

A Review of Artificial Intelligence and Knowledge Management in the 21st Century: Pros and Cons

¹*Nwankwe, Dimie, ²V.I.E Anireh, ³Boye, Aziboledia F.

^{1,2}Department of Computer Science, Rivers State University, Rivers State, Port Harcourt, Nigeria

³Department of Computer Science, Informatics & Cybersecurity, Otuoke, Bayelsa State, Nigeria

*Corresponding Author's E-mail: dimie.nwankwe@gmail.com

Co-Author E-mails: ²anireh.ike@ust.edu.ng; ³boyeaf@fuotuoike.edu.ng

Abstract: Artificial intelligence (AI) and knowledge management (KM) have converged more rapidly in the first quarter of the twenty-first century than at any time since both fields first emerged. Early rule-based expert systems sought to codify scarce expertise; machine learning then automated the classification, search and recommendation of organisational knowledge; and large language models and other generative AI systems have begun to build, summarise and apply knowledge in natural language at scale since 2022. This paper reviews the literature on AI-enabled KM from scientific, technical and policy perspectives, mainly from 2000 to 2026. This review traces the development of AI for KM in three stages, examines how it has been applied to each of these processes, and critically analyses the advantages and disadvantages in literature. AI has shown that it can significantly increase employees' productivity (especially among those with less experience in the field), bring knowledge back into the system, help re-use knowledge that would otherwise be lost, and lead to faster decisions. It also shows that AI brings serious risks: fake or unverified outputs, poor understanding of tacit and contextual knowledge, embedded bias, opacity, privacy and security exposure, over-reliance and deskilling, homogenisation of organisational knowledge, workforce displacement, environmental costs and widening digital divides - especially in developing countries. The paper argues that AI is a net value when KM is being done, but not just the technology, but how organisations make decisions about it. It proposes a human-centred AI-KM integration framework with five layers (strategy, knowledge and data foundations, technology architecture, people and culture, governance and assurance), identifies research gaps and suggests methods for practitioners, policymakers and researchers.

Keywords: Artificial intelligence; Knowledge management; Large language models; Knowledge sharing; AI governance; Human-AI collaboration; Digital transformation; Developing economies.

1. INTRODUCTION

It is known that knowledge has long been regarded as the most strategic resource [30]. Those who are capable of creating, storing, sharing and applying knowledge better than their competitors are better able to innovate, adjust to change and remain successful [4]; [49]. Knowledge management (KM) was a formal discipline introduced in the 1990s to provide the strategies, processes and technologies to do so [17]. KM has been based on information technology since it started in the 1990s: from document repositories and intranets to groupware, enterprise portals and expert directories [4]; [33]. Artificial intelligence (AI) (that is, systems that perceive, reason, learn and act in a way that would be considered intelligent if done by humans (Russell & Norvig, 2020)) has, however, had a longer, less harmonious relationship with KM. In the 1970s and 1980s, rule-based expert systems like MYCIN aimed to capture the knowledge of specialists in explicit rules [15]. By the turn of the

21st century, AI and KM were already partners for good because AI and KM are both concerned with knowledge and with representation, acquisition, reasoning and use [44]. But AI's contribution in organisational KM is limited to text mining, search and specialised decision support applications for the first decade of the 21st century.

That situation has changed greatly in the last 20 years. Deep learning [41], transformer architecture [67] and large language models (LLMs) trained on vast text corpora (Brown *et al.*, 2020) have helped to build powerful systems for summarising documents, answering questions, drafting reports, developing software code, and communicating in natural language. These capabilities touch every stage of the knowledge life cycle at every stage. Generative AI is now a technology that is changing the way knowledge gets generated, stored, transferred and applied in organisations [5]; [8]. At the same time, experimental and field studies have begun to quantify the

productivity and the hidden costs of AI for knowledge work [13]; [19]; [42]; [50]. Despite the rapid growth of research at the intersection of AI and KM, the literature is fragmented. Technical research is on algorithms and architecture; management studies are on adoption, productivity and organisational design; ethics and policy studies on fairness, accountability and regulation. A number of reviews have examined aspects of the picture: the benefits and implementation of AI in KM [64], the human–AI partnership in KM [37], the KM implications of generative AI [5] and the implementation challenges of AI in KM [57]. Few studies have come together, in a balanced way, to give an account of AI-enabled KM in the twenty-first century and how organisations can capture the former while containing the latter. This is particularly true for developing economies where infrastructure, skills and regulatory capacity are very different from those in the countries where most of the work has been done.

1.2 Objectives and Research Questions

The aim of this paper is to provide a critical, integrative review of AI and KM in the twenty-first century, with explicit attention to the pros and cons of AI adoption for organisational knowledge. The review addresses four research questions:

- i. RQ1: How has the role of AI in knowledge management evolved since 2000?
- ii. RQ2: How is AI currently applied across the core knowledge-management processes of creation, storage and retrieval, transfer and application?
- iii. RQ3: What advantages and disadvantages of AI for knowledge management are documented in the literature?
- iv. RQ4: What governance arrangements and research directions can help organisations realise the benefits of AI for knowledge management while mitigating its risks?

1.2 Structure of the Paper

The remainder of the paper is organised as follows. Section 2 sets out the conceptual foundations of knowledge, KM and AI. Section 3 describes the review methodology. Section 4 traces the evolution of AI in KM across three phases. Section 5 maps AI applications onto the KM processes. Sections 6 and 7 examine the advantages and disadvantages respectively, and Section 8 synthesises them. Section 9 proposes a human-centred AI–KM integration framework. Section 10 considers implications for developing economies, with reference to Nigeria. Section 11 identifies research gaps, Section 12 states the contribution to knowledge, and Section 13 concludes with recommendations.

II. CONCEPTUAL FOUNDATIONS

2.1 Data, Information and Knowledge

A useful starting point is the distinction between data, information and knowledge. Data are raw symbols or observations; information is organised data and given meaning in context; knowledge is information combined with experience, judgement and values so that it can guide action [2]; [17]. The distinction matters for AI because most AI systems operate directly on data and information, and their capacity to embody genuine knowledge, in the sense of justified, contextual and actionable understanding, is precisely what is contested in the literature [59].

2.2 Tacit and explicit knowledge

[55] Observation that people know more than they are able to articulate introduced the idea of tacit knowledge: personal, experience-based know-how that is difficult to articulate or formalise. Explicit knowledge, by contrast, can be expressed in words, numbers, procedures and documents. [48] and [49] built on this distinction to propose the SECI model, in which organisational knowledge is created through four modes of conversion: socialisation (tacit to tacit), externalisation (tacit to explicit), combination (explicit to explicit) and internalisation (explicit to tacit). Most information technology, including most AI, is strongest in the combination mode, where existing explicit knowledge is collected, reorganised and recombined. Whether AI can meaningfully support socialisation, externalisation and internalisation is a central question in this review.

2.3 Knowledge Management Processes and Strategies

[4] Conceptualised organisations as knowledge systems composed of four interdependent processes: knowledge creation, knowledge storage and retrieval, knowledge transfer and knowledge application. This framework remains among the most widely used in KM research and has recently been applied directly to AI and generative AI [5]; [38]; [57]. It is adopted here as the organising lens. [33] distinguished two broad KM strategies. A codification strategy stores knowledge in databases where it can be accessed and reused by anyone in the organisation, whereas a personalisation strategy relies on direct person-to-person contact and uses technology mainly to help people find and communicate with each other. Communities of practice, groups of people who share a concern and learn through regular interaction, are a key vehicle of personalisation [69]. Knowledge transfer between units is in turn a source of competitive advantage [6]. As later sections show, AI has

historically reinforced codification but is increasingly being used to support personalisation as well. The international standard ISO 30401:2018 now specifies requirements for KM systems, defining KM as a management discipline focused on the systematic and holistic approach to optimising the use of knowledge [35]. This standard provides a reference point for integrating AI into an organisation's KM system rather than treating AI tools as standalone technologies.

2.4 Artificial Intelligence: From Symbolic Systems to Foundation Models

AI encompasses a family of approaches. Symbolic AI represents knowledge explicitly, for instance in rules, logic or ontologies, and reasons over it; expert systems are its best-known organisational application [15]; [58]. Machine learning (ML) instead infers patterns from data, and deep learning uses multi-layered neural networks to learn hierarchical representations directly from raw inputs such as text, images and speech [41]. The transformer architecture, based on attention mechanisms, made it practical to train very large language models [67], and scaling such models produced systems able to perform many language tasks from a few examples or instructions [12]. [11] coined the term "foundation models" to describe large models trained on broad data that can be adapted to a wide range of downstream tasks, noting that their generality creates both opportunities and systemic risks. Two further technologies are especially relevant to KM. Knowledge graphs represent entities and their relationships in a structured, machine-readable network and are widely used for enterprise search, recommendation and question answering [34]. Retrieval-augmented generation (RAG) combines a language model with a retrieval component that fetches relevant documents from an external knowledge store, so that generated answers can be grounded in, and traced to, specific sources [28]; [43]. Researchers increasingly argue that the combination of LLMs and knowledge graphs can offset the weaknesses of each, pairing the fluency and generality of the former with the structure and verifiability of the latter [51].

2.5 The Convergence of AI and KM

AI and KM converge because both concern knowledge, but they approach it differently. KM is primarily a managerial and social discipline concerned with people, processes and culture; AI is primarily a technical discipline concerned with computational representation and inference. [44] argued that AI could supply KM with tools for knowledge acquisition, representation and discovery, while KM could supply AI with organisational purpose. Later work reframed the relationship as a

partnership in which AI handles complexity and scale while humans contribute judgement, context and the ability to deal with uncertainty and ambiguity [37]; [38]. [10] characterise contemporary AI as the frontier of computational advancement, marked by increasing autonomy, learning and inscrutability, three properties that make it simultaneously powerful and difficult to manage. These properties frame both the advantages and the disadvantages examined in this review.

III. RESULTS AND DISCUSSION

3.1 Review Design

This study adopts an integrative literature review design. Integrative reviews are suited to mature and emerging topics that span several disciplines, because they critically synthesise representative literature in order to generate new frameworks and perspectives rather than to estimate a single pooled effect [62]; [65]. The review is concept-centric rather than author-centric, organising the literature around the concepts and processes under study [68]

3.2 Sources and search strategy

Literature was identified through searches of major scholarly databases and indexes, including Scopus, Web of Science, IEEE Xplore, the ACM Digital Library, ScienceDirect and Google Scholar, supplemented by the websites of standards bodies, regulators and research institutions. Search strings combined terms for AI (for example "artificial intelligence", "machine learning", "deep learning", "generative AI", "large language model", "expert system", "knowledge graph", "retrieval-augmented generation") with terms for KM (for example "knowledge management", "knowledge sharing", "knowledge creation", "knowledge transfer", "organisational memory", "tacit knowledge"). Backward and forward citation tracking from seminal works [4]; [44]; [49] was used to identify further sources.

3.3 Inclusion and Exclusion Criteria

Sources were included if they (a) were peer-reviewed journal articles, peer-reviewed conference papers, scholarly books, working papers from recognised research institutions, international standards or official policy and regulatory documents; (b) addressed AI in relation to knowledge, knowledge work or KM processes, or provided foundational theory for either field; and (c) were published mainly between 2000 and 2026, with earlier seminal works retained where they underpin current theory. Opinion pieces without evidential

grounding, vendor marketing material and sources whose bibliographic details could not be verified were excluded. Every source cited in this paper was checked against its publisher, journal or institutional record.

3.4 Analysis and Synthesis

The selected literature was analysed thematically. Each source was coded by period, AI technology, KM process addressed (creation, storage and retrieval, transfer, application), type of evidence (conceptual, experimental, field, survey, technical) and nature of the findings (benefit, risk, mitigation). The coded material was then synthesised in three steps: a chronological synthesis addressing RQ1, a process-based mapping addressing RQ2, and a comparative pros-and-cons synthesis addressing RQ3. The governance framework and research agenda addressing RQ4 were developed inductively from the patterns in the synthesis.

3.5 Limitations of the method

Integrative reviews involve judgement in source selection and interpretation, and their findings are less reproducible than those of protocol-driven systematic reviews. The literature on generative AI is also evolving quickly, so some findings may be superseded. The evidence base is concentrated in North America, Europe and East Asia; findings may not transfer directly to other contexts, a limitation that Section 10 addresses explicitly. Finally, a number of influential recent studies are working papers or conference proceedings rather than journal articles; these are identified as such in the reference list.

3.6 The Evolution of AI in Knowledge Management in the 21st Century

The literature suggests three broads, overlapping phases in the evolution of AI for KM since 2000. Table 1 summarises them.

Table 1: Three phases of AI in knowledge management, 2000 to the present

Dimension	Phase 1: Codification and search (c. 2000–2010)	Phase 2: Data-driven intelligence (c. 2011–2021)	Phase 3: Generative and agentic AI (c. 2022–present)
Dominant AI approach	Rule-based and expert systems, early text mining, ontologies and the semantic web	Machine learning, deep learning, natural language processing, knowledge graphs	Large language models, multimodal foundation models, retrieval-augmented generation, AI agents
Typical KM tools	Document repositories, enterprise portals, keyword search, expert directories	Intelligent search, recommender systems, chatbots, automated tagging and classification, predictive analytics	Conversational assistants over enterprise content, automated summarisation and drafting, copilots embedded in work tools
Primary KM process supported	Storage and retrieval of explicit knowledge	Retrieval, transfer and application of explicit knowledge at scale	All four processes, including content creation
View of the human	Knowledge source to be codified	User and validator of algorithmic outputs	Collaborator, editor and steward of AI-generated knowledge
Characteristic limitation	Knowledge-acquisition bottleneck; brittle rules; low user adoption	Need for large labelled datasets; opacity; narrow task scope	Fabricated outputs; difficulty of verification; governance, privacy and security risks
Key references	Liebowitz (2001); Alavi & Leidner (2001); Hansen et al. (1999)	LeCun et al. (2015); Davenport & Ronanki (2018); Hogan et al. (2022)	Brown et al. (2020); Lewis et al. (2020); Alavi et al. (2024)

Source: Author's synthesis of the reviewed literature.

3.6.1 Phase 1: Codification and search (c. 2000–2010)

At the start of the century, KM practice was dominated by codification. Organisations invested heavily in repositories, portals and "lessons learned" databases to capture explicit knowledge [33]. AI contributed mainly through rule-based reasoning, case-based reasoning, intelligent agents and early text mining, techniques that [44] saw as the bridge between the two fields. Semantic technologies and ontologies attempted to make organisational content machine-interpretable. The persistent difficulty was what expert-systems researchers had called the knowledge-acquisition bottleneck: extracting tacit expertise from people and encoding it in explicit rules was slow, expensive and fragile [15]. Many KM systems of this period suffered from low adoption because they demanded effort from contributors without delivering commensurate value to users [17].

3.6.2 Phase 2: Data-driven intelligence (c. 2011–2021)

The second phase was driven by big data, cheaper computing and deep learning [41]. Instead of encoding knowledge by hand, organisations increasingly learned patterns from data. [18], drawing on a study of 152 projects, classified business AI into three types: process automation, cognitive insight (pattern detection and prediction from data) and cognitive engagement (interaction with employees and customers through natural language). Each has KM implications: automation of document handling, analytics that generate new insights from operational data, and chatbots that answer employees' questions. Knowledge graphs became the backbone of enterprise search and recommendation [34]. Organisational scholars began to examine how AI changes decision-making structures, distinguishing

situations in which decisions should be delegated to algorithms, retained by humans or shared between them [60]. The second phase also brought sustained attention to algorithmic bias, opacity and accountability [46]; [26], concerns that remain central today.

3.6.3 Phase 3: Generative and agentic AI (c. 2022–present)

The third phase began when LLM-based assistants became publicly available at scale in late 2022. Unlike earlier tools, generative AI can produce new text, code and images in response to natural-language instructions, which allows it to participate directly in knowledge creation rather than merely retrieving what humans have already produced [5]; [23]. The authors [8] observe that generative AI holds both promises and perils for knowledge and creative work, because the same capabilities that accelerate content production also make it harder to judge the provenance and accuracy of that content. Enterprises have responded by connecting LLMs to their own document stores through retrieval-augmented generation, so that answers are grounded in approved organisational sources [28]; [43], and by embedding AI assistants in everyday productivity software. More recently, AI agents capable of planning and executing multi-step tasks have begun to appear, raising further questions about autonomy, oversight and accountability. Regulators have also responded: the European Union adopted the AI Act in 2024 as the first comprehensive, risk-based legal framework for AI [24].

IV. AI APPLICATIONS ACROSS THE KNOWLEDGE MANAGEMENT PROCESSES

This section maps AI applications to the four KM processes identified by [4]. Table 2 provides an overview.

Table 2: AI applications mapped to knowledge-management processes

KM process	Representative AI applications	Primary SECI mode supported	Illustrative sources
Knowledge creation	Pattern discovery in operational and sensor data; idea generation and drafting with generative AI; simulation and predictive modelling; literature and patent analysis	Combination; partial support for externalisation	Davenport & Ronanki (2018); Alavi et al. (2024); Jarrahi et al. (2023)

KM process	Representative AI applications	Primary SECI mode supported	Illustrative sources
Knowledge storage and retrieval	Automated classification and tagging; semantic and conversational search; knowledge graphs; retrieval-augmented generation over enterprise content; summarization	Combination	Hogan et al. (2022); Lewis et al. (2020); Pan et al. (2024)
Knowledge transfer and sharing	Expert finding and recommendation; chatbots and virtual assistants; machine translation; personalised learning and onboarding	Combination; partial support for internalisation	Jarrahi (2018); Taherdoost & Madanchian (2023); Rezaei (2025)
Knowledge application	Decision support and prediction; copilots for writing, analysis and coding; process automation; agents executing workflows	Internalisation (via use and learning)	Brynjolfsson et al. (2025); Noy & Zhang (2023); Peng et al. (2023)

Source: Author's synthesis based on Alavi and Leidner (2001) and Nonaka (1994).

4.1 Knowledge creation

AI supports knowledge creation in two distinct ways. Analytical AI discovers patterns, correlations and anomalies in large datasets that human analysts would struggle to detect, thereby generating new explicit knowledge, for example predictive models of equipment failure, demand or risk [18]. Generative AI supports creation more directly by producing drafts, alternatives and syntheses that humans can refine. [5] argue that generative AI most clearly strengthens the combination mode of knowledge creation, because it recombines vast amounts of existing explicit knowledge, and can also support internalisation by helping individuals learn from that knowledge. Its support for socialisation, which depends on shared experience among people, is much weaker. [38] therefore describes AI's role in knowledge creation as augmenting, rather than replacing, human sense-making.

4.2 Knowledge storage and retrieval

Storage and retrieval is the process in which AI has the longest track record. Machine learning automates the classification, tagging and de-duplication of documents, reducing the manual effort that undermined earlier repositories (Taherdoost & Madanchian, 2023). Knowledge graphs link people, documents, projects and concepts so that related knowledge can be found by navigation as well as by search (Hogan *et al.*, 2022). Semantic search interprets the meaning of a

query rather than matching keywords. Retrieval-augmented generation goes further by allowing users to ask questions in natural language and receive synthesised answers that cite the underlying sources [43]; [28]. Because LLMs alone tend to produce plausible but unsupported statements, grounding them in curated organisational knowledge stores, and increasingly in knowledgegraphs, is regarded as essential for enterprise use [51].

4.3 Knowledge transfer and sharing

AI facilitates knowledge transfer by connecting people with relevant content and with each other. Expert-finding systems infer who knows what from documents, messages and project records, supporting the personalisation strategy of KM [33]. Conversational assistants answer routine questions immediately, freeing experts for more complex problems. Machine translation reduces language barriers in multinational organisations. Adaptive learning systems tailor training to individual needs, supporting the internalisation of explicit knowledge. [37] emphasises that AI can extend the reach of human knowledge sharing but cannot replace the trust, reciprocity and shared context on which effective sharing within communities of practice depends [69]. [5] add a caution: if employees consult AI instead of colleagues, the interpersonal exchanges through which tacit knowledge is transferred may weaken.

4.4 Knowledge Application

Knowledge has value only when it is applied to decisions and actions [4]. AI supports application through decision-support systems, predictive analytics, process automation and, most recently, copilots embedded in the tools that knowledge workers use. Field and experimental evidence on these applications is reviewed in Section 6.1. Shrestha et al. (2019) note that the appropriate allocation of decisions between humans and AI depends on factors such as the specificity of the decision space, the interpretability of the decision process, the size of the alternative set, the speed required and replicability. [56] show that automation and augmentation are interdependent rather than mutually exclusive, and that organisations that over-emphasise either can harm both performance and society.

V. THE PROS: ADVANTAGES OF AI FOR KNOWLEDGE MANAGEMENT

5.1 Productivity and efficiency in knowledge work

The most robust evidence for the benefits of AI in knowledge work concerns productivity. In a preregistered online experiment with college-educated professionals performing occupation-specific writing tasks, [50] found that access to ChatGPT reduced the time taken by about 40 per cent and raised the quality of output by about 18 per cent. In a large field study of more than 5,000 customer-support agents, [13] found that a generative AI assistant increased the number of issues resolved per hour by 14 per cent on average. In a controlled experiment, software developers with access to an AI pair programmer completed a programming task 55.8 per cent faster than those without [54]. In a field experiment with 758 consultants at a global management consulting firm, [11] found that, for tasks within the capability frontier of the AI, consultants using GPT-4 completed 12.2 per cent more tasks, worked 25.1 per cent faster and produced results of more than 40 per cent higher quality than a control group. These gains arise from the automation of routine information processing, faster retrieval and synthesis of knowledge, and the reduced effort of producing first drafts.

5.2 Levelling Up Less-Experienced Workers

A notable and consistent finding is that AI benefits less-experienced and lower-performing workers most. [13] report that novice and lower-skilled agents improved their productivity by roughly a third, while the most experienced agents saw little gain, and they interpret this as evidence that the AI tool disseminated the tacit best practices of high performers, embedded in its training data, to newer workers. [50] found that

AI compressed the productivity distribution, raising weaker writers more than stronger ones, and [19] observed the largest improvements among consultants with lower baseline performance. From a KM perspective, this is significant: AI can act as a channel through which accumulated organisational knowledge is transferred to people who have not yet acquired it through experience, accelerating onboarding and reducing the dependence of quality on individual seniority.

5.3 Improved Access, Retrieval and Reuse of Knowledge

Organisations routinely fail to reuse what they already know because relevant knowledge is scattered across systems, formats and languages. AI-enabled search, knowledge graphs and RAG-based assistants reduce the cost of finding and synthesising that knowledge [34]; [43]. Natural-language interfaces lower the skill barrier to querying complex knowledge bases, which helps to democratise access across the workforce [5]. Automated summarisation makes long technical documents, contracts and reports usable in less time. These capabilities address a long-standing weakness of the codification strategy, namely that stored knowledge was too hard to find to be worth storing [17].

5.4 Preservation of Organisational Memory and Expertise

Many industries face the loss of expertise as experienced employees retire or move on. AI can help to preserve organisational memory by indexing and interlinking project records, incident reports and technical documentation, and by making them queryable long after their authors have left. Although AI cannot fully capture tacit knowledge (see Section 7.2), it can support externalisation by prompting experts to articulate their reasoning, transcribing and structuring interviews and debriefs, and drawing out patterns from the records of past decisions [38]; [64]. In this way, AI can partially alleviate the knowledge-acquisition bottleneck that limited earlier expert systems.

5.5 Better and Faster Decision Support

AI-based analytics can detect weak signals, forecast outcomes and evaluate alternatives at a speed and scale beyond human capacity. When combined with human judgement, these capabilities can improve decision quality, particularly where decisions involve large volumes of structured data [60]. [37] characterises the ideal arrangement as human-AI symbiosis: AI contributes analytical power in conditions of complexity, while humans contribute holistic, intuitive judgement in conditions of uncertainty and equivocality. [18] found that the most successful

corporate AI projects were those that augmented, rather than replaced, human capabilities.

5.6 Innovation and Knowledge Creation

By recombining existing knowledge in novel ways, generating alternatives and accelerating experimentation, AI can stimulate innovation [5]; [8]. In research-intensive domains, AI-enabled literature analysis helps knowledge workers identify connections across large bodies of publications. In product and process development, generative tools allow more options to be explored at lower cost. Brynjolfsson and McAfee (2014) argue that the greatest economic value of digital technologies arises from such recombinant innovation, in which existing ideas are combined to create new ones.

5.7 Personalisation, Scalability and Continuous Availability

AI systems can tailor knowledge delivery to the role, experience and preferences of each user, and they are available continuously across time zones. A single AI-enabled knowledge service can support thousands of employees simultaneously, which allows organisations to scale expert support in ways that would be prohibitively expensive with human staff alone [64]. [13] also report improvements in the experience of work, including more courteous customer interactions and fewer requests for escalation to a manager, suggesting that well-designed AI support can reduce some of the strain of knowledge work.

VI. THE CONS: DISADVANTAGES AND RISKS OF AI FOR KNOWLEDGE MANAGEMENT

6.1 Hallucination, accuracy and verification burden

Generative AI systems can produce fluent statements that are factually wrong or unsupported by any source, a phenomenon commonly termed hallucination [39]. Because such outputs are linguistically plausible, they can contaminate organisational knowledge bases, mislead decision-makers and propagate errors at scale. [9] warn that language models reproduce the statistical patterns of their training data without any grounding in meaning or communicative intent, which explains why their confidence is not a reliable indicator of their accuracy. [19] demonstrated the practical consequence: for a task deliberately chosen to lie outside the AI's capability frontier, consultants using AI were 19 percentage points less likely to produce a correct solution than those working without it. The capability frontier is "jagged", meaning that tasks of apparently similar difficulty can fall on different sides of it, so users cannot easily anticipate when the AI

will fail. As a result, AI shifts effort from producing knowledge to verifying it, and that verification burden is often underestimated [42].

6.2 Limited Grasp of Tacit and Contextual Knowledge

The second fundamental limitation concerns tacit knowledge. [59] argue that AI is constrained by the tacit dimension: much of the knowledge that matters most in organisations is embodied, situated and social, and cannot be fully captured in the data on which AI systems are trained. AI can process what has been written down, but not what experts know without being able to say [55]. AI-generated recommendations may therefore omit crucial local context, cultural nuance and practical judgement. Over-emphasis on AI may also reinforce a codification bias, in which organisations value what can be digitised and neglect the personal networks and communities through which tacit knowledge is created and shared [33]; [69].

6.3 Bias, Fairness and Discrimination

AI systems learn from historical data and can therefore reproduce or amplify the biases embedded in that data. [45] catalogue numerous sources of bias, from unrepresentative sampling and historical prejudice to measurement and evaluation biases and show that biased models can systematically disadvantage particular groups. In a KM context, biased systems may privilege the knowledge and perspectives of dominant groups, under-represent minority viewpoints, and embed discriminatory assumptions into decisions on recruitment, promotion, credit or service provision [46]. [9] note that very large web-scraped training corpora tend to over-represent hegemonic viewpoints, which may be particularly problematic for organisations operating in non-Western contexts.

6.4 Opacity, explainability and accountability

Many high-performing AI models, particularly deep neural networks, are "black boxes" whose internal reasoning cannot be readily inspected [3]. Lack of explainability undermines the KM goal of making knowledge understandable and transferable, because a recommendation that cannot be explained cannot easily be learned from, challenged or improved. It also complicates accountability: when an AI-informed decision causes harm, it may be unclear who is responsible [46]. Explainable AI techniques have been developed to address this problem, but there is still no consensus on how interpretability should be defined and evaluated [21]. [10] identify inscrutability

as one of the defining management challenges of contemporary AI.

6.5 Privacy, Security and Intellectual Property

AI-enabled KM systems aggregate large volumes of potentially sensitive organisational and personal information, creating new exposures. [16] showed that it is possible to extract verbatim sequences of training data, including personal information, from a large language model, demonstrating that models can memorise and leak the data on which they were trained. LLM-integrated applications are also vulnerable to indirect prompt injections, in which malicious instructions hidden in retrieved documents or web pages manipulate the system's behaviour [31]. For KM, this means that connecting AI to enterprise knowledge stores can expose confidential information, trade secrets or personal data if access controls and security measures are inadequate. Employees who paste proprietary content into public AI services may also leak intellectual property inadvertently. [70] situates such concerns within a broader critique of data extraction in the digital economy. [57], in a Delphi study and confirmatory factor analysis across the four KM processes, security and privacy found to be prominent technological challenges in all of them.

6.6 Over-Reliance, Deskilling and Erosion of Critical Thinking

Automation has long been known to produce ironies: the more a system is automated, the less practice human operators get, yet the more critical their skills become when the automation fails [7]. [52] distinguished misuse of automation, including over-reliance, from disuse and abuse. These insights apply directly to knowledge work. In a survey of 319 knowledge workers, [42] found that higher confidence in generative AI was associated with less critical thinking, whereas higher self-confidence was associated with more, and that the nature of critical thinking shifted towards verifying information, integrating responses and stewarding tasks. The risk for KM is that organisations may gradually lose the expertise required to evaluate AI outputs, to handle novel situations and to create genuine new knowledge. [4] note a related risk that junior employees may miss the learning that comes from performing entry-level tasks themselves.

6.7 Trust Calibration and Algorithm Aversion

Effective human-AI collaboration requires appropriately calibrated trust. [29] show that trust in AI depends on factors such as tangibility, transparency, reliability and immediacy, and

that cognitive and emotional trust can develop differently. Too much trust leads to over-reliance; too little leads to algorithm aversion, in which people abandon algorithms after seeing them make errors, even when the algorithms outperform humans [20]. Both extremes reduce the value of AI-enabled KM, and achieving calibrated trust requires transparency, training and experience.

6.8 Homogenisation of Knowledge and Model Collapse

When many people rely on the same AI systems, their output may converge, reducing the diversity of ideas on which innovation depends. A more technical concern is model collapse: [61] demonstrated that training generative models indiscriminately on content produced by earlier models causes irreversible defects, with the tails of the original data distribution disappearing. As AI-generated text increasingly populates the web and organisational repositories, there is a risk that future systems, and the organisations that use them, will lose access to rare but valuable knowledge. For KM, this underscores the importance of preserving and labelling human-generated, authoritative sources.

6.9 Workforce Displacement and Job Insecurity

AI can automate tasks that were previously performed by knowledge workers. [27] estimated that around 47 per cent of US employment was at substantial risk of computerisation, although this estimate has been widely debated. [1] emphasise that the net effect of automation depends on whether the displacement of labour from existing tasks is offset by the creation of new tasks in which humans have a comparative advantage. In the KM literature, fear of job loss is a significant barrier to knowledge sharing: employees who believe that codifying their expertise will make them replaceable may withhold it. [57] found that lack of job security was loaded strongly as an ethical challenge across all four KM processes.

6.10 Environmental and Economic Costs

Training and operating large AI models consumes substantial energy and computing resources. [63] drew attention to the energy and carbon costs of training large natural language processing models, and [53] estimated that training GPT-3 consumed about 1,287 MWh of electricity and emitted roughly 552 tonnes of carbon dioxide equivalent. Such costs are significant for organisations with sustainability commitments. AI-enabled KM also requires investment in data infrastructure, integration, licensing, skills and ongoing maintenance, and these

costs may not be recovered if the underlying knowledge is poorly organised or if adoption is weak [57]; [64].

6.11 Implementation, Governance and Regulatory Complexity

Finally, the integration of AI into KM raises organisational and governance challenges: poor data quality, legacy systems, skills shortages, resistance to change, budget constraints and the need to comply with an evolving regulatory landscape (Dwivedi *et al.*, 2021; Rezaei, 2025). More than 80 sets of AI ethics guidelines had been published by 2019, converging on principles such as transparency, justice and fairness, non-maleficence, responsibility and privacy, but differing substantially in how

those principles should be interpreted and implemented [40]. Organisations operating across jurisdictions must now navigate binding regulation such as the EU AI Act [24], voluntary frameworks such as the NIST AI Risk Management Framework [47], management-system standards such as ISO/IEC 42001:2023 [36], and global normative instruments such as the UNESCO Recommendation on the Ethics of Artificial Intelligence [66].

VII. SYNTHESIS: WEIGHING THE PROS AND CONS

Table 3 consolidates the advantages and disadvantages identified in Sections 6 and 7 and pairs each with the mitigation mechanisms most frequently recommended in the literature.

Table 3: Summary of the pros and cons of AI for knowledge management and associated mitigation mechanisms

Dimension	Pros (advantages)	Cons (disadvantages)	Mitigation mechanisms
Productivity	Faster drafting, retrieval and analysis; large gains on in-frontier tasks	Lower accuracy on out-of-frontier tasks; hidden verification costs	Task-level assessment of AI suitability; human review of critical outputs
Knowledge quality	Synthesis of large bodies of knowledge; source-grounded answers via RAG	Hallucination; contamination of knowledge bases; model collapse	Retrieval grounding with citations; curated authoritative sources; provenance labelling
Tacit knowledge	Support for externalisation and learning; transfer of best practice to novices	Weak capture of embodied and contextual knowledge; codification bias	Sustain communities of practice and mentoring alongside AI
Equity and fairness	Democratised access to expertise; levelling up of less-experienced staff	Embedded bias; under-representation of minority and local perspectives	Bias testing; diverse and local data; impact assessments
Transparency	Explicit links between answers and sources in well-designed systems	Black-box reasoning; unclear accountability	Explainability tools; documentation; clear decision rights
Security and privacy	AI-assisted threat detection and access monitoring	Data leakage; prompt injection; IP exposure	Access controls; data classification; secure deployment; staff guidance
Human capability	Augmentation of judgement; more time for complex work	Over-reliance; deskilling; reduced critical thinking	AI literacy; designs that prompt reflection; deliberate skills practice
Workforce	New roles and tasks; improved experience of work	Displacement; job insecurity; reduced willingness to share knowledge	Reskilling; augmentation-oriented job design; transparent change management

Dimension	Pros (advantages)	Cons (disadvantages)	Mitigation mechanisms
Sustainability and cost	Scalable, continuously available knowledge services	Energy use and emissions; infrastructure and licensing costs	Right-sized models; efficiency measures; business-case discipline
Governance	Opportunity to formalise KM and data governance	Regulatory complexity; fragmented ethical guidance	ISO 30401, ISO/IEC 42001, NIST AI RMF and applicable regulation

Source: Author's synthesis of Sections 6 and 7.

Three cross-cutting insights emerge from this synthesis. First, many of the advantages and disadvantages are two sides of the same capability. The fluency that makes generative AI useful for drafting also makes its errors persuasive; the ability to transfer best practice to novices is the same ability that can erode the learning they would otherwise gain through practice; and the scale that allows one system to serve thousands of users also allows one error or bias to propagate to all of them. This duality is consistent with the automation–augmentation paradox described by [56]. Second, the balance between pros and cons depends heavily on the task. The evidence from [19] shows that the same tool can improve or worsen performance depending on whether the task lies inside or outside the AI's capability frontier. Organisational value therefore depends on the ability of managers and workers to judge which knowledge tasks should be delegated to AI, which should be shared, and which should remain with humans [60]. Third, the decisive variable is governance rather than technology. Where AI is grounded in

curated knowledge, embedded in clear processes, overseen by people with the skills and confidence to challenge it, and governed by explicit policies, the documented benefits tend to dominate. Where AI is deployed as a standalone tool without these conditions, the documented risks become more likely. This insight motivates the framework presented in the next section.

VIII. A HUMAN-CENTRED AI-KM INTEGRATION FRAMEWORK

Drawing on the synthesis, this paper proposes a human-centred framework for integrating AI into organisational KM. The framework has five interdependent layers, shown in Figure 1, and is anchored in the four KM processes of [4]. It is intended to be compatible with ISO 30401:2018 for KM systems [36], ISO/IEC 42001:2023 for AI management systems (ISO/IEC, 2023) and the Govern, Map, Measure and Manage functions of the NIST AI Risk Management Framework [47].




Figure 1: Human-centred AI-KM integration framework (author's proposal)

8.1 Layer 1: Strategy and Value

AI initiatives in KM should begin with knowledge problems rather than with technology. Organisations should identify where knowledge gaps, losses or bottlenecks most constrain performance, such as onboarding, retention of retiring expertise, reuse of project lessons or response to customer queries, and define measurable objectives for each. [18] recommend a portfolio of projects that augment existing capabilities, starting with achievable applications and scaling successful ones. The strategy should explicitly state the intended balance between automation and augmentation [56].

8.2 Layer 2: Knowledge and Data Foundations

AI is only as good as the knowledge it draws upon. Organisations need curated, authoritative and up-to-date knowledge sources; metadata and taxonomies; knowledge graphs that link entities and relationships; and clear ownership of content. Human-generated authoritative sources should be preserved and labelled to guard against contamination and model collapse [61]. Data quality, provenance, classification and retention policies are prerequisites for safe AI use [57].

8.3 Layer 3: Technology Architecture

For most enterprise KM purposes, literature favours architectures that ground generative models in organisational knowledge through retrieval-augmented generation, ideally combined with knowledge graphs, and that cite sources so that users can verify answers [28]; [43]; [51]. Security by design is essential, including role-based access controls that respect existing permissions, protection against prompt injections, logging and monitoring, and secure deployment environments [31]. Model selection should balance capability against cost and environmental impact [53].

8.4 Layer 4: People and Culture

People remain the ultimate source and judge of organisational knowledge. Organisations should build AI literacy so that employees understand what AI can and cannot do, how to verify its outputs and when not to use it [42]. Workflows should preserve opportunities for learning by doing, mentoring and communities of practice so that tacit knowledge continues to be created and transferred (Wenger, 1998). Job design should emphasise augmentation and reskilling to address fears of displacement, which otherwise discourage knowledge sharing [1]; [57]. Trust should be calibrated through transparency about system performance and limitations [29].

8.5 Layer 5: Governance and Assurance

Governance should define decision rights, accountability for AI-assisted decisions, acceptable-use policies and risk-assessment procedures. Organisations should conduct impact assessments for higher risk uses, test for bias, document models and data sources, monitor performance over time and provide channels for users to report errors. These activities correspond closely to the Govern, Map, Measure and Manage functions of the NIST AI Risk Management Framework [47] and to the requirements of ISO/IEC 42001 [35], and they should reflect ethical principles such as beneficence, non-maleficence, autonomy, justice and explicability [26]. Where applicable, compliance with binding regulations such as the EU AI Act must also be assured [24].

8.6 The Human-In-The-Loop Principle Across the Knowledge Processes

Across all layers, the framework applies a human-in-the-loop principle calibrated to the stakes of each knowledge process. In knowledge creation, AI generates options and humans select and refine them. In storage and retrieval, AI organises and retrieves while humans curate and validate authoritative content. In transfer, AI extends reach while humans sustain relationships through which tacit knowledge flows. In application, AI recommends while humans decide, at least wherever decisions are consequential, novel or ethically sensitive. This allocation operationalises the human-AI partnership advocated by [38].

IX. IMPLICATIONS FOR DEVELOPING ECONOMIES: THE NIGERIAN CONTEXT

Most of the empirical evidence reviewed above comes from high-income economies. Developing economies face distinctive opportunities and constraints. On the opportunity side, AI-enabled KM could help organisations compensate for shortages of specialist expertise, preserve knowledge in sectors with high staff turnover, and make knowledge available in local languages. The finding that AI benefits less-experienced workers most [13]; [50] suggests that it could accelerate capability development where experienced professionals are scarce. On the constraint side, [32] identify limited digital infrastructure, data scarcity, skills gaps and weak regulatory frameworks as key challenges for AI deployment in Africa. Models trained predominantly in Western data may perform poorly on local languages, contexts and knowledge, and may embed perspectives that do not reflect local realities [9]. Unreliable electricity and connectivity, the cost of computing and dependence on foreign

technology providers can widen rather than narrow the digital divide. Nigeria illustrates both sides. In 2024 the Federal Ministry of Communications, Innovation and Digital Economy, working with the National Information Technology Development Agency (NITDA) and its National Centre for Artificial Intelligence and Robotics, published a draft National Artificial Intelligence Strategy that sets out principles such as responsible conduct, inclusivity, innovation, sustainability, collaboration and data ethics, and that emphasises the value of locally developed solutions [25]. For Nigerian organisations, including those in knowledge-intensive sectors such as oil and gas, banking, telecommunications and public administration, the framework proposed in Section 9 implies several priorities: investing in the curation and digitisation of local organisational knowledge; building AI literacy and data skills; adopting secure, cost-appropriate architectures; aligning with the national strategy and emerging data-protection requirements; and ensuring that AI augments rather than displaces scarce local expertise.

X. RESEARCH GAPS AND FUTURE DIRECTIONS

The review identifies several gaps that merit further research. Longitudinal evidence: Most studies of AI in knowledge work measure short-term effects. Longitudinal research is needed on how sustained AI use affects expertise development, organisational memory, innovation and the quality of knowledge bases over years rather than weeks. Tacit knowledge: There is limited empirical work on whether and how AI can support the socialisation and externalisation of tacit knowledge, and on how organisations can protect tacit knowledge flows as AI use grows [17]. Measurement of KM outcomes: The field lacks validated metrics for the quality, reuse and value of AI-mediated knowledge, beyond task-level productivity. Knowledge of homogenisation: The organisational consequences of convergent AI-assisted outputs and of model collapse for long-term knowledge diversity remain largely unexplored [61]. Developing-economy contexts: Empirical studies of AI-enabled KM in Africa and other under-represented regions, including work on local languages and indigenous knowledge, are scarce [32]. Agentic AI: As AI agents begin to execute multi-step tasks, research is needed on their implications for accountability, oversight and the allocation of knowledge work between humans and machines. Governance effectiveness: Little is known about which governance mechanisms, standards and regulatory approaches are most effective in practice at reducing AI-related KM risks while preserving benefits. Sector-specific studies: Industry-level evidence is needed from safety-critical and asset-intensive sectors, such as energy, health care

and infrastructure, where errors in applied knowledge carry high consequences.

XI. CONTRIBUTION TO KNOWLEDGE

This paper makes five contributions. First, it provides an integrated account of the evolution of AI in KM across the twenty-first century, organised into three phases (codification and search, data-driven intelligence, and generative and agentic AI), which clarifies how the human role in KM has shifted from knowledge source to validator to collaborator and steward. Second, it maps contemporary AI applications systematically onto the four KM processes and the SECI modes, showing that AI is strongest in combination and weakest in socialisation. Third, it offers a balanced, evidence-based synthesis of the pros and cons of AI for KM, pairing each with mitigation mechanisms, and identifies three cross-cutting insights: the duality of AI capabilities, the task dependence of AI value, and the primacy of governance. Fourth, it proposes a five-layer human-centred AI-KM integration framework that is aligned with ISO 30401, ISO/IEC 42001 and the NIST AI Risk Management Framework, providing practitioners with an actionable structure. Fifth, it extends the discussion to developing economies, using Nigeria as an illustration, and sets out a research agenda that addresses gaps in longitudinal, tacit-knowledge, measurement, contextual and governance research.

XII. CONCLUSION AND RECOMMENDATIONS

12.1 Conclusion

Over the first quarter of the twenty-first century, AI has moved from the margins to the center of knowledge management. Expert systems promised to capture expertise but foundered on the difficulty of codifying it; machine learning made knowledge easier to find and apply at scale; and generative AI now participates directly in creating knowledge. The evidence reviewed here shows that AI can make knowledge work substantially faster and better, spread expertise to less-experienced workers, improve access to organisational knowledge and support better decisions. It also shows that AI can generate convincing falsehoods, miss tacit and contextual knowledge, embed bias, obscure accountability, expose sensitive information, erode human skills and critical thinking, homogenise knowledge, threaten jobs and impose environmental and financial costs. The central conclusion of this review is that AI is neither inherently beneficial nor inherently harmful to organisational knowledge. Its net value depends on the purposes for which it is used, the quality of the knowledge on which it draws, the capabilities and judgement of the people who use it,

and the governance arrangements within which it operates. Organisations that treat AI as a partner within a deliberately designed, human-centred KM system are most likely to realise its benefits while containing its risks.

12.2 Recommendations

For organisational leaders and KM practitioners: Start from knowledge problems and measurable objectives and adopt a portfolio approach that favours augmentation over wholesale automation. Secondly, invest in the curation, ownership and quality of authoritative knowledge sources before scaling AI, and ground generative AI in these sources with visible citations. Thirdly, assess each knowledge task for AI suitability and require human review where outputs are consequential, novel or ethically sensitive. Fourthly, Build AI literacy across the workforce, including skills in verification, and design workflows that preserve learning, mentoring and communities of practice. In another we address job-security concerns openly through reskilling and transparent management changes so that employees continue to share their knowledge. In another, we Implement security by design, data classification and acceptable-use policies to protect confidential and personal information. For policymakers and regulators: Develop proportionate, risk-based AI governance that is interoperable with international frameworks and translate national AI strategies into practical guidance for organisations. Also, invest in digital infrastructure, local-language data resources and AI skills, particularly in developing economies, to prevent the widening of digital divides. Encourage transparency about the energy use and environmental impact of AI systems. For researchers: Pursue longitudinal, field-based and sector-specific studies of AI-enabled KM, with particular attention to tacit knowledge, expertise development and knowledge diversity. Develop validated measures of the quality and value of AI-mediated knowledge. Expand empirical research in under-represented regions, including Africa.

REFERENCES

- [1] Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- [2] Ackoff, R. L. (1989). From data to wisdom. *Journal of Applied Systems Analysis*, 16, 3–9.
- [3] Adadi, A., & Berrada, M. (2018). Peeking inside the black box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6, 52138–52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
- [4] Alavi, M., & Leidner, D. E. (2001). Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues. *MIS Quarterly*, 25(1), 107–136. <https://doi.org/10.2307/3250961>
- [5] Alavi, M., Leidner, D. E., & Mousavi, R. (2024). A knowledge management perspective of generative artificial intelligence. *Journal of the Association for Information Systems*, 25(1), 1–12. <https://doi.org/10.17705/1jais.00859>
- [6] Argote, L., & Ingram, P. (2000). Knowledge transfer: A basis for competitive advantage in firms. *Organizational Behavior and Human Decision Processes*, 82(1), 150–169. <https://doi.org/10.1006/obhd.2000.2893>
- [7] Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- [8] Benbya, H., Strich, F., & Tamm, T. (2024). Navigating generative artificial intelligence promises and perils for knowledge and creative work. *Journal of the Association for Information Systems*, 25(1), 23–36. <https://doi.org/10.17705/1jais.00861>
- [9] Bender, E. M., Gebru, T., McMillan-Major, A., & 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* (pp. 610–623). ACM. <https://doi.org/10.1145/3442188.3445922>
- [10] Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433–1450.
- [11] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., . . . Liang, P. (2021). On the opportunities and risks of foundation models (arXiv:2108.07258). arXiv. <https://doi.org/10.48550/arXiv.2108.07258>
- [12] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., . . . Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
- [13] Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of*

- Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- [14] Brynjolfsson, E., & McAfee, A. (2014). The second machine age: Work, progress, and prosperity in a time of brilliant technologies. W. W. Norton.
- [15] Buchanan, B. G., & Shortliffe, E. H. (Eds.). (1984). Rule-based expert systems: The MYCIN experiments of the Stanford Heuristic Programming Project. Addison-Wesley.
- [16] Carlini, N., Tramèr, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, Ú., Oprea, A., & Raffel, C. (2021). Extracting training data from large language models. In *Proceedings of the 30th USENIX Security Symposium* (pp. 2633–2650). USENIX Association.
- [17] Davenport, T. H., & Prusak, L. (1998). Working knowledge: How organizations manage what they know. Harvard Business School Press.
- [18] Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
- [19] Dell'Acqua, F., McFowland, E., III, Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality (Harvard Business School Technology & Operations Management Unit Working Paper No. 24-013). Harvard Business School. <https://doi.org/10.2139/ssrn.4573321>
- [20] Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126. <https://doi.org/10.1037/xge0000033>
- [21] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning (arXiv:1702.08608). arXiv. <https://doi.org/10.48550/arXiv.1702.08608>
- [22] Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P. V., Janssen, M., Jones, P., Kar, A. K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., . . . Williams, M. D. (2021). Artificial intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57, 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
- [23] Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., . . . Wright, R. (2023). "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
- [24] European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L, 2024/1689. <http://data.europa.eu/eli/reg/2024/1689/oj>
- [25] Federal Ministry of Communications, Innovation and Digital Economy. (2024). National artificial intelligence strategy (Draft). National Information Technology Development Agency & National Centre for Artificial Intelligence and Robotics. https://ncair.nitda.gov.ng/wp-content/uploads/2024/08/National-AI-Strategy_01082024-copy.pdf
- [26] Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- [27] Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- [28] Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, M., & Wang, H. (2023). Retrieval-augmented generation for large language models: A survey (arXiv:2312.10997). arXiv. <https://doi.org/10.48550/arXiv.2312.10997>
- [29] Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627–660. <https://doi.org/10.5465/annals.2018.0057>
- [30] Grant, R. M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(S2), 109–122. <https://doi.org/10.1002/smj.4250171110>

- [31] Greshake, K., Abdelnabi, S., Mishra, S., Endres, C., Holz, T., & Fritz, M. (2023). Not what you've signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection. In *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security* (pp. 79–90). ACM. <https://doi.org/10.1145/3605764.3623985>
- [32] Gwagwa, A., Kraemer-Mbula, E., Rizk, N., Rutenberg, I., & de Beer, J. (2020). Artificial intelligence (AI) deployments in Africa: Benefits, challenges and policy dimensions. *The African Journal of Information and Communication*, 26, 1–28. <https://doi.org/10.23962/10539/30361>
- [33] Hansen, M. T., Nohria, N., & Tierney, T. (1999). What's your strategy for managing knowledge? *Harvard Business Review*, 77(2), 106–116.
- [34] Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., de Melo, G., Gutierrez, C., Kirrane, S., Labra Gayo, J. E., Navigli, R., Neumaier, S., Ngonga Ngomo, A.-C., Polleres, A., Rashid, S. M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., & Zimmermann, A. (2022). Knowledge graphs. *ACM Computing Surveys*, 54(4), Article 71, 1–37. <https://doi.org/10.1145/3447772>
- [35] International Organization for Standardization. (2018). Knowledge management systems—Requirements (ISO Standard No. 30401:2018). ISO.
- [36] International Organization for Standardization & International Electrotechnical Commission. (2023). *Information technology—Artificial intelligence—Management system* (ISO/IEC Standard No. 42001:2023). ISO.
- [37] Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577–586. <https://doi.org/10.1016/j.bushor.2018.03.007>
- [38] Jarrahi, M. H., Askay, D., Eshraghi, A., & Smith, P. (2023). Artificial intelligence and knowledge management: A partnership between human and AI. *Business Horizons*, 66(1), 87–99. <https://doi.org/10.1016/j.bushor.2022.03.002>
- [39] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), Article 248, 1–38. <https://doi.org/10.1145/3571730>
- [40] Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- [41] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- [42] Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25) (pp. 1–22). ACM. <https://doi.org/10.1145/3706598.3713778>
- [43] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474.
- [44] Liebowitz, J. (2001). Knowledge management and its link to artificial intelligence. *Expert Systems with Applications*, 20(1), 1–6. [https://doi.org/10.1016/S0957-4174\(00\)00044-0](https://doi.org/10.1016/S0957-4174(00)00044-0)
- [45] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2022). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), Article 115, 1–35. <https://doi.org/10.1145/3457607>
- [46] Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2). <https://doi.org/10.1177/2053951716679679>
- [47] National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- [48] Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37. <https://doi.org/10.1287/orsc.5.1.14>
- [49] Nonaka, I., & Takeuchi, H. (1995). The knowledge-creating company: How Japanese companies create the dynamics of innovation. *Oxford University Press*.
- [50] Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- [51] Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J., & Wu, X. (2024). Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge*

- and *Data Engineering*, 36(7), 3580–3599. <https://doi.org/10.1109/TKDE.2024.3352100>
- [52] Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230–253. <https://doi.org/10.1518/001872097778543886>
- [53] Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). Carbon emissions and large neural network training (arXiv:2104.10350). arXiv. <https://doi.org/10.48550/arXiv.2104.10350>
- [54] Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot (arXiv:2302.06590). arXiv. <https://doi.org/10.48550/arXiv.2302.06590>
- [55] Polanyi, M. (1966). The tacit dimension. *Doubleday*.
- [56] Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://doi.org/10.5465/amr.2018.0072>
- [57] Rezaei, M. (2025). Artificial intelligence in knowledge management: Identifying and addressing the key implementation challenges. *Technological Forecasting and Social Change*, 217, 124183. <https://doi.org/10.1016/j.techfore.2025.124183>
- [58] Russell, S., & Norvig, P. (2020). Artificial intelligence: A modern approach (4th ed.). Pearson.
- [59] Sanzogni, L., Guzman, G., & Busch, P. (2017). Artificial intelligence and knowledge management: Questioning the tacit dimension. *Prometheus*, 35(1), 37–56. <https://doi.org/10.1080/08109028.2017.1364547>
- [60] Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4), 66–83. <https://doi.org/10.1177/0008125619862257>
- [61] Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. *Nature*, 631(8022), 755–759. <https://doi.org/10.1038/s41586-024-07566-y>
- [62] Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104, 333–339. <https://doi.org/10.1016/j.jbusres.2019.07.039>
- [63] Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 3645–3650). *Association for Computational Linguistics*. <https://doi.org/10.18653/v1/P19-1355>
- [64] Taherdoost, H., & Madanchian, M. (2023). Artificial intelligence and knowledge management: Impacts, benefits, and implementation. *Computers*, 12(4), 72. <https://doi.org/10.3390/computers12040072>
- [65] Torraco, R. J. (2005). Writing integrative literature reviews: Guidelines and examples. *Human Resource Development Review*, 4(3), 356–367. <https://doi.org/10.1177/1534484305278283>
- [66] UNESCO. (2021). Recommendation on the ethics of artificial intelligence. *United Nations Educational, Scientific and Cultural Organization*. <https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence>
- [67] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- [68] Webster, J., & Watson, R. T. (2002). Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly*, 26(2), xiii–xxiii.
- [69] Wenger, E. (1998). Communities of practice: Learning, meaning, and identity. *Cambridge University Press*.
- [70] Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier of power. *Public Affairs*.

Citation of this Article:

Nwankwe, Dimie, V.I.E Anireh, & Boye, Azibolelia F. (2026). A Review of Artificial Intelligence and Knowledge Management in the 21st Century: Pros and Cons. *Journal of Artificial Intelligence and Emerging Technologies (JAJET)*. 3(10), 33-49. Article DOI: <https://doi.org/10.47001/JAIET/2026.310005>

*** End of the Article ***