

Meta-Books: Restructuring Organisational Knowledge for Agentic AI Reasoning

Sima Fatemipour

Metis Co-Intelligence, UK

sfatemipour@gmail.com

Abstract: Organisations hold significant knowledge capital that their AI systems cannot access as knowledge. The knowledge bases most organisations maintain — manuals, onboarding documents, process documentation, and accumulated expertise — are structurally incompatible with AI reasoning: there are no dependency structures to traverse, no role classifications to honour, no relational context to surface. This paper introduces meta-books as a knowledge management design response to this structural problem. A meta-book transforms a conventional organisational document into a typed, dependency-linked knowledge graph. Source content is decomposed into typed knowledge objects — concepts, procedures, examples, controversies, and reflective insights — connected by semantically meaningful edges such as prerequisite-of, illustrates, contradicts, and applies-to. Every node carries pedagogical metadata: depth layer, role classification, engagement time, and prerequisite links. Unlike retrieval-augmented generation, meta-books restructure source content itself, making dependency, context, and role visible to AI at the point of reasoning, and support cross-source integration allowing knowledge nodes to migrate across individual, team, and organisational graphs without losing relational context. Meta-books reposition organisational knowledge as a compound, AI-traversable asset with implications across the full employee lifecycle — from recruitment and adaptive onboarding to knowledge retention beyond individual tenure. The paper positions meta-books against learning object metadata, adaptive hypermedia, semantic web standards, and enterprise knowledge graphs, demonstrating that no existing standard simultaneously occupies the intersection of typed knowledge structure, pedagogical metadata, and AI queryability. Five design principles are derived using design science research methodology and evaluated through illustrative decomposition of a representative organisational document. The paper’s most distinctive argument is that meta-books are the first organisational KM framework explicitly aligned with the cognitive architecture of human learning — integrating schema theory, cognitive load theory, dual coding, retrieval practice, the spacing effect, meaningful learning, and connectionist neuroscience — with implications for universities, professional certification, and all structured knowledge delivery systems.

Keywords: Knowledge graphs, Agentic AI, Organisational learning, Knowledge representation, Cognitive load, Learning sciences

1. Introduction

The knowledge organisations hold is largely inaccessible to the AI systems now deployed to reason over it — organisational documents were simply never designed to support machine reasoning. Despite decades of investment in knowledge management (KM) systems, manuals, onboarding materials, and accumulated expertise remain structurally incompatible with the artificial intelligence (AI) systems now querying them.

The rise of large language models (LLMs) has transformed expectations of what AI can do with organisational knowledge. Systems employing retrieval-augmented generation (RAG) — a technique that grounds AI responses in external document collections — allow LLMs to query internal repositories. In practice, however, these systems retrieve text fragments. They surface sentences matching a query but cannot traverse relationships between concepts, identify what an employee already knows, or distinguish a foundational principle from an advanced application. The knowledge exists in the document; the intelligence cannot reach it.

The gap reflects a failure of knowledge architecture, not of AI capability. Organisational documents are written for sequential human reading: linear, undifferentiated, and untyped. A concept and a procedure sit in the same paragraph; a foundational principle and an advanced application occupy the same section. There are no machine-readable signals indicating what type of knowledge a passage represents, what it depends upon, or who it is for. Retrieval is not the problem. The source material was never structured to enable reasoning in the first place.

This paper introduces meta-books as a design-science response — an intervention grounded in established theory and evaluated against the requirements it claims to meet. A meta-book transforms a conventional organisational document into a typed, dependency-linked knowledge graph, enabling agentic AI to traverse rather than merely search it. The contribution is threefold: articulating the structural incompatibility as a distinct KM problem requiring a design-level response; introducing the meta-book framework with its components and design principles; and advancing a novel theoretical argument that meta-books are the first KM framework explicitly aligned with the cognitive architecture of human learning, drawing on schema theory, cognitive load theory, dual coding, and connectionist neuroscience. The paper proceeds from literature review and

methodology through the framework, design principles, and organisational implications, before developing the learning-science argument in the discussion.

2. Literature Review

2.1 Organisational Knowledge Management

Organisational knowledge is relational, contextual, and dynamic — properties that linear document repositories structurally suppress. Alavi and Leidner (2001) established that KM systems must support knowledge creation, transfer, and application rather than merely information storage, while Davenport and Prusak (1998) positioned knowledge as a strategic asset that resists the formalisation information systems require. Nonaka and Takeuchi (1995) identified tacit knowledge as the most valuable and least transferable organisational resource, arguing through their SECI model (Socialisation, Externalisation, Combination, Internalisation) that knowledge creation is a dynamic process that linear document management fundamentally disrupts. Senge (1990) demonstrated that continuous knowledge flow is a precondition of the learning organisation — a condition static repositories inhibit. Cohen and Levinthal (1990) showed that absorptive capacity is path-dependent and cumulative, requiring structured knowledge access, and Wenger (1998) confirmed that knowledge flows most richly through communities of practice where it is enacted rather than stored. A knowledge base that suppresses relational structure cannot support relational reasoning.

2.2 Retrieval-augmented Generation and its Structural Limitations

Lewis et al. (2020) introduced RAG as a mechanism for grounding LLM responses in external documents, reducing hallucination and improving factual accuracy. Gao et al. (2024) demonstrate in a comprehensive survey that RAG retrieves text chunks by semantic proximity, without regard for epistemic type, dependency structure, audience, or knowledge state — so an employee querying an escalation procedure may receive the relevant paragraph but not the prerequisite concept, the role classification, or the reflective insight that qualifies when procedure should yield to judgement. Pan et al. (2024) identify the complementarity of LLMs and knowledge graphs — LLMs provide fluency, graphs provide relational precision — and argue their integration is necessary for enterprise knowledge tasks requiring both. Meta-books are designed at exactly this intersection.

2.3 Existing Standards and the gap Meta-books Fill

Meta-books must be positioned against existing standards that address related problems. Five architectural requirements emerge: epistemic node typing (R1); semantically typed dependency relationships (R2); learner-state-responsive pedagogical metadata (R3); AI queryability as primary design objective (R4); and cross-source node portability (R5). Table 1 maps each prior approach against these requirements: full tick = substantively met; partial tick = approached but not fulfilled; cross = not addressed.

Table 1: Gap analysis — prior standards against five meta-book requirements

Standard	R1Epistemic typing	R2Dep. edges	R3Ped. metadata	R4AI queryability	R5Portability	Designed to meet all 5?
LOM (IEEE, 2002)	No	No	Partial	No	No	No
IMS Learning Design (2003)	No	Partial	Partial	No	No	No
SKOS (Miles & Bechhofer, 2009)	Partial	Partial	No	No	Partial	No
OWL Ontologies (Hogan et al., 2021)	Partial	Yes	No	Partial	No	No
RAG (Lewis et al., 2020; Gao et al., 2024)	No	No	No	Partial	No	No
Adaptive Hypermedia (Brusilovsky, 2001)	No	No	Partial	No	No	No
Meta-Books (this paper)	Yes	Yes	Yes	Yes	Yes	Yes

The Learning Object Metadata standard (LOM) (IEEE, 2002) introduced content metadata for cataloguing but encodes no relational structure and no epistemic typing — cataloguing and graph traversal are categorically different operations. IMS Learning Design (IMS LD) (IMS Global, 2003) added pedagogical sequencing at course level, but sequences are pre-authored rather than generated by AI reasoning, and individual knowledge objects remain untyped. The Simple Knowledge Organization System (SKOS) (Miles and Bechhofer, 2009) provides taxonomic relationships in the Resource Description Framework (RDF) but cannot distinguish a prerequisite from an illustrative relationship or encode depth layers and role classifications. Web Ontology Language (OWL) ontologies (Hogan et al., 2021) offer the richest prior architecture — formally typed relationships supporting logical inference — but encode world knowledge rather than learning knowledge; pedagogical metadata is absent. RAG retrieves by semantic proximity, rendering epistemic type and dependency structure invisible. Adaptive hypermedia (Brusilovsky, 2001; Conlan et al., 2002) achieves personalised delivery but locates the intelligence in a separate delivery engine, requiring specialist authoring and maintained learner models. The meta-book framework, introduced in Section 4, is designed to occupy this intersection — combining OWL’s typed relational rigour, LOM’s metadata aspiration, IMS LD’s learner-state responsiveness, and an explicit AI queryability objective no prior standard was designed to meet.

3. Methodology

This paper employs design science research (DSR) methodology as articulated by Hevner et al. (2004) and operationalised through the framework of Peffers et al. (2007). DSR is appropriate where the research objective is the creation of a novel artefact addressing an identified organisational problem. The meta-book is positioned as a design artefact at the construct and model level — defining a new vocabulary for knowledge representation and a structural model for its implementation. Design principles are derived through literature synthesis across KM theory, knowledge representation standards, adaptive learning systems, and AI reasoning architectures, combined with design argumentation — each principle justified by demonstrating that its absence would prevent satisfaction of one or more of the five requirements in Section 2. Evaluation takes the form of analytical demonstration — an illustrative decomposition of a hypothetical onboarding document — consistent with DSR’s acceptance of scenario-based evaluation at the conceptual contribution stage (Hevner et al., 2004). Section 4 presents the resulting artefact; Section 5 derives the design principles; the illustrative evaluation appears in Section 4.5. The comparison in Table 1 reflects the authors’ own design intentions rather than independent empirical verification; demonstrating that meta-books produce measurably superior AI reasoning outcomes compared to RAG in live settings remains the primary empirical objective for future research.

4. The Meta-book Framework

4.1 Definition

A meta-book is a conventional organisational document transformed into a typed, dependency-linked knowledge graph. The transformation operates at the content architecture level: source material is decomposed, typed, related, and annotated — not summarised or replaced. Transformation rather than replacement preserves existing knowledge assets. Three interdependent components comprise the framework: nodes (typed knowledge objects), edges (typed relationships), and attributes (pedagogical metadata).

4.2 Nodes: Typed Knowledge Objects

Five node types are defined, reflecting the epistemic diversity of organisational knowledge — declarative, procedural, illustrative, contested, and experiential — grounded in the cognitive architecture developed in Section 7. Concepts are definitional nodes establishing foundational vocabulary. Procedures encode how something is done — sequences, decision points, conditions. Examples ground abstract knowledge in concrete instances. Controversies represent contested claims or competing approaches. Reflective insights encode integrative knowledge — lessons learned, failure patterns, and contextual judgements developed through practice — the closest approximation of tacit knowledge that explicit structure can achieve, and the primary vehicle for capturing experienced practitioners’ knowledge before departure.

4.3 Edges: Typed Relationships

Five typed edges are defined: prerequisite-of (A must precede B); illustrates (an example grounds a concept or procedure); extends (progressive depth); contradicts (competing claims or practices); and applies-to (connecting theoretical nodes to practical contexts). Typed edges are the mechanism by which AI reasoning differs from retrieval: a prerequisite edge enables knowledge-state assessment; a contradicts edge surfaces epistemic tension rather than false consensus.

4.4 Pedagogical Metadata Attributes

Every node carries: depth layer (Core / Applied / Integrative); role classification (all staff / specific roles / senior practitioners); prerequisite links (queryable dependency enumeration); and estimated engagement time (enabling time-constrained pathway composition). Core nodes are designed for initial onboarding; Applied nodes for employees in early active practice; Integrative nodes for experienced practitioners consolidating advanced understanding. Depth layer and role classification combine to generate personalised pathways without human curation: an AI querying for a new graduate holds senior-classified nodes in reserve; one querying for a Core-mastered practitioner advances to Applied. These attributes transform the graph from a knowledge representation into a reasoning substrate for personalised, role-sensitive delivery.

4.5 Illustrative Decomposition

Table 2 demonstrates meta-book decomposition applied to a hypothetical onboarding extract on client engagement protocols. A new graduate analyst receives nodes N01 → N02 → N05 → N08 in order; an experienced lateral hire enters at N09; a senior manager conducting a governance review receives N11 and N12 as primary nodes. These pathways are generated through AI reasoning over graph structure — a capability unavailable to RAG-based systems operating on untyped linear text.

Table 2: Illustrative meta-book node decomposition — client engagement onboarding extract

Node	Content summary	Type	Depth	Role	Prereqs	Primary edges
N01	Definition of a client engagement	Concept	Core	All	None	prereq-of N02, N05, N09
N02	Engagement tier classification (T1/T2/T3)	Concept	Core	All	N01	prereq-of N05; extends N01
N05	Procedure: raising a client status report	Procedure	Applied	All	N01, N02	prereq-of N06; illustrates N08
N08	Worked example: RAG status for T2 engagement	Example	Applied	All	N02, N05	illustrates N05
N11	Debate: formalise informal escalation handling?	Controversy	Integrative	Senior/Mgrs	N09, N10	contradicts N10 assumption
N12	Insight: relationship drift as primary failure	Reflective insight	Integrative	Senior	N05, N09, N10	contradicts N05; applies-to N09

5. Design Principles

Five design principles govern the construction of a valid meta-book, each derived by design argumentation — its absence demonstrably prevents satisfaction of one or more of the five architectural requirements identified in Section 2.

P1 — Modularity. Every node must be self-contained enough to be accessed independently, enabling AI traversal without requiring sequential reading. Node authors must therefore write for independent access — each node must carry sufficient context to be understood without reading what precedes it, satisfying R1 and enabling R2.

P2 — Typed relationships. Edges must carry semantic meaning beyond sequence, distinguishing prerequisite-of from illustrates from extends from contradicts from applies-to. Untyped connections reduce the meta-book to a sophisticated hyperlink structure, forfeiting the reasoning capacity typed edges enable. Typed edges directly satisfy R2 and are the mechanism through which R4 becomes possible.

P3 — Pedagogical metadata completeness. Every node must declare its type, depth layer, role classification, engagement time, and prerequisite links. Metadata is not supplementary — it is constitutive of the reasoning substrate. Incomplete metadata produces incomplete reasoning, directly violating R3.

P4 — AI queryability. The graph must be implemented in formats natively traversable by AI reasoning systems, not stored as annotated narrative text. A document with metadata remains a document; a graph with typed nodes and edges is a reasoning substrate. This principle directly satisfies R4.

P5 — Cross-level graph compatibility. Nodes designed for a single graph context cannot migrate without loss of meaning. Nodes must be importable into individual and team knowledge graphs without losing relational context, consistent with Wenger's (1998) communities of practice model, directly satisfying R5.

Together these five principles constitute a validity specification for meta-book construction — conditions that, if satisfied, produce an artefact capable of AI traversal, personalised delivery, and cross-level knowledge integration.

6. Implications

The meta-book framework has implications across the full employee lifecycle — from recruitment through onboarding, tacit knowledge retention, and strategic advantage — and beyond into broader knowledge delivery systems.

6.1 Recruitment and Knowledge State Assessment

Current recruitment relies on CVs, interviews, and credential proxies — indicators of probable knowledge rather than direct assessment. The outcome is binary: hire or not hire, based on inference. Where role-relevant knowledge is encoded as a typed meta-book graph, each node carries a role classification and depth layer specifying what is required at each level. A structured task mapped against the graph identifies which nodes a candidate traverses independently, which require scaffolding, and which are absent — generating a precise knowledge-state profile. An organisation can then determine the onboarding pathway that closes the gap most efficiently, transforming recruitment from binary selection into knowledge-state assessment.

6.2 Adaptive Training and Onboarding

Current onboarding delivers fixed content sequences treating all new employees as epistemically equivalent. A meta-book graph traversed by an AI system identifies each employee's knowledge state, determines the appropriate entry point, generates a personalised dependency-ordered pathway, and accelerates through mastered nodes — delivering the personalised, dependency-ordered progression that fixed content sequences cannot.

6.3 Tacit Knowledge Capture, Retention, and Transfer

Davenport and Prusak (1998) identified employee departure as a fundamental knowledge vulnerability — expertise accumulated over years exits with the practitioner. Reflective insight nodes provide a structured mechanism for externalising this tacit knowledge — the evaluative judgements and contextual reasoning practitioners hold — the knowledge Nonaka and Takeuchi (1995) identified as most valuable and least transferable in the SECI externalisation process. An organisation that systematically elicits reflective insights from departing practitioners builds an organisational memory that AI systems can surface for successors — making tacit knowledge transfer not merely aspired to but structurally possible.

6.4 Strategic Advantage and AI Training Substrates

Organisations that structure knowledge for AI reasoning gain a compounding advantage: as Haenlein and Kaplan (2019) note, AI value is determined by knowledge quality, not model sophistication. Meta-books function as organisation-specific training substrates, enabling AI models grounded in verified institutional knowledge rather than generic pre-training data.

6.5 Application Beyond Organisational KM

The meta-book framework extends beyond its organisational origins. In higher education, meta-book graphs would allow AI tutoring systems to assess prior knowledge and generate personalised learning pathways calibrated to each student's zone of proximal development. In professional certification — medical training, legal education, accounting qualifications — meta-books enable AI to map candidate knowledge against competency graphs and identify the most efficient path to certification readiness. In school-based education, AI traversal of curriculum meta-books surfaces only nodes for which each student has the prerequisite foundation, implementing the spacing effect at the system level. Across all contexts, meta-books reposition knowledge delivery from a broadcast model to a precision model: the right knowledge, in the right order, at the right depth, for the right learner, at the right time.

7. Discussion

7.1 Alignment with Cognitive Architecture and Learning Sciences

We argue that meta-books are the first organisational KM framework explicitly aligned with the cognitive architecture of human learning — an alignment that is structural rather than incidental, grounded in eight converging theoretical traditions that suggest why graph-structured knowledge delivery may produce superior learning outcomes to linear document delivery.

Schema theory (Bartlett, 1932; Rumelhart, 1980) establishes that new information is meaningfully encoded when anchored to prior cognitive structures — the brain builds knowledge as interconnected schemas of related concepts, procedures, and contexts. Bartlett's foundational work on reconstructive memory demonstrated that unanchored content degrades and distorts over time. Meta-book prerequisite links formalise this schema-building process computationally, ensuring every node encountered is anchored to the prior knowledge it requires, mirroring the associative architecture through which long-term memory is constructed.

Cognitive load theory (Sweller, 1988; Sweller, van Merriënboer and Paas, 1998; 2019) identifies extraneous load — unnecessary cognitive effort from poor design — as the primary obstacle to learning. Linear documents impose extraneous load by forcing employees to infer relevance and sequencing; meta-books eliminate it through modularity, depth-layer filtering, and prerequisite ordering, redirecting capacity toward the germane processing that produces genuine knowledge acquisition.

Dual coding theory (Paivio, 1986) demonstrates that information encoded through both verbal and non-verbal channels produces more durable retention than single-channel encoding. The meta-book's illustrates edge — connecting abstract concept nodes to concrete example nodes — encodes this dual-coding relationship structurally, ensuring AI traversal automatically implements the bimodal encoding strategy that memory science identifies as optimal.

The testing effect (Roediger and Karpicke, 2006; Adesope, Trevisan and Sundararajan, 2017) is among the most replicated findings in cognitive psychology: active retrieval from memory strengthens retention far more than passive re-reading. Meta-book controversy nodes and reflective insights create structural conditions for active retrieval, requiring employees to generate evaluative positions on contested claims rather than passively absorb content.

The spacing effect (Ebbinghaus, 1885; Cepeda et al., 2006) demonstrates that distributed practice produces dramatically superior long-term retention over massed learning. Cepeda et al.'s (2006) synthesis of 254 studies confirmed that spacing exposure across multiple sessions counteracts the exponential decay of the forgetting curve most effectively. The meta-book depth layer architecture — Core nodes during onboarding, Applied during early practice, Integrative as expertise develops — implements optimal spacing by design, ensuring progressive access to complex nodes only after prior encodings have consolidated, rather than delivering the entire knowledge base simultaneously.

The generation effect (Slamecka and Graf, 1978) demonstrates that self-generated knowledge is encoded more durably than passively received information, because generation requires deeper elaborative processing. Meta-book controversy nodes operationalise this by presenting contested organisational claims without resolution, requiring employees to generate evaluative positions — transforming knowledge delivery from a broadcast model into a generative cognitive act.

The zone of proximal development (Vygotsky, 1978) defines the cognitive space between independent competence and scaffolded capability. Ausubel's (1968) advance organiser principle establishes that new knowledge must be anchored to existing cognitive structures. Meta-book prerequisite links function as a dynamic advance organiser; the depth layer system implements the ZPD computationally — Core nodes representing current competence, Applied and Integrative nodes scaffolding progressive development — with AI traversal acting as the scaffolding expert.

Connectionist neuroscience (Rumelhart and McClelland, 1986; Dehaene, 2020) understands knowledge as distributed patterns of neural activation and learning as the strengthening of weighted connections. Dehaene (2020) argues the brain's knowledge architecture is fundamentally graph-like rather than linear — making meta-books structurally consistent with the brain's own cognitive organisation, suggesting why graph-structured delivery may produce superior learning outcomes to linear text, though this alignment remains a theoretical argument rather than an empirically demonstrated structural match.

Together, these eight traditions form a unified argument: linear document delivery is cognitively suboptimal by design — imposing extraneous load, suppressing schema construction, preventing retrieval practice, and misaligning with the brain's connectionist architecture. Meta-books are the structural correction: simultaneously AI-traversable and cognitively aligned.

8. Limitations and Future Research

Several limitations warrant acknowledgement. Meta-book construction requires significant initial investment: decomposing documents into typed nodes and annotating metadata is labour-intensive, and the combination of domain expertise and knowledge architecture competence required may not be widely available. LLM-assisted semi-automated decomposition is a promising mitigation, with current models demonstrating sufficient comprehension to support node extraction with human review. The five node types and edge categories are not exhaustive — clinical, legal, and creative domains may require alternative taxonomies, and the adequacy of the proposed typology across diverse knowledge cultures remains an open question. These limitations frame an active research agenda without diminishing the conceptual contribution of the framework, which stands on its theoretical grounding and design argumentation independently of empirical deployment. The most pressing research need is empirical evaluation of AI reasoning over meta-book graphs compared to RAG in live settings — the test that would move meta-books from theoretical proposition to evidenced practice. Further priorities include: development of semi-automated construction tools; node type adequacy across diverse domains; longitudinal study of learning outcomes; and governance frameworks addressing data ownership of AI-generated graph extensions, epistemological bias in node typing and classification, equity of access across roles, and knowledge decay — challenges Mittelstadt et al. (2016) identify as inherent to any algorithmic system embedding normative assumptions in its design. Addressing these governance questions is the research that should accompany meta-book adoption rather than a precondition for it.

9. Conclusion

This paper has argued that the structural incompatibility between conventional organisational documents and AI reasoning systems represents a distinct and undertheorised KM problem. Existing approaches — RAG, OWL ontologies, SKOS, LOM, IMS Learning Design, and adaptive hypermedia — each address components of the problem but none simultaneously provides epistemic node typing, dependency graph structure, learner-state-responsive metadata, AI queryability, and cross-source portability. Meta-books address this gap by restructuring source content itself as a transformation of its architecture, not as a layer added to existing documents.

The paper advances three contributions. First, a gap analysis mapped across five architectural requirements and six prior standards establishes the novelty of the meta-book intersection. Second, the meta-book framework is introduced as a design artefact at the construct and model level, with five design principles derived through design argumentation. Third, the framework is grounded in eight converging traditions from learning science and cognitive neuroscience, demonstrating that meta-books are simultaneously AI-traversable and cognitively aligned — correcting the fundamental misalignment between linear knowledge delivery and the graph-structured architecture of the human brain. As agentic AI becomes embedded in organisational workflows and educational systems, the quality of the knowledge structures those systems can reason over will increasingly determine the intelligence they can deliver. What makes meta-books more than an AI infrastructure decision is the alignment of their design with the cognitive architecture of the human brain — the same structure that AI systems must ultimately serve. Meta-books offer a design response grounded in established theory, derivable from existing content, and applicable across every domain where knowledge must be transferred from those who hold it to those who need it.

Ethics Declaration: This paper is a conceptual and design science contribution. No primary data were collected and no human participants were involved. Ethical clearance was not required.

AI Declaration: Claude (Anthropic) was used during the preparation of this paper as a writing, editorial, and research development assistant. Specific uses included: drafting and iterative refinement of all sections; reference verification and correction; acronym and terminology auditing; and document formatting to conference specification. The theoretical framework, conceptual architecture of meta-books, the five-requirement gap analysis, the identification of eight converging neuroscience traditions, all design arguments, intellectual claims, and decisions about content and structure are the author's own. The author directed all stages of development and takes full responsibility for the paper as submitted.

References

- Adesope, O.O., Trevisan, D.A. and Sundararajan, N. (2017) "Rethinking the Use of Tests: A Meta-Analysis of Practice Testing", *Review of Educational Research*, Vol 87, No. 3, pp 659–701, <https://doi.org/10.3102/0034654316689306>.
- Alavi, M. and Leidner, D.E. (2001) "Knowledge management and knowledge management systems: Conceptual foundations and research issues", *MIS Quarterly*, Vol 25, No. 1, pp 107–136, <https://doi.org/10.2307/3250961>.
- Ausubel, D.P. (1968) *Educational Psychology: A Cognitive View*, Holt, Rinehart and Winston, New York.
- Bartlett, F.C. (1932) *Remembering: A Study in Experimental and Social Psychology*, Cambridge University Press, Cambridge.
- Brusilovsky, P. (2001) "Adaptive hypermedia", *User Modelling and User-Adapted Interaction*, Vol 11, No. 1–2, pp 87–110, <https://doi.org/10.1023/A:1011143116306>.
- Cepeda, N.J., Pashler, H. and Vul, E. et al. (2006) "Distributed practice in verbal recall tasks: A review and quantitative synthesis", *Psychological Bulletin*, Vol 132, No. 3, pp 354–380, <https://doi.org/10.1037/0033-2909.132.3.354>.
- Cohen, W.M. and Levinthal, D.A. (1990) "Absorptive capacity: A new perspective on learning and innovation", *Administrative Science Quarterly*, Vol 35, No. 1, pp 128–152, <https://doi.org/10.2307/2393553>.
- Conlan, O., Wade, V. and Bruen, C. et al. (2002) "Multi-model, metadata driven approach to adaptive hypermedia services for personalized eLearning", in De Bra, P., Brusilovsky, P. and Conejo, R. (Eds.), *Adaptive Hypermedia and Adaptive Web-Based Systems*, Lecture Notes in Computer Science, Vol 2347, Springer, Berlin and Heidelberg, pp 100–111.
- Davenport, T.H. and Prusak, L. (1998) *Working Knowledge: How Organisations Manage What They Know*, Harvard Business School Press, Boston.
- Dehaene, S. (2020) *How We Learn: The New Science of Education and the Brain*, Allen Lane, London.
- Ebbinghaus, H. (1885) *Über das Gedächtnis: Untersuchungen zur experimentellen Psychologie*, Duncker and Humblot, Leipzig.
- Gao, Y., Xiong, Y. and Gao, X. et al. (2024) "Retrieval-augmented generation for large language models: A survey", *arXiv preprint*, arXiv:2312.10997, available at <https://arxiv.org/abs/2312.10997>.
- Haenlein, M. and Kaplan, A. (2019) "A brief history of artificial intelligence: On the past, present, and future of artificial intelligence", *California Management Review*, Vol 61, No. 4, pp 5–14, <https://doi.org/10.1177/0008125619864925>.
- Hevner, A., March, S. and Park, J. et al. (2004) "Design science in information systems research", *MIS Quarterly*, Vol 28, No. 1, pp 75–105, <https://doi.org/10.2307/25148625>.
- Hogan, A., Blomqvist, E. and Cochez, M. et al. (2021) "Knowledge graphs", *ACM Computing Surveys*, Vol 54, No. 4, pp 1–37, <https://doi.org/10.1145/3447772>.
- IEEE LTSC (2002) *Draft Standard for Learning Object Metadata (IEEE 1484.12.1)*, IEEE, New York.
- IMS Global Learning Consortium (2003) *IMS Learning Design Information Model, Version 1.0 Final Specification*, IMS Global Learning Consortium, available at <http://www.imsglobal.org/learningdesign/index.html>.
- Lewis, P., Perez, E. and Piktus, A. et al. (2020) "Retrieval-augmented generation for knowledge-intensive NLP tasks", *Advances in Neural Information Processing Systems*, Vol 33, pp 9459–9474.
- Miles, A. and Bechhofer, S. (Eds.) (2009) *SKOS Simple Knowledge Organization System Reference*, W3C Recommendation, World Wide Web Consortium, available at <https://www.w3.org/TR/skos-reference>.
- Mittelstadt, B.D., Allo, P. and Taddeo, M. et al. (2016) "The ethics of algorithms: Mapping the debate", *Big Data and Society*, Vol 3, No. 2, <https://doi.org/10.1177/2053951716679679>.
- Nonaka, I. and Takeuchi, H. (1995) *The Knowledge-Creating Company*, Oxford University Press, New York.
- Paivio, A. (1986) *Mental Representations: A Dual Coding Approach*, Oxford University Press, New York.
- Pan, S., Luo, L. and Wang, Y. et al. (2024) "Unifying large language models and knowledge graphs: A roadmap", *IEEE Transactions on Knowledge and Data Engineering*, Vol 36, No. 7, pp 3580–3599, <https://doi.org/10.1109/TKDE.2024.3352100>.
- Peffer, K., Tuunanen, T. and Rothenberger, M. et al. (2007) "A design science research methodology for information systems research", *Journal of Management Information Systems*, Vol 24, No. 3, pp 45–77, <https://doi.org/10.2753/MIS0742-1222240302>.
- Roediger, H.L. and Karpicke, J.D. (2006) "Test-enhanced learning: Taking memory tests improves long-term retention", *Psychological Science*, Vol 17, No. 3, pp 249–255, <https://doi.org/10.1111/j.1467-9280.2006.01693.x>.
- Rumelhart, D.E. (1980) "Schemata: The building blocks of cognition", in Spiro, R.J., Bruce, B.C. and Brewer, W.F. (Eds.), *Theoretical Issues in Reading Comprehension*, Erlbaum, Hillsdale, NJ, pp 33–58.
- Rumelhart, D.E. and McClelland, J.L. (1986) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol 1, MIT Press, Cambridge, MA.
- Senge, P.M. (1990) *The Fifth Discipline: The Art and Practice of the Learning Organisation*, Doubleday, New York.
- Slamecka, N.J. and Graf, P. (1978) "The generation effect: Delineation of a phenomenon", *Journal of Experimental Psychology: Human Learning and Memory*, Vol 4, No. 6, pp 592–604, <https://doi.org/10.1037/0278-7393.4.6.592>.
- Sweller, J. (1988) "Cognitive load during problem solving: Effects on learning", *Cognitive Science*, Vol 12, No. 2, pp 257–285, https://doi.org/10.1207/s15516709cog1202_4.
- Sweller, J., van Merriënboer, J.J.G. and Paas, F. (1998) "Cognitive architecture and instructional design", *Educational Psychology Review*, Vol 10, pp 251–296, <https://doi.org/10.1023/A:1022193728205>.
- Sweller, J., van Merriënboer, J.J.G. and Paas, F. (2019) "Cognitive architecture and instructional design: 20 years later", *Educational Psychology Review*, Vol 31, No. 2, pp 261–292, <https://doi.org/10.1007/s10648-019-09465-5>.

- Vygotsky, L.S. (1978) *Mind in Society: The Development of Higher Psychological Processes*, Harvard University Press, Cambridge, MA.
- Wenger, E. (1998) *Communities of Practice: Learning, Meaning and Identity*, Cambridge University Press, Cambridge.