

From Documents to Decisions: AI-Augmented Knowledge Management for Cycling Advocacy

A Technical White Paper on the Cyclox Platform

Gub McNicoll
15 March 2026

Abstract

Non-profit advocacy organisations operate under persistent resource constraints yet must engage with complex policy processes that demand deep institutional knowledge. This paper describes **Cyclox**, an open-source platform that combines a structured ontological knowledge base with a suite of specialised AI agents to support a cycling advocacy campaign based in Oxford, UK. I present the platform's architecture through the lens of **knowledge management theory**, examining how explicit and tacit organisational knowledge is captured, structured, and operationalised through human–AI collaboration. The paper discusses the design of a domain-specific ontology that encodes the organisation's positions, evidence base, and conceptual vocabulary; a multi-tier AI agent system that grounds its outputs in this ontology; and the governance mechanisms that maintain knowledge quality while enabling AI-assisted augmentation. Drawing on both classical knowledge management frameworks (Nonaka and Takeuchi, 1995; Davenport and Prusak, 1998) and contemporary research on human–AI collaboration (Jarrahi et al., 2023; Raisch and Krakowski, 2021), I argue that large language models, when grounded in organisational knowledge structures, can function as active participants in knowledge creation processes — a claim with significant implications for organisational epistemology, the democratisation of knowledge work in civil society, and the transferability of the approach to other advocacy domains.

Table 1: The Seven Knowledge Agents

Agent	Knowledge Function	Model Tier	SECI Phase
Ontology Q&A;	Institutional memory retrieval	Advanced	Application
Evidence Finder	Document synthesis & citation	Standard	Combination
Summariser	Triage & relevance filtering	Standard	Combination
Proposal Analyser	Policy evaluation & alignment	Premium	Combination

Position Checker	Consistency assurance	Premium	Application
Response Drafter	Formal advocacy output	Advanced	Externalisation
Blog Writer	Public communication	Premium	Externalisation

1. Introduction

Cycling advocacy in the UK is carried out predominantly by volunteer-led organisations that must respond to planning applications, transport consultations, and policy proposals with limited time and expertise. Cyclox, an Oxford-based cycling campaign, faces a challenge common to many such groups: its institutional knowledge is distributed across hundreds of consultation responses, position papers, meeting notes, and the memories of long-serving members. When a new planning application arrives, the relevant organisational positions may exist but are not easily retrievable.

This paper presents the Cyclox platform as a case study in **AI-augmented knowledge management** for civil-society organisations. The platform's contribution is not merely technical; it represents an attempt to address fundamental questions in organisational epistemology: How can an organisation's collective knowledge be made explicit without losing the nuance of context? How can AI systems be constrained to operate within an organisation's established positions rather than generating plausible but unauthorised views? And how can knowledge governance be maintained when AI agents are actively suggesting additions to the knowledge base?

I situate the platform within the theoretical framework of Nonaka and Takeuchi's knowledge creation model (1995), extending it to consider how large language models function as a new kind of actor in the SECI (Socialisation, Externalisation, Combination, Internalisation) knowledge spiral. The platform's architecture is designed not to replace human knowledge workers but to lower the barriers to knowledge retrieval, synthesis, and application — enabling a small volunteer organisation to engage with policy processes at a level that would otherwise require professional staff.

Recent advances in generative AI have prompted a renewed wave of scholarship on how organisations create, share, and apply knowledge. Von Krogh (2018) argues that AI presents opportunities for 'phenomenon-based theorizing' in management research — that is, AI systems exhibit behaviours and capabilities so qualitatively different from prior technologies that they cannot be adequately explained by simply extending existing management theories. Instead, the distinctive ways in which AI generates, retrieves, and recombines knowledge demand new theoretical frameworks built from careful observation of what these systems actually do in organisational contexts.

Jarrahi et al. (2023) go further, proposing that AI should be understood as a **knowledge management partner** rather than a tool, capable of participating in knowledge creation, curation, and transfer. Their framework identifies a qualitative shift: where previous knowledge management technologies (databases, wikis, search engines) could store and retrieve explicit knowledge, AI partners can *synthesise* across knowledge artefacts — combining evidence from multiple documents, relating new information to existing organisational positions, and producing

contextualised outputs that reflect an understanding of how different pieces of knowledge relate to one another. The Cyclox platform provides a concrete illustration of this shift. When its proposal analyser cross-references a 40-page planning document against dozens of ontology positions and produces a structured alignment assessment, it is performing a knowledge synthesis task that no prior technology could accomplish: not merely retrieving relevant documents (as a search engine would) or matching keywords (as a database query would), but interpreting the semantic content of a proposal in light of the organisation's established stance and producing a reasoned, citation-backed evaluation. Similarly, when the ontology suggestion mechanism identifies a concept that appears across multiple documents but has not yet been codified in the knowledge base, it is performing a form of knowledge *creation* — surfacing latent organisational knowledge that existed implicitly in the document corpus but had never been made explicit. These are the types of phenomena that demand the new theoretical frameworks Von Krogh calls for: an AI system that does not merely store or transmit knowledge but actively participates in its organisation and application.

This paper takes up that challenge, using the Cyclox platform as a concrete case through which to examine how classical KM theory and contemporary AI capabilities interact in practice.

2. Theoretical Foundations

2.1 The Knowledge Management Challenge in Advocacy

Davenport and Prusak (1998) define knowledge as 'a fluid mix of framed experience, values, contextual information, and expert insight that provides a framework for evaluating and incorporating new experiences and information.' This definition captures precisely the resource that cycling advocacy depends upon: not raw data about traffic volumes or accident statistics, but the contextualised understanding of how those facts relate to the organisation's values, prior positions, and strategic objectives.

In Nonaka and Takeuchi's SECI model, organisational knowledge creation proceeds through four modes: **Socialisation** (tacit to tacit, e.g. experienced campaigners mentoring newcomers), **Externalisation** (tacit to explicit, e.g. writing a position paper), **Combination** (explicit to explicit, e.g. synthesising multiple evidence sources into a consultation response), and **Internalisation** (explicit to tacit, e.g. a volunteer learning the organisation's stance by reading its outputs).

The Cyclox platform intervenes primarily in the **Combination** and **Externalisation** phases. Its AI agents perform combination by synthesising evidence from documents, ontology concepts, and positional data into structured outputs. Its ontology suggestion mechanism supports externalisation by prompting the capture of new concepts when the AI detects gaps in the knowledge base.

2.2 Ontology as Organisational Memory

The choice of an ontology — rather than a simple tag system or folder structure — as the primary knowledge representation is grounded in research on organisational memory systems (Walsh and Ungson, 1991). An ontology captures not merely *what* the organisation knows but the *relationships* between concepts: that 'protected cycle lanes' *require* 'physical separation', that 'Dutch-style roundabouts' *mitigate* 'junction risk', and that Cyclox holds a 'strong position' on both.

This relational structure enables a form of **institutional reasoning**: when an AI agent analyses a planning proposal that mentions a shared-use path, it can traverse the ontology to find that Cyclox's position on shared-use infrastructure is nuanced — acceptable in low-traffic contexts but strongly opposed on arterial routes. The ontology thus functions as what Weick (1995) terms a 'sensemaking' device, providing the interpretive framework through which new information is evaluated.

2.3 Contemporary Perspectives: AI as Knowledge Partner

The emergence of large language models (LLMs) has prompted a fundamental reappraisal of the relationship between AI and organisational knowledge. Jarrahi et al. (2023) propose a framework in which AI is not merely a knowledge *tool* but a knowledge *partner*, capable of participating in each phase of knowledge management: creation (identifying patterns humans might miss), curation (organising and structuring knowledge artefacts), sharing (translating specialist knowledge into accessible forms), and application (deploying knowledge in context-specific tasks). The Cyclox platform instantiates this partnership model: its agents create (ontology suggestions), curate (evidence synthesis), share (blog writing, summaries), and apply (proposal analysis, position checking) organisational knowledge.

Raisch and Krakowski (2021) identify a critical tension in AI deployment: the **augmentation–automation paradox**. Organisations that initially deploy AI to augment human capabilities frequently drift toward automation as the technology proves reliable — yet the value of augmentation depends on maintaining meaningful human involvement. The Cyclox platform addresses this paradox through its governance architecture: the ontology suggestion workflow ensures that AI-generated knowledge claims must pass through human review, and the explicit grounding of all agent outputs in organisational data prevents the system from operating autonomously.

Berente et al. (2021) argue that managing AI in organisations requires attention to four dimensions: agency (the degree of autonomous action granted to AI), transparency (the ability to understand AI outputs), representational fidelity (how well AI models correspond to reality), and speed of adaptation. The Cyclox platform's design choices map directly onto these dimensions: constrained agency (agents cannot modify the ontology without approval), transparent provenance (citation-backed outputs with cost tracking), representational fidelity (the ontology as an auditable knowledge model), and adaptive learning (the suggestion-review loop enables the knowledge base to evolve).

3. Platform Architecture

The platform comprises four integrated subsystems, each addressing a distinct phase of the knowledge management lifecycle.

3.1 The Document Repository

The document repository ingests consultation responses, position papers, council proposals, government guidance, and campaign materials. Each document undergoes a processing pipeline: text extraction (supporting PDF, DOCX, and plain text), chunking into semantically coherent segments, and embedding via OpenAI's text-embedding-3-small model. The resulting

1,536-dimensional vectors are stored in PostgreSQL using the pgvector extension, enabling cosine-similarity search across the corpus.

This architecture implements the **retrieval-augmented generation** (RAG) pattern described by Lewis et al. (2020), which has become foundational to knowledge-grounded AI systems. Rather than relying solely on the language model's parametric knowledge (its training data), RAG systems retrieve relevant passages from an external corpus and present them alongside the user's query, ensuring that the model's outputs are anchored in verifiable sources. In the Cyclox context, RAG ensures that agent responses reflect the organisation's documents rather than generic internet-derived knowledge.

This architecture supports two modes of retrieval: **semantic search** (finding passages conceptually related to a query, even when the exact terminology differs) and **keyword search** (precise term matching for known phrases or document titles). The hybrid approach ensures that both exploratory and targeted queries are well served.

3.2 The Ontology Knowledge Graph

The ontology is stored as a directed graph in PostgreSQL (concepts and typed relationships), with an in-memory NetworkX cache rebuilt at startup for fast traversal. Each concept carries: a canonical name, a definition, the organisation's position (if any), a position-strength indicator (strong, moderate, or emerging), and links to supporting evidence in the document repository.

Critically, the ontology is not a static artefact. The platform includes an **ontology suggestion mechanism**: when an AI agent encounters a concept during analysis that does not exist in the knowledge base, it can propose an addition. These suggestions enter a review queue for human approval, creating a feedback loop between AI-assisted knowledge discovery and human governance. This design reflects Alavi and Leidner's (2001) observation that effective knowledge management systems must support not just storage and retrieval but knowledge *creation*.

3.3 The Agent System

The AI capability is delivered through a system of seven specialised agents (termed 'skills'), each designed for a specific knowledge task. All agents share a common persona — 'Spokes' — and are grounded in the same ontology and document base, but differ in their analytical approach, output format, and computational requirements.

The skill architecture implements what Malone et al. (2018) describe as 'superminds' — systems where human and machine intelligence are combined to achieve outcomes neither could accomplish alone. The agents do not generate positions from their training data; they retrieve and synthesise the organisation's own documented positions. When positions do not exist, the agents explicitly state this — and may suggest that the organisation develop one, triggering the ontology suggestion workflow.

This agentic pattern — where LLMs are equipped with structured tools for querying external knowledge sources and are given specific roles within a collaborative workflow — reflects what Shrestha, Ben-Menahem and von Krogh (2021) identify as a shift from **centralised** to **distributed** decision-making architectures in AI-augmented organisations. Each skill functions as a specialised decision-support node, routing queries through domain-specific analytical lenses while sharing a common knowledge substrate.

3.4 The Model Tier System

A distinctive feature of the platform is its **cost-aware model routing** system. Tasks are classified into three tiers — Standard, Advanced, and Premium — each mapped to an appropriate language model. Simple queries and summarisation tasks use a smaller, more efficient model; complex analytical tasks such as proposal analysis and position checking use more capable (and more expensive) models. This design reflects the operational reality of non-profit organisations: AI capabilities must be accessible but financially sustainable. Per-session cost tracking provides transparency and helps prevent resource wastage.

4. The Seven Knowledge Agents

Each agent implements a specific knowledge management function. I describe them here in terms of their epistemic role rather than their technical implementation.

4.1 Ontology Q&A — Acts as the institutional memory interface. When a user asks 'What is Cyclox's position on 20mph zones?', this agent queries the ontology graph, retrieves the relevant concept and its relationships, and presents the organisation's documented position with its strength indicator. It performs what Alavi and Leidner (2001) call 'knowledge application' — making stored knowledge available for decision-making.

4.2 Evidence Finder — Performs retrieval-augmented search across the document repository. Rather than simply returning matching passages, it synthesises findings into a structured evidence brief with precise citations. This agent supports the 'combination' phase of the SECI model, bringing together evidence from multiple sources.

4.3 Summariser — Condenses documents for a Cyclox audience, highlighting aspects most relevant to cycling advocacy. The agent cross-references the ontology to identify which organisational concepts appear in the source material, enabling rapid triage of incoming documents.

4.4 Proposal Analyser — The most complex agent, designed for evaluating planning applications and transport proposals. It identifies design choices within a proposal, classifies each against Cyclox's known positions (aligned, partially aligned, conflicts, or no position held), and provides a structured analysis with an evidence hierarchy distinguishing between design standards, research evidence, best practice, and organisational opinion. This agent exemplifies what Simon (1947) termed 'bounded rationality' augmentation — extending the cognitive capacity of volunteer reviewers who lack the time to cross-reference every policy position manually.

4.5 Position Checker — Validates external statements against the organisation's positions with a structured verdict: Aligned, Partially Aligned, Conflicts, or No Position Held. Each verdict includes a confidence level, cited positions, and reasoning. This agent serves a quality-assurance function, helping ensure that the organisation's public communications remain consistent with its documented stances.

4.6 Response Drafter — Generates formal advocacy responses to consultations, planning applications, and policy proposals. The agent structures outputs according to established advocacy practice: acknowledgement, evidentiary arguments, policy references, and constructive recommendations. It draws on both the document repository and the ontology to ground every claim in the organisation's evidence base.

4.7 Blog Writer — Creates accessible articles for public communication. Unlike the formal response drafter, this agent adapts tone and length to user specifications, supporting the organisation's public engagement goals. It grounds content in ontology positions and evidence while maintaining an engaging, human-readable style.

5. Knowledge Governance and Quality

A persistent challenge in AI-augmented knowledge systems is maintaining the boundary between what the organisation *knows* (its documented positions) and what the AI can *generate* (plausible but potentially unauthorised claims). The platform addresses this through several mechanisms.

Grounding constraints. All agents are instructed to base their outputs on the ontology and document repository. When relevant positions exist, agents cite them explicitly. When they do not, agents state that the organisation has no formal position — they do not fabricate one. This design principle distinguishes the platform from general-purpose chatbot deployments where the AI speaks with its own 'voice'.

Position strength indicators. The ontology assigns each position a strength level: strong (deeply held, well-evidenced), moderate (established but with acknowledged nuance), or emerging (recently adopted, subject to evolution). Agents present these indicators alongside positions, enabling users to judge the epistemic weight of each claim.

The suggestion-review workflow. When an agent detects a gap in the ontology — a concept referenced in documents but not yet captured — it can propose an addition. These suggestions are queued for human review, preserving the principle that the ontology represents the *organisation's* knowledge, not the AI's. This mechanism implements what Grant (1996) describes as the 'integration' of specialist knowledge: the AI identifies candidates for inclusion; the humans decide what meets the threshold of organisational endorsement.

Transparent provenance. Every agent output includes citations linking back to specific documents, ontology concepts, and relationships. Session transcripts record which skill was used, which model tier, and the estimated cost — creating a full audit trail that supports accountability.

Epistemic boundaries in human–AI systems. Berente et al. (2021) observe that a core challenge in managing AI is determining the scope of *agency* afforded to AI systems. The CycloX platform draws a clear epistemic boundary: agents can *retrieve*, *synthesise*, and *present* organisational knowledge, and they can *suggest* additions — but they cannot unilaterally *create* or *modify* the organisation's knowledge base. This mirrors what Raisch and Krakowski (2021) term a 'structured augmentation' approach, where the division of labour between human and AI is deliberately designed rather than allowed to emerge organically.

6. Implications for Knowledge Management Theory

6.1 AI as a SECI Participant

The platform's architecture suggests that large language models can function as active participants in Nonaka and Takeuchi's knowledge spiral, not merely as tools. The ontology suggestion mechanism is a form of AI-assisted **Externalisation**: the model reads organisational documents

(which contain tacit knowledge made partially explicit through writing) and proposes structured ontology entries that make the knowledge fully explicit and navigable. The proposal analyser performs **Combination** at a scale and speed that would not be feasible for a volunteer organisation: cross-referencing a 40-page planning document against dozens of ontology positions in seconds rather than hours.

However, the platform deliberately excludes AI from the **Socialisation** and **Internalisation** phases, which remain fundamentally human. The warm, approachable persona of 'Spokes' is a design choice that facilitates interaction, not a claim that the AI understands or shares the organisation's values.

6.2 Democratisation of Analytical Capacity

A significant implication of the platform is the **redistribution of analytical capacity**. Previously, only experienced campaigners with deep knowledge of Cyclox's historical positions could draft effective consultation responses or quickly identify whether a proposal conflicted with organisational policy. The platform enables newer volunteers to access the same institutional knowledge, mediated by AI agents that retrieve and synthesise it on demand. This represents what Zuboff (1988) anticipated as 'informating' — using technology not just to automate tasks but to make information (and in this case, organisational knowledge) more broadly accessible within the organisation.

6.3 The Ontology as a Boundary Object

Star and Griesemer's (1989) concept of the **boundary object** — an artefact that is shared across communities of practice and maintains enough common identity to be recognisable to all, while being adaptable enough to serve each community's needs — maps well to the Cyclox ontology. The same ontology graph serves as: a reference for AI agents grounding their outputs; a navigable knowledge map for volunteers exploring the organisation's positions; and a governance tool for administrators maintaining the organisation's knowledge base. The ontology viewer, with its interactive graph visualisation, makes this boundary object tangible and explorable.

6.4 Human–AI Augmentation: Evidence and Implications

Emerging empirical research supports the theoretical claims for AI-augmented knowledge work. Noy and Zhang (2023) demonstrated in a randomised experiment that access to generative AI reduced the time required for professional writing tasks by 40% while improving output quality, with the largest gains accruing to less experienced workers. Brynjolfsson, Li and Raymond (2023) found similar patterns in a large-scale study of customer support agents: AI assistance improved productivity by 14% on average, with novice workers benefiting most dramatically — a 34% improvement. These findings resonate directly with the Cyclox use case, where newer volunteers stand to gain the most from AI-mediated access to institutional knowledge.

The platform's design embodies what Raisch and Krakowski (2021) call **structured augmentation**: a deliberate architecture in which AI capabilities complement rather than replace human judgement. The agents perform computationally intensive tasks (cross-referencing dozens of positions against a planning document) while humans retain authority over strategic decisions (which positions to hold, how to weigh competing concerns, and when to evolve the organisation's stance). This structured division of labour avoids the 'automation drift' that Raisch and Krakowski warn against — the tendency for augmentation systems to gradually displace the human role they were designed to

support.

6.5 Knowledge Graphs Meet Large Language Models

The Cyclox architecture represents a specific instance of a broader trend in AI research: the integration of structured knowledge graphs with large language models (Pan et al., 2024). Knowledge graphs provide the factual grounding, consistency, and explainability that LLMs lack; LLMs provide the natural language understanding, generative capability, and flexibility that knowledge graphs lack. In the Cyclox platform, the ontology (a knowledge graph) supplies authoritative organisational facts, while the LLM translates those facts into context-appropriate outputs — consultation responses, evidence briefs, blog posts. This synergy exemplifies what Pan et al. describe as 'KG-enhanced LLM applications', where the knowledge graph serves as both a retrieval source and a constraint mechanism.

7. Transferability and Future Directions

While developed for cycling advocacy, the platform's architecture is domain-agnostic. Any non-profit organisation that maintains positions, engages with policy consultations, and needs to synthesise evidence from a document corpus could adopt the same pattern. Environmental campaigns, tenants' rights organisations, patient advocacy groups, and community planning bodies face structurally similar knowledge management challenges.

The key transferable components are: the ontology-grounded agent pattern (ensuring AI outputs reflect organisational rather than generic knowledge); the suggestion-review governance loop (balancing AI knowledge discovery with human editorial control); and the cost-aware model routing (making AI affordable for resource-constrained organisations).

Future development directions include: fine-tuning the intent classification system using collected interaction data; adding longitudinal position tracking to capture how organisational stances evolve over time; and exploring federated knowledge sharing between allied advocacy organisations — allowing, for example, cycling campaigns across different cities to share ontology structures while maintaining local positions.

The convergence of knowledge management and AI research also opens avenues for empirical investigation. Following the experimental methodology of Noy and Zhang (2023) and Brynjolfsson, Li and Raymond (2023), future work could measure the platform's impact on volunteer productivity, response quality, and knowledge retention. Such studies would contribute to the growing evidence base on how AI augmentation affects knowledge work in practice, extending findings from commercial settings to the under-studied civil society context.

8. Conclusion

The Cyclox platform demonstrates that AI-augmented knowledge management is not the exclusive preserve of large enterprises with dedicated knowledge engineering teams. By combining a relatively simple but well-structured ontology with specialised AI agents that are constrained to operate within the organisation's documented knowledge, a small advocacy group can achieve a level of institutional memory and analytical capability that was previously out of reach.

The theoretical contribution of this work lies in the observation that large language models, when appropriately constrained and grounded, can function as active participants in organisational knowledge processes — not replacing human judgement but extending human capacity for knowledge retrieval, synthesis, and application. The practical contribution is a working, open-source platform that embodies these principles and is deployable by any organisation with modest technical resources.

I invite researchers and practitioners to examine, extend, and challenge the platform's assumptions. Knowledge management in civil society is an underexplored domain; the intersection with AI capabilities makes it a fertile area for both theoretical inquiry and practical impact.

References

- Alavi, M. and Leidner, D.E., 2001. 'Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues.' *MIS Quarterly*, 25(1), pp.107-136.
- Berente, N., Gu, B., Recker, J. and Santhanam, R., 2021. 'Managing artificial intelligence.' *MIS Quarterly*, 45(3), pp.1433-1450.
- Brynjolfsson, E., Li, D. and Raymond, L.R., 2023. 'Generative AI at work.' *NBER Working Paper*, No. 31161.
- Davenport, T.H. and Prusak, L., 1998. *Working Knowledge: How Organizations Manage What They Know*. Harvard Business School Press.
- Grant, R.M., 1996. 'Toward a knowledge-based theory of the firm.' *Strategic Management Journal*, 17(S2), pp.109-122.
- Jarrahi, M.H., Askay, D., Eshraghi, A. and Smith, P., 2023. 'Artificial intelligence and knowledge management: A partnership between human and AI.' *Business Horizons*, 66(1), pp.87-98.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S. and Kiela, D., 2020. 'Retrieval-augmented generation for knowledge-intensive NLP tasks.' *Advances in Neural Information Processing Systems*, 33, pp.9459-9474.
- Malone, T.W., Laubacher, R. and Dellarocas, C., 2018. 'Harnessing crowds: Mapping the genome of collective intelligence.' *MIT Sloan Management Review*.
- Nonaka, I. and Takeuchi, H., 1995. *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford University Press.
- Noy, S. and Zhang, W., 2023. 'Experimental evidence on the productivity effects of generative artificial intelligence.' *Science*, 381(6654), pp.187-192.
- Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J. and Wu, X., 2024. 'Unifying large language models and knowledge graphs: A roadmap.' *IEEE Transactions on Knowledge and Data Engineering*, 36(7), pp.3580-3599.
- Raisch, S. and Krakowski, S., 2021. 'Artificial intelligence and management: The automation–augmentation paradox.' *Academy of Management Review*, 46(1), pp.192-210.
- Shrestha, Y.R., Ben-Menahem, S.M. and von Krogh, G., 2021. 'Organizational decision-making structures in the age of artificial intelligence.' *California Management Review*, 63(4), pp.66-83.
- Simon, H.A., 1947. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization*. Macmillan.
- Star, S.L. and Griesemer, J.R., 1989. 'Institutional ecology, translations and boundary objects.' *Social Studies of Science*, 19(3), pp.387-420.
- von Krogh, G., 2018. 'Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing.' *Academy of Management Discoveries*, 4(4), pp.404-409.
- 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*. Sage.
- Zuboff, S., 1988. *In the Age of the Smart Machine: The Future of Work and Power*. Basic Books.