Generative AI and Organisational Collective Intelligence: A Dependency-Structured Framework

Abstract

Generative AI (GenAI) is now part of everyday organisational knowledge work, affecting how information is gathered, interpreted, argued over, and stored. It should not be treated as a neutral aid: in many settings it can function as an active participant in collective intelligence, shaping what groups notice and how issues are framed and settled. This paper sets out a dependency-structured framework that connects information acquisition, sensemaking and framing, shared reasoning, coordinated action, and organisational memory. The framework is intended as a diagnostic and explanatory lens for organisational analysis, rather than a predictive or causal model. A key implication of the framework is that weaknesses in early stages can propagate upward and distort later decisions, even when outputs appear faster or more coherent. A hypothetical case of a mid-sized financial advisory firm illustrates how GenAI can strengthen performance while risking increased epistemic fragility when foundational processes are compressed or bypassed. The paper ends with diagnostic prompts and governance principles for information professionals, arguing that GenAI tends to amplify existing organisational tendencies rather than reliably augment intelligence.

Keywords: Generative Artificial Intelligence (GenAI), Collective Intelligence, Organisational Sensemaking, Human–AI Collaboration, Information Governance

Introduction

The rapid diffusion of generative artificial intelligence (GenAI) is changing organisational information practices in ways few earlier technologies have managed. Earlier AI systems were typically built to automate well-defined tasks or to analyse structured data. Generative models are different. They operate directly on language and are now used for work that sits close to judgement: summarising documents, synthesising evidence, framing issues, and drafting arguments that feed into how organisational decisions are prepared and justified. As a result, GenAI is no longer confined to technical support roles at the margins of organisations. It is increasingly embedded in the everyday cognitive work through which organisations scan their environments, interpret signals, deliberate about options, and retain knowledge. This raises a practical question for information professionals and organisational leaders: what happens to collective intelligence when a non-human system participates directly in sensemaking and reasoning? Here, sensemaking refers to the collective work of interpreting ambiguous signals and stabilising shared frames that guide attention and action. Throughout the paper, speed and surface coherence are treated as properties of outputs, whereas collective intelligence refers to the robustness of the underlying epistemic processes that produce them. Specifically, in this paper, epistemic fragility refers to a weakening of those underlying processes of information acquisition, interpretation, and reasoning, such that outputs appear coherent while resting on a narrowed or insufficient evidential base. Fragility does not necessarily produce immediate error; rather, it reflects reduced resilience in how collective intelligence is formed and sustained over time.

Existing research offers partial answers to this question. Studies of collective intelligence show that group performance depends on how well teams coordinate attention and integrate diverse perspectives to sustain productive interaction (Malone, Laubacher and Dellarocas, 2010; Woolley et al., 2010). Research on human–AI collaboration further suggests that AI systems can extend attention and analytic capacity in team settings (Gupta et al., 2025). At the same time, emerging empirical work points to countervailing effects, including reduced exploration and erosion of shared knowledge resources when generative systems are used without care (Bender et al., 2021; Burtch et al., 2024; del Rio-Chanona et al., 2024).

What remains underdeveloped is a way of explaining how these opposing effects arise inside organisations. Much existing investigation focuses on individual tasks or isolated interactions. In practice, organisations commonly report faster decisions and more polished outputs alongside concerns that over-reliance on AI summaries can drive premature convergence and shallow organisational learning. Without a framework that spans the full set of processes involved in collective intelligence, these tensions are difficult to diagnose or govern. To address this gap, a five-layer framework of organisational collective intelligence is set out. The framework emphasises how higher-order performance depends on the integrity of more foundational processes. It brings together established insights from information science, organisational theory, and collective intelligence research within a dependency-structured model. A central aim is to treat GenAI not simply as a tool, but as an active participant operating across information acquisition, sensemaking, shared reasoning, coordinated action, and organisational memory. The analysis focuses on knowledge-intensive organisations, particularly those operating in information-rich or regulated environments, where GenAI is embedded in judgement-adjacent work.

Collective Intelligence in Organisations

Collective intelligence (CI) refers to the enhanced capacity that emerges when groups combine knowledge and effort to solve problems or make decisions. It is not reducible to individual intelligence, nor to the simple aggregation of opinions. Rather, it depends on how structured interaction and shared representations are supported by mechanisms that allow information from different sources to be integrated and acted upon. Across disciplines, collective intelligence has been approached from several complementary perspectives. In cognitive science, it is closely associated with distributed cognition, in which thinking is spread across people, artefacts, and environments (Hutchins, 1995). In organisational and management research, CI is examined through group problem-solving and coordination processes (Woolley et al., 2010). From information science and sociological perspectives, it is understood as a set of social knowledge processes through which information is produced, validated, and retained over time (Lévy, 1997). From an information perspective, organisational collective intelligence depends on several interrelated capacities. These include the ability to acquire relevant signals, interpret them collectively, reason about options, coordinate action, and retain what is learned. Breakdowns in any of these areas can degrade CI even when individuals are capable and well intentioned. For this reason, CI is best understood as a socio-technical system shaped by information infrastructures and organisational routines and norms, as well as by human cognition.

Recent scholarship converges on the view that collective intelligence cannot be reduced to individual cognition or simple aggregation. Instead, it emerges from structured interaction among people, artefacts, and institutional processes. Systematic reviews describe CI as a multi-level phenomenon spanning information gathering, sensemaking and framing, deliberation, coordination, and learning, shaped by both social and technical infrastructures (Søilen, 2019; Suran, Pattanaik and Draheim, 2020; Berditchevskaia et al., 2022). More recent theoretical work develops this view further by treating CI as an architectural system, in which higher-order coordination and decision quality depend on the integrity of lower-level epistemic processes (Xiao, 2024; Friston et al., 2024). From this perspective, CI is less a property of groups alone than of the socio-technical ecosystems that structure how information is acquired and interpreted, then contested and retained over time. This framing also draws attention to vulnerability: weaknesses at lower levels can propagate upward, producing outcomes that appear coherent but are fragile.

AI-Enabled Collective Intelligence

AI-enabled collective intelligence (AI-CI) refers to socio-technical configurations in which GenAI participates in the processes through which groups acquire information, interpret signals, coordinate action, and generate shared knowledge. Earlier generations of AI tended to support collective intelligence indirectly, for example by improving data access, or automating workflows, or producing analytic outputs that were then interpreted by humans. GenAI alters this relationship in more fundamental ways. Contemporary models engage directly in linguistic and cognitive activities that previously took place through human collaboration. They are used for everything from summarising complex materials to proposing interpretations, generating options, and simulating stakeholder perspectives. As a result, GenAI shapes not only what information is available, but also how it is framed and discussed over time. A growing body of work characterises AI-enabled collective intelligence as a hybrid arrangement in which humans and AI jointly participate in sensemaking, reasoning, and decision-making (Peeters et al., 2021; Cui et al., 2024; Riedl and De Cremer, 2025). Some studies emphasise the potential of generative models to expand search spaces, scaffold collaboration, and support collective problem-solving (Heyman et al., 2024; Woolley, 2025; Page, 2025). Others point to risks when these systems are left unguided, including compressed deliberation, reduced independent contribution, and premature convergence (Halpin, 2025; del Rio-Chanona et al., 2024).

Taken together, GenAI is treated here not as a neutral accelerator but as an active participant whose influence is systemic, shaping patterns of contribution and stabilised interpretations, which in turn influence organisational memory (Casebourne et al., 2025; Zhang et al., 2025). GenAI systems can augment collective intelligence by extending attention, memory, and reasoning capacity. They can also distort it by narrowing exploration, which encourages premature convergence and displaces human contributions to shared knowledge resources. Understanding these effects requires moving beyond task-level analysis toward a layered view of how AI-CI is constructed and sustained within organisations. This shift brings questions of governance, scaffolding, and institutional design to the foreground (Suran et al., 2022; Gambrell, 2025).

A Layered Framework of Organisational Collective Intelligence

This section sets out a layered architecture for understanding organisational collective intelligence. The framework is visualised in Figure 1 as a pyramid, using dependency rather than hierarchy to organise the layers. Breakdowns in collective intelligence are most visible at higher levels of performance but most often originate in lower-level epistemic processes. Accordingly, higher-order outcomes in AI-enabled collective intelligence, such as coordinated execution and organisational learning, depend on the reliability of more foundational capacities, including information acquisition, sensemaking and framing, and shared reasoning. GenAI can accelerate activity at upper layers even when lower layers are weak, creating what can appear as an illusion of intelligence. Over time, the layers also shape one another: what organisations remember and institutionalise influences future scanning, interpretation, and reasoning; thereby generating feedback loops across the architecture. The annotations in Figure 1 indicate how common GenAI activities give rise to characteristic CI failure modes when underlying dependencies are weak. Each layer is grounded in established research across information science, organisational studies, cognitive science, and human–computer interaction, and is increasingly affected by the integration of GenAI into everyday knowledge work.

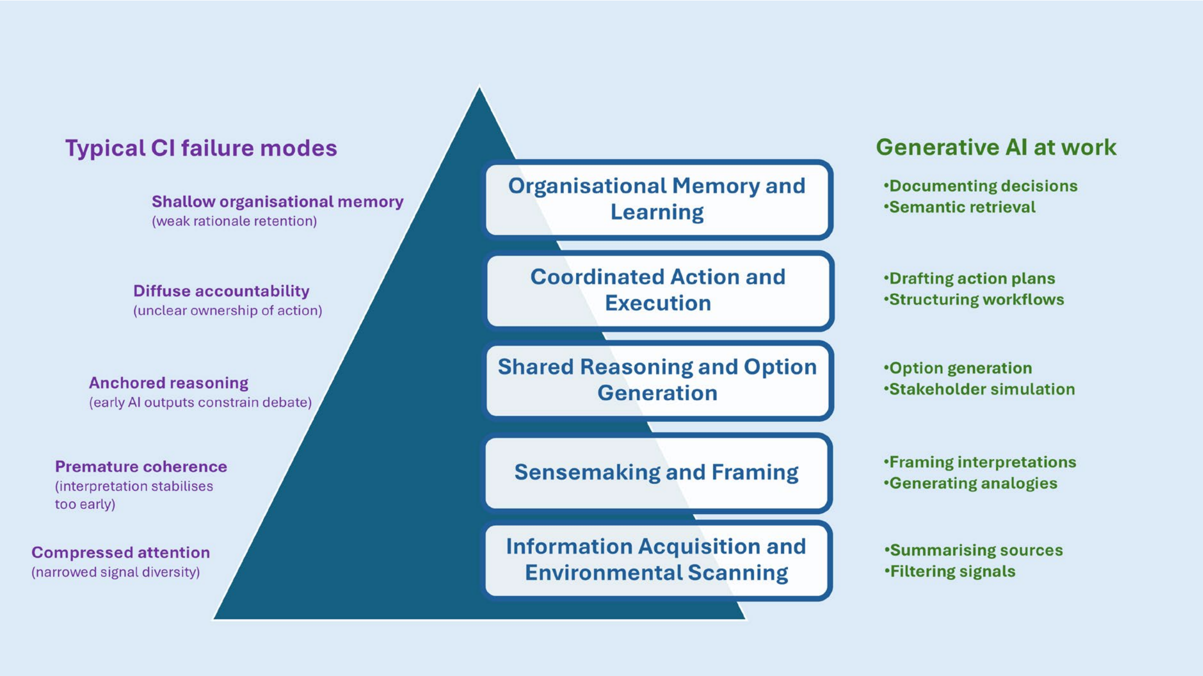

Figure 1. Conceptual framework of organisational collective intelligence and GenAI participation

Layer 1: Information Acquisition and Environmental Scanning

Collective intelligence depends first on how organisations notice and prioritise signals from their environments. Work on environmental scanning (Aguilar, 1967) and organisational attention (Ocasio, 1997) shows that this process is selective and constrained rather than comprehensive. From a sensemaking perspective, what is noticed and acted upon is shaped by shared interpretations rather than by information availability alone (Weick, 1995). At this layer, limits on attention and exploration set the conditions under which collective intelligence can develop.

GenAI can widen access to information by accelerating search and synthesis across sources. At the same time, it can narrow how organisations explore those sources. When AI-generated summaries become the default entry point, engagement with primary material often declines, and exploration becomes more shallow. In this way, increases in speed and volume can coexist with a loss of signal diversity.

Layer 2: Sensemaking and Framing

Sensemaking refers to the social processes through which organisations interpret ambiguous information and stabilise shared frames that guide attention and action (Weick, 1995). These processes do not aim for objective accuracy. They privilege plausibility and coherence, allowing groups to move forward under conditions of uncertainty.

GenAI can intervene directly in this work. It can generate alternative interpretations, surface assumptions, and make it easier to compare competing frames. At the same time, its fluency creates risks. AI-generated narratives often appear confident and internally consistent, even when they rest on weak or incomplete evidence. When such narratives are accepted without challenge, shared understanding can stabilise quickly while remaining fragile. What looks like clarity may mask unresolved ambiguity.

Layer 3: Shared Reasoning and Option Generation

Shared reasoning concerns how groups evaluate interpretations and test assumptions as they generate options through collective deliberation. This work builds on distributed cognition, in which reasoning is spread across people and artefacts (Hollan, Hutchins and Kirsh, 2000), as well as on theories of team reasoning and collective problem-solving (Sugden, 1993; Woolley et al., 2010). The quality of outcomes depends less on individual insight than on how disagreement and subsequent revision are handled over time.

GenAI can shape these dynamics in several ways. It can support deliberation, from moderating discussion and surfacing alignment or disagreement to simulating stakeholder perspectives (Argyle et al., 2023; Park et al., 2023). In some settings, this expands the space of reasoning. In others, AI-generated options become anchors that structure subsequent discussion. When outputs are treated as authoritative starting points rather than provisional contributions, disagreement weakens and exploration narrows. Reasoning becomes faster, but also thinner.

Layer 4: Coordinated Action and Execution

Coordinated action concerns how shared reasoning is translated into aligned execution across roles, teams, and time. Research on organisational routines and coordination shows that action depends not only on plans, but on how attention is allocated, responsibilities are understood, and work is synchronised across settings (Feldman and Pentland, 2003; Ocasio, 1997). Effective coordination relies on shared situational awareness and clear ownership, especially when work unfolds under time pressure.

GenAI can accelerate coordination, from producing action plans to structuring communication and supporting handoffs across teams. These capabilities reduce friction, but they also introduce new failure modes. When plans are generated quickly and circulate widely, responsibility can become diffuse. Tasks framed as “AI-recommended” are sometimes treated as provisional rather than owned. In such cases, execution speeds up while alignment weakens. Coordination appears smoother, but gaps emerge when plans encounter contextual constraints that were never surfaced.

Layer 5: Organisational Memory and Learning

Organisational memory refers to the systems and practices through which organisations retain, retrieve, and reuse knowledge over time (Walsh and Ungson, 1991). Memory is not simply a repository of past decisions. It shapes future sensemaking by influencing what is recalled or ignored and which lessons are treated as settled.

GenAI affects organisational memory in opposing ways. It lowers the cost of producing documentation and makes knowledge easier to retrieve through conversational interfaces. At the same time, it can weaken incentives to curate shared knowledge. When large volumes of AI-generated material accumulate, repositories grow quickly but lose clarity. Decision rationales and contextual judgement are often thinly captured. Emerging research also points to broader effects on public knowledge commons, with indirect consequences for organisational learning ecosystems (Bender et al., 2021; Burtch et al., 2024; del Rio-Chanona et al., 2024). Memory expands in volume, but becomes harder to reuse well.

Positioning and Contribution

This paper's contribution can be situated in relation to three strands of work on collective intelligence and related fields. The first is classic research on collective intelligence, which focuses on how groups combine people, incentives, structures, and processes to produce intelligent outcomes. Work such as Malone, Laubacher and Dellarocas (2010) offers a useful "genome" for describing these combinations. That perspective remains valuable, but it largely treats technology as a component within coordination systems. It is less attentive to what changes when technology operates inside the cognitive workflow itself.

The second strand centres on team and group intelligence. Research in this tradition, including Woolley et al. (2010) and more recent work on hybrid human–AI teams (Woolley et al., 2023; Woolley, 2025), has clarified the conditions under which groups reason effectively and adapt to complexity. This paper complements that work by shifting the level of analysis. Rather than focusing on teams in isolation, it examines how collective intelligence is shaped across organisational processes that precede and outlast individual episodes of collaboration. This perspective helps explain a pattern reported in practice: GenAI can increase speed and surface coherence while weakening the conditions for robust judgement, especially contested interpretation and productive disagreement.

The third strand concerns broader accounts of artificial intelligence and collective intelligence that emphasise large-scale human–machine collaboration and the emergence of CI fields and infrastructures since the late 1990s (Mulgan, 2018; Lévy, 1997). These accounts highlight the systemic nature of collective intelligence, but they often operate at a high level of abstraction. Building on this work, the present framework is designed for organisational diagnosis. It integrates distributed cognition (Hutchins, 1995; Hollan, Hutchins and Kirsh, 2000), organisational sensemaking (Weick, 1995), and organisational memory (Walsh and Ungson, 1991) within a single dependency-structured model. The contribution lies not in asserting that these processes matter, but in showing how GenAI increasingly acts across all of them at once, with consequences that depend on how it is governed and embedded.

Taken together, the framework offers (1) a clearer account of where GenAI reshapes collective intelligence, (2) a dependency logic that explains how brittle intelligence can emerge alongside rising productivity, and (3) a diagnostic bridge between CI theory and information governance concerns in organisational practice.

Diagnosing Collective Intelligence Bottlenecks

The diagnostic guidance that follows is intended as a prompt for information professionals and organisational leaders to reflect on where collective intelligence may be weakening across the CI architecture. The focus is not on evaluating GenAI systems in isolation, rather it is on diagnosing how GenAI interacts with the organisation's information, sensemaking, and coordination arrangements. Accordingly, what appear as failures of reasoning or judgement at the output level often originate in earlier breakdowns that propagate upward through the CI architecture, producing brittle intelligence masked by efficiency. Table 1 summarises the diagnostic questions used to locate where breakdowns are occurring across the CI architecture.

Table 1. Collective Intelligence Diagnostic Guidance Matrix

CI Layer (function)	Breakdown Symptoms	GenAI Risk Signals	Diagnostic Questions
Learn (retention & reuse)	• Documentation without reuse • Knowledge loss with turnover • Defaulting to AI over repositories	Multiple AI "truths" coexist Low-value content crowds repositories Declining knowledge contributions	Do we know what the organisation knows? What AI outputs are retained, and why? Are we weakening future learning capacity?
Execute (alignment & accountability)	• Plans ignored or weakly enacted • Ambiguous ownership • Faster execution, weaker alignment	Generic or context-blind plans Superficial compliance Failures blamed on "implementation"	Are actions clearly human-owned? Do plans reflect tacit constraints? Where does accountability sit if AI is wrong?
Reason (options & trade-offs)	• Refinement replaces generation • Conflict avoidance framed as objectivity • Declining exploratory reasoning	First AI output anchors discussion Simulations replace perspective-taking Disagreement becomes procedural	Is AI diversifying or shortcutting reasoning? Is productive disagreement visible? When do teams override AI logic?
Frame (interpretation & narrative)	• Incompatible mental models • Ambiguous risks remain unowned • Low contestation of interpretations	Early AI drafts become default stories Declining interpretive disagreement Coherent but fragile narratives	Are multiple frames compared before convergence? Does AI accelerate sensemaking or force closure? Where do alternative interpretations originate?
Acquire (signals & sources)	• Persistent overload • Narrow or habitual scanning • Fragmented data sources	Reliance on AI summaries over originals Convergent framings across units Weak signals absent from outputs	Who defines information salience? Is AI expanding or compressing attention? Which sources never surface, and why?

The diagnostic matrix should be read vertically as well as horizontally. Apparent failures in reasoning or execution often originate in earlier breakdowns in information acquisition or framing. In organisations that adopt GenAI rapidly, a recurring concern is that outputs become faster and more coherent while underlying collective intelligence degrades. The matrix is therefore meant to redirect attention away from isolated tool performance and toward the interaction between GenAI and organisational practices over time.

Rather than applying every diagnostic question exhaustively, practitioners can begin with a small number of high-leverage probes. At the level of information acquisition, the key issue is whether GenAI expands organisational attention or compresses it by privileging familiar sources and dominant framings. At the level of sensemaking, leaders should examine whether multiple interpretations are surfaced before convergence, or whether early AI-generated drafts become the default narrative through inertia. In shared reasoning, the signal to watch is whether disagreement remains substantive or whether appeals to AI output dampen exploration. Persistent failures in execution or learning should prompt scrutiny of upstream responsibility and knowledge stewardship, rather than being dismissed as implementation problems.

Hypothetical Case Study: A Mid-Sized Financial Advisory Firm

The following hypothetical case is used to illustrate how GenAI-related frictions can emerge across the CI architecture. The case is not intended to be representative, but to surface mechanisms that may remain latent in empirical accounts of GenAI adoption. Arden Finance is a mid-sized financial advisory and risk consulting firm operating across multiple regulatory jurisdictions. The firm has established cloud infrastructure and data governance, but practices related to analytics, automation, and knowledge management vary across teams. In response to growing regulatory complexity, particularly in financial crime compliance, the firm introduces GenAI tools to support scanning, analysis, and advisory work. Leadership expects improvements in responsiveness and workload reduction, along with better decision quality. What emerges instead is a pattern of mixed outcomes. Performance improves in visible ways, but epistemic fragility can increase beneath the surface. These dynamics reflect how GenAI interacts with existing collective intelligence practices rather than the capabilities of the tools themselves.

Layer 1: Information Acquisition: *Compressed Attention*

Before adopting GenAI, analysts manually monitored sources ranging from regulatory updates to enforcement actions and industry guidance. With GenAI, these sources are unified and summarised into near real-time briefings. Coverage increases, but several unintended effects follow. Junior analysts rely heavily on AI-generated summaries and consult primary sources less often. Region-specific nuances are missed, and scanning narrows toward sources that are easiest to summarise. Leadership later identifies blind spots in emerging risks that were weakly represented in training data.

CI implication: The result is higher information throughput paired with reduced signal variety. Weak signals and local deviations become easier to miss.

Layer 2: Sensemaking and Framing: *Premature Coherence*

Weekly compliance briefings increasingly begin with AI-generated draft narratives. These drafts accelerate preparation but also shape the boundaries of discussion. Interpretations converge early, and minority viewpoints that previously sustained debate become less visible. GenAI frequently frames new developments by analogy to past cases, flattening contextual differences. Ambiguity remains, but the narrative appears coherent enough to close discussion.

CI implication: Shared understanding forms quickly, but it is less resilient. Alternative framings are filtered out before they can be tested.

Layer 3: Shared Reasoning: *Anchored Deliberation*

In cross-functional meetings, teams work from AI-generated option sets that structure subsequent discussion. Deliberation focuses on refining or rejecting predefined options rather than exploring the problem space. Although GenAI is sometimes used to simulate regulator or client perspectives, these simulations often substitute for direct debate. Where teams instruct GenAI to generate conflicting analyses, reasoning improves. Elsewhere, option spaces narrow early.

CI implication: Deliberation becomes anchored to initial AI outputs, making reasoning quality highly sensitive to how GenAI is introduced.

Layer 4: Coordinated Action: *Diffuse Accountability*

GenAI-generated plans and communication drafts accelerate execution. At the same time, ownership becomes less clear. Some tasks are treated as AI recommendations rather than managerial commitments. Teams comply procedurally but disengage when plans fail to account for tacit constraints or client-specific judgement.

CI implication: Execution speeds up, but alignment weakens. Responsibility becomes harder to trace.

Layer 5: Organisational Memory: *Shallow Learning*

GenAI tools automatically generate documentation for decisions and advice. Over time, stored material grows rapidly while curation declines. Multiple AI-generated versions of similar documents coexist, and decision rationales are weakly captured. Analysts increasingly consult GenAI rather than internal repositories. Contributions to shared knowledge decline.

CI implication: Organisational memory expands in volume, but loses context. Knowledge is easier to retrieve than to reuse well.

Integrating GenAI into Collective Intelligence: Design and Governance Principles

To operationalise the framework, this section offers a set of design and governance principles for integrating GenAI into collective intelligence. While many of these principles may appear familiar in isolation, the framework shows why their effects are amplified or muted depending on where they operate within the CI architecture. Further, they are intended as actionable prompts for shaping workflows, incentives, and accountability across the CI architecture, not as an evaluation of tools in isolation.

Decouple exploration from summarisation.

When GenAI is used to summarise large volumes of information, exploration can be unintentionally compressed. Treating exploration and consolidation as distinct phases helps preserve scanning breadth and reduces the risk of missing weak signals.

Preserve contested sensemaking.

Introducing AI-generated narratives too early can stabilise meaning before alternatives are explored. Sensemaking practices should ensure that human interpretations are surfaced before AI-generated frames and that GenAI is used to generate competing perspectives rather than a single account.

Design for divergence before convergence.

GenAI is often deployed to accelerate convergence. Deliberative workflows should instead require divergence before evaluation, using GenAI to surface alternatives that challenge dominant assumptions.

Guard against AI-anchored groupthink.

Fluent AI outputs can act as cognitive anchors. Treating them as provisional and requiring explicit critique helps preserve productive disagreement.

Maintain human ownership of action.

When GenAI generates plans or recommendations, responsibility can blur. Clear validation, assignment of ownership, and documentation of judgement help maintain accountability.

Curate organisational memory.
Lower documentation costs increase the risk of accumulation without learning. Clear criteria for retention and annotation are needed to ensure that organisational memory supports future sensemaking rather than simply storing outputs.
Design GenAI as a visible participant.
When GenAI operates invisibly, its influence on attention, framing, and reasoning is harder to govern. Making AI participation explicit supports accountability and more deliberate use.
Govern collective intelligence at the system level.
Governance should attend to how GenAI reshapes information flows and reasoning dynamics, as well as coordination and learning over time.

These principles set out the conditions under which AI-enabled collective intelligence is less likely to become brittle as GenAI use scales. The following section turns to what it takes to put these conditions into practice and to recognise early signs of breakdown when they are not being met.

Practical Implications and Governance Considerations

The implications below focus on how organisations can operationalise the preceding principles and monitor whether collective intelligence is being sustained in practice. They are derived from the framework and the illustrative case and are intended as practical considerations rather than prescriptive rules.

One implication is the value of explicit stewardship for AI-enabled collective intelligence that spans technical, informational, and organisational domains. Small cross-disciplinary groups can fulfil this role by maintaining shared standards, templates, and guardrails. The aim is coherence and learning, not centralised decision-making. A second implication is the need to monitor collective intelligence itself, not only AI system performance. Periodic reviews can examine changes in information breadth, interpretive diversity, reasoning dynamics, coordination, and learning across teams. Warning signals include narrowing source diversity and declining disagreement, which can contribute to reduced reuse of organisational knowledge even when outputs remain fast and polished. A further implication concerns capability development beyond technical tool use. Teams need capabilities to interpret uncertainty, question apparent coherence, and engage constructively with disagreement. These capabilities help ensure that GenAI supports judgement rather than displacing it.

Finally, sustaining AI-enabled collective intelligence depends on deliberate stewardship of organisational memory. As GenAI lowers the cost of producing documentation, clearer criteria are needed for what is retained and how AI-generated artefacts are labelled. It also matters which materials capture context, rationale, and uncertainty rather than only polished outputs. Without this, organisational memory grows quickly but contributes little to learning.

Conclusion

Generative AI is reshaping organisational information work by operating directly within interpretive and reasoning processes. As a result, its effects on collective intelligence are systemic rather than task-specific. They cannot be understood by evaluating tools or workflows in isolation. GenAI participates in the processes through which organisations acquire information, make sense of uncertainty, coordinate action, and retain knowledge. This paper introduced a five-layer architectural framework that conceptualises organisational collective intelligence as a dependency-structured system spanning the CI architecture, from information acquisition through sensemaking and shared reasoning to coordinated action and organisational memory. The framework shows how GenAI can enhance performance across these layers while also amplifying existing weaknesses. The illustrative case of Arden Finance demonstrated how faster decisions and more coherent outputs can coexist with growing epistemic fragility when foundational processes are compressed or bypassed. The central implication is that effective GenAI adoption is in large part a governance challenge. The diagnostic guidance and design principles developed here highlight how collective intelligence can be strengthened or eroded depending on how GenAI is embedded within information practices and accountability structures. Interventions at lower layers of the architecture, particularly those that preserve scanning breadth, contested interpretation, and productive disagreement, are especially consequential for execution quality and organisational learning. Taken together, the framework reframes GenAI not as a neutral accelerator or a substitute for judgement, but as a participant in organisational collective intelligence whose influence must be actively shaped. GenAI will not make organisations intelligent on its own. It will amplify existing organisational tendencies, for better or for worse.

Limitations and Future Work

This paper has several limitations that also point to directions for future research. First, the framework is conceptual rather than empirically validated. While the hypothetical case study is grounded in realistic organisational practices and informed by existing research, future work should test the framework through empirical methods, including field studies, surveys, and longitudinal analyses of GenAI adoption. Second, the framework operates primarily at a meso level, focusing on organisational processes and information infrastructures that shape collective intelligence across teams and over time. It does not model micro-level interactions, such as moment-to-moment prompting strategies, conversational dynamics between individuals and GenAI systems, or fine-grained cognitive effects in specific encounters. As a result, the framework is intended to explain structural patterns and governance implications rather than predict outcomes in particular interactions. Future studies could examine how micro-level practices shape transitions between layers, particularly in sensemaking and shared reasoning. Third, the analysis centres on knowledge-intensive organisations operating in regulated or information-rich environments. The applicability of the framework to other contexts, such as frontline service work, creative industries, public-sector decision-making, or very small organisations, remains an open question. Comparative studies across organisational settings would help clarify the framework's scope and boundary conditions. Finally, the paper highlights potential risks to organisational and public knowledge commons arising from widespread GenAI use, but it does not assess long-term effects on knowledge production ecosystems. Future research could explore how organisations might contribute to sustaining shared knowledge infrastructures while also benefiting from generative systems.

References

Aguilar, F.J. (1967) *Scanning the business environment*. New York: Macmillan.

Argyle, L.P., Bail, C.A., Busby, E.C., Gubler, J.R., Howe, T., Rytting, C., Sorensen, T. and Wingate, D. (2023) ‘Leveraging AI for democratic discourse: chat interventions can improve online political conversations at scale’, *Proceedings of the National Academy of Sciences*, 120(41), p. e2311627120.

Bender, E.M., Gebru, T., McMillan-Major, A. and Shmitchell, S. (2021) ‘On the dangers of stochastic parrots: can language models be too big?’, in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. New York: ACM, pp. 610–623.

Berditchevskaia, A., Maliaraki, E. and Stathoulopoulos, K. (2022) ‘A descriptive analysis of collective intelligence publications since 2000, and the emerging influence of artificial intelligence’, *Collective Intelligence*, 1(1), p. 26339137221107924.

Burtch, G., Lee, D. and Chen, Z., 2024. The consequences of generative AI for online knowledge communities. *Scientific Reports*, *14*(1), p.10413.

Casebourne, I., Shi, S., Hogan, M., Holmes, W., Hoel, T., Wegerif, R. and Yuan, L. (2025) ‘Using AI to support education for collective intelligence’, *International Journal of Artificial Intelligence in Education*, 35(3), pp. 1597–1629.

Cui, H. *et al.* (2024) ‘AI-enhanced collective intelligence’, *Patterns*, 5(11), p. 101074.

del Rio-Chanona, R.M., Laurentsyeva, N. and Wachs, J. (2024) ‘Large language models reduce public knowledge sharing on online Q&A platforms’, *PNAS Nexus*, 3(9), p. pgae400.

Feldman, M.S. and Pentland, B.T. (2003) ‘Reconceptualizing organizational routines as a source of flexibility and change’, *Administrative Science Quarterly*, 48(1), pp. 94–118.

Friston, K.J., Ramstead, M.J., Kiefer, A.B., Tschantz, A., Buckley, C.L., Albarracin, M. *et al.* (2024) ‘Designing ecosystems of intelligence from first principles’, *Collective Intelligence*, 3(1), p. 26339137231222481.

Gambrell, D. (2025) ‘AI for eGovernance: combining artificial intelligence and collective intelligence to develop evidence-based AI policy’, in *Proceedings of the Eleventh International Conference on eDemocracy & eGovernment (ICEDEG)*. IEEE, pp. 317–324.

Gupta, P., Nguyen, T.N., Gonzalez, C. and Woolley, A.W., 2025. Fostering collective intelligence in human–AI collaboration: laying the groundwork for COHUMAIN. *Topics in cognitive science*, *17*(2), pp.189-216

Halpin, H. (2025) ‘Artificial intelligence versus collective intelligence’, *AI & Society*, 40, pp. 4589–4604.

Heyman, J.L., Rick, S.R., Giacomelli, G., Wen, H., Laubacher, R., Taubenslag, N. *et al.* (2024) ‘Supermind ideator: how scaffolding human–AI collaboration can increase creativity’, in *Proceedings of the ACM Collective Intelligence Conference*, pp. 18–28.

Hollan, J., Hutchins, E. and Kirsh, D. (2000) ‘Distributed cognition: toward a new foundation for human–computer interaction’, *ACM Transactions on Computer-Human Interaction*, 7(2), pp. 174–196.

Hutchins, E. (1995) *Cognition in the wild*. Cambridge, MA: MIT Press.

Lévy, P. (1997) *Collective intelligence: mankind's emerging world in cyberspace*. New York: Perseus Books.

Malone, T.W., Laubacher, R. and Dellarocas, C. (2010) 'The collective intelligence genome', *MIT Sloan Management Review*, 51(3), pp. 21–31.

Mulgan, G. (2018) 'Artificial intelligence and collective intelligence: the emergence of a new field', *AI & Society*, 33, pp. 631–632.

Ocasio, W. (1997) 'Towards an attention-based view of the firm', *Strategic Management Journal*, 18(S1), pp. 187–206.

Page, S.E. (2025) 'Everyone, everywhere, all at once: LLMs and the new physics of collective intelligence', *Collective Intelligence*, 4(3), p. 26339137251367733.

Park, J.S., O'Brien, J., Cai, C.J., Morris, M.R., Liang, P. and Bernstein, M.S. (2023) 'Generative agents: interactive simulacra of human behavior', in *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–22.

Peeters, M.M., Van Diggelen, J., Van den Bosch, K., Bronkhorst, A., Neerincx, M.A., Schraagen, J.M. and Raaijmakers, S. (2021) 'Hybrid collective intelligence in a human–AI society', *AI & Society*, 36(1), pp. 217–238.

Riedl, C. and De Cremer, D. (2025) 'AI for collective intelligence', *Collective Intelligence*, 4(2), p. 26339137251328909.

Søilen, K.S. (2019) 'Making sense of the collective intelligence field: a review', *Journal of Intelligence Studies in Business*, 9(2), p 6.

Sugden, R. (1993) 'Thinking as a team: toward an explanation of non-selfish behaviour', *Social Philosophy and Policy*, 10(1), pp. 69–89.

Suran, S., Pattanaik, V. and Draheim, D. (2020) 'Frameworks for collective intelligence: a systematic literature review', *ACM Computing Surveys*, 53(1), pp. 1–36.

Suran, S., Pattanaik, V., Kurvers, R., Hallin, C.A., De Liddo, A., Krimmer, R. and Draheim, D. (2022) 'Building global societies on collective intelligence: challenges and opportunities', *Digital Government: Research and Practice*, 3(4), pp. 1–6.

Walsh, J.P. and Ungson, G.R. (1991) 'Organizational memory', *Academy of Management Review*, 16(1), pp. 57–91.

Weick, K.E. (1995) *Sensemaking in organizations*. Thousand Oaks, CA: Sage.

Woolley, A.W. (2025) 'Generative AI and collaboration: opportunities for cultivating collective intelligence', *Journal of Organization Design*, pp. 1-6.

Woolley, A.W., Chabris, C.F., Pentland, A., Hashmi, N. and Malone, T.W. (2010) 'Evidence for a collective intelligence factor in the performance of human groups', *Science*, 330(6004), pp. 686–688.

Woolley, A.W., Gupta, P. and Glikson, E. (2023) ‘Using AI to enhance collective intelligence in virtual teams: augmenting cognition with technology to help teams adapt to complexity’, in *Handbook of virtual work*. Cheltenham: Edward Elgar Publishing, pp. 67–88.

Xiao, R. (2024) ‘Four development stages of collective intelligence’, *Frontiers of Information Technology & Electronic Engineering*, 25(7), pp. 903–916.

Zhang, T.J., Yang, Y., Huang, Y., Lu, S., Schölkopf, B. and Jin, Z. (2025) ‘Collective intelligence: on the promise and reality of multi-agent systems for AI-driven scientific discovery’, *Preprints.org*, posted August 2025