

Final Project 2

PMG 497

Thinking Systems in Organizations: Frameworks, Architectures, Metrics, Governance, and Cross-Sector Evidence

Hajer aldossari

Project 2

2026 Feb

Organizations “think” through **repeatable, governed configurations of people, processes, and technologies** that convert signals into interpretations, decisions, actions, and learning. In this report, a *thinking system* is treated as an **organizational capability** rather than an individual trait: cognition is **distributed** across teams, routines, artifacts (documents, dashboards, decision logs), and platforms—consistent with distributed cognition and sociotechnical systems theory. [\[1\]](#)

Several mature research streams converge on this view and jointly explain why effective thinking systems produce more reliable outcomes:

Organizational cognition emphasizes *bounded rationality*: decision-makers face attention and computational limits, rely on simplifying heuristics, and “satisfice.” This makes decision quality highly sensitive to information design, routines, and decision rights. [\[2\]](#)

Collective intelligence shows groups can exhibit a measurable general performance factor (“c”) across diverse tasks, driven substantially by **interaction patterns** (e.g., balanced participation) rather than only by individual IQ. [\[3\]](#)

Sensemaking explains how organizations construct shared “maps” of ambiguous environments via retrospective interpretation, social negotiation, and iterative action; plausibility and power dynamics shape which narratives become “reality.” [\[4\]](#)

Learning organizations and organizational learning research emphasize feedback loops that update routines and memory; progress depends on disciplined reflection (single- and double-loop learning) and institutionalization into standards, training, and governance rather than episodic “lessons learned.” [\[5\]](#)

From these foundations, the report compares three architectures of thinking systems:

1) **Human** architectures (judgment-centered; reliant on meetings, narratives, facilitation),

- 2) **Socio-technical** architectures (joint optimization of social and technical subsystems; routinized knowledge/decision workflows), and
- 3) **AI-augmented** architectures (socio-technical systems where AI materially shapes interpretation/option generation/decisions, requiring explicit lifecycle risk governance and monitoring). [\[6\]](#)

A central analytic conclusion is that AI does not automatically improve organizational thinking: meta-analytic evidence indicates human–AI combinations are often **worse than the best of humans or AI alone**, especially for decision tasks, implying that thoughtful role allocation, calibration, and governance gates are design requirements—not optional add-ons. [\[7\]](#)

Implementation success is most strongly associated with treating the thinking system as an engineered management system: **diagnose** → **design** → **pilot** → **scale** → **institutionalize**, supported by layered governance (knowledge management system governance + AI risk governance + model risk governance where relevant). This aligns with recognized standards and guidance: ISO 30401 for KM management systems; NIST AI RMF for AI risk governance; and SR 11-7 for model risk management in regulated, model-driven decision contexts. [\[8\]](#)

Unspecified constraints noted: the user did not specify organization size, industry, regulatory jurisdiction(s), budget, labor arrangements, existing tool stack, data

classification posture, or risk tolerance. These are treated as unspecified; design choices and timelines are therefore presented as cross-sector defaults with explicit assumptions.

Definitions and theoretical frameworks

A working definition of “thinking systems” in organizations

A thinking system in an organization is the **governed socio-technical arrangement** through which the organization:

- detects signals (internal/external), •
- constructs shared meaning (frames, hypotheses, narratives), •
- makes commitments (decisions with explicit tradeoffs), •
- executes coordinated action, and •
- learns by updating memory (routines, standards, knowledge assets) based on feedback. •

This definition treats the unit of cognition as extending beyond individuals—consistent with distributed cognition accounts that conceptualize cognition as spanning people plus artifacts and practices. [\[9\]](#)

Organizational cognition: bounded rationality and routines

Bounded rationality. Herbert A. Simon [\[10\]](#) argued that real-world choice is constrained by limited information, limited computation, and limited time; organizations thus

rationalize by simplifying and satisficing rather than optimizing. [\[11\]](#)

Routines as cognitive scaffolding. Organizational learning and cognition are strongly routine-based: routines stabilize behavior and enable coordination, but are also vehicles for change when enacted and adapted. The ostensive (abstract) and performative (enacted) aspects of routines explain why routines create both stability and flexibility. [\[12\]](#)

Implication for thinking systems: because attention and computation are scarce resources, high-performing thinking systems reduce cognitive load by designing **decision rights, information pathways, and routines** that make “good behavior easy” and “bad behavior hard.” [\[13\]](#)

Collective intelligence: the c-factor and interaction patterns
Work by Anita Woolley [\[14\]](#) and colleagues reported converging evidence for a **general collective intelligence factor (“c”)** that predicts group performance across diverse tasks; measured c is associated with interaction properties such as conversational turn-taking equality and social sensitivity, not simply with the maximum or average intelligence of members. [\[15\]](#)

A later PNAS meta-analytic synthesis across multiple studies provides strong support for a general collective intelligence factor and emphasizes that **collaboration process** is a major predictor of collective intelligence. [\[16\]](#)

Implication for thinking systems: organizations should treat *interaction design* (facilitation norms, participation equality, psychological safety, decision protocols) as a first-class design surface, not merely “culture.” [\[17\]](#)

Sensemaking: shared meaning under uncertainty

Karl E. Weick[\[18\]](#)’s sensemaking tradition emphasizes how groups structure ambiguity to enable action: people build plausible “maps,” test them via action and conversation, and refine or abandon them depending on credibility. [\[19\]](#)

Sensemaking is not purely cognitive; power and politics can shape which interpretations dominate. Research on power in sensemaking underscores how narratives and interpretive frames can become locked-in, excluding alternative readings and suppressing weak signals. [\[20\]](#)

Implication for thinking systems: organizations need explicit practices for (a) surfacing disconfirming evidence, (b) documenting assumptions, and (c) revisiting frames at defined triggers, rather than relying on spontaneous debate. [\[21\]](#)

Learning organizations and organizational learning:
feedback to memory

Peter M. Senge[\[22\]](#) popularized the learning organization as a disciplined system of building shared vision, surfacing mental models, and practicing systems thinking—framing learning as an organizational design problem rather than individual training alone. [\[23\]](#)

Single- and double-loop learning, associated with Chris Argyris[24] and Donald Schön[25], distinguishes between correcting errors within existing governing variables (single-loop) versus questioning and changing those governing variables (double-loop). [26]

In a foundational review of organizational learning, Levitt and March characterize organizational learning as **routine-based, history-dependent, and target-oriented**, highlighting how organizations encode experience into routines and memory structures. [27]

Implication for thinking systems: learning requires mechanisms that (1) generate feedback, (2) convert it into validated knowledge, and (3) reliably embed it into standards, tools, and decision routines so it affects future choices. [28]

Comparative table of frameworks

Framework	Core construct	Unit of analysis	What “good” looks like in practice	Characteristic failure mode	Key primary/peer-reviewed sources
Organizational cognition	bounded rationality; satisficing;	individuals-structure;	decisions are scoped/structured to scarce attention;	local optimization; information overload	Simon (1955) [29]; Levitt & March (1988)

Framework	Core construct	Unit of analysis	What “good” looks like in practice	Characteristic failure mode	Key primary/ peer-reviewed sources
	attention limits	routines	routine design reduces errors	d; procedural drift	[30] ; routines as stability/change [31]
Collective intelligence	general group performance factor (“c”)	group/team	interaction norms produce consistently high performance across tasks	dominance; unequal participation; poor coordination	Woolley et al. (Science, 2010; may be paywalled on Science) [15] ; PNAS meta-analysis [32]
Sensemaking	plausibility-driven	group-in-	assumptions explicit; weak signals	narrative lock-in; power-	Weick (2005) [33] ; power in

Framework	Core construct	Unit of analysis	What “good” looks like in practice	Characteristic failure mode	Key primary/peer-reviewed sources
	map-making	context	tracked; frames revisited	distorted meaning	sensemaking [20]
Learning organization / organizational learning	single- vs double-loop learning; systems thinking	multi-level (team –org)	feedback loops update routines; learning institutionalized	“lessons observed” but not adopted; ritualized review	Senge (1990) [34] ; Argyris (1977) [35] ; Levitt & March review [30]

Architectures of thinking systems

Architectural lens: where cognition is instantiated and governed

Thinking-system architecture can be modeled as a choice about **where and how cognition is distributed**:

Human architecture: cognition primarily in human judgment and conversation; artifacts are supportive but not deeply integrated into workflows. •

Socio-technical architecture: cognition distributed across people plus designed routines, documentation, and platforms; emphasizes joint optimization of social and technical subsystems. [\[36\]](#)

AI-augmented architecture: a socio-technical system in which AI models materially shape interpretation, retrieval, option generation, prioritization, or decision execution; introduces lifecycle risks requiring explicit governance and monitoring (e.g., drift, privacy, automation bias). [\[37\]](#)

This report assumes no specific industry; therefore, architectural “best fit” is presented as conditional by decision criticality, regulation, and scale (all unspecified in the user request).

Comparative table of architectures

Architecture	Primary capabilities	Typical failure modes	Governance emphasis	Best fit (conditional; scale/sector or unspecified)
Human	nuanced judgment; tacit knowledge use; rapid	inconsistent memory; meeting-driven bottlenecks;	decision rights; facilitation; conflict-	early-stage exploration; high ambiguity

Architecture	Primary capabilities	Typical failure modes	Governance emphasis	Best fit (conditional; scale/sector or unspecified)
	reframing; flexible exceptions	bias/politics; weak auditability	of-interest controls; documentation minimal standard	; low repeatability; small-to-moderate scale
Socio-technical	scalable reuse of knowledge; repeatable decision cycles; auditable learning loops; consistent handoffs	“process theater”; metric gaming; tool overload; slow adaptation if routines fossilize	KM system governance; lifecycle for knowledge assets; measurement and continuous improvement	repeated decisions; regulated operations; high coordination complexity; multi-site organizations

Architecture	Primary capabilities	Typical failure modes	Governance emphasis	Best fit (conditional; scale/sector or unspecified)
AI-augmented	faster synthesis/retrieval; pattern detection; decision support at scale; automated monitoring and triage	automation bias/overreliance; opaque errors; data leakage; drift; human–AI mismatch (hybrid underperformance)	AI risk governance; model monitoring; validation; incident response; privacy/security controls	high-volume decisions; complex search; personalization; forecasting—especially where monitoring is feasible

Empirical caution: human–AI combinations frequently underperform the best component (human or AI alone), especially in decision tasks; design must explicitly target complementarity (role allocation, calibration, risk gates).

[\[7\]](#)

Architectural boundary conditions (explicitly unspecified)

The following determinants are unspecified and materially change design choices:

- organization size and geographic distribution, •
- regulatory jurisdiction(s) and whether “high-risk” AI categories apply, •
- data sensitivity/classification and cross-border transfer constraints, •
- existing enterprise platforms and identity/access controls, •
- budget, staffing, and time-to-value expectations, •
- baseline KM maturity and documentation norms. •

These are explicitly treated as unspecified; tables and roadmaps therefore present cross-sector reference patterns rather than organization-specific prescriptions.

[\[38\]](#)

Processes and components

A canonical process model: sense → interpret → decide → act → learn

At a high level, thinking systems can be modeled as a closed-loop cycle that converts signals into learning embedded in memory and routines. This loop integrates sensemaking’s “map-building,” organizational learning’s routine updating, and KM’s capture/curation/access/reuse lifecycle. [\[39\]](#)

flowchart LR

A[Sense: signals, anomalies, weak cues] --> B[Interpret: framing, hypotheses, narratives]

B --> C[Deliberate: tradeoffs, options, commitments]

C --> D[Decide: decision rights exercised; decision log created]

D --> E[Act: coordinated execution in workflow]

E --> F[Observe: outcomes, metrics, incidents, feedback]

F --> G[Learn: after-action review, postmortem, retrospective]

G --> H[Institutionalize: update playbooks, standards, training, knowledge base]

H --> B

Knowledge capture: tacit/explicit conversion and organizational memory

A robust thinking system treats knowledge as both **tacit** (experience, intuition, contextual know-how) and **explicit** (documents, rules, models). Nonaka's theory of organizational knowledge creation formalizes knowledge creation as conversion between tacit and explicit modes and provides a conceptual basis for structuring capture practices (e.g., externalization into decision logs, playbooks, and lessons learned). [\[40\]](#)

Practical capture mechanisms (cross-sector patterns):

Decision logs (what was decided, why, assumptions, • evidence, dissent).

After-action reviews/postmortems (structured • reflection that converts experience into reusable knowledge). [\[41\]](#)

Lessons learned repositories (curated, reviewed • lessons with visibility and reuse pathways). [\[42\]](#)

Communities of practice knowledge sharing • (ongoing tacit exchange that precedes formal codification). [\[43\]](#)

Curation, access, and reuse: from documents to actionable memory

Curation is the bridge between “information stored” and “knowledge reused.” Without curation, knowledge bases degrade into duplication and low-trust content.

Curation includes: taxonomy/ontology, metadata, • review workflows, deprecation policies, and authorship accountability—consistent with KM maturity approaches that treat governance and measurement as core dimensions of program capability. [\[44\]](#)

Access includes: enterprise search and relevance • ranking evaluation; without evaluation, search becomes a perceived “black box” and users revert to informal channels. [\[45\]](#)

Reuse includes: embedding knowledge into • workflows (templates, checklists, default decision pathways), rather than relying on voluntary reading.

Decision rights and transactive memory as system components

Decision quality depends not only on “better data” but on **decision rights** (who decides, who provides input, and escalation logic). In teams, the ability to locate expertise is itself a cognitive asset: transactive memory systems describe “who knows what” and the team’s retrieval processes; validated measures exist for assessing such systems in the field. [\[46\]](#)

Design implication: formalizing “who owns what” (knowledge domains, decisions, models) reduces overload and ambiguity and improves retrieval speed—especially in socio-technical and AI-augmented architectures.

Entities and artifacts in a thinking system (ER model)

erDiagram

```
PERSON ||--o{ CONTRIBUTION : creates
CONTRIBUTION }o--|| KNOWLEDGE_ASSET :
becomes
PLATFORM ||--o{ KNOWLEDGE_ASSET : stores
KNOWLEDGE_ASSET }o--o{ DECISION : informs
MODEL ||--o{ DECISION : recommends
GOVERNANCE_POLICY ||--o{ DECISION : constrains
DECISION ||--o{ ACTION : triggers
ACTION ||--o{ OUTCOME : produces
OUTCOME }o--|| FEEDBACK_EVENT : generates
FEEDBACK_EVENT }o--o{ KNOWLEDGE_ASSET :
updates
```

This model aligns with distributed cognition: cognition is enacted through interactions among people, artifacts, and constraints, not only through individual reasoning. [\[9\]](#)

AI-augmented decision sequence: an accountable workflow

AI augmentation is safest and most effective when embedded into a workflow that enforces **governance gates**, provenance, and monitoring. NIST emphasizes lifecycle risk management and notes AI systems may require more frequent maintenance triggers due to data/model/concept drift and can introduce privacy risks through aggregation. [\[47\]](#)

sequenceDiagram

participant DM as Decision Maker

participant WS as Workflow System

participant KB as Knowledge Base

participant AI as AI Model (RAG/Analytics)

participant GOV as Risk & Governance Gate

DM->>WS: Define decision question + constraints

WS->>KB: Retrieve policies, precedents, evidence

WS->>AI: Generate options + rationale + uncertainty

AI-->>WS: Options + citations/provenance + confidence

notes

WS->>GOV: Gate: privacy, scope, bias, model validity,
logging

GOV-->>WS: Approve / require controls / block

WS-->>DM: Decision brief + tradeoffs + audit trail

DM->>WS: Commit decision; record justification

WS->>KB: Store decision & outcome signals for learning

This design directly addresses documented hybrid failure risks: overreliance (automation bias), miscalibration, and governance gaps. [\[48\]](#)

Metrics and evaluation methods

Why measurement is structurally necessary

Thinking systems compete against two predictable degradations:

- 1) **Memory degradation** (content becomes outdated, duplicated, or distrusted), and
- 2) **Loop integrity failure** (reviews occur but do not update standards/routines, producing repeated failures).

KM maturity frameworks emphasize measurement progression from activity indicators to business-aligned ROI and sustained improvement; without measurement, programs stagnate and become compliance or “portal” theater. [\[49\]](#)

IR metrics for enterprise search and retrieval (precision, recall, nDCG)

Enterprise knowledge retrieval is an information retrieval problem and can be evaluated using established IR methods:

Precision: proportion of retrieved items that are • relevant.

Recall: proportion of relevant items in the collection • that are retrieved.

nDCG (normalized discounted cumulative gain): • evaluates ranked retrieval quality by assigning larger credit to relevant results appearing earlier in the rank list. [\[45\]](#)

The IR evaluation literature also stresses the need for **test collections**, relevance judgments, and empirical evaluation to demonstrate improvements, not only anecdotal user satisfaction. [\[45\]](#)

KM maturity progression: APQC and ISO 30401 alignment

APQC [\[50\]](#)'s KM Capability Assessment Tool describes a multi-level maturity progression for KM programs and highlights dimensions such as governance, processes, technology, and measurement sophistication. [\[49\]](#)

International Organization for Standardization [\[51\]](#)'s ISO 30401:2018 specifies requirements and guidelines for establishing, implementing, maintaining, reviewing, and improving a KM management system and states applicability across organization types and sizes (the standard text is proprietary; the ISO abstract is publicly accessible). [\[52\]](#)

Interpretive alignment: APQC maturity can be used as an assessment and benchmarking instrument, while ISO 30401 provides a formal management-system backbone (plan–do–check–act logic) supporting systematic improvement. [\[44\]](#)

Decision-quality proxies and monitoring metrics

Because “decision quality” is hard to measure directly (counterfactual outcomes are unobserved), well-run thinking systems use **proxy metrics** that are difficult to game and can be tied to organizational outcomes.

Key proxy families include:

Cycle time: time from signal detection to decision and execution; useful where speed matters (incidents, customer escalations). •

Reversal rate: frequency of major decision reversals within a defined window (e.g., 90 days), adjusted for changing conditions. •

Forecast calibration: alignment of predicted probabilities with observed frequencies (e.g., measured via calibration curves or proper scoring rules), useful for risk, demand, and incident prediction contexts. •

Post-event recurrence: repeated incidents or repeated policy exceptions indicating learning-loop failure (used in SRE and safety cultures). [\[53\]](#) •

For AI-augmented systems, monitoring must additionally track:

Data/model/concept drift and triggers for corrective maintenance, •

Bias checks and subgroup performance, •

Privacy/security incidents, and •

Automation bias signals (e.g., override rates, •
reliance patterns). [\[54\]](#)

Comparative table: metric families and how they are used

Metric family	Definition	Applies most strongly to	Example operation alization	Risks of misuse
Adoption/en gagement	Whether the system is used	all architec tures	active contributo rs; decision logs completed	vanity metrics; “usage without value”
Knowledge quality	Trustwort hiness and usability of knowled ge assets	socio-technic al; AI-augme nted	review pass rate; freshness; duplicatio n rate	over-bureaucra tization; slow publishing
IR/retrieval effectiveness	Ability to retrieve relevant knowled	socio-technic al; AI-	precision/r ecall; nDCG;	optimizing offline metrics

Metric family	Definition	Applies most strongly to	Example operation alization	Risks of misuse
Decision proxies	ge quickly Observable correlates of decision quality	augme nted all; especia lly human/socio-technic al	time-to-find [45] cycle time; reversal rate; calibration in forecastin g domains	that miss real tasks confoundi ng (environm ental change), perverse incentives
Learning-loop integrity	Whether feedback updates memory and practice	socio-technic al; AI-augme nted	% incidents with postmorte m; action item closure; recurr enc e rate [55]	ritualized reviews; action items not institution alized
AI assurance & monitoring	Safety and performa	AI-augme nted	drift triggers; bias checks;	monitoring overload; false

Metric family	Definition	Applies most strongly to	Example operation alization	Risks of misuse
	nce stability		audit trails; override behavior [56]	sense of control

Illustrative metric emphasis by architecture (xychart)

The chart below uses **illustrative** (non-empirical) weightings to show how metric emphasis typically shifts as organizations move from human → socio-technical → AI-augmented architectures.

xychart-beta

title "Illustrative metric emphasis by architecture (0–5)"

x-axis ["Human","Socio-technical","AI-augmented"]

y-axis "Emphasis" 0 --> 5

bar "Interaction/process quality" [5,4,3]

bar "KM maturity & governance" [2,5,5]

bar "IR/search effectiveness" [1,4,5]

bar "AI assurance & monitoring" [0,2,5]

Evaluation methods: a rigorous portfolio

A rigorous evaluation program uses multiple methods matched to the artifacts and risks:

Field quasi-experiments: compare pre/post changes •
or matched units (teams/sites) on cycle time,
recurrence, or quality proxies; credible when
confounders are monitored and effects are sustained.

Task-based retrieval tests: build relevance-labeled •
test sets for high-value “golden queries,” then
evaluate with IR metrics including ranked measures.

[\[45\]](#)

Team cognition measures: use structured •
assessments of collective intelligence or process
measures where feasible; evidence supports the
predictive role of collaboration process. [\[57\]](#)

AI assurance: follow lifecycle risk management— •
mapping context, measuring risks, and managing
them continuously—consistent with NIST’s AI RMF
framing. [\[58\]](#)

Governance, implementation strategy, and roadmap

Governance stack: KM + AI risk + model risk (where
relevant)

Effective thinking systems require layered governance
because they combine:

1) **Knowledge assets** (documents, playbooks,
taxonomies),

2) **Decision processes** (decision rights and
accountability), and

3) **Models** (statistical/ML/LLM systems that influence decisions).

Key governance references:

ISO 30401:2018 defines KM management system • requirements and guidelines applicable to any organization type and size; the ISO synopsis is publicly accessible, while the detailed standard text is proprietary. [\[52\]](#)

NIST AI RMF 1.0 is a voluntary, sector-agnostic • framework emphasizing lifecycle risk management; it identifies risks such as drift and privacy risk from aggregation and frames AI as socio-technical. [\[47\]](#)

SR 11-7 (Federal Reserve supervisory guidance) • frames model risk as potential adverse consequences from decisions based on incorrect or misused models and describes pillars spanning development/use, validation, and governance/controls; this is most relevant where models materially affect high-stakes or regulated decisions. [\[59\]](#)

In addition, U.S. banking agencies have adopted aligned supervisory expectations for model risk management (useful as a high-assurance benchmark even outside banking). [\[60\]](#)

Regulatory timelines relevant to AI governance (as of 2026-03-07)

For organizations operating in or serving the EU (jurisdiction unspecified), the **EU AI Act** timeline is material for planning:

The European Commission's official AI Act timeline • states the AI Act entered into force on **1 August 2024** and is fully applicable **two years later (2 August 2026)**, with staged application: prohibitions and AI literacy obligations from **2 February 2025**; obligations for general-purpose AI models from **2 August 2025**; and extended transition for high-risk AI embedded in regulated products until **2 August 2027**. [\[61\]](#)

The binding regulation text is Regulation (EU) • 2024/1689. [\[62\]](#)

Status nuance as of early 2026: EU institutional materials and reporting discuss proposals (e.g., “Digital Omnibus”) that may adjust implementation details or timelines for certain high-risk obligations; these are proposals and should be treated as subject to legislative change. [\[63\]](#)

For broader principles, the Organisation for Economic Co-operation and Development [\[64\]](#) Recommendation on AI (revised in 2023 to update the definition of AI systems) provides widely cited norms for trustworthy, human-centered AI and responsible stewardship. [\[65\]](#)

Implementation strategy: staged and measurable

A staged implementation reduces risk and builds legitimacy through demonstrated wins:

Diagnose: evaluate current-state decision cycles, •
knowledge flows, and pain points; establish baselines.

Design: target operating model (decision rights, •
roles), knowledge model (taxonomy, asset types),
instrumentation plan (metrics), and governance.

Pilot: implement in 1–2 high-value decision loops; •
instrument, learn, and iterate.

Scale: expand to additional teams/functions; harden •
platforms; standardize practices.

Institutionalize: embed in policy, training, •
performance management, and continuous
improvement.

This approach aligns with APQC maturity logic
(progressive capability build) and management system
design logic (review and improve). [\[44\]](#)

Roles and responsibilities (governance and
implementation)

The table below is a cross-sector role template; specific
job titles, reporting lines, and staffing levels are
unspecified.

Role (template)	Core responsibilities	Notes on governance alignment
Executive sponsor	legitimacy, funding, cross- functional prioritization	enables institutionaliz

Role (template)	Core responsibilities	Notes on governance alignment
		ation; removes barriers consistent with change leadership principles [66]
Knowledge manageme nt (KM) owner	KM management system design; lifecycle rules; content standards; adoption plan	aligns to ISO 30401 management system constructs (public synopsis) [52]
Decision architecture lead	decision rights, escalation logic, decision logs, review cadence	reduces bounded- rationality failure modes through structure [29]
Platform/pro duct owner	roadmap for KM/search/collaboration	prevents “portal drift”;

Role (template)	Core responsibilities	Notes on governance alignment
	tooling; integration into workflows	drives usability and uptake
AI risk lead / trust lead	AI RMF implementation; risk mapping/measurement/management; incident response	aligns with NIST AI RMF functions and lifecycle risk management [58]
Model risk/validation lead (where relevant)	independent validation; challenge; model inventory; usage controls	aligns with SR 11-7 pillars for high-stakes model use [67]
Data governance & privacy	data classification; retention; access controls; privacy compliance	addresses privacy risk from aggregation in AI systems [68]
Security engineering	threat modeling; logging; secrets management; vendor risk	critical for AI/knowledge platform

Role (template)	Core responsibilities	Notes on governance alignment
		safety and leakage prevention
Change & learning lead	training; psychological safety interventions; community enablement	aligns with evidence on psychological safety and learning behavior [69]
Community of practice leads	tacit knowledge exchange; “sensemaking forums”; curation input	aligns with communities of practice as learning systems [43]

Eighteen-month roadmap (Gantt)

Dates are set relative to the current context date (2026-03-07). Organization bandwidth and budget are unspecified; therefore, the roadmap is a reference plan.

gantt

title 18-Month Thinking System Roadmap (Reference Plan)

dateFormat YYYY-MM-DD

axisFormat %b %Y

section Diagnose

Baseline mapping (decisions, knowledge, risks) :a1,
2026-03-15, 45d

Baseline metrics + instrumentation plan :a2, 2026-
03-15, 60d

section Design

Target operating model (decision rights, roles) :b1,
2026-05-01, 60d

KM governance design (ISO 30401-aligned) :b2,
2026-05-01, 60d

AI risk governance design (NIST AI RMF-aligned) :b3,
2026-05-15, 75d

section Pilot

Pilot loop #1 (decision log + knowledge capture) :c1,
2026-07-01, 90d

Pilot loop #2 (review loop: AAR/postmortem) :c2,
2026-07-15, 90d

Pilot evaluation (IR tests, cycle time, recurrence):c3,
2026-09-15, 45d

section Scale

Scale to 3–5 additional workflows :d1, 2026-
11-01, 120d

Standardize taxonomy + curation workflow :d2,
2026-11-01, 120d

Enterprise search evaluation program :d3, 2026-
12-01, 120d

section AI augmentation (if applicable)
 RAG/AI decision support MVP with governance gates
 :e1, 2027-02-01, 120d
 Monitoring & drift/bias dashboards :e2, 2027-03-01, 120d

section Institutionalize
 Policy embed + training + performance integration :f1, 2027-06-01, 90d
 Re-baseline assessment + continuous improvement :f2, 2027-08-15, 45d

Milestones (tabular)

Month	Milestone	Evidence of completion
2	Baseline and scope agreed	metrics baseline; decision-loop inventory; target outcomes
4	Governance stack defined	documented ISO 30401-aligned KM governance; NIST AI RMF-aligned AI risk process; model inventory approach if relevant [8]
7	Pilot loops operating	decision logs in workflow; AAR/postmortem cadence established [41]

Month	Milestone	Evidence of completion
9	Pilot evaluation complete	task-based IR test set; baseline-to-pilot comparisons on cycle time/recurrence [45]
13	Scale wave delivered	standardized curation workflow; expanded adoption; search quality improving
16	AI augmentation monitored (if used)	drift/bias/override metrics; governance gates operational [70]
18	Institutionalization	policies updated; training embedded; continuous improvement loop scheduled

Culture and change-management factors

Psychological safety as a structural enabler

Team learning behavior depends on a climate where speaking up is safe. Amy Edmondson [\[71\]](#) introduced team psychological safety as a shared belief that the team is safe for interpersonal risk-taking and empirically linked it to learning behavior and performance in work teams. [\[72\]](#)

In thinking systems, psychological safety is not an HR slogan; it is a **necessary condition** for accurate sensemaking because weak signals, dissent, and error

reporting are key inputs to learning loops. This logic is consistent with blameless postmortem practices in reliability engineering cultures. [\[73\]](#)

Communities of practice as social infrastructure

Etienne Wenger [\[74\]](#) conceptualizes communities of practice as learning systems in which shared practice and identity enable knowledge flow that formal structures often cannot replicate. [\[43\]](#)

In a thinking system, communities of practice function as:

mechanisms for tacit knowledge exchange, •

peer review channels for curation, and •

“sensemaking arenas” where interpretations are •
compared and refined.

Change management: Kotter’s steps as implementation scaffolding

John Kotter [\[75\]](#)’s work emphasizes that transformation often fails due to missing steps (e.g., insufficient urgency, weak coalition, lack of short-term wins). His articulation of change failure modes and an eight-step model is widely used in practice (original HBR formulations may be paywalled; accessible PDF summaries and reprints exist).

[\[76\]](#)

Mapping cultural interventions to roadmap milestones

Culture interventions should be timed to the technical and governance build:

Diagnose/design (months 0–4): establish “speak- •
up” norms for problem framing; train facilitators for
AAR/postmortems; publish principles (blameless
learning, evidence-based debate). [\[77\]](#)

Pilot (months 4–9): institutionalize structured review •
ceremonies; measure participation equality and
learning-loop closure; celebrate early wins to build
credibility. [\[78\]](#)

Scale (months 9–16): formalize communities of •
practice; establish peer-review rotations for
knowledge assets; integrate with performance
systems cautiously to avoid gaming. [\[79\]](#)

Institutionalize (months 16–18): codify practices in •
policy and training; maintain auditability and
continuous improvement cadence (management-
system approach). [\[80\]](#)

Technology enablers, risks, and cross-sector case
evidence

Technology enablers: categories, evaluation criteria, and
engineering requirements

Technology is cognitive infrastructure; its job is to
strengthen capture, retrieval, deliberation, and learning-
loop closure. Critical to rigor is the distinction between
capability (what the tool can do) and **reliability** (how it
performs over time under governance constraints).

Comparative table of technology categories

Technology category	Primary purpose in thinking systems	Evaluation criteria	Engineering/governance notes	Example vendors (optional; non-exhaustive)
KM platforms / knowledge bases	structured content, governance workflows	authoring UX; versioning; review workflow; metadata; retention policies	align to KM lifecycle governance and auditability [81]	Examples include common enterprise suites (vendor choice unspecified)
Collaboration tools	deliberation, coordination, decision capture	integration into workflows; searchability; governance controls	risk: knowledge trapped in chat; require extraction/capture routines	Examples include common chat/wiki/task ecosystems
Enterprise search	discoverability and retrieval	IR metrics (precision/recall/nDCG); latency; explainability	requires relevance test sets and ongoing	Examples include major search stacks

Technology category	Primary purpose in thinking systems	Evaluation criteria	Engineering/governance notes	Example vendors (optional; non-exhaustive)
	reliability	ity; security trimming	evaluation [45]	(unspecified)
Knowledge graphs	shared semantic structure; entity/relationships reasoning	schema quality; coverage; queryability; validation	supports provenance and structured retrieval; aligns to KG definitions and practices [82]	graph DB and KG toolchains (unspecified)
RAG (retrieval-augmented generation)	grounded generation using external	factuality/grounding; citation quality; retrieval success; red-	combines parametric and non-parametric memory; requires retrieval	LLM tooling ecosystems (unspecified)

Technology category	Primary purpose in thinking systems	Evaluation criteria	Engineering/governance notes	Example vendors (optional; non-exhaustive)
	Internal corpora	teaming results	evaluation [83]	
Analytics/dashboards	make patterns visible; monitor outcomes	metric validity; update cadence; alerting	must connect to learning-loop closure; avoid vanity dashboards	common BI stacks (unspecified)
MLOps	reliable ML/AI deployment and monitoring	CI/CD for ML; reproducibility; monitoring coverage; incident response	mitigates ML technical debt and supports lifecycle governance [84]	pipeline/registry/monitoring stacks (unspecified)

Technical debt, provenance, and MLOps needs

AI-augmented thinking systems incur risks beyond traditional IT. “Hidden technical debt” in ML systems includes entanglement, feedback loops, and maintainability burdens that compound silently if not managed. [\[85\]](#)

MLOps practices (automation, reproducibility, monitoring, feedback loops) are emphasized as the operational response to such complexity, including continuous delivery patterns for ML. [\[86\]](#)

RAG architectures specifically address the limitation that model parameters are not a reliable organizational memory by coupling generation with retrieval from external corpora, enabling more updateable knowledge and potential provenance links—while creating new dependencies on retrieval quality and corpus governance. [\[87\]](#)

Risks and ethics: key failure modes and governance controls

Empirically grounded human–automation risks

Automation can improve performance but introduces predictable misuse patterns. In “Humans and Automation: Use, Misuse, Disuse, Abuse,” Raja Parasuraman [\[88\]](#) and Victor Riley [\[89\]](#) synthesize how automation can be misused (overreliance), disused, or abused (inappropriate function allocation). [\[90\]](#)

Automation bias (excessive deference to automated recommendations) remains a central risk in AI decision

support; a recent review in *AI & Society* defines and analyzes automation bias in human–AI collaboration and links it to overreliance dynamics. [\[91\]](#)

A systematic review and meta-analysis in *Nature Human Behaviour* reports that, on average, human–AI combinations perform worse than the best of humans or AI alone, and that decision tasks show performance losses more often than creative tasks—highlighting the ethical and operational risk of deploying hybrid systems without careful design. [\[7\]](#)

NIST’s generative AI profile explicitly notes automation bias and how overreliance can exacerbate risks such as confabulation and bias. [\[92\]](#)

Model risk, privacy, and regulatory compliance

Model risk: SR 11-7 frames the risk of adverse • consequences from incorrect or misused models and calls for robust validation and governance, which becomes directly relevant when AI models influence decisions with financial, safety, or compliance implications. [\[93\]](#)

Privacy/data leakage: NIST AI RMF highlights • privacy risks due to enhanced data aggregation and notes maintenance triggers due to drift; combined with enterprise knowledge and RAG, this necessitates access controls, logging, and data minimization. [\[94\]](#)

Regulatory compliance: for EU-exposed • organizations (jurisdiction unspecified), staged AI Act

obligations create timeline pressure; as of early 2026, official EU sources provide staged application dates, while proposals and debate about adjustments underscore the need for monitoring regulatory change. [\[95\]](#)

Narrative lock-in: sensemaking research shows how power and shared narratives can suppress alternative interpretations, creating systematic blind spots. [\[20\]](#)

Risk matrix (cross-sector baseline)

The following matrix is qualitative and assumes unspecified organization context (sector, scale, and risk appetite are not provided).

Risk	Likelihood (baseline)	Impact (baseline)	Primary mitigations (system-level)	Key references
Automation bias / overreliance	High	High	role allocation; override mechanisms; training; monitoring reliance patterns; blameless review of	[96]

Risk	Likelihood (baseline)	Impact (baseline)	Primary mitigations (system-level)	Key references
Model drift / performance decay	Medium –High	Medium –High	AI-driven errors drift monitoring; triggers; retraining governance ; continuous evaluation	[97]
Privacy/data leakage	Medium	High	data classification; access controls; logging; minimization; vendor risk management	[94]
Model misuse / scope creep	Medium	High	model inventory; usage	[67]

Risk	Likelihood (baseline)	Impact (baseline)	Primary mitigations (system-level)	Key references
			controls; validation; governance gates	
Knowledge base degradation (stale/duplicated)	High	Medium	curation workflows; deprecation ; measurement; ownership	[44]
Narrative lock-in / weak-signal suppression	Medium	High	structured dissent; assumption logging; periodic reframing triggers	[98]

Cross-sector case studies: evidence and transferable lessons

The cases below were selected because they provide public evidence of specific thinking-system mechanisms:

structured capture, review, reuse, and institutionalization into routines.

Case-study summary table

Sector	Case	Elements of thinking system	Evidence (official/primary)	Transferable lesson
Aerospace/engineering	National Aeronautics and Space Administration [99] Lessons Learned Information System	reviewed lessons repository; agency-wide access; supports reuse	NASA describes official reviewed lessons and an online system for collecting/sharing lessons learned [100]	knowledge reuse must be designed as a service with visibility and review
Public accountability	U.S. Government Accountability Office [10]	evaluation of LLIS role; focus on preventing	GAO states LLIS goal is preventing NASA	learning must be embedded into policy/process

Sector	Case	Elements of thinking system	Evidence (official/primary)	Transferable lesson
	11 survey of NASA lessons learned	ing “relearning”	from “relearning” lessons and notes sharing via policy/guidance revisions [102]	documents, not only stored
Military/training	United States Army [103] 1 After Action Reviews	guided analysis with facilitator; conducted at every echelon; immediate standard	FM 7-0 Appendix K defines AAR as guided analysis to improve future performance; includes	disciplined, scaled reflection creates reliable learning loops

Sector	Case	Elements of thinking system	Evidence (official/primary)	Transferable lesson
		s-based feedback	facilitator and occurs across echelons [104]	
Reliability engineering	Google [105] SRE postmortem culture	blameless postmortems; documentation; preventive actions; recurrence reduction	SRE guidance states postmortems document incidents, understand causes, and prioritize preventive actions; blamelessness as	psychological safety plus strong follow-up action management improves reliability

Sector	Case	Elements of thinking system	Evidence (official/primary)	Transferable lesson
			a tenet [53]	
Manufacturing /operations	Toyota Motor Corporation [106] Toyota Production System and Kaizen	standard work + continuous improvement; elimination of waste; learning through experimentation	Toyota describes TPS as eliminating waste and improving efficiency; kaizen described as continuous improvement and standardized work enabling problem detection [107]	standardization can serve learning by making deviations visible

Sector	Case	Elements of thinking system	Evidence (official/primary)	Transferable lesson
Health systems	National Academies of Sciences, Engineering, and Medicine [108] learning health system vision	continuous learning and improvement; feedback loops from care to evidence to practice	“Best Care at Lower Cost” articulates a continuously learning health care system and highlights feedback loops for learning and improvement [109]	feedback loops must be designed at system level, not left to heroics
National statistical agency	Statistics Canada [110] high-	transparency, reproduc	StatsCan describes investme	operational AI requires

Sector	Case	Elements of thinking system	Evidence (official/primary)	Transferable lesson
	level MLOps practice	ibility, rigor, scalability; production ML governance	nts and plans for MLOps emphasizing transparency, reproducibility, rigor, scalability, and efficiency in production ML [111]	lifecycle discipline as a public service capability

Short narrative analysis (transfer patterns)

NASA and GAO: memory systems plus application mechanisms.

NASA positions its lessons learned system as a curated database of official reviewed lessons. [\[100\]](#) GAO’s review emphasizes the system’s objective of preventing “relearning,” and highlights that effective learning also

requires integrating lessons into policies, guidance, and procedural documents—not merely storing them. [\[102\]](#) This pairing illustrates a core thinking-system principle: *organizational memory must connect to decision routines.*

Army AAR and Google postmortems: structured reflection plus psychological safety.

Army doctrine defines AARs as guided analyses with facilitators, conducted at every echelon to improve future performance with immediate, standards-based feedback. [\[112\]](#) Google’s SRE culture frames postmortems as learning opportunities, emphasizing blamelessness and preventive actions. [\[53\]](#) These practices converge on a generalizable design: episodic events become learning only when reflection is structured, blame is minimized to preserve truth-telling, and follow-up actions are tracked to completion.

Toyota, learning health systems, and Statistics Canada: routinized learning as an operating principle.

Toyota publicly frames TPS around waste elimination and efficiency and describes kaizen as continuous improvement supported by standardized work. [\[107\]](#) The National Academies’ learning health system vision emphasizes continuous improvement and feedback loops. [\[113\]](#) Statistics Canada’s MLOps practice similarly frames transparency, reproducibility, and rigor as operational practices—showing that, across sectors, mature thinking systems are not “knowledge portals” but **operating**

philosophies implemented through routines and governance. [\[114\]](#)

Recommendations and an implementation roadmap

Because organization context is unspecified (sector, scale, budget, tool stack, risk appetite), recommendations are provided as **prioritized cross-sector defaults**.

Short-term recommendations (0–6 months)

1) **Choose two high-value decision loops** (e.g., incident response, customer escalation, credit/risk review, clinical triage—specific choice unspecified) and instrument them end-to-end (cycle time, recurrence, decision logs). This creates a measurable nucleus for scaling. [\[115\]](#)

2) **Implement decision traceability as a minimum standard** (decision logs with assumptions and evidence). This mitigates sensemaking lock-in and supports later audits and learning. [\[21\]](#)

3) **Stand up a lightweight governance stack**: KM owner + decision architecture lead + security/privacy lead, with a documented policy for knowledge asset lifecycle (review, deprecate, ownership). [\[81\]](#)

Medium-term recommendations (6–12 months)

4) **Establish an IR evaluation program** for enterprise search: build relevance sets for “golden queries,” track precision/recall/nDCG, and pair with user time-to-find measures. [\[45\]](#)

5) **Institutionalize learning loops** (AAR/postmortems) with closure discipline: action items must be tracked to

completion and policy/playbook updates must be verified. [\[41\]](#)

6) **Build communities of practice** for critical domains to stabilize tacit knowledge exchange and supply peer review for knowledge quality. [\[43\]](#)

Long-term recommendations (12–18 months)

7) **If deploying AI augmentation, enforce governance gates and monitoring:** adopt NIST AI RMF lifecycle practices and define drift/bias/override monitoring; in regulated decisions, adopt SR 11-7-like validation and model inventory controls. [\[116\]](#)

8) **Invest in MLOps as reliability infrastructure** to avoid hidden ML technical debt and operational brittleness; treat ML systems as long-lived products, not experiments. [\[117\]](#)

Roadmap note

The Gantt roadmap and milestone table provided earlier in this report constitute the requested 18-month plan; organization-specific resourcing and prioritization are unspecified and should be set during the “Diagnose” phase. [\[44\]](#)

Methodology note

Search strategy prioritized (in order): (1) official standards and government guidance (NIST AI RMF; ISO 30401 synopsis; SR 11-7), (2) peer-reviewed and high-citation foundational research (Simon 1955; Levitt & March 1988; Woolley et al. 2010; PNAS collective intelligence synthesis; Weick sensemaking), and (3) official

organizational case evidence (NASA, GAO, Army doctrine, Google SRE, Toyota, National Academies, Statistics Canada). [\[118\]](#)

Paywall/access limitations were handled as follows: when primary sources were potentially paywalled (e.g., Science DOI landing pages, some journal portals), accessible versions (author-hosted PDFs, institutional repositories, official summaries, or publicly available reprints) were used where available and cited alongside the canonical reference when feasible. [\[119\]](#)

A user-provided academic report template (Arabic) indicates a conventional five-chapter research structuring approach; integration into that Word template is **unspecified**, and this report therefore uses Word-compatible headings, tables, and mermaid code blocks without modifying user files. □filecite□turn1file0□

[\[1\]](#) [\[9\]](#) [\[14\]](#) [\[74\]](#) Distributed Cognition Edwin Hutchins ...
[https://arl.human.cornell.edu/linked%20docs/Hutchins Distributed Cognition.pdf?utm_source=chatgpt.com](https://arl.human.cornell.edu/linked%20docs/Hutchins_Distributed_Cognition.pdf?utm_source=chatgpt.com)

[\[2\]](#) [\[11\]](#) [\[13\]](#) [\[29\]](#) a behavioral model of rational choice
https://iif.library.cmu.edu/file/Simon_box00063_fld04838_bdl0001_doc0001/Simon_box00063_fld04838_bdl0001_doc0001.pdf?utm_source=chatgpt.com

[\[3\]](#) [\[15\]](#) [\[105\]](#) [\[119\]](#) Evidence for a Collective Intelligence Factor in the ...

https://www.science.org/doi/10.1126/science.1193147?utm_source=chatgpt.com

[4] [19] [21] [33] [39] Organizing and the Process of Sensemaking

https://www.sietmanagement.fr/wp-content/uploads/2016/04/Weick2005.pdf?utm_source=chatgpt.com

[5] [27] [28] [30] Organizational Learning

https://sjbae.pbworks.com/f/levitt_march_1988.pdf?utm_source=chatgpt.com

[6] [36] Some Social and Psychological Consequences of the ...

https://journals.sagepub.com/doi/10.1177/001872675100400101?utm_source=chatgpt.com

[7] When combinations of humans and AI are useful

https://www.nature.com/articles/s41562-024-02024-1?utm_source=chatgpt.com

[8] [52] [80] [81] [103] ISO 30401:2018 - Knowledge management systems

https://www.iso.org/standard/68683.html?utm_source=chatgpt.com

[10] [53] [55] [73] Blameless Postmortem for System Resilience

https://sre.google/sre-book/postmortem-culture/?utm_source=chatgpt.com

[\[12\]](#) [\[31\]](#) [\[101\]](#) Reconceptualizing Organizational Routines as a Source of ...

https://socialecology.uci.edu/sites/socialecology.uci.edu/files/users/feldmanm/Feldman_and_Pentland_2003.pdf?utm_source=chatgpt.com

[\[16\]](#) [\[32\]](#) [\[57\]](#) Quantifying collective intelligence in human groups

https://www.pnas.org/doi/pdf/10.1073/pnas.2005737118?utm_source=chatgpt.com

[\[17\]](#) [\[78\]](#) Evidence for a Collective Intelligence Factor in the ...

https://ofew.berkeley.edu/sites/default/files/evidence_for_a_collective_intelligence_factor_in_the_performance_of_human_groups_woolley_et_al.pdf?utm_source=chatgpt.com

[\[18\]](#) [\[69\]](#) [\[72\]](#) [\[77\]](#) Psychological Safety and Learning Behavior in Work Teams

https://web.mit.edu/curhan/www/docs/Articles/15341_Readings/Group_Performance/Edmondson%20Psychological%20safety.pdf?utm_source=chatgpt.com

[\[20\]](#) [\[75\]](#) [\[98\]](#) Power in Sensemaking Processes

https://repub.eur.nl/pub/126878/0170840619847718.pdf?utm_source=chatgpt.com

[\[22\]](#) [\[26\]](#) [\[35\]](#) [\[110\]](#) Double Loop Learning in Organizations

<https://www.avannistelrooij.nl/wp/wp-content/uploads/2017/11/Argyris-1977-Double-Loop->

[Learning-in-Organisations-HBR.pdf?utm_source=chatgpt.com](#)

[\[23\]](#) [\[34\]](#) The fifth discipline: the art and practice of the learning ...

[https://www.e-education.psu.edu/geog468/sites/www.e-education.psu.edu/geog468/files/TheFifthDiscipline.pdf?utm_source=chatgpt.com](#)

[\[24\]](#) [\[111\]](#) [\[114\]](#) Adopting a high Level MLOps Practice for the Production ...

[https://www.statcan.gc.ca/en/data-science/network/high-level-mlops-practice?utm_source=chatgpt.com](#)

[\[25\]](#) [\[42\]](#) [\[100\]](#) NASA Lessons Learned

[https://www.nasa.gov/nasa-lessons-learned/?utm_source=chatgpt.com](#)

[\[37\]](#) [\[47\]](#) [\[54\]](#) [\[56\]](#) [\[58\]](#) [\[68\]](#) [\[70\]](#) [\[88\]](#) [\[89\]](#) [\[94\]](#) [\[97\]](#) [\[116\]](#) [\[118\]](#) Artificial Intelligence Risk Management Framework (AI ...

[https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf?utm_source=chatgpt.com](#)

[\[38\]](#) Knowledge Management Capability Assessment Tool

[https://www.apqc.org/what-we-do/benchmarking/assessment-survey/knowledge-management-capability-assessment-tool?utm_source=chatgpt.com](#)

[\[40\]](#) A Dynamic Theory of Organizational Knowledge Creation ...

https://www.svilendobrev.com/1/Nonaka_1994-Dynamic_theory_of_organiz_knowledge_creation.pdf?utm_source=chatgpt.com

[\[41\]](#) [\[104\]](#) [\[112\]](#) After Action Reviews

https://www.first.army.mil/Portals/102/FM%207-0%20Appendix%20K.pdf?utm_source=chatgpt.com

[\[43\]](#) [\[79\]](#) [\[106\]](#) Communities of practice:

https://www.wenger-trayner.com/wp-content/uploads/2022/06/1998-EWT-Article-for-the-Systems-Thinker.pdf?utm_source=chatgpt.com

[\[44\]](#) [\[49\]](#) [\[50\]](#) [\[115\]](#) How Mature Is Your KM Program?

https://www.apqc.org/sites/default/files/files/K03302_Understanding_the_APQC_KM_CAT.pdf?utm_source=chatgpt.com

[\[45\]](#) Evaluation in information retrieval

https://nlp.stanford.edu/IR-book/pdf/08eval.pdf?utm_source=chatgpt.com

[\[46\]](#) [\[64\]](#) (PDF) Measuring Transactive Memory Systems in the Field

https://www.researchgate.net/publication/10597070_Measuring_Transactive_Memory_Systems_in_the_Field_Scale_Development_and_Validation?utm_source=chatgpt.com

[\[48\]](#) [\[90\]](#) [\[96\]](#) Humans and Automation: Use, Misuse, Disuse, Abuse

https://web.mit.edu/16.459/www/parasuraman.pdf?utm_source=chatgpt.com

[\[51\]](#) [\[60\]](#) Adoption of Supervisory Guidance on Model Risk
...

https://www.fdic.gov/news/financial-institution-letters/2017/adoption-supervisory-guidance-model-risk-management-revised?utm_source=chatgpt.com

[\[59\]](#) [\[67\]](#) [\[93\]](#) Supervisory Letter SR 11-7 on guidance on Model Risk ...

https://www.federalreserve.gov/supervisionreg/srletters/sr1107.pdf?utm_source=chatgpt.com

[\[61\]](#) [\[95\]](#) AI Act | Shaping Europe's digital future - European Union

https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai?utm_source=chatgpt.com

[\[62\]](#) Regulation - EU - 2024/1689 - EN - EUR-Lex - European Union

https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng?utm_source=chatgpt.com

[\[63\]](#) Digital Omnibus on AI | Think Tank - European Parliament

https://www.europarl.europa.eu/thinktank/en/document/EP_RS_BRI%282026%29782651?utm_source=chatgpt.com

[65] Recommendation of the Council on Artificial Intelligence

https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449?utm_source=chatgpt.com

[66] [76] Leading Change - Why Transformation Efforts Fail

https://heeo.e.hee.nhs.uk/sites/default/files/leading_change_why_transformation_efforts_fail.pdf?utm_source=chatgpt.com

[71] [107] Toyota Production System | Vision & Philosophy | Company

https://global.toyota/en/company/vision-and-philosophy/production-system/index.html?utm_source=chatgpt.com

[82] Knowledge Graphs

https://aidanhogan.com/docs/knowledge-graphs-computing-surveys.pdf?utm_source=chatgpt.com

[83] [87] Retrieval-Augmented Generation for Knowledge-Intensive ...

https://proceedings.neurips.cc/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf?utm_source=chatgpt.com

[\[84\]](#) [\[85\]](#) [\[117\]](#) Hidden Technical Debt in Machine Learning Systems

https://papers.neurips.cc/paper/5656-hidden-technical-debt-in-machine-learning-systems.pdf?utm_source=chatgpt.com

[\[86\]](#) [\[108\]](#) Continuous Delivery for Machine Learning

https://martinfowler.com/articles/cd4ml.html?utm_source=chatgpt.com

[\[91\]](#) [\[99\]](#) Exploring automation bias in human–AI collaboration - Springer

https://link.springer.com/article/10.1007/s00146-025-02422-7?utm_source=chatgpt.com

[\[92\]](#) Artificial Intelligence Risk Management Framework

https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf?utm_source=chatgpt.com

[\[102\]](#) GAO-01-1015R Survey of NASA's Lessons Learned Process

https://www.gao.gov/assets/gao-01-1015r.pdf?utm_source=chatgpt.com

[\[109\]](#) [\[113\]](#) The Path to Continuously Learning Health Care in America

https://professionalismvalue.org/wp-content/uploads/2021/02/Better_Care_at_Lower_Cost_The_Path_to_Continuously_Learning_Health_Care_in_America.pdf?utm_source=chatgpt.com

