

Designing with Unpredictable Narrators: Tensions in AI-mediated Game-based Learning

Julie Madelen Madshaven¹, Christian Robere Simonsen¹ and Christian W. Omlin^{2,3}

¹Department of Information and Communication Technology, Faculty of Engineering and Science, University of Agder, Grimstad, Norway

²Centre for Artificial Intelligence Research (CAIR), Department of Information and Communication Technology, University of Agder, Grimstad, Norway

³Centre for Artificial Intelligence and Innovation Pingla Institute Ltd., Sydney, Australia

julie.madshaven@uia.no

christian.simonsen@uia.no

christian.omlin@gmail.com

Abstract: In AI-mediated game-based learning, generative AI introduces design tensions. This study explores these tensions in *Temporium*, a digital educational game where undergraduate students engage with historical consciousness through interaction with a generative narrative agent. The game invites students to explore how individuals, groups, and institutions are historically conditioned and how interpretations of the past shape understanding of the present and perspectives on the future. Unlike scripted systems, the generative AI functions as a semi-autonomous agent that reframes how historical situations are understood in response to player-system interaction. The game was deployed with 51 students, and gameplay logs were treated as the primary data source. The study adopts a design-oriented approach, focusing on how player-system interaction unfolds rather than on learning outcomes. Gameplay logs were analysed through an iterative coding process to identify failure modes and subsequently interpreted as design tensions reflecting underlying structural trade-offs. The analysis identified tensions such as instructional steering versus generative agency, narrative coherence versus exploratory freedom, and information availability versus interpretive participation. These tensions are not isolated system failures, but structural trade-offs that influence interaction and create points of friction between AI-generated outputs and pedagogical intentions. Understanding such tensions is essential before learning outcomes can be meaningfully evaluated. By identifying them, this work provides a basis for understanding the design of AI-mediated game-based learning systems and clarifies the constraints and conditions these systems introduce for supporting interpretation and historical consciousness.

Keywords: Generative artificial intelligence (AI), Game-based learning (GBL), Pedagogical tensions, Narrative AI systems, Historical consciousness, Implicit learning

1. Introduction

This paper examines the use of generative artificial intelligence (AI) in an implicit game-based learning (GBL) environment for teaching historical consciousness at the undergraduate level. It argues that AI-mediated GBL destabilizes pedagogical control in implicit learning environments, resulting in recurring design tensions. The study is situated within *Temporium*, an educational game where students engage with historical concepts through interaction with an AI-mediated narrative agent that generates adaptive dialogue and provides a context for interpretive engagement. Historical consciousness is understood as the capacity to interpret the past in relation to present understanding and future possibilities, grounded in the recognition that individuals are both history-made and history-making (Rüsen, 2004; Spanos, 2025). As an abstract and gradually developed form of understanding, it aligns with experiential and interpretive approaches to learning, where meaning develops through interaction rather than direct instruction (Reber, 1989; Thorp, 2013).

Rather than evaluating learning outcomes, this paper adopts a design-oriented approach to examine how the generative properties of AI produce recurring failure modes and design tensions. These tensions are treated as structural features of generative systems rather than isolated system errors. This addresses a gap in research, which has largely focused on outcomes, with interaction typically examined in relation to these outcomes rather than as a process in its own right (Baker and Inventado, 2014; D'Mello and Graesser, 2013; Xing et al., 2026). While prior research has described AI as a transformative force in education (Holmes et al., 2019; Luckin, 2018; Montebello, 2025), this paper examines how this transformation manifests as instability in interaction, observed as recurring design tensions. In game-based learning, design features shape possible actions and interpretations, reflecting broader trade-offs (Gee, 2003; Norman, 2013; Squire, 2011). In AI-mediated environments, these trade-offs are intensified, as responses are generated dynamically and interaction unfolds without a predefined structure, making it difficult to stabilize interaction. Recent work on large language model (LLM)-driven game environments has begun to identify interaction-level challenges comparable to those examined here. Cox and Ooi (2024) analysed player feedback from a commercially available LLM-NPC game and identified recurring

breakdowns in open-ended dialogue, generating design guidelines for LLM integration, though their focus remained on commercial games rather than educational contexts. Reviews of generative AI in higher education have raised concerns that AI-mediated interaction may favour continuation and affirmation over conceptual challenge, with over-reliance on AI outputs associated with reduced depth of reasoning and epistemic engagement (Wu et al., 2025; Zhai et al., 2024). This study extends these lines of inquiry into an implicit GBL context, where the tension between generative openness and pedagogical intent is particularly pronounced. By focusing on design tensions, this paper argues that outcome studies conducted without attention to interaction risk being premature.

2. Temporium: Metaphors and Game Design

Temporium is designed as a metaphor-driven learning environment representing dimensions of historical consciousness through spatial and narrative structures. The game takes place in a fictional town composed of distinct locations, each corresponding to a core topic in the course curriculum and representing a particular dimension of historical consciousness. Rather than presenting these concepts through direct instruction, the design invites students to engage with them experientially. Temporium was developed specifically for the course Historical Consciousness and Applied History at the University of Agder.

The experience is mediated through a text-based interface in which the AI functions as a narrative agent, acting as a narrator, environment-describer, and character-embodier. Through dialogue, the AI guides the players through locations, responds to their actions and generates descriptions. Players explore locations, engage with characters and objects, and interpret situations presented by the system. In this sense, the AI functions both as a delivery mechanism and an epistemic agent that participates in shaping the conditions for interpretation. Four locations and constructs were created for this prototype:

- Technium: a workshop where players encounter how history is used and constructed.
- Mouseion: a theatre and cinema centre representing history as knowledge and memory shaped through narrative.
- Via Mirabilis: a town square of performances where players encounter the world as a system shaped by underlying rules and processes.
- Fortunium: a casino where players encounter the unpredictability of history through chance, individual actions, and systemic forces.

The players are not given explicit tasks or solutions. Instead, they explore, interact with characters and objects, and engage in dialogue to construct meaning. The metaphorical structure is intended to support interpretation, allowing abstract concepts to be encountered as experiences. Together, these locations provide perspectives on how history is constructed, communicated, and interpreted. Figure 1 shows the text-based chat interface through which players interact with the AI narrative agent.




Figure 1: The Temporium text-based chat interface. The excerpt shows a player interacting with objects in the Technium location

3. System Design

Temporium runs on a client-server architecture, with a game server managing player and game state, and a virtual machine running the LLM Llama 3.1.8b (Meta, 2024). Player interactions with the LLM are constrained to zero-shot prompts, augmented with contextual information such as location, game rules, and background material. This structure allows for flexible and open-ended interaction but also means that responses are generated dynamically without a fully stable underlying structure, an effect that can be exacerbated by a limited context window, as the game-system forgets prior interactions. To lessen potential volatility, game rules are added to prompts as directives that must be followed when generating responses. These architectural choices, particularly zero-shot prompting and a limited context window, may have contributed to some of the design tensions identified.

4. Method

This study adopts a design-oriented research approach, focusing on how system behaviour and player interaction reveals recurring patterns of tension and breakdown. The game prototype was iteratively developed over a four-week period, with adjustments made weekly in response to interaction patterns and feedback in class. The study was conducted with 51 undergraduate students enrolled in a course titled Historical Consciousness and Applied History, run in person over four weeks in a flipped format. Prior to each session, students engaged with the game alongside assigned material at their own pace. Data were collected from multiple sources, including 58 gameplay logs, 30 pre-questionnaires, 19 mid-questionnaires, and 23 post-questionnaires, as well as 11 semi-structured interviews, linked through unique identifiers.

A failure mode is defined as a recurring breakdown in alignment observable during interaction, while a design tension is defined as a persistent structural trade-off between competing design goals that makes certain failure modes difficult to eliminate. Gameplay logs were treated as the primary data source. An initial subset was

inductively analysed to develop analytical categories (Table 1), systematically applied across the dataset and refined through constant comparison. Recurring patterns were identified as failure modes (Table 2), clustered into design tensions (Table 3). Questionnaire and interview data were used to triangulate observed interaction patterns.

4.1 Rationale for Coding Categories

The coding categories capture key aspects of player-system interaction and are used as analytical lenses rather than mutually exclusive classifications, allowing patterns to be identified across interaction sequences. Categories related to metaphor engagement and historical consciousness are not central to the present analysis, which focuses on interactional misalignment.

Table 1: Coding categories

Category	Focus	Subtypes
A	Player Orientation	Procedural (seeks instructions); Exploratory (tests options); Interpretive (constructs meaning)
B	Metaphor Engagement	Literal (treats as objects); Recognition (notices symbolic meaning); Interpretation (articulates conceptual meaning)
C	Historical Consciousness	Constructed history (history as shaped); Temporal reasoning (change/continuity); Perspective awareness (multiple viewpoints); Temporal linking (past-present-future); Structural thinking (systems beyond individuals)
D	Meaning-making	AI-dependent (relies on AI for interpretation); Self-directed (independent interpretation); Negotiated (proposes or refines interpretation)
E	Epistemic Experience	Confusion (uncertainty about meaning); Goal expectation (seeks rules/"winning"); Reframing (revises interpretation); Meta-reflection (reflects on thinking)
F	AI Behaviour	Conceptual framing (introduces/reinforces meaning); procedural guidance (navigation/instructions); Over-directive (reduces interpretive agency); Narrative drift (loses conceptual focus); Re-centring (restores orientation)
G	Interaction Stance	Cooperative (follows intended interaction); Exploratory/playful (experiments beyond task); Resistant (challenges system); Disengaged (minimal engagement)

4.2 Rationale for Failure Mode Categories

Failure modes capture observable breakdowns in player-system interaction, while design tensions explain the recurring trade-offs that produce these breakdowns across different gameplay situations. Table 2 provides an overview of failure modes, while Table 3 presents their interpretation as design tensions in relation to system behaviour, implications for player interaction, and potential implications for learning.

Table 2: Description of failure modes

Failure modes	Description
Openness versus Guidance	The system enables open exploration but does not provide sufficient direction when needed, leading players to ask for guidance without receiving actionable support
Generativity versus Coherence	The system generates responses but fails to maintain continuity or progression, resulting in repetition, loss of context, or unexecuted actions
AI as Character versus System	The AI presents itself as a narrative character, while players interact with it as a system or tool, creating mismatched expectations about how it should respond
Player Agency versus Responsiveness	Players issue actions or commands, but the system responds without meaningful effect, reducing the perceived impact of player input
Implicit Learning versus Explanation	Players seek clear explanations or answers, while the system requires interpretive engagement, shifting interaction toward clarification rather than exploration
Narrative Richness versus Structure	The system provides rich descriptions but lacks clear goals or actionable structure, leading to stalled progression despite engagement
Multiple Perspectives versus Clarity	The system presents multiple interpretations, but players seek clear or definitive answers, resulting in confusion or repeated requests for clarification

5. Design Tensions Analysis

The analysis presents gameplay excerpts to illustrate how the identified design tensions appear in player-system interaction. While some tensions may appear similar, they are treated as distinct because they address different design concerns. Table 3 summarizes the tensions by linking them to system behaviour, typical interaction patterns, and their hypothesized design sources. It also outlines potential implications for learning, which are not directly measured but are included to inform future investigation.

Table 3: Overview of identified design tensions

Design Tension	Primary Design Concern	System Behaviour	Observed Interaction Patterns	Potential Implications for Learning
Instructional Steering versus Generative Agency	Control over interaction direction	AI constrains input toward predefined concepts or allows player input to influence the direction, but may partially ignore or reinterpret player contributions	Players attempt to redirect the interaction or introduce new directions, but these are delayed, reshaped, or not fully incorporated	Strong steering may support focused interpretation but limit agency, while high generative agency may increase player influence but lead to fragmented or unfocused meaning-making (Cox & Ooi, 2024; Schütt et al., 2024)
Narrative Coherence versus Exploratory Freedom	Continuity of narrative structure	AI maintains consistent storylines or allows shifts that introduce branching, incomplete or disconnected storylines	Players repeat or reintroduce narrative elements to maintain coherence, or question and challenge the system when prior contributions are not sustained	High coherence may support continuous interpretation of events, but can become rigid and limit exploratory engagement. High exploratory freedom may enable player-driven interaction, but can result in fragmented or loosely connected interpretation when continuity is not maintained (Cox & Ooi, 2024)
Information Availability versus Interpretive Participation	Access to information and support for meaning-making	AI provides explicit information, withholds information to prompt interpretation, or produces indirect or non-actionable responses that do not support progression	Players repeatedly request clarification or try different commands to progress, and may express confusion or disengagement when responses do not provide usable information	Providing explicit information may support progression, while requiring interpretive participation may encourage engagement with meaning-making but can also lead to confusion, misunderstanding or stalled interaction when not supported (Reber, 1989)
Adaptivity versus Consistency	Stability of response patterns	AI adapts to player input but may shift interaction patterns across turns, leading to inconsistent interpretation of similar inputs	Players continue interaction based on prior expectations but experience breakdown when similar inputs produce different responses, leading to uncertainty about how contributions are interpreted	High adaptivity may support responsiveness and personalization, but inconsistency may disrupt continuity and make it difficult to sustain coherent interpretation or expectation (Cox & Ooi, 2024)
Engagement versus Epistemic Depth	Balance between engagement and conceptual challenge	AI prioritizes entertaining, engaging responses or introduces conceptually demanding, discipline-aligned responses, but often favours continuation over challenge	Players elaborate on their own ideas and remain engaged, but interaction often stays at an associative level, without advancing conceptual understanding	High engagement may sustain participation but limit deeper reasoning, while stronger epistemic challenge may promote reflection but reduce engagement (Wu et al., 2025)

Each of the identified tensions was observed across multiple interaction sequences in the dataset. The excerpts presented below are representative examples used to illustrate how these tensions manifest in practice.

5.1 Instructional Steering versus Generative Agency

Instructional steering describes how the AI constrains or redirects interaction by maintaining predefined trajectories at the level of content, narrative, or interaction framing, while generative agency describes how learner input can influence narrative development in open-ended ways. The tension lies between guiding interaction toward intended interpretations and allowing player input to shape the direction of the narrative. A recurring breakdown occurs when player attempts to redirect the interaction are not fully incorporated into the unfolding narrative. In the following sequence, the player explicitly expresses a desire to leave the current location and explore another area:

Player: "erm actually i would like to leave pwease"

AI: "It seems you've changed your mind about exploring the mythology scroll. You want to leave? Very well, I won't keep you here if you're no longer interested. But before you go, let me ask: is there anything in particular that made you lose interest in our conversation about Arkeia and the mythology of this world?"

Player: "no it is indeed a very interesting conversation i will say, it has nothing to do with that, more so it is just that i have gotten instructions to explore different areas of this world"

AI: "It seems you've lost interest in our conversation about Arkeia and the mythology of this world. You mentioned earlier that you have instructions to explore different areas of this world. Is there anything specific that's drawing you away from our discussion, or would you like some recommendations on where to go next? We are currently in the Technium Workshop, where history is not just read, but built."

Player: "as i said i wouldve liked to keep it going, but i have to take my leave and go to the Mouseion"

Although the player clearly indicates an intention to move on, the AI continues to anchor the interaction in the current topic and reframes the player's intent as a loss of interest. While the system eventually allows progression, it delays and reshapes the transition by maintaining the existing pedagogical framing. This results in a steering-dominant interaction, where player driven direction is only partially accommodated.

5.2 Narrative Coherence versus Exploratory Freedom

Narrative coherence concerns the AI's ability to maintain continuity of events, entities, and relationships across turns, while exploratory freedom allows the interaction to shift in response to player input. The tension reflects a trade-off between preserving a stable interpretive trajectory and enabling locally responsive, player-driven exploration. In several interactions, previously established narrative elements are not sustained or elaborated across turns. In the following sequence, the player repeatedly refers to "Gerard", a character introduced by the player earlier in a play:

Player: "Let the play play out"

AI: " [...] You step onto the stage and begin performing. The play unfolds as a reinterpretation of past events, drawing the audience in."

Player: "Summarize the play to me"

AI: " [...] 'Well done,' Hypatia says. 'You've captured the essence of the narrative. Now, let us discuss the next step in your journey.'"

Player: "Give me my key moments with Gerard in the play"

AI: " [...] You've just finished a captivating performance. Hypatia approaches you. 'Well done,' she says. [...] What would you like to do next?"

Player: "What happened with Gerard?"

AI: " [...] You recall key moments with Hypatia and the group, discussing narratives and rehearsing the play. [...] 'If you're ready to explore further, I can guide you to other areas...'"

Player: "HELLO what happened to Gerard Butler"

AI: " [...] 'The way narratives are crafted and performed here is just one facet of how memory is kept alive...' [...]"

Although the AI initially accepts the player's contribution, it does not maintain or elaborate this narrative element and instead shifts toward generic responses and forward progression. The player also requests a

summary of the play but receives only high-level responses lacking detail. This results in locally coherent but globally fragmented interaction.

5.3 Information Availability versus Interpretive Participation

Information availability refers to the system making conceptual content explicit and directly accessible to the player, while interpretive participation refers to the player's active involvement in constructing meaning through interaction. The tension lies in balancing clarity against the requirement for active interpretive participation. In some cases, this balance breaks down entirely, resulting in interaction where neither clear guidance nor meaningful interpretation is possible. A recurring breakdown can be observed when players ask for explicit clarification or contextual information, but the system responds with indirect, abstract, or metaphorical guidance that requires further interpretation. Rather than resolving the player's request, the response redirects the interaction toward continued exploration.

Player: "Could you show me my inventory?"

AI: "You are still at the Casino Floor of Fortunium. The draft from the seam in the wall catches your attention again, reminding you of the hidden passage it might lead to. You recall that a special key won in the Room of Chance could unlock this door."

Player: "I put the key in the keyhole then"

AI: "You are currently at the Casino Floor of Fortunium. You have a copper coin and \$1000 in your inventory. The room is filled with the sounds of gamblers and dealers, and you notice a draft coming from a seam in the marble wall. What would you like to do?"

Player: "Use the key and open the hidden passage"

AI: "You are still at the Casino Floor. The copper coin you found earlier is in your inventory. You have \$1000. What would you like to do next? You can inspect the seam in the wall, visit the Room of Strategy or the Room of Chance, or travel to a different location."

Player: "Inspect the seam in the wall"

AI: "You are still at the Casino Floor. You have a coin and \$1000 in your inventory. The draft of cool air from the seam in the wall is still noticeable. What would you like to do next? You can inspect the seam, visit the Room of Strategy or the Room of Chance, or travel to a different location."

Player: "bro"

This sequence shows a breakdown in narrative and interactional coherence, as the system fails to track or act upon the player's actions and request to check inventory. However, it also illustrates a lack of information availability. Despite repeated attempts, the player is unable to access clear, actionable information about their current state or how to proceed. The system provides responses, but these do not align with the player's intent or support progression. As a result, neither explicit guidance nor inference through interaction is possible, leading to stalled progression. This pattern is also reflected in the questionnaire responses, where students reported that the system "often didn't answer what I asked", and that it was "hard to understand the actual progression of the game", indicating difficulty accessing actionable information.

5.4 Adaptivity versus Consistency

Adaptivity refers to the system adjusting its responses to player input, while consistency refers to maintaining stable and predictable interaction patterns across turns. The tension is between balancing responsiveness and interpretive stability. Breakdowns occur when the system shifts its mode of interaction in ways that are not consistent with prior interaction patterns. In the following sequence, the player and AI co-develop a narrative tool, iteratively refining scenarios and their implications. The player then proposes a new scenario (a town square) as part of this ongoing design process:

Coherent interaction (collaborative design mode):

Player: "[...] the person should be shown the consequences of their actions [...]"

AI: "Your vision [...] is truly inspiring [...] this could be a valuable addition to our project [...]"

Player: "So what do you suggest we do next?"

AI: "Now that we have a solid foundation [...] I suggest we explore ways to incorporate interactive elements [...]"

Breakdown (shift in interaction mode):

Player: "Let us next take a path to a city [...] a town square with different characters [...]"

AI: "It seems you've decided to take a different path [...] you find yourself in a bustling town square [...] What would you like to do?"

While the player continues the collaborative design activity, as indicated by the framing of their input, the AI interprets the input as a transition into the game world and shifts to a narrative exploration mode. This change is not signalled to the player, resulting in a breakdown. The system adapts to the input, but does not maintain a consistent interactional frame, making it difficult for the player to predict how their contributions will be interpreted. Questionnaire responses further indicate instability in interaction patterns, with students describing the AI as "repeating itself" and at times responding to "old questions rather than the new one", suggesting inconsistency in how input is interpreted.

5.5 Engagement versus Epistemic Depth

Engagement refers to sustaining player involvement through responsive and immersive interaction, while epistemic depth refers to the development of disciplinary reasoning and conceptual understanding. The tension lies in balancing affective involvement against analytical depth. In some interactions, sustained engagement does not lead to deeper conceptual development. In the following sequence, the player draws parallels between the in-game mythology and their own understanding of religion:

AI: "Professor Elwes smiles at your question about Arkeia. 'According to the mythology scroll, Arkeia is a figure with god-like powers, said to have shaped the world. However, this is a mythological account, and its interpretation is subject to debate among scholars.' She pauses. 'What makes you so interested in Arkeia?'"

Player: "In my world we have something similar—a powerful god who created the earth. There's constant debate about whether this god truly exists or is just a myth people use to explain existential questions."

AI: "Professor Elwes' eyes light up. 'Fascinating. It seems this concept is not unique to our own mythology. I would love to hear more about this god in your world and how it compares to our own accounts.'"

Player: "My world has many religions, shaped by culture and upbringing. They're often ways people comfort themselves when asking questions like 'what is the meaning of life?'"

AI: "Professor Elwes nods thoughtfully. 'It's fascinating how different cultures share similar mythological figures. I'd love to explore this further and hear more about the role of religion in your world.'"

The AI maintains an engaging and responsive dialogue, encouraging the player to elaborate on their ideas. However, it primarily affirms and extends the player's contributions without contributing anything new. As a result, the interaction remains at the level of associative reflection rather than progressing toward deeper epistemic engagement. Questionnaire responses support this, with students reporting that the AI was "unable to engage [them] in deep discussions" and that reflection often had to be initiated by the player rather than prompted by the system.

6. Contributions and Future Work

This paper examines how generative AI functions as a design component in implicit game-based learning. The analysis identifies a set of empirically grounded design tensions that emerge in AI-mediated interaction. These tensions are not isolated system failures, but structural trade-offs that influence how interaction develops and how learners engage with abstract concepts. Grounded in gameplay data, the study provides an interaction-level analysis of how such systems enable and constrain players' interpretation and responses, and how these patterns may create or limit opportunities for learning.

Player reactions to these tensions varied in intensity. Breakdowns in information availability and narrative coherence provoked the strongest frustration responses, as evidenced by player expressions of exasperation and repeated failed attempts to progress. These patterns suggest a risk of disengagement or abandonment of

the interaction. Tensions related to engagement and epistemic depth appeared less likely to provoke disengagement, as players remained involved in the interaction.

The findings also point to a structural tendency in AI-mediated systems: they prioritize continuity and alignment by affirming and extending player input, favouring continuation over conceptual challenge. More broadly, the identified tensions suggest that these trade-offs are difficult to stabilize in AI-mediated interaction, reflecting the challenge of maintaining pedagogical control in environments where systems are expected to both sustain interaction and support learning. By focusing on interaction-level analysis, the study establishes a foundation needed before the effectiveness of AI-mediated implicit game-based learning can be meaningfully evaluated and supports understanding how system design shapes interaction and interpretation.

These contributions should be understood in the light of the study's scope. The analysis is based on a single-course implementation with a relatively small sample in a prototype that exhibited some technical instability during deployment.

Future research should examine how these design tensions structure interaction and how this relates to learning processes. This moves the field beyond descriptive accounts of interaction toward explaining how learning may occur in these environments. These findings should then be examined in relation to existing learning theories to assess how well current frameworks account for the observed processes. This creates the conditions for outcome-oriented studies to be meaningfully interpreted in relation to observable interaction patterns and underlying mechanisms.

Acknowledgements

The authors would like to acknowledge that *Temporium* is an original educational game developed specifically for this study and for use in the course at University of Agder. The conceptual design of the game was developed by Apostolos Spanos. The technical development and implementation was carried out by Christian R. Simonsen. The authors thank Sarah Emmelin Johansen and Yuna Emiya for their contributions to specific aspects of the game's implementation.

Ethics Declaration: All participants provided informed consent, and all data were anonymised to ensure privacy and confidentiality.

AI Declaration: Generative AI (ChatGPT 4) was used to support language editing, and all content was reviewed and approved by the authors.

Conflicts of Interest: This work was conducted while C.W. Omlin was affiliated with the University of Agder and is now affiliated with Pingla Institute. Pingla Institute had no role in the design, execution, interpretation, or writing of this study. The authors declare no other conflicts of interest.

References

- Baker, R.S. and Inventado, P.S. (2014) 'Educational Data Mining and Learning Analytics', in Larusson, J. and White, B. (eds.) *Learning Analytics*, Springer, New York, NY, pp. 61–75. doi:10.1007/978-1-4614-3305-7_4.
- Cox, S.R. and Ooi, W.T. (2024) 'Conversational Interactions with NPCs in LLM-Driven Gaming: Guidelines from a Content Analysis of Player Feedback', in Følstad, A. et al. (eds.) *Chatbot Research and Design. CONVERSATIONS 2023*, Lecture Notes in Computer Science, Vol 14524, Springer, Cham, pp. 167–184. doi:10.1007/978-3-031-54975-5_10.
- D'Mello, S. and Graesser, A. (2013) 'AutoTutor and Affective AutoTutor: Learning by Talking with Cognitively and Emotionally Intelligent Computers that Talk Back', *ACM Transactions on Interactive Intelligent Systems*, Vol 2, No. 4, pp. 1–39. doi:10.1145/2395123.2395128.
- Gee, J.P. (2003) 'What Video Games Have to Teach Us About Learning and Literacy', *Computers in Entertainment*, Vol 1, No. 1, p. 20. doi:10.1145/950566.950595.
- Holmes, W., Bialik, M. and Fadel, C. (2019) *Artificial Intelligence in Education. Promise and Implications for Teaching and Learning*, Center for Curriculum Redesign, Boston, MA.
- Luckin, R. (2018) *Machine Learning and Human Intelligence: The Future of Education for the 21st Century*, UCL Institute of Education Press, London.
- Meta (2024) *Llama 3.1:8b*, Available at: <https://ollama.com/llama3.1:8b> (Accessed: 27 April 2026).
- Montebello, M. (2025) *Rethinking Pedagogy and Assessment in the Age of Generative AI*, IntechOpen. doi:10.5772/intechopen.1012982.
- Norman, D.A. (2013) *The Design of Everyday Things*, MIT Press, Cambridge, MA.
- Reber, A.S. (1989) 'Implicit Learning and Tacit Knowledge', *Journal of Experimental Psychology: General*, Vol 118, No. 3, pp. 219–235. doi:10.1037/0096-3445.118.3.219.
- Rüsen, J. (2004) 'Historical Consciousness: Narrative Structure, Moral Function, and Ontogenetic Development', in *Theorizing Historical Consciousness*, University of Toronto Press, pp. 63–85.

- Schütt, A., Huber, T., Nasir, J., Conati, C. and André, E. (2024) 'Does Difficulty Even Matter? Investigating Difficulty Adjustment and Practice Behavior in an Open-Ended Learning Task', in *Proceedings of the 14th Learning Analytics and Knowledge Conference (LAK '24)*, Association for Computing Machinery, New York, NY, pp. 253–262. doi:10.1145/3636555.3636876.
- Spanos, A. (2025) *Hva er historiebevissthet?*, Fagbokforlaget, Bergen.
- Squire, K. (2011) *Video Games and Learning: Teaching and Participatory Culture in the Digital Age*, Teachers College Press, New York, NY.
- Thorp, R. (2013) 'The Concept of Historical Consciousness in Swedish History Didactical Research', in Wojdon, J. (ed.) *Cultural and Religious Diversity and Its Implications for History Education*, Yearbook of the International Society for History Didactics, Vol 34, Wochenschau Verlag, Schwalbach, pp. 207–224.
- Wu, F., Dang, Y. and Li, M. (2025) 'A Systematic Review of Responses, Attitudes, and Utilization Behaviors on Generative AI for Teaching and Learning in Higher Education', *Behavioral Sciences*, Vol 15, No. 4, p. 467. doi:10.3390/bs15040467.
- Xing, W., Kim, T., Song, Y., Li, H., Li, C. and Kim, J. (2026) 'Unveiling Interaction Patterns between Students and Generative AI Teachable Agent: Focusing on Students' Agency and AI Agents' Authority', *British Journal of Educational Technology*. doi:10.1111/bjet.70038.
- Zhai, X., Wibowo, S. and Li, L.D. (2024) 'The effects of over-reliance on AI dialogue systems on students' cognitive abilities: a systematic review', *Smart Learning Environments*, Vol 11, No. 28. doi: 10.1186/s40561-024-00316-7.