

A Taxonomy and Reference Architecture for AI-enabled Industrial Knowledge Management

Mohamed Amine Guedria and Maximilian Dommermuth

Institute of Human and Industrial Engineering (ifab), Department of Mechanical Engineering,
Karlsruhe Institute of Technology (KIT), Germany

mohamed.guedria@student.kit.edu

dommermuth@kit.edu

Abstract: Much of the knowledge that keeps an industrial operation running is never written down: it sits with individual employees, scattered across incompatible systems, and cannot be found in time—a risk that becomes acute whenever people change roles or retire. Artificial intelligence (AI) is widely proposed as a remedy, but existing approaches concentrate on documentary, white-collar work; embodied, blue-collar work is comparatively underserved, particularly by large language models with no native grounding in physical activity. Existing work also treats AI as a single undifferentiated capability, leaving practitioners without a principled basis for choosing among technologies or integrating their outputs. This paper proposes a taxonomy of five AI paradigms (perceptive, dialogic, interpretive, structural and contextual), each defined by its contribution to one of three knowledge processes (capture, structuring, transfer) and by the type of work it serves, with a boundary marked where collective tacit knowledge resists codification. A separate orchestration layer, realised by autonomous agents and enabling technologies (knowledge graphs, retrieval-augmented generation, augmented reality), connects the paradigms into a pipeline. Applied diagnostically to three tools in one manufacturer’s training programme, an assembly-guidance system, a structured interview system, and a RAG-based conversational assistant, the taxonomy shows all three occupy the capture or transfer columns while structuring goes unserved, leaving each tool’s knowledge inaccessible to the others. A pump-assembly scenario shows how an agent-orchestrated pipeline over a shared knowledge graph, with human validation, could unify their outputs. The tools’ reported gains, a 29 per cent onboarding-time reduction and an over 90 per cent retrieval-time reduction for 3,000+ daily users, are taken at face value; whether integration compounds them, and for whom, is a working hypothesis, not a demonstrated result. The paper concludes with a staged evaluation strategy measuring cross-tool retrieval coverage and validation throughput to isolate the structural layer’s impact.

Keywords: Knowledge management, Knowledge transfer, Human-centred AI, Industrial AI, Reference architecture

1. Introduction

Industrial organisations depend on knowledge that is unevenly distributed and difficult to retain. A large share is never documented: it resides with individual employees and cannot be found quickly enough when needed. The risk spikes whenever a role changes hands or an employee retires, acute enough that ISO 9001:2015 obliges organisations to identify and retain the knowledge needed to operate their processes (Sumbal et al., 2017), while in German manufacturing the demographic transition has moved the problem from forecast to present concern (Ottersböck, Dander and Stowasser, 2026). Production-worker and manager departures damage performance through different mechanisms (Eckardt, Skaggs and Youndt, 2014); retention must therefore attend to both the type of work and the specific role.

A further share is nominally documented (work instructions, procedures, ticket systems) but scattered across incompatible systems, so locating the right version in time is itself the problem, and outdated content propagates into errors, the administrative counterpart of the searching-and-waiting waste that lean production systems target (Verein Deutscher Ingenieure, 2012). Capturing knowledge before it is lost, structuring it so it remains usable, and delivering it to those who need it are central industrial challenges; retention practice nonetheless remains inconsistent (Sumbal et al., 2017).

AI is widely proposed to address these tasks (Jarrahi et al., 2023; He and Burger-Helmchen, 2025; Olan et al., 2022; Dommermuth, 2026). Two difficulties follow. First, a coverage bias: approaches concentrate on indirect, white-collar work—not because manual work lacks judgement (assembly and maintenance require diagnosis and improvisation) but because the dominant paradigm of the last decade, the LLM, has no native grounding in physical activity. Purpose-built worker-guidance and AR systems do exist, but even where deployed their outputs are rarely organised into a form other tools can reuse. Second, an absence of shared classification: each technology is described in its own terms, with no basis for deciding which fits which task or worker. Recent skill-management work confirms that integrating individually successful AI pilots into a unified, agent-orchestrated platform with worker-differentiated delivery remains an open research problem (Dommermuth, 2026).

This paper contributes: a taxonomy of five AI paradigms with a separate orchestration layer distinguishing AI from the enabling technologies it employs (Section 3); a diagnostic application to three deployed tools that

identifies a systematic gap in structuring (Section 4); and an agent-orchestrated reference architecture resolving the gap by design (Section 5), while stating plainly that its integration benefit is a hypothesis to be tested, not an established finding (Section 6).

2. Background and Related Work

2.1 The Knowledge-loss Problem, and its Limits

Knowledge loss through staff departures is among the best-documented risks in the literature: the knowledge most at risk is experiential and tacit, accumulated over years and rarely written down, and organisations typically begin retention too late (Sumbal et al., 2017). Eckardt, Skaggs and Youndt (2014) show production-worker and manager turnover damage performance through distinct mechanisms; exposure is smallest where standardised work is strongest, since well-maintained work plans already absorb part of the loss.

A further complication is currency: workarounds become obsolete as processes or equipment change, so treating all captured knowledge as equally durable risks preserving advice that is no longer correct—a boundary condition on the capture mechanism of Section 5. Two anchors follow: support must span capture, structuring and transfer before holders depart, and the relevant unit is not just work type but the specific role within it.

2.2 Knowledge Types and the Limits of Codification

This paper collapses Nonaka and Takeuchi's (1995) SECI processes into three practitioners recognise: capture, structuring, transfer. Collins (2010) distinguishes somatic tacit knowledge (embodied, not reachable through text or motion-time notation), relational tacit knowledge (articulable but usually left unarticulated), and collective tacit knowledge (CTK, acquired only through participation in a social milieu). Sanzogni, Guzman and Busch (2017) argue CTK cannot be codified by any machine; treating it as capturable is a category mistake that undermined knowledge-management systems since the expert-systems era. AI therefore addresses explicit, externalisable knowledge; the collective residue remains the preserve of human, social transfer—a hard constraint in the taxonomy of Section 3, not a caveat added afterward.

2.3 AI in Knowledge Management: The Literature and the gap

Jarrahi et al. (2023) frame AI-in-KM as human–AI symbiosis and identify “know-who” as a contribution beyond direct codification; He and Burger-Helmchen (2025) stress that socialisation stays a human, community-situated process; Olan et al. (2022) show performance gains only when AI is combined with human knowledge sharing; Haefner et al. (2021) show the value of structured classification in innovation management. Dommermuth (2026) surveys multiple production units and reports capture and provision as by far the most frequently cited gaps, not structuring—the direct empirical anchor for the tool selection in Section 4. None of this literature supplies the bridge, at the level of Skill Management, between where a knowledge holder sits and where capture should be targeted.

3. Taxonomy: AI Paradigms, Enabling Technologies and the Orchestration Layer

Industrial knowledge management often groups machine learning, LLMs, knowledge graphs, RAG and AR under the broad label of “AI”. This paper distinguishes AI paradigms, which learn from data, recognise patterns or generate outputs (computer vision, deep learning, LLMs, vision-language models), from enabling technologies, which operationalise them without themselves learning or generating: knowledge graphs provide structured knowledge, RAG grounds generative models in external sources, and AR delivers outputs to users. The distinction underpins the taxonomy, while acknowledging these technologies may be considered AI in a broader historical sense; all three are integral to the architecture in Section 5.

3.1 How the Taxonomy was Built

The taxonomy follows the iterative development method of Nickerson, Varshney and Muntermann (2013). The meta-characteristic was fixed first: the functional contribution a paradigm makes to the industrial knowledge flow. The object set comprised the three deployed tools of Section 4, decomposed into their constituent AI functions (detection, step-state tracking, interviewing, transcription, annotation, triple extraction, grounded retrieval and generation), plus the technologies named in the gap survey of Dommermuth (2026). An empirical-to-conceptual iteration grouped these functions by shared characteristics into candidate paradigms; a conceptual-to-empirical iteration checked the candidates against the two dimensions of Section 3.2, which are exhaustive by construction: any industrial knowledge activity captures, structures or transfers knowledge, for either embodied or documentary work. Development ended when every object was classified, a further pass

added no new characteristic, and no paradigm was redundant—Nickerson et al.'s objective ending conditions. Theofanos, Choong and Jensen (2024) apply the same classification-by-function principle. The five are not claimed as exhaustive, only as the smallest set covering every populated cell in this case material.

Because the paradigms were partly abstracted from the same three tools that Section 4 diagnoses, an empty cell recovered from that diagnosis is not, by itself, an independent discovery (Section 4.1). No independent study has tested these boundaries beyond the case material; validating the taxonomy against independently sampled deployments is future work (Section 6.3).

3.2 The two Classification Dimensions

The taxonomy classifies AI paradigms along two dimensions. The first is the knowledge process served: capture (acquiring knowledge from holders or physical activity), structuring (organising captured material into reusable form), or transfer (delivering knowledge when and where needed). The second is the type of industrial work: predominantly manual, embodied (blue-collar) work, as distinct from predominantly cognitive, documentary (white-collar) work.

Worker type is a coarse proxy for knowledge type (Collins, 2010): blue-collar work is dominated by somatic knowledge best surfaced by observing activity; white-collar work by relational and documentary knowledge best surfaced through dialogue and documents. It says nothing about role-level differences within a type (a line operator and a shift lead are both “blue-collar” but hold different knowledge) and should not imply manual work involves no judgement. Both types contain collective tacit knowledge no paradigm can codify.

3.3 Five AI Paradigms

Each paradigm is defined by its functional contribution to the knowledge flow, not by a specific product, vendor or model family.

Perceptive AI captures embodied knowledge by observing physical activity: computer vision and deep learning detect objects, recognise actions and track step states, yielding raw artefacts such as video segments and event logs.

Dialogic AI captures verbalisable but undocumented knowledge through structured conversation and, redeployed as a conversational assistant, transfers knowledge back through natural language.

Interpretive AI bridges raw multimodal captures and usable structure, aligning video, transcripts, audio and sensor logs into semantically annotated form, occupying the boundary between capture and structuring.

Structural AI organises annotated knowledge into a relational, navigable representation: knowledge-graph construction and ontology alignment turn heterogeneous sources into one queryable body.

Contextual AI transfers knowledge by grounding generated answers in that structured base; trustworthiness at the point of transfer is consequently a function of Structural AI upstream, not of the language model alone.

3.4 The Orchestration Layer

A sixth role, orchestration, functions distinctly from the five. AI agents act as orchestrators: chaining paradigms, routing requests, monitoring confidence and triggering human review when uncertain. Because an agent coordinates rather than itself perceiving, conversing, interpreting, structuring or delivering knowledge, it is not a knowledge paradigm and is therefore depicted separately in Table 1, spanning every process column.

3.5 The Capability map

Table 1 positions the five paradigms and the orchestration layer against the two dimensions of Section 3.2.

Table 1: Five AI paradigms plus the orchestration layer mapped to three knowledge processes and two worker types. Em-dashes mark inactive pairs; enabling technologies appear in italics. [B] = blue-collar; [W] = white-collar

Role	Capture	Structure	Transfer
Perceptive AI	Object detection, step-state tracking → raw event log [B]	—	—
Dialogic AI	Expert interview, audio capture [B,W]	—	Conversational Q&A, LLM guidance [B,W]

Role	Capture	Structure	Transfer
Interpretive AI	Multimodal annotation, transcription [B,W]	Cross-modal indexing → annotated records [W]	—
Structural AI	—	KG construction, ontology alignment → <i>knowledge graph</i> [W]	—
Contextual AI	—	—	RAG-grounded answers, AR-delivered guidance [B,W]
Orchestration	Monitors capture sources	Triggers structuring; routes low-confidence items to review	Selects delivery channel (text, audio, AR)

Boundary: collective tacit knowledge is AI-resistant in both worker types; the collective residue requires human-to-human transfer (Collins, 2010; Sanzogni, Guzman and Busch, 2017).

3.6 Reading the map

The taxonomy situates widely discussed technologies within one structure: RAG is a mechanism through which Contextual AI operates, conditional on Structural AI upstream; AR is a delivery channel; knowledge graphs are the output of Structural AI and the input to Contextual AI; agents are connective tissue, not a sixth kind of knower.

Table 1 also exposes a dependency chain: Perceptive and Dialogic AI produce raw captures → Interpretive AI annotates → Structural AI builds the knowledge graph → Contextual AI retrieves → orchestration sustains confidence and routes exceptions to human review. Breaking any link, particularly Structural AI, disconnects capture from transfer—which Section 4 shows is what happens today across three tools built by the same organisation.

4. Diagnosis: Applying the Taxonomy to three Deployed Tools

4.1 Why These Three Tools, and a Caveat on Independence

The three tools are not presented as the best instances of their category; they are an existing empirical case flagged by a prior, independent gap analysis as addressing the two most frequently reported problems. Dommermuth (2026) found that, across production units of a large manufacturer, capture and provision were named as deficient far more often than structuring. The three tools below are that organisation's deployed answers to those gaps, each built for one problem and one user group rather than for interoperability—precisely the pattern this paper diagnoses.

A methodological caveat belongs here. Since the five paradigms were abstracted from these same three tools (Section 3.1), finding the Structure column empty when the taxonomy is applied back onto them is not an independent discovery. What licenses the diagnosis is the prior, independent gap analysis (Dommermuth, 2026), which identified capture and provision as the dominant gaps before the taxonomy existed, from practitioner interviews. The taxonomy's role is to show why, structurally, that gap recurs across independently commissioned tools, not to establish that it exists (Section 6.3).

The diagnosis is a single-organisation, embedded case design: each tool is a unit of analysis, classified against Table 1 using system documentation, deployment reports and the prior gap survey as data sources. Because the three are a diagnostic instance rather than a recommendation, the argument does not depend on their being optimal; a different organisation might select different tools, and the taxonomy is meant to remain applicable there. What generalises is the diagnostic pattern, not the tool names. Table 2 positions each on the capability map.

The AR-guided assembly-guidance system is a computer-vision and augmented-reality system at the assembly workstation. Its Perceptive AI component detects parts and tools as oriented bounding boxes in a workstation-fixed frame and drives a per-step state machine; the detector is fine-tuned on assembly-specific imagery, since generic detectors do not reliably distinguish similar fasteners under shop-floor lighting. This yields a structured event log and drives in-situ AR guidance directly, since no shared structured base exists. A single-area pilot reported a 29 per cent reduction in onboarding time (Dommermuth, 2026)—an indicative one-workstation result whose sample size and comparison group are unavailable (Section 6.3). Records are produced every cycle,

but without a Structural AI step recurring error modes remain in the raw log. (Figure 1's "AR overlay" channel refers to the proposed architecture of Section 5, not this tool's current operation.)

The structured interview system is a dialogic capture tool eliciting expert knowledge through conversations with a human interviewer or a protocol-following AI avatar, over 900 questions designed to surface reasoning and exceptions (Dommermuth, 2026). Each interview yields a transcript and person–concept–step triples, but the system stops at capture, so the record waits to be found. No throughput or completion figures are available; it is reported as deployed but not yet evaluated.

The RAG-based conversational assistant is a multi-agent system returning grounded, source-referenced answers with a trustworthiness indicator to both worker types, reporting over 90 per cent reduction in typical retrieval time across more than 3,000 daily users (Dommermuth, Riesch and Guedria, 2025; Dommermuth, 2026). This organisation-wide aggregate, with no variance by role or query type, is a ceiling on what already-indexed content delivers, not a role-disaggregated effect size. It transfers only what it can reach: no shared structured base connects the three tools.

Table 2: Three deployed tools positioned on the capability map.

Tool	Capture	Structure	Transfer	Result	Evaluation basis
AR-guided assembly system (blue-collar)	Vision, step-state tracking	—	In-situ AR guidance (state machine, not graph retrieval)	29% onboarding time reduction	Single-area pilot; N and comparison group unreported
Structured interview system (white-collar)	AI-avatar interviews, 900+ items	—	—	Deployed; not yet evaluated	No quantitative data available
RAG-based assistant (blue- and white-collar)	—	—	RAG retrieval, trust indicator	>90% retrieval time reduction	Org-wide, 3,000+ users; period and variance unreported

Diagnosis: all three tools occupy the Capture or Transfer columns; the Structure column is empty in every case—corroborated by, not independently re-derived from, the survey of Dommermuth (2026); see Section 4.1.

4.2 A Systematic gap, and What it Costs Concretely

The pattern is systematic: three independently commissioned tools all stopped short of structuring, which the prior gap analysis suggests is not unique to them (Dommermuth, 2026). Structuring has no payoff visible to a point-solution project—it neither captures from a holder nor delivers at the moment of need—so it is the first thing cut from scope.

The cost is concrete. An operator's error, once resolved through the AR overlay, exists only in that system's log; a similar error elsewhere cannot benefit, because the conversational assistant has no route to it. Symmetrically, a senior planner's interview may contain the exact nuance a worker is asking about, and the assistant will not find it, since it was never indexed anywhere it looks. The missing structuring step and the inter-tool silo problem are the same fact.

4.3 A Concrete Scenario: Pump-assembly Error Capture, Validation and Reuse

The scenario below is a composite of workflows already in place; the integration layer of Section 5 does not yet exist. It builds on Deviation Management, already present in most assembly operations: the missing piece is a systematic route from the record to knowledge that prevents recurrence.

A less-experienced operator assembles a hydraulic pump at a workstation equipped with the assembly-guidance system. At step 7 (torquing the end-cap bolts) it detects a deviation and raises an AR alert; the operator corrects the sequence and the alert clears. This event—exactly what Deviation Management would already log—is step one.

The proposed pipeline additionally captures step two—how and why it was resolved—via a short hands-free audio comment through the AR headset: "What made this step difficult, and what helped you fix it?" Interpretive AI transcribes the note and appends it to the record with the AR frames; a process expert then assesses whether the fix is correct and generalisable. If accepted, Structural AI ingests it into the shared knowledge graph as a node linked to the assembly step, pump model, remediation sequence and commentary.

When the same deviation recurs elsewhere, or a new operator trains on this assembly, the conversational assistant can retrieve the validated fix and the AR overlay can display it in situ—knowledge previously held only by those who had seen the error before.

5. An Agent-orchestrated Reference Architecture

5.1 From Diagnosis to Design

The diagnosis motivates the architecture directly: fill the Structure column, connect the tools through what fills it, and govern the quality of what enters it. Three failure modes traceable to the empty Structure column are resolved by design: inter-tool silos (knowledge captured by one tool is unreachable by another), ungrounded transfer (Contextual AI without Structural AI upstream must rely on curated content or risk unsupported output), and absent expertise metadata (no tool can answer who knows what organisation-wide).

None of this is new: knowledge graphs, ontologies and expert systems targeted the same problem two decades ago and largely failed, chiefly because manual ontology curation scaled linearly with domain size and no natural-language interface let non-specialists query the result. What has changed is the cost: LLM-based agents now propose graph updates from unstructured material at a fraction of manual-curation cost (Zhu et al., 2024), and RAG lets the graph be queried in natural language.

RAMI 4.0 (Deutsches Institut für Normung, 2016) organises Industry 4.0 assets along hierarchy, life-cycle and layer axes, but treats its information and functional layers generically, without distinguishing knowledge processes, worker types or a codification boundary; the present architecture is best read as a knowledge-specific refinement of those layers, not a competing model. Equally, the design is deliberately socio-technical (Baxter and Sommerville, 2011): the validation gate, the CTK boundary and the retention of accountability with human experts treat the social subsystem as a design object, not an afterthought.

5.2 Architecture

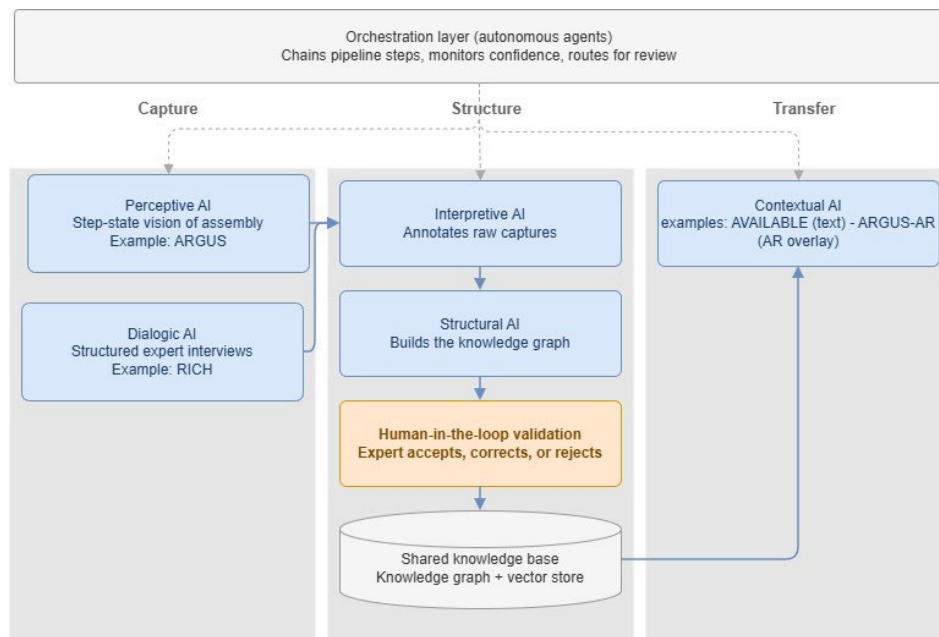


Figure 1: Reference architecture. Capture: the assembly-guidance system (Perceptive AI) and the interview system / operator voice note (Dialogic AI) generate raw records. Structure: Interpretive AI annotates; Structural AI builds the knowledge graph; a human-in-the-loop gate routes low-confidence proposals to expert review. Transfer: the conversational assistant (text) and the AR overlay (in-situ) retrieve from the shared base. Connection to systems of record (e.g. SAP PM) is a required integration outside the pipeline shown (Section 6.3). The figure retains internal project codenames (ARGUS, RICH, AVAILABLE), which map to the assembly-guidance system, the structured interview system and the RAG-based conversational assistant respectively

A knowledge graph with a vector store is the single source of truth: Interpretive and Structural AI write to it, Contextual AI reads from it (Figure 1). A pump-assembly error record and a senior engineer's interview record land in the same graph with provenance intact and both become retrievable by the conversational assistant,

whose trustworthiness indicator (Dommermuth, Riesch and Guedria, 2025) is thereby grounded in structured evidence rather than model weights alone. In production it must also reconcile its entities with enterprise systems of record such as SAP Plant Maintenance work orders and formal work plans—a requirement, not a solved problem (Section 6.3).

LLM-based agents extract entities and relations from heterogeneous material (event logs, video frames, transcripts, voice notes), propose graph updates with confidence scores, and align new material with the existing schema—the pattern whose feasibility Zhu et al. (2024) demonstrate. Because multimodal extraction adds latency and per-event agents multiply compute cost, the design treats structuring as an asynchronous step queued after the event closes rather than blocking the AR guidance loop; sizing that queue, like prompts and routing rules, is left to implementation.

A primary deployment challenge is the prevalence of unverified, outdated or conflicting information—“knowledge junk”—whose direct ingestion into a RAG system leads to hallucinations and incorrect instructions reaching shop-floor employees (Dommermuth, 2026). The human-in-the-loop gate of Figure 1 counters this: agent-proposed extractions below a confidence threshold go to a domain expert, who sees the proposal alongside source evidence and can accept, correct or reject it; accepted items enter the shared base, and corrections improve agent calibration. Operationally, the gate prevents a wrongly documented step from being retrieved as guidance by every subsequent operator; organisationally, it keeps accountability with humans, so AI augments judgement rather than replacing it (Jarrahi et al., 2023; Olan et al., 2022). This three-stage pattern—capture, expert validation, single curated base—is already implicit in the tools’ workflows; the architecture makes it explicit.

The shared base dissolves the silos, grounds transfer in curated content, and functions as cross-tool expertise metadata supporting “know-who” queries (Jarrahi et al., 2023). Two boundaries remain: it cannot codify CTK (Section 2.2), though the graph can identify which human holds it; and nothing substitutes for keeping work plans current—where a documentation problem is really a maintenance problem, the fix is to update the work plan.

6. Discussion, Limitations and Next Steps

6.1 Contributions and their Value

The taxonomy offers three kinds of value. Descriptively, it separates AI paradigms from enabling technologies, preventing category errors that inflate claims. Diagnostically, applied to three independently built tools, it located a systematic blind spot—the unserved structuring step—grounded in a prior, independent gap analysis rather than the three tools alone. Prescriptively, it turns technology selection into a structured decision: identify the work type and knowledge process, locate the paradigm in Table 1, and check the enabling technologies it requires.

6.2 Scope of the Hypothesis

Both headline figures in Table 2 are aggregate averages—from a single-area pilot and an organisation-wide rollout (Dommermuth, Riesch and Guedria, 2025; Dommermuth, 2026); neither reports how the effect varies by knowledge type, tool-use pattern or user segment, so it is unknown whether the benefit accrues mainly to the worker or to those managing onboarding and quality around them. They are the baseline from which the hypothesis is stated, not evidence for the architecture. Its components are grounded in deployment results, the integration mechanism has demonstrated feasibility elsewhere (Zhu et al., 2024), and the failure of the prior generation of knowledge-graph systems (Section 5.1) is a caution built into the design.

6.3 Limitations

The taxonomy is conceptual, derived from the same three tools and one gap survey in a single organisation that Section 4 diagnoses; whether it holds elsewhere is open (Sections 3.1 and 4.1). Worker-type tags are coarse: a role-level analysis distinguishing a line operator from a shift lead could shift which paradigm is most relevant, and is left to future work.

On the architecture side, the validation gate checks correctness at review time only, not durability. The confidence threshold routing items to it needs calibration by assembly type and criticality; access control—who may read or validate what—is not yet addressed; and reconciliation with systems of record such as SAP PM and formal work plans remains a required integration not yet designed.

Tool selection, while grounded in a prior gap analysis, is one defensible choice among several, and the claim that the architecture is agnostic to it is untested against a second tool set. Neither headline figure includes sample size, comparison-group composition or variance by role; both are reported at face value from their sources.

6.4 Proposed Evaluation Strategy and Next Steps

Evaluation must isolate the structuring layer's contribution from the tools' baselines. The proposed design: (1) a staged rollout—a baseline period on the standalone tools, then a treatment period after the shared structuring layer is added, holding the capture tools unchanged, so any measured change is attributable to structuring; (2) cross-tool retrieval coverage as primary KPI—the share of validated fixes captured at one tool that the assistant can retrieve from a different workstation; (3) validation throughput—expert accept/correct/reject rates and time-to-validate at the gate; and (4) disaggregation of every KPI by knowledge type, tool and user role, directly answering who benefits from integration.

A second, upstream direction is the coupling this diagnosis points to but does not test—between Skill Management, which locates knowledge holders and roles needing support, and the capture tools then applied to them. Investigating it across a broader set of tools and organisations would address whether point solutions should keep being built in isolation, consistent with the call for a unified, agent-orchestrated AI platform (Dommermuth, 2026).

Ethics Declaration: Ethical clearance was not required as this research focuses on conceptual development and uses anonymised, aggregated operational data. This independent doctoral research at the Karlsruhe Institute of Technology involves no conflicts of interest and received no employer funding.

AI Declaration: AI was used solely for editorial assistance with proofreading and formatting. All core concepts, architecture, and intellectual property are the work of the human authors, who bear full responsibility for the content.

References

- Baxter, G. and Sommerville, I. (2011) "Socio-technical systems: from design methods to systems engineering", *Interacting with Computers*, Vol 23, No. 1, pp 4–17.
- Collins, H. (2010) *Tacit and Explicit Knowledge*, University of Chicago Press, Chicago.
- Deutsches Institut für Normung (2016) *Referenzarchitekturmodell Industrie 4.0 (RAMI4.0)*, DIN SPEC 91345:2016-04, Beuth Verlag, Berlin.
- Dommermuth, M. (2026) "Building a resilient and competitive workforce: an AI-based skill management framework for adaptive production and work organization", in *IEEE Conference Proceedings*, in press.
- Dommermuth, M., Riesch, L. and Guedria, A. (2025) "Menschzentrierter KI-Chatbot zur individuellen Bereitstellung vertrauenswürdiger Lerninhalte", in *GfA Frühjahrskongress 2025*, Aachen, pp 963–968.
- Eckardt, R., Skaggs, B.C. and Youndt, M. (2014) "Turnover and knowledge loss: an examination of the differential impact of production manager and worker turnover in service and manufacturing firms", *Journal of Management Studies*, Vol 51, No. 7, pp 1025–1057.
- Haefner, N., Wincent, J., Parida, V. and Gassmann, O. (2021) "Artificial intelligence and innovation management: a review, framework, and research agenda", *Technological Forecasting and Social Change*, Vol 169, 120392.
- He, X. and Burger-Helmchen, T. (2025) "Evolving knowledge management: artificial intelligence and the dynamics of social interactions", *IEEE Engineering Management Review*, Vol 53, No. 5, pp 215–231.
- Jarrahi, M.H., Askay, D., Eshraghi, A. and Smith, P. (2023) "Artificial intelligence and knowledge management: a partnership between human and AI", *Business Horizons*, Vol 66, No. 1, pp 87–99.
- Nickerson, R.C., Varshney, U. and Muntermann, J. (2013) "A method for taxonomy development and its application in information systems", *European Journal of Information Systems*, Vol 22, No. 3, pp 336–359.
- Nonaka, I. and Takeuchi, H. (1995) *The Knowledge-Creating Company*, Oxford University Press, New York.
- Olan, F., Arakpogun, E.O., Suklan, J., Nakpodia, F., Damij, N. and Jayawickrama, U. (2022) "Artificial intelligence and knowledge sharing: contributing factors to organizational performance", *Journal of Business Research*, Vol 145, pp 605–615.
- Ottersböck, N., Dander, H. and Stowasser, S. (eds) (2026) *Wissensmanagement mit Künstlicher Intelligenz: Erfahrungswissen effizient sichern und transferieren*, ifaa-Edition, Springer Vieweg, Berlin.
- Sanzogni, L., Guzman, G. and Busch, P. (2017) "Artificial intelligence and knowledge management: questioning the tacit dimension", *Prometheus*, Vol 35, No. 1, pp 37–56.
- Sumbal, M.S., Tsui, E., See-to, E. and Barendrecht, A. (2017) "Knowledge retention and aging workforce in the oil and gas industry: a multi perspective study", *Journal of Knowledge Management*, Vol 21, No. 4, pp 907–924.
- Theofanos, M.F., Choong, Y.-Y. and Jensen, T. (2024) *AI Use Taxonomy: A Human-Centered Approach*, Technical Report NIST AI 200-1, National Institute of Standards and Technology, Gaithersburg, MD.

- Verein Deutscher Ingenieure (2012) *Ganzheitliche Produktionssysteme – Grundlagen, Einführung und Bewertung*, VDI 2870 Blatt 1, Verein Deutscher Ingenieure e.V., Düsseldorf.
- Zhu, Y., Wang, X., Chen, J., Qiao, S., Ou, Y., Yao, Y., Deng, S., Chen, H. and Zhang, N. (2024) "LLMs for knowledge graph construction and reasoning: recent capabilities and future opportunities", *World Wide Web*, Vol 27, No. 5, Article 58.