctt-v72i4p112>
- Gkintoni, E., Antonopoulou, H., Sortwell, A., & Halkiopoulos, C. (2025). Challenging Cognitive Load Theory: The Role of Educational Neuroscience and Artificial Intelligence in Redefining Learning Efficacy. In *Brain Sciences* (Vol. 15, Issue 2). Multidisciplinary Digital Publishing Institute (MDPI).
<https://doi.org/10.3390/brainsci15020203>
- Grájeda, A., Burgos, J., Córdova, P., & Sanjinés, A. (2023). Assessing student-perceived impact of using artificial intelligence tools: Construction of a synthetic index of application in higher education. *Cogent Education*, 11(1).
<https://doi.org/10.1080/2331186X.2023.2287917>
- Gul, R., Leong, K., Mubashar, A., Al-Faryan, M. A. S., & Sung, A. (2023). The Empirical Nexus between Data-Driven Decision-Making and Productivity: Evidence from Pakistan's Banking Sector. *Cogent Business and Management*, 10(1).

- <https://doi.org/10.1080/23311975.2023.2178290>
- Harnowo, A. S. (2022). Blending a MOOC course into a Business School's Course to Introduce Python for Data Analytics. *Business Education Innovation Journal*, 14(2), 31–36.
- Johnson, M. E., Albizri, A., & Jain, R. (2020). Exploratory analysis to identify concepts, skills, knowledge, and tools to educate business analytics practitioners. *Decision Sciences Journal of Innovative Education*, 18(1), 90–118. <https://doi.org/10.1111/dsji.12195>
- Joshi, S. (2025a). Review of Gen AI Models for Financial Risk Management. <https://doi.org/10.32628/CSEIT2511114>
- Joshi, S. (2025b). The Transformative Role of Agentic GenAI in Shaping Workforce Development and Education in the US. <https://ssrn.com/abstract=5133376>
- Kok Cha, W., & Daud, P. (2025). Enhancing Early Education with Artificial Intelligence: A Comparative Study of AI-Powered Learning Versus Traditional Methods. *International Journal of Academic Research in Business and Social Sciences*, 15(2). <https://doi.org/10.6007/IJARBS/v15-i2/24690>
- Kovari, A. (2025). A systematic review of AI-powered collaborative learning in higher education: Trends and outcomes from the last decade. In *Social Sciences and Humanities Open* (Vol. 11). Elsevier Ltd. <https://doi.org/10.1016/j.ssaho.2025.101335>
- Layne, R. (2023, May 31). With predictive analytics, companies can tap the ultimate opportunity: Customers' routines. Harvard Business School Working Knowledge. <https://www.library.hbs.edu/working-knowledge/with-predictive-analytics-companies-can-tap-the-ultimate-opportunity-customers-routines>
- Lin, H., & Chen, Q. (2024). Artificial intelligence (AI) -integrated educational applications and college students' creativity and academic emotions: students and teachers' perceptions and attitudes. *BMC Psychology*, 12(1). <https://doi.org/10.1186/s40359-024-01979-0>
- Luo, X., & Adelopo, I. (2024). Exploring pedagogies, opportunities and challenges of teaching and learning programming in business school. *Journal of International Education in Business*. <https://doi.org/10.1108/JIEB-05-2024-0060>
- Microsoft. (2025). Overview of Copilot for Power BI. <https://learn.microsoft.com/en-us/power-bi/create-reports/copilot-introduction>
- Porubän, J., Nosál', M., Sulír, M., & Chodarev, S. (2024). Teach Programming Using Task-Driven Case Studies: Pedagogical Approach, Guidelines, and Implementation. *Computers*, 13(9), 221. <https://doi.org/10.3390/computers13090221>
- SAP PRESS. (2021). An Overview of SAP Analytics Cloud Smart Assist and Smart Predict Services. <https://blog.sap-press.com/an-overview-of-sap-analytics-cloud-smart-assist-and-smart-predict-services>
- Simonet, K. (2020, June 9). Why data literacy is a key ingredient to success in the age of data and analytics: And how to unlock its value: The essential roles of measurement and data culture. Deloitte Belgium. <https://www.deloitte.com/be/en/services/consulting/analysis/data-literacy-is-key-to-success.html>
- Software Sustainability Institute. (2021). GUIs for Research Software: Why Are They Relevant? (Part One). <https://www.software.ac.uk/blog/guis-research-software-why-are-they-relevant-part-one>
- Tu, X., Zou, J., Su, W., & Zhang, L. (2024). What Should Data Science Education Do With Large Language Models? License: Creative Commons Attribution 4.0 International License (CC-BY 4.0). *Harvard Data Science Review* • Issue, 6(1), 2024. <https://doi.org/10.1162/99608f92.bff007ab>
- Tucker, M. C., Shaw, S. T., Son, J. Y., & Stigler, J. W. (2023). Teaching statistics and data analysis with R. *Journal of Statistics and Data Science Education*, 31(1), 18–32. <https://doi.org/10.1080/26939169.2022.2089410>
- Vieriu, A. M., & Petrea, G. (2025). The Impact of Artificial Intelligence (AI) on Students'

Academic Development. *Education Sciences*,
15(3).
<https://doi.org/10.3390/educsci15030343>

Wymbs, C. (2016). Managing the innovation
process: Infusing data analytics into the

undergraduate business curriculum. *Journal
of Information Systems Education*, 27(1),
61–74.
[https://jise.org/Volume27/n1/JISEv27n1p61
.html](https://jise.org/Volume27/n1/JISEv27n1p61.html).

Appendix 1: Survey Questions

GenAI Survey

Have you used SAP Analytics Cloud prior to this course?

Have you used any other tools that can auto generate several pages of visualization reports including key influencers, simulation, and descriptive charts, with a few clicks?

If this SAC was not provided, how long do you think it will take you to complete a similar BI report that can answer all the questions? Reminder: those questions are related to Key Influences, Simulations, and descriptive summary.

How difficult is it to implement the Smart Discovery in this lab?

How difficult is it to use the story created by Smart Discovery to answer data and business questions in the assignment?

To what extent does the GenAI functionality, i.e., Smart Discovery in SAC, help you feel confident that you are able to provide an insightful BI report that goes beyond just descriptive summary and charts, as compared to using tools without it?

You can choose multiple answers. What are the challenges in using Smart Discovery in creating meaningful and in-depth BI reports?

After this experience, has your attitude towards using GenAI tool in data mining changed?

Will you embrace GenAI tools for data mining in the future? Choose the best that describes your attitude.

Do you agree that GenAI tools help prepare people qualified for various data mining tasks with considerably less technical barriers and better efficiency?

R and SAC Comparison Survey

How do you think learning these two tools, SAC and R GUI, help you learn designing, applying, and comparing predictive models?

If we only had time to learn one tool, which one would you select? Why?

What did you like about SAC in this module?

What did you dislike about SAC in this module?

What did you like about R GUI in this module?

What did you dislike about R GUI in this module?

What is(are) your major(s), and track if any?

Have you taken any programming related courses? If yes, what is the course and the programming language used?

Appendix 2 (hyperlink)

Note: [SAC predictive scenarios](#) were used for regression model task.

Integrating Artificial Intelligence into Undergraduate Business Programs: An Analysis of AACSB-Accredited Schools

Kiku Jones
Kiku.Jones@qu.edu
Quinnipiac University
Hamden, CT 06518

Lori Leonard
lori-leonard@utulsa.edu
University of Tulsa
Tulsa, OK 74104

Wendy Ceccucci
Wendy.Ceccucci@qu.edu
Quinnipiac University
Hamden, CT 06518

Guido Lang
guido.lang@quinnipiac.edu
Quinnipiac University
Hamden, CT 06518

Abstract

Artificial intelligence (AI) is increasingly shaping how organizations operate, and business schools play an important role in preparing students for these changes. This study reviews undergraduate AI programs at AACSB-accredited business schools to understand their prevalence and thematic focus. The results indicate that undergraduate AI programs are far more common in computing and engineering schools than in business schools. Although there is some overlap in program themes, the emphasis differs by discipline. Computing and engineering programs tend to focus on technical skills, while business programs concentrate on value creation and managerial decision-making. Because both perspectives are necessary and the technical dimension is already well established, business schools need to expand their AI offerings to address the business-oriented aspects that are currently underserved.

Keywords: AI programs, AI skills, business education, undergraduate business curriculum, curriculum analysis

Recommended Citation: Jones, K., Leonard, L., Ceccucci, W., Lang, G., (2026). Integrating Artificial Intelligence into Undergraduate Business Programs: An Analysis of AACSB-Accredited Schools *Information Systems Education Journal*, v24(n6) pp 21-30. DOI# <https://doi.org/10.62273/NCJK8904>

Integrating Artificial Intelligence into Undergraduate Business Programs: An Analysis of AACSB-Accredited Schools

Kiku Jones, Lori Leonard, Wendy Ceccucci and Guido Lang

1. INTRODUCTION

Traditional Artificial Intelligence (AI) courses and programs (e.g., machine learning, robotics, etc.) have been offered for quite some time in engineering and computer science programs. However, in the last few years, specifically generative AI has become an essential part of people's lives (Gartner, 2025). It is impacting how we work every day, everywhere. Generative AI can write, design images, generate code, and create audio. It is being used in offices around the world, transforming how businesses operate. It is quickly becoming an essential technology. As a result, employers are looking to hire people who not only know how to use these tools, but who can think strategically about how to apply them. This represents a major shift in the skills that matter most in today's job market.

While engineering and computer science programs have been teaching AI (including offering some generative AI courses) as part of their curriculum, business schools have fallen behind (AACSB, 2025). Teaching AI has historically required teaching difficult mathematics, finding materials that make sense to non-tech students, and keeping up with a field that is constantly changing. Some business schools are offering a few AI-related electives at the graduate level. However, business schools need to realize that AI, specifically generative AI, isn't just a technical tool. It's changing how we do business, from marketing products to analyzing finances, managing supply chains, and more.

Nearly one in three American workers is already using generative AI on the job (Microsoft, 2025). Business students need to be prepared to utilize this technology. This paper examines where business schools stand compared to other schools and substantiates the need for the development of generative AI programs across the business curriculum to meet industry needs.

2. LITERATURE REVIEW

AI

Generative AI refers to artificial intelligence systems capable of creating new content across various modalities including text, images, audio,

video, code, and 3D models. These systems learn patterns from existing data and use this knowledge to generate novel outputs that weren't explicitly programmed (Bommasani, et al., 2022).

Generative AI models are primarily built on neural network architectures, with transformer-based models currently dominating the field. These systems function by learning statistical patterns from massive datasets during training. Then these learned patterns are used to predict and generate new content that exhibits similar characteristics. This provides the ability to produce outputs that can adapt to specific prompts or instructions (LeCun, et al., 2023)

The foundation of modern generative AI is large language models (LLMs) and neural network models, which provide generative AI model capabilities in content generation across different domains (Brown, et al., 2020; Ramesh, et al., 2022).

AI in Academia

The integration of AI in education presents significant challenges, particularly "when introducing this complex and very important area of technology to audiences beyond the students of computing and engineering disciplines" (Xu & Babaian, 2021, p. 1). While AI is commonly expected within computer science and engineering curricula, business programs have been slow to adopt it (Southworth, et al., 2023). Even introductory AI courses in technical programs are recognized as challenging to teach (Harris & Kiefer, 2004), largely due to AI's broad scope and rapidly evolving nature.

Business schools face an additional disadvantage, lacking the pedagogical resources available to engineering programs (Xu & Babaian, 2021). Most AI textbooks rely heavily on complex mathematical or computer algorithms and provide few real-world examples or case studies. An analysis of AI courses across the top 46 business schools (based on 2020 U.S. News rankings) revealed that machine learning forms the core of business school AI curriculum, which typically combines AI and data analytics topics (Chen, 2022). Chen also found that AI is primarily

introduced in graduate-level information systems courses. However, since AI impacts every aspect of business, it should be integrated across all disciplines, not just technical ones, and given its complexity, multiple courses should be offered within each discipline (Sollosy & McInerney, 2022; Southworth, et al., 2023).

A survey of U.S. higher education leaders indicates that while students are perceived to use AI tools extensively, faculty adoption lags. More than a third of these leaders believe their institutions are below average (28%) or behind (7%) compared to other universities in generative AI tool implementation (Watson & Rainie, 2025). Although generative AI is rapidly transforming education, its complexity creates significant challenges for educators. "AI education has yet to catch up with the challenge of introducing this complex and very important area of technology to audiences beyond the students of computing and engineering disciplines" (Xu & Babaian, 2021). However, AI provides value and is necessary in university education. "AI training equips you with the knowledge and skills to harness the technology's power, allowing you to innovate and solve complex problems traditional methods can't address" (Gibson, 2024). AI also adds value specifically in business schools. Students can enhance decision making and increase efficiency and innovation (Farrugia, 2023; Gibson, 2024).

AI in Industry

The impact of generative AI on the workforce is expected to be striking, transforming both routine and non-routine tasks (Poba-Nzaou, et al., 2021). Between 2010-2019, AI-related job postings showed significant growth in both absolute numbers and as a percentage of total job listings (Alekseeva, et al., 2021).

A recent survey of 4,278 U.S. workers revealed that 30.1% currently use generative AI in their workplace. Users of these AI tools tend to be younger, more educated, and earn higher incomes. Industry adoption varies considerably, with information services and management companies reporting 60% AI utilization, while real estate, construction, and education sectors hover around 40% (Hartley, et al., 2024).

Currently, approximately 12% of U.S. workers are in occupations where tasks are expected to be automated by generative AI (Fuller, et al., 2025). This automation is already reducing or eliminating many entry-level positions, particularly in software engineering (Fuller, et al., 2025).

However, in other fields, generative AI may create new opportunities as the landscape of employable skills evolves (Poba-Nzaou, et al., 2021).

An analysis of UK job postings from 2018 to 2024 found that employers increasingly value candidates' AI skills over their educational degrees (Gonzalez, et al., 2024). This trend strongly suggests that there is an increasing need for universities to adequately equip graduates with the AI competencies demanded by today's job market.

Skills Needed in Industry

Employers are increasingly seeking hybrid talent with a skill set that combines technical capabilities such as data manipulation and training data development alongside strategic competencies including advanced problem-solving and strategic decision-making.

According to research by Jeppson (2025), the AI revolution is driving substantial demand for professionals with expertise in data analytics, STEM disciplines, and information technology.

Organizations integrating AI are rapidly transitioning toward more educated and specialized workforces, with a particular emphasis on technical domains. Hazan, et al. (2024) further highlights this trend, noting a significant increase in job postings requiring specialized skills like robotics engineering, advanced data analysis, and complex system integration.

The AI professional landscape requires a sophisticated and multifaceted skill set that exceeds the traditional technological boundaries. Technical competencies form the foundational framework for success in this rapidly evolving field, with programming languages like Python and R serving as critical tools for innovation.

Data engineering knowledge has become increasingly important, reflecting the growing complexity of AI systems. Professionals must demonstrate advanced capabilities in data manipulation, analysis, and interpretation. This requires not only technical proficiency but also the ability to extract meaningful insights from intricate datasets using advanced visualization tools like Power BI and Tableau.

Cloud computing platforms including Azure and AWS have emerged as crucial infrastructure for AI development, necessitating strong familiarity

and strategic implementation skills. These platforms enable scalable, flexible technological solutions that are fundamental to modern AI research and application.

Beyond technical skills, employers are placing increased emphasis on complementary professional attributes. The World Economic Forum's (2025) *The Future of Jobs Report 2025* highlights the critical importance of soft skills such as strategic problem-solving, advanced critical thinking, professional communication, and effective collaboration. These capabilities enable AI professionals to translate complex technological concepts into actionable business strategies.

Educational pathways for AI professionals have become increasingly specialized. Typical credentials include bachelor's degrees in computer science, master's degrees in AI-related disciplines, and doctoral research focusing on advanced technological applications. Emerging specialized roles like Generative AI Engineer, Computer Vision Engineer, and Remote AI Training Specialist underscore the field's growing complexity and segmentation.

Researchers like Shen (2024) emphasize that success in the AI arena requires continuous learning and adaptive professional development. Professionals must consistently update their knowledge frameworks, cultivate innovative thinking, and maintain a dynamic approach to technological engagement. This involves not only mastering existing technologies but also developing the capacity to anticipate and integrate emerging computational paradigms.

The rapidly transforming labor market demands that AI professionals remain agile, proactively updating their skills to match evolving technological demands. This requires a holistic approach that balances deep technical expertise with broader strategic understanding, positioning professionals to navigate the complex and dynamic landscape of artificial intelligence.

3. METHODOLOGY

At the time of data collection, January to May 2024, there were 550 AACSB accredited US business schools. Data was collected from each of the universities that housed these schools to determine 1) if they had some type of AI program at the university, 2) which school/college was it located, and 3) how the business school compared to the other schools/colleges. AI programs/certificates were identified by use of

Artificial Intelligence in the name of the program. Ninety-nine (18%) universities had some type of AI program, either major, minor, certificate or concentration. Of these, there were fifty-five at the undergraduate level (56%).

Due to the differences in how universities name their schools, a classification needed to be created for the data collected to be mapped. Appendix A shows the classification of school and the list of specific school names that fell into each one (this only included the school area and not any specific "named" schools). This study utilized a similar process used in a prior study which surveyed university IS program curriculum described on the websites of the universities (Yang & Wen, 2017). Multiple authors went through the school areas and coded them based on the classification item that best described the school area. Where there were differences, these were discussed, and a decision was made. The school classifications were Arts & Sciences (A&S), Business, Computing & Engineering (C&E), Health, and Science & Technology (S&T).

Each of the AI programs identified the number of required courses needed to complete the program. Course names and descriptions were obtained for the required courses. The elective courses were not included because these courses would not be taken by every student in a program. These courses were coded by the school classification. Microsoft Co-Pilot identified common themes among all courses, as well as within schools and by program type. These themes were then reviewed by the authors in a similar process as the school classification. The themes were iteratively modified by the authors. The full list of themes can be found in Appendix B.

4. RESULTS

The size of the schools based on undergraduate enrollment and the program types that are offered are given in Table 1. Universities were classified as small if they had less than 10,000 undergraduate students. They were classified as medium sized if they had between 10,000 and less than or equal to 19,999. Universities with 20,000 or more undergraduate students were classified as large. Small size universities had the highest percentage of AI majors (45%), minors (72%), and concentrations in AI (54%). Large size universities had the greatest number of certificate programs (64%). A chi-square test of independence showed a significant association between university size and AI program type ($\chi^2(6) = 16.91, p < 0.01$).

Size	Major	Minor	Cert.	Conc.	Total
Large	36%	11%	64%	8%	28%
Medium	18%	17%	27%	38%	23%
Small	45%	72%	9%	54%	48%
Total	34%	28%	17%	20%	100%

Table 1. University Size and AI Program Offerings (Cert. = Certificate, Conc. = Concentration)

Sixty-six percent of the universities were public. Table 2 shows the breakdown of university type and AI program types. Public universities had the highest percentage of AI major programs (59%), while private universities had the highest percentage of AI minor programs (56%). AI certifications were only offered by public universities, and they offered the greatest number of concentrations in AI (77%). A chi-square test of independence showed a significant association between university type and AI program type ($\chi^2(3) = 8.52, p < 0.05$).

Type	Major	Minor	Cert.	Conc.	Total
Private	11%	14%	0%	5%	30%
Public	23%	14%	17%	15%	70%
Total	34%	28%	17%	20%	100%

Table 2. University Type and AI Program Offerings (Cert. = Certificate, Conc. = Concentration)

Figure 1 shows that most of the program types offered are majors (34%). This is followed by minors (28%). Finally, concentrations and certificate programs, with 20% and 17% respectively, follow the minor programs.

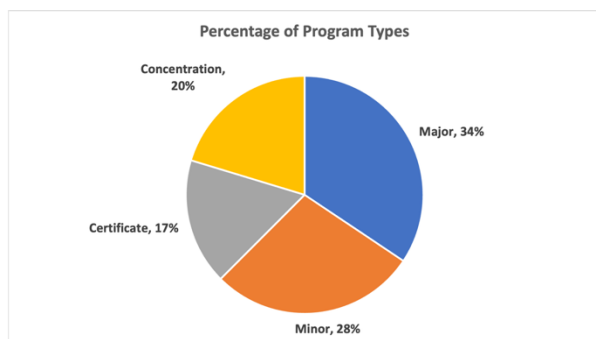


Figure 1. Percentage of Program Types

Figure 2 indicates the percentage of programs (regardless of type) that each school classification holds. C&E schools make up 47% of where the AI programs are housed. Following them is A&S

schools with 24%. S&T schools make up 15% of the programs. Business schools make up just 13% followed only by Health schools with 2%.

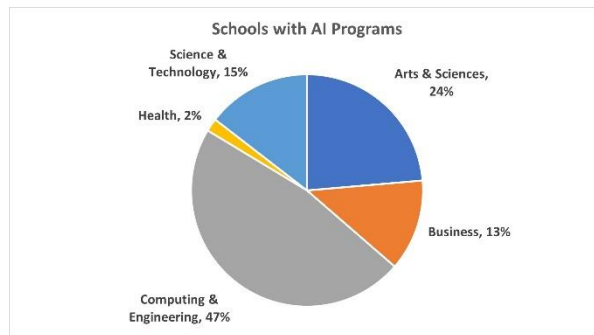


Figure 2. Schools with AI Programs

Table 3 provides information on the different AI program types within the schools. C&E schools not only offer the greatest number of AI programs overall, but they also have the highest percentage of each type of AI program. They have 45% of the major programs, 56% of the minor programs, 45% of the certificate programs, and 46% of the AI concentrations. A&S has the second highest percentage of the major programs (27%), certificate programs (27%), and concentrations (31%). Business has the second highest percentage of the minor programs (28%). Health schools only have certificate programs. However, a chi-square test of independence did not detect a statistically significant association between schools and AI program type ($\chi^2(12) = 20.02, p = .067$).

School	Major	Minor	Cert.	Conc.	Total
Arts & Sciences	27%	6%	27%	31%	22%
Business	14%	28%	9%	0%	14%
Computing & Engin.	45%	56%	45%	46%	48%
Health	0%	0%	18%	0%	3%
Science & Technology	14%	11%	0%	23%	13%
Total	100%	100%	100%	100%	100%

Table 3. Program Types within Schools (Cert. = Certificate, Conc. = Concentration)

Figure 3 provides unique and overlapping program themes from the Business school and the C&E school classifications. The C&E school themes were selected to compare with the Business school themes because the C&E classification makes up the highest percentage of AI programs. The themes that overlapped between the two areas were Algorithms & Data Structures, Cloud Computing, Deep Learning,

Ethics, Foundations of AI, Machine Learning, and Programming. Algorithms & Data Structures, Ethics, Foundations of AI, Machine Learning, and Programming are all fundamental themes in AI which would be expected to be seen in most programs. Deep Learning makes sense for both areas as this is important with processing large amounts of data as well as creating complex models. When working with AI solutions, leveraging Cloud Computing can be important for its scalability and efficiency.

Regarding the unique themes found within each classification, there appears to be clear differences in the AI application and focus. The Business themes are more focused on using AI to drive business value and competitive advantage, make data-driven decisions, and enhance business processes with specific applications. These are more application-oriented seeking to solve business problems. The C&E themes are more focused on the technical side of AI emphasizing the development and optimization of AI technologies such as computer vision and robotics.

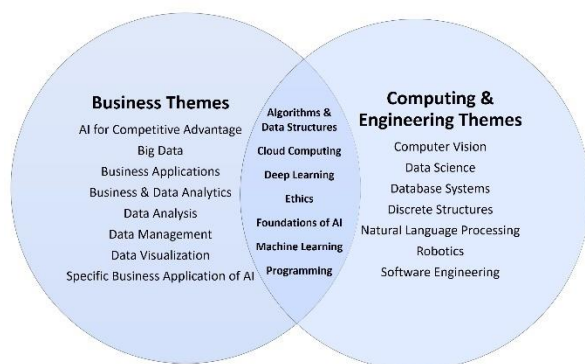


Figure 3. Business Themes vs. C&E Themes

5. LIMITATIONS

This research is limited by its focus on US universities with AACSB accredited business schools. Only universities that met that requirement were included in this assessment. In addition, only identifying programs which specifically have artificial intelligence in the program name may create a biased representation of current AI education in universities. Another limitation is data was collected at a point in time. It is possible universities that were assessed have made changes to their curriculum since the data collection.

6. CONCLUSION

Understanding what programs are offering AI (majors, minors, certificates, and concentrations) and what concepts/themes are being covered is vital to the future of generative AI in education. Five hundred and fifty US universities that have AACSB accredited business schools were assessed. Of those universities, 99 had some type of AI program, with 55 of those being at the undergraduate level. Thirty-four percent of the program types were majors in AI, with computing and engineering housing 47% of all AI programs. Business schools only offered 13%. Additionally, computing and engineering had the highest percentage of each type of program (major, minor, certificate, and concentration). Business only had the second highest in minors. It is clear from this analysis that business programs in AI are lagging.

AI programs offered in C&E schools focus on technical aspects, such as AI development and optimization. When AI programs are offered in business schools, they focus on using AI to propel business value, improve decision making, and boost business processes. There is some overlap in concepts, but still a distinct approach to each area. Given the differences in the themes, business school education in AI cannot rely on computing and engineering course offerings. More development of AI programs must occur across business schools to meet the industry needs.

7. REFERENCES

AACSB (2025). GenAI Adoption in Business Schools: Deans and Faculty Respond. Retrieved June 3, 2025 from <https://www.aacsb.edu/insights/reports/2025/genai-adoption-in-business-schools-deans-and-faculty-respond>

Alekseeva, L., Azar, J., Gine, M., Samila, S., & Taska, B. (2021). The demand for AI skills in the labor market, *Labour Economics*, 71, pp. 1-27. <https://doi.org/10.1016/j.labeco.2021.102002>

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2022). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*. <https://doi.org/10.48550/arXiv.2108.07258>

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Amodei, D.

- (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877-1901. <https://doi.org/10.48550/arXiv.2005.14165>
- Gartner (2025). Gartner Experts Answer the Top Generative AI Questions for Your Enterprise. Retrieved June 3, 2025 from <https://www.gartner.com/en/topics/generative-ai>
- Gonzalez Ehlinger, E., and Stephany, F., (February 2024). Skills or degree? The rise of skill-based hiring for AI and green jobs. Available at SSRN: <https://ssrn.com/abstract=4603764> or <http://dx.doi.org/10.2139/ssrn.4603764>
- Chen, L. (Fall 2022). Current and future artificial intelligence (AI) curriculum in business school: a text mining analysis, *Journal of Information Systems Education*, 33:4, pp. 416-426.
- Farrugia, C. (2023). Five digital skills for graduates to boost their employability. Available July 28, 2025, from: <https://www.businessnews.ie/education/five-digital-skills-for-graduates-to-boost-their-employability/>.
- Fuller, J., Sigelman, M., & Fenlon, M. (March 10, 2025). How Gen AI could change the value of expertise, *Harvard Business Review*.
- Gibson, K. (August 28, 2024). How AI training can help you compete in your industry, Harvard Business School. Retrieved July 28, 2025, from: <https://online.hbs.edu/blog/post/ai-training>.
- Harris, H. D., & Kiefer, S. M. (2004). The pedagogy of artificial intelligence: a survey of faculty who teach introductory AI. *Proceedings of the 17th International Florida Artificial Intelligence Research Society Conference*, Menlo, California.
- Hartley, J., Jolevski, F., Melo, V., & Moore, B. (December 18, 2024). The labor market effects of generative artificial intelligence. Available at SSRN: <https://ssrn.com/abstract=5136877> or <http://dx.doi.org/10.2139/ssrn.5136877>
- Hazan, E., Madgavkar, A., Chui, M., Smit, S., Maor, D., Dandone, G., & Huyghues-Despointes, R. (2024). A new future of work: The race to deploy AI and raise skills in Europe and beyond. McKinsey Global Institute. Retrieved March 27, 2025 from <https://www.mckinsey.com/mgi/our-research/a-new-future-of-work-the-race-to-deploy-ai-and-raise-skills-in-europe-and-beyond>
- Jeppson, T. (2025). How Is Artificial Intelligence Impacting the Job Market? LINQTO. Retrieved March 27, 2025 from <https://www.linqto.com/blog/how-is-artificial-intelligence-impacting-the-job-market/>
- KPMG (2025). KPMG AI Quarterly Pulse Survey, Q1 2025. Retrieved June 3, 2025 from <https://kpmg.com/us/en/articles/2025/ai-quarterly-pulse-survey.html>
- McKinsey (2025). The State of AI: How Organizations are Rewiring to Capture Value. Retrieved June 3, 2025, from <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- Microsoft (2025). Work Trend Index Annual Report 2025: The Year the Frontier Firm Is Born. Retrieved June 3, 2025, from <https://www.microsoft.com/en-us/worklab/work-trend-index/2025-the-year-the-frontier-firm-is-born>
- Poba-Nzaou, P., Galani, M., Uwizeyemungu, S., & Ceric, A. (2021). The impacts of artificial intelligence (AI) on jobs: an industry perspective, *Strategic HR Review*, 20:1, pp. 60-65. <https://doi.org/10.1108/SHR-01-2021-0003>
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, F., & Chen, M. (2022). Hierarchical text-conditional image generation with CLIP latents. <https://doi.org/10.48550/arXiv.2204.06125>
- Shen, Y. (2024). Artificial Intelligence and the Economy . The Impact of Artificial Intelligence on the Job Market, *Advances in Economics Management and Political Sciences*, 92(1) 71-74. <https://doi.org/10.54254/2754-1169/92/20231275>
- Sollosy, M., & McInerney, M. (2022). Artificial intelligence and business education: what should be taught, *The international Journal of Management Education*, 20, pp. 1-7. <https://doi.org/10.1016/j.ijme.2022.100720>
- Southworth, J., Migliaccio, K., Glover, Jo., Glover, Ja., Reed, D., McCarty, C., Brendemuhl, J., & Thomas, A. (2023). Developing a model for AI across the curriculum: transforming the higher education landscape via innovation in AI literacy, *Computers and Education: Artificial Intelligence*, 4, pp. 1-10. <https://doi.org/10.1016/j.caeai.2023.100127>

World Economic Forum (2025). The Future of Jobs Report 2025. Retrieved March 27, 2025 from <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>

Watson, C. E., & Rainie, L. (January 2025). Leading through disruption: higher education executives assess AI's impacts on teaching and learning, AAC&U (The American Association of Colleges and Universities), pp. 1-22.

Xu, J. J., & Babaian, T. (2021). Artificial intelligence in business curriculum: the pedagogy and learning outcomes, The

International Journal of Management Education, 19, pp. 1-18. <https://doi.org/10.1016/j.ijme.2021.100550>

Yang, S. C., & Weng, B. (2017). Toward a cybersecurity curriculum model for undergraduate business schools: A survey of AACSB-accredited institutions in the United States. Journal of Education for Business, 92(1), 1-8. <https://doi.org/10.1080/08832323.2016.1261790>

APPENDIX A: SCHOOL CLASSIFICATIONS

Arts & Sciences

College of Arts and Sciences
College of Liberal Arts
College of Liberal Arts and Sciences
School of Arts and Sciences

Business

School of Business

Computing & Engineering

College of Computing
College of Computing and Informatics
College of Engineering

College of Engineering and Computer Science
College of Engineering and Polymer Science
College of Engineering, Technology, and Physical Sciences
College of Science, Technology, Engineering & Mathematics
Herbert Wertheim College of Engineering
School of Computer Science
School of Computing
School of Engineering
School of Engineering and Applied Science
School of Engineering and Computer Science
School of Engineering, Computer Science and Artificial Intelligence
School of Informatics, Computing, and Engineering
Tickle College of Engineering

Health

College of Health and Human Performance
College of Public Health and Health Professions

Science & Technology

College of Information Science & Technology

College of Innovation and Technology

College of Natural Science

College of Science

College of Science and Technology

College of Sciences

College of Sciences and Technology

School of Science and Mathematics

APPENDIX B: THEMES

Final List of Themes

Algorithms and Data Structures	Database Systems
AI for Competitive Advantage	Deep Learning
Big Data	Discrete Structures
Business and Data Analytics	Ethics
Business Applications	Foundations of AI
Cloud Computing	Machine Learning
Computer Vision	Natural Language Processing
Data Analysis	Programming
Data Management	Robotics
Data Science	Software Engineering
Data Visualization	Specific Business Application of AI

Business and C&E Themes

Business Themes	C&E Themes
AI for Competitive Advantage	Algorithms and Data Structures
Algorithms and Data Structures	Cloud Computing
Big Data	Computer Vision
Business and Data Analytics	Data Science
Business Applications	Database Systems
Cloud Computing	Deep Learning
Data Analysis	Discrete Structures
Data Management	Ethics
Data Visualization	Foundations of AI
Deep Learning	Machine Learning
Ethics	Natural Language Processing
Foundations of AI	Programming
Machine Learning	Robotics
Programming	Software Engineering
Specific Business Application of AI	

Teaching Case

Developing Critical Thinking in Data Analytics Education: A Teaching Case Evaluating ChatGPT Responses to a Visualization

Benjamin Larson
blarson@aum.edu

Jeffrey Bohler
jbohler1@aum.edu

Nandini Bolekar
nbolekar@aum.edu

Auburn University at Montgomery
Montgomery, AL, 36117, USA

Hook

Can you trust what AI tells you? This case challenges students to analyze AI outputs and develop critical thinking in data-driven decision-making.

Abstract

Data analytics education increasingly incorporates generative AI tools, yet AI output can be inconsistent or misleading. This teaching case develops critical thinking by challenging students to evaluate AI-generated interpretations of data visualizations through Robert Ennis' three dimensions: logical, criterial, and pragmatic. Using a wage-by-tenure visualization that subtly violates linear regression assumptions, students prompt ChatGPT with structured queries, assess responses, and refine their prompts while engaging in peer reflection. The case targets the Data Understanding phase of CRISP-DM, where students must determine whether relationships are linear or require feature engineering. Structured rubrics guide assessment of prompt engineering and reasoning. Designed for undergraduate analytics, MIS, and statistics courses (adaptable to graduate programs), this case prepares students for AI-augmented workplaces by strengthening both technical skills and professional judgment in evaluating AI outputs. Supporting materials include complete datasets, sample deliverables, and validated implementation guides.

Keywords: Critical Thinking, Generative AI in teaching, Data Analytics, Data Visualization, Prompt Engineering, CRISP-DM

Recommended Citation: Larson, B.E., Bohler, J.A., Bolekar, N., (2026). Developing Critical Thinking in Data Analytics Education: A Teaching Case Evaluating ChatGPT Responses to a Visualization. *Information Systems Education Journal*, v24(n6) pp 31-43. DOI# <https://doi.org/10.62273/RLDU1198>

Developing Critical Thinking in Data Analytics Education: A Teaching Case Evaluating ChatGPT Responses to a Visualization

Benjamin Larson, Jeffrey Bohler, Nandini Bolekar

1. INTRODUCTION

Data-driven decision-making is essential for modern businesses and requires a wide range of technical skills (Collier & Powell, 2024). Connecting details to higher-level concepts requires education and practice (Holmes, Bialik, & Fadel, 2019). However, there is a shift towards low- and no-code solutions that enable a wider range of students to tackle more technical challenges, such as machine learning (Sundberg & Holmström, 2024). Generative AI is making many complex processes more straightforward, and we may see a future in which most code and programs are generated from prompts using natural language and multimodal inputs such as images, audio, and video. If this is the case, an emphasis on critical thinking and problem-solving is needed when using AI (Kasneci et al., 2023). Akhtar (2025) and many others argue that careers are less likely to be replaced by AI itself, but rather by individuals who effectively leverage AI tools, in other words, AI literate. This highlights the importance of exposure to AI while enhancing “people” skills (Gonzalez, 2023). Ultimately, generative AI can serve as a critical component in scaffolding, enabling students to tackle complex problems. This case allows students to improve their critical thinking skills by applying scaffolding to the Data Exploration step of the Data Understanding phase of the CRISP-DM framework. However, students still need to evaluate AI responses for relevance and accuracy.

2. CASE AS A PART OF A LARGER PROJECT

The assignment described here is part of a much larger semester project. The semester project is intended to use most or all of the CRISP-DM model and complete many of its tasks (Table 1). This assignment is designed to provide scaffolding for the data understanding phase, specifically to support data exploration. The overall goal of the semester-long project is to create a model that will allow Widgets Inc to control for and statistically evaluate wage fairness. It would be appropriate to provide students with portions of the Business Understanding Phase to help them understand the problem if this phase and its tasks have not

been completed earlier (Appendix A). The students may also be asked to generate the graph evaluated in the assignment, using the sample data provided in Appendix B and the complete data set available upon request. The assignment involves evaluating one scatter plot; an example is provided in Appendix C. As this assignment is intended to be performed during data understanding or preparation, additional slides and materials are available to supplement those phases.

CRISP-DM Phase	Tasks	Weeks
Business Understanding	Determine Objectives, Assess Situation, Determine Goals, Produce Plan	2
Data Understanding	Collect Data, Describe Data, Explore Data, Verify Quality	3
Data Preparation	Select Data, Clean Data, Construct Data, Integrate Data, Format Data	3
Modeling	Select Model, Generate Test Design, Build Model, Assess Model	4
Evaluation	Evaluate Results, Review Process, Determine Future Steps	1
Deployment	Develop Strategy, Integrate Systems, Maintain, Collect Feedback, Document	1

Table 4. CRISP-DM Phases and Selected Outputs

3. TARGET AUDIENCE AND PREREQUISITES

This case is designed for undergraduate or graduate students in a statistics, data mining, or data analytics course. Students should be introduced to the basics of analytics and statistics. However, this assignment is limited to

the relationship between two variables and between two groups, so advanced knowledge is not required. No prior experience with AI is required, although prompt engineering may be discussed either before or after the assignment is completed.

4. IMPLEMENTATION

This assignment may be implemented in several different ways. Semester project groups of two to five students should be assigned or formed in advance. The assignment may then be given to each individual to complete, after which they discuss it within their groups. Completion of the assignment and initial group discussions should take roughly 20-30 minutes. Depending on prior knowledge, a discussion of linear regression assumptions, with an emphasis on linearity, should be conducted (10-15 minutes). The instruction should also cover critical thinking skills (10-15 minutes) before attempting the assignment. The discussion of the implications of the responses should then be held or presented as discussion boards (20-30 minutes).

Activity	Time
Critical Thinking	10-15 minutes
Model Assumptions	10-15 minutes
Assignment	20-30 minutes
Assignment Discussion	20-30 minutes

Table 5. Assignment Implementation Timeline

5. TEACHING CASE FOUNDATION

5.1 Learning Objectives

Upon completion of this case, students will be able to:

1. Critically evaluate AI-generated statistical interpretations using structured analytical frameworks (Ennis' logical, criterial, and pragmatic dimensions).
2. Identify violations of linearity assumptions in scatterplot visualizations through visual pattern recognition.
3. Construct and refine effective prompts through conversational AI interaction and chain-of-thought techniques.
4. Assess the reliability and consistency of generative AI outputs across different users and models.
5. Apply CRISP-DM Data Understanding phase concepts to justify data preparation decisions.
6. Articulate when human domain expertise should override AI recommendations in data analysis.
7. Demonstrate improved AI literacy by

recognizing model training effects and prompt engineering importance.

5.2 Theoretical Framework

This case draws on three complementary pedagogical frameworks. In the Critical Thinking Framework, Ennis (1962) identifies three dimensions of critical thinking that provide structure for this case:

- Logical dimension: Students identify patterns and relationships in data visualizations
- Criterial dimension: Students apply statistical knowledge to evaluate AI claims against known assumptions (e.g., linearity in regression)
- Pragmatic dimension: Students make practical decisions about when to override AI recommendations based on domain expertise

Additionally, Constructivist Learning (Piaget, 1954; Vygotsky, 1978) is used, with students actively constructing understanding through AI interaction rather than passively receiving instructor explanations.

5.3 Summary

This assignment is a critical thinking exercise that also encourages students to interact with generative AI during either the Data Understanding or Data Preparation phase of the Cross-Industry Standard Process for Data Mining (CRISP-DM) or a similar phase in other machine learning or data mining models. Using Robert Ennis' three dimensions of critical thinking (logical, criterial, and pragmatic) (Ennis, 1962), it may also be delivered as an introductory assignment that encourages the use of generative AI in the course while discussing its limitations. The students need to have sufficient knowledge to evaluate the critical dimension of the response and enough confidence in themselves to override the AI tool's response.

Data visualizations are a popular tool for evaluating relationships between variables. For this study, students will upload an image that explores the relationship between years of service and wages for two positions at a fictitious company. Specifically, they will evaluate whether the relationship between years of service and wages is linear, which is an assumption of linear regression.

The visualization presented (Figure 1) is not intended to blatantly violate the linearity assumption, but rather to indicate that the relationship may be changing. This is something that not all AI models may pick up on, given a small number of records. However, when

combined with a user who critically evaluates visualization with knowledge of the problems and model assumptions, there is potential to improve the model by engineering features based on business principles rather than aiming for the best fit, which may lead to overfitting. Within the figure, Position 1 Wages become non-linear between Year 6 and 7. This would violate the linearity required for a linear regression and should be taken into consideration. Position 1 represents workers, and Position 2 represents supervisors, both of whom are paid hourly.

AI will also not necessarily provide consistent responses to similar prompts, which can cause an AI divide in the classroom, as students who have previously worked with AI on similar problems may receive better responses to their prompts. Through peer discussion, students will discover the need to train their model appropriately and to improve their ability to prompt AI to receive appropriate responses. To help facilitate this, they will be given a series of prompts to gauge the performance of their generative AI.

The objective of the case is to develop critical thinking skills by interacting with AI, analyzing responses, and questioning assumptions in data interpretation. Students will also learn to generate AI responses and critically evaluate their accuracy, biases, and limitations. Finally, through interactions with their classmates, they will also evaluate if their models are similarly trained and learn the importance of training their models to improve future results.

To maximize the pedagogical impact of this AI-enhanced critical thinking exercise, instructors should establish clear evaluation criteria that assess both technical competency and analytical reasoning. Students benefit from structured rubrics that evaluate their ability to craft effective prompts, critically assess AI-generated interpretations, and identify when generative AI fails to recognize subtle patterns in data visualization.

The assignment's success hinges on creating collaborative learning environments where students compare AI responses and identify inconsistencies, fostering discussions about variations in model training and the importance of human expertise in validating AI conclusions. By incorporating peer review sessions and reflective components, educators can ensure students develop not only technical skills in AI interaction but also the professional judgment necessary to leverage these tools effectively in real-world business contexts, ultimately

preparing them for a future where AI augments rather than replaces human decision-making capabilities.

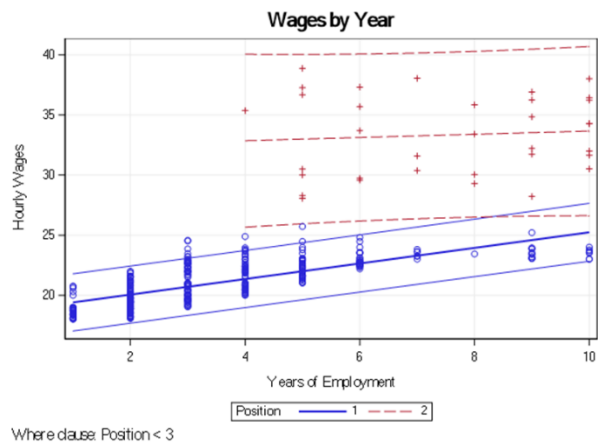


Figure 1. Wages by Year Visualization

6. CRISP-DM

This case employs the Cross-Industry Standard Process for Data Mining (CRISP-DM) (Wirth & Hipp, 2000; IBM, n.d. -a) cycle framework, which comprises six phases: business understanding, data understanding, data preparation, modeling, evaluation, and deployment (Wirth & Hipp, 2000), see Figure 2.

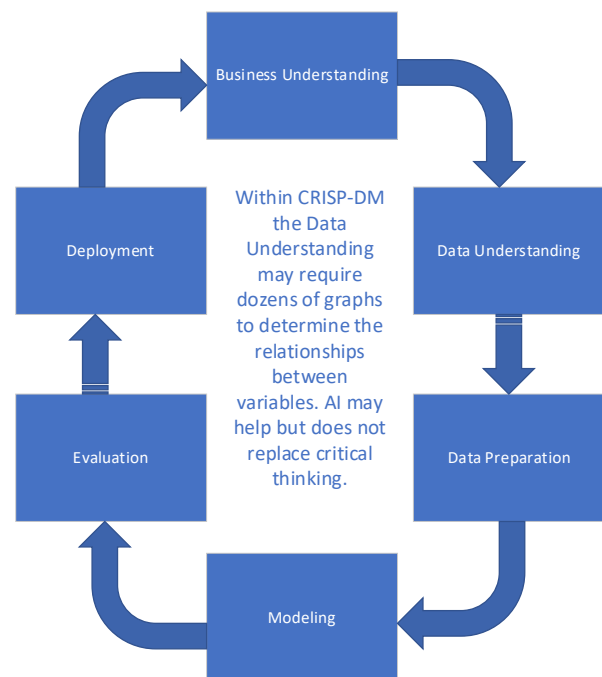


Figure 2. CRISP(DM) Data Mining Life Cycle

Business Understanding defines the project objectives and requirements from a business

perspective, converting them into a data mining problem definition that includes project management, situational assessment, and identification of business assumptions. A key assumption of the business may be that, for some positions, there is a rapid rise to equitable compensation early on, leading to lower salary increases once an employee has established seniority.

During the data understanding phase, the team collects initial data, describes its characteristics, and explores it to identify data quality issues, patterns, or initial hypotheses. In this phase, we may evaluate assumptions, such as linearity among the variables, to determine if a model is appropriate or if feature engineering is required. The case study is a part of the Data Understanding phase of CRISP-DM. The results are included in the example memo for the larger case used in the class, where it may be appropriate to engineer a feature that maximizes the influence of years of service. This may be something discovered at this phase, but it would be more appropriate to note it as a business assumption in the prior phase. The exploration should then lead to action in the Data Preparation phase.

In the Data Preparation phase, the team selects, cleans, engineers, features, and formats the data for modeling. The decision may be made to engineer a feature that caps the years of service used in the linear model.

Modeling involves selecting and applying various modeling techniques. These techniques include adjusting model parameters and evaluating each model's performance to determine its effectiveness.

Evaluation is the process of reviewing model results to ensure they meet business objectives and to determine whether additional steps or revisions are needed. In this case, the students would likely create a model that caps the number of years of service and then evaluate whether this improves the model fit.

Deployment is the phase that delivers the results, including reports, dashboards, or integrated models or systems, into the business environment. A model can be created to establish statistical control over employee salaries, which can then be used to determine whether employees are being paid equitably.

Each phase is meant to be connected to the others; however, without examples, students

may lose that connection. This case serves as an example: business understanding provides the business goals and the project goals, while also assessing the situation, which may include listing assumptions.

One of the assumptions is that employee salaries for certain positions may rise earlier in an employee's tenure and then stabilize. The need to engineer a feature should be noted during the data understanding phase, as the data is explored, then in the data preparation phase, and finally modeled and assessed in the subsequent phases.

7. AI RESPONSE EVALUATION

Information systems education should incorporate AI into our courses, including allowing students to utilize AI in projects to enhance their learning (Zhang, 2025). By having students prompt AI in various ways to understand its limitations, they will develop skills essential in the modern workplace. Prompt engineering is the writing of prompts in a way that enables you to get the desired results from AI. The more knowledgeable students are about the problem's context and terminology, the more effective they will be at evaluating AI responses and adjusting their prompts. However, not all students will start at the same point with their AI. Different out-of-the-box models may yield different results or have been modified by the individual user through training. This represents one or more "shots of training" of the model by the individual student. In this way, the model may be improved to provide better results (Li et al., 2023).

For users who lack the knowledge or experience to create fully developed prompts, conversational prompting may be a more suitable approach (Ein-Dor et al., 2024). This allows students to engage AI conversationally rather than structuring a formal prompt, enabling AI to provide scaffolding for individual concepts, as students may focus on one area of a complex problem rather than creating a complex prompt. Here, having a conversation with AI about the relationships between a few variables and whether they meet an assumption would allow students to understand a smaller section of the analysis, rather than evaluating the whole process and model at once. This can be combined with a chain-of-thought prompt, which requires the user to understand the reasoning or train of thought that leads the model to make its response (IBM, n.d.-b; PromptingGuide.ai, 2025; Mitra et al.,

2024).

Another trend in AI prompting is its growing multimodality, allowing images and other media alongside text to be used in prompts (Mitra et al., 2024). Utilizing visualizations in prompts and discussing the results in conversations with AI will provide personalized scaffolding that instructors would not otherwise be able to provide easily. However, it is something the instructor needs to monitor, as prompt responses will not be consistent if they are too complicated or if the prompters use different models.

By having assignments that prompt discussions of their results, students may become aware of potential limitations as their prompts become more complex. It will also help train less-developed models, especially if students have varying levels of exposure to statistics and/or AI. Students can use AI for reflective learning. Reflective learning using AI refers to the process by which learners *critically assess* their experiences, decision-making processes, emotions, outcomes, and assumptions with the support of AI tools (such as large language models or generative AI), enabling them to identify strengths and weaknesses, deepen self-awareness, and improve future performance (Yuan & Hu, 2024).

8. CASE SETUP

The case presents a novel approach that enables students to interact with AI and with one another in the classroom. In this way, it promotes the use of novel technology, develops critical thinking skills, and helps build a learning community in the classroom. The assignment may be given as a formal assignment to be completed in advance of class or uploaded to a discussion board (Appendix C). Alternatively, the assignment may be given with a class period, and the class may discuss the concepts as they complete each prompt.

The students should be made aware, prior to the assignment, that some positions may experience a rapidly flattening wage-increase curve and that the linear regression assumption is based on a linear relationship between the dependent and independent variables. For this example, the dependent variable would be wages, and we want the AI to recognize that the relationship between years of employment and wages flattens between six and seven years of employment.

Students should be directed to use a generative AI tool, such as ChatGPT or Copilot, and to prompt the AI as detailed below. Between each prompt, students should discuss the assigned

discussion question in small groups. Alternatively, suppose the student is in an online course. In that case, they should be directed to save their responses and to respond to each question in a single post, so the rest of the class can see and respond if they receive similar results. Critical Thinking tasks are abbreviated as CT in the list of prompts below.

Prompt 0: Upload the image with no additional comments.

Discussion Question 0.1 Did AI provide an adequate description of what the image was (CT – Pragmatic)?

Discussion Question 0.2 - Did AI note any change in the relationship of wages and years of employment for position 1 (CT – Logical)?

Prompt 1: Do the years of employment appear to have a linear relationship with wages?

Discussion Question 1.1 - Did AI note that the relationship was linear (CT – Logical)?

Discussion Question 1.2 - Did AI note any caution related to Position 1 having a changing relationship? If caution was noted, what was its chain of thought? Make an argument about whether the chain of thought is appropriate if provided (CT Pragmatic & Criterial).

Discussion Question 1.3 - Evaluate whether additional prompting is needed to elicit the model's reasoning process and justify your decision (CT Pragmatic & Criterial).

Prompt 2: Are you concerned about the values after year six being below the regression line?

Discussion Question 2.1 - Did AI note that there is a concern that the relationship may be changing (CT Logical)?

Discussion Question 2.2 - Did AI provide a reason why the relationship may be changing? If yes, what was it (CT Logical)?

Discussion Question 2.3 - Did AI provide a reason why the results may still be linear? If yes, what was it (CT Logical)?

Discussion Question 2.4 - Explain whether you agree with the AI and why (CT Pragmatic & Criterial).

Prompt 3: Would it make sense to cap the years of employment at seven? Why or why not?

Discussion Question 3.1 - Did AI provide a recommendation? If yes, what was it (CT Logical)?

Discussion Question 3.2 - Did AI justify? If yes, what was it (CT Logical)?

Discussion Question 3.3 Did AI provide alternatives? If yes, what were they (CT Logical)?

Discussion Question 3.4 - Make an argument as to why you would agree or disagree with the recommendation (CT Pragmatic & Criterial).

Prompt 4: What formula would I use in Excel to generate a variable with a maximum of seven years?

Discussion Question 4.1 - Did AI provide the correct formula (CT Criterial)?

Discussion Question 4.2 - How would you describe the value of AI and the human operator in this process (CT Logical, Criterial, and Pragmatic)?

9. CONCLUSIONS

This case provides a practical example aligned with CRISP-DM phases while introducing students to multimodal prompting. Given current AI capabilities, students should also be aware of inconsistencies in the responses generated by generative AI. The case mainly received positive feedback from students for its simple execution and relevance as a discussion topic. For many, the change in the relationship between wages and years of employment would not have been apparent at all. This should leave students feeling they need to critically analyze AI-generated output to add value to the analysis they are asked to perform.

10. CASE LIMITATIONS

This case provides a straightforward introduction to critical thinking with AI in data analytics, but instructors should be aware of several limitations when implementing it.

AI Model Evolution and Variability

This case was developed and tested with specific versions of generative AI (primarily ChatGPT 3.5, ChatGPT 4, and Microsoft Copilot) during 2023-2025. AI models are continuously updated, and their responses to the same prompts can change over time. This creates several challenges:

Model Updates: A strength of this case is that it reveals AI's inconsistencies and limitations. However, if future AI versions consistently detect the non-linearity issue, the pedagogical value may shift. Instructors should pilot the exercise with current AI versions before each semester to understand likely student experiences. If AI has improved significantly, the case still works. The discussion can shift from AI's limitations to its

growing capabilities. The critical thinking framework remains valuable in either scenario.

Response Variability Within Sessions: Even with the same AI version, students may receive different responses depending on their account history and prior interactions with the AI. This variability is pedagogically functional, as it prompts discussions about model training and the importance of consistent analytical processes. However, it can be frustrating for students who expect deterministic outputs. Instructors should frame this variability as part of the learning objective rather than a problem.

Technology Access and Equity

AI Account Requirements: Students need access to a generative AI tool capable of processing images. As of this writing, several options exist with free tiers (ChatGPT, Microsoft Copilot, Google Gemini, Claude). However, free-tier limitations may affect the student experience. Some students may have paid subscriptions that provide access to more advanced models, potentially creating inequity in the quality of AI responses. Instructors should discuss this openly and frame it as an example of how AI access itself can create workplace advantages.

Prerequisite Knowledge Assumptions

Statistical Background: This case assumes students have encountered linear regression and understand the concept of assumptions, even if they cannot articulate them formally. Students with no prior exposure to statistics may struggle to evaluate whether the visualization violates the linearity assumptions. Instructors should assess whether students need a brief review of linear regression concepts before the exercise.

CRISP-DM Familiarity: The case is designed for courses using the CRISP-DM framework. Students unfamiliar with this process model may not fully appreciate how the Data Understanding phase connects to subsequent phases. Instructors using different frameworks (e.g., KDD, SEMMA) should adapt the language and connections accordingly.

Data Visualization Literacy: Students should be able to interpret scatterplots with trend lines. Those with limited exposure to data visualization may need additional scaffolding to identify patterns visually.

Despite these and other implementation limitations (e.g., Face-to-Face or Online), instructors can adapt the case to different class sizes, exercise durations, and student experience

levels. The case's core pedagogical value, to develop critical thinking about AI outputs, remains relevant even as specific AI capabilities evolve. Instructors should view these limitations as opportunities to adapt and contextualize rather than barriers to implementation.

11. REFERENCES

- Akhtar, R. (2025, March 27). *AI will not replace you. A human using AI will.* *Forbes*. Retrieved August 15, 2025, from <https://www.forbes.com/sites/reeceakhtar/2025/03/27/ai-wont-replace-you-a-human-using-ai-will/>
- Collier, C. A., & Powell, A. L. (2024). Data Analyst Competencies: A Theory-Driven Investigation of Industry Requirements in the Field of Data Analytics. *Journal of Information Systems Education*, 35(3), 325–376. <https://doi.org/10.62273/SPYC4248>
- Ein-Dor, L., Toledo-Ronen, O., Spector, A., Gretz, S., Dankin, L., Halfon, A., & Slonim, N. (2024). Conversational prompt engineering. <https://doi.org/10.48550/arXiv.2408.04560>
- Ennis, R. H. (1962). *A concept of critical thinking* (p. 84). Harvard Educational Review.
- Gonzalez, A. (2023, August 22). *IBM study: AI won't replace people, people who use AI will replace people who don't.* *Going Concern*. Retrieved August 15, 2025, from <https://www.goingconcern.com/ibm-study-ai-wont-replace-people-people-who-use-ai-will-replace-people-who-dont/>
- Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial Intelligence in Education: Promises and Implications for Teaching and Learning*. Boston: Center for Curriculum Redesign.
- IBM. (n.d.-a). Modeling. In IBM SPSS Modeler SaaS documentation. Retrieved August 15, 2025, from <https://www.ibm.com/docs/en/spss-modeler/saas?topic=guide-modeling>
- IBM. (n.d.-b). What is chain of thought (CoT) prompting? IBM Think. Retrieved August 15, 2025, from <https://www.ibm.com/think/topics/chain-of-thoughts>
- Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groß, C., Günemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Larissa, N., Schlag, B., & Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Li, Y., Hui, B., Xia, X., Yang, J., Yang, M., Zhang, L., & Li, Y. (2023). One-shot learning as instruction data prospector for large language models. <https://doi.org/10.48550/arXiv.2312.10302>
- Mitra, C., Huang, B., Darrell, T., & Herzig, R. (2024). *Compositional chain-of-thought prompting for large multimodal models*. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 14420–14431). IEEE. <https://doi.org/10.1109/CVPR.2024.01445>
- Piaget, J. (1954). *The construction of reality in the child*. Basic Books.
- PromptingGuide.ai. (2025, July 5). *Chain-of-Thought (CoT) prompting*. Retrieved August 15, 2025, from <https://www.promptingguide.ai/techniques/cot>
- Sundberg, L., & Holmström, J. (2024). Teaching Tip: Using No-Code AI to Teach Machine Learning in Higher Education. *Journal of Information Systems Education*, 35(1), 56–66. <https://doi.org/10.62273/CYPL2902>
- Wirth, R., & Hipp, J. (2000, April). CRISP-DM: *Towards a standard process model for data mining*. In Proceedings of the 4th International Conference on the Practical Applications of Knowledge Discovery and Data Mining (Vol. 1, pp. 29–39).
- Yuan, B., & Hu, J. (2024, November 19). *Generative AI as a tool for enhancing reflective learning in students*, <https://doi.org/10.48550/arXiv.2412.02603>
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.
- Zhang, X. (2025). Teaching Tip: Incorporating AI Tools Into Database Classes. *Journal of Information Systems Education*, 36(1), 37–52. <https://doi.org/10.62273/GKZI247>

APPENDIX A

Relevant Parts of the Business Understanding Phase

Business Understanding

Determine Business Objectives

Background

Widget Inc. has several locations where it provides goods and services to local consumers. The executives at Widget Inc. are concerned about labor cost equity and want to establish statistical controls to monitor whether employees are paid fairly. They have had employees complain of discrimination based on height and gender, and may be facing a lawsuit.

Objectives

The primary objective of this analysis is to determine what factors influence salary at Widgets Inc. The analysis will evaluate how performance, location, position, and experience influence salary. Furthermore, the constructed model is intended to enable the company to test for discrimination based on height or gender and to provide human resources with a tool to identify outliers that may indicate salary issues within the organization.

Business Objectives Success Criteria

Success will be determined by the ability to create a model accurate enough to establish statistical controls to monitor fair pay within the organization. An essential metric the company wants to establish is a performance-based salary increase baseline.

Data Mining Tools

Our team has access to Excel, SAS Studio, and SAS Viya

Data Sources

Human resources has provided us with a CSV file containing 400 records and 8 attributes, including salary, position, location, performance review, on-the-job errors, gender, height, and years of employment.

Other Resources

This could be money, or other teams or organizations with unique skills.

Requirements

Schedule: Analysis will be complete by the Final Documents due date (look up the due date). Final recommendations will be completed by the due date of the presentation (look up the due date).

Quality: The project will detail all steps to prepare the data, verify the model's appropriateness, assess its accuracy, provide a detailed analysis of the results, and make recommendations based on the analysis. (NOTE: We will revisit this section as we conduct our analysis to add detail.)

Security: No results should be shared with anyone outside the team or human resources, as employees may be inferred from even minimal analysis.

Legal: The organization is legally required to protect personal information and to ensure equal pay for equal work regardless of gender, race, etc.

Data Mining Goals

Our analysis should provide a model to estimate an employee's salary, allow for statistical control of salary, and estimate the influence of years of service and performance on salary.

Data Mining Success Criteria

A model explaining more than 80% of the variance and estimates with 95% confidence of the impact of the years of service and performance. The model needs to be simple enough to allow continual review at a nominal cost.

APPENDIX B

Sample Data

EmpID	Location	Position	Hourly Wages	On the Job Errors	Performance Review	Years of Employment	Height	Gender
3464555	10	1	23.21	1	19	9	63	M
3511683	2	2	33.69	1	19	6	70	M
3517969	11	3	72.71	1	19	9	74	F
3519233	5	1	20.85	13	7	2	68	F
3526791	8	1	18.55	11	9	2	77	M
3533586	11	6	139.29	3	17	8	64	F
3538924	10	1	19.9	11	9	3	61	F
3552242	5	1	19.77	14	6	2	65	M
3571609	3	1	22.93	2	18	6	73	M
3597660	6	1	21.1	7	13	2	72	M
3602023	1	1	18.77	16	4	1	75	F
3680725	8	1	20.7	7	13	4	69	F
3681183	11	4	46.85	2	18	10	67	F
3698777	10	1	19.95	15	5	3	65	M
3717787	3	1	18.12	15	5	2	78	M
3735292	11	3	63.01	0	20	7	76	F
3745587	5	1	23.74	1	19	4	64	F
3791596	7	1	20.53	8	12	3	67	F
3793117	1	1	19.79	9	11	2	73	M

APPENDIX C

Example of an AI Assignment

Assignment: Critical Thinking in AI-Assisted Data Interpretation

Objective

This assignment will help you develop critical thinking skills by interacting with AI, analyzing responses, and questioning assumptions in data interpretation. You will generate your own AI responses and critically evaluate their accuracy, biases, and limitations. You will also interact with your classmates to evaluate the consistency of your AI's responses to the prompts.

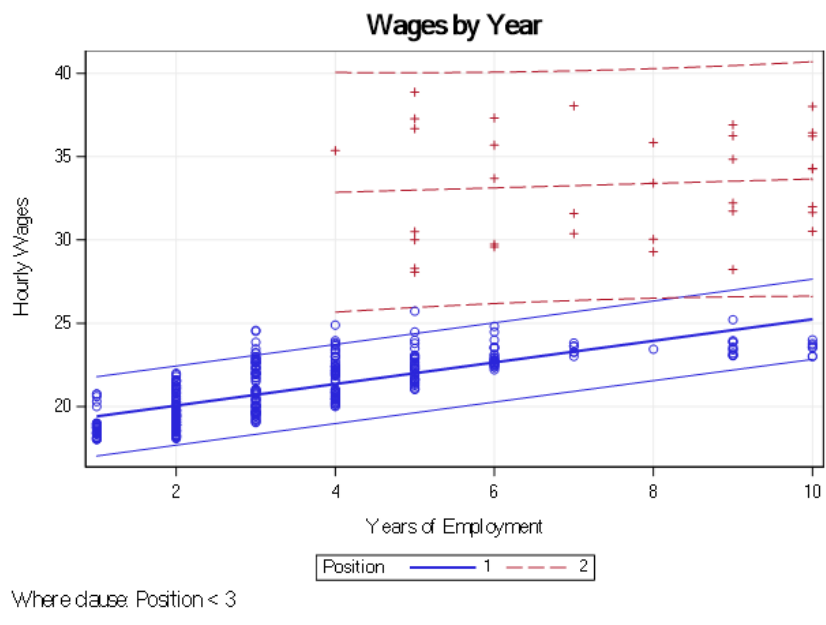
Instructions

Part 1: Engage with AI

1. Use the following prompts to interact with <https://chatgpt.com> one at a time and record the responses.
2. Do **not** guide the AI in any direction—let it generate its own reasoning.
3. Copy and paste the AI's responses into your assignment document.

Prompts to Use:

- Copy and paste the following image.



- Prompt 1: Do the years of employment appear to have a linear relationship with wages?
- Prompt 2: Are you concerned about the values after year 6 being below the regression line?
- Prompt 3: Would it make sense to cap the years of employment at 7? Why or why not?
- Prompt 4: What formula would I use in Excel to generate this variable?

Part 2: Critical Analysis

For each AI response, answer the following:

1. **Does your AI response match your group?**
2. **Did the AI response express concern over linearity before prompt 2?**
3. **What recommendation did AI provide for capping the years of employment?**
4. **Do you agree with the recommendation? Why or why not?**

Part 3: Reflection & Summary

Write a **brief reflection** on what you learned from this exercise. Consider:

- Did the AI responses align with your expectations?
- Did you find any biases or limitations in the AI's reasoning?
- How did questioning the AI responses deepen your understanding of data interpretation?
- What do you feel your role is in data understanding, given the rise of generative AI?

Grading Criteria:

- **Engagement (30%)** – Did you record and present the AI's responses accurately?
- **Critical Analysis (40%)** – Did you effectively evaluate and critique the AI's reasoning?
- **Reflection (30%)** – Did you demonstrate thoughtful learning from the exercise?

Prompt	Response
Image	
1	
2	
3	
4	

Teaching Activity

Introducing Data Analytics and AI Collaboration to Novice Students: An Assignment Using Real Business Data

Thilini Ariyachandra
ariyachandrat@xavier.edu
Xavier University
Cincinnati, OH USA

Hook

"How do you teach data analytics and AI collaboration together in introductory IS courses? This assignment shows students developing both competencies simultaneously through real world business scenarios."

Abstract

Information systems instructors face the challenge of teaching both data analytics and artificial intelligence (AI) collaboration to novice students in introductory courses where students have limited experience with either tool. Mostly, these skills are taught separately, leaving students unprepared for modern workplaces where data analysis and AI collaboration happen together as standard business operations.

This assignment provides a structured multi-stage assignment to apply both foundational data analytics and AI collaboration skills through a real-world business scenario. Students assume the role of junior analysts at Starbucks Corporation directed to find stakeholder focused data insights. They work with nutritional data for beverages to develop actionable insights while learning to collaborate with large language models as analytical assistants. The assignment is designed for introductory IS courses where students have basic familiarity with visualization tools and AI platforms from earlier course modules.

The four-stage assignment structure leads students through data exploration, dashboard creation in Tableau, AI supported revision and reflective analysis. The business context requires students to move from exploring their own interests in the dataset to developing insights that serve specific organizational stakeholders such as Marketing or Nutrition Policy teams.

Assignment adoption across course sections demonstrate that novice students gain both analytical thinking and AI collaboration skills through the assignment. Student reflections indicate a shift in learning from a focus on completing tasks toward a more iterative and reflective approach to analytical work. The assignment contributes to IS education by demonstrating that both analytical reasoning and collaborative engagement with AI can be developed through structured, experiential learning.

Keywords: Data Analytics Education, AI Collaboration, Generative AI, Introductory Information Systems Courses, Data Visualization

Recommended Citation: Ariyachandra, T., (2026). Introducing Data Analytics and AI Collaboration to Novice Students: An Assignment Using Real Business Data. *Information Systems Education Journal*, v24(n?) pp 44-62. DOI# <https://doi.org/10.62273/IJIV9938>

Introducing Data Analytics and AI Collaboration to Novice Students: An Assignment Using Real Business Data

Thilini Ariyachandra

1. INTRODUCTION

The business world has changed dramatically. Organizations now generate massive amounts of data and artificial intelligence (AI) tools have become commonplace in workplace analytics. The U.S. Bureau of Labor Statistics projects that data scientist roles will grow 36% from 2023 to 2033, with over 113,000 new positions expected (U.S. Bureau of Labor Statistics, 2024). At the same time, a 2024 McKinsey survey found that 65 percent of organizations regularly use generative AI tools for business functions, including data analysis (Singla et al., 2024). This creates a clear need in industry for both these skills. Business graduates must understand both basic data analytics and how to work with AI tools effectively.

Most students entering information systems (IS) courses have little experience with either data analytics or AI collaboration. They may have never worked with a dataset or used tools like Copilot or Chatgpt for analytical tasks. Traditional introductory courses often treat these skills separately, if they address AI at all. This leaves a significant gap in preparing students for modern workplaces where data analysis and AI collaboration happen together as part of normal business operations.

Students must learn to explore data, identify patterns and create visualizations while also learning to use AI tools as analytical assistants. The key challenge is teaching these skills simultaneously without overwhelming novice learners. Students need to develop their own analytical thinking while learning to collaborate with AI tools effectively and critically.

The assignment described in this paper addresses this challenge through a series of structured tasks that introduces students to data analytics using real business data while teaching them to collaborate with AI tools. Students work with Starbucks beverage nutritional data to apply the data analytics process from data exploration through dashboard creation to insight refinement. During this process, they use large language models (LLMs) to get feedback and improve their

work. The assignment assumes students have prior exposure to both data visualization concepts and LLMs through introductory IS coursework. Although implemented using Tableau, the assignment structure can be adapted to other visualization platforms depending on course context and instructor goals. The assignment focuses on applying these existing skills to business analytics scenarios.

This approach fills an important gap in current IS education. Students often encounter AI tools informally without guidance on appropriate professional use. The assignment provides an opportunity to engage with an AI tool meaningfully within a data analytics scenario. Students are exposed to both the capabilities and limitations of AI assistance while developing basic data analysis skills. This prepares them for the AI integrated analytical environments they will encounter in their careers.

The assignment contributes to IS education by offering a practical framework that instructors can implement immediately. The four-stage structure provides clear guidance for applying both data analytics concepts and AI collaboration skills to a real-world business scenario. The familiar Starbucks business context allows students to focus on learning analytical concepts and AI collaboration without being distracted by unfamiliar domains.

The learning outcomes of the assignment show that students successfully develop both analytical thinking and AI collaboration skills. They gain confidence working with data, learn to create effective visualizations and understand how AI tools may be used as analytical assistants rather than replacements for critical thinking. These combined competencies prepare students for the data driven, AI enhanced business environments they will enter upon graduation. It responds to an urgent need in undergraduate business education where collaboration with AI is increasingly becoming an expected skill across data analytics and other functional domains.

First this paper describes current challenges in teaching data analytics to novice students and examines emerging practices for AI integration in

business education. Second, it describes the four-stage assignment structure, including the business context, dataset design and specific learning activities at each stage. Third, it presents student learning outcomes and reflections that demonstrate how the assignment develops both analytical thinking and AI collaboration skills. Finally, it provides implementation guidance and teaching notes to help instructors adapt the assignment for their own courses. This paper presents the assignment as both a teaching activity (for instructors) and a complete student handout (Appendix A), offering a structure that is practical and easy to adopt.

2. LITERATURE REVIEW

Challenges in Introductory Data Analytics Education

Today, instructors face challenges in engaging novice students with data analytics. Traditional approaches often emphasize statistical theory over practical application. This can be overwhelming or frustrating to students who lack prior experience with datasets (Campos et al., 2012). Experiential learning can be a response to this challenge which can be a means of engaging with data early in the data analytics learning process. The National Academies (2018) suggest experiential learning approaches, which involve direct interaction with data, are effective in preparing students before introducing complex analytical concepts.

Knaflic (2015) argues that effective data visualization education must begin with hands on practice rather than theoretical frameworks. Kolb's (2014) experiential learning theory reinforces this view and highlights the role of experience in the development of analytical skills. However, most business students enter analytics courses with limited mathematical preparation and very limited experience working with real datasets (Diamant, 2024; Stratopoulos et al., 2019). Few (2012) emphasizes the importance of real-world business contexts in education, noting that students engage more effectively when they understand how analytical work serves organizational needs.

AI Integration in Business Education

The integration of large language models into educational settings has gained considerable attention recently. Still, research on effective implementation remains limited. Mollick and Mollick (2023) found that AI tools can enhance learning when used as collaborative partners rather than replacement technologies. Van Allen and Katz (2023) call for strategies that encourage

students to engage in the evaluation of AI generated output rather than passive acceptance of AI suggestions.

However, Zawacki-Richter et al. (2019) and Tan et al. (2024) note that most educational AI research focuses on preventing misuse rather than teaching appropriate business use cases. This creates a gap between student informal AI use and workplace expectations. As a result, practical guidance for integrating AI into business analytics education remains limited.

Data Visualization and Iterative Learning

Developing effective data visualizations as part of the analytics process requires students to build both technical proficiency and visual literacy. Design thinking offers a helpful foundation by encouraging students to think intentionally about how data is presented, not just how it is processed (Krekhov, Michalski, & Krüger, 2019). Few (2012) emphasizes that strong visualizations depend on understanding both the capabilities of the tools and the principles of clear communication. Yet students often focus more on technical execution. They overlook the importance of clearly conveying analytical meaning. Diamant (2024) observes that this can result in polished visuals or models that appear sophisticated but fail to convey meaningful insights or support sound business understanding.

Helping students learn to communicate through data takes time, iteration and feedback. The design education literature points to the importance of formative and iterative feedback (Joshi et al., 2021). Delivering that kind of support at scale remains a major challenge in most classroom settings (Jamal & Renzella 2024). Within this assignment, iteration is supported through cycles of feedback generated by AI tools. Students first create a visualization, receive structured commentary on clarity and design, and then revise their work in response. This process enables multiple rounds of refinement without placing an undue grading burden on the instructor. Faculty oversight is still part of the process as the professor monitors progress and offers targeted feedback at key points in the process. AI tools offer a potential way forward by providing structured and timely feedback on both visual and communication elements of assignments. This remains a promising and developing area of practice with growing relevance as instructors look for practical ways to support student learning in data storytelling.

As the literature discussion above suggests, there are several tensions that impact learning basic analytics skills and AI collaboration at the introductory level in IS business education. Most analytics education efforts within IS programs are concentrated in specialized or graduate level courses. Limited attention is given to how novice business students begin developing analytical skills in general education settings. This makes it difficult to align foundational IS skill development with industry expectations. (Lyytinen, Topi, & Tang, 2023). At the same time, most current research on AI integration has focused on tools or technical adoption without considering how these technologies influence the learning process (Zawacki-Richter et al., 2019). In addition, no previous studies have examined approaches for introducing data analytics and AI collaboration simultaneously to novice students.

This assignment responds directly to these tensions by offering a structured, experiential approach that supports both data analysis and AI collaboration in an introductory IS course setting. By engaging students with real world business data and guiding them through the process of interpreting, visualizing and refining their work, the assignment offers a practical model for integrating AI meaningfully into early stage analytics education.

3. ASSIGNMENT DETAILS - THE STARBUCKS ANALYTICS CHALLENGE

Students take on the role of junior analysts on Starbucks Corporation's Consumer Insights Team to work on this assignment. This business context provides familiar reference points while engaging students with realistic analytical challenges that are similar to professional data analysis workflows. Instructors can find the complete student assignment description, which serves as a ready to distribute assignment handout in Appendix A.

Business Context and Stakeholder Framework

The assignment establishes a realistic business problem centered on leadership concerns about beverage nutritional content and the need for clearer insights to support strategic decisions. Students must develop analytical insights that serve one or more cross functional stakeholder groups in Marketing, Product Development and Nutrition Policy. Each group has distinct priorities and decision-making contexts. This forces students to move beyond general data exploration toward insights that serve specific audiences.

Marketing teams require insights that support consumer communication strategies, competitive positioning and customer segment targeting. Product Development may focus on ingredient alternatives or new product concepts. Nutrition Policy teams may address regulatory compliance, transparency requirements and corporate wellness initiatives. This stakeholder diversity helps students understand that effective analysis requires consideration of organizational needs rather than just technical patterns. Focusing on the organizational context helps in developing students' analytical maturity. The data set used in the assignment allows for business relevant analysis that is stakeholder focused.

Dataset Design and Educational Value

The Starbucks beverage dataset was adapted from a Kaggle dataset on nutrition facts for Starbucks menu items. All nutritional information is standardized to 12oz serving sizes, enabling meaningful comparisons across beverage types and preparation methods. The dataset includes beverage category, name, preparation method, calories, fats, carbohydrates, protein, vitamins and caffeine content. This comprehensive nutritional data supports analysis from different stakeholder perspectives while providing a real-world business context for student learning. Appendix A also presents the dataset variables and sample data. An Excel file with the full dataset is also provided.

The data structure encourages students to explore relationships between customized beverage choices and nutritional outcomes, identify healthier alternatives within product categories and understand how ingredient substitutions affect overall nutritional profiles. The data fields support meaningful analysis while being simple and straight forward to students with little to no data analysis experience.

Students approach the dataset with varying skill levels and analytics experience. Some begin with personal relevance and consumer concerns, while others immediately identify business implications or technical patterns. This diversity provides rich material for student peer learning. It also suggests that effective analytical thinking can come from multiple starting points. Students work with this dataset through a structured four stage process that builds analytical and AI collaboration skills progressively. Although formal peer feedback was not incorporated in the current design, the variety of student approaches suggests potential for future iterations to

integrate group interaction as a means of enhancing collaborative learning.

Four Stage Learning Process

The assignment progresses through four stages, with each stage building on the last to guide students toward the development of meaningful analytical insights. Along the way, students are exposed to the analytics process with an AI collaboration aspect in a structured manner that is similar to professional practice.

Stage 1: Data Exploration and Problem Framing

The first stage requires students to conduct initial data exploration and identify an interesting analytical question(s) that is relevant to the stakeholder group(s). This stage emphasizes critical thinking about stakeholder needs rather than technical execution.

Stage 2: Dashboard Development and Insight Visualization

The second stage focuses on transforming the results of data analysis and exploration into clear, professional visualizations using a data visualization tool such as Tableau. Students demonstrate technical skills in dashboard creation and data storytelling skills while maintaining focus on stakeholder needs identified in Stage 1.

Students often begin with an "everything is important" approach, attempting to include all available data in their visualizations. However, the constraint of fitting effectively on one screen forces difficult choices on analytical focus and prioritizing key data results. This constraint encourages more intentional design choices and greater sensitivity and awareness of audience needs in data visualization.

Stage 3: AI Supported Analysis Refinement

The third stage introduces structured interaction with large language models to obtain feedback and refine both visualizations and analytical insights. Through the assignment, students develop the ability to critically evaluate AI generated suggestions, use emerging technologies responsibly and engage in collaborative improvement of their analytical work. Rather than substituting for human judgment, the interaction with AI supports deeper reflection. It increases students' awareness of how they communicate insights through data.

Stage 4: Reflection and Analytical Process Analysis

In the final stage, students are asked to reflect on their overall learning experience and assess how

AI contributed to their development of analytical thinking and insight generation. This stage develops metacognitive awareness about data analysis thinking and prepares students for ongoing learning in rapidly evolving technological environments.

This stage helps students reflect on their learning experiences that shape their growth as emerging analysts. Many described a meaningful shift in mindset. Rather than focusing solely on task completion, they began to appreciate the value of revision and iterative improvement as integral to the analytical process. This four-stage learning process requires careful timing and course integration to ensure students can apply their existing tool knowledge effectively.

Course Integration and Timeline

The assignment has been successfully implemented in two undergraduate introductory information systems course sections focused on business applications of technology. This assignment is most effective when scheduled after students have gained initial exposure to tools like Tableau and AI platforms earlier in the course. It then offers a meaningful opportunity to apply those foundational skills in an integrated business analytics scenario.

Table one summarizes the timeline for the four stages of the assignment along with the tasks, learning objectives and expected outcomes.

Week	Stage	Task	Learning Purpose	Expected Outcome
1	Part 1	Explore Starbucks dataset, identify a story	Learn problem framing, data exploration and question formulation	Clear research question and narrative focus
2-3	Part 2	Create Tableau dashboard	Practice visualization, narrative development and tool fluency	Clear dashboard aligned to a central insight or theme
4	Part 3	Revise dashboard and story using LLM feedback	Develop iteration skills and stakeholder framing awareness	Improve visual clarity and communication effectiveness
5	Part 4	Reflect on AI's role in shaping insights and revisions	Build metacognitive thinking and ethical awareness in AI use	Insightful written reflection and learning

Table One: Timeline for the four-part assignment

The modular design also enables instructors to emphasize either analytics fundamentals or AI

collaboration based on course objectives. A key component of successful implementation involves establishing clear guidelines for appropriate AI use throughout the assignment.

The structured use of AI in this assignment supports both academic integrity and responsible professional practice. Students were required to submit transcripts of their AI interactions, creating transparency and allowing the instructor to evaluate the purpose and quality of their AI collaboration.

Clear guidelines emphasized that AI could be used for feedback, alternative perspectives, and design suggestions, but that students remained responsible for all analytical decisions and interpretations. Class time was also used to discuss how AI can function as a collaborator rather than a shortcut.

To support this process, Appendix B provides a categorized prompt bank with example prompts to help students improve visual design, explore alternative approaches, and clarify their thinking. These resources encouraged students to use AI more intentionally by asking productive questions rather than relying on generic requests for help. Together, these guidelines reinforced the idea that AI should support, not replace, human analysis.

Assessment Approach

Student work was evaluated across four components including quality of analytical insights, effectiveness of dashboard design, extent of AI collaboration and depth of reflective analysis. This structure supports both technical skill development and critical thinking through deliverables at each stage.

Each stage was assessed using criteria aligned with the learning goals for that part of the assignment. The criteria and point values for all four deliverables are included in Appendix A. Taken together, these stage level criteria served as the grading rubric for the assignment as a whole. Stronger submissions showed clear insight development, thoughtful dashboard design, purposeful use of AI, and meaningful reflection on the learning process.

The reflection component played an important role in encouraging students to examine how their thinking evolved over time and how AI influenced their process. Most students showed clear awareness of this progression, making the reflection a valuable learning consolidation tool.

4. STUDENT LEARNING AND OUTCOMES

Student reflections offer valuable insights into how the assignment influenced learning across the four stages. Student voices reveal not only skill development but also shifts in analytical thinking, metacognitive awareness and professional readiness. These reflections show how the assignment develops both analytical and AI collaboration skills at the same time. Appendix D provides an example of one student's progression from Part 2 (i.e., initial dashboard development) through AI supported revision in Part 3. It shows the iterative improvement process and gradual development of business focused analytical thinking.

Analytical Growth and Problem Framing

Students showed meaningful growth in their approach to data analysis. Some moved from surface level pattern recognition toward more thoughtful problem framing and increased attention to stakeholder perspectives. Most students reported that the structured tasks helped them develop more focused analytical thinking. As one student noted, "I chose to analyze beverage types based on cholesterol instead of focusing on something broad like calories. That's a change I made after AI looked at my visualization." This progression from broad, unfocused exploration to targeted, business relevant analysis shows the assignment can help increase analytical maturity.

Design Thinking and Iterative Improvement

One of the more important learning outcomes was students' growing awareness of design and the value of revision. Many students had little experience revising analytical work in a structured way. The AI supported feedback process encouraged them to look again at their dashboards, rethink their design choices, and improve how they presented data. As a result, students began to see dashboard development as an iterative process rather than a onetime task.

The iterative nature of the assignment proved informative for students who previously approached analytical work as a onetime submission process. A particularly revealing reflection stated, "I usually submit and forget it, but using AI made me go back and really think through the story." This shift from completion oriented to revision-oriented thinking represents an important change in analytical mindset that extends beyond individual assignments to professional skill development.

Students also developed an appreciation for the complexity of effective visual communication. One student noted, "AI really helped me organize my thoughts better. It also gave advice on how to declutter my dashboard and make it clearer and easier by suggesting things like getting rid of grid lines and fixing titles." This recognition shows understanding of design principles and audience considerations.

AI Integration and Critical Evaluation

Students also developed thoughtful approaches to using AI feedback while maintaining control over their analytical decisions. Many students found AI most helpful for gaining fresh perspective on their work. One student noted, "The LLM gave a unbiased perspective from an outside source that allowed me to make my visualization overall improve." This shows students' ability to selectively apply suggestions while keeping ownership of their conclusions.

Another student reflected, "I was honestly shocked with the number of helpful suggestions it gave and how it seemed to understand what my goal for the dashboard was." Students also recognized AI limitations. One reflected, "Sometimes the AI suggestions were too generic for my specific stakeholder focus." This awareness shows students learned to evaluate AI feedback contextually rather than accepting all suggestions.

Professional Preparation and Metacognitive Awareness

The reflection component produced evidence of metacognitive development with students showing awareness of their own learning process and analytical development. Students often described how the experience changed their understanding of analytical work itself.

Students developed understanding of the iterative nature of analytical work and its broader applications. One student reflected, "This experience helped me learn how to break down a data set and to tell a story about a data set. It taught me how to tell a story using visualizations in a manner that is clean and concise."

Students also connected their learning experience to professional contexts and future applications. Many students described how the experience prepared them for professional analytical work. One reflection stated, "I was able to experience a sort of life as a business analyst and what the everyday job could involve." This awareness of professional roles represents preparation for

evolving workplace environments and expectations.

5. TEACHING NOTES AND IMPLEMENTATION GUIDANCE

Instructor Preparation and Setup

This assignment is intended to be implemented after students have already had introductory exposure to a data visualization tool such as Tableau and to current AI tools. Rather than functioning as a tutorial on Tableau itself, the activity is designed to help students apply previously learned visualization skills in a business analytics context while also learning how to collaborate thoughtfully with AI. In preparation for the assignment, instructors should provide the cleaned Starbucks dataset, ensure that students have access to a visualization platform and an approved AI tool, and share AI prompting guidelines. If students need additional support with Tableau, instructors may provide basic tutorial resources or cover those concepts earlier in the course. Institutional policies regarding AI use should also be reviewed in advance to ensure that all students can participate appropriately. Appendix C contains a quick implementation guide that can be used along with the guidance provided in this section.

Managing the Dual Learning Challenge

Students often find the combination of data analytics and AI collaboration overwhelming at first. Instructors can help by breaking the assignment into smaller parts and reminding students that experimentation, revision and learning from mistakes are expected. The business context helps reduce cognitive load by giving students a familiar, everyday setting. At the same time, students need regular reminders to focus on stakeholder needs and practical insight development rather than technical sophistication alone.

Supporting Student Development

Students often struggle to move from finding patterns in the data to developing business relevant insights. Instructors can support this process by providing examples of analytical questions that connect to stakeholder needs. It is also helpful to show students the difference between describing what the data shows and developing insights that can support decision making. Encouraging students to think about audience needs, rather than only focusing on the available data, can further strengthen their work.

AI Integration Best Practices

This support is especially important when students begin working with AI. Clear guidelines should explain that AI can be used for feedback, alternative perspectives and design suggestions, but that all analytical insights and final decisions must reflect the student's own thinking. The transcript requirement helps create transparency. It allows the instructor to evaluate the quality and purpose of the AI interaction, and gives students a way to reflect on how they used the tool. Instructors should introduce the AI prompt bank in Appendix B during the transition to Stage 3, since students often struggle at first to write effective prompts. Showing one or two examples from the prompt bank can help students see the difference between detailed questions and general requests for feedback.

Assessment and Feedback

Assessment should focus on student learning and progress rather than technical perfection alone. The reflection is especially useful because it provides insight into student learning that may not be fully visible in the technical deliverables. If time allows, peer feedback sessions can also be helpful, especially during dashboard review, because students can learn from seeing different approaches. Evaluation of AI collaboration should consider both the quality of student prompts and how thoughtfully students reviewed and used AI suggestions. Evidence of this can be seen in the AI transcript, the revised dashboard and the reflection.

6. CONCLUSION

This assignment addresses an important need in undergraduate business education by providing a practical means for introducing both data analytics fundamentals and AI collaboration skills to novice students. The four-stage assignment structure successfully helps apply technical skills while building the critical thinking abilities students need to do analytical work in AI integrated business environments.

The assignment works by combining a practical business context with structured technology interaction and reflective learning processes. Students work with real world business data while learning to collaborate with AI tools as analytical assistants rather than relying on them to replace human judgement or insight.

The focus on both analytics and AI collaboration responds to evolving workplace expectations where both skills are increasingly becoming a standard. Students enter the workforce better prepared for AI integrated environments

equipped with strong foundational skills in data analysis and insight communication.

The assignment contributes to information systems education by demonstrating that introductory students can successfully develop both analytical and AI collaboration capabilities when given proper structure and real business contexts. Future development could explore integration with advanced analytics concepts, expansion to team based analytical projects or adaptation for graduate level applications.

The assignment represents a timely response to the challenge of preparing students for data rich business environments while introducing them to appropriate AI collaboration. By providing students experience applying both skills in working with a dataset, the assignment offers a practical solution for IS instructors seeking to prepare students for the technological landscape they will encounter upon graduation.

7. REFERENCES

- Campos, C. R., Jacobini, O. R., Lorenzetti, M. L., Wodewotzki, M. L. L., & Ferreira, D. H. L. (2012). Social-political interfaces in teaching statistics. In C. Batanero, G. Burrill, & C. Reading (Eds.), *Proceedings of the Eighth International Conference on Teaching Statistics (ICOTS8)* (pp. [page numbers]). International Statistical Institute. http://cerme8.metu.edu.tr/wgpapers/WG10/WG10_Campos.pdf.
- Diamant, A. (2024). Introducing prescriptive and predictive analytics to MBA students with Microsoft Excel. *INFORMS Transactions on Education*, 24(2), 152-174. <https://doi.org/10.1287/ited.2023.0286>.
- Few, S. (2012). *Show me the numbers: Designing tables and graphs to enlighten* (2nd ed.). Analytics Press.
- Jamal, R., & Renzella, J. (2024). Enhancing Formative Feedback at Scale with the Intelligent Feedback Assistant. *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 2*. <https://doi.org/10.1145/3626253.3635482>.
- Joshi, R., Ghosh, S., Simileysky, A., & Bhanot, M. (2021). Structuring Formative Feedback in an Online Graphics Design Course in BME. *Biomedical Engineering Education*, 1, 325 - 333. <https://doi.org/10.1007/s43683-021-00046-z>.

- Knaflic, C. N. (2015). *Storytelling with data: A data visualization guide for business professionals*. John Wiley & Sons.
- Kolb, D. A. (2014). *Experiential learning: Experience as the source of learning and development*. FT press.
- Lyytinen, K., Topi, H., & Tang, J. (2023). MaCuDE IS task force phase II report: views of industry leaders on Big Data analytics and AI. *Communications of the Association for Information Systems*, 52(1), 429-464. <https://doi.org/10.17705/1CAIS.05217>.
- Mollick, E., & Mollick, L. (2023). Assigning AI: Seven approaches for students, with prompts. *arXiv preprint arXiv:2306.10052*, <https://doi.org/10.48550/arXiv.2306.10052>
- National Academies of Sciences, Engineering and Medicine. (2018). *Data science for undergraduates: Opportunities and options*. The National Academies Press. <https://doi.org/10.17226/25104>.
- Singla, A., Sukharevsky, A., Yee, L., Chui, M., & Hall, B. (2024, May). *The state of AI in early 2024: Gen AI adoption spikes and starts to generate value*. McKinsey & Company. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024>.
- Stratopoulos, T. C., & Vanden Bosch, N. (2019, December 12). A first-year data analytics course for accounting students. <https://doi.org/10.2139/ssrn.3618697>.
- U.S. Bureau of Labor Statistics. (2024). *Occupational outlook handbook: Data scientists*. <https://www.bls.gov/ooh/math/data-scientists.htm>.
- Van Allen, J., & Katz, S. (2023). Viewing open education within the Technological, Pedagogical, Content Framework: illustrating educator knowledge, skills and dispositions. *Research in Learning Technology*, 31. <https://doi.org/10.25304/rlt.v31.2829>.
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education – where are the educators? *International Journal of Educational Technology in Higher Education*, 16, Article 39. <https://doi.org/10.1186/s41239-019-0171-0>.

APPENDIX A

Student Assignment Instructions (Handout)

This appendix provides the complete student assignment description for the Starbucks Analytics Challenge described in the main manuscript. Instructors can use this material directly as a student handout or adapt it for their specific course context. The assignment implements the four-stage learning progression outlined in Section 3 of the main manuscript.

Data to Dashboard: A Starbucks Analytics Challenge

Background

You have just been hired as a junior analyst on Starbucks Corporation's Consumer Insights Team. The company leadership has expressed growing concerns about the nutritional content of their beverage menu, particularly regarding sugar, calories and caffeine levels across different drink categories and customizations.

A cross functional team comprising members from Marketing, Product Development and Nutrition Policy has requested clearer insights about the current beverage portfolio. They need data driven insights to support upcoming decisions about menu positioning, health focused marketing campaigns and/or potential product reformulations.

The Challenge

Your supervisor has tasked you with analyzing Starbucks' beverage nutritional data and creating a dashboard that communicates actionable insights to one or more of these internal stakeholders in the cross functional team.

Assignment Overview

Stage/Part	Deliverable	Timeline	Points
1	Data exploration summary	Week 1	10
2	Tableau dashboard	Weeks 2-3	20
3	AI revised dashboard	Week 4	30
4	Reflection essay	Week 5	10

Deliverables: Always use clear file naming conventions: LastName_FirstName_PartNumber_FileName

Stage 1: Data Exploration and Story Identification

Deliverable: One-page summary identifying your analytical focus
Timeline: Week 1; Points: 10

Explore the Starbucks dataset (more details provided later). The dataset includes details like calories, fat content, caffeine and sugar levels for various Starbucks drinks. Download the data and review all the fields (e.g., "Calories," "Total Fat," "Sugars," "Caffeine"). Spend a few minutes exploring the data and thinking about what aspects might be interesting to compare. You can use either Excel or Tableau to do this exploration.

What story do you want to tell members of the stakeholder team? Consider what patterns, comparisons or insights would be most valuable for business decision making. Write a brief summary of your observations from this initial data exploration (one paragraph) and include these insights in your submission.

Your submission should include:

- The specific question or story angle you will pursue
- Brief explanation of why this focus is relevant to the stakeholder team
- Any initial observations about the data that support your chosen direction
- Identification of your primary stakeholder audience (Marketing, Product Development and/or Nutrition Policy)

Assessment Criteria:

- Question Focus (3 pts): Clear, specific question with business relevance
- Stakeholder Relevance (3 pts): Demonstrates understanding of stakeholder needs and business context
- Data Observations (3 pts): Shows evidence of data exploration
- Communication Quality (1 pt): Professional writing and formatting

Stage 2: Visualization Development (i.e., Dashboard)

Deliverable: Tableau dashboard with 2-4 visualizations and insight summary

Timeline: Weeks 2-3; Points: 20

- a) Use Tableau to create between 2 and 4 visualizations (e.g., charts, graphs) that help communicate interesting insights from the data. Make sure that the visualizations fit on one screen, like a dashboard and that the text is legible. Come up with a main objective or theme for your dashboard and make it your title. What are you trying to show with your visualizations?
- b) What story are you trying to tell?
Write a paragraph explaining the insights or story your visualizations are trying to communicate. Why did you choose these particular aspects of the data to visualize?

Requirements:

- 2-4 related visualizations that support your central story
- One paragraph insight summary explaining your key findings
- Layout that fits on one screen

Assessment Criteria:

- Visualization Design and Insights (8 pts): Effective charts with meaningful insights
- Overall Layout and Presentation (4 pts): Professional, one screen design
- Storytelling and Explanation (6 pts): Clear narrative connecting visualizations
- Complete Submission (2 pts): All deliverables submitted as instructed

Submission Format:

- Tableau workbook file (.twbx)
- PDF export/screen shot of dashboard
- One paragraph insight summary (Word document or PDF)

Stage 3: AI Supported Revision

Deliverable: Revised dashboard, AI transcript and updated insights

Timeline: Week 4; Points: 30

Use a LLM (a large language model such as Chatgpt, Claude or Copilot) to obtain feedback on your dashboard and insights. Apply this feedback to improve your work.

The Process:

a) Ask the LLM for feedback - Use an AI tool to help you review or enhance your visualizations. Ask the LLM questions like:

“How can I improve my visualizations of Starbucks data?”

“What interesting trends can I show using this data?”

b) Record the conversation with the LLM - Copy and paste the LLM prompts you used and the responses you received. You will submit this along with your visuals and write up.

c) Refine your visuals and story - Based on the LLM’s feedback, revise your visualizations. Make sure the visualizations fit on one screen and that it has a main theme or objective (e.g., dashboard title and other relevant text) for your intended audience.

Write a paragraph explaining the insights or story your visualizations are trying to communicate. Why did you choose these particular aspects of the data to visualize?

Assessment Criteria:

- LLM Interaction and Documentation (9 pts): Evidence of thoughtful AI engagement with complete transcripts
- Revised Visualization Design and Insights (9 pts): Meaningful improvements based on feedback
- Storytelling Explanation (9 pts): Revised narrative with discussion of changes based on AI interaction
- Complete Submission (3 pts): All deliverables submitted as instructed

Submission Requirements:

- Original dashboard (Stage 2 version) – Screenshot or PDF
- Revised dashboard incorporating AI feedback - Tableau file and PDF/Screenshot
- Complete transcript of your AI interaction (prompts and responses)
- Updated insight paragraph reflecting your revisions

Stage 4: Reflection and Analysis

Deliverable: Reflection essay (500-750 words)

Timeline: Week 5; Points: 10

Reflect on your analytical process and the role of AI in shaping your work. Consider these questions:

- How did using AI feedback influence your thinking about the data story?
- What did you learn about effective data communication through this process?
- How did the iterative revision process affect the quality of your final insights?
- What are the benefits and limitations of using AI for analytical feedback?
- How might this experience apply to real world business analytics scenarios?

Assessment Criteria:

- Critical Reflection: (9 pts): Reflection on AI Use, learning insights and professional application
- Writing Quality (1 pt): Clear, professional writing within word count requirements

Data Set Information

You will work with an Excel dataset ([click here for link to data set](#)) containing nutritional information for Starbucks beverages. All nutritional information for drinks is for a 12oz serving size. Current research states that adults should consume no more than 30% of their calories in the form of fat, they need about 50 grams (women) or 63 grams (men) of protein daily and should provide for the remainder of their caloric intake with complex carbohydrates. One gram of fat contains 9 calories and carbohydrates and proteins contain 4 calories per gram. A "good" diet should also contain 20-35 grams of dietary fiber.

The data includes the following key variables:

Beverage category	Cholesterol (mg)
Beverage (Name)	Dietary Fiber (g)
Beverage prep	Sugars (g)
Calories	Protein (g)
Total Fat (g)	Vitamin A (% DV)
Trans Fat (g)	Vitamin C (% DV)
Saturated Fat (g)	Calcium (% DV)
Sodium (mg)	Iron (% DV)
Total Carbohydrates (g)	Caffeine (mg)

The following is 5 sample rows from the dataset.

Beverage Category	Beverage	Beverage Prep	Calories	Total Fat (g)	Trans Fat (g)	Saturated Fat (g)	Sodium (mg)	Total Carbohydrates (g)	Cholesterol (mg)	Dietary Fibre (g)	Sugars (g)	Protein (g)	Vitamin A (% DV)	Vitamin C (% DV)	Calcium (% DV)	Iron (% DV)	Caffeine (mg)
Coffee	Brewed Coffee	Short	3	0.1	0	0	0	5	0	0	0	0.3	0	0	0	0	175
Coffee	Brewed Coffee	Tall	4	0.1	0	0	0	10	0	0	0	0.5	0	0	0	0	260
Coffee	Brewed Coffee	Grande	5	0.1	0	0	0	10	0	0	0	1	0	0	0	0	330
Coffee	Brewed Coffee	Venti	5	0.1	0	0	0	10	0	0	0	1	0	0	0.02	0	410
Classic Espresso Drinks	Caff� Nonfat Latte	Short Milk	70	0.1	0.1	0	5	75	10	0	9	6	0.1	0	0.2	0	75

Stakeholder Profiles

Consider how your analysis might serve one or more of these audiences when crafting your story and recommendations:

Marketing Team: Interested in consumer facing messaging about healthier options, competitive positioning and customer segment targeting.

Product Development: Focused on reformulation opportunities, ingredient alternatives and new product concepts.

Nutrition Policy: Concerned with meeting regulatory requirements, transparency in nutritional information and supporting corporate wellness initiatives.

Academic Integrity Note

While you are encouraged to use AI tools for feedback and revision, all analytical thinking, interpretation and final insights must be your own work. Document all AI interactions transparently and critically evaluate suggestions rather than unconditionally accepting them. *This assignment is designed to simulate real world business analytics scenarios while developing your skills in data analysis and storytelling, iterative improvement and responsible AI use. Approach each stage thoughtfully and use the stakeholder context to guide your analytical choices.*

APPENDIX B AI Prompt Bank

Students may struggle with formulating useful prompts when engaging with large language models. This prompt bank provides categorized options to help students engage more effectively with AI tools during Stage 3 of the assignment. Instructors should introduce these prompts or discuss examples in class as students begin working on their AI supported revisions.

Students can adapt these prompts to their specific dashboard context and use them with any approved AI platform (Chatgpt, Claude, Copilot or other tools). The key is asking specific questions that provide context rather than generic requests for help.

Dashboard Design and Layout Feedback

- "I've created a dashboard about [your topic]. How can I improve the layout and visual clarity?"
 - "What are the strengths and weaknesses of this dashboard design?"
 - "My dashboard looks cluttered. What elements should I prioritize or remove?"
 - "Which of these design changes would have the biggest impact on clarity?"
- ### Storytelling and Narrative Development
- "I'm analyzing [your topic] for a [stakeholder type] team. Is my story compelling, or am I just showing data?"
 - "What's a clearer or more impactful title for this dashboard about [your topic]?"
 - "How can I better connect my visualizations to tell a unified story?"
 - "My insight is: [state your insight]. How can I make this more specific and actionable?"

Stakeholder Focused Revision

- "How should I revise this dashboard for Starbucks' marketing team focused on health-conscious messaging?"
 - "What would a Starbucks product development team want to see in this nutritional analysis?"
 - "Would a non-technical stakeholder understand my main message from this dashboard?"
 - "How can I make this dashboard more useful for [specific stakeholder type]?"
- ### Technical Improvement and Chart Selection
- "I'm comparing [variables] across [categories]. What chart types would best communicate this relationship?"
 - "My dashboard has [describe current charts]. Are these the most effective choices for my message?"
 - "How can I better use color, labels, or formatting to improve readability?"

Critical Evaluation and Reflection

- "What are the risks of relying too heavily on AI feedback for data analysis and interpretation?"
- "How should I balance AI suggestions with my own analytical thinking?"
- "What limitations might AI have in understanding my specific business context?"
- "Which AI suggestions improved my work and why?"

APPENDIX C

Quick Implementation Guide

Use this guide alongside the detailed student assignment description in Appendix A and the comprehensive implementation guidance in the main manuscript. This appendix focuses on key timing, common student challenges and practical considerations during assignment delivery.

Pre-Assignment Setup – Make sure to schedule the assignment across a 5-week period.

During the Assignment Week by Week

- Week 1: Introduce stakeholder framework; Stress the importance of business relevance over individual student personal interests in exploring the data set.
- Week 2-3: Remind students of the one screen constraint; Provide design feedback during work sessions.
- Week 4: Demonstrate effective AI prompting (use prompt bank – Appendix B); Remind students of the importance of evaluating AI suggestions prior to adoption.
- Week 5: Facilitate reflection discussion; connect to professional practice.

Common Student Questions

"What story should I tell?" = Guide toward stakeholder needs, not specific topics

"My dashboard is too cluttered" = Remind about audience focus and screen constraint

"Should I use this AI suggestion?" = Ask them to justify how it improves their analysis

"Is this insight good enough?" = Push for strong insights that are actionable and business relevant

Things to consider during implementation

- Grade Part 1 quickly for early feedback
- Focus on the progression of learning rather than technical perfection
- Look for evidence of critical thinking about AI suggestions
- Save strong examples (with permission) for future semesters

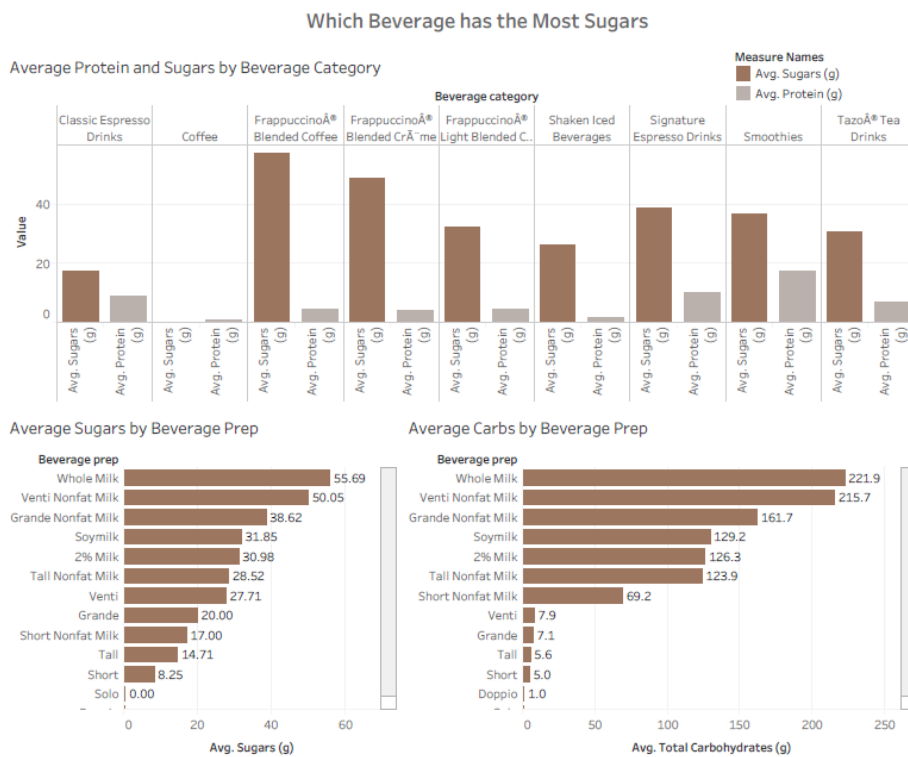
APPENDIX D Sample Student Artifact Summary

These artifacts are one student’s submissions for Stage 2 and Stage 3 of the assignment to present sample student work of how a student developed insights, applied AI feedback and revised their dashboards.

Part 2 - Initial narrative

In my dashboard the main insight I am trying to communicate is the great amounts of sugar and carbs related to the way the drink is prepared. This analysis is targeted at the Product Development team who could use this information to reformulate beverages with healthier options. Whole milk has on average nearly 100 more grams of carbs and over 20 more grams of sugar than the next option. I also looked at the side by side of average protein and sugars in different drink types. The proportion between the two is vastly different between the coffee drinks and the noncoffee beverages. I chose these visualizations to show some of the health comparisons in the options at Starbucks that might not always be as easy to look at.

Part 2 - Initial dashboard screen shot



Part 3 - AI Engagement

Prompt:

Look at this dashboard along with the data it is from and tell me how I can improve my visualizations of Starbucks data

Response:

Looking at the provided dashboard, here are some tips to improve the visualization of the Starbucks data to make it more insightful and visually appealing:

1. Clarify Labeling and Titles

- The titles could be more descriptive. For example, instead of "Average Protein and Sugars by Beverage Category," consider "Average Protein and Sugar Content by Starbucks Beverage Category (g)." This makes it clear what units are being used.
- Remove ambiguous abbreviations or label issues (like "Frappuccino® Blended Crème"). These characters likely come from encoding issues and should be corrected to ensure readability.

2. Improve Visual Hierarchy

- Use consistent color schemes across charts to represent categories. For instance, use one color palette for "Beverage Category" and another for "Beverage Prep." This will create a clear visual connection for the viewer.
 - Increase font sizes for the main titles and reduce clutter in the labels to help with readability.
3. Highlight Key Insights
- Use visual markers (like highlighting the bar in a different color) to draw attention to the drink with the highest sugar content or protein content in each chart.
 - Add annotations on bars to highlight significant data points, such as the highest or lowest averages.
4. Add Comparison Metrics or Additional Context
- Including percentage comparisons or benchmarks, such as how each category compares to an "average" Starbucks drink, can add more context.
 - For example, next to the "Average Sugars by Beverage Prep," you could add a reference line indicating the average sugar content across all beverages.
5. Refine Chart Types and Layout
- Consider using a heat map or a stacked bar chart for comparisons across categories or preparation types if the values are clustered.
 - For high-cardinality metrics like "Beverage Prep," a horizontal bar chart might enhance readability.
6. Improve Interactivity if Possible
- If using a tool like Tableau, consider adding interactive elements such as filters to toggle between metrics (e.g., sugars vs. protein vs. carbs).
 - Enable sorting to allow users to easily identify the top and bottom items within each chart.
- These enhancements can make the data clearer and help the audience quickly grasp the main insights about Starbucks beverages in terms of sugars, carbs, and protein.

Prompt:

What are interesting trends related to the relation of protein and sugar that I can show using this data

Response:

Here are some engaging trends that explore the relationship between protein and sugar in Starbucks beverages:

1. Low-Protein, High-Sugar Beverages

- Trend: Highlight beverages that are high in sugar but low in protein, which are typically more indulgent options that may have a lower satiety value.
- Insight: Drinks like Frappuccinos might fall into this category, where sugar levels are high, but protein content is low.
- Visualization Idea: Use a scatter plot where sugar content is on the x-axis and protein is on the y-axis. This would show a cluster of high-sugar, low-protein drinks in one section.

2. High-Protein, Moderate Sugar Options

- Trend: Identify beverages that provide a balance of protein and sugar, which may be perceived as more "nutritionally balanced" or suitable for those looking for a filling option.
- Insight: Certain smoothies or milk-based drinks (such as those with soy or whole milk) could provide a balance, with moderate sugar and higher protein.
- Visualization Idea: A bar or bubble chart comparing different drink categories, with size/color representing sugar levels and a secondary axis for protein content, highlighting the more balanced drinks.

3. Comparing Protein and Sugar by Beverage Preparation Type

- Trend: Show how different preparation types impact protein and sugar levels. For instance, beverages made with nonfat milk might have lower sugars and moderate protein, while those with whole milk or added flavors have higher sugar and protein levels.
- Insight: Preparation type significantly affects both protein and sugar, with certain choices increasing one nutrient more than the other.
- Visualization Idea: A grouped bar chart showing different preparation types along with their average protein and sugar content could reveal these differences.

4. Protein-to-Sugar Ratios Across Categories

- Trend: Calculate and compare the protein-to-sugar ratio to highlight the most "protein-rich" drinks relative to their sugar content.

- Insight: This could reveal that espresso-based drinks have a higher protein-to-sugar ratio, making them a preferred option for people who want protein without much added sugar.
 - Visualization Idea: Display the protein-to-sugar ratio across beverage categories, perhaps as a ranked list or sorted bar chart. This will let viewers quickly identify beverages with the best (and worst) protein-to-sugar ratios.
5. Impact of Size on Protein and Sugar Balance
- Trend: Explore how increasing beverage size affects the balance between sugar and protein. Larger sizes often have a disproportionate increase in sugar compared to protein.
 - Insight: This can highlight how the nutritional content scales up in a way that might make larger sizes less balanced in terms of protein-to-sugar ratios.
 - Visualization Idea: A line or step chart where each drink type's sugar and protein increase by size, showing how the ratio changes as the drink size increases (Short, Tall, Grande, Venti).
6. Milk Type's Influence on Protein and Sugar
- Trend: Show how different milk options (whole, nonfat, soy, etc.) affect the protein and sugar balance.
 - Insight: Non-dairy options like soy milk could increase both protein and sugar, while nonfat milk may increase sugar without contributing as much protein.
 - Visualization Idea: A side-by-side comparison of each milk type's average protein and sugar content, displayed as grouped bars or a small multiples chart for easy comparison.
7. Protein and Sugar in 'Healthier' vs. 'Indulgent' Drinks
- Trend: Create a comparison of "healthier" options (like teas and black coffee) versus "indulgent" options (like Frappuccinos) to see how protein and sugar content vary across these two ends of the spectrum.
 - Insight: Healthier options generally have lower sugar and protein, while indulgent ones may have higher sugar with varied protein levels depending on milk content.
 - Visualization Idea: A dual-axis chart where one axis is sugar content and the other is protein, showing "healthier" drinks clustered near the origin and indulgent drinks toward higher sugar values. These trends offer insights into how protein and sugar interplay in Starbucks drinks, highlighting balance or lack thereof for viewers with specific dietary goals.

Prompt:

Give me a strong title for my dashboard

Response:

Here are some compelling title ideas for your Starbucks dashboard:

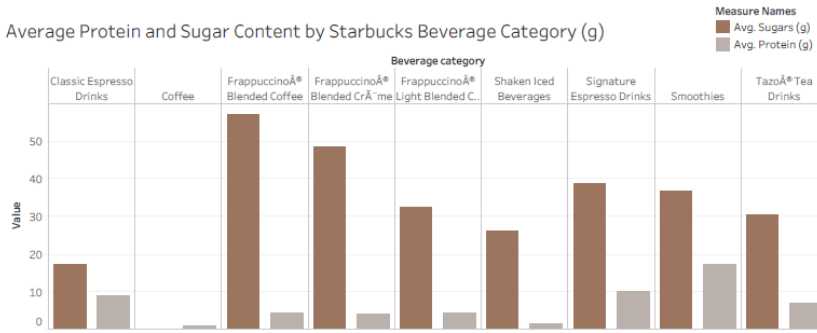
1. "Unpacking Starbucks: A Nutritional Breakdown of Your Favorite Drinks"
2. "Inside the Cup: Nutritional Insights on Starbucks Beverages"
3. "Sips & Stats: A Deep Dive into Starbucks Drink Nutrition"
4. "From Frappuccinos to Teas: Understanding the Nutrition Behind Starbucks Beverages"
5. "Calories, Sugars, and More: The Nutritional Profile of Starbucks Drinks"

Each title aims to capture the attention of viewers by hinting at the deeper insights and comparisons within the data. Let me know if you'd like any variations or further focus!

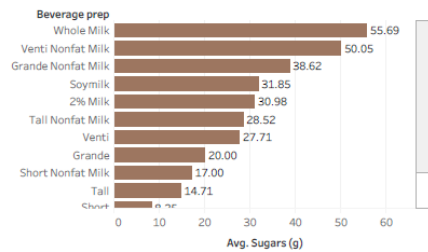
Part 3 - Revised Dashboard Screenshot

Inside the Cup: Nutritional Insights on Starbucks Beverages

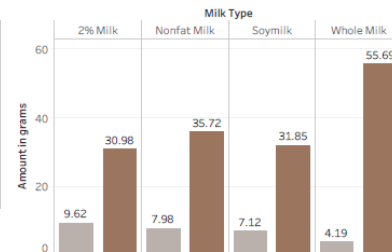
Average Protein and Sugar Content by Starbucks Beverage Category (g)



Average Sugar Content by Starbucks Beverage Prep (g)



Average Protein and Sugar Content by Milk Type



Part 3 - Revised narrative

AI helped me realize to direct my focus towards one thing. I decided that I wanted that story to be the relation between average protein and sugars in the drinks. This insight is particularly valuable for the Marketing team because it shows which drinks offer better nutritional balance for health conscious customers. I believe that this is a very important insight and it gave me the idea to incorporate milk type into this which provided an extra insight. The Marketing team could use this information to promote drinks with better protein to sugar ratios or highlight milk alternatives that provide more balanced nutrition. I think this is a stronger story and helps get my original stance about looking into the nutritional values across better.

Teaching with Generative AI: Ethics and Analytical Confidence in Undergraduate IS Education

Asish Satpathy
Asish.Satpathy@asu.edu

Jayaram Balachander
Jayaram.Balachander@asu.edu

W. P. Carey School of Business,
Department of Information Systems
Arizona State University
Tempe, AZ, USA

Abstract

This teaching tip presents a scalable model for integrating Generative AI (GenAI) and ethical reasoning into WPC 300 (Problem Solving & Actionable Analytics), a required core course taken by all undergraduate business students at Arizona State University. By embedding responsible GenAI use within established course modules, the intervention aims to strengthen students' analytical confidence, deepen their understanding of ethical issues in data analysis, and prepare them for real-world business challenges. The model also supports instructors in delivering relevant, industry-aligned skills.

The course redesign introduced GenAI tools, including OpenAI's ChatGPT and Google Notebook, alongside core analytical methods such as regression analysis and data visualization. These tools were selected for their user-friendly interfaces and their ability to extend analytical capabilities beyond traditional Excel-based methods, supporting exploratory analysis, pattern recognition, and interpretation. Ethical considerations were scaffolded through a dedicated "Mindful AI" module emphasizing fairness, accountability, and societal implications. A mixed-methods evaluation, combining paired t-tests with sentiment and thematic analysis of student reflections, demonstrated statistically significant gains in students' analytical confidence and understanding of ethical issues. Qualitative feedback further indicated increased appreciation for AI-assisted learning and ethical reflection.

This model serves as a practical guide for educators aiming to modernize analytics instruction while maintaining core learning objectives. It provides students with a pathway to greater analytical confidence and ethical awareness and supports instructors by enabling the staged introduction of AI tools within existing course materials, reinforced by concrete applied examples.

Keywords: GenAI, Business Analytics, Teaching Innovation, AI Ethics, Undergraduate Education, Excel, Data Analytics

Recommended Citation: Satpathy, A., Balachander, J., (2026). Teaching with Generative AI: Ethics and Analytical Confidence in Undergraduate IS Education. *Information Systems Education Journal*, v24(n6), pp 63-74. DOI# <https://doi.org/10.62273/ZYUC4609>

Teaching with Generative AI: Ethics and Analytical Confidence in Undergraduate IS Education

Asish Satpathy and Jayaram Balachander

1. INTRODUCTION

The growing demand for analytical competencies in business underscores the urgent need for educational innovation that aligns with the evolving workforce driven by artificial intelligence (AI). In response, the W. P. Carey School of Business at Arizona State University (ASU) developed an undergraduate course, Problem Solving & Actionable Analytics (WPC 300), an introductory analytics course uniquely positioned as a foundational element for undergraduate business students across diverse disciplines (Darcy & Satpathy, 2023). Initially implemented to meet broad industry requirements, WPC 300 emphasizes developing problem-solving and analytical skills within a hybrid, online learning model that integrates synchronous and asynchronous teaching methodologies. The course uniquely combines practical applications, rigorous assessments based on Bloom's taxonomy, and analytical frameworks inspired by the modified CRISP-DM (Cross Industry Standard Process for Data Mining) model to maximize student engagement and learning outcomes.

The importance of analytical skills in business education has been extensively documented, highlighting a growing gap between industry demands and workforce preparedness (Deloitte & Quinnipiac University, 2025). Despite its robust framework, traditional analytics education, which relies predominantly on Excel and established methods, faces significant limitations, including reduced student engagement and difficulty in connecting theoretical skills with real-world applications (Frost et al., 2021). Recent studies emphasize that traditional teaching methods may inadequately prepare students for emerging analytical challenges in an AI-driven economy (Weng et al., 2024).

Furthermore, the integration of Generative AI (GenAI) tools has been increasingly recognized for its transformative potential in education, particularly in analytics courses (Gastineau, 2024; Sekli et al., 2024; Wahba et al., 2024; Satpathy & Balachander, 2025). These AI tools extend analytical capabilities beyond traditional methods, offering new opportunities for predictive analytics, pattern recognition, and

deeper insights. However, their widespread adoption also introduces ethical complexities that must be addressed within educational contexts (Cotton et al., 2023; Garcia-Lopez and Trujillo-Linan, 2025; Sywelem et al., 2024). The literature emphasizes the importance of integrating ethical discussions into AI-driven curricula to ensure students develop a comprehensive understanding of responsible AI practices, encompassing fairness, accountability, and transparency (Britchenko, 2025; Luckin & Holmes, 2016; Shirzad & Darazi, 2025).

To address these gaps, this teaching tip outlines how to integrate GenAI tools and ethical frameworks into the WPC 300 course. The goal is to bridge traditional analytics methods with innovative AI-driven techniques, thereby supporting reflective learning and enhancing practical competencies, and preparing students ethically and technically for the AI-driven workforce. Specifically, we illustrate the approach of blending traditional analytical tools with beginner-friendly AI platforms, such as ChatGPT and Google Notebook, alongside structured ethical discussions. By detailing practical implementation methods and sharing insights from pre- and post-implementation assessments, we offer a scalable model for educators to replicate and adapt across similar analytics courses.

This teaching tip contributes to the growing literature on AI in education by offering a scalable and discipline-agnostic model that aligns with evolving workforce demands. We provide a detailed account of the course implementation, assessment methods, and student learning outcomes, both quantitative and qualitative, to support adoption in similar instructional contexts.

2. ASSIGNMENT & COURSE IMPLEMENTATION

Each module in the WPC 300 course is designed to build progressively, beginning with foundational problem identification and disaggregation (Modules 1-3), analytical execution (Modules 4-5), data storytelling (Module 6), and culminating in ethical reflection and synthesis (Module 7). This modular structure

provided a natural scaffold for integrating Generative AI and ethics into the curriculum, particularly in Modules 5 through 7, where regression analysis, visualization, and ethical reasoning align closely with emerging AI applications. The table (Appendix A: Table 1) summarizes the course flow and highlights where GenAI and ethics enhancements were embedded.

The integration of GenAI and ethical considerations in the WPC 300 course followed a structured, phased approach to ensure a seamless educational experience for students. This began with an initial pre-implementation survey, strategically administered immediately following the due date for Module 4 assignments. This survey was designed to establish baseline data on student engagement levels, technical competencies, and ethical awareness, and remained open for student responses for over one week.

Following the initial survey, new instructional content was introduced in Module 5, extending foundational analytics to include Generative AI-supported interpretation and evaluation. Students first performed correlation, ANOVA, and regression analysis using standard analytical procedures, and then used GenAI tools to generate explanations of their results and propose business insights. They were required to critically evaluate these AI-generated outputs by comparing them with their own analyses, assessing accuracy, identifying unsupported conclusions, and recognizing missing assumptions. This two-stage approach emphasized interpretation and analytical judgment while preserving core statistical skills. In addition, students explored core AI concepts and practical applications through guided videos and assignments, including the assisted visualizations for data storytelling. These activities accounted for 3% of the course grade.

Module 6 further strengthened visual analytics skills through focused instruction on data storytelling. Module 7 introduced the concept of Mindful AI, emphasizing accountability, fairness, transparency, and bias mitigation in AI-supported analytics. In this module, students examined GenAI outputs through these lenses and discussed how data, model design, and prompt choices can introduce or reduce harm. A representative activity asked students to analyze a short business case (e.g., AI-assisted customer prioritization), generate an AI-based recommendation, and then respond to guided questions about who benefits or is disadvantaged,

where bias might arise, and what transparency and accountability measures would be needed. This module contributed 2% to the course grade. A post-implementation survey was conducted to assess changes in engagement, technical skills, and ethical awareness, mirroring the initial survey. To accommodate varying participation levels, a supplementary assignment offered 3% extra credit to students who were unable to participate in the experiment; it involved a business case study that required students to interpret data visualizations through a 10-question quiz on the Canvas LMS.

3. METHODS

Before conducting statistical analysis, all survey response data were cleaned and preprocessed to ensure consistency and accuracy. After all cleaning and preprocessing steps, a total of 178 student responses were selected for further analysis.

Survey Instruments and Data Sample

Two matched surveys, administered pre- and post-intervention, captured data on student confidence, tool familiarity, engagement, and ethical awareness. Each survey included ten Likert-scale questions and five open-ended reflection questions. All closed-ended questions used a 4-point Likert scale (1 = strongly disagree to 4 = strongly agree) without a neutral midpoint. This format was chosen to reduce central tendency bias, simplify response demands for a large introductory cohort, and provide sufficient sensitivity to detect pre-post changes in broad constructs such as confidence and ethical awareness. After data cleaning and removing incomplete entries, 178 paired responses (same student pre-post) out of 230 were retained for analysis. Students represented a broad mix of business majors (marketing, finance, supply chain, etc.). Before the GenAI and ethics integration, most students reported low familiarity with AI tools and only slight to moderate confidence in using Excel and data analysis techniques. Engagement levels were modest, though students widely recognized the relevance of ethics in analytics. Comfort with and willingness to use AI tools varied, indicating uncertainty and limited prior exposure. These responses underscored the need for structured AI training and ethical framing in the later course modules.

Quantitative Analysis

To evaluate changes in key learning outcomes, we conducted paired-sample t-tests on matched pre- and post-course responses. Because responses were paired for the same students before and after the intervention, the paired t-test was appropriate for evaluating within-subject changes. Assumptions of normality were assessed using Q-Q plots. Survey items used a standardized four-point Likert scale to support valid comparison. The internal consistency of the Likert-scale items was acceptable (Cronbach's $\alpha = 0.64$), supporting the instrument's reliability (Cohen et al., 2002). The Cronbach's alpha is appropriate given that the items were intentionally designed to span multiple domains, such as engagement, ethics, familiarity, and confidence, so a somewhat lower coefficient is expected for this multi-dimensional construct. The t-tests examined changes in confidence levels regarding data analysis skills, familiarity with Excel, engagement in course activities, understanding of ethical concerns in data analysis, and comfort in using generative AI tools. Statistically significant improvements were observed (Appendix A: Table 2 & Appendix B: Figure 1) in students' confidence in applying data analysis skills and their understanding of ethical concerns, as indicated by p-values below 0.05. Other test items did not show significant changes.

Qualitative Analysis

Open-ended responses were analyzed thematically and supplemented by sentiment analysis. Sentiment polarity was calculated using VADER, a rule-based model optimized for short social text, to measure shifts in student tone before and after the intervention (Hutto & Gilbert, 2014).

This comprehensive approach provided both depth and breadth in assessing the pedagogical impact of GenAI and the integration of ethics, allowing triangulation of student performance, perceptions, and engagement.

4. EVIDENCE OF EFFECTIVENESS

To evaluate the effectiveness of integrating AI and ethics in WPC 300, we designed a mixed-methods assessment framework centered on pre- and post-intervention surveys (Kilinc, S., 2024). These surveys measured changes in student engagement, analytical confidence, familiarity with AI and Excel-based tools, and ethical understanding. The pre-intervention survey, launched at the end of Module 4, consisted of ten multiple-choice questions and five open-ended questions. It established baseline metrics for

student familiarity with data analytics tools, frequency of use of advanced Excel features, engagement in analytics coursework, exposure to AI tools, and understanding of ethical issues, including bias and data privacy. Students were also asked to reflect on their learning expectations, anticipated challenges with AI tools, and how ethics and AI could enhance their understanding of business analytics.

The post-intervention survey, deployed after the completion of Module 7, mirrored the structure of the initial study, focusing on changes observed during the intervention. It captured students' reflections on the perceived usefulness of GenAI tools, their confidence in applying statistical analysis, their comfort with AI technologies, and the impact of ethics discussions on their critical thinking. The open-ended questions encouraged students to articulate how their perspectives on AI evolved, the benefits of ethical training, and suggestions for improving future iterations of the course. By comparing the two surveys, we gathered both qualitative and quantitative evidence to assess the pedagogical impact of incorporating GenAI and ethics into a foundational analytics course.

5. RESULTS

The paired t-test results reveal statistically significant improvements in students' understanding of ethical concerns in data analysis and a statistically significant, though nuanced, change in confidence in applying data analysis skills following the intervention (Figure 1). Although Figure 1 shows a modest decline in self-reported confidence among some students immediately after the intervention, this may reflect recalibration rather than an actual reduction in skill. As students moved from procedural, Excel-based work to judgment-intensive, AI-assisted analysis that foregrounded uncertainty and ethical risk, they were exposed to the limits of their prior understanding. This pattern is consistent with research on self-efficacy calibration and related work on the Dunning-Kruger effect, which suggests that deeper exposure to complex tasks can temporarily reduce perceived competence as learners develop more realistic self-assessments (Dunning, 2011; Jansen et al., 2021). Accordingly, the observed dip is interpreted as a possible indicator of more accurate self-evaluation, not necessarily diminished learning. Instructors can view this short-term decline in confidence as a healthy learning signal instead of a setback.

Other measured outcomes, including familiarity with Excel, engagement in course activities, and comfort using generative AI, did not show statistically significant changes (Wang & Guo, 2025). This stability is consistent with students' relatively high baseline familiarity and engagement, the short duration of the intervention, and the modest grading weight assigned to the GenAI activities. Overall, these results suggest that the observed gains were concentrated in learning outcomes most directly and intensively targeted by the AI- and ethics-focused instructional components. To provide deeper context for these findings, we next examine qualitative student reflections to understand how perceptions and attitudes evolved during the intervention.

Qualitative Analysis of Student Reflections

To complement and contextualize the quantitative findings, we conducted a qualitative analysis of student reflections collected before and after the intervention. Text responses were analyzed for sentiment and topic frequency to identify shifts in student attitudes and conceptual focus.

Pre-intervention reflections often expressed skepticism or uncertainty about using AI tools, highlighting concerns about data accuracy, potential misuse, and a lack of familiarity with them. The sentiment analysis of these reflections indicated a predominantly neutral to slightly negative tone, emphasizing difference (Lāma & Lastovska, 2025).

Post-intervention reflections showed a clear shift toward more positive sentiment and deeper understanding. As shown in Appendix B: Figure 3, the proportion of positive responses increased across all categories, particularly in students' comfort with AI tools and their perceptions of ethics in analytics. This shift is further supported by the thematic patterns observed in student reflections, where more specific and action-oriented language, such as "insights," "real-world," "explained," and "projects", replaced earlier vague or uncertain expressions (Appendix B: Figure 2). While word clouds provide a high-level visualization of frequently occurring terms, they are used here as a supplementary tool; the primary insights are derived from thematic coding of student responses. Students consistently highlighted the value of GenAI tools for data analysis and visualization, as well as awareness of ethical considerations in AI use. This overall shift in sentiment and thematic depth, also reflected in Appendix B: Figure 4, suggests not only improved technical confidence but also

stronger ethical awareness, reinforcing the impact of the integrated instructional approach (Balart et al., 2026).

Appendix A: Table 3 summarizes qualitative findings from open-ended student responses reflecting on their experiences with generative AI tools and ethics instruction in the course. Each row presents the overall sentiment polarity, key themes, and representative quotes for specific prompts.

The analysis reveals a predominantly positive perception of integrating AI and ethics, with students appreciating AI as a learning aid, recognizing the importance of ethical thinking, and valuing the combined use of AI and Excel. Some challenges were noted, particularly around prompt formulation and the reliability of AI outputs. The strongest positive sentiment was expressed in student suggestions for improving future courses, emphasizing the value of practical, ethically grounded examples.

6. DISCUSSION

The integration of Generative AI and ethical reasoning into WPC 300 produced pedagogical benefits, addressing key limitations of traditional analytics instruction. Students entered the course with limited exposure to real-world data tools and minimal understanding of how GenAI intersects with business decision-making. While Excel has served as a valuable foundation for teaching, it often fails to inspire broader conceptual thinking or address the nuanced demands of modern data literacy. By embedding GenAI tools and ethical frameworks into key modules, the course strengthened students' analytical confidence and supported more reflective, ethically informed approaches to real-world data analysis. Students reported increased appreciation for data interpretation, visualization, and the ethical implications of technology use, consistent with AACSB learning goals related to responsible and analytical thinking.

Several factors contributed to this model's success. First, its placement within the existing modular framework of WPC 300 allowed for seamless curricular integration. We infused AI into core analytical practices, such as regression analysis, visualization, and statistical interpretation, making its relevance immediately apparent. Second, the phased introduction of GenAI and ethics ensured that students had time to build foundational skills before applying them through AI-enhanced tools. This scaffolded

approach reduced cognitive overload and eased students into unfamiliar territory.

However, implementation was not without challenges. Some students initially expressed resistance to using GenAI tools, citing concerns about reliability, complexity, and academic integrity. To mitigate these concerns, the instructional team emphasized transparency around tool capabilities and limitations and positioned AI as a thinking partner, not a replacement, for critical thinking. Technological issues, such as occasional limitations in access to specific platforms (e.g., Google Notebook), were addressed through alternative resources and on-demand support. The ethical modules encouraged students to articulate these concerns and constructively reflect on them.

The observed gains in ethical awareness, supported by qualitative indicators of reflective engagement, underscore the value of explicitly incorporating GenAI into analytics education, but perhaps more importantly, doing so through a lens of responsibility and critical inquiry. The qualitative and quantitative evidence demonstrate that this model is not only pedagogically sound but also scalable and adaptable. Other business analytics courses, especially those with applied focuses, can draw on this framework to integrate emerging technologies without compromising core learning objectives. Future iterations could benefit from even more interactive AI assignments, integration with industry datasets, and reflective exercises that require students to confront ethical dilemmas using real-world business contexts.

Ultimately, this teaching innovation points to a broader shift in how we prepare students for an AI-mediated future. By weaving technical fluency with ethical literacy, we create better analysts and more thoughtful, socially responsible professionals.

This study has several limitations. First, it was conducted at a single U.S. university and focused on business majors, so the results may not generalize to other institutions or disciplines. Second, participation in the pre- and post-surveys was voluntary, which may introduce response bias if more motivated students were more likely to respond. Third, the course relied on access to specific online tools that may not be available in all settings. Even with these constraints, the overall approach, adding new analytical activities and ethics discussions into existing modules, can be adapted in non-business or non-U.S. programs by using locally available

tools and tailoring the cases and examples to their own students and contexts.

7. ACTIONABLE TIPS

The statistically significant gains in analytical confidence and ethical awareness, supported by strong qualitative feedback, affirm the value of integrating Generative AI tools into data analytics education. To further enhance effectiveness and scalability, we recommend the following instructional strategies:

- Start with guided AI prompting tutorials to build foundational prompting skills and reduce initial confusion.
- Use real-world, imperfect datasets to encourage critical evaluation of AI-generated outputs and deepen ethical reasoning.
- Retain structured analytical workflows (e.g., regression in spreadsheets) before introducing AI-assisted interpretation to preserve foundational skills.
- Design two-stage assignments in which students first perform analysis independently and then compare their results with AI-generated interpretations.
- Incorporate case-based activities that require students to evaluate AI recommendations in terms of fairness, bias, and stakeholder impact.
- Embed short reflection prompts that ask students to assess when AI is helpful, when it fails, and how it should be used responsibly.
- Gradually shift from AI-assisted support toward critical evaluation, where students are expected to question and validate AI outputs.
- Provide clear rubrics that reward reasoning, interpretation, and critique of AI outputs, not just correct answers.
- Use low-stakes activities early in the course to reduce student anxiety and build confidence in using AI tools.
- Offer example prompts and model responses to establish expectations for both analytical quality and ethical reasoning.

The model is designed to be resource-light. Most activities are embedded into existing modules and require roughly 1–2 extra hours of preparation per week during the first implementation. It relies on widely available GenAI tools (or free/educational tiers), the LMS, and Excel, but instructors should align their use with institutional AI and academic integrity policies. To mitigate

access inequities, we recommend providing lab or shared-device access and offering instructor-generated AI outputs as an alternative when students cannot use specific platforms outside class.

8. CONCLUSION

As artificial intelligence continues to transform the business landscape, analytics education must evolve to equip students not only with technical proficiency but also with ethical discernment. The WPC 300 course redesign demonstrates a practical and scalable model for integrating Generative AI and ethics into an undergraduate business curriculum. By embedding AI tools and ethical reasoning within existing analytical modules, the intervention maintained instructional continuity while enhancing student engagement, confidence, and critical thinking.

One of the strengths of this approach lies in its replicability. The modular integration strategy, use of accessible AI platforms like ChatGPT, and alignment with standard learning outcomes make it feasible for instructors at other institutions to adopt or adapt this model without requiring major curricular overhauls. Additionally, the inclusion of guided tutorials, real-world datasets, and structured reflection offers a flexible blueprint that can be tailored to various course formats, learning levels, or delivery modes.

Looking ahead, AI will play an increasingly central role in shaping analytics education. As technologies advance, teaching methods must evolve to help students use AI effectively while critically interpreting and applying it. Embedding AI literacy and ethical reasoning into core courses prepares students to become thoughtful, capable professionals ready to lead in an AI-driven world.

9. ACKNOWLEDGMENTS

This study was supported by the Research on Teaching and Learning Grant Program at the W. P. Carey School of Business, Arizona State University.

10. REFERENCES

Balart, T., Díaz, B., & Shryock, K. (2026). A Systematic Literature Review on the Pedagogical Implications and Impact of GenAI on Students' Critical Thinking. *Algorithms*, 19(3), 179.

Britchenko, I. (2025). Ethical AI in Education: Principles, Governance, and Responsible

Implementation. *Pedagogy and Education Management Review*, (4(22)), 17–30. <https://doi.org/10.36690/2733-2039-2025-4-17-30>

Cohen, L., Manion, L., & Morrison, K. (2002). *Research methods in education*. Routledge. <https://doi.org/10.4324/9780203224342>

Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 61(2), 228–239. <https://doi.org/10.1080/14703297.2023.2190148>

Darcy, D. P., & Satpathy, A. (2023). Teaching Tip: A Scalable Hybrid Introductory Analytics Course. *Journal of Information Systems Education*, 34(4), 360–369.

Deloitte & Quinnipiac University. (2025). Closing the skills gap with dynamic partnerships. AACSB Insights. <https://www.aacsb.edu/insights/reports/2025/closing-the-skills-gap-with-dynamic-partnerships>

Dunning, D. (2011). The Dunning–Kruger effect: On being ignorant of one's own ignorance. In *Advances in experimental social psychology* (Vol. 44, pp. 247–296). Academic Press.

Frost, R., Matta, V., & Kenyp, L. (2021). A System to Automate Scaffolding and Formative Assessment While Preventing Plagiarism: Enhancing Learning in IS and Analytics Courses That Use Excel. *Journal of Information Systems Education*, 32(4), 228–243.

García-López IM and Trujillo-Liñán L (2025) Ethical and regulatory challenges of Generative AI in education: a systematic review. *Front. Educ.* 10:1565938. doi: 10.3389/educ.2025.1565938

Gastineau, J. (2024). Exploring the Impact of Generative AI Prompt Engineering in Higher Education: A Study in Undergraduate and Graduate Business Analytics Courses (Doctoral dissertation, University of Arkansas).

- Gupta, B., Goul, M., & Dinter, B. (2015). Business Intelligence and Big Data in Higher Education: Status of a Multi-Year Model Curriculum Development Effort for Business School Undergraduates, MS Graduates, and MBAs. *Communications of the Association for Information Systems*, 36(23). <https://doi.org/10.17705/1CAIS.03623>
- Hutto, C., & Gilbert, E. (2014, May). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAAI Conference on Web and Social Media (Vol. 8, No. 1, pp. 216-225)*.
- Jansen, R. A., Rafferty, A. N., & Griffiths, T. L. (2021). A rational model of the Dunning-Kruger effect supports insensitivity to evidence in low performers. *Nature Human Behaviour*, 5(6), 756-763.
- Kılınc, S. (2024). Comprehensive AI assessment framework: Enhancing educational evaluation with ethical AI integration. *Journal of Educational Technology and Online Learning*, 7(4-ICETOL 2024 Special Issue), 521-540. <https://doi.org/10.31681/jetol.1492695>
- Lāma, G., & Lastovska, A. (2025). AI competence and sentiment: a mixed-methods study of attitudes and open-ended reflections. *Frontiers in Artificial Intelligence*, 8, 1658791.
- Luckin, R., & Holmes, W. (2016). *Intelligence Unleashed: An Argument for AI in Education*. Pearson.
- Manyika, J. (2011). *Big data: The next frontier for innovation, competition, and productivity*. McKinsey Global Institute, 1.
- Satpathy, A., & Balachander, J. (2025). Future Workforce Evolution - Impact of Artificial Intelligence Across Industries. *Journal of Information Systems Applied Research and Analytics*. v18, n3, pp 44-60. DOI# <https://doi.org/10.62273/DHUV7230>
- Sekli, G. M., Godo, A., & Véliz, J. C. (2024). Generative AI solutions for faculty and students: A review of literature and roadmap for future research. *Journal of Information Technology Education: Research*, 23, 014. <https://doi.org/10.28945/5304>
- Shirzad, S., & Darazi, M. A. (2025). AI literacy in education: Balancing innovation, ethics, and equity in the digital age. *Journal of New Trends in English Language Learning*, 4 (Special Issue), Article si-02. <https://doi.org/10.57647/JNTELL.2025.si-02>
- Sywelem, M. M. G., & Mahklouf, A. M. E. S. (2024, June). Ethical considerations in the integration of Artificial Intelligence in education: An overview. In *CS & IT Conference Proceedings (Vol. 14, No. 12)*. *CS & IT Conference Proceedings*. <https://doi.org/10.5121/csit.2024.141201>
- Wahba, F., Ajlouni, A. O., & Abumosa, M. A. (2024). The impact of ChatGPT-based learning statistics on undergraduates' statistical reasoning and attitudes toward statistics. *Eurasia Journal of Mathematics, Science and Technology Education*, 20(7), em2468. <https://doi.org/10.29333/ejmste/14726>
- Wang, K., & Guo, Z. (2025). Can Learners' Use of GenAI Enhance Learning Engagement?—A Meta-Analysis. *Education Sciences*, 15(12), 1578.
- Weng, X., Qi, X. I. A., Gu, M., Rajaram, K., & Chiu, T. K. (2024). Assessment and Learning Outcomes for Generative AI in Higher Education: A Scoping Review on Current Research Status and Trends. *Australasian Journal of Educational Technology*. <https://doi.org/10.14742/ajet.9540>

APPENDIX A

Module	Title	Core Focus	GenAI Integration	Ethics Integration
1	Introduction to Problem Solving	Problem-solving frameworks, Excel basics	None	None
2	Define the Problem	Refining problem statements, Pivot Tables	None	None
3	Disaggregation & Prioritization	Logic trees, EDA, probability	None	None
4	Conduct Analysis	T-tests, ANOVA, root cause analysis	Pre-implementation survey administered	None
5	Additional Analysis	Regression, model building	AI for regression, EDA using ChatGPT & Google Notebook	Introduced a reflection on AI usage
6	Data Visualization & Storytelling	Dashboarding, storytelling, and advanced Excel visuals	AI-generated insights & visualizations	None
7	Mindful AI & Responsible Use	Bias, fairness, transparency, and ethics frameworks	AI + ethics reflections Post-implementation survey administered	Core ethical discussion module

Table 1: Course Module Overview with GenAI and Ethics Integration

Test Question	t-Statistic	p-Value	Statistical Significance (p-value < 0.05)
Confidence in applying data analysis skills to real-world problems	-2.1095	0.0363	Yes
Familiarity with Excel	0.0767	0.9389	No
Engagement in course activities	-0.3068	0.7598	No
Confidence in understanding ethical concerns in data analysis	6.3261	0	Yes
Comfort level in using generative AI	0.6382	0.5248	No

Table 2: Paired t-Test results on student learning outcomes before and after GenAI and Ethics instruction

Question	Overall Sentiment	Mean Polarity	Key Themes	Representative Quote
Q11: Perspective on using Generative AI tools after course completion	Positive	0.214	Appreciation of AI tools; Awareness of limitations; Importance of prompt engineering	"I believe that it is even more relevant than I initially thought, but it needs careful guidance."
Q12: Challenges faced while using Generative AI tools	Neutral to slightly negative	0.064	Inaccuracy of AI outputs; Prompt framing difficulties; Instructor-AI mismatch; Data privacy concerns	"The AI tools did not understand the expectations clearly, which led to wrong outputs."
Q13: Benefits from learning and applying ethical principles	Positive	0.214	Awareness of responsible data use; Critical thinking on AI bias; Ethics as core to analytics	"Learning and applying ethical principles helped me see how real-world analytics can go wrong."
Q14: Impact of combining Generative AI and Excel on understanding	Positive	0.171	AI for flexible learning; Excel for structured problem-solving; Complementary methods	"Using both AI tools and Excel helped me understand the concepts in different ways."
Q15: Suggestions for improving future courses	Most positive	0.229	Desire for practical case studies; Early exposure; Better prompts; Clearer ethical frameworks	"I would give more realistic examples of data scenarios to apply AI ethically."

Table 3: Qualitative sentiment analysis of student reflections on Generative AI and Ethics integration

APPENDIX B

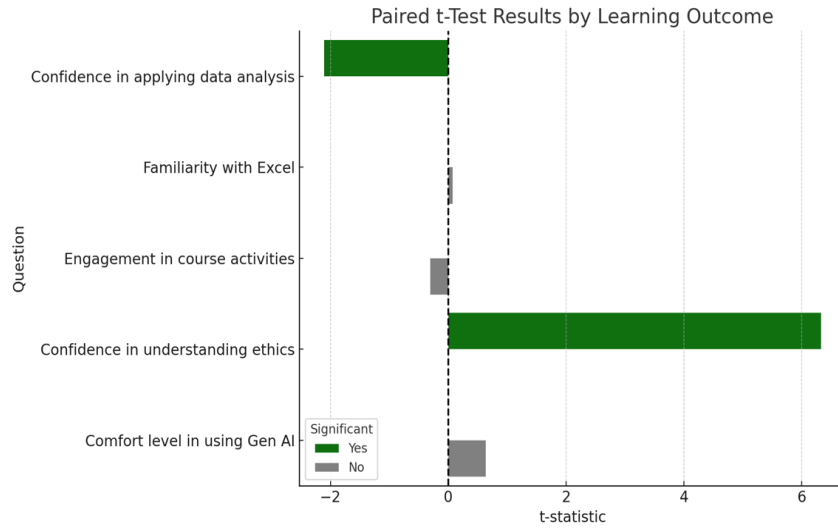


Fig 1: t-Test evaluation of changes in student confidence and skills before and after the introduction of GenAI and Ethics modules in the curriculum.



Figure 2: Thematic word clouds from student reflections, Post-GenAI integration in business analytics

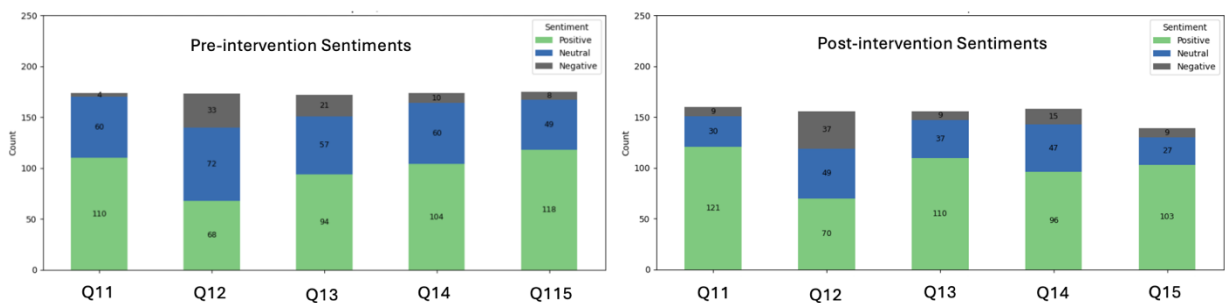


Figure 3: Shift in student sentiment before (left) and after (right) GenAI and Ethics integration to the curriculum.

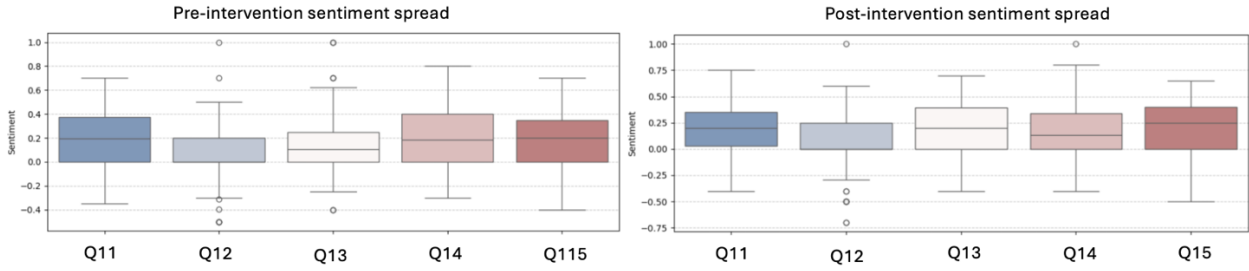


Figure 4: Sentiment distribution shift in student reflections: Pre (left) vs. Post-GenAI-Ethics (right) intervention.

Invited Article

Designing for Tension: Balancing the Pace of Trust and the Pace of Innovation in AI-Integrated Higher Education

David Firth

David.Firth@umontana.edu
Management Information Systems Department

Jason Triche

Jason.Triche@umontana.edu
Management Information Systems Department

Isho Tama-Sweet

Isho.Tama-Sweet@umontana.edu
Accounting and Finance Department

College of Business
University of Montana
Missoula, MT, 59812, USA

ABSTRACT

The rapid emergence of generative artificial intelligence (GenAI) has created a structural crisis in higher education that is as much organizational as it is pedagogical. As large language models (LLMs) and agentic systems evolve at an exponential rate, academic institutions find themselves caught between moving at the pace of technological innovation and the pace required to build institutional trust. A flagship university in the Rocky Mountain West developed a university-wide AI commitment framework that can be mapped to four robust organizational frameworks (i.e., organizational ambidexterity, relational coordination theory, the automation-augmentation paradox, and psychological safety). By applying these frameworks to a successful AI policy development context, the analysis articulates a strategy for designing for tension. This proposed strategy is a deliberate approach to governance that requires a pace of trust to enable the pace of innovation for success in GenAI infusion across college campuses.

Keywords: Generative Artificial Intelligence, AI Policy, Pace of Trust, Pace of Innovation, Literacies, Competencies

Recommended Citation: Firth, D., Triche, J., Tama-Sweet, I. Designing for Tension: Balancing the Pace of Trust and the Pace of Innovation in AI-Integrated Higher Education. *Information Systems Education Journal*, v24(n6), pp-75-88, DOI# # <https://doi.org/10.62273/MONF8056>

Designing for Tension: Balancing the Pace of Trust and the Pace of Innovation in AI-Integrated Higher Education

David Firth, Jason Triche, and Isho Tama-Sweet

1. INTRODUCTION

For Information Systems (IS) educators, the current AI revolution is not merely an external disruption but a core disciplinary challenge. Historically, the IS field has served as the bridge between technological capability and organizational value. Sharp et al. (2025) tracked this evolution since the inception of ISEDJ in 2003, covering shifts from cloud computing and data analytics to the current generative AI era. While GenAI tools like ChatGPT-5 offer unparalleled possibilities for knowledge acquisition and personalized learning, they also introduce profound risks regarding academic integrity, ethical use, and the potential displacement of critical thinking.

The pedagogical value of GenAI is being navigated in real time across campuses globally. IS faculty are uniquely positioned to lead this endeavor, as the discipline's focus on model curriculum, outcomes assessment, and socio-technical systems provides the necessary tools to evaluate AI integration. ISEDJ and its sister publication, the Journal of Information Systems Education (JISE), have become primary venues for documenting how students interact with these tools, and for suggesting curricula updates. However, the rate of technological change often outpaces the traditional curriculum update process, creating a relevance gap that requires a proactive, strategic response. This relevance gap is exacerbated by the tensions that exist on university campuses. In one camp are those who see their university as a place where students learn to think critically, with an emphasis on ethics, epistemology, and caution. They see the university as a place to teach established and trusted literacies. In the other camp are those who advocate for emphasizing rapid experimentation and workforce relevance. They see the university as a place to teach students competencies.

This paper introduces an AI commitment framework called the Future Project, developed at our university. The Future Project was built using existing frameworks and theories. The Future Project also leans into the tension between the pace of technological change and the pace of academia and uses it as a feature, not a defect.

In this paper, we first discuss literacies and competencies and introduce terms such as the pace of trust and the pace of innovation. We then discuss how moving either at the pace of trust or the pace of innovation creates an inherent tension on university campuses. Next, we introduce four theoretical frameworks to help understand how to think about this tension. We then present an exemplar of how an AI policy was developed at our university. We apply the four frameworks to this exemplar to help understand why it is successful and conclude with a prescription for AI policy development at other universities.

2. LITERACIES AND COMPETENCIES

To effectively navigate the integration of artificial intelligence in higher education, it is important to distinguish between the two often conflated constructs of literacies and competencies. In our view, both are essential to student development. They represent different dimensions of learning, engagement, and application within academic and professional contexts. The following discussion clarifies these concepts and highlights their complementary roles in preparing students for an AI-augmented world.

In an academic setting, the concept of literacies is best understood as a set of foundational, context-sensitive capabilities that enable individuals to critically engage with, interpret, and make meaning from information, tools, and systems within a given domain (Long & Magerko, 2020). Unlike discrete skills, literacies are inherently reflective, metacognitive, and situated. They require not only the ability to use a tool or interpret information, but also an understanding of the underlying assumptions, limitations, and implications of that use. In the context of emerging technologies such as artificial intelligence, literacies extend beyond technical familiarity to include epistemological awareness, including how knowledge is generated, validated, and contested. For example, AI literacy involves recognizing the probabilistic nature of outputs, identifying bias or hallucination, and understanding the socio-technical systems that shape the generative AI's behavior (Long & Magerko, 2020). Within higher education, literacies are closely tied to the institution's

traditional role as a site of critical inquiry and knowledge validation. Universities support students in becoming discerning participants in knowledge ecosystems, capable of questioning sources, evaluating evidence, and situating their understanding within broader intellectual and ethical frameworks. As such, literacies are less about performance and more about interpretive capacity, forming the cognitive and ethical foundation upon which more applied abilities are built.

In contrast, competencies in an academic setting refer to the demonstrable ability to apply knowledge and skills to effectively perform tasks or solve problems in specific contexts, often aligned with professional or disciplinary expectations (Voorhees, 2001). Competencies are action-oriented and outcome-focused (Boyatzis, 2008). They emphasize what learners can do, particularly in real-world or simulated environments. In the context of artificial intelligence, competencies might include integrating AI tools into workflows, generating and refining outputs, analyzing data with AI-assisted methods, and designing solutions that leverage AI capabilities. These competencies are increasingly central to workforce readiness, as employers seek graduates who can operate effectively in technology-rich environments (Dwivedi et al., 2023).

Competencies are not merely technical, however. They also encompass judgment, adaptability, and the capacity to collaborate with both humans and machine agents. Importantly, competencies are most effective when grounded in underlying literacies. Without critical understanding, the application of AI tools risks becoming superficial, uncritical, or even harmful. Thus, within higher education, competencies represent the operationalization of knowledge, translating abstract understanding into practical capability. The relationship between literacies and competencies is therefore complementary rather than hierarchical. Literacies provide the critical and ethical grounding, while competencies enable purposeful action. Together, they define a comprehensive approach to preparing students for participation in complex, AI-augmented academic and professional environments.

2.1 The Pace of Trust: AI Literacies

A literacy focus involves maintaining academic rigor, ensuring standardized assessment, and protecting the reputation of the university's degrees. This is what we term the pace of trust. It moves slowly because changes must be vetted to ensure they do not undermine the value of the

academic credential. While the phrase pace of trust is not formally codified in the literature, it reflects a well-established body of work emphasizing the need to work with ethical responsibility, transparency, and institutional legitimacy (Stilgoe et al., 2013; Floridi et al., 2018; Raisch & Krakowski, 2021). The pace of trust perspective reflects higher education's longstanding role in safeguarding knowledge, intellectual rigor, and ethical practice. This perspective aligns closely with emerging frameworks around AI literacies which emphasize critical engagement with AI systems (Long & Magerko, 2020).

Key elements of these AI literacies include understanding how AI systems generate outputs, recognizing bias, hallucinations, and limitations, evaluating outputs critically, and preserving authorship and academic integrity. Concerns around these key elements reflect deeper epistemological questions, such as what constitutes knowledge in an AI-mediated environment, and how learning should be assessed when AI can generate high-quality outputs. Scholars have noted that generative AI challenges traditional models of assessment and authorship, requiring new approaches to evaluation and pedagogy (Kasneji et al., 2023).

2.2 The Pace of Innovation: AI Competencies

In contrast, a competency perspective emphasizes workforce readiness and adaptability. We term this the pace of innovation. From this viewpoint, AI is already embedded in professional contexts, and higher education must prepare students accordingly (Dwivedi et al., 2023). This leads to a focus on AI competencies, including integrating AI into workflows, using AI to solve real-world problems, rapidly iterating on ideas, and enhancing productivity and creativity.

2.3 Why This Tension Matters

The emergence of artificial intelligence in higher education has amplified a structural tension between the pace of trust and the pace of innovation, reflecting competing institutional priorities that are difficult to reconcile in practice. In one respect, universities are historically grounded in norms of rigor, peer review, ethical accountability, and epistemic validation. These processes inherently require time, deliberation, and consensus-building (Stilgoe et al., 2013; Floridi et al., 2018). Alternatively, the rapid

evolution of AI technologies, coupled with external pressures from industry and workforce demands, incentivizes institutions to adopt tools and pedagogies more rapidly to remain relevant and competitive (Dwivedi et al., 2023). This divergence creates friction across faculty, administrative, and disciplinary lines, as stakeholders differ in their tolerance for uncertainty, risk, and change. As a result, universities become sites of contested adaptation where decisions about AI integration are shaped not only by technical capability but also by underlying values regarding knowledge, authority, and responsibility. Rather than being a temporary disruption, this tension represents a persistent and defining feature of AI adoption in higher education, requiring institutions to intentionally navigate and, ultimately, design for the coexistence of both imperatives.

The tension between the pace of trust and pace of innovation reflects deeper differences in how higher education conceptualizes knowledge, teaching, and assessment, and the transformation AI imposes upon them. In contrast to how universities traditionally conceptualize knowledge, AI shifts knowledge from static content to dynamic content by continuously generating outputs. This challenges traditional models of content mastery and emphasizes the evaluation and synthesis of information (Luckin et al., 2016). In terms of teaching, AI repositions instructors from content deliverers to facilitators of critical thinking and application (Zawacki-Richter et al., 2019). Assessment is also challenged under the new paradigm. Traditional product-based (e.g. "write an essay", or "create a PowerPoint deck") assessment becomes less meaningful when AI can generate outputs. This necessitates process-based and reflective assessment models (Kasneci et al., 2023). Lastly, the pace of innovation challenges how we evaluate authorship of scholarship. AI complicates notions of authorship, raising questions about originality and intellectual ownership (Floridi & Chiriatti, 2020). As agentic systems contribute increasingly sophisticated reasoning, the primary tension lies in determining whether "originality" resides in the initial prompt, the algorithmic synthesis, or the human refinement of the output. Consequently, institutions must reconcile these shifting definitions of ownership with existing frameworks of academic integrity to prevent institutional paralysis and maintain the pace of trust.

In the next section, we present four theoretical frameworks to help us understand and reconcile this tension.

3. THEORETICAL FRAMEWORKS TO UNDERSTAND AND RECONCILE THE TENSION

3.1 Theoretical Framework I: Organizational Ambidexterity

The central challenge for any university in the age of AI is a classic management dilemma of how to leverage and maintain existing pedagogical and research strengths while exploring radically new, disruptive technologies. This is the essence of organizational ambidexterity, which is the ability to be equally adept in both camps. In the context of higher education, the theory of organizational ambidexterity (Tushman & O'Reilly, 1996) is the ability for universities to exploit current capabilities like existing curriculum, academic standards, and administrative processes, while simultaneously exploring new opportunities, like developing and implementing AI-centric curricula, and automating administrative systems. We refer to the former part of this theory (i.e., exploiting current capabilities) as the pace of trust. What the theory refers to as exploration, conversely, involves experimentation with new technologies and radical innovation. We refer to this as the pace of innovation. Tushman and O'Reilly (1996) define ambidexterity as the ability to simultaneously exploit existing capabilities such as stability and trust, and explore new frontiers through innovation. Exploitation (i.e., pace of trust) is about safeguarding academic integrity and human authorship. Exploration (i.e., pace of innovation) is about rapidly piloting AI-augmented curricula.

Scholars suggest that the success of digital transformation in knowledge-intensive organizations depends on the balance between these two modes (March, 1991). If an institution focuses too heavily on exploitation, doing what it has done before, and advancing only when a new idea has proved it can be completely trusted, it risks becoming obsolete as the workforce demands new AI competencies. If it focuses solely on exploration without a foundation of trust, it may suffer from theoretically unsound thinking, fragmented governance, and shadow AI adoption.

3.2 Theoretical Framework II: Relational Coordination Theory

While ambidexterity provides the structure, relational coordination theory (RCT) provides the social mechanism for successful AI integration. Developed by Gittell (2002), RCT posits that effective coordination of complex, interdependent tasks requires high-quality communication supported by relationships of shared goals, shared knowledge, and mutual respect. AI integration in a university setting is a quintessential example of reciprocal interdependence (Thompson, 1967), in which the actions of IT departments, faculty, administrators, and students influence one another in a continuous loop.

In many institutions, the introduction of AI has exposed fragmented and poorly coordinated relationships across functional groups, reflecting low levels of relational coordination (Gittell, 2002; Gittell, 2011). Faculty may view AI as a threat to academic integrity, while administrators may view it primarily as a tool for efficiency and cost reduction. RCT offers a way to bridge these silos by fostering a learning network that mirrors the

network of care found in healthcare settings (Gittell, 2002).

3.2.1 The Seven Dimensions of Relational Coordination as Applied to AI

To navigate the tension between innovation and trust, institutions must operationalize the seven dimensions of RCT within their AI strategy. The seven dimensions are shared goals, shared knowledge, mutual respect, and communication that is frequent, timely, accurate, and problem-solving (Gittell, 2002) that collectively enable effective coordination in highly interdependent environments. This is detailed in Table 1.

3.3 Theoretical Framework III: The Automation-Augmentation Paradox

The third pillar of our theoretical framework is the automation-augmentation paradox (Raisch & Krakowski, 2021). This theory challenges the simplistic narrative that organizations should always prioritize augmentation (i.e., humans collaborating with machines) over automation (i.e., machines replacing human tasks).

RCT Dimension	Role in AI strategy	Outcomes
Frequent Communication	Moving beyond ad-hoc emails to regular, structured updates on AI capabilities and policies.	Reduces uncertainty and the use of shadow AI tools.
Timely Communication	Addressing ethical dilemmas and integrity violations in real-time rather than waiting for semester-end reviews.	Prevents the buildup of faculty resentment and anxiety.
Accurate Communication	Providing clear data on model biases, privacy risks, and hallucination rates.	Builds trust through transparency and honesty.
Problem-Solving Communication	Shifting from a punitive model of cheating detection to a collaborative model of authentic assessment.	Fosters a shared responsibility for student learning.
Shared Goals	Aligning the campus under a common vision of human-centered AI for the public good.	Overcomes subgoal optimization where departments compete for resources.
Shared Knowledge	Understanding how AI affects different disciplines differently (e.g., coding vs. creative writing).	Enables cross-boundary coordination and mutual support.
Mutual Respect	Valuing the perspectives of both the tech-forward early adopters and the cautious, skeptical faculty.	Overcomes status barriers and encourages inclusive dialogue.

Table 1: The seven dimensions of RCT (Gittell, 2002) and their role in AI strategy

The paradox suggests that as AI systems become more capable and automate more complex cognitive tasks, the need for high-level augmentation (i.e., human oversight, ethical judgment, and verification) actually increases. In higher education, this paradox will become evident in a shift from grading content to grading process. As LLMs become capable of generating A-level essays (i.e., automation), the role of the professor shifts toward evaluating the student's process, prompt engineering logic, and source verification (i.e., augmentation). If an institution focuses too heavily on automation (e.g., using AI for all grading and administrative triage), it risks reinforcing cycles that lead to a loss of human expertise and institutional memory. Conversely, overemphasizing augmentation without sufficient automation can lead to cognitive overload and burnout for faculty who are forced to manage increasingly complex student-AI, faculty-AI, and administration-AI workflows.

3.4 Shadow AI and Theoretical Framework IV: Psychological Safety

A recurring theme in IS research is the emergence of shadow AI, referring to the unauthorized use of consumer-grade AI tools by faculty, students, and staff seeking efficiency outside of sanctioned institutional channels, a phenomenon conceptually aligned with prior research on shadow IT (e.g. Behrens, 2009; Silic & Back, 2014). When the pace of innovation (e.g., the availability of easy-to-use tools) exceeds the pace of trust (e.g., the development of clear, safe policies), users will inevitably bypass sanctioned tools in favor of personal solutions that offer measurable value. The risks of shadow AI are:

- **Data Privacy:** Sensitive student or research data may be input into models without appropriate security protocols.

- **Academic Integrity:** Unauthorized use creates an opaque environment where it is impossible to distinguish between original work and machine-generated output.
- **Digital Inequity:** Students who can afford premium, private AI tools may gain an unfair advantage over those relying on limited, institutional alternatives.

3.4.1 The Role of Psychological Safety

To mitigate shadow AI, institutions must cultivate an environment of psychological safety. Edmondson (1999) defines this as a shared belief that a team is safe for interpersonal risk-taking. In the context of AI, this means that faculty and students must feel safe to experiment with new tools without fear of judgment or a competence penalty.

Research suggests that individuals who feel psychologically safe are more likely to speak up with ideas, report errors, and seek feedback (Edmondson, 1999). For AI adoption, this is critical because AI outputs are often uncertain, ambiguous, or imperfect. If users fear negative evaluation, they will hide their use of AI, perhaps leading to the shadow AI problem. Conversely, a psychologically safe environment allows the community to fail forward and learn from the technology's mistakes in a transparent manner. Psychological safety operates at each level of analysis within higher education. Table 2 articulates the impact of psychological safety at each of these levels and a strategy the institution can take to improve safety at that level.

Level of Interaction	Impact of Psychological Safety on AI Adoption	Institutional Strategy
Individual Faculty	Willingness to admit lack of fluency and seek training.	Providing stipends and professional development for AI exploration.
Classroom (Student)	Willingness to disclose AI use and discuss its ethical implications.	Engaging students as partners in policy formation.
Administrative Teams	Willingness to report hallucinations or biased outcomes in automated workflows.	Establishing open, inclusive dialogue that welcomes skepticism.
Institutional Level	Ability to pivot strategy based on feedback from the front lines.	Maintaining principled adaptation and regular quarterly surveys.

Table 2: Impact of Psychological Safety on AI Adoption at different levels within higher education and strategies to improve safety

Research suggests that individuals who feel psychologically safe are more likely to speak up with ideas, report errors, and seek feedback (Edmondson, 1999). For AI adoption, this is critical because AI outputs are often uncertain, ambiguous, or imperfect. If users fear negative evaluation, they will hide their use of AI, perhaps leading to the shadow AI problem. Conversely, a psychologically safe environment allows the community to fail forward and learn from the technology's mistakes in a transparent manner. Psychological safety operates at each level of analysis within higher education. Table 2 articulates the impact of psychological safety at each of these levels and a strategy the institution can take to improve safety at that level.

4. FUTURE PROJECT AT OUR UNIVERSITY

Our university provides a useful longitudinal case study for how a flagship public university can navigate the tension between the pace of trust and the pace of innovation. Launched in early 2025, with support from a locally headquartered technology consulting company, this endeavor, labeled the Future Project, was designed as a university-wide initiative that prioritizes community input, academic values, and thoughtful integration over pace. The project proceeded in phases, which are described below.

4.1 Phase I: Exploration and Engagement (Spring 2025)

The first phase of the Future Project focused on the relational and structural foundations of AI governance. This phase was not about selecting tools, but about orienting faculty and staff coordinators and conducting cross-campus conversations. By engaging over 200 participants, the university established a shared knowledge base, which reflects a core tenet of relational coordination theory (see Table 1).

The identification of key opportunities and challenges during this phase allowed the university to build its absorptive capacity, which is the organization's ability to recognize the value of new knowledge, assimilate it, and apply it effectively (Cohen & Levinthal, 1990). Rather

than a top-down mandate, the project coordinators, who represent 55 departments, served as boundary spanners who could translate the university's AI vision into discipline-specific practices. Boundary spanners are individuals who link internal and external organizational environments and facilitate information flow (Tushman & Scanlan, 1981).

4.2 Phase II: Strategy Development (Summer 2025)

In Phase II, the university transitioned from inquiry to formal strategy. This involved defining the university's AI Commitments (described in Section 4.5) and building a framework for sustainable AI governance. A key structural outcome was the development of a visual representation of the university's statewide AI leadership role, which positioned the university as a central node in our state's workforce development ecosystem.

4.3 Phase III: Implementation and Community Dialogue (Fall 2025 – Spring 2026)

The current phase involves soliciting campus-wide feedback on the draft commitments. This iterative process is a hallmark of principled adaptation (Stilgoe et al. 2013), where policies are refined based on real-world usage patterns and community critique. This phase also includes the launch of pilot programs and early action items, which allow for controlled exploration within a safe, institutional pace of trust framework.

4.4 The Four Lines of Effort: A Comprehensive Information Systems Perspective

Each phase of the Future Project is organized around four interconnected lines of effort. These map directly to the core concerns of IS education. These lines of effort functioned as subgroups within the Future Project. Each subgroup was composed of self-selecting individuals who wanted to work on issues related to this domain.

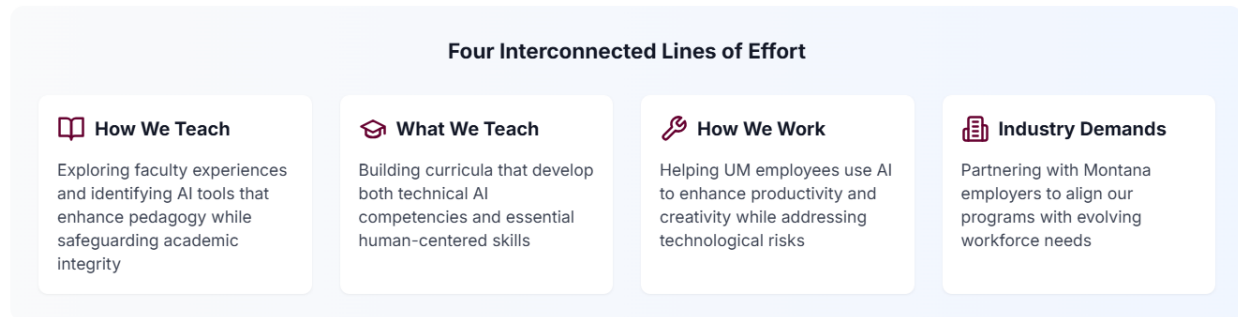


Figure 1: The Future Project is split into four lines of effort

#	AI Commitment	Description
#1	Preserve Human Judgment	Preserve human judgment and authorship in teaching, research, creative work, and decision-making. AI should augment, not replace, the distinctly human dimensions of life and work on campus.
#2	Engage Students as Partners	We engage students as partners in shaping the ethical, creative, and civic possibilities of AI, preparing them for thoughtful leadership in an AI-augmented world.
#3	Uphold Transparency	We uphold transparency and accountability in how AI tools are selected, applied, and evaluated, ensuring clarity in both capabilities and limitations.
#4	Ensure Equitable Access	We ensure equitable access to AI tools, training, and learning opportunities across all disciplines, backgrounds, and stages of life, bridging digital divides rather than reinforcing them.
#5	Support Faculty Development	Support faculty and instructors as they navigate evolving pedagogies, assessment practices, and ethical dilemmas introduced by AI. This includes professional development, peer communities, and flexible policy guidance.
#6	Equip Staff and Administrators	Equip staff and administrators with the resources, training, and policy clarity needed to integrate AI into their work responsibly, efficiently, and in alignment with our values and goals.
#7	Invest in Sustainable Infrastructure	Invest in sustainable infrastructure, responsible governance, and inclusive policies that promote innovation while protecting rights, privacy, and intellectual integrity.
#8	Environmental Responsibility	Pursue creative, sustainable solutions to reduce the environmental impact of AI development and deployment, aligning our technological progress with climate responsibility.
#9	Foster Inclusive Dialogue	Foster an open, inclusive campus dialogue that welcomes curiosity, critique, and collaboration from all perspectives, including those skeptical of AI's role in education and society.

Table 3: The nine elements of the AI Commitment

The four lines of effort are:

1. **How We Teach:** This effort explores faculty experiences and identifies AI tools that enhance pedagogy while safeguarding academic integrity. It addresses the automation-augmentation paradox (see Section 3.3) by seeking tools that augment the instructor's capability rather than automating the educational experience out of existence.
2. **What We Teach:** This focuses on building AI literacy across disciplines and developing critical thinking about AI capabilities and limitations. It recognizes that students must not only be savvy users but ethical and innovative creators.
3. **How We Work:** This investigates how our employees can use AI to enhance their work while addressing risks of technological displacement. This line of effort applies organizational ambidexterity (see Section 3.1) to the university's administrative side, seeking administrative efficiency without sacrificing human oversight in decision-making.
4. **Industry Demands:** By engaging directly with the employers of our graduates, our university aligns its academic programs with evolving workforce needs. This ensures that the relevance gap in the IS curriculum is closed through direct industry partnerships.

4.5 Our University's AI Commitment

During Phase II, the strategy-development phase of the Future Project, the university developed the AI Commitment. This is the university's foundational approach to artificial intelligence, rooted in human learning, creativity, and community flourishing (see <https://Umontana.AI/commitment>). A screenshot of the web page for this is in Appendix A.

Our AI commitment represents a more adaptive and institutionally appropriate approach than a traditional AI policy because it prioritizes guiding principles over prescriptive rules in a rapidly

evolving technological landscape. Policies, by design, tend to be static, compliance-oriented, and quickly outdated when applied to technologies like AI that are advancing at an exponential pace (Stilgoe et al., 2013). In contrast, a commitment framework enables principled adaptation, allowing institutions to remain anchored in core values, such as transparency, equity, and human judgment, while still encouraging experimentation and innovation. A principled adaptation is one in which institutions evolve in response to technological change while remaining grounded in enduring normative commitments (Floridi et al., 2018). This closely aligns with the organizational ambidexterity framework (see Section 3.1). This is particularly important in higher education, where rigid policies risk either constraining pedagogical creativity or being widely ignored as impractical. By articulating shared values rather than fixed prohibitions, as is common in policy, an AI Commitment fosters a culture of trust, dialogue, and collective responsibility, aligning more effectively with the university's mission of inquiry and intellectual exploration. In doing so, it provides a flexible yet ethically grounded foundation for navigating the tension between the pace of trust and the pace of innovation. Table 3 describes the nine AI Commitments at our university.

We map the AI Commitments to the theoretical frameworks of organizational ambidexterity (Section 3.1), relational coordination (Section 3.2), the automation-augmentation paradox (Section 3.3), and psychological safety (Section 3.4) in Table 4.

Our university explicitly engages the shadow AI problem by fostering psychological safety. It does this by welcoming curiosity, critique, and collaboration from all perspectives, including those skeptical of AI's role under the Foster Inclusive Dialogue element of our AI commitment. By providing clear, sanctioned, secure enterprise alternatives like Grammarly and NotebookLM, we reduce the incentive for users to turn to unauthorized tools.

#	AI Commitment	Theoretical Lens	Institutional Application
1	Preserve Human Judgment	Ambidexterity (Exploit/Pace of Trust) and Automation-Augmentation	Ensures AI augments rather than replaces cognition. This commitment ensures the university continues to exploit its core competency of critical thinking. It signals that while we explore AI, we will not abandon the stable human elements that define higher education, thereby maintaining institutional trust
2	Engage Students as Partners	Relational Coordination (Shared Goals and Mutual Respect)	Co-creating ethical frameworks with learners. Gittel (2011) emphasizes mutual respect as a driver of pace. Treating students as partners rather than subjects of an AI Commitment allows the university to move at the pace of trust by co-creating the rules of engagement.
3	Uphold Transparency	Psychological Safety (Safety)	Openly discussing tool selection and evaluation. As Edmondson (1999) notes, transparency reduces the perceived risk of innovation. By being clear about how AI is selected and evaluated, our university creates a safe zone for faculty and students to operate without fear of arbitrary punishment.
4	Ensure Equitable Access	Relational Coordination (Shared Knowledge)	Closing the digital divide to maintain social cohesion. Access is the baseline for shared knowledge. If only a portion of the student body has access to advanced AI, the relational coordination of the entire campus is fractured. Equitable access ensures everyone is moving at the same pace of innovation.
5	Support Faculty Development	Psychological Safety (Capability Building)	Provide training to reduce fear and increase agency.
6	Equip Staff and Administrators	Psychological Safety (Capability Building)	Training is the antidote to the fear of obsolescence. By supporting development, we increase the psychological safety of our staff, allowing them to lead innovation rather than resist it.
7	Invest in Sustainable Infrastructure	Ambidexterity (Explore/Pace of innovation)	Building infrastructure for responsible integration.
8	Environmental Responsibility	Ambidexterity and Psychological Safety	Being environmentally responsible encompasses the pace of innovation. Ambidexterity, as well as providing the psychological safety that the academic community values, is important for addressing environmental issues.
9	Foster Inclusive Dialogue	Relational Coordination (Mutual Respect) and Psychological Safety	Encouraging skepticism and critique as a form of trust. Inclusive dialogue serves as the bridge between exploration and exploitation. Leading with integrity ensures that the pace of innovation never outruns the pace of trust, creating a sustainable, long-term ambidextrous culture.

Table 4: Mapping AI commitments to the theoretical foundations

Our university explicitly engages the shadow AI problem by fostering psychological safety. It does this by welcoming curiosity, critique, and collaboration from all perspectives, including those skeptical of AI's role under the Foster Inclusive Dialogue element of our AI commitment. By providing clear, sanctioned, secure enterprise alternatives like Grammarly and NotebookLM, we reduce the incentive for users to turn to unauthorized tools.

5. DESIGN FOR PRODUCTIVE TENSION

The central argument of this paper is that our academic institutions should not resolve the tension between moving at the pace of trust and moving at the pace of innovation, but instead design for it. Dynamic tension is a strength. Innovation without trust leads to risk, inequity, and institutional rejection. Trust without innovation leads to stagnation and irrelevance for the workforce. Effective institutions maintain a dual framing strategy in which AI competency is taught, and every AI tool used must be paired with an AI literacy reflection.

To successfully design for productive tension, higher education institutions should follow a multi-level strategy derived from effective frameworks and organizational theories. Specifically, we recommend the following strategies of establishing structural ambidexterity, prioritizing relational coordination, cultivating psychological safety, re-centering the human in the automation paradox, and closing the relevance gap in the curriculum.

5.1. Establish Structural Ambidexterity

Create a dedicated exploration unit (e.g., a Future Project) that is spatially separated from the day-to-day operations but functionally integrated through a network of department-specific coordinators. This allows for rapid experimentation without destabilizing the core exploitative pace of trust mission of the university.

5.2. Prioritize Relational Coordination

Invest in human architecture before technical architecture. Use the seven dimensions of RCT to build trust between IT, faculty, students, and administration. Move beyond top-down mandates

to inclusive dialogue that values skepticism and critique.

5.3. Cultivate Psychological Safety

Develop a safe-to-fail environment where faculty, students, and administrators can disclose their AI use and learn from mistakes. Provide sanctioned, secure enterprise alternatives to mitigate the risks of shadow AI.

5.4. Re-Center the Human in the Automation Paradox

Apply a human-in-the-loop framework to all AI-augmented processes. Ensure that AI augments and never replaces the distinctly human dimensions of education. Prioritize authentic, process-led assessment over content-based grading.

5.5. Close the Relevance Gap in Curriculum

Regularly update the IS curriculum to include AI literacy and technical AI competencies. Engage directly with industry partners to ensure that graduates possess both the technical skills and the human-centered skills (e.g., creativity, ethics, reasoning) that the future of work demands.

6. CONCLUSION

The integration of artificial intelligence into higher education is not a technical problem to be solved, but a social and organizational challenge to be managed. The pace of trust is the foundational constraint that determines whether innovation strengthens or diminishes the academic mission. Based on our experience, the most effective path forward is a deliberate, human-centered approach that is grounded in academic freedom, critical thinking, and transparency. This allows faculty, staff, and students to move safely and sensibly at the pace of innovation.

By synthesizing organizational ambidexterity, relational coordination, the automation-augmentation paradox, and psychological safety, IS educators can provide the theoretical rigor and empirical evidence needed to guide their institutions through this transition. The goal is not to reach a final state of AI adoption but to build a framework, an AI Commitment for ongoing growth and principled adaptation. In this new landscape, the university's value lies not in its ability to generate content, but in its ability to foster the wisdom, curiosity, courage, and care

needed to wield that content for the public good. As we design for tension, we ensure that higher education remains a vital, human-centered institution in an increasingly automated world.

7. REFERENCES




- Behrens, S. (2009). Shadow systems: The good, the bad and the ugly. *Communications of the ACM*, 52(2), 124–129. <https://doi.org/10.1145/1461928.1461960>
- Boyatzis, R. E. (2008). Competencies in the 21st century. *Journal of Management Development*, 27(1), 5–12. <https://doi.org/10.1108/02621710810840730>
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128–152. <https://doi.org/10.2307/2393553>
- Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., ... & Wright, R. (2023). So what if ChatGPT wrote it? *International Journal of Information Management*. Advance online publication. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
- Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), 350–383. <https://doi.org/10.2307/2666999>
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chila, G., Dignum, V., Duffy, B., Madelin, B., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30, 681–694. <https://doi.org/10.1007/s11023-020-09548-1>
- Gittell, J. H. (2002). Coordinating mechanisms in care provider groups: Relational coordination as a mediator and input uncertainty as a moderator of performance effects. *Management Science*, 48(11), 1408–1426. <https://doi.org/10.1287/mnsc.48.11.1408.268>
- Gittell, J. H. (2012). New directions for relational coordination theory, in Gretchen M. Spreitzer, and Kim S. Cameron (eds), *The Oxford Handbook of Positive Organizational Scholarship*. 2011; online edn, Oxford Academic, 21 Nov. 2012), <https://doi.org/10.1093/oxfordhb/9780199734610.013.0030>.
- Kasneji, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerb, J., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., Stadler, M., Weller, J., Kuhn, J., Kasneji, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). <https://doi.org/10.1145/3313831.3376727>
- Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. B. (2016). *Intelligence unleashed: An argument for AI in education*. Pearson Education
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87. <https://doi.org/10.1287/orsc.2.1.71>
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://doi.org/10.5465/amr.2018.0072>
- Sharp, J., Anderson, J., & Lang, G. (2025). Visualizing the eras of IS education: A co-occurrence analysis of the *Information Systems Education Journal*. *Information Systems Education Journal*, 23(4), 4–21. <https://doi.org/10.62273/FYDF2042>

- Silic, M., & Back, A. (2014). Shadow IT—A view from behind the curtain. *Computers & Security*, 45, 274–283. <https://doi.org/10.1016/j.cose.2014.06.007>
- Stilgoe, J., Owen, R., & Macnaghten, P. (2013). Developing a framework for responsible innovation. *Research Policy*, 42(9), 1568–1580. <https://doi.org/10.1016/j.respol.2013.05.008>
- Thompson, J. D. (1967). *Organizations in action: Social science bases of administrative theory*. McGraw-Hill.
- Tushman, M. L., & O'Reilly, C. A. (1996). Ambidextrous organizations: Managing evolutionary and revolutionary change. *California Management Review*, 38(4), 8–30. <https://doi.org/10.2307/41165852>
- Tushman, M. L., & Scanlan, T. J. (1981). Boundary spanning individuals: Their role in information transfer and their antecedents. *Academy of Management Journal*, 24(2), 289–305. <https://doi.org/10.2307/255842>
- Voorhees, R. A. (2001). Competency-Based Learning Models: A Strategic Opportunity. *New Directions for Institutional Research*, 2001(110), 5–12. <https://doi.org/10.1002/ir.7>
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education – where are the educators? *International Journal of Educational Technology in Higher Education*, 16, 39. <https://doi.org/10.1186/s41239-019-0171-0>





APPENDIX A: The University of Montana “AI Commitment”





Our Commitments

-  **Preserve Human Judgment**
Preserve human judgment and authorship in teaching, research, creative work, and decision-making. AI should augment—not replace—the distinctly human dimensions of life and work on campus.
-  **Engage Students as Partners**
We engage students as partners in shaping the ethical, creative, and civic possibilities of AI, preparing them for thoughtful leadership in an AI-augmented world.
-  **Uphold Transparency**
We uphold transparency and accountability in how AI tools are selected, applied, and evaluated, ensuring clarity in both capabilities and limitations.



-  **Ensure Equitable Access**
We ensure equitable access to AI tools, training, and learning opportunities across all disciplines, backgrounds, and stages of life—bridging digital divides rather than reinforcing them.
-  **Support Faculty Development**
Support faculty and instructors as they navigate evolving pedagogies, assessment practices, and ethical dilemmas introduced by AI. This includes professional development, peer communities, and flexible policy guidance.
-  **Equip Staff & Administrators**
Equip staff and administrators with the resources, training, and policy clarity needed to integrate AI into their work responsibly, efficiently, and in alignment with UM's values and goals.
-  **Invest in Sustainable Infrastructure**
Invest in sustainable infrastructure, responsible governance, and inclusive policies that promote innovation while protecting rights, privacy, and intellectual integrity.



-  **Environmental Responsibility**
Pursue creative, sustainable solutions to reduce the environmental impact of AI development and deployment, aligning our technological progress with climate responsibility.
-  **Foster Inclusive Dialogue**
Foster an open, inclusive campus dialogue that welcomes curiosity, critique, and collaboration from all perspectives—including those skeptical of AI's role in education and society.

From: <https://Umontana.AI/commitment>