

Pathways to AI-Ready Entry-Level Talent for Industry Domains: Modular, Adaptive Design Principles for Business-School Curriculum

Completed Research Paper

Arun Rai

Georgia State University
Atlanta, GA, USA
arunrai@gsu.edu

Cynthia Breazeal

Massachusetts Institute of Technology
Cambridge, MA, USA
cynthiab@media.mit.edu

Benjamin M. Ampel

Georgia State University
Atlanta, GA, USA
bampel@gsu.edu

Xinyu Fu

Georgia State University
Atlanta, GA, USA
xinyufu@gsu.edu

Balasubramaniam Ramesh

Georgia State University
Atlanta, GA, USA
bramesh@gsu.edu

Eric Klopfer

Massachusetts Institute of Technology
Cambridge, MA, USA
klopfer@mit.edu

Amrita George

Georgia State University
Atlanta, GA, USA
ageorge12@gsu.edu

Madhu Kota

Georgia State University
Atlanta, GA, USA
mkota@gsu.edu

Abstract

Generative Artificial Intelligence (GenAI) is reconfiguring the entry pathway into business work faster than university curricula can adapt. This disruption erodes the apprenticeship layer through which graduates historically developed expertise. This study employs a Design Science Research approach to develop and instantiate the AIX Hub, an industry-partnered AI curriculum co-designed by two R1 hub institutions. The design follows a no-code to low-code to full-code progression grounded in cognitive apprenticeship theory, aligns learning objectives to a 35-skill labor market taxonomy, and embeds work-based learning through a stage-based industry partnership model. Spring 2026 deployment reached 1,060 students across 12 sections of four courses at a minority-serving institution, with eight work-based learning opportunities spanning eight industry partners. From this instantiation, we formalize three preliminary design principles (derived progression, adaptive industry co-design, and equity-relevant deployment) for curricula that prepare AI-ready entry-level talent.

Keywords: *AI literacy, curriculum design, design science research, work-based learning, no-code/low-code development*

Introduction

The rapid proliferation of Generative Artificial Intelligence (GenAI) and autonomous agents across industries has heightened demand for business graduates who can understand, evaluate, and direct AI systems (Brynjolfsson et al., 2025a; Sajadieh et al., 2026; OECD, 2024). A systematic review of 89 empirical GenAI-in-education studies highlights GenAI's extensive impact on student learning but notes that research has focused primarily on individual tools and course-level interventions rather than institution-level curriculum frameworks (Liu et al., 2025). Institutional responses remain fragmented, ranging from prohibition to uncritical adoption (Kishore et al., 2023). However, business school curricula continue to rely on traditional structures that do not prepare students to collaborate with, supervise, or critically evaluate contemporary AI systems. This gap is especially pronounced in information systems (IS) education, where curricula must balance technical rigor with accessibility for students with limited

computing backgrounds across diverse business disciplines. This gap reflects a deeper shift in how entry-level expertise is developed in AI, rather than a simple lag in curriculum updates.

The urgency stems from a structural disruption of the apprenticeship layer of knowledge work. Recent evidence shows GenAI is reconfiguring entry-level business work by automating tasks through which novices gain practical experience (Tomlinson et al., 2025; Brynjolfsson et al., 2025b; Gmyrek et al., 2023). Younger workers face disproportionate employment declines in AI-exposed occupations while employment for more experienced workers remains stable or grows, suggesting compression of the apprenticeship layer (Brynjolfsson et al., 2025a; Sajadieh et al., 2026). Task-level analyses show that GenAI is especially applicable to drafting, summarizing, coding scaffolds, and scripted customer interaction work (Tomlinson et al., 2025). These activities have traditionally helped entry-level employees build domain familiarity and skill. Firms gain productivity by substituting AI for low-stakes cognitive work but reduce opportunities for novices to acquire experience. Thus, the disruption represents a breakdown in how workers become skilled.

This breakdown reveals a fundamental issue in curriculum design. When the entry pathway into business work evolves faster than a university cohort can progress, preparing students for predefined roles becomes impractical. At the same time, this disruption creates a new frontier of entry-level work focused on collaborating with, evaluating, validating, and orchestrating AI systems. These roles are not yet fully stabilized and require active definition and shaping. Addressing this challenge requires moving beyond reskilling as mere tool adoption toward redesigning entry pathways as socio-technical systems. The proposed curriculum design is an ecosystem approach integrating cognitive apprenticeship with work-based learning (WBL), embedding learners in authentic production contexts while maintaining guided participation, feedback, articulation, and reflection (Collins et al., 1989). Under AI-driven conditions, the goal shifts to enabling learners, educators, and employers to collaboratively determine the division of tasks between humans and AI, identify essential capabilities, and establish performance evaluation criteria. The curriculum design functions as an artifact consistent with a design science perspective (Hevner et al., 2004), treating workforce development as an iterative process of designing and evaluating new educational pathways alongside evolving technologies.

These challenges motivate three design requirements stemming from the disruption of the apprenticeship layer. First, curricula must respond to skill and role volatility at a granularity that updates faster than full courses. Second, curricula must recreate the developmental experiences through which novices learn judgment in real work contexts. Third, institutions serving disadvantaged students must implement these curricula in ways that remain accessible, including pedagogical strategies that lower barriers to AI learning.

To meet these requirements, this study develops and instantiates the AIX Hub. The AIX Hub was co-designed by two R1 institutions and supported by an industry-sponsored workforce-development initiative. It follows a hub-and-spoke architecture, where each R1 hub anchors a regional ecosystem of 4-year and 2-year-spoke institutions. It features a four-course progression grounded in cognitive apprenticeship (Collins et al., 1989). The progression scaffolds students from no-code AI use as practitioners, through low-code configuration of model pipelines as technologists, to full-code design and deployment as builders. The Spring 2026 deployment enrolled over 1,000 undergraduates across four newly developed AI courses implementing this modular progression. It engaged eight industry partners in a regional FinTech corridor who provided capstone projects, internships, and curriculum input. The focal regional hub's ecosystem also includes a Community College Partner, a Regional University Partner, and an HBCU partner.

From this instantiation, the contribution is a set of preliminary design principles for AI curricula in business schools, formally stated as DP1–DP3 in Table 4 and articulated through three interrelated components. The first is a structured no-code-to-low-code-to-full-code skill progression grounded in cognitive apprenticeship, which lowers entry barriers while building technical competence and meta-cognitive judgment about appropriate tool use. The second is a bidirectional demand-signals-to-curriculum alignment infrastructure, designed for adaptive realignment in response to rapid skill drift. The third is a hub-and-spoke deployment model that situates the curriculum in equity-relevant institutional contexts and uses course-embedded capstones to provide authentic AI work experience to all enrolled students.

Collectively, these three design principles define the proposed curriculum's purpose, structure, deployment logic, and theoretical foundation, as demonstrated through the initial Spring 2026 implementation. The study is situated in a design science research framework (Hevner et al., 2004; Rai, 2017; Hevner & Chatterjee, 2010) and aligns with the design-theory anatomy outlined by Gregor and Jones (2007). While

prior IS education research has explored single-course AI instruction (Grøder et al., 2022; Chen et al., 2023), AI literacy tools for non-experts (Pinski et al., 2023), and student competency gaps (Ogundipe et al., 2023), it has not addressed multi-course AI progression, industry-partnered demand-signal alignment, or equity-relevant scaling in business school curricula. The present study contributes a transferable and evaluable design for institutions encountering similar challenges under varying local conditions.

Literature Review

AI Literacy and IS Education

The concept of AI literacy (the competency to critically evaluate and engage with AI technologies) has gained increasing attention in IS and educational technology research (Long & Magerko, 2020; Ng et al., 2021). Scholars have distinguished multiple dimensions of AI literacy, including understanding how AI systems work, recognizing the social and ethical implications of AI deployment, and being able to use AI tools purposefully in domain contexts (Kandlhofer et al., 2016). For business students, this framing is particularly important. The goal is not to produce highly technical AI engineers, but to develop business professionals who can collaborate effectively with AI systems, critically assess AI outputs, and reason about problem-solving processes in technology-enabled environments—capabilities closely aligned with computational thinking (Wing, 2006). In IS education, Grøder et al. (2022) demonstrate that responsible AI can be taught through experiential learning using public welfare case studies, establishing that IS courses can simultaneously develop technical understanding and ethical reasoning. However, these single-course interventions do not fully address the larger challenge of how to build AI competency systematically across an entire business school cohort.

In the IS education literature, calls for modernizing curricula to incorporate machine learning and data science have grown since 2015 (Duan et al., 2019). However, business schools often still offer elective, upper-division courses targeted at students who have already self-selected into quantitative tracks. This elective model does little to develop broad AI competency across the student body (Zawacki-Richter et al., 2019). For example, a study engaged 120 undergraduate MIS students in prototyping AI chatbots to address social problems and produced measurable gains in AI conceptual understanding, ethical awareness, and social impact orientation (Chen et al., 2023).

A systematic review of 89 empirical studies found that GenAI research in education focuses on tool-level adoption and individual-course interventions, with few studies on multi-course, institution-level curriculum design (Liu et al., 2025). The review identified seven future research directions, with pedagogical integration and inclusive adoption most directly addressed by our proposed design principles (Liu et al., 2025). Multi-site experimental evidence reinforces the urgency. When students have structured access to GenAI, many cannot evaluate AI outputs or construct effective prompts (Ramos et al., 2025). This suggests that AI literacy is not a single skill but a developmental competency that requires progressive, scaffolded instruction. Therefore, AI curricula need accessible entry points that reduce initial technical barriers while supporting progression toward deeper understanding.

No-Code and Low-Code Approaches in IS Education

No-code and low-code (NCLC) platforms have been proposed as mechanisms for democratizing software development and data analysis by reducing syntactic barriers that often discourage non-programming majors (Waszkowski, 2019; Bock & Frank, 2021). In IS education, NCLC tools have been applied in introductory programming courses, database management, and in machine learning instruction. The argument for NCLC approaches is based on cognitive load theory (Sweller et al., 1998). By offloading syntactic concerns to visual or template-based interfaces, learners can focus their working memory on conceptual understanding (i.e., what a model does and why) rather than on debugging. Design science research on AI literacy applications confirms this principle, demonstrating that implementing a no-code AI literacy tool significantly improved non-expert understanding of AI concepts (Pinski et al., 2023).

However, NCLC-only pedagogies have limits. Students who never progress beyond drag-and-drop interfaces may struggle to adapt to novel tools or diagnose failures (Touretzky et al., 2019). This progression creates an AI teaching problem in which accessibility is necessary at the entry point but insufficient at the endpoint. Prompt engineering helps address this gap in AI scenarios because it requires students to begin intentionally shaping system behavior without demanding full programming fluency. Prior work suggests

that prompt quality affects AI output quality and that worked-example instruction improves prompting performance (Tolzin et al., 2024). Prior research suggests that prompt engineering can serve as an intermediate layer between no-code interaction and full-code system building. This body of work suggests the need for a progressive learning ladder, in which students move from accessible no-code tools (e.g., Google Gemini, Teachable Machine, n8n) toward progressively lower-level abstractions (e.g., Google Colab notebooks and Python frameworks) as their conceptual foundations deepen.

We note that a progressive technical sequence alone is not sufficient to adequately prepare students for the AI workforce. If students are to develop judgment about when and how AI should be used, that progression must be embedded in authentic organizational contexts where tools, constraints, and performance expectations reflect real work. This makes industry-academic partnership and WBL central to AI curriculum design.

Industry-Academic Partnerships and Work-based Learning

WBL (i.e., embedding authentic workplace activities in academic programs) is well-established for improving graduate employability and student engagement (McQuaid & Lindsay, 2005; Jackson & Collings, 2018). In IS education, authentic workplace activities provide up-to-date use cases, datasets, and problem framings that keep the curriculum current in fast-moving fields. A key challenge is that academic curriculum cycles (typically annual or semiannual) do not align with the industry's agility requirements (Rowe et al., 2012). Grounded theory research similarly suggests that students often perceive a disconnect between their IS education and professional readiness, and they address this disconnect through self-directed development rather than curricular support (Ogundipe et al., 2023). This gap calls for WBL designs that embed AI skill scaffolding in authentic workplace contexts to produce competent practitioners. Scaffolding in our context means structuring the complexity of AI tool mastery, explicitly guiding students from using high-level, no-code applications to achieving technical competence in full-code implementation.

Simulation-based learning research identifies Practice-Ready Transformation (moving beyond skill acquisition to real practitioners) as the aspirational ceiling of experiential IS education (Faisal et al., 2025). Curriculum co-design models have shifted from one-directional industry input (e.g., curriculum advisory boards) toward iterative co-creation approaches in which employers are embedded as project sponsors, dataset providers, or capstone evaluators (Billett, 2011; Cooper et al., 2010). This model reflects an emerging strand of IS education research focused on employer-engaged curriculum design (Topi et al., 2017).

Organizational training research points to a similar set of competency priorities. Schaetzle and Sohrabi (2025), for example, propose a GenAI training model for employees whose design principles align with the progression described here. This convergence suggests that the competencies emphasized in academic AI curriculum design are not unique to higher education but are also emerging as priorities in organizational training contexts. It therefore strengthens the case for curriculum designs that align course content with employer demand signals and embed that alignment in ongoing partnership structures rather than one-time advisory input.

Literature Observations

Three observations emerge from the literature review. First, AI literacy is a developmental competency that demands progressive, scaffolded instruction. Single-course interventions establish what is possible (Grøder et al., 2022; Chen et al., 2023; Pinski et al., 2023) but do not address competence at scale (Liu et al., 2025; Ramos et al., 2025). Second, NCLC pedagogies provide an entry point but are limited if students never advance beyond drag-and-drop interfaces (Touretzky et al., 2019). This suggests a sequence of abstractions rather than a fixed tool tier (Tolzin et al., 2024). Third, the literature on industry-academic partnerships increasingly emphasizes co-design over consultation (Billett, 2011; Cooper et al., 2010), with employer-sourced data and authentic problems driving deeper skill development (Aniceto & Corbett, 2025). Further, corporate training models point to the same skill hierarchy that academic programs must produce (Schaetzle & Sohrabi, 2025). The IS 2017 curriculum guidelines (Topi et al., 2017) specify model competency areas for IS programs but predate the GenAI shift, leaving open how these observations should be integrated into a single, transferable artifact for institutions adopting AI curricula at scale. The AIX Hub is positioned as that artifact. It contributes a set of design principles whose form-and-function choices, deployment principles, and theoretical grounding are articulated in the sections that follow, with the Spring 2026 deployment as the first expository instantiation.

Theoretical Framework: Cognitive Apprenticeship and the NCLC Progression

The kernel theory anchoring the AIX Hub design is cognitive apprenticeship (Collins et al., 1989). Cognitive apprenticeship extends traditional apprenticeship to cognitive tasks through six instructional mechanisms: (1) modeling, (2) coaching, (3) scaffolding, (4) articulation, (5) reflection, and (6) exploration. The framework has been applied in domains such as writing, mathematics, and computer science, but its use in AI education remains underdeveloped. We adopt cognitive apprenticeship because its explicit instructional sequence provides the structural granularity needed to derive a multi-course progression in a way that broader learning theories do not.

Cognitive apprenticeship provides the rationale for a no-code-first sequence. Structuring the curriculum solely around cognitive load reduction, such as beginning with worked examples in a full-code environment, misaligns with the apprenticeship sequence in two key respects. First, modeling must precede coaching; learners benefit from observing expert performance before attempting tasks independently. In AI education, no-code interfaces enable students to examine system outputs without the syntactic complexity of code. Second, articulation must precede exploration; learners should be able to explain system behavior before engaging in novel problem-solving. No-code environments emphasize conceptual understanding before implementation details. Consequently, the NCLC progression is derived from cognitive apprenticeship principles.


Each course instantiates a different phase of the apprenticeship sequence. Course 1 (Foundations in AI) emphasizes modeling and scaffolding. Instructors demonstrate AI workflows, and students evaluate outputs in a no-code environment that minimizes syntactic load. Course 2 (Data Science Foundations) shifts toward coaching with faded scaffolds. Students implement models in low-code Python, interpret results, and compare alternative approaches to the same business problem. Prompt engineering functions here as an intermediate stage between observation and implementation (Tolzin et al., 2024). Course 3 (Agentic AI) foregrounds articulation and reflection in the context of multi-agent systems, with students designing AI ventures and evaluating system reliability. Course 4 (Deep Learning for Business) emphasizes exploration. Students define their own problems across data modalities (e.g., images, audio, text) and build full-code AI systems with minimal scaffolding. The progression is a sequenced instantiation of cognitive apprenticeship for AI tasks.

Two design implications emerge. First, students who complete only Course 1, which is intentionally the largest cohort, should still acquire the early-phase competencies anticipated by cognitive apprenticeship. These students are expected to develop the ability to evaluate AI outputs critically and to articulate their trustworthiness. The intended outcome for this group is the modeling- and articulation-oriented judgment established in the initial stages of the sequence. Second, the progression maintains coherence only if the tools employed in each course align with the instructional requirements of that stage. Tool selection should support the relevant phase of apprenticeship, rather than merely representing a no-code, low-code, or full-code tier. This principle informs the subsequent curriculum design choices.

Research Design

Approach and Empirical Context

This study employs the Design Science Research (DSR) paradigm (Hevner et al., 2004; Rai, 2017) to develop the AIX Hub (see Figure 1) as a set of instantiated design principles for AI curricula in business schools. The paper is structured according to the design-theory anatomy, specifying the artifact's purpose and scope, articulating its form and function, grounding it in cognitive apprenticeship as the kernel theory, and documenting the Spring 2026 deployment as the initial instantiation (Gregor & Jones, 2007). The primary contribution is the design itself, rather than the testing of hypotheses regarding a preexisting phenomenon. This approach aligns with prior IS education research that has utilized design science to develop AI literacy interventions (Pinski et al., 2023) and responsible AI course designs (Grøder et al., 2022), but extends these efforts to a multi-course, industry-partnered curriculum at collegiate scale.



The AIX Hub artifact was co-designed by the focal regional hub institution (a large minority-serving R1 university with over 50,000 students and one of the world’s largest IS programs) and the focal national hub institution, a world-leading tech institution with strong leadership in digital technologies and AI for education. The empirical context for this paper is the focal regional hub’s College of Business and Department of Information Systems, which deployed the four-course progression in Spring 2026. The focal regional hub’s student body is 83% BIPOC, 44% Pell-eligible, 45% first-generation college students, and 60% female, making it an important test case for equitable AI education at scale. The Spring 2026 implementation enrolled 1,060 students across four courses and 12 sections. Data were collected through the AIX Hub platform, an internally developed curriculum management system that aggregates enrollment, skills, WBL, and labor market data.

Data were collected using a mixed-methods approach. Quantitative metrics, including enrollment, completion rates, and grade distributions for all 1,060 enrolled students, were aggregated from the focal regional hub’s student information system and the AIX Hub platform. The curriculum design process and outcomes were analyzed through a review of curriculum artifacts, specifically course syllabi, the formalized 35-skill taxonomy, and industry-provided capstone project charters. A recognized limitation of the empirical evidence reported here is that, although the artifact was designed across two R1 contexts (the focal regional hub and the focal national hub), only the focal regional hub has instantiated the design at the time of this writing. The focal national hub’s instantiation and the planned Fall 2026 expansion to the Community College Partner, Regional University Partner, and HBCU Partner will provide the first cross-instantiation evidence for the design principles’ transferability.

The AIX Hub: Curriculum Framework

Program Architecture and Institutional Scope

The AIX Hub is organized as a hub-and-spoke model in which the focal regional hub institution anchors a regional ecosystem and supports three partner institutions in the first expansion cohort. Each spoke serves a distinct student population and integrates selected curriculum modules into its existing programs, subject to skill-precedence constraints, adopting modules that fit local prerequisite structures rather than replicating the full hub curriculum. The spoke institutions draw on shared industry partnerships, the AIX skills taxonomy, and support from the industry-sponsored workforce-development initiative (see Table 1).

Institution	Students	Demographics	Status	Planned Courses
Hub	52,400 total; 1,700 IS	83% BIPOC, 44% Pell, 45% 1st-gen, 60% female	Active (Spring 2026)	Courses 1–4
Community College	18,000	Community college	Planned (Fall 2026)	AI 101 (modules)
Regional University	12,000	Suburban population	Planned (Fall 2026)	IT 3XX, BUS 3XX
Historically Black College/University	4,000; ~1,000 business	HBCU	Planned (Fall 2026)	IS 3XX (AI Innovation)

Table 1: The AIX Hub Institutional Network (Spring-Fall, 2026)

The curriculum is designed to avoid the Turing Trap by emphasizing human augmentation and new value creation rather than simple task substitution (Brynjolfsson, 2022). Its underlying premise is that AI-readiness depends not only on technical proficiency, but on the ability to work with AI in ways that expand human contribution across domains. Accordingly, the curriculum combines technical training across the four core AI courses with broader capabilities, including creative thinking, analytical reasoning, resilience, adaptability, agility, and technological literacy.

The No-Code to Low-Code Progression

The curriculum follows a four-course progression that introduces AI through accessible applications before moving students toward deeper technical understanding and implementation. Grounded in cognitive apprenticeship theory, the sequence moves students from practitioner-oriented tool use to technologist-oriented configuration, to builder-oriented design and deployment of AI systems (Collins et al., 1989). Each course is composed of skill-cluster modules. A module is a one- to three-week instructional unit organized around related skills from the AIX Hub taxonomy and aligned with a specific phase of the apprenticeship sequence. Modules, rather than courses, are the primary unit of curriculum change, allowing the curriculum to respond to shifting labor-market signals by revising, replacing, or adding targeted modules while preserving overall course stability. This skills-to-modules-to-courses architecture supports adaptive alignment on a single-semester cycle and enables spoke institutions to adopt selected modules without replicating the full hub curriculum. Table 2 summarizes the four-course progression.

Course	Title	Level	Enrollment	Primary Tools	Anchor Outcome
Course 1	Foundations in AI	No-code	770 (6 sections)	Excel, Gemini, Tableau, Colab (read-only observation of code execution)	AI literacy & critical evaluation
Course 2	Data Science Foundations	Low-code	120 (2 sections)	Python/Colab, scikit-learn, prompting	Modeling & interpretation for business
Course 3	Agentic AI	No- and Low-code	87 (2 sections)	LangChain, n8n, RAG, vellum.ai, Zapier	Design & pitch an AI venture
Course 4	Deep Learning for Business	Full-code	90 (2 sections)	PyTorch, TensorFlow, Hugging Face	Deploy DL models on real business data

Table 2: The AIX Hub Curriculum Progression - Spring 2026

Course 1 (Foundations in AI) serves as the on-ramp and is required for all students, enrolling approximately 770 students per semester across six sections. The course relies primarily on no-code tools, including Excel for data manipulation, Google Gemini for generative AI interaction, Tableau for data visualization, and Google Colab notebooks in which students run pre-written Python code to observe model behavior. Its learning objectives emphasize conceptual understanding. Students must be able to trace data flow through an AI pipeline, distinguish between supervised and unsupervised learning, evaluate when human-in-the-loop decision-making is appropriate, and critically assess AI-generated outputs for bias and accuracy. Course 1's modules include Data-Pipeline Tracing, Generative-AI Interaction, Output Evaluation and Bias Detection, and Human-in-the-Loop Decision-Making, each delivered in a one- to three-week segment and aligned with the modeling-phase emphasis on observation and articulation.

Course 2 (Data Science Foundations) moves students to low-code learning. Students use Python in Google Colab with provided scaffolding to implement and interpret regression models, time-series forecasts, clustering algorithms, decision trees, and neural networks. Prompt engineering with LLMs is introduced as a distinct technical skill, positioning students to use GenAI tools as programming co-pilots. The instruction specifically teaches detailed prompting techniques, building on evidence that prompt engineering skill causally predicts AI output quality and that this skill requires progressive, scaffolded instruction. Course 2's modules include Regression and Time-Series Forecasting, Clustering and Decision-Tree Algorithms, Neural-Network Foundations, and Prompt Engineering for Programming. Each model is calibrated to the coaching-with-faded-scaffolds emphasis on guided implementation, articulation, and reflection. The capstone project requires students to select a business-domain dataset, apply modeling approaches, and present their findings in a format appropriate for a non-technical business audience.

Course 3 (Agentic AI) bridges no-code/low-code agent design with full-code development. The course is organized around the full agent lifecycle: designing prompts, connecting tools and data, orchestrating workflows, evaluating trust and reliability, and iterating toward deployable artifacts. Students first build agentic workflows through visual and no-code/low-code platforms, then extend these systems through customized functions, working code, debugging, and evaluation. Course 3's modules include Prompting and Agent Foundations, Visual and Low-Code Agent Building, Tool and Data Integration, Customized Functions and Full-Code Extension, Agent Evaluation and Trust, and Project Development. These modules are aligned with the articulation-and-reflection phase of the apprenticeship sequence because students must explain system behavior, diagnose reliability issues, compare performance, and justify design choices. The course culminates in a capstone project in which students design, develop, and pitch an AI venture that integrates agentic workflows with business applications. These ventures require students to translate technical capabilities into value-creating solutions, producing stakeholder-ready deliverables such as presentations, reports, and working artifacts (e.g., code, dashboards, or prototypes). In Spring 2026, ten course ventures were submitted to the university entrepreneurship accelerator (two were selected).

Course 4 (Deep Learning for Business) is the only course in the full-code tier. It prepares students to function as builders by designing, training, and deploying AI systems in production contexts. The course centers on the full-stack deployment of deep learning models using real business data, with hands-on experience in industry-standard tools such as PyTorch, TensorFlow, and Hugging Face. The curriculum spans multiple data modalities, including text, image, audio, and video, and emphasizes business applications such as fraud detection, recommender systems, and deepfake detection. Course 4's modules include Foundations of Deep Learning, Applied Data & Models, Business Domains & Applications, and Deployment & Strategy. These modules are aligned with the exploration phase of the apprenticeship sequence, emphasizing student-defined problems and full-code mastery. The course's anchor outcome is students' ability to deploy deep learning models on real business data.

Industry Partnership Model

The four-course progression was explicitly designed to align student capabilities with real organizational problems. Meeting this goal required an external partnership structure through which employers help shape WBL opportunities, skill priorities, and applied project contexts. The focal institution's preexisting employer relationships were primarily with human resources departments recruiting for traditional, non-AI internship roles on fixed hiring timelines. AI-focused WBL required the creation of new roles aligned with the skills developed in the curriculum and, therefore, new relationships with executive and functional leaders in organizations actively adopting AI. In this program, WBL opportunities are sourced through direct company relationships, a regional industry-focused talent development organization, and national workforce development organizations. These partnerships are coordinated through an industry consortium that advises on strategy, skills, and curriculum alignment while also supporting WBL placements. The overall process, supported by this governance structure, is illustrated in Figure 2.




Figure 2. Industry Partnership Model Flow

Figure 2 presents a stage-based model with bidirectional feedback for industry partnership development, showing the transition required to support AI-focused WBL. This model pivots the relationship with industry partners from a transactional, HR-driven approach to an actively governed, co-created one, supported by a formal Industry Consortium (Advisory Board). The model outlines four distinct stages:

- **Starting Point:** Existing relationships are primarily with HR departments and focused on traditional, non-AI internships. This stage is characterized by misalignment among pre-established hiring timelines, roles that are not aligned with AI skill development, and current AI workforce needs.
- **New Partnerships Required:** This stage marks a structural pivot from transactional recruiting to strategic co-design. Engagement moves beyond HR to executive leadership and functional leaders in AI-adopting organizations to co-create new AI-focused roles aligned with student IS and AI capabilities.
- **WBL Opportunity Sourcing:** A multi-channel sourcing model is established to scale opportunity creation, including Direct Company Relationships with senior leadership; leveraging the Regional Talent Partner (FinTech Academy) networks, and utilizing National Workforce Development Organizations to access broader pipelines.
- **Outcomes:** The model results in high-quality, AI-relevant WBL opportunities, leading to enhanced student learning and career readiness, and a stronger alignment between IS education and industry needs, forming a talent pipeline. The structure is bidirectional: partner companies and student capstone teams co-shape what entry-level AI work looks like in each industry domain—through capstone charters that frame AI problems with real business constraints, datasets that surface domain-specific challenges, and apprenticeship and placement positions that operationalize emerging role definitions—with the curriculum as the medium through which that co-shaping happens.
- The **Governance Layer**, composed of industry leaders from IS-intensive firms, is a cross-cutting mechanism that advises on strategy, guides curriculum alignment, and supports the creation and expansion of WBL opportunities across all stages.

Phase 1 of the industry partnership model focuses on the regional FinTech sector, selected because of the region's position as the U.S. payment transactions leader and the presence of major employers in proximity. Phase 1 partners include a payroll technology company, a media and technology conglomerate, a payment services company, a credit data company, a payment technology company, a global payment solutions firm, a prepaid payment technology company, a POS technology company, an AI consulting firm, the regional FinTech Academy, and a professional services firm (shown in Table 3). Partners span the full FinTech value chain to provide diverse project contexts across course levels.

Partner	WBL Type	Title / Project	Course(s)	Format
Credit data company	Internship	AI Data & Analytics Internship	Course 2, Course 3	Summer/fall, paid internship
Payment technology company	Capstone	AI Product Capstone (Fraud Intelligence)	Course 3	Semester-long capstone project
Global payment solutions firm	Internship	ML Engineer Internship	Course 2, Course 3	Summer, paid internship
Payroll technology company	Capstone	AI Automation Capstone	Course 4	Semester-long capstone project
POS technology company	Internship	POS AI Innovation Internship	Course 3, Course 4	Summer, paid internship
Media and technology conglomerate	Apprenticeship	Data Science Apprenticeship	Course 2	52-week structured apprenticeship
Regional FinTech Academy	Placement	AI Track Placement Program	All courses	10 positions per cycle, cohort placement
AI consulting firm	Project	AI Product Development Project	Course 4	Semester-long industry project

Table 3: Active Work-based Learning Opportunities - Spring 2026

Eight formalized WBL opportunities were active in Spring 2026, ranging from short capstone projects (one semester) to a longer data science apprenticeship. The WBL taxonomy distinguishes between internships (partner-hosted, compensated), capstones (project-based, course-embedded), apprenticeships (extended,

structured workplace learning), and placements (cohort-based partner hiring pipelines). The regional FinTech Academy's AI Track maintains ten dedicated placement positions per cycle. These differentiated formats allow students with varying scheduling constraints and financial needs to access authentic workplace experience.

Skills Taxonomy and Labor Market Alignment

A distinctive feature of the AIX Hub curriculum design process is its grounding in a multi-source, living skills taxonomy. We do not define course learning objectives based on faculty expertise or textbook structure. Instead, we constructed a 35-skill taxonomy spanning four domains: (1) Foundational AI (5 skills), (2) Technical AI (13 skills), (3) Agentic/Emerging AI (7 skills), and (4) Domain-Specific (10 skills across FinTech, Healthcare, IT/Cybersecurity, and Logistics). We cross-referenced each skill against four external labor market data sources: (1) a commercial labor market intelligence skill taxonomy, (2) O*NET occupation importance ratings, (3) ESCO v1.2 (European Skills/Competencies/Qualifications framework), and (4) BLS Occupational Employment Statistics wage and growth data¹.

This cross-referencing process served two purposes. First, it anchored curriculum decisions to demonstrate employer demand. Skills with high posting frequency on a commercial labor market intelligence platform and high O*NET importance ratings were prioritized for inclusion in required courses. Second, it enabled the team to assess the AI exposure of the occupations that students are likely to enter (Eloundou et al., 2023; Felten et al., 2023) and to design course content that specifically prepares students to work in high-exposure roles rather than be displaced by them. For example, the Foundations in AI course explicitly addresses Human-in-the-Loop (HITL) decision-making frameworks because HITL oversight is identified as a high-demand competency in occupations with elevated AI exposure (e.g., data analysts, financial analysts, compliance officers).

The skills taxonomy is maintained as a queryable database linked to both course learning objectives and industry partner demand signals, enabling continuous curriculum monitoring rather than relying on periodic, one-time crosswalks. Faculty review coverage gaps on a recurring basis, and the system flags skills with rapid increases in job posting frequency as candidates for curriculum updates. This infrastructure supports adaptive curriculum alignment in the face of skill volatility, defined as the capacity to detect labor-market shifts and incorporate them into course content within a single academic cycle, rather than waiting for the next scheduled curriculum review. The underlying premise is that AI capabilities and the associated skills required in business contexts evolve rapidly and continuously (Eloundou et al., 2023; Felten et al., 2023), rendering static, one-time mappings of skills to courses obsolete within the timeframe of a graduating cohort.

We present three illustrative skills elevated through this process between Fall 2025 and Spring 2026. First, HITL decision-making was integrated through the HITL Decision-Making module in Course 1: Foundations in AI, triggered by AI-exposure and employer-demand signals showing growing need for oversight and critical evaluation in data analyst, financial analyst, and compliance roles; Second, Prompt Engineering for Programming was integrated through the Prompt Engineering for Programming module in Course 2: Data Science Foundations, triggered by evidence that prompt quality predicts GenAI output quality and by observed student gaps in moving beyond one-sentence queries; Third, AI Workflow Orchestration was integrated through the Multi-Agent module into Course 3, driven by partner demand for systems that integrate multiple models, tools, and data sources into end-to-end business processes. In each case, the elevation moved a skill from advisory-only mention to required-module coverage on a single-semester cycle, demonstrating the pipeline operating end-to-end at the skills → modules → courses three-tier architecture.

Spring 2026 Outcomes

Enrollment and Reach

¹ Lightcast Skills Taxonomy, <https://lightcast.io/open-skills>; O*NET OnLine, <https://www.onetonline.org/>; ESCO v1.2, https://esco.ec.europa.eu/en/classification/skill_main; BLS Occupational Employment and Wage Statistics, <https://www.bls.gov/oes/>.

Spring 2026 marked the initial deployment of all four courses simultaneously, an ambitious schedule made possible by support from the industry-sponsored workforce-development initiative and the pre-existing infrastructure of the focal institution's IS department. Total enrollment reached 1,060 students across 12 course sections. Course 1 (Foundations in AI) represented the largest reach, with 770 enrolled students in six sections, establishing AI literacy as a de facto requirement for all IS-track business students. The two advanced courses (Course 3 and Course 4) each enrolled 80-90 students, consistent with upper-division elective capacity and reflecting genuine student demand for AI specialization. Causal evidence that large-scale GenAI integration interventions substantially improve student academic performance (Tsekouras et al., 2024) supports the case for deploying AI-integrated curricula at this scale rather than through limited single-section pilots. The outcomes reported in this section are implementation metrics that establish deployment feasibility at scale. Evaluation of the design against its three design requirements is planned for the Fall 2026 multi-site deployment (see Limitations and Future Directions).

Work-based Learning Pipeline

Eight WBL opportunities were active in Spring 2026, spanning four formats (internship, capstone, apprenticeship, placement) and eight industry partners. The capstone-format WBL projects were course-embedded, meaning all students in the linked courses participated in the project rather than a competitive subset. Competitive internship selection processes tend to disadvantage students who lack professional networks, whereas course-embedded capstones provide authentic employer exposure to all enrolled students simultaneously. This equity design principle aligns with experimental evidence that AI-powered teaching assistants reduce help-seeking avoidance among underrepresented students by providing anonymous, non-judgmental support that overcomes self-presentation and stereotype threat barriers (Gaskin et al., 2024).

A WBL opportunity pipeline was created for Summer 2026 for students who complete the AI courses in the Spring 2026 pilot, so they can acquire experiential learning before graduating. Each of the 87 students in Course 3, Agentic AI, completed a project with a real-world use case that demonstrates how the AI skills learned in class are applied to solve a business problem. Ten of these projects have been submitted to the university entrepreneurship venture program, and two have been selected.

Qualitative Project Insights

Qualitative analysis of course-embedded capstones, supported by faculty and industry partners, revealed a clear functional division of labor in the advanced courses (Agentic AI and Deep Learning for Business). Students primarily assumed the HITL role, focusing on high-level cognitive tasks such as problem framing, model selection, and critical evaluation of AI outputs. Conversely, the Agentic AI systems, utilizing orchestration frameworks, handled low-level, repetitive tasks, including data fetching, simple reasoning, and API tool selection. This division was intentional, aiming to move beyond simple prompting by focusing on translating complex business constraints into effective prompting strategies. Observed performance gaps centered on student difficulty in debugging multi-agent systems and evaluating system reliability.

This observation is presented as a preliminary, deployment-grounded proposition consistent with the design-principles framework established in the Research Design section: when AI curricula are structured around HITL principles in advanced courses, a predictable functional division of cognitive labor emerges, with learners specializing in problem framing and evaluation while the agentic system manages execution. The Spring 2026 capstones provide initial evidence for this proposition; observed performance gaps in debugging and reliability evaluation identify boundary conditions, specifically that the predicted division emerges when learners have meta-cognitive scaffolds for evaluating system reliability and diminishes when such scaffolds are absent. Future empirical research, including the Fall 2026 multi-site expansion, will be necessary to test this proposition across varying institutional contexts and to specify the metacognitive scaffolds that support it.

Research University Collaboration and Faculty Development

Each course in the AIX Hub has one or more research university collaborators who participated in curriculum design and co-taught or provided guest instruction. Collaborators include researchers specializing in AI education, machine learning, and human-AI interaction, assigned across the four courses based on domain expertise. This research university collaboration structure provided access to cutting-edge AI research and tools that would otherwise be difficult for a business school department to incorporate at

scale, and it created a professional development mechanism for the focal institution faculty in rapidly evolving technical areas.

The collaboration with two key research university partners follows a partner-led but hub-advised strategy. Engagement focuses on four key areas: (1) curriculum co-design, (2) knowledge transfer, (3) faculty development, and (4) student mentorship. Research collaborators from the partner universities worked directly with focal institution faculty to co-develop modules for courses such as Foundations in AI, Agentic AI, and Deep Learning for Business. This process included bi-weekly coordination calls and workshops to align the curriculum with a national AI skills taxonomy. Faculty professional development is further supported through just-in-time weekly training for each module and an instructor's tools repository in the program portal. Key innovations originating from this partnership include the development of a four-course "NCLC Ladder" progression and the integration of specialized frameworks such as Human-in-the-Loop decision-making and Agentic AI orchestration.

Discussion

Formalized Design Principles

Table 4 states the three design principles in the anatomy proposed by Gregor et al. (2020), specifying each principle's aim, mechanism, rationale, and boundary conditions. The principles are preliminary. They are grounded in kernel theory and one at-scale instantiation, and the Fall 2026 multi-site expansion will provide the first cross-context test of their transferability.

Design Principle	Aim, Mechanism, Rationale, and Boundary Conditions
DP1: Apprenticeship-Sequenced NCLC Progression	Aim: Build AI competence and meta-cognitive judgment in student cohorts with heterogeneous technical backgrounds. Mechanism: Structure instruction as a multi-course no-code-to-low-code-to-full-code progression in which each course instantiates a successive phase of cognitive apprenticeship (modeling, coaching, articulation and reflection, exploration). Rationale: Modeling must precede coaching, and articulation must precede exploration (Collins et al., 1989); no-code tools operationalize the modeling phase by exposing system behavior without syntactic load. Boundary conditions: Holds when tool tiers align with apprenticeship phases; weakens when learners exit the progression without scaffolds toward code-level understanding.
DP2: Bidirectional Demand-Signal Alignment	Aim: Keep curriculum content aligned with volatile AI skill demand within a single academic cycle. Mechanism: Maintain a skills-to-modules-to-courses architecture linked to a living, multi-source skills taxonomy and employer demand signals, with modules (not courses) as the unit of change and WBL partners co-shaping entry-level role definitions. Rationale: AI skills evolve faster than course- and degree-level governance cycles (Eloundou et al., 2023; Felten et al., 2023); the module is the largest curriculum unit that can change within one semester. Boundary conditions: Requires sustained data-engineering capacity and partnership governance; evidenced in one institutional context to date.
DP3: Equity-Relevant Hub-and-Spoke Deployment	Aim: Extend authentic AI work experience to all enrolled students rather than a self-selected subset. Mechanism: Deploy through a hub-and-spoke model in which spoke institutions adopt modules subject to skill-precedence constraints, and embed WBL in required courses through course-embedded capstones. Rationale: Competitive internship selection structurally disadvantages students without professional networks; course embedding democratizes access to authentic employer exposure (cf. Gaskin et al., 2024). Boundary conditions: Spoke adoption depends on local prerequisite structures; differential subgroup outcomes remain unmeasured pending the Fall 2026 expansion.

Table 4: Formalized Design Principles for AI Curricula in Business Schools

Theoretical Contributions

This study contributes to IS education research in three primary ways. First, it extends the no-code/low-code pedagogy literature by articulating a multi-course NCLC progression as a set of design principles for AI curricula in business schools, anchored in cognitive apprenticeship (Collins et al., 1989) as the kernel theory. The progression is derived from, rather than asserted by, the theoretical framework: the no-code-first sequence follows from the cognitive-apprenticeship requirement that modeling precedes coaching and

articulation precedes exploration, with each course representing a distinct phase of the apprenticeship sequence. This approach addresses critiques of NCLC-only pedagogies (Touretzky et al., 2019) by positioning no-code tools as the modeling-phase entry point, not the endpoint, and providing explicit scaffolding toward code-level understanding for students who pursue the full progression. While prior IS education research has demonstrated the value of individual AI-integrated courses (Grøder et al., 2022; Chen et al., 2023) and informal AI literacy tools (Pinski et al., 2023), the AIX Hub represents the first formalization of a multi-course NCLC progression as a set of design principles anchored in a kernel learning theory and implemented at collegiate scale across an entire business school curriculum.

Second, this study contributes to the industry-academic partnership literature by articulating a bidirectional demand-signals-to-curriculum infrastructure for adaptive curriculum alignment in the face of skill volatility. This infrastructure is a structured, technology-mediated process that continuously translates labor-market signals and employer demand into course content within a single academic cycle and engages partner companies through WBL to shape the entry-level AI roles to which the curriculum aligns, rather than passively receiving fixed role definitions. The contribution lies in the infrastructure itself, not in a single instantiation: the volatility premise is that AI capabilities and the skills required for their deployment in business contexts evolve rapidly and continuously (Eloundou et al., 2023; Felten et al., 2023), rendering periodic, one-time curriculum reviews structurally inadequate within the timeframe of a graduating class. Activity Theory research demonstrates that authentic, employer-sourced data and problems are essential for deeper WBL skill development (Aniceto & Corbett, 2025); the adaptive-alignment infrastructure systematically creates and updates these conditions at scale. Grounded theory evidence indicates that IS students currently address competency gaps through self-directed development due to curricular shortcomings (Ogundipe et al., 2023); the infrastructure is designed to proactively close these gaps. The skill-elevation examples in the Skills Taxonomy section illustrate the pipeline operating end-to-end within a single-semester cycle.

Third, this study contributes to the emerging literature on equity-relevant AI education by demonstrating that a progressive, industry-partnered AI curriculum can be deployed at scale in institutional contexts serving student populations historically underrepresented in technology fields, and by identifying course-embedded capstones as a specific equity-design feature. The focal-institution context (83% BIPOC, 45% first-generation) serves as the deployment site, but the primary contribution is the design decision to embed WBL in required courses, thereby extending authentic AI work experience to all enrolled students rather than restricting it to a self-selected subset typically favored by competitive internship processes (cf. Gaskin et al., 2024 on stereotype-threat barriers in self-presentation contexts). The inclusion of an HBCU and a community college system in the Fall 2026 expansion cohort represents a deliberate extension of this design choice beyond research-intensive institutions. Differential outcomes by student subgroups have not yet been measured and are identified as a boundary condition for future evaluation of these design principles.

Practical Implications

For IS educators pursuing similar initiatives, the AIX Hub design offers several practical implications. First, the industry partnership model necessitates sustained relationship management, including structured mechanisms for partner companies to participate in shaping entry-level AI role definitions through capstone charters and skills alignment—not only to host students. This relational work is often undervalued in faculty workload and promotion criteria; institutions should establish explicit mechanisms to recognize and support it. Second, the transition from no-code to low-code instruction requires curricular articulation grounded in cognitive apprenticeship logic—modeling-phase modules must precede coaching-phase modules—so that students completing only Course 1 acquire the modeling-and-articulation competencies (critical evaluation of AI outputs) that the kernel theory predicts are achievable in that phase. Third, operating an adaptive-alignment infrastructure under skill volatility demands ongoing data-engineering resources at the skills-to-modules layer, either through internal platform development or through commercial labor-market intelligence tools, which many business school IS departments currently lack. Finally, the alignment between the AIX Hub design and organizational GenAI competency training models (Schaetzle & Sohrabi, 2025) and emerging work on scalable and personalized AI-supported learning environments (Rai et al., 2024) indicates that these design principles may be transferable to corporate learning and executive-education contexts, presenting opportunities for future industry collaboration.

For administrators and accreditation bodies, the AIX Hub design suggests that AI literacy can be meaningfully integrated into business school curricula without displacing existing core content, provided

courses are designed with appropriate scope and sequencing (Course 1 focuses on AI concepts and critical evaluation, not programming). The hub-and-spoke multi-institutional model also offers a mechanism for extending AI education access to partner institutions: spokes integrate hub modules into their existing curricula subject to skill-precedence constraints, rather than developing AI courses from scratch—a path that fits institutions whose faculty capacity, program structure, or student-prerequisite profile does not support full hub-course adoption.

Limitations and Future Directions

Several limitations of this study should be acknowledged. The Spring 2026 data represent a single-semester deployment at the focal regional hub. Longer-term outcomes, including graduate employment rates, employer satisfaction with AI competencies, and student retention in AI career pathways, require longitudinal follow-up. Although the artifact has been designed across two R1 contexts (the focal regional hub and the focal national hub), only the focal regional hub has instantiated the design at the time of this writing; cross-instantiation evidence at the focal national hub and at the spoke institutions is not yet available, limiting empirical generalizability. The focal national hub's instantiation and the Fall 2026 expansion to a Community College Partner, a local Regional University Partner, and an HBCU Partner will together provide the first cross-instantiation evidence for the design principles' transferability.

Additionally, the paper reports design and early process outcomes rather than rigorous student learning assessments. Future work should incorporate supported AI literacy assessment instruments to measure actual learning gains. The Fall 2026 multi-site deployment will evaluate the design against its three design requirements using grade distributions by course tier, completion rates by prior-programming background, capstone assessment rubric results, structured partner feedback, and student AI-usage and classroom-experience data collected through the AIX Hub platform. Qualitative research on students' use of ChatGPT reveals that sociocultural factors and academic skepticism shape how students engage with AI tools (Katavic et al., 2023), suggesting that students' experiences in the AIX Hub curriculum may vary in ways not captured by enrollment and completion metrics. For example, future studies can include ethnographic or interview-based methods. Finally, Faisal et al. (2025) propose Practice-Ready Transformation (the internalization of resilience, initiative, and collaborative agility) as a fifth stage of experiential learning. AIX Hubs should empirically validate if the curricula achieve this level of professional identity formation, which will require longitudinal follow-up well beyond the Spring 2026 deployment examined here.

Furthermore, several additional limitations warrant consideration. First, selection effects may be present in participation in WBL, as students who opt into competitive internships or advanced capstones may have higher baseline motivation or pre-existing professional networks. Second, there are potential confounds in attributing student outcomes solely to the curriculum, as individual student characteristics—such as prior technical experience or quantitative aptitude—may substantially influence their success. Finally, measuring "soft" outcomes, such as a student's capacity for critical AI evaluation and the quality of human-in-the-loop judgment, poses ongoing challenges, as these competencies are more difficult to quantify than technical skills.

Conclusion

This study presents a set of preliminary design principles for industry-partnered AI curriculum development in business schools, instantiated in the AIX Hub—co-designed by two R1 hub institutions and supported by an industry-sponsored workforce-development initiative. The design articulates a no-code to low-code to full-code learning progression (the NCLC Ladder) anchored in cognitive apprenticeship, a demand-signals-to-curriculum infrastructure for adaptive curriculum alignment under skill volatility operating at a skills-to-modules-to-courses three-tier architecture, and a hub-and-spoke deployment model that situates the curriculum in equity-relevant institutional contexts and employs course-embedded capstones to extend authentic AI work experience to all enrolled students. The Spring 2026 deployment at the focal regional hub enrolled 1,060 undergraduate students across four courses and engaged eight FinTech industry partners in curriculum co-design and WBL. The focal national hub's instantiation and Fall 2026 expansion to spoke partner institutions are underway, with broader diffusion to additional R1 hubs in progress.

The AIX Hub offers a potential solution to a central challenge facing IS education in the coming decade: preparing business students for an AI-permeated workplace without sacrificing technical depth or

excluding non-programming majors. This design provides a model of derived progression, adaptive industry co-design, and equity-relevant deployment. The AIX Hub design serves as an institutional mechanism for bridging this gap and enabling transformation at scale—not through a single course or AI tool, but through a coordinated, industry-partnered, labor-market-aligned curriculum architecture accessible to all business students, regardless of prior technical background. These design principles are presented as a working prototype, intended for critique, replication, and refinement by the IS education community.

References

- Aniceto, J., & Corbett, J. (2025). Data literacy development in hackathons: An activity theory perspective. *Proceedings of the International Conference on Information Systems*.
- Billett, S. (2011). Vocational education: Purposes, traditions, and prospects. Springer.
- Bock, A. C., & Frank, U. (2021). Low-code platform. *Business & Information Systems Engineering*, 63(6), 733–740.
- Brynjolfsson, E. (2022). The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence. *Dædalus*, 151(2), 272–287.
- Brynjolfsson, E., Chandar, B., & Chen, R. (2025a). *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Stanford Digital Economy Lab.
- Brynjolfsson, E., Li, D., & Raymond, L. (2025b). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942.
- Chen, Y., Albert, L. J., & Macias, H. (2023). Prototyping AI-powered social innovation in an undergraduate MIS course. *Proceedings of the International Conference on Information Systems*.
- Collins, A., Brown, J. S., & Newman, S. E. (1989). Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. In L. B. Resnick (Ed.), *Knowing, learning, and instruction: Essays in honor of Robert Glaser* (pp. 453–494). Lawrence Erlbaum Associates.
- Cooper, L., Orrell, J., & Bowden, M. (2010). *Work integrated learning: A guide to effective practice*. Routledge.
- Duan, Y., Edwards, J. S., & Dwivedi, Y. K. (2019). Artificial intelligence for decision making in the era of Big Data: Evolution, challenges and research agenda. *International Journal of Information Management*, 48, 63–71.
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. arXiv preprint arXiv:2303.10130.
- Faisal, N., Chadhar, M., Gorris-Hunter, A., & Stranieri, A. (2025). Towards a fifth stage of the experiential learning cycle: Extending ELT through simulation to foster practice-ready transformation in IS education. *Proceedings of the International Conference on Information Systems*.
- Felten, E., Raj, M., & Seamans, R. (2023). How will language modelers like ChatGPT affect occupations and industries? arXiv preprint arXiv:2303.01157.
- Gaskin, J., Blondeel, E., Schuetzler, R., Serre, R., Steffen, J., & Wood, D. (2024). Chatbots mitigate help-seeking avoidance. *Proceedings of the International Conference on Information Systems*.
- Gmyrek, P., Berg, J., & Bescond, D. (2023). *Generative AI and jobs: A global analysis of potential effects on job quantity and quality*. International Labour Organization.
- Gregor, S., Chandra Kruse, L., & Seidel, S. (2020). Research perspectives: The anatomy of a design principle. *Journal of the Association for Information Systems*, 21(6), 1622–1652.
- Gregor, S., & Jones, D. (2007). The anatomy of a design theory. *Journal of the Association for Information Systems*, 8(5), 312–335.
- Grøder, C. H., Schmäger, S., Parmiggiani, E., Vassilakopoulou, P., Pappas, I., & Papavlasopoulou, S. (2022). Educating about responsible AI in IS: Designing a course based on experiential learning. *Proceedings of the International Conference on Information Systems*.
- Hevner, A. R., & Chatterjee, S. (2010). *Design research in information systems: Theory and practice*. Springer.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105.
- Jackson, D., & Collings, D. (2018). The influence of work-integrated learning and paid work during studies on graduate employment and underemployment. *Higher Education*, 76(3), 403–425.

- Kandlhofer, M., Steinbauer, G., Hirschmugl-Gaisch, S., & Huber, P. (2016). Artificial intelligence and computer science in education: From kindergarten to university. *Proceedings of the IEEE Frontiers in Education Conference*, 1-9.
- Katavic, R., Pahuja, A., & Syed, T. A. (2023). Navigating the use of ChatGPT in classrooms: A study of student experiences. *Proceedings of the International Conference on Information Systems*.
- Kishore, S., Hong, Y., Nguyen, A., & Qutab, S. (2023). Should ChatGPT be banned at schools? Organizing visions for generative artificial intelligence in education. *Proceedings of the International Conference on Information Systems*.
- Liu, Y., Cui, W., Wu, S., & Tan, W. K. (2025). Generative AI in education: A review of applications, impacts, and future research directions. *Proceedings of the International Conference on Information Systems*.
- Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1-16.
- McQuaid, R. W., & Lindsay, C. (2005). The concept of employability. *Urban Studies*, 42(2), 197-219.
- Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, 100041.
- OECD. (2024). *OECD employment outlook 2024*. OECD Publishing.
- Ogundipe, T. V., Brown, I., & Johnston, K. (2023). Competency maturing grounded theory. *Proceedings of the International Conference on Information Systems*.
- Pinski, M., Haas, M., & Franz, A. (2023). AiLingo: A design science approach to advancing non-expert adults' AI literacy. *Proceedings of the International Conference on Information Systems*.
- Rai, A. (2017). Editor's comments: Diversity of design science research. *MIS Quarterly*, 41(1), iii-xviii.
- Rai, A., Chen, L., Breazeal, C., Ramesh, B., Long, Y., & Aria, A. (2024). Design and evaluation attributes for scalable, cost-effective personalization of LLM tutors in programming education. *Proceedings of the International Conference on Information Systems*.
- Ramos, I., Mesquita, A., Lopes, A. L., Moreira, F., Barros, V., Sousa, R. D., Regadas, G., Gomes, L., & Leite, P. (2025). Studying the GenAI's impact on learning and critical thinking: A multi-site field experiment. *Proceedings of the International Conference on Information Systems*.
- Rowe, A. D., Mackaway, J., & Winchester-Seeto, T. (2012). 'But I thought WIL was just work experience': Clarifying the terminology and investigating understanding of work-integrated learning in the workplace. *Asia-Pacific Journal of Cooperative Education*, 13(4), 203-217.
- Sajadieh, S., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Santarlasci, L., Pava, J., Maslej, N., Altman, R., Brynjolfsson, E., Brodley, C., Clark, J., Dignum, V., Kumar, V., Landay, J., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., ... Weld, D. (2026). The AI Index 2026 Annual Report. Stanford Institute for Human-Centered Artificial Intelligence, Stanford University.
- Schaetzle, A., & Sohrabi, S. (2025). Empowering employees: Designing a training concept for GenAI competence. *Proceedings of the International Conference on Information Systems*.
- Sweller, J., Van Merriënboer, J. J., & Paas, F. G. (1998). Cognitive architecture and instructional design. *Educational Psychology Review*, 10(3), 251-296.
- Tolzin, A., Knoth, N., & Janson, A. (2024). Leveraging prompting guides as worked examples for advanced prompt engineering strategies. *Proceedings of the International Conference on Information Systems*.
- Tomlinson, K., Jaffe, S., Wang, W., Counts, S., & Suri, S. (2025). *Working with AI: Measuring the applicability of generative AI to occupations*. Microsoft Research.
- Topi, H., Valacich, J. S., Wright, R. T., Kaiser, K. M., Nunamaker, J. F., Sipior, J. C., & de Vreede, G. J. (2017). IS 2017: Curriculum guidelines for baccalaureate degree programs in information systems. *Association for Computing Machinery and Association for Information Systems*.
- Touretzky, D., Gardner-McCune, C., Martin, F., & Seehorn, D. (2019). Envisioning AI for K-12: What should every child know about AI? *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(1), 9795-9799.
- Tsekouras, D., Belo, R., & Cornelius, P. (2024). Generative AI and student performance: Evidence from a large-scale intervention. *Proceedings of the International Conference on Information Systems*.
- Waszkowski, R. (2019). Low-code platform for automating business processes in manufacturing. *IFAC-PapersOnLine*, 52(10), 376-381.
- Wing, J. M. (2006). Computational thinking. *Communications of the ACM*, 49(3), 33-35.
- 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), 1-27.