

Orchestrating Scaffolding AI Agents: Design Principles for Mechanism-Specific Learner Support

Diana Kozachek^{[0000–0001–8302–1007]*} and Andreas Janson^[0000–0003–3149–0340]

Institute of Information Systems, University of St.Gallen, St.Gallen, Switzerland
`{diana.kozachek, andreas.janson}@unisg.ch`

Abstract. AI agents are increasingly adopted in higher education, yet current systems handle requests uniformly, promoting cognitive offloading over sustained skill development. With the rise of orchestrated multi-agent systems, learning goals can be targeted by specialized AI agents that each address a distinct scaffolding mechanism: conceptual, procedural, strategic, or metacognitive. This study follows a Design Science Research approach to derive design requirements from 32 student interviews, iteratively refine a prototype with 22 IS experts, 6 educators, and 34 students, and computationally analyze scaffold effectiveness. We contribute three design principles: (1) multi-agent coordination through a lead orchestrator that delegates to mechanism-specific sub-agents, (2) adaptive learner profiling that enables cross-session scaffolding fading, and (3) continuous institutional knowledge integration grounding scaffolds in verified course materials. Our effectiveness analysis reveals that delivery order predicted learning behavior, positioning orchestration as a key design concern for AI-assisted educational systems.

Keywords: AI Agents · Orchestration · Design Science Research.

1 Introduction

Consider a novice analyzing a business case: they start understanding frameworks, navigate multiple resources, develop a solution and a strategy towards it, and ideally, they reflect on their approach. Each of these steps requires distinct support, rather than generic answers [25].

Nowadays, learners often consult Generative AI agents, autonomous, goal-directed systems that can adapt to user needs and coordinate tools across tasks, to identify key information and generate results [9,14]. AI agent-assisted learning has emerged as the fourth most prevalent use case globally [58], with recent implementations in higher education supporting foreign language acquisition [60], innovation management [12], and software development [40]. And while many tools already hold the potential to personalize learning experiences, especially when tutors are scarce [32,43], generic AI tools are not built with human learning

* ✉ Corresponding author

as their core functionality [46]. Recent studies suggest that this lack of pedagogic grounding may harm the learning process by motivating cognitive offloading and over-reliance [26,59], indicating that mechanism-specific support has yet to transcend knowledge-transfer to foster sustained skill-development [46].

AI agents that adapt to human learning are designed to follow pedagogic frameworks, such as the Zone of Proximal Development (ZPD) [50], defined as the space between a learners’ independent skills and what they can achieve with targeted support from a more knowledgeable other [56]. Within a learners’ individual ZPD, scaffolding is a method that structures temporary and progressively adjusted support by educators [41]. Aiming to gradually reduce a learners’ dependence to facilitate skill development, conceptual, procedural, strategic, and metacognitive scaffolding mechanisms each activate a distinct cognitive process [3,23]. While a human teacher can easily adapt support mechanisms throughout a learning session [41], scaffolding AI agents need to leverage architectural possibilities to move beyond isolated conversational interventions, into dynamic, mechanism-specific tutoring encounters [8,9]. One of these architectural approaches consider orchestration as framework to streamline multiple AI agents’ internal planning logic into a unified workflow [1], capable of supporting humans in a targeted manner [15]. As multiple specialized AI agents are needed [57] to scaffold the right cognitive process at the right time [5], orchestration poses a useful mechanism to prioritize and sequence tasks tied to learning [15,21]. This leads to the following research question **(RQ): How can AI agents be designed to orchestrate scaffolding mechanisms for sustained learning support?**

To address this, we follow a Design Science Research (DSR) approach by Pefers et al. [39], emphasizing iterative artifact development, grounded in real-world problems and user needs [28]. Scaffolding poses as a kernel theory that informs implementation, the design rationale is derived from orchestrating distinct scaffolding AI agents, targeting cognitive processes [24]. Orchestration becomes the novel technical lens to improve the issue with AI-agent assisted scaffolding towards mechanism-specific tutoring [16]. This study elaborates on the first design cycle: the introduction covers problem identification and objective definition. The design and development is discussed in the method section, informed by management student interviews (N=32), a demonstration with experts from IS research (N=22), and a system evaluation by educators (N=6) and a focus group of students (N=34), as well as a computational analysis of the effectiveness of AI agent scaffolds. This study contributes three design principles that specify how orchestration coordinates scaffolding AI agents: (1) multi-agent coordination through a lead orchestrator that delegates to mechanism-specific sub-agents, (2) adaptive learner profiling that enables cross-session fading of support, and (3) continuous institutional knowledge integration that grounds scaffolds in verified course materials. These principles, operationalized through five design features, offer reusable design knowledge for AI-agent assisted educational systems that need to coordinate multiple support mechanisms.

2 Conceptual Background

2.1 AI Agents

Artificial intelligence refers to artifacts that demonstrate reasoning, learning, or decision-making [4]. With the introduction of agentic capabilities into Generative AI, AI agents are regarded as the next step in the evolution of the frontier technology [62]. We define AI agents through four capabilities that distinguish them from static prompt-response systems [42, 57]: **(1)** operating across domains and time horizons (environmental complexity), **(2)** prioritizing understanding over task completion (goal complexity), **(3)** responding to shifting user needs (adaptability), and **(4)** coordinating tools and planning to sequence tasks effectively (independent task execution). While generic AI agents display an absence of coordination toward overarching goals [57], this architectural shortcoming has prevented differentiated scaffolding aligned with educational theory [9]; novel orchestration mechanisms for multi-agent communication, organization, and task coordination, including tool usage and sustained memory [57], indicate that this deficit is solvable and may present a solution to a known challenge [16].

Orchestration frameworks manage AI agents’ planning logic within a workflow [1], addressing delegation of tasks and roles in Multi-Agent-Systems [20, 33]. Systems containing multiple AI agents pose design challenges regarding task and role delegation [1, 7, 29]. Since scaffolding requires the coordinated delivery of distinct support mechanisms, the delegation provided by orchestration frameworks offers a way to assign mechanisms to specialized agents and coordinate their deployment based on learner state [6, 21]. We refer to this as *mechanism-specific* support: the targeted activation of a particular scaffold type, combined with orchestration logic that determines delivery order and coordinates across agents.

2.2 Scaffolding AI Agents

Tutoring encounters, whether human- or AI-mediated, deliver scaffolding interventions [23, 41], catalyzing external support mechanisms into self-regulatory competencies [2]. AI agents have demonstrated effectiveness in delivering instruction that adapt support frequency support based on domain knowledge [9, 10, 23], but have failed to implement multiple support mechanism at once [3]. Facilitating mechanism-specific learner-AI interaction necessitates a scaffolding framework that is attuned to orchestrated AI agents, since advanced systems introduce novel possibilities such as sustained learner-adaption [8, 32]. Literature identifies four scaffold types to structure scaffolds [3, 24]:

1. Conceptual scaffolds highlight key knowledge and relationships between frameworks and information [5, 43].
2. Procedural scaffolds guide navigation of learning environments, including tools, resources and functionality [22].
3. Strategic scaffolds help learners to select and sequence the learning process in terms of planning and problem-solving [5, 36].
4. Metacognitive scaffolds prompt learners to plan, monitor, and regulate thinking along the process [2], mediating not only knowledge transmission [54], but also motivational support [10].

While each scaffold mechanism addresses a distinct cognitive need, effective AI agent-assisted educational systems requires a coordinated deployment: a learner struggling with a concept may first need metacognitive prompting to recognize their gap, then conceptual explanation to address it, and finally strategic guidance to integrate their ideas. The rise of AI agent orchestration enables novel directions to develop design knowledge for such mechanism-specific support. The following section presents our methodological approach to deriving design requirements, principles, and features for orchestrated scaffolding AI agents.

3 Method and Design

3.1 Setting, Participants and Procedure

We aim to develop an artifact that supports management students during the two-semester qualification phase of a leading European business school. This phase challenges novices to unify and conceptualize information across subjects and contexts. To ensure a user-centered design, we recruited 32 students (13 female, 19 male; ages 19–25; $M=23.2$) who had successfully completed this phase, and were enrolled in the third semester or beyond of their bachelor’s program. Prior to the study, participants completed a beginner’s course on AI agent fundamentals to ensure a shared conceptual baseline. Each unpaid, semi-structured interview lasted approximately 30 minutes. The approach follows the informative and consultative user involvement steps proposed by [28], emphasizing relevance and user-centricity as guiding principles in DSR.

The interview protocol consisted of three parts: (1) five questions on AI usage and observed limitations when working on business cases, (2) nine theory-informed use cases derived from AI agent theory, and (3) five open-ended prompts to elicit additional needs. Interviews were transcribed and analyzed using focused coding, based on thematic analysis of orchestrated AI agents and scaffolding theory [3, 23], and refined based on emergent themes. The coding was conducted by the first author and reviewed by another researcher to ensure consistency. Themes were mapped to design requirements through iterative comparison with theory, ensuring that they reflected conceptual grounding and user relevance [13].

3.2 Deriving Scaffolding AI Agent Design Requirements

Building on the scope of AI agents, each orchestration functionality has been mapped to a corresponding scaffold mechanism, divided into strategic, conceptual, procedural, and metacognitive scaffolds. Table 1 summarizes the mapping rationale, linking AI agent requirements to learner-facing use cases and scaffold types. To complement the theoretical mapping with empirical exploration, we conducted user interviews ($N=32$) on predefined technical functionalities to validate relevance [19]. By combining theory-driven mapping with user-centered validation, we aim to ensure that the resulting scaffolding structures align with the experiences and expectations of learners navigating continuous, personalized learning journeys. The interviews result in two core observations:

Observation 1: Interaction with non-educational AI agents invite cognitive offloading [46, 47]. Generic AI agents (e.g., ChatGPT) cannot infer which cognitive mechanisms support effective learning from prompts alone, especially novice students default to simple copy-pasting usage without reflection on their broader learning journey. This user behavior promotes cognitive offloading and undermines problem-solving and critical thinking [3, 26, 59].

Observation 2: AI agent implementation forces continuous re-prompting. Even if students intend to engage with resources more deeply, current AI agents operate as standalone session without domain specialization, contextual memory, or coordination across support mechanisms [46, 61], resulting in generic task-completion [42].

Table 1. Learner-facing design requirements mapped to scaffold types.

Scaffold Type	Learner Needs (Interviews)	Design Requirement	Testability Criteria
Conceptual: Clarifying domain knowledge and concept relationships	“Explains better than professors; adapts to my speed” (I1, I4)	DR1: The AI agent shall provide domain-specific explanations that adapt in complexity to the learner’s demonstrated understanding level.	Explanation accuracy; adaptation frequency
Strategic: Supporting planning and task decomposition	“Don’t know where to start; helps me structure” (I4, I5, I8)	DR2: The AI agent shall support learners in decomposing complex tasks into manageable steps and sequencing activities.	Plan quality ratings; task completion rates
Procedural: Guiding resource and tool usage	“Knows what was covered in lecture; connected to old exams” (I4, I5)	DR3: Learners shall receive scaffolds grounded in verified institutional materials relevant to their current task.	Retrieval accuracy; guidance uptake
Metacognitive: Sustaining self-monitoring and reflection	“Tracks where I’m stuck; recognizes my mistake patterns” (I5, I8)	DR4: Learners shall receive reflection prompts adapted to their demonstrated learning patterns and progress over time.	Profile-accuracy; scaffolding fading rate

We derive four learner-facing design requirements (DR1–DR4) from the interview findings, following de Weck’s criteria for clarity, testability, and feasibility [52]. DR1 and DR2 address the absence of context-specific cognitive support identified in Observation 1: scaffolding theory requires that support be contingent on the learner’s current understanding [41], which generic AI agents cannot assess from prompts alone [3]. DR3 and DR4 address the system-level deficits identified in Observation 2: effective scaffolding depends on continuity across learning sessions [56] and alignment with institutional curricula [24], neither of which standalone AI agents provide [42]. The orchestration of the scaffolding mechanisms, describing how and when they are delivered, is addressed below.

3.3 Developing Design Principles and Features

The four learner-facing design requirements inform three system-level design principles (DP1–DP3) (Table 2), formulated according to Gregor et al.’s anatomy

of design principles [17]. Each principle specifies its mechanism, aim, context, and theoretical rationale. We then derive five design features (DF1–DF5) that translate these principles into an orchestrated AI agent architecture (Figure 1).

DP1: Orchestrated Multi-Agent Coordination *Derived from DR1 and DR2.* For scaffolding AI agents to achieve coherent support across diverse cognitive tasks in higher education contexts, implement a lead orchestrator agent that delegates to specialized sub-agents based on scaffold type, as learners require seamless transitions between conceptual, strategic, procedural, and metacognitive support without context loss [57]. This principle enacts the scaffold functions of *contingency* (responding to learner state) and *intersubjectivity* (maintaining shared understanding across agents). DP1 is implemented through DF1 (Lead Orchestrator Agent) and DF2 (Specialized Sub-Agents).

DP2: Adaptive Learner Profiling *Derived from DR1 and DR4.* For scaffolding AI agents to achieve appropriate support intensity in sustained learning activities, implement persistent learner profiles that track demonstrated understanding and learning patterns, because scaffolding effectiveness depends on accurate assessment of the learner’s zone of proximal development [50]. This principle enacts the scaffold functions of *fading* (reducing support as competence grows) and *transfer* (applying insights across domains) [56]. DP2 is implemented through DF3 (Learner Memory Store).

DP3: Continuous Institutional Knowledge Integration *Derived from DR3 and DR4.* For scaffolding AI agents to achieve curriculum-aligned support in institutional learning environments, implement verified connections to course materials and educator oversight mechanisms [9, 55], because learners require scaffolds grounded in approved content, while instructors require visibility into AI-mediated learning [24]. This principle enacts the scaffold functions of *procedural guidance* (directing resource use) and *recruitment* (maintaining orientation toward institutional goals) [41]. DP3 is implemented through DF4 (Institutional Resource Connector) and DF5 (Educator Oversight Dashboard).

Table 2. Mapping of design requirements to principles and features.

Design Requirement	Design Principle	Design Feature	Scaffold Function
DR1: Adaptive domain explanation	DP1: Multi-Agent Coordination	DF1: Lead Orchestrator	Contingency
DR2: Task decomposition support		DF2: Specialized Sub-Agents	Intersubjectivity
DR1: Adaptive domain explanation	DP2: Adaptive Learner Profiling	DF3: Learner Memory Store	Fading
DR4: Pattern-based reflection			Transfer
DR3: Curriculum-grounded scaffolds	DP3: Institutional Knowledge Integration	DF4: Resource Connector	Procedural guidance
DR4: Pattern-based reflection		DF5: Educator Dashboard	Recruitment




Fig. 1. Design features of the scaffolding AI agent architecture.

Scaffolding theory anticipates gradual *fade-out* as learners develop competence within their ZDP [56]. DF3 enables this by tracking skill progression across sessions, allowing the orchestrator (DF1) to reduce intervention intensity for mastered topics. Interview participants emphasized the importance of human oversight during this transition, as they value the expertise of an educator over the generic explanations they would receive otherwise: DF5 ensures educators can monitor fade-out timing and intervene when the AI misjudges learner readiness.

3.4 Demonstrating the Scaffolding AI Agent Prototype

System Description and Orchestration Logic The system [27] implements a hierarchical multi-agent architecture (Figure 1). A **lead agent (DF1)** serves as the orchestrator, managing the workflow, sub-agent activation, and learner state. Upon each concept map submission, the orchestrator follows a fixed delegation sequence: it first dispatches input to the content ingestion agent, which parses and tracks concept map changes across rounds; then to the example map agent, which compares the learner’s map against an expert reference without revealing it; and finally to the scaffolding agent assigned to the current round. The active **scaffolding agent (DF2)** is determined by the experimental condition, which assigns one of four specialized agents per round. Each scaffolding agent adjusts its support intensity (high, medium, low) based on the **learner’s memory store (DF3)**, which the profiling component updates by analyzing concept map complexity relative to previous submissions. This enables scaffolding fading: as the learner’s assessed competence increases, support intensity decreases. Agent activation and deactivation are controlled through the **educators’ configuration (DF5)**, allowing design features to be enabled or disabled for experimental conditions, which is directly tied to the **resource connector (DF4)** on a task content-level.

The environment is built on Python 3.12 [48] with a Streamlit interface [45], and integrates LLM-based natural language generation for agent responses, MongoDB [34] for data persistence, and Cytoscape [38] for real-time concept map interaction. A progress tracking dashboard provides visual feedback on learner development across rounds.

Task Description The concept mapping exercise focused on a fictional market entry challenge for a foreign software startup to prevent external knowledge contamination and ensure experimental control (Figure 2). "*Adaptive Market Gatekeeping (AMG)*" could neither be found online, nor did generic AI agents provide suitable information on it. The task required participants to construct concept maps illustrating the interconnections between four core domains: Market analysis, resources and capital, entry strategies, and the fictional AMG mechanisms: dynamic adaptation, rule-changing, network control, and resource blocking by incumbent firms. Concept mapping externalizes learners' knowledge structures and misconceptions in a machine-parseable graph [37] and has strong evidence for learning benefits and assessment validity [22, 53]. Scaffolding in concept mapping tasks demonstrates feasibility to provide structured support [31]

Demonstration Procedure We recruited 22 experts from IS research (9 female, 13 male) with experience in design and implementation of GenAI-based systems from a leading European Business School. Experts tested the prototype individually by working through the case task, providing think-aloud commentary on system interaction (Figure 3). Sessions lasted one hour and were screen-recorded. During this phase, experts identified more than 20 usability barriers at the interface level, which were resolved prior to proceeding to the evaluation.

Stage 1: IS Expert Evaluation of Design Principles The same 22 experts participated in a structured evaluation session assessing principle validity and feature utility. Our evaluation follows FEDS [49], adopting a Human Risk & Effectiveness strategy appropriate for educational AI deployment. Experts received descriptions of each design principle (DP1–DP3) and their corresponding features (DF1–DF5), then rated: (1) whether features adequately operationalized their parent principles (*principle–feature traceability*); (2) whether the orchestration logic produced appropriate scaffold sequencing (*DP1 validity*); (3) whether learner profiles enabled meaningful adaptation (*DP2 validity*); and (4) whether institutional resource integration maintained curriculum alignment (*DP3 validity*). Feedback was collected in a follow-up discussion. This assessment identified two refinement categories: *system persistence* to improve memory, so learner profiles persisted across sessions, and *scaffolding effectiveness*, clarifying the lead agents' orchestration logic was transparent to evaluators.

Stage 2: Educators Evaluation of Scaffolding AI Agent Output To ensure that the AI agents delivered the intended scaffold types, six educators (all female), each experienced in scaffolding novice learners, blindly classified agent outputs by mechanism (conceptual, procedural, metacognitive, strategic) through focused coding. Ratings were conducted in rotating groups of three;

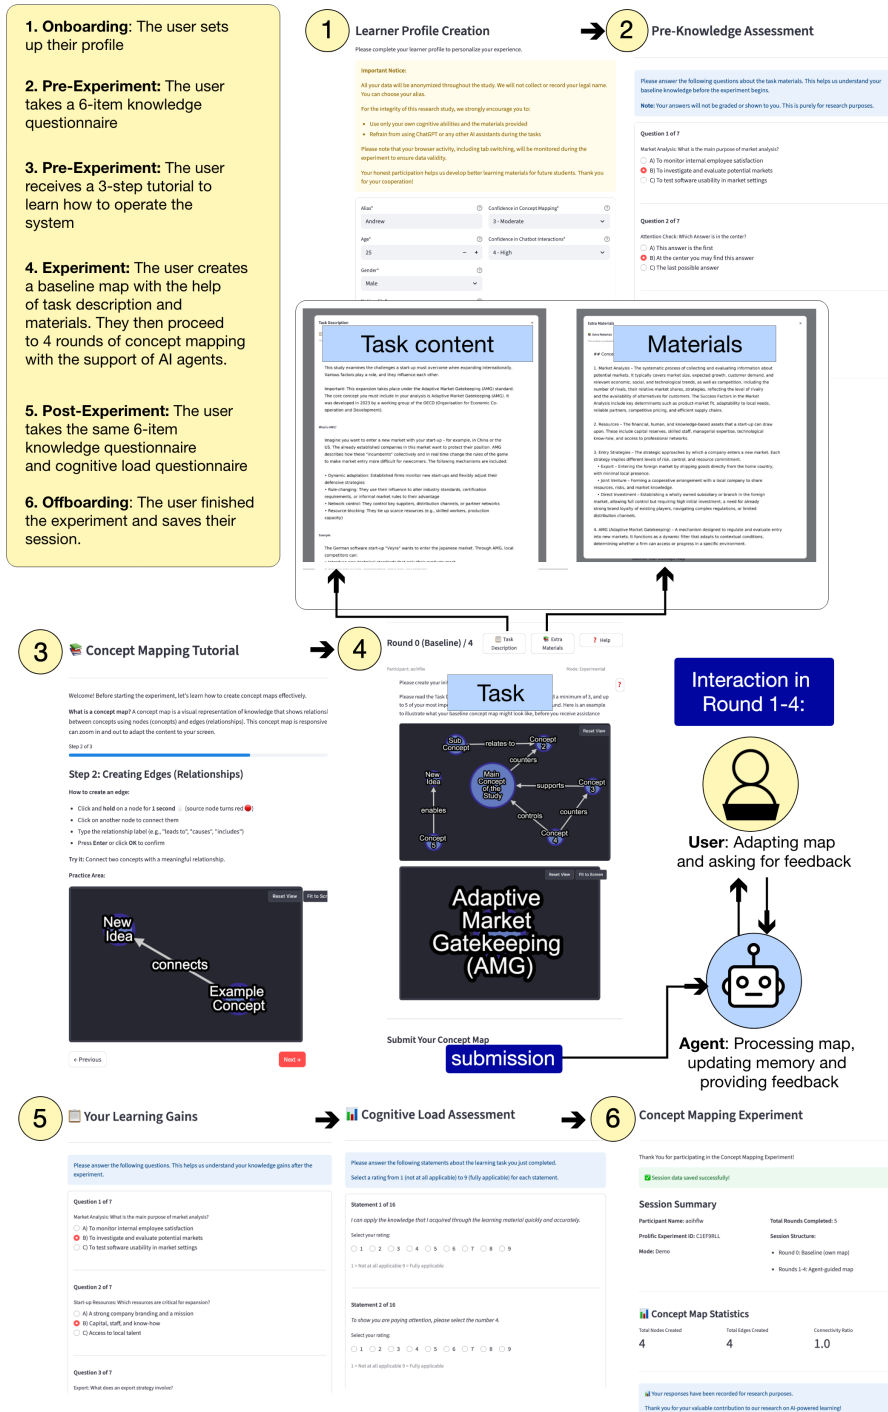


Fig. 2. Implemented design features within a workflow

only prompts with consistent classifications were implemented. Educators then provided coded edits to the prompt-templates [47] to increase contextualization for the problem-solving task (Table 3). Their adaptations aim to cater towards more contextualized learner feedback during the problem-solving task.

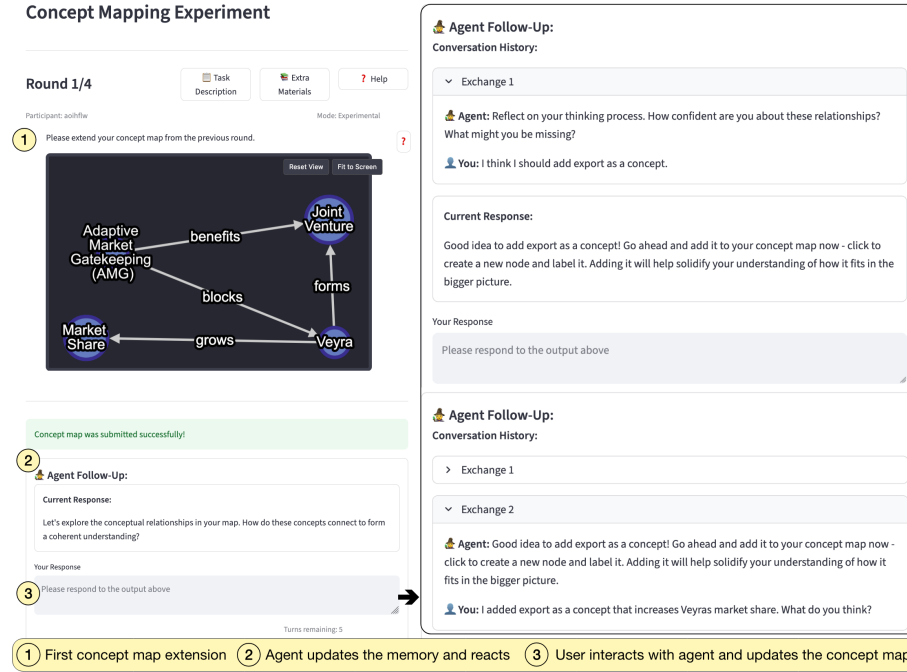


Fig. 3. Prototype displaying the design features within in an interaction

Table 3. Prompt templates of the scaffold types iteratively refined by educators (N=6).

Prompt before edits	Prompt after edits
Conceptual: Can you explain the difference between {concept A} and {concept B}?	The concept of {concept A} relates to {concept B} in international markets. How does this theoretical understanding change how you see these relationships?
Procedural: Have you considered alternative ways to organize these concepts? What might work better?	I notice {observation}. Let me walk you through a method for adding complex relationships: (1) identify the placement of {concept A}, (2) determine the relationship with {node A}, (3) create the linking phrase.
Strategic: What's your strategy for identifying the most critical relationships first? How do you prioritize connection-building?	Looking at your challenge with {observation}, what strategy would be most effective here? Have you considered breaking {observation} down into smaller concepts as with {concept A}?
Metacognitive: What was challenging about creating this concept map?	How has your understanding of {concept A} changed since you started working on this map? What specific insights have you gained?

Stage 3: Student Focus Group Evaluation (N=34) A focus group with 34 first-semester management students (20 male, 14 female), none with prior experience in business case analysis, revealed findings along two dimensions. Regarding *usability*, simultaneous novelty of the interface and multi-agent interaction caused information overload, and unfamiliarity with the task domain confounded task difficulty with system complexity. Regarding *scaffolding reception*, expectations shaped by ChatGPT usage limited some students’ engagement with reflective scaffolding, while others noted that the metacognitive agent prompted deeper engagement with their own learning process, particularly in priming self-reflection early in the task. These findings indicate that summative evaluation requires incremental feature introduction and task familiarization before assessing scaffolding effectiveness in the full orchestrated system.

Stage 4: Scaffold Effectiveness Analysis (N=153)

To assess whether the AI agents delivered scaffolding as designed and whether these scaffolds translated into observable learning behavior, we conducted a computational analysis of 153 agent–learner interactions from the metacognition-first condition using two complementary methods.

LLM-as-a-Judge Classification. Following recent advances in automated content analysis [62], we employed a custom LLM-based classifier to evaluate each agent response along two dimensions: (1) *scaffolding act* (elaboration, prompting, or feedback), reflecting whether the agent provided explanatory content, posed reflective questions, or offered corrective input; and (2) *scaffold mechanism* (conceptual, strategic, procedural, or metacognitive), indicating which cognitive process the response targeted. Additionally, elaborations were rated for *quality* (helpful, neutral, or unhelpful) based on whether they provided actionable guidance aligned with the learner’s current task state. The classifier achieved mean confidence of 0.93 across these classifications.

Effectiveness Measure. We operationalized scaffold effectiveness through concept map growth: for each agent interaction, we computed $\Delta\text{concepts}$ as the difference in node count between consecutive rounds. An elaboration was coded as *behaviorally effective* if the learner added at least one concept following the interaction ($\Delta\text{concepts} > 0$). This measure captures whether scaffolds translated into externalized knowledge construction, the creation of new nodes and edges.

Analysis Procedure. We fitted mixed-effects regression models with random intercepts by session to account for repeated measures within learners. The dependent variable was $\Delta\text{concepts}$ per round; predictors included agent type (delivery position), helpful elaboration count, total interactions, and baseline concept map size, which represents the first round, where students did not receive AI agent support. A Bootstrap resampling (1,000 iterations) at the session level provided robust confidence intervals [18].

Results. Table 4 presents scaffold effectiveness by agent type in the early metacognitive group. Delivery position strongly predicted concept addition: the metacognitive agent delivered first and produced the highest growth ($M = 2.66$, $SD = 3.50$), while the conceptual agent, who delivered last, showed minimal effect ($M = 0.43$, $SD = 1.15$). Bootstrap analysis confirmed that only the

metacognitive agent effect was robust (95% CI [1.24, 1.91], excluding zero). The effect size between best- and worst-performing agents was large (Cohen’s $d = 0.96$), indicating that position accounted for the variance in learning behavior.

Table 4. Scaffold effectiveness by agent type (N=153).

Agent Type	Pos.	M	SD	n	95% CI	Sig.
Metacognitive	1st	2.66	3.50	31	[1.24, 1.91]	*
Strategic	2nd	1.40	2.81	23	[-0.02, 1.20]	—
Procedural	3rd	0.84	2.79	48	[-0.45, 1.53]	—
Conceptual	4th	0.43	1.15	51	ref.	—

*CI excludes zero. Reference: conceptual. Δ concepts = concepts added per round.

These findings reveal a primacy effect: metacognitive scaffolds delivered early produced greatest immediate behavioral change. This supports DP1 by demonstrating that orchestration—specifically, delivery coordination is a consequential design parameter. The pattern suggests that early scaffolds capture learner attention and cognitive resources, while later scaffolds may encounter diminishing reception. Practitioners designing multi-agent systems should therefore consider not only *which* scaffold to deploy but *when* in the sequence to deploy it.

4 Discussion

4.1 Findings and Contributions

This study contributes design knowledge for orchestrated scaffolding AI agents in higher education, as well as an accessible prototype [27]. The contribution is threefold: (1) design principles that specify how a lead orchestrator coordinates mechanism-specific sub-agents, with evidence that orchestration affects scaffold effectiveness (DP1); (2) a persistent learner profiling approach that enables cross-session scaffolding fading based on assessed ZPD (DP2); and (3) an institutional knowledge integration architecture that grounds AI-generated scaffolds in course materials (DP3). These principles, grounded in scaffolding theory and validated through learner interviews (N=32), IS expert evaluation (N=22), educator refinement (N=6), and a student focus group (N=34), are formulated following Gregor et al. [17] and reusable design knowledge for multi-agent systems in education. Our scaffold effectiveness analysis revealed that delivery position predicted learning behavior: early metacognitive scaffolds produced the largest concept map growth (Cohen’s $d = 0.96$), suggesting a primacy effect that orchestration policies must account for.

A Framework Linking Scaffolding and AI Agent Orchestration.

Prior educational AI has treated scaffolding as monolithic: Systems either scaffold or they do not, utilizing a single support mechanism [3, 23, 47]. Through novel possibilities introduced by orchestrated AI agents, our framework decomposes this known issue into coordinated, distinct mechanisms and maps each to specialized agents, enabling researchers to study *which* scaffold affects *which* outcome [16]. Our empirical analysis demonstrates that mechanism-specific support alone is insufficient: the same scaffold type (metacognitive) produced high

effectiveness when delivered first. This finding elevates orchestration from implementation detail to first-order design concern. While prior work explores AI-assisted education [32,35,43], these single-agent systems cannot address position-dependent effectiveness. DP1 responds by establishing coordination logic, how agents sequence interventions, as consequential as the scaffolds themselves, addressing the reusability gap [11].

Validated Design Knowledge for mechanism-specific Support. The design requirements (DR1–DR4) translate scaffold types [3,24] into testable AI agent orchestration following de Weck [52]. The design principles (DP1–DP3) follow Gregor et al.’s anatomy [17], supporting cumulative knowledge building.

DP1: Orchestrated Multi-Agent Coordination. While Lange and Schlimbach [30] differentiate cognitive levels, they do not address cross-level coordination. Our primacy finding concerning early scaffolds as capturing learner receptivity regardless of type, demonstrates why orchestration matters: practitioners must sequence scaffolds strategically, not just select appropriate types. This extends Belland et al.’s [3] observation that scaffolding requires mechanism-specific adaptation by showing that *timing* of adaptation is equally critical. Designers should: (1) implement scaffold-type classifiers; (2) design position-aware sequencing policies; (3) log orchestration decisions for pattern analysis.

DP2: Adaptive Learner Profiling. Wambsganss and Schmitt [51] enhance learning through post-hoc trace analysis. Our contribution enables real-time profiling: persistent learner profiles allow orchestrated AI agents to calibrate intensity and delivery order based on assessed understanding. While Neis et al. [35] adapt their system within sessions, it resets across conversations. DP2 maintains cross-session continuity, enabling the fading function central to scaffolding [56]. Designers could: (1) implement lightweight diagnostics; (2) store profiles with privacy safeguards; (3) design fading policies informed by position effects.

DP3: Continuous Institutional Knowledge Integration. Steinherr et al. [44] address infrastructure for human-designed materials, not AI agents, other AI agents operated independently of institutional content [35,43]; orchestration now enables dedicated resource-retrieval agents that validate curriculum alignment before scaffolds reach learners. The orchestrated AI agents coordinate when and what to retrieve, and how to integrate it into ongoing scaffolding sequences, leading to the following suggestions for Designers: (1) implementation of retrieval agents within the orchestrated architecture; (2) designing validation protocols for curriculum alignment; (3) providing educator dashboards for oversight.

4.2 Limitations and Future Work

Our evaluation remains formative-artificial [49]: experts and students assessed validity and utility, while LLM-as-a-judge analysis validated scaffold delivery, but summative outcome measurement is pending. The focus group revealed that simultaneous novelty (interface, task, agents) obscured effects, requiring feature-by-feature evaluation. The primacy effect we observed warrants replication across sequences, as our analysis examined early metacognitive scaffolds. The interview sample (N=32) from one European institution limits generalizability. Institutional integration (DP3) was demonstrated in controlled conditions; production

embedding remains unaddressed. The next cycle will compare feature-ablated versions (planned N=200) and test adaptive scaffolding orchestration policies to isolate position from mechanism effects.

Subsequent work addresses cross-domain generalization and longitudinal deployment to assess scaffolding persistence. By establishing AI agent orchestration as the architectural foundation for mechanism-specific scaffolding (DP1 and DP2), this study demonstrates how advances in multi-agent coordination can now address a long-standing challenge in AI-assisted educational systems: delivering the right support, at the right time, adapted to the learner’s evolving needs.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Allmendinger, S., Bonenberger, L., Endres, K., et al.: Multi-agent ai. *Electronic Markets* **36** (2026). <https://doi.org/10.1007/s12525-025-00862-z>
2. Azevedo, R., Hadwin, A.F.: Scaffolding self-regulated learning and metacognition: Implications for the design of computer-based scaffolds. *Instructional Science* **33**, 367–379 (2005). <https://doi.org/10.1007/s11251-005-1272-9>
3. Belland, B.R., Lee, E., Zhang, A.Y., Kim, C.: Characterizing the most effective scaffolding approaches in engineering and technology education: A clustering approach. *Computer Applications in Engineering Education* **30**(6), 1795–1812 (2022). <https://doi.org/10.1002/cae.22556>, <https://onlinelibrary.wiley.com/doi/10.1002/cae.22556>, accessed April 30, 2025
4. Berente, N., Gu, B., Recker, J., Santhanam, R.: Managing artificial intelligence. *MIS quarterly* **45**(3) (2021)
5. Borchers, C., Fleischer, H., Schanze, S., Scheiter, K., Alevi, V.: High scaffolding of an unfamiliar strategy improves conceptual learning but reduces enjoyment compared to low scaffolding and strategy freedom. *Computers & Education* **236** (2025). <https://doi.org/10.1016/j.compedu.2025.105364>
6. Charisi, V., Malinverni, L., Rubegni, E., Schaper, M.M.: Empowering children’s critical reflections on ai, robotics and other intelligent technologies. *NordiCHI ’20*, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3419249.3420090>, <https://doi.org/10.1145/3419249.3420090>
7. Constantinescu, M., Kaptein, M.: Responsibility gaps, llms & organisations: Many agents, many levels, and many interactions. *Science and Engineering Ethics* **31** (2025). <https://doi.org/10.1007/s11948-025-00560-1>
8. Dai, L., Jiang, Y.H., Chen, Y., Guo, Z., Liu, T.Y., Shao, X.: Agent4edu: Advancing AI for education with agentic workflows. In: *Proceedings of the 3rd International Conference on Artificial Intelligence and Education (ICAIE 2024)*. p. 6. ACM, Xiamen, China (nov 2024). <https://doi.org/10.1145/3722237.3722268>
9. D’Mello, S.K., Graesser, A.: Intelligent tutoring systems: How computers achieve learning gains that rival human tutors. In: Schutz, P.A., Muis, K.R. (eds.) *Handbook of Educational Psychology*, pp. 603–629. Routledge, 4 edn. (2024)
10. Duffy, M.C., Azevedo, R.: Motivation matters: Interactions between achievement goals and agent scaffolding for self-regulated learning within an intelligent tutoring system. *Computers in Human Be-*

- havior **52**, 338–348 (2015). <https://doi.org/10.1016/j.chb.2015.05.041>, <https://doi.org/10.1016/j.chb.2015.05.041>
11. Elshan, E., Engel, C., Ebel, P., Siemon, D.: Assessing the reusability of design principles in the realm of conversational agents. In: *The Transdisciplinary Reach of Design Science Research: 17th International Conference on Design Science Research in Information Systems and Technology, DESRIST 2022*, St Petersburg, FL, USA, June 1–3, 2022, Proceedings. p. 128–141. Springer-Verlag, Berlin, Heidelberg (2022). https://doi.org/10.1007/978-3-031-06516-3_10, https://doi.org/10.1007/978-3-031-06516-3_10
 12. Füller, J., Tekic, Z., Hutter, K.: Rethinking innovation management—how ai is changing the way we innovate. *The Journal of Applied Behavioral Science* **60**(4), 603–612 (2024). <https://doi.org/10.1177/00218863241287323>, <https://doi.org/10.1177/00218863241287323>, accessed April 30, 2025
 13. Gioia, D.: A systematic methodology for doing qualitative research. *The Journal of Applied Behavioral Science* **57**(1), 20–29 (2021). <https://doi.org/10.1177/0021886320982715>
 14. Goddard, Q., Moton, N., Hudson, J., He, H.A.: A chatbot won’t judge me: An exploratory study of self-disclosing chatbots in introductory computer science classes. In: *Proceedings of the 26th Western Canadian Conference on Computing Education (WCCCE ’24)*. vol. 9, pp. 1–7. Association for Computing Machinery (2024). <https://doi.org/10.1145/3660650.3660662>
 15. Gonzalez, C., Donahue, K., Goldstein, D.G., Heidari, H., Jalali, M.S., Schelle, B., Singh, A., Woolley, A.W.: Toward a science of human-ai teaming for decision making: A complementarity framework. *PNAS Nexus* **5**(3), pgag030 (02 2026). <https://doi.org/10.1093/pnasnexus/pgag030>, <https://doi.org/10.1093/pnasnexus/pgag030>
 16. Gregor, S., Hevner, A.R.: Positioning and presenting design science research for maximum impact. *Management Information Systems Quarterly* **37**(2), 337–355 (06 2013). <https://doi.org/10.25300/MISQ/2013/37.2.01>
 17. Gregor, S., Kruse, L., Seidel, S.: The anatomy of a design principle. *Journal of the Association for Information Systems* **21**(6), 1622–1652 (2020). <https://doi.org/10.17705/1jais.00649>, <https://aisel.aisnet.org/jais/vol21/iss6/2/>
 18. Henderson, A.R.: The bootstrap: A technique for data-driven statistics. using computer-intensive analyses to explore experimental data. *Clinica Chimica Acta* **359**(1), 1–26 (2005). <https://doi.org/https://doi.org/10.1016/j.cccn.2005.04.002>, <https://www.sciencedirect.com/science/article/pii/S0009898105002433>
 19. Hevner, A.R.: A three cycle view of design science research. *Scandinavian Journal of Information Systems* **19**(2), 87–92 (2007), <http://aisel.aisnet.org/sjis/vol19/iss2/4>
 20. Holldack, F., Banh, L., Strobel, G.: Agentic information systems. *Electronic Markets* **36** (2026). <https://doi.org/10.1007/s12525-025-00861-0>
 21. Hunke, F., Seebacher, S., Thomsen, H.: Please tell me what to do – towards a guided orchestration of key activities in data-rich service systems. In: *Designing for Digital Transformation. Co-Creating Services with Citizens and Industry: 15th International Conference on Design Science Research in Information Systems and Technology, DESRIST 2020*, Kristiansand, Norway, December 2–4, 2020, Proceedings. p. 426–437. Springer-Verlag, Berlin, Heidelberg (2020). https://doi.org/10.1007/978-3-030-64823-7_41, https://doi.org/10.1007/978-3-030-64823-7_41
 22. Janson, A., Söllner, M., Leimeister, J.M.: Ladders for learning: Is scaffolding the key to teaching problem-solving in technology-mediated learning contexts? *Academy of Management Learning & Education* **19**(4), 439–468 (2020). <https://doi.org/10.5465/amle.2018.0078>

23. Kim, M.C., Hannafin, M.J.: Scaffolding problem solving in technology-enhanced learning environments (teles): Bridging research and theory with practice. *Computers & Education* **56**(2) (2011). <https://doi.org/10.1016/j.compedu.2010.08.024>
24. Kim, N.J., Belland, B.R., Walker, E.A.: Effectiveness of computer-based scaffolding in the context of problem-based learning for stem education: Bayesian meta-analysis. *Educational Psychology Review* **30**(3) (2018). <https://doi.org/10.1007/s10648-017-9419-1>
25. Koedinger, K.R., Aleven, V.: Exploring the assistance dilemma in experiments with cognitive tutors. *Educational Psychology Review* **19**, 239–264 (2007). <https://doi.org/10.1007/s10648-007-9049-0>
26. Kosmyna, N., Hauptmann, E., Yuan, Y.T., Situ, J., Liao, X., Beresnitzky, A.V., Maes, P.: Your brain on chatgpt: Accumulation of cognitive debt when using an ai assistant for essay writing task (2025)
27. Kozachek, D., Göldi, A., Schwarz, K., Kiafard, A.: Multi-Agent Scaffolding System for Higher Education Research (Aug 2026), https://github.com/koizachek/scaffolding_multiagentsystem
28. Kujala, S.: Effective user involvement in product development by improving the analysis of user needs. *Behaviour & Information Technology* **27**(6), 457–473 (2008). <https://doi.org/10.1080/01449290601111051>
29. Kumar, N., Wei, X., Zhang, H.: Agentic artificial intelligence as a new frontier in information systems: Promise, peril, and research opportunities. *Information Management* **63**(3), 104317 (2026). <https://doi.org/https://doi.org/10.1016/j.im.2026.104317>, <https://www.sciencedirect.com/science/article/pii/S0378720626000170>
30. Lange, T.C., Schlimbach, R.: Designing knowledge for conversational ai applications: A bloom’s taxonomy perspective. In: *Local Solutions for Global Challenges. Proceedings of the 20th International Conference on Design Science Research in Information Systems and Technology (DESRIST 2025)*. Lecture Notes in Computer Science, vol. 15703. Springer, Cham (2025). https://doi.org/10.1007/978-3-031-93976-1_14
31. Luchini, K., Quintana, C., Krajcik, J., Farah, C., Nandihalli, N., Reese, K., Wiczorek, A., Soloway, E.: Scaffolding in the small: designing educational supports for concept mapping on handheld computers. In: *CHI ’02 Extended Abstracts on Human Factors in Computing Systems*. p. 792–793. CHI EA ’02, Association for Computing Machinery, New York, NY, USA (2002). <https://doi.org/10.1145/506443.506600>, <https://doi.org/10.1145/506443.506600>
32. Maedche, A., Legner, C., Benlian, A., Berger, B., Gimpel, H., Hess, T., Hinz, O., Morana, S., Söllner, M.: Ai-based digital assistants. *Business & Information Systems Engineering* **61**, 535–544 (2019). <https://doi.org/10.1007/s12599-019-00600-8>
33. Masters, C., Vellanki, A., Shangguan, J., Kultys, B., Gilmore, J., Moore, A., Albrecht, S.: Orchestrating human-ai teams: The manager agent as aunifying research challenge. In: *Proceedings of the 2025 7th International Conference on Distributed Artificial Intelligence*. p. 91–107. DAI ’25, Association for Computing Machinery, New York, NY, USA (2025). <https://doi.org/10.1145/3772429.3772439>, <https://doi.org/10.1145/3772429.3772439>
34. MongoDB, Inc.: *Mongodb community server (version 8.0)*. <https://www.mongodb.com/docs/> (2025)
35. Neis, N., Spleth, P., Kudlek, C., Zarnekow, R., Winkelmann, A.: Designing a large language model based conversational agent for language acquisition. In: *Local Solutions for Global Challenges. Proceedings of the 20th International Conference on Design Science Research in Information Systems and Technology (DESRIST 2025)*. Lecture Notes in Computer Science, vol. 15703. Springer, Cham (2025). https://doi.org/10.1007/978-3-031-93976-1_15, https://doi.org/10.1007/978-3-031-93976-1_15

36. NGoon, T.J.: Inventive scaffolds catalyze creative learning. In: Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA '19). pp. 1–5. Association for Computing Machinery (2019). <https://doi.org/10.1145/3290607.3299073>
37. Novak, J.D., Cañas, A.J.: The theory underlying concept maps and how to construct and use them. Tech. Rep. IHMC CmapTools 2006-01 Rev 01-2008, Florida Institute for Human and Machine Cognition, Pensacola, FL (2008)
38. Otasek, D., Morris, J.H., Bouças, J., Pico, A.R., Demchak, B.: Cytoscape automation: empowering workflow-based network analysis. *Genome Biology* **20**(1), 185 (2019). <https://doi.org/10.1186/s13059-019-1758-4>, doi.org/10.1186/s13059-019-1758-4
39. Peffers, K., Tuunanen, T., Rothenberger, M., Chatterjee, S.: A design science research methodology for information systems research **24**(3) (2007). <https://doi.org/10.2753/MIS0742-1222240302>
40. Petrovska, O., Clift, L., Moller, F., Pearsall, R.: Incorporating generative ai into software development education. In: Waite, J., Crosby, R. (eds.) CEP '24: Proceedings of the 8th Conference on Computing Education Practice. pp. 37–40. Association for Computing Machinery, Durham, United Kingdom (2024). <https://doi.org/10.1145/3633053.3633057>
41. van de Pol, J., Volman, M., Beishuizen, J.: Scaffolding in teacher–student interaction: A decade of research. *Educational Psychology Review* **22**, 271–296 (2010). <https://doi.org/10.1007/s10648-010-9127-6>
42. Shavit, Y., O’Keefe, C., Eloundou, T., McMillan, P., Agarwal, S., Brundage, M., Robinson, D.G.: Practices for governing agentic ai systems. Technical report, OpenAI (2023), <https://cdn.openai.com/papers/practices-for-governing-agentic-ai-systems.pdf>
43. Sjöström, J., Dahlin, M.: Tutorbot: A chatbot for higher education practice. In: Designing for Digital Transformation: Co-Creating Services with Citizens and Industry: 15th International Conference on Design Science Research in Information Systems and Technology, DESRIST 2020, Kristiansand, Norway, December 2–4, 2020, Proceedings. Lecture Notes in Computer Science, vol. 12388, pp. 93–98. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64823-7_10, https://link.springer.com/chapter/10.1007/978-3-030-64823-7_10
44. Steinherr, V.M., Brehmer, M., Stöckl, R., Reinelt, R.: Design science research as a guide for innovative higher education teaching: Towards an application-oriented extension of the proficiency model. In: Design Science Research for a Resilient Future: 19th International Conference on Design Science Research in Information Systems and Technology, DESRIST 2024, Trollhättan, Sweden, June 3–5, 2024, Proceedings. Lecture Notes in Computer Science, vol. 14621, pp. 213–228. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-61175-9_15, https://link.springer.com/chapter/10.1007/978-3-031-61175-9_15
45. Streamlit, Inc.: Streamlit. <https://github.com/streamlit/streamlit> (2025), gitHub-Repository; Commit COMMIT; 2025-09-11
46. Tankelevitch, L., Glassman, E.L., He, J., Kazemitabaar, M., Kittur, A., Lee, M., Palani, S., Sarkar, A., Ramos, G., Rogers, Y., Subramonyam, H.: Tools for thought: Research and design for understanding, protecting, and augmenting human cognition with generative ai. In: Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems. CHI EA '25, Association for Computing Machinery, New York, NY, USA (2025). <https://doi.org/10.1145/3706599.3706745>
47. Thomann, H., Deutscher, V.: Scaffolding through prompts in digital learning: A systematic review and meta-analysis of effectiveness on learning achievement. *Educational Research Review* **47**, 100686 (2025). <https://doi.org/10.1016/j.edurev.2025.100686>

48. Van Rossum, G., Drake Jr, F.L.: Python tutorial. Centrum voor Wiskunde en Informatica Amsterdam, The Netherlands (1995)
49. Venable, J., Pries-Heje, J., Baskerville, R.: Feds: a framework for evaluation in design science research. *European journal of information systems* **25**(1), 77–89 (2016)
50. Vygotsky, L.S.: *Mind in Society: Development of Higher Psychological Processes*. Harvard University Press (1978), <http://www.jstor.org/stable/j.ctvjf9vz4>
51. Wambsganss, T., Schmitt, A.: Enhancing personalized learning through process mining. *Business & Information Systems Engineering* (2024). <https://doi.org/10.1007/s12599-024-00901-7>
52. de Weck, O.L.: Requirements definition. <https://ocw.mit.edu/courses/16-842-fundamentals-of-systems-engineering-fall-2015> (2015), lecture notes, Session 2, 16.842 Fundamentals of Systems Engineering, MIT OpenCourseWare
53. Whitelock-Wainwright, A., Laan, N., Wen, D., Gašević, D.: Exploring student information problem solving behaviour using fine-grained concept map and search tool data. *Computers & Education* **145**, 103731 (2020). <https://doi.org/https://doi.org/10.1016/j.compedu.2019.103731>
54. Winkler, R., Hobert, S., Fischer, T., Salovaara, A., Soellner, M., Leimeister, J.M.: Engaging learners in online video lectures with dynamically scaffolding conversational agents. In: *Proceedings of the 28th European Conference on Information Systems (ECIS)*. Association for Information Systems, Virtual Conference (2020), https://aisel.aisnet.org/ecis2020_p/97_june15-17_2020
55. Winkler, R., Roos, J.: Bringing ai into the classroom: Designing smart personal assistants as learning tutors. In: *ICIS 2019 Proceedings* (2019), https://aisel.aisnet.org/icis2019/learning_environ/learning_environ/10/
56. Wood, D., Bruner, J.S., Ross, G.: The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry* **17**(2), 89–100 (1976). <https://doi.org/10.1111/j.1469-7610.1976.tb00381.x>
57. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K.R., Cao, Y.: React: Synergizing reasoning and acting in language models. In: *Proceedings of the Eleventh International Conference on Learning Representations (ICLR 2023)* (2023)
58. Zao-Sanders, M.: How people are really using gen ai in 2025. *Harvard Business Review* (April 2025), <https://hbr.org/2025/04/how-people-are-really-using-gen-ai-in-2025>
59. Zhai, C., Wibowo, S., Li, L.D.: The effects of over-reliance on ai dialogue systems on students’ cognitive abilities: a systematic review. *Smart Learning Environments* **11**(28) (2024). <https://doi.org/10.1186/s40561-024-00316-7>
60. Zhai, C., Wibowo, S.: A systematic review on artificial intelligence dialogue systems for enhancing english as foreign language students’ interactional competence in the university. *Computers and Education: Artificial Intelligence* **4**, 100134 (2023). <https://doi.org/https://doi.org/10.1016/j.caeai.2023.100134>, <https://www.sciencedirect.com/science/article/pii/S2666920X23000139>
61. Zhang, Y., Cai, Y., Zuo, X., Luan, X., Wang, K., Hou, Z., Dong, J.S.: Position: Trustworthy AI agents require the integration of large language models and formal methods. In: *Proceedings of the 42nd International Conference on Machine Learning*. pp. 1–19. Vancouver, Canada (2025)
62. Zhiheng, X., Wenxiang, C., Xin, G., Wei, H., Yiwen, D., Boyang, H., Ming, Z., Junzhe, W., Senjie, J., Enyu, Z., Rui, Z., Xiaoran, F., Xiao, W., Limao, X., Yuhao, Z., Weiran, W., Changhao, J., Yicheng, Z., Xiangyang, L., Zhangyue, Y., Shihan, D., Rongxiang, W., Wensen, C., Qi, Z., Wenjuan, Q., et al.: The rise and potential of large language model based agents: A survey (2023)