Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Toward a Technological Theory of Learning Experience Design: New Learning Science Colliding with Artificial Intelligence

Andrew S. Gibbons

andygibbonsster@gmail.com

Brigham Young University

Research Article

Keywords: learning experience design, technological theory, artificial intelligence, prospection, instructional architecture, design layers

Posted Date: November 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11310255/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Learning experience design (LXD) has reached a crisis point. Two powerful revolutions are colliding: artificial intelligence is rapidly transforming what learning technologies can do, while advances in neuroscience and cognitive science are reshaping what we understand learning itself to be. Generative and conversational systems can now respond, adapt, explain, simulate, and interact in ways that earlier instructional technologies could not. Research on prospection, prediction, memory, emotion, and social cognition is portraying learning in a new way as increasingly dynamic, forward-looking, and affectively charged. The problem is the lack of an adequate technological theory for bridging the gap between these two developments. Scientific theories explain learning, and intelligent technologies supply unprecedented capabilities, but neither by itself tells designers how adaptive learning experiences should be structured, governed, and realized in this new regime. This paper proposes a technological theory of learning experience design intended to bridge that widening gap. It proposes conversation as the fundamental structural form of learning interaction. The resulting framework positions AI not as an autonomous instructional authority but as a designed participant operating within humanly defined purposes and constraints. The collision of new learning science and new intelligent technology therefore represents not merely another technological change, but a turning point in the theory and practice of learning experience design.

Submission Synopsis

Problem/Innovation Studied. Learning experience design is approaching a crisis point created by the collision of two revolutions. Artificial intelligence is rapidly expanding what learning technologies can do: converse, generate, adapt, interpret, simulate, and respond in real time. At the same moment, neuroscience and cognitive science are revising long-standing accounts of learning by emphasizing prediction, prospection, memory reconstruction, emotion, and social interaction. These revolutions are advancing faster than the field's design theory. This paper proposes a technological theory for bridging the widening gap between new accounts of human learning and the unprecedented capabilities of intelligent learning systems. It explains how the normally verbose characteristics of AI can be regulated to serve within conversational interaction governed by new strategic rules.

Importance for Educational Technologies. Educational technology has entered a period in which the old division between instructional design and delivery technology is breaking down. Generative systems can now make consequential instructional decisions during runtime rather than merely execute predetermined scripts. That power exposes a fundamental weakness: capability is arriving faster than the principles needed to govern it. Designers must now learn how goals, learner control, emotional conditions, conversational roles, data, adaptive rules, and AI authority are coordinated. Without an adequate technological theory, the field risks allowing the capabilities of the tool to determine the form of instruction. The proposed framework returns to the designer control over strategic choices.

Key Findings. This conceptual paper argues that the present collision of a science and a technology can be organized through a synthesis of technological design theories. The new view of learning through prospection frames learning in terms of continuous cycles of anticipation, action, evaluation, and repair. Conversation provides the structural form in which these cycles can be deliberately shaped. Design layers separate major classes of design decision; dynamic models supply manipulable subject matter; goals regulate the direction of experience across time; runtime rules transform design commitments into adaptive behavior; and governance constrains and focuses the participation of artificial intelligence. Together these elements form a traceable architecture capable of generating responsive learning experiences without surrendering instructional authority to AI.

Implications. This framework implies a substantial shift in the identity and level of work of the learning experience designer. As AI assumes more of the mechanical production of messages, media, code, and variations, the designer's most important work moves upward toward architecture, orchestration, governance, and the explicit specification of intentions and adaptive decision rules. The field is therefore not facing the disappearance of design but a redefinition of it. Research can test the proposed components individually and in combination, while implementations can be judged by responsiveness, learner agency, emotional integrity, traceability, and instructional coherence. This moment provides the opportunity to establish a stronger technological and theoretical foundation for the field.

A Theory Crisis Emerges

Educational technology has seen disruptive inventions before. The blackboard, audiovisual media, programmed instruction, computers, and the Internet each disturbed established practice and forced designers to rethink what instruction could be. The present disruption is different in kind as well as degree. For the first time, a revolution in the science of learning is colliding directly with a revolution in the machinery available for creating learning experiences. One revolution is changing the picture of the learner. The other is changing the powers of the instructional medium. Their trajectories have now crossed.

Artificial intelligence can generate explanations, conduct extended conversations, construct examples, interpret open-ended responses, simulate roles and environments, produce software, and adapt its behavior while an interaction is underway. These capabilities are arriving with extraordinary speed. At the same time, research across neuroscience, cognitive science, memory, emotion, and social cognition increasingly portrays the brain not as a passive receiver of information but as a prospective system continually anticipating, acting, comparing, valuing, and revising.

That convergence should be exhilarating to a field devoted to the design of learning. It should also be unsettling. The conceptual machinery used to design instruction was largely formed under very different technological conditions. Print, film, broadcast media, programmed instruction, and early computer-based instruction required designers to determine most messages, representations, paths, and

responses before instruction began. Even when interaction was available, the range of possible exchanges was normally anticipated and scripted. The new medium breaks that assumption.

A generative system can now make consequential choices at runtime. It can choose wording that the designer never wrote, follow a line of inquiry the designer did not foresee, infer something about the learner, propose a next step, change tone, generate a representation, or produce a response from a vast knowledge resource. The old boundary between design time and delivery time has become porous. Decisions that once belonged unmistakably to the designer can now be made while instruction is occurring.

This is the crisis point. The danger is not simply that AI may make mistakes. The deeper problem is that computational capability is advancing faster than the technological theory needed to govern it. Scientific theory can tell us more about how learning occurs. Artificial intelligence can supply remarkable new services. Neither, however, tells us what a learning experience ought to become, which decisions may properly be delegated, how those decisions should be coordinated, or how the system should remain accountable to human purposes. Without a bridging theory to supply design discipline, the tool itself can begin to dictate the future forms of instruction.

The central argument of this paper is that the bridge can be built. It requires a synthesis of theories rather than a new instructional recipe: a prospective account of learning, a theory of technological knowledge, conversation as a structural form, a layered architecture for design decisions, model-centered content, dynamically managed goals, dynamic management of emotional atmospherics, runtime rules, and explicit governance of AI. Taken together, these ideas describe a learning system that can be adaptive without becoming unbounded, conversational without becoming aimless, and intelligent without surrendering instructional authority.

The Learner Has Changed in Our Theories

The first revolution concerns the learner. A growing body of work treats the brain as fundamentally prospective. Schacter, Addis, and Buckner (2007) argued that remembering the past and imagining the future draw upon much of the same neural machinery. Memory, in this account, is not merely a record of what happened; it is a resource for constructing *what might happen next.* Work on predictive processing and prospection extends this idea across perception, action, emotion, and decision making (Seligman et al., 2016; Barrett, 2017; A. Clark, 2023).

This changes the backdrop against which instruction occurs. The learner is never simply waiting to receive the next message. Before the designer intervenes, learners are already knee-deep in a stream of their own anticipations. Prior experience biases perception. Attention is being selectively allocated. Possible actions are being evaluated. Expected consequences are being compared with incoming evidence. Errors are prompting revision. Learning can be viewed as continuing cycles of anticipation, action, evaluation, and repair. Design can be seen as entering those cycles in hopes of recruiting and influencing their energies.

The new complementary learning systems (CLS) account of memory deepens this picture. The rapid episodic learning associated with hippocampal systems and the slower abstraction and integration associated with neocortical systems each contribute differently to learning (McClelland et al., 1995). Experience is selected, replayed, consolidated, generalized, and reconstructed by these cooperating systems rather than being simply deposited intact. For learning designers, this means that the organization of experience matters as much as the organization of information. Contrast, surprise, action, error, repetition, reflection, and opportunities to revisit experience can influence what is remembered and what becomes stable and usable knowledge.

Emotion adds another force to the prospective cycle. Research shows that feeling and emotion are inseparable from perception, attention, memory, choice, and action (Damasio, 2012; Barrett, 2017; Siegel, 2020). Learners do not merely anticipate what will happen. They anticipate whether what might happen is desirable, threatening, interesting, embarrassing, rewarding, or worth the effort. These valuations give force to goals and shape persistence. They also place emotional climate, trust, uncertainty, agency, and the handling of failure squarely inside the design problem.

The implications are substantial. If learners are already engaged in forward-looking cycles of expectation and adjustment, then instruction does not initiate cognition; it enters cognition already in motion. Design acts within an unfolding landscape. The question becomes not only What information should be presented? but also What is the learner expecting now? What possible future is being made salient? What action will anticipation awaken? How should the system respond when expectation and consequence diverge?

Yet this scientific account does not specify an instructional architecture. It tells us why certain phenomena may matter, not how to build a system that manages them. The first revolution therefore creates the need for an additional revolution.

The Missing Bridge: Technological Theory

Instructional design has long been urged to become more firmly grounded in theory (Schmidt et al., 2020). The advice is sound, but the word theory often carries an implicit assumption: that scientific theory is the principal source from which good design should be derived. That assumption becomes increasingly inadequate as the artifacts being designed grow more complex.

Scientific theories explain and predict phenomena. Design, however, must transform existing situations into preferred ones (Simon, 1996). The historical development of flight provides a useful analogy. Understanding lift did not specify a workable aircraft. Sustained flight also required technological knowledge of structure, control, propulsion, materials, instrumentation, fabrication, and the coordination of all these elements. Vincenti (1990) showed that engineers accumulate distinctive forms of knowledge through designing and making: design criteria, operational principles, quantitative relationships, practical considerations, instrumentalities, and other knowledge whose validity is demonstrated in artifacts that work.

Learning experience design requires its own technological theory of this kind. Such theory is organized knowledge for conceiving, structuring, constructing, operating, and improving learning systems. It is informed by scientific theory but not reducible to it. It grows through design, implementation, testing, comparison, failure, revision, and the public articulation of principles. It supplies languages for describing structures and relationships that scientific accounts of cognition do not provide (Gibbons & Rogers, 2009; Gibbons, 2014).

The arrival of AI makes this distinction impossible to ignore. A language model may be capable of explaining, questioning, role-playing, summarizing, classifying, generating examples, or altering tone. This capability, however, does not confer educational authority nor establish design discipline. It does not determine what goals deserve pursuit, which model of subject matter should count as authoritative, how much productive struggle should be retained, what inferences about a learner are legitimate, when initiative should shift, or what ethical limits should constrain the exchange. These are technological design judgments.

The collision now becomes clear. New science is providing a more dynamic account of learning at precisely the moment new technology is providing a more dynamic medium of instruction. Between them lies an underdeveloped space of advanced design technology. A technological theory of learning experience design must occupy that space, describing how scientific insight translates indirectly through design rules into structures that intelligent systems can enact without allowing those systems to define their own purposes.

Conversation Becomes the Structural Kernel

The first structural move is to treat conversation as more than a teaching technique. Conversation can be understood as a general form for coordinated human learning activity—arguably the most iconic form. Participants continually form actions, interpret intentions, take turns, project what is relevant next, detect breakdowns, and repair misunderstandings. Conversation analysis (Sidnell & Stivers, 2013) and pragmatics (H. Clark, 2012) make visible a fact that traditional message-centered design often hides: meaning is produced not by isolated utterances but through sequences of mutually interpreted action (Levinson, 2013).

Instructional events display the same structure. Explaining, asking, prompting, demonstrating, challenging, confirming, redirecting, scaffolding, and repairing are complex conversational actions. Even media that appear one-directional can be understood as constrained modes of a larger conversation. A text anticipates a reader. A demonstration directs attention. A simulation invites action and returns consequences. A test asks and evaluates. Conversation is therefore not confined to spoken dialogue; it is a structural way of describing how meaning and action are coordinated over time through joint action (Gibbons & Boling, 2021).

Prospection and conversation fit tightly together. Every conversational move sets up new possibilities for what may happen next. Participants interpret present actions partly by anticipating their likely

continuation. A misunderstanding becomes visible when the projected trajectory breaks, and repair restores coordination. In a designed learning conversation, these same dynamics can be exploited deliberately: invite a prediction, expose it to consequence, detect a mismatch, draw attention to the discrepancy, and support repair.

Generative AI suddenly makes rich conversational exchange technologically abundant. Unharnessed, that abundance can easily be mistaken for instructional adequacy. A fluent system can still wander, over-explain, answer too quickly, conceal uncertainty, accommodate misconceptions, or pursue a locally plausible path that undermines a larger instructional purpose. Conversational power therefore intensifies rather than removes the need for structure and strategic discipline. But that discipline, administered within an unpredictable conversational flow, must be designed at a higher level of abstraction.

The design for learning conversation still requires a governing architecture. It must know what project is underway, what models define the subject matter, what goal currently has priority, what forms of initiative are legitimate, what learner actions are available, what evidence may alter the system's judgment, and what kinds of repair are authorized. Conversation supplies the living form of the experience. Architecture prevents that form from dissolving into improvisation.

An Architecture for the Collision

The remaining components of this proposed theory can be understood as an architectural response to the crisis. They separate design decisions that become dangerously entangled when a generative system is allowed to perform many functions at once. The architecture does not prescribe a particular software platform. It identifies decision spaces that must remain visible to the designer if intelligent learning systems are to be understandable, governable, and improvable.

Design Layers

Design-layer theory (Gibbons & Rogers, 2009) treats a learning system as the integrated and coordinated operation of multiple classes of decision. A prototypical set of design layers includes content, message, representation, strategy, control, data management, relational-emotional, and governance concerns. That number is not sacred; the principle is separation with coordination.

Content design determines what meanings, models, relationships, and processes are available, independent of their mediation. Message design concerns communicative intentions. Representation design makes those intentions perceptible through language, graphics, sound, action, or other mediated forms. Strategy design organizes trajectories of interaction. Control design determines what actions the learner and system can take and how. Data management design preserves relevant state across time. Relational-emotional design addresses climate, trust, challenge, support, and the interpretation of learner state. Governance design determines authority, constraints, escalation, accountability, permissible AI behavior, and the ultimate governing role of the human decision maker.

The urgency of this separation becomes obvious in AI-supported systems. If one generative process is allowed simultaneously to determine content, strategy, wording, learner interpretation, and progression, then design commitments disappear inside the output of the content model. Layering pulls those commitments back into perspective. A decision about authoritative content need not be confused with the wording of an explanation. A decision about learner control need not be confused with strategy. A decision about what AI can generate need not be confused with what AI is authorized to decide.

Layering does not make a complex learning system simple. It makes the complexity inspectable and manageable. That distinction matters. The theory crisis is not solved by reducing complexity but by organizing it so that designers can see where consequential choices are being made and can alter one class of choices without blindly disturbing all the others.

In prospective terms, every layer shapes a field of possible futures. Content defines what can become meaningful; messages propose communicative actions; representations make those actions perceivable; strategies frame possible interaction trajectories; controls permit or constrain action; data collection preserves the past so future decisions can differ from earlier ones. A learning product is no longer merely a fixed artifact. It becomes a coordinated system for *generating bounded possibilities*.

Models as the Substance of Adaptive Interaction

The content layer deserves special emphasis because generative technology makes it dangerously easy to equate fluent language with reliable subject matter. Model-centered instruction takes the opposite position. Content is represented in terms of systems, structures, relationships, processes, and behaviors that can be examined, manipulated, and reasoned about (Gibbons, 2001). A dynamic model changes state in response to action and therefore permits the learner to encounter consequences rather than merely read descriptions of them.

This orientation extends a long tradition in intelligent tutoring systems, where subject-matter, learner, pedagogical, and communication models are coordinated to produce adaptive interaction (Wenger, 1987; Collins & Brown, 1988; Johnson et al., 2000; VanLehn, 2011; Woolf, 2009; Rus et al., 2013; Nye et al., 2014; Kulik & Fletcher, 2016). The present proposal brings that modeling tradition to the designer within a broader architecture of conversation, prospection, layers, goals, and governance.

The distinction becomes critical with generative AI. If the language model itself is treated as the content model, source of explanation, evaluator, strategist, and arbiter of progress, then the instructional system becomes difficult to inspect. A model-centered architecture allows AI to assist without becoming epistemically sovereign. The system can ask AI to explain a defined relationship, generate an analogy, interpret a learner's utterance, or propose a representation while retaining explicit models of what states and relationships count as valid.

This restores a crucial design principle: intelligent generation may vary the form of the experience without silently redefining its substance.

Dynamic Goals

If conversation gives the system form and models give it substance, goals give it direction. Traditional instructional goals are often written before instruction and revisited primarily during assessment. An adaptive system requires something more dynamic. Long-range intentions must be related to intermediate and momentary goals; goals may need to be suspended, decomposed, repaired, negotiated, or replaced as the interaction unfolds.

A goal layer and associated assessments can maintain current priorities, evaluate whether a state is sufficiently stable to count as attainment, diagnose the need for repair, select subgoals after failure, manage transitions, and reopen learner initiative. This turns goals from static declarations into runtime regulators.

The prospective account of learning makes this shift especially important. Human action is drawn toward valued possible futures (Baumeister et al., 2016). A learner's goal ecology therefore includes a much broader curriculum that targets curiosity, confidence, ownership of learning, desire for learning, social connection ability, and many other outcomes. Designers cannot measure all of these perfectly, and they should not pretend to. They must instead specify which inferences are justified, how uncertainty is handled, and what actions are legitimate when evidence is weak. The point is not to make the system psychologically omniscient. It is to prevent hidden psychological assumptions from being made irresponsibly by an intelligent engine.

AI as a Governed Participant

The most important architectural principle is also the simplest: artificial intelligence participates in the learning interaction but does not own its purposes. Authority to define goals, constraints, responsibilities, and permissible actions remains with the design and with the humans and institutions accountable for it.

This rejects an increasingly seductive metaphor: the AI tutor as an autonomous instructional mind hidden behind the interface. In a governed architecture, AI is a serviceable participant. It may interpret language, propose messages, create representations, classify possibilities, search resources, or suggest candidate actions. But main system logic determines when AI is called, what context is supplied, what constraints apply, what kind of output is requested, and how the result is bound back into the learning system.

The distinction is more than technical. It is a defense of educational purpose. The more capable AI becomes, the easier it is to allow its fluency to conceal transfers of authority that were never consciously designed. Governance makes those transfers explicit. It addresses decision rights, limits of adaptation, transparency, oversight, data use, ethical boundaries, escalation, error handling, and accountability.

Data management supports this governance by preserving the basis for consequential decisions. If the system changes feedback, progression, or support, designers should be able to reconstruct why. More data do not automatically improve this process. Excessive collection can create new risks and unjustified inference. The governing principle is sufficiency for the design, not maximal capture.

The field's challenge is therefore not to make AI powerful enough to teach. It is already becoming powerful enough to do many teaching-like things. The challenge is to make its participation legible, bounded, and answerable to instructional intentions.

From Architecture to Action

A theory that remains at the architectural level cannot resolve the crisis. The design must survive contact with runtime. The proposed bridge is a system of explicit rules that converts relatively stable design commitments into moment-by-moment action. Rules can consult model states, learner actions, conversation history, goals, data, and governance constraints, then authorize a next operation. Some operations are deterministic. Others may invoke AI for bounded interpretation or generation.

This changes the meaning of adaptive instruction. Adaptation need not mean that an intelligent engine is free to improvise. It can mean that designers specify a structured space of legitimate variation. The experience is generated rather than completely scripted, but its generation occurs inside explicit constraints.

Gibbons (2026) illustrates this progression from concept to architecture, operational specification, rules, and executable software. Learner actions alter a dynamic model; the system observes state changes, applies constraints, updates goals, and produces consequences and instructional responses. Its significance lies less in the simulation domain than in its traceability. Working behavior can be traced backward to design concepts, while design intentions can be traced forward into runtime operations.

Traceability should become a central criterion for intelligent learning systems. When a system performs a consequential act, a designer should be able to ask: Why was this action permitted? What goal did it serve? What learner evidence influenced it? What model state was consulted? What rule authorized it? Did AI contribute, and if so, under what constraints?

A system that cannot answer those questions may still be dazzling. It may even appear highly adaptive. But educationally, it remains difficult to govern, investigate, improve, or defend. At the very moment AI increases the apparent intelligence of the interface, technological theory must increase the intelligibility of the design.

What the Crisis Demands of Designers

The collision of new science and new technology changes the unit of design. Instruction can no longer be understood primarily as a lesson, module, presentation, or predetermined sequence. Increasingly it

must be understood as an architecture of coordinated processes capable of generating experience. Designers still make familiar decisions about content, message, representation, activity, feedback, and assessment, but those decisions become finer-grained, more explicit, and more tightly coupled to runtime behavior.

This shift also changes the professional role of the designer. As AI becomes increasingly able to generate prose, media, code, examples, and preliminary design alternatives, professional value will migrate away from the manual production of every surface feature. It will concentrate more heavily in deciding what should exist, what futures should be made possible, what models should be trusted, what kinds of interaction are educationally desirable, how much agency learners should retain, how emotional conditions should be protected, and what authority an intelligent service should possess.

This is not the disappearance of the designer. It is an elevation of the design problem. The mechanical burden of production may decline while the burden of judgment increases. Designers will be asked to specify intentions, limits, and designs at a level of detail that previous media often allowed them to leave implicit. There were always design factors that needed to be included; now we are more aware of them and have a way to deal with them.

This framework also opens a research agenda. Individual elements can be tested: conversational repair rules, goal-management procedures, representations of learner state, forms of model coordination, governance constraints, layer decision patterns, or strategies for binding AI services into rule-governed interactions. Complete systems can be compared for responsiveness, transfer, learner agency, emotional integrity, instructional coherence, and auditability. Because the architecture separates design concerns, failures may be localized rather than disappearing into a single generative black box. Terminologies become more coherent and stable.

The framework is deliberately independent of any particular AI engine. Today's dominant systems will change. The underlying crisis will remain: computational capacity will continue to outrun our ability to decide how it should be used unless technological theory advances with it. So long as learning remains a human enterprise directed toward valued futures, designers will need explicit goals, roles, models, constraints, and responsibilities.

Conclusion: A Turning Point for the Field

Educational technology has crossed turning points before, but the present one is unusually consequential. A new science of learning is portraying the learner as prospective, predictive, affective, and socially engaged at the same moment that a new class of technology is becoming capable of participating in those processes dynamically. The two revolutions are no longer approaching one another. They have collided.

That collision creates both danger and opportunity. The danger is that the field will mistake capability for theory and permit the properties of generative technology to determine the architecture of learning. The

opportunity is to use this disruption to articulate a stronger technological foundation for learning experience design.

The synthesis proposed here is one attempt at that foundation. Prospection describes the forward-looking dynamics of learning. Conversation supplies a structural form through which anticipation, action, interpretation, breakdown, and repair can be coordinated. Design layers expose major classes of decision. Dynamic models anchor subject matter. Goals regulate direction across time. Runtime rules convert architecture into behavior. Governance places AI inside humanly established purposes and constraints.

The claim, therefore, is not merely that conversational AI may improve instruction. It is that AI-supported learning experiences must be designed as governed systems whose adaptive behavior emerges from explicit architectures. Artificial intelligence becomes a participant in those systems, not their source of purpose. The designer's work moves upward toward intentions, structures, relationships, rules, boundaries, and accountability.

That movement marks the crisis in its second meaning: a point of judgment and a turning point. The field can attempt to fit an unprecedented medium into design concepts inherited from earlier technologies, or it can build the technological theory required by the new conditions. The choice will influence whether AI becomes merely another instructional device or a catalyst for a more mature conception of learning experience design.

The theory offered here should remain open, revisable, and contestable. Its value will be determined by whether its concepts help designers create systems that are more coherent, more responsive, more humane, and more accountable. But one conclusion already seems difficult to escape: the science has changed, the technology has changed, and the design theory that joins them must now change as well.

Declarations

Author Contribution

This paper was authored entirely by the principal author.

References

1. Barrett, L. F. (2017). *How emotions are made: The secret life of the brain.* Houghton Mifflin Harcourt. ISBN 978-0-544-13331-0 (hardcover); 978-0-544-12996-2 (eBook).
2. **Baumeister, R. F., Vohs, K. D., & Oettingen, G.(2016).** Pragmatic prospection: How and why people think about the future. *Review of General Psychology, 20*(1), 3–16. https://doi.org/10.1037/gpr0000060

3. Clark, A. (2023). *The experience machine: How our minds predict and shape reality.* Pantheon. ISBN 978-1-5247-4845-6 (hardcover).
4. Clark, H. H. (2012). *Using language.* Cambridge, UK: Cambridge University Press. ISBN 9780511620539
5. **Collins, A., & Brown, J. S. (1988).**The computer as a tool for learning through reflection. In H. Mandl & A. Lesgold (Eds.), *Learning issues for intelligent tutoring systems,* 1–18. Springer. https://doi.org/10.1007/978-1-4684-6350-7. ISBN 978-0-387-96616-8.
6. Damasio, A. (2012). *Self comes to mind: Constructing the conscious brain.* Vintage. ISBN 978-0-307-47495-7 (paperback).
7. Gibbons, A. S. (2001). Model-centered instruction. *Journal of Structural Learning and Intelligent Systems,* 14(4), 511–540.
8. Gibbons, A. S. (2014). *An architectural approach to instructional design.* New York: Routledge. ISBN 978-0-415-80739-5
9. Gibbons, A. S. (2020). What is instructional strategy? Seeking hidden dimensions. *Educational Technology Research and Development*, 68(6), 2799–2815. https://doi.org/10.1007/s11423-020-09820-2
10. ERIC record: https://eric.ed.gov/?id=EJ1277911
11. Gibbons, A. S. (2026). *A technological theory of learning experience design: The synthesis of learning theory and design technology.* EdTechBooks. Available online at: https://edtechbooks.org/edtechnological_theory_of_learning_experience_design
12. Gibbons, A. S., & Boling, E. (2021). *Conversational forms of instruction and message layer design.* New York: Springer. https://doi.org/10.1007/978-3-030-84220-8, ISBN 978-3-030-84219-2 (softcover); 978-3-030-84220-8 (eBook).
13. Gibbons, A. S., & Rogers, P. C. (2009). The architecture of instructional theory. In C. M. Reigeluth & A. A. Carr-Chellman (Eds.), *Instructional-design theories and models, Volume III: Building a common knowledge base* (pp. 305–326). Routledge. ISBN 978-0-8058-6456-4.
14. **Johnson, W. L., Rickel, J. W., & Lester, J. C. (2000).** Animated pedagogical agents: Face-to-face interaction in interactive learning environments. *International Journal of Artificial Intelligence in Education, 11*(1), 47–78.
15. **Kulik, J. A., & Fletcher, J. D. (2016).** Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research, 86*(1), 42–78. https://doi.org/10.3102/0034654315581420
16. Levinson, S. C. (2013). Action formation and ascription. In J. Sidnell & T. Stivers (Eds.), *The handbook of conversation analysis* (pp. 103–130). Wiley-Blackwell. https://doi.org/10.1002/9781118325001.ch6
17. McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3), 419–457. https://doi.org/10.1037/0033-295X.102.3.419

18. **Nye, B. D., Graesser, A. C., & Hu, X. (2014).**AutoTutor and family: A review of 17 years of natural language tutoring. *International Journal of Artificial Intelligence in Education, 24*(4), 427–469. https://doi.org/10.1007/s40593-014-0029-5
19. **Rus, V., D'Mello, S., Hu, X., & Graesser, A. C. (2013).** Recent advances in conversational intelligent tutoring systems. *AI Magazine, 34*(3), 42–54. https://doi.org/10.1609/aimag.v34i3.2485
20. Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience*, 8(9), 657–661. https://doi.org/10.1038/nrn2213
21. Schmidt, M., Tawfik, A. A., Jahnke, I., Earnshaw, Y. & Huang, R. (2020). Learner and user experience research: An introduction for the field of learning design and technology. EdTechBooks. Online at: https://edtechbooks.org/ux
22. Seligman, M. E. P., Railton, P., Baumeister, R. F., & Sripada, C. (Eds.). (2016). *Homo prospectus.* Oxford University Press. ISBN 978-0-19-937447-2.
23. Sidnell, J. & Stivers, T. (Eds.), (2013). *The handbook of conversation analysis*. New York: Wiley-Blackwell. ISBN:9781444332087
24. Siegel, D. J. (2020). *The developing mind: How relationships and the brain interact to shape who we are (3rd ed.)*. Guilford Press. ISBN 978-1-4625-4275-8 (hardcover).
25. Simon, H. A. (1996). *The sciences of the artificial (3rd ed.)*. MIT Press. ISBN 978-0-262-69191-8 (paperback); 978-0-262-26449-5 (eBook).
26. VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist, 46*(4), 197–221. https://doi.org/10.1080/00461520.2011.611369
27. Vincenti, W. G. (1990). *What engineers know and how they know it: Analytical studies from aeronautical history*. Johns Hopkins University Press. ISBN 978-0-8018-3974-0. https://doi.org/10.56021/9780801839740
28. Wenger, E. (1987). *Artificial intelligence and tutoring systems: Computational and cognitive approaches to the communication of knowledge.* Morgan Kaufmann. ISBN 978-0-934613-26-2. https://doi.org/10.1016/C2013-0-07697-9
29. Woolf, B. P. (2009). *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning.* Morgan Kaufmann. ISBN 978-0-12-373594-2. https://doi.org/10.1016/B978-0-12-373594-2.X0001-9