



# Interdisciplinary Journal of Information, Knowledge, and Management

An Official Publication  
of the Informing Science Institute  
[InformingScience.org](http://InformingScience.org)

[IJIKM.org](http://IJIKM.org)

Volume 19, 2024

## IS KNOWLEDGE MANAGEMENT (FINALLY) EXTRACTIVE? – FULLER’S ARGUMENT REVISITED IN THE AGE OF AI

Norman Mooradian

School of Information,  
San José State University,  
San José, California, USA

[norman.mooradian@sjsu.edu](mailto:norman.mooradian@sjsu.edu)

### ABSTRACT

Aim/Purpose	The rise of modern artificial intelligence (AI), in particular, machine learning (ML), has provided new opportunities and directions for knowledge management (KM). A central question for the future of KM is whether it will be dominated by an automation strategy that replaces knowledge work or whether it will support a knowledge-enablement strategy that enhances knowledge work and uplifts knowledge workers. This paper addresses this question by re-examining and updating a critical argument against KM by the sociologist of science Steve Fuller (2002), who held that KM was extractive and exploitative from its origins.
Background	This paper re-examines Fuller’s argument in light of current developments in artificial intelligence and knowledge management technologies. It reviews Fuller’s arguments in its original context wherein expert systems and knowledge engineering were influential paradigms in KM, and it then considers how the arguments put forward are given new life in light of current developments in AI and efforts to incorporate AI in the KM technical stack. The paper shows that conceptions of tacit knowledge play a key role in answering the question of whether an automating or enabling strategy will dominate. It shows that a better understanding of tacit knowledge, as reflected in more recent literature, supports an enabling vision.
Methodology	The paper uses a conceptual analysis methodology grounded in epistemology and knowledge studies. It reviews a set of historically important works in the field of knowledge management and identifies and analyzes their core concepts and conceptual structure.
Contribution	The paper shows that KM has had a faulty conception of tacit knowledge from its origins and that this conception lends credibility to an extractive vision supportive of replacement automation strategies. The paper then shows that recent

Accepting Editor Ahmad Samed Al-Adwan | Received: May 3, 2024 | Revised: July 29, August 2, August 6, 2024 | Accepted: August 8, 2024.

Cite as: Mooradian, N. (2024). Is knowledge management (finally) extractive? – Fuller’s argument revisited in the age of AI. *Interdisciplinary Journal of Information, Knowledge, and Management*, 19, Article 23.

<https://doi.org/10.28945/5354>

(CC BY-NC 4.0) This article is licensed to you under a [Creative Commons Attribution-NonCommercial 4.0 International License](https://creativecommons.org/licenses/by-nc/4.0/). When you copy and redistribute this paper in full or in part, you need to provide proper attribution to it to ensure that others can later locate this work (and to ensure that others do not accuse you of plagiarism). You may (and we encourage you to) adapt, remix, transform, and build upon the material for any non-commercial purposes. This license does not permit you to use this material for commercial purposes.

scholarship on tacit knowledge and related forms of reasoning, in particular, abduction, provide a more theoretically robust conception of tacit knowledge that supports the centrality of human knowledge and knowledge workers against replacement automation strategies. The paper provides new insights into tacit knowledge and human reasoning vis-à-vis knowledge work. It lays the foundation for KM as a field with an independent, ethically defensible approach to technology-based business strategies that can leverage AI without becoming a merely supporting field for AI.

Findings

Fuller's argument is forceful when updated with examples from current AI technologies such as deep learning (DL) (e.g., image recognition algorithms) and large language models (LLMs) such as ChatGPT.

Fuller's view that KM presupposed a specific epistemology in which knowledge can be extracted into embodied (computerized) but disembedded (decontextualized) information applies to current forms of AI, such as machine learning, as much as it does to expert systems.

Fuller's concept of expertise is narrower than necessary for the context of KM but can be expanded to other forms of knowledge work. His account of the social dynamics of expertise as professionalism can be expanded as well and fits more plausibly in corporate contexts.

The concept of tacit knowledge that has dominated the KM literature from its origins is overly simplistic and outdated. As such, it supports an extractive view of KM. More recent scholarship on tacit knowledge shows it is a complex and variegated concept. In particular, current work on tacit knowledge is developing a more theoretically robust and detailed conception of human knowledge that shows its centrality in organizations as a driver of innovation and higher-order thinking. These new understandings of tacit knowledge support a non-extractive, human enabling view of KM in relation to AI.

Recommendations for Practitioners

Practitioners can use the findings of the paper to consider ways to implement KM technologies in ways that do not neglect the importance of tacit knowledge in automation projects (which neglect often leads to failure). They should also consider how to enhance and fully leverage tacit knowledge through AI technologies and augment human knowledge.

Recommendations for Researchers

Researchers can use these findings as a conceptual framework in research concerning the impact of AI on knowledge work. In particular, the distinction between replacement and enabling technologies, and the analysis of tacit knowledge as a structural concept, can be used to categorize and analyze AI technologies relative to KM research objectives.

Impact on Society

The potential of AI on employment in the knowledge economy is a major issue in the ethics of AI literature and is widely recognized in the popular press as one of the pressing societal risks created by AI and specific types such as generative AI. This paper shows that KM, as a field of research and practice, does not need to and should not add to the risks created by automation-replacement strategies. Rather, KM has the conceptual resources to pursue a (human) knowledge enablement approach that can stand as a viable alternative to the automation-replacement vision.

Future Research

The findings of the paper suggest a number of research trajectories. They include:

Further study of tacit knowledge and its underlying cognitive mechanisms and structures in relation to knowledge work and KM objectives.

Research into different types of knowledge work and knowledge processes and the role that tacit and explicit knowledge play.

Research into the relation between KM and automation in terms of KM's history and current technical developments.

Research into how AI arguments knowledge works and how KM can provide an enabling framework.

Keywords            knowledge management, knowledge work, artificial intelligence, automation, knowledge enablement, tacit knowledge, abductive reasoning

## INTRODUCTION

---

From its inception, knowledge management was positioned as a boon to workers and associated with the concept of knowledge work. Knowledge work and the knowledge economy held out the promise of higher wages and greater professional autonomy for individuals based on their human intellectual capital (Drucker, 1988). Knowledge management practices seemed to support this vision. Not everyone, however, saw its objectives in a positive light. Steve Fuller, in particular, developed an extended critique aimed to show that knowledge management has an exploitative, extractive goal. In Fuller's (2002) view, knowledge management (KM) represents a business strategy the goal of which is to extract, codify, and automate human knowledge. Fuller saw expert systems as the core technology animating knowledge management. However, expert systems have seen a decline since the publication of his book, and his argument is not widely acknowledged.

Nevertheless, the technologies associated with knowledge management have changed since 2002 (Ribiere et al., 2021). Machine learning (ML) and deep learning (DL) have developed rapidly and pose new threats to knowledge work because of their association with hyperautomation and their integration into KM technology suites at an increasing rate (Gartner, n.d.). Most recently, Large Language Models and their interfaces (OpenAI's ChatGPT) burst onto the scene, showing the potential of AI to produce human-like text presenting encyclopedic knowledge. The power of these technologies and their unfolding integration into the KM technology stack makes Fuller's argument newly urgent. It raises the question of whether Fuller's argument was simply ahead of its time and whether he diagnosed KM accurately before its extractive symptoms were fully evident. Further, if Fuller was right, not only was KM oriented toward extraction and exploitation from its origins, but the rise of AI today should be expected to set those tendencies free. Theorists such as Ribiere et al. (2021) wonder whether the next era of KM will be "... purely techno-centric or human-techno powered ..." (p. 45). The findings of this paper will help adumbrate these two potential trajectories.

This paper offers a theoretical/conceptual re-examination of Fuller's argument in light of current developments in AI and KM technologies. It begins by reviewing the critique in its original context. It explicates the original argument as being based on an identification of KM with knowledge engineering and focused on expert systems. Further, according to Fuller, KM's vision was bolstered by epistemological assumptions operating within a social dynamic. The epistemological assumption was that knowledge consists of propositions that can be abstracted and transferred from human minds and social contexts into computer systems. The social dynamic is explained by the power relations and social capital of professionals. As a result, Fuller argued that KM could be seen as a knowledge engineering project aiming to capture expert knowledge into expert systems, which systems could be positioned in social relations to allow them to replace human knowledge workers. In other words, KM was exploitative from its origins.

After reviewing the original argument, the paper reconsiders its main premises in light of current developments in AI, in particular, machine learning (ML)/deep learning (DL), including LLMs. It shows that developments in these areas do support Fuller's original argument insofar as they succeed in performing certain expert tasks within a supporting social context. The tentative conclusion is that the updated version of Fuller's argument does have force and should be taken seriously. Expert systems may have presented a false start, but Fuller's argument, properly updated in light of current technologies, raises questions about the relation of KM to AI. If KM, from its inception, had extractive aims, its current efforts to incorporate AI in the KM technical stack may suggest that it is finally realizing its extraction/automation aims now that the technical infrastructure is coming into place. The future of KM, therefore, at least as it relates to information technology, may be to take a supporting role for AI-based automation, as opposed to asserting a distinct technology-based business strategy of its own. Besides limiting KM, this development would have ethical implications, as the question of the automation of knowledge work is a matter of society concern.

In order to resist this pessimistic direction for KM, the paper argues that Fuller's argument depends upon assumptions about the relation of tacit knowledge to explicit knowledge, and explicit knowledge to automation, that are outdated and incomplete. The faulty conception of tacit knowledge is found in its original absorption into KM by pioneers such as Nonaka and Takeuchi (1995). They held a conception in which tacit knowledge could be converted, in whole or part, into explicit knowledge. Such a conception, however, supports an extraction/automation vision. The paper argues against this view of tacit knowledge by re-examining its original explication in the work of Polanyi (1958, 1966). Polanyi saw tacit knowledge as structurally supportive of and essential to explicit knowledge.

Using Polanyi (1958) as a foundation, the paper reviews more recent theories of tacit knowledge, in particular, the work of Harry Collins (2010). Collins introduced a threefold classification of tacit knowledge that helps clarify its relation to AI. His categories of relational tacit knowledge (RTK), somatic tacit knowledge (STK), and collective tacit knowledge (CTK) are used to explicate differences between procedural and strategic knowledge. The goal of this analysis is to show the importance and essentiality of tacit knowledge in organizations, especially where innovation and higher-order thinking are at issue. The paper then argues that Collins' theory, while providing a more sophisticated view of tacit knowledge, fails to fully account for its cognitive dimensions. To address this gap, the paper introduces the concept of abductive inference, originally introduced in the work of Charles Sanders Peirce (1931), and shows that it closely answers Polanyi's view of the role of tacit knowledge in scientific thinking. In particular, abduction provides a logical structure for hypothesis formation that is distinct from deductive and inductive inference. Since the latter two forms of inference exhaust the reasoning space of AI, a further argument is given that tacit knowledge should be seen as not simply a resource to convert into automatable explicit knowledge but as the driving force behind innovative thinking in organizations.

The paper concludes by showing that the more theoretically sound theories of tacit knowledge reviewed provide a solid grounding for a human-enabling vision of KM that stands in stark contrast to Fuller's extractive vision. The paper then provides some recommendations for research that call for greater integration of cognitive science, epistemology, and related fields for future attempts to understand the relationship between the KM and AI fields.

## **FULLER'S ARGUMENT**

---

In the early 2000s, Fuller laid out an extended critique of knowledge management, arguing that its central objective is to enable commercial organizations to extract knowledge from knowledge workers and convert it into intellectual capital. Along with the converted knowledge, the value it generates is transferred from the individual to the organization, thereby devaluing human knowledge from a social and economic perspective. This project of knowledge alienation has two targets. The first is

the capture of knowledge workers' intellectual capital through converting tacit knowledge into an explicit form and, accordingly, capturing it into a system. This is the epistemological objective. The second target is the appropriation of the knowledge worker's social capital, accomplished by positioning the system into the individual's role and social network (Fuller, 2002, p. 106-109).

In Fuller's view, the paradigm of a knowledge worker and knowledge work is, respectively, the expert and professional expertise. Likewise, the paradigmatic knowledge system is an expert system, and the principal KM practice is knowledge engineering. Knowledge engineering aligns with the epistemological project of converting tacit knowledge into an explicit form amendable to computer processing. Expert systems align with the social role of answering clients' questions within restricted domains.

The epistemological project of transferring intellectual capital is based on classical epistemology, which views knowledge as consisting of extractable, abstract structures that can be stored or represented in computers with automated reasoning as the end result (Fuller, 2002). Knowledge is "... a structure of abstract entities (e.g., propositions) that are the bearers of truth values and that are possessed or grasped by individual knowers" (Mooradian, 2007, p. 4). Two persons have the same knowledge "... when they grasp (and have the same mental dispositions in relation to) the same structure of propositions" (Mooradian, 2007, p. 4). By this interpretation, knowledge can be transferred to expert systems and IT systems in general (Fuller, 2002, pp. 117-125). In other words, knowledge is a structured set of propositions that can be transferred from one container (human brain) to another (computer system).

The transfer of structured information into expert systems, along with the ability of such systems to automate deductive inferences, creates intellectual capital for the organization. However, to use intellectual capital, it must be possible for expert systems to stand in for experts and be accorded those experts' intellectual and social capital. Fuller (2002) argues that this is possible due to the sociological factors that shape and constitute expertise as a social function. Expertise and professional competence are defined by their role in answering questions within a subject-specific domain in a way that generally satisfies the goals or expectations of inquiring clients. Experts (professionals), however, define the standards of interrogatory success through their profession in such a way as to guarantee the adequacy of their answers. They do this in many ways, which include maintaining a level of collegiality that prevents mutual criticism, defining their domain and the standards of performance, preselecting clients, and shifting the burden of failed advice to clients (pp. 143 -150).

Together, Fuller's (2002) two strategies – knowledge extraction and transfer – offer an avenue for expert systems to replace human experts. Replacement is accomplished by transferring knowledge into automated systems that answer client questions successfully as defined by organizational standards. Using this knowledge engineering strategy, organizations can codify the expertise (intellectual capital) of professionals into expert systems and position those systems in the expert's social role, thereby assuming and reducing their social capital concurrently.

Fuller's argument seriously challenged the KM movement of the early 2000s and presaged recent concerns about the impact of AI-driven automation on knowledge work. The argument was flawed, however, by (1) a narrow focus on expertise and expert systems and (2) the reduction of KM to a replication of the reasoning processes of individuals. Big data and ML did not gain prominence until the early 2000s, while expert systems gained prominence during the same time when KM was forming as a field. Consequently, a number of articles promoted expert systems and knowledge engineering as a core technology for and approach to KM. (See Russell and Norvig [2021, pp. 22-26] on expert systems and AI; for an example of knowledge engineering and KM, see Nissen [2003, pp. 189-190]). Mooradian (2007, p. 301) criticizes Fuller's argument on these grounds, noting that KM technologies are a broad range of systems that include enterprise database applications, web portals, content management systems, search solutions, and collaboration suites. Further, Mooradian argues that knowledge management has the broader objective of creating organizational knowledge that can be

shared among knowledge workers in a way that is enabling rather than substitutive. Document repositories, for example, hold rich sources of collective knowledge that enable knowledge sharing between human experts (p. 303).

Effectively, Fuller limited his argument by hitching his wagon to expert systems. Despite this shortcoming, he identified a direction for knowledge management that is potentially realizable by a wider range of technologies, including the ones iterated above. Mooradian (2007, pp. 303-304) noted that search technologies, for example, could evolve to provide specific, authoritative answers dynamically and serve some of the same objectives of expert systems. Deep learning and other recent developments in AI provide new material in support of this argument. These developments offer new, more effective paths to cognitive automation, which are highly relevant to the question of tacit knowledge and its conversion into corporate assets. Further, they connect many of the traditional components of KM systems, such as content repositories and search. The following section updates Fuller's argument in light of this broader understanding of knowledge work and recent AI developments.

## UPDATING AND EXPANDING THE ARGUMENT

---

### *EXPERTISE*

Fuller's argument focuses on expertise, making it paradigmatic of knowledge work and largely understood in terms of professionalism and professional fields. From this perspective, knowledge work is identified with professionalism, and the features of professionalism explain the social capital of expertise and its transfer to expert systems. Fuller's identification of knowledge work and expertise/professionalism is overly narrow, as is his conception of expertise. Nevertheless, there are elements of his characterization of knowledge work and expertise that can be updated or adapted to apply to contemporary KM and AI. I will draw attention to the limitations of his account (both the identification of expertise and knowledge work and his account of expertise) and then show how they can be adapted.

First, Fuller's identification of knowledge work and expertise is unduly restrictive. While expertise as a form of knowledge is critical to organizations, knowledge falling outside of the definition of expertise – referred to hereafter as “non-expert knowledge” – is equally, if not more, important. Experts are “... people who produce clearly above average (outstanding) performances on a regular basis” and who obtain their expertise through extensive study and practice in stable domains (Bilalić, 2017, p. 3). So characterized, expert knowledge is critical to many organizations, but it cannot be fully cultivated from within, as the existence of a stable domain implies an external source.

By contrast, non-expert knowledge may lack a stable domain, and therefore, the creation of stable knowledge may be unfeasible if the environment is constantly evolving. However, where such dynamic environments have some level of continuity, non-expert knowledge is possible and holds high value to organizations. Moreover, it will require generation from within organizations and their industries in large part. As such, non-expert knowledge is a target for organizations to capture and is much broader and more variegated than the class of traditionally recognized professions.

Further, within the category of non-expert knowledge, there are forms of knowledge that share with expertise the structural characteristic of requiring training\study in a stable domain but lack the level of mastery. Here, we are talking about competence. As with expertise, competence requires practice and training in a domain, and it answers to a standard of evaluation. However, by definition, competence falls short of expertise in relation to the relevant standard. Still, for many organizational roles and types of knowledge work, competence is perfectly adequate and often cost-effective in comparison with expertise. An example where competence is both adequate and desired is in customer service. Customer service agents typically do not need to demonstrate mastery of a domain in their work supporting customers. They simply need to provide basic information or perform highly learnable tasks. The typical “domains” of customer service, that is, types of products, do not require years of

study or practice to perform the customer service role. By contrast, the developers of products (computer programmers, electrical engineers, etc.) do exhibit mastery of domains based on years of study and practice.

Second, Fuller's sociological characterization of expertise is too narrow and fails to account for other perspectives, such as those in cognitive and neuroscience, as well as philosophy. According to Fuller, a defining characterization of expertise is the relationship between the putative expert and the client, in which the expert determines the conditions of success (see above). From a cognitive scientific view, expertise requires high-level performance in a domain, and it is based on cognitive strategies that differ from those of non-experts. Unlike novices, experts apply knowledge stored in long-term memory (LTM) to new situations. This cognitive strategy holds for the three domains of expertise studied in cognitive and neuroscience, namely perceptual, cognitive, and motor. What does differ across these areas is the neural implementation of LTM. For example, brain imagery of radiologists shows higher activity in the fusiform gyri in the inferotemporal cortex. By contrast, motor expertise shows activation of mirror neurons or the action observation network (AON) (Bilalić, 2017, pp. 22, 25). Complementing the cognitive scientific perspective, philosophers such as Goldman (2018) focus on propositions held by experts in comparison with non-experts, with a threshold quantity of relevant true and/or justified beliefs being the criterion.

Finally, Fuller's conception of expertise falls within the category of subject expertise as opposed to practical expertise, as elaborated by Winch (2012). According to Winch, subject expertise is associated with academic work and is "... related to declarative or propositional knowledge ..." while practical knowledge relates to "... ability or skill" (p. 1). Medicine, which Fuller includes in his examples of expertise, falls on the practical side of the divide in Winch's classification. However, Fuller's characterization of medical and other forms of expertise focuses on informational activities such as answering questions. He neglects to discuss the panoply of motor skills involved in many forms of expertise, such as medicine, that come into play in the use of instrumentation, physical manipulation of patients, and, of course, surgery.

Despite the shortcomings of Fuller's account of expertise, it still has relevance and applicability to important contexts in KM. First, while Fuller's sociological definition of expertise as a kind of power dynamic between experts and their clients may have, at best, partial application to professions such as law and medicine, in many organizational contexts, institutional aspects of professionalism are much less relevant than market forces. Market forces often include the type of dynamic that Fuller attributes to professions, namely, the ability of commercial organizations to shape consumer expectations regarding what counts as satisfactory answers to questions or information requests. This ability to shape expectations certainly applies to outputs from automation technologies. Anyone who has performed self-checkout in a big-box retail outlet or spoken to an automated phone system knows that commercial and governmental organizations often dictate the terms of service and shape the criteria of knowledgeability ascriptions adopted by their customers. Organizations using expert systems (or successor AI technologies, more realistically) set the conditions of their use and acceptability as much as do expert practitioners and their professional associations, or more so.

Further, many, if not the majority, of interactions between commercial organizations do not require professional-level expertise. Rather, they require non-expert forms of knowledge, including competence, to provide information about products and services, as well as solutions to customer problems. Such knowledge can be stored in a variety of traditional information systems such as databases, hyper-text systems, document repositories, and expert systems. Further, as such knowledge work consists of informational interactions, motor skills and expertise are less relevant.

To summarize, a useful update to Fuller's argument expands the scope of knowledge work to non-experts (knowledge workers who would not be traditionally classified as experts) within a broader commercial market context untethered to professional norms. The next section considers how Fuller's argument can be updated and applied to the wide range of modern AI technologies.

## *FROM EXPERT SYSTEMS TO MACHINE LEARNING*

Fuller's argument relates knowledge management to the conversion of tacit knowledge into explicit knowledge and, in turn, to the specific approach of knowledge engineering and expert systems. Of course, knowledge management is a broader strategy than just the conversion of tacit knowledge into systems; it includes managing knowledge workers and knowledge work in ways that facilitate knowledge creation and sharing. Nevertheless, capturing tacit knowledge in IT systems has always been an integral part of knowledge management.

The full strength of Fuller's argument was forestalled by the failure of expert systems to deliver on their promise. By applying a domain-based knowledge approach to solve problems of tractability and relevance, these systems appeared effective with early general logic-based problems. They also contributed to numerous industrial applications and advanced research on knowledge-based reasoning, with prominent successors being ontologies and description logics. Their success was hampered, however, by the knowledge acquisition problem at the center of knowledge engineering: extracting knowledge from experts is expensive and slow, and even in the best circumstances, the extracted knowledge is only a partial specification of that knowledge domain and the expert's understanding of it.

Ironically, the problem resides in the nature of tacit knowledge. Eliciting knowledge is a time-consuming, arduous process. Even in specialized domains, the amount of knowledge needed to reason is vast, as these complex domains are part of a larger causal system. Further, while much tacit knowledge is simply implicit or unstated, much also consists of mental operations that remain hidden from us. How we select premises to reason from or to is a yet-to-be-understood process, and the selection decision likely depends on tacit knowledge, which is difficult to surface or impossible to make explicit. Additionally, expert systems rely on deductive inference and are brittle in the face of uncertainty. Due to this inability to learn from data and their inferential fragility, early AI systems – in the form of expert systems – failed to meet their objectives. As a result, interest in expert systems evaporated, and an “AI winter” prevailed (Russell & Norvig, 2021, p. 24).

The story of the resurgence of AI in tandem with the rise of big data and the symbiosis of ML techniques with big data is now quite familiar. DL, in particular, revived connectionism, a long-standing research program, and is among the most successful and attention-getting categories of algorithm types and approaches. ML/DL technologies are ubiquitous across applications, and a hurried effort to integrate them into the KM technology stack is underway (AIIM, 2019; Patton, 2022). Moreover, ML and DL are poised to encroach upon human knowledge work and expertise.

Deep learning technologies such as large language models (LLMs) (e.g., ChatGPT) have been described as capturing tacit knowledge (Brynjolfsson et al., 2023). Many LLM task areas coincide with characteristics of expertise, such as the ability to make fine-grained discriminations, provide high-quality recommendations, and answer questions within complicated domains. The power and variegated nature of DL applications raises new possibilities for knowledge management that a single approach (e.g., knowledge engineering, expert systems) does not. Moreover, ML/DL technologies combine with contemporary knowledge-based systems such as ontologies and knowledge graphs in ways that breathe new life into knowledge engineering approaches by allowing knowledge systems to learn from data. Ontology engineering is an example of such a combination.

Considering the above, the idea that knowledge management may be extractive cannot be dismissed simply because its argument was based on a single, earlier approach to knowledge automation in the form of expert systems. A panoply of powerful approaches to creating knowledge from data and enhancing semantic or knowledge based-approaches currently exist. While it may appear that these approaches are not extractive per se, as they go directly to data and information and leave out the human expert (with the exception of hybrid knowledge systems), this observation will hold only in some cases. In many others, the extraction is less direct or even unnoticeable.



In place of interviewing identifiable experts within an organization or field, deep learning techniques use data sets generated by, or collected from, processes and activities of knowledge work of various kinds. Where supervised learning is used, data sets include the target variable of which the values in rows represent correct answers, which, in essence, are the coding of human decisions. The prior selection of data samples and their proper tagging may involve expert knowledge (e.g., radiology). In the case of unsupervised learning, even though the target variable is not supplied, the selection of attributes reflects prior knowledge. Finally, in the case of LLMs, while large swaths of the Internet may be scrapped, specialized language models trained on knowledge products may make up the training base. These are not data sets but text corpora. Databases like PubMed containing scientific articles are good examples. The articles in PubMed are produced by experts in various areas of health research, contain the knowledge of their authors, and indirectly include the knowledge of supporting experts such as peer reviewers, ethics reviewers, funding entities, etc. (Singhal et al., 2023). The upshot is that ML and DL generally depend on prior human knowledge, including expert knowledge. Some forms of ML, especially deep learning, perform tasks associated with expertise including perceptual discrimination and answering questions. The following subsections elaborate on these two task areas.

### **Experts recognize: Image recognition**

As alluded to above, perceptual discrimination is the basis for numerous knowledge areas. Examples of perceptual and cognitive expertise include visual expertise, auditory expertise, and gustatory expertise (Bilalić, 2017, pp. 34-99). Fuller (2002) traces the etymology of ‘expertise’ to ‘experience,’ noting that the first persons to be denominated as “experts” were “... witnesses in trials to detect handwriting forgeries” (p. 143). The visual expertise of radiologists is a more recent focus of study in the cognitive science of expertise. Radiologists exhibit expert performance through quick and accurate detection of abnormalities (lesions) in radiographic images. Experiments have shown that expert radiologists (like other visual experts) focus their visual scanning of an image to the most relevant areas as per their holistic perceptions of the entire image. To do this, they draw from their domain knowledge in such a way as to automatically detect patterns relevant to the potential abnormalities (Bilalić, 2017, pp. 7-8, 58-65).

Visual discrimination is an active area in the research and development of deep learning systems, for which convolutional neural networks (CNNs) are a principal architecture. CNNs are often used with labeled images (for supervised learning) and have been highly successful in image discrimination and classification (e.g., types of animals, products, etc.) (Russell & Norvig, 2021, p. 896). A plethora of use cases in business contexts exist, ranging from product image labeling to facial recognition for identification purposes. There are also use cases in areas involving professional-level expertise, including radiology and medical imaging in general.

For decades, the use of ML in medical imaging has been an ongoing project. Deep learning provides significant advancements in this area, much as it has in other fields (El Naqa & Murphy, 2022; Mayo & Leung, 2018). General use areas include computer-aided detection (CADe) and computer-aided diagnosis (CADx) (El Naqa & Murphy, 2022, p. 10). Current uses of CADe and CADx are assistive but show the possibility that DL technologies become capable of performing tasks normally handled by human experts. They also reflect the importance of prior expert knowledge when tagging and preparing training sets composed of medical images. This knowledge can be described as tacit as it requires the integration of years of practical experience with theoretical knowledge. DL reflects this tacit dimension because it learns from data (guided by prior knowledge) instead of relying on explicit, pre-formulated rules to detect specific lesion types (El Naqa & Murphy, 2022, p. 175).

### **Experts provide answers: Conversational AI**

Fuller correlates expertise with professionalism and professionalism with answering questions. This characterization is too narrow, as expert performance varies greatly in type and includes both percep-

tual expertise and physical activities (e.g., performing surgery). Further, answering questions, explaining things, making arguments, and conducting other verbal performances comprise a large part of expertise.

LLMs are a form of deep learning that has recently gained prominence. ChatGPT, a conversational agent based on the GPT LLM family, demonstrates the power of LLMs in producing coherent text in the form of answers, explanations, arguments, software coding, and other language-based tasks. Trained on vast corpora of text from the web, including but not limited to articles, web pages, blogs, and social media, ChatGPT appears to provide encyclopedic knowledge in a manner equivalent to an interdisciplinary assembly of experts. However, the data sets (corpora), though large, are not complete or totally current. More importantly, they are built on statistics and probabilistic methods (probability distributions for succeeding tokens are calculated from text embeddings using a neural network). As a result, “hallucinations” are common. As opposed to a proposition with semantic content, a false answer as a predicted string of text may be more probable than a correct answer. Despite these limitations, LLMs like ChatGPT serve as foundations for domain-specific conversational systems that provide serviceable answers and summaries in a form similar to those provided by experts from a range of industries, business functions, and professional fields. Trained with information created by an organization’s knowledge workers or within a scientific discipline, LLMs can potentially provide relevant, high-quality answers to domain-organization-specific questions.

Brynjolfsson et al. (2023) provide an early case study on ChatGPT and customer support, a common functional area targeted by KM technology vendors. Their working paper describes a pilot project for the customer service function of a software developer. The project consisted of a rollout of a conversational assistant tool to aid low-experience customer service agents with finding resolutions to technical support issues and answering customers’ questions. The conversational assistant tool used ChatGPT as a foundation, providing natural language processing (NLP) functionality. The tool was then trained using the texts from successful customer service calls handled by more experienced customer service agents. In this way, a supervised learning approach creates a domain-specific model capable of improving the success of less experienced workers by embodying “... the best practices of high-skill workers ...” (p. 24). Of particular interest to the authors was the “... model’s ability to encode the potentially tacit knowledge of high performers ...” (p. 24).

The medical domain, a target for early expert systems (e.g., MYCIN), is a key area of investigation into the viability of LLM applications in professional/scientific domains (Cawsey, 1998, pp. 54-58; Russell & Norvig, 2021, p. 23). For example, Singhal et al. (2023) developed Med-PaLM, an instruction prompt-tuned version of Flan-PaLM, with itself being an LLM specialized for the medical domain. Using their robust evaluation benchmark framework, the researchers report impressive results compared to other LLMs and professional clinicians themselves. While these results show great promise, the authors identified limitations resulting from the complexity and high-risk nature of the medical domain.

Karim et al. (2023) provide another example of the use of LLMs in the medical domain. Their paper details a project to build a knowledge graph based on OncoNetOntology (ONO), an oncology ontology constructed from existing biological ontologies as the knowledge base. It also incorporates LLM (NLP) technologies to update the base and take advantage of the large volumes of medical knowledge contained in repositories such as PubMed. NLP techniques, including entity extraction, entity linking, and relation extraction, were employed to further populate the ontology of knowledge structures. These cases exemplify how machine learning approaches can be combined with knowledge representation approaches to create AI systems supportive of knowledge discovery and question-answering functionality.

## AI AND KM STACK

---

The technologies reviewed above are examples of recent AI approaches in specific domains. Increasingly, these technologies are being integrated into the KM technology stack with the goal of re-energizing knowledge management technologies. Describing a KM technology stack, however, requires further elucidation because many KM-associated technologies exist that support multiple application types besides KM. As Lytras et al. (2008) noted, “the justification of an application as a knowledge management one has to be based on a context” (p. xvi). Technologies may be organized conceptually in different ways in terms of whether they support knowledge processes or knowledge products (p xvii). Applications in the first group include but are not limited to net conferencing, discussion groups, and messaging. The second group includes but is not limited to document repositories, intranets, and full-text retrieval.

According to McKendrick’s (2023) survey, “selected knowledge management capabilities” (ordered by funding level) include content management, search, process guidance, document management, metadata management, conversational guidance in workflows, and text/data mining and analysis. Enterprise content management (ECM) includes document management, metadata management, workflow, and other components (Mooradian, 2021). As a platform stack, content management is allied with KM technology and has its own tab alongside AI on the KM World website (<https://www.kmworld.com/>). According to the trending capabilities identified by the McKendrick (2023) survey, ECM vendors are actively integrating AI technologies into their platforms. Also, a white paper from AIIM (2019) lists the following capabilities in its AI roadmap: topic identification and classification, entity extraction, automated image tagging, recommendation engines, conversational agents, speech recognition, other natural language processing capabilities, and other functionalities. Microsoft’s content management strategy, originally based in SharePoint, was rebranded as “content AI” and incorporates NLP capabilities, including machine translation and recommendation functionality (Patton, 2022). Thus, to the extent that KM and ECM technologies overlap, the KM stack can be considered capable of quickly incorporating AI capabilities.

Further, the integration of AI into the KM technology stack and strategy is the subject of active research. For example, Majumder and Dey (2022, p. 88) explored “AIKM” in their book *AI-empowered Knowledge Management*. They reviewed use cases similar to those above, including those addressing the integration of recommendation engines into workflow technologies and also the integration of big data analytics as a source of knowledge capture. Additionally, Müller et al. (2020) provide a case study on the integration of recommendation engine capabilities into a technical support case management system. Finally, Brynjolfsson et al.’s (2023) case study, discussed above, concerns a customer support function where conversational AI supports customer service agents. While Müller et al. (2020) and Brynjolfsson et al. (2023) do not use the label “knowledge management,” their use cases align well with those on KM technologies focused on customer service and those discussing the incorporation of conversational AI into KM and ECM stacks (AIIM, 2019).

The success of machine learning and the development of intelligent solutions that combine ML-related technologies have reinvigorated both artificial intelligence and knowledge management. Once central to a KM technology vision, the decline of expert systems forestalled strong automation strategies in knowledge management. With the ascension of ML/DL, the question is whether the knowledge management strategy will now shift toward automation.

## RESPONSE TO ARGUMENT

---

### *TWO PARADIGMS OF KM: KNOWLEDGE AND INTELLIGENCE*

An argument against an ineluctable march toward extractive automation may be found in the idea that knowledge management has always had two technology paradigms: one of knowledge products

captured and managed in repositories and one of artificial intelligence (originally represented by expert systems). These two paradigms correspond to different ways that systems manage knowledge as opposed to information. Knowledge products such as documents are vehicles or containers of human knowledge that synthesize thought using explanatory, argumentative, or narrative structures. These knowledge products rise above mere information, which, while representing basic cognition, does not reflect the tacit knowledge and cognitive skills that can produce an explanation, extrapolate a generalization, or conceive of a series of steps to accomplish a task. Likewise, artificial intelligence has a claim to represent knowledge over mere information insofar as it simulates reasoning steps. The reasoning tasks go by many names (e.g., classification, prediction, etc.), but they all are types of deductive or inductive inference.

Given the two paradigms, one could make an argument that it is only the AI paradigm that has an extractive and replacement tendency. By contrast, the knowledge product paradigm does not. Mooradian (2007) identified these paradigms and described them in terms of competing models in KM: the expert/customer model and the expert/expert model. The expert/customer model corresponds to the use of expert systems that interface directly with customers and clients, therefore bypassing experts in the way that Fuller envisions. In contrast, the expert/expert model corresponds to knowledge repositories built from the contributions of knowledge workers (e.g., problem resolution notes) and accessed by other knowledge workers (pp. 302-304).

The characterizations of these two models are relevant today, though they are not exhaustive in the way put forward by Mooradian (2007). While documents contain knowledge that requires interpretation based on the background knowledge of persons (experts, professionals), the scope of interpretation depends on the nature of the document. A scientific or technical paper contains a high degree of knowledge, requiring expertise to understand and apply. However, a document may also represent a stepwise process that can be easily read, followed, and applied in a mechanical fashion. The knowledge it took to develop that process may have depended on human expert knowledge, much of which was tacit, but the procedural format allows users to understand and apply it without access to that original tacit knowledge. Therefore, within the expert/customer model, the document is usable in a way similar to an expert system, which Mooradian acknowledges (p. 304).

Conversely, within the expert/expert model, AI-based systems can share knowledge among knowledge workers to enhance performance. Brynjolfsson et al.'s (2023) article is an example of how a conversational AI agent supports knowledge workers. Instead of the system interfacing directly with customers, knowledge workers (in this case, customer service agents) used the system as an information tool to learn and improve their performance. Likewise, the case study by Müller et al. (2020) involves the use of recommendation engines that find and present relevant problem-resolution information to technical support workers as they work on their cases.

Knowledge-based reasoning engines, the successor to expert systems, can provide summaries and visualization of their reasoning steps. For example, knowledge graphs, a growing area of research, are utilized in industry to provide insights for knowledge workers to interpret and act upon. In the case of machine learning, explainable AI (XAI) provides an understandable reconstruction of the reasoning process that, in turn, can be a source of learning for knowledge workers. The implication of this review of Mooradian's models (expert/expert vs. customer/expert) is that the distinction between the two paradigms of knowledge products and artificial intelligence described at the beginning of this section does not perfectly support the distinction between an enabling and automation approach to knowledge management.

### ***TACIT KNOWLEDGE AND KM***

If the above account is correct, modern AI technologies present concerns similar to those posed by expert systems. Automation in the form of advanced AI, combined with rules-based processing and procedural documentation, provides a hyper-automated trajectory for organizations. Organizations can pursue knowledge engineering strategies where they are feasible and likely to bear fruit (e.g.,

Collins' (2010) concept of relational tacit knowledge) or use deep learning approaches when knowledge engineering is not effective. With one foot in automation, KM can easily evolve into a supporting discipline, though it would shed large swaths of theory and practice by doing so.

To avoid this direction, it is necessary to re-examine some underlying epistemological assumptions fundamental to knowledge management and its technical paradigms. These assumptions have to do with conceptions of tacit knowledge and its conversion into explicit knowledge. The conceptions were introduced into knowledge management by Nonaka and Takeuchi (1995) and other pioneering thinkers inspired by the philosopher of science Michael Polanyi (1966). While recent scholarship reflects a departure from Polanyi's original theory, revisiting both his original conception and more recent works on tacit knowledge promises to provide a foundation for positioning knowledge management as an enabling (rather than exploitative) strategy in relation to the recent development of AI and future technological developments.

Nonaka and Takeuchi (1995) made the concept of tacit knowledge central to knowledge management. In their seminal work, they distinguish between tacit and explicit knowledge – tacit knowledge is personal, context-specific, and therefore hard to formalize and communicate. Explicit or “codified” knowledge, on the other hand, refers to knowledge that is transmittable in formal, systematic language.

Importantly, Nonaka and Takeuchi (1995) make the development of methods for converting tacit knowledge into explicit knowledge a central goal of knowledge management. According to Mooradian's (2005) interpretation, Nonaka and Takeuchi (1995) hold a robust view of tacit knowledge, seeing it as different in kind from explicit knowledge and, therefore, in need of transformation in order to become explicit. Other early researchers, such as Nissen (2003), offer a less robust characterization of the distinction, stating that tacit knowledge is not necessarily different in kind from explicit knowledge. Rather, tacit knowledge is simply knowledge in people's minds, requiring only elicitation and codification to become explicit. Nissen's (2003) characterization corresponds to Collins' (2010) conception of relational tacit knowledge (RTK) in his tripartite division.

Whether robust or weak, both views present tacit knowledge as something that can be converted into explicit knowledge with varying degrees of transformation and loss. Once converted, however, tacit knowledge would be extracted from people's heads and codified externally within organizational systems. While there might be some loss during conversion, once the conversion is complete, the tacit knowledge in question can be considered exhausted or consumed as a resource, and the organization can then pursue more tacit knowledge from the same or other human sources.

It does need to be acknowledged that in Nonaka and Takeuchi's (1995) model, explicit knowledge, including that converted from tacit knowledge, becomes a source for the ongoing creation of new tacit knowledge. A focus on sharing explicit knowledge with knowledge workers so that they can both internalize it and add value to organizational objectives is fundamental to the KM enterprise as a whole, as it provides an enabling strategy for AI. Nevertheless, this strategy is one option for organizations. Many organizations will choose a replacement strategy instead of an enabling strategy because replacement objectives are more easily assessable in terms of return on investment (ROI). With ROI, it is easy enough to envision whether or not the estimations are actually borne out. In contrast, re-organizing work to optimize the benefits of human interaction with AI technologies requires a greater feat of imagination than does a replacement strategy, and it is more difficult to assess. For this reason, Nonaka and Takeuchi's conception of tacit knowledge as a convertible resource to be mined combines well with an ROI-driven replacement strategy of automation based on contemporary AI.

In sum, from its inception, knowledge management has maintained two contrary tendencies: one toward extraction and the other toward (human) knowledge creation or enablement. Further, the tendencies trace back to Nonaka and Takeuchi's (1995) founding theory of tacit knowledge, even though it envisions a harmonization of these tendencies. To support the enablement tendency as a

core dimension of knowledge management, a re-examination of tacit knowledge and its place in KM is needed. To this end, the next section takes a deeper look at Nonaka and Takeuchi's conception of tacit knowledge and envisions a more robust theory consistent with that found in Polanyi's (1966) original conception and developed by contemporary theorists.

### ***RETHINKING TACIT KNOWLEDGE IN KM***

Nonaka and Takeuchi's (1995) conception of tacit knowledge was criticized by Mooradian (2005). Mooradian argued that Nonaka and Takeuchi fundamentally misapplied the concept, in large part because they ignored its central role in Polanyi's philosophy of science. In particular, he argued that they failed to recognize the central feature of tacit knowledge as conceived by Polanyi, namely, that tacit knowledge has a structural relationship to explicit knowledge in that an individual's tacit knowledge enables explicit knowledge. Polanyi (1958, pp. 55-65) elaborated his theory in *Personal Knowledge: Towards a Post-Critical Philosophy*, in which he describes the explicit dimension of human knowledge as "focal" and the implicit dimension as "subsidiary." In *The Tacit Dimension*, Polanyi (1966, p. 13) describes this same relationship as "proximate and distal forms of knowing." Regardless of the exact terminology, according to Mooradian (2005), the nature of this structural relationship is that, in any act of explicitly knowing something (Ke), there is a set of enabling tacit elements of knowledge (Kt1 ... Ktn) that is not explicit knowledge. For Polanyi, these tacit elements (subsidiary or proximate) are elemental and integrate into an unanalyzable whole. Mooradian's (2005, pp. 105-107) account includes this part/whole structure as a special case while broadening the approach to include a wide range of supporting structural relationships extant between explicit and tacit knowledge.

Another aspect of Polanyi's (1966) theory not appreciated in Nonaka and Takeuchi's (1995) work is that it was developed as a criticism of a formalist/objectivist conception of knowledge, namely, logical positivism, which held that knowledge is purely objective and representable by rules of logical inference. This foundational context was lost in its importation into knowledge management (Mooradian, 2005, p. 2). This argument, however, is central to Polanyi's theory and its end goal. In *The Tacit Dimension*, he makes this clear:

The declared aim of modern science is to establish a strictly detached, objective knowledge ... But suppose that tacit thought forms an indispensable part of all knowledge, then the ideal of eliminating personal elements of knowledge would, in effect, aim at the destruction of knowledge. (Polanyi, 1966, p. 20)

The idea that tacit knowledge is part of all acts of knowing was lost in its appropriation into the knowledge management literature. Instead, a conversion or convertibility paradigm predominated. In this paradigm, domains or subdomains of knowledge are seen as convertible into explicit knowledge, even if imperfectly. Since explicit knowledge can be codified in information systems, it is potentially usable for automation purposes. Fuller's argument against knowledge management presupposes the conversion paradigm. It makes expert systems a representative technology. While expert systems require explicit knowledge structures, machine learning is able to work with data and unstructured content. Therefore, assuming that converted tacit knowledge accounts for the (usable) majority of human knowledge in a domain or subdomain, automation of knowledge domains would seem feasible.

Polanyi's (1966) conception of tacit knowledge, however, supports a different, human-enabling vision. In his view, explicit knowledge is enabled by tacit knowledge. Consequently, knowledge domains are not fully automatable. A tacit element is required to make explicit knowledge function as knowledge in both its comprehension and application because the relationship of tacit to explicit knowledge is structural. Further, he finds this structural relation both in everyday cognitive acts and in the highest achievements of science. For example, the ability to recognize faces (focal awareness) is a mundane cognitive skill dependent upon an integrated and subsidiary awareness of the features of the face. If we do find some rules that can transfer this knowledge by surfacing the subsidiary elements as police sketch artists might by matching feature representations identified by witnesses, this

still requires tacit knowledge on the part of the sketch artist and the witnesses (Polanyi, 1966, pp. 4-5). While some tacit knowledge is made explicit, other tacit elements continue to support and shape explicit knowledge.

The case of the scientist choosing a problem also illustrates the relationship between tacit and explicit knowledge and forms the arc of Polanyi's (1966) argument that scientific knowledge (and hence all knowledge) is based on tacit knowledge and not fully formalizable. All research begins with the formulation of a problem. The ability to formulate the problem, however, is based on prior tacit knowledge. This tacit knowledge grounds a type of recognition that the problem is potentially interesting and will lead to an undiscovered truth. This intimation cannot be fully justified with an explicit argument (pp. 21-25). This description of hypothesizing is also relevant to the process of innovation, which is an important concern in KM.

The practical upshot of Mooradian's (2005) attempt to reconnect tacit knowledge with its philosophical origins in Polanyi (1966) was to identify risks and limitations in the knowledge transfer process and to suggest a strategy for addressing them. Under the convertibility paradigm, tacit knowledge was seen as capturable and transferable in part or whole. However, the theory did not match practice, and tacit knowledge proved to be "sticky," as Szulanski (2003) argued in his book, *Sticky Knowledge: Barriers to Knowing in the Firm*. Recognizing the structural dependency of explicit and tacit knowledge allowed for a strategy of identifying areas of tacit knowledge needed by the receivers of explicit knowledge or the users of systems containing it to make the transfer successful (Mooradian, 2005, pp. 110-112). The tacit knowledge required to understand and use explicit knowledge may not be as extensive as that which produced it, but some tacit knowledge will be required. Failing to take this into account and to identify the means of creating the requisite tacit knowledge (e.g., through training, further communications, etc.) could easily lead to failure in the transmission.

While Mooradian's (2005) analysis addressed failures in the transmission of knowledge, the point can be generalized to automation contexts. When automation projects attempt to automate a process from end-to-end, they often fail precisely because one or more steps have not been accounted for in the requirements analysis. Expert systems are a prominent but special case of this kind of failure. Knowledge engineers attempt to elicit all the necessary rules in the form of if-then conditionals but invariably fail to capture the tacit knowledge of experts or process knowledge holders. The resulting knowledge base ends up being incomplete, and all relevant cases, especially exceptions, are unable to be handled. The type of tacit knowledge missed in the requirements building or knowledge engineering process may vary as to its location on the difficult-to-convert scale (Collins, 2010; Mooradian, 2005). It may be relatively superficial information that was simply missed or deliberately hidden (Collins' relational tacit knowledge or RTK). It may also be deep tacit knowledge that cannot be fully articulated (Collins' collective tacit knowledge or CTK).

Whether shallow or deep, missing tacit knowledge is a critical failure point for automation projects and is even more severe than the problems that arise from knowledge transmission. This is because, in the case of human-to-human knowledge transmission, the human recipients of faulty knowledge can often mitigate the failure and find workarounds by identifying the missing information and seeking solutions. Computer systems rely entirely on pre-programmed rules and do not have the kind of autonomy and meta-knowledge required to work with partial knowledge.

For this reason, when automation succeeds, it is often (or perhaps always) because those interacting with the system supply the supporting knowledge. Humans make up for the gaps in, or limitations of, an automated system by supplying the missing information or knowledge, the latter of which often involves the application of (outputted) information to a context. In Collins' (2010) view, systems can be compared to prostheses that humans make work:

In medicine, prostheses rarely work in exactly the same way as the part they replace. What happens is that the relevant elements of the embedding organism makes [sic]

up for and metaphorically “repair” the difference between the original and the prosthesis. It is the same with the calculator. The calculator can only work as a social prosthesis, the deficiencies of which are made up for and repaired by the surrounding organism. (p. 71).

The deficiency of the calculator, as explained by Collins, lies in the simple fact that it is automating arithmetical functions. Humans, however, put the functions to use in context to obtain answers to questions that fit a purpose. The human-computer interaction appears seamless when it is successful. Trouble arises when automation ambitions attempt to cut out the “human from the loop” (to borrow a phrase from a different context).

## REPOSITIONING TACIT KNOWLEDGE IN KNOWLEDGE WORK AND AI

---

If the above is correct, it suggests a direction for knowledge management to reposition tacit knowledge within automation strategies in a way that preserves the centrality of knowledge workers and that is non-extractive (contra Fuller) while at the same time fully leveraging AI technologies within the KM stack. The first step consists of recognizing the structural relationship between explicit and tacit knowledge and using this reconceptualization as the basis for configuring knowledge work. The second step requires analyzing different types of tacit knowledge and their relation to knowledge work and AI automation. This second step will be an ongoing project for KM and will require a deeper engagement with epistemology, cognitive science, and related fields in order to better situate human knowledge in relation to artificial knowledge, and human intelligence in relation to artificial intelligence. It will include a taxonomy of types of knowledge work and knowledge activities and the role tacit knowledge plays in these types of knowledge work and activities. Such an analysis can use Polanyi’s work as a point of reference but will require current theoretical developments in the study of tacit knowledge.

This section sketches a rough example of how recent work on tacit knowledge as well as human and AI reasoning, bear on knowledge work. It starts with a distinction between procedural knowledge and strategic knowledge and then analyzes these two categories of knowledge activity from the perspective of Collins’ (2010) threefold classification of tacit knowledge. It then examines the concept of abductive inference and its relation to tacit knowledge and AI. Finally, it reflects more broadly on the relation between tacit knowledge, AI, and KM.

The distinction between procedural knowledge and strategic knowledge is familiar. Procedural knowledge requires following steps to achieve an outcome. It is rules-based, though not every detail is captured by the rules, especially exceptions (as noted above). Procedures are knowledge artifacts critical to the success of organizations, as they support repeatable, consistent, and efficient organizational actions and operations (Bercerra-Fernandez et al., 2024, p. 28; LaFayette et al., 2019, pp. 121-122). Strategic knowledge, by contrast, has to do with goal-based problems and principles. Such goal-based problems lack stepwise solutions and require instead the formulation of strategies based on the nature of the goal and relevant principles (Clark, 2008, pp. 145-147). Strategic knowledge relates to expertise, as experts invariably possess it, though it may also be found outside of what may be considered expert domains. An example of strategic knowledge is that of a chef (Clark, 2008, p.146). The chef’s knowledge includes culinary principles, food design, the preferences and interests of customers, local supply of ingredients, etc. The chef’s objectives are to create unique dishes that meet high standards of quality. The chef integrates different types of knowledge, discovers solutions to diverse problems, and innovates.



### ***COLLINS' CLASSIFICATION OF TACIT KNOWLEDGE***

Collins' (2010) tripartite classification of tacit knowledge can be used to analyze the role of tacit knowledge in both procedural and strategic knowledge. Collins divided tacit knowledge into the categories of relational tacit knowledge (RTK), somatic tacit knowledge (STK), and collective tacit knowledge (CTK). RTK is tacit for contingent reasons. For example, it may take too much time and effort to codify, the holder may not want to share, or the holder and others may not recognize its value (pp. 91-97). CTK is based on the socialization and language acquisition process and cannot be made explicit (pp. 119-133). Somewhere in the middle is somatic tacit knowledge (STK), which cannot so much be captured as replicated in a mechanism (pp. 99-113).

Starting with procedural knowledge, we can first note that much of it falls into the category of explicit knowledge and, as such, is codified in a natural language or a programming language embedded within a workflow engine (Collins, 2010, pp. 16-20). Much of the tacit knowledge, on the other hand, falls into the category of relational tacit knowledge, or RTK. It is tacit not because it cannot, in principle, be articulated but because there are difficulties in articulating everything that can be articulated or because there are motivational headwinds to full articulation. RTK is articulable knowledge that lies below the surface but does not bubble up for multiple, contingent reasons (pp. 91-98). Exceptions to rules will often fall into the category of RTK. They are often missed in requirements analysis because they are too infrequent to remember or because too much detail escapes recall in a given requirements elicitation session.

Process knowledge may also include tacit knowledge in the category of somatic tacit knowledge, or STK if it requires physical skill. STK consists of skillful bodily movement, which is a rule that persons cannot articulate. Polanyi's (1966) bicycle riding is the most famous example in the literature (Collins, 2010, pp. 99-101). Making sourdough bread involves following explicit instructions, but it also requires manipulating the bread artfully. Medicine, which, as noted above, Fuller treats in purely informational terms, involves high levels of somatic skill. Finally, but to a lesser extent, process knowledge may involve collective tacit knowledge or CTK. CTK is linguistic-cultural knowledge acquired through socialization in a particular social-linguistic community. It includes the whole of linguistic and social rules or understandings that individuals possess but cannot articulate (Collins, 2010, pp. 99-113). CTK comes into play in process knowledge insofar as it involves background linguistic knowledge and common knowledge, as well as social norms. For example, a process that involves communicating with people (which many will) presupposes rules of social etiquette and shared values. Here, Collins and Evans' (2007, pp. 13-14) periodic table of expertise is relevant. They posit a spectrum of types of expertise from common to specialist and hold that all are undergirded by common tacit knowledge.

Turning to strategic knowledge, we can anticipate that Collins' (2010) classification applies somewhat in reverse. Understanding how to realize goals based on principles requires a wide range of diverse concepts that draw on social knowledge. This social knowledge will include the relevant expert domain as well as common background social knowledge. CTK is, therefore, a principal constituent. To take Clark's (2008) chef as an example, the chef will have been trained in a social environment consisting of culinary experts of varying sorts (master chefs, connoisseurs, growers, etc.) but will also learn the practice within cultural/national milieus. Clark's chef can also be expected to have STK to a significant degree, along with other experts and competent practitioners who engage in skillful performance, such as doctors and nurses. Finally, RTK will be present as an aid in the interpretation and completion of process knowledge.

### ***ABDUCTION AND TACIT KNOWLEDGE***

While Collins' (2010) CTK provides an account of sorts of strategic knowledge in terms of conceptual and social knowledge, it does not describe specific reasoning methods that are required in the application and development of strategic knowledge. Strategic knowledge, among other things, in-

volves identifying solutions to complex, multi-dimensional problems from a potentially infinite solution space. Strategic knowledge manifests itself in innovation, whether the innovation is the development of or improvement in a business process, a product design, or a business model. It is also similar to the vision of scientific thinking at the heart of Polanyi's (1966) project. For Polanyi, scientific thinking is quintessentially about selecting a hypothesis to explain a phenomenon or a research question out of an open-ended set of options based on inarticulable intuitions and commitments. This intuitive grasping for hypotheses and the latching on to research trajectories is only later formalized into logically structured theories.

This line of Polanyi's (1966) thought is not fully accounted for in Collins' (2010) classification, as his deepest category of tacit knowledge is social, not cognitive. Other sources of insight are needed to complement the picture. The cognitive neuroscientific accounts of expertise cited earlier in the paper can be expected to shed light on the sort of inarticulable knowledge behind innovative and scientific thinking by identifying the relevant mechanisms operating within the brain. As they do, we can envision adding another category to Collins' classification, neural tacit knowledge or NTK. It could be situated between STK and CKT in the classificatory scheme for a number of reasons. First, intelligence is not only based in the brain but is constituted by the nervous system as well, while STK itself is not simply a matter of the body but also the brain and nervous system. Second, our brains are formed in part by our socialization (Feldman Barret, 2017). They are both the interface to and instantiation of social knowledge for each of us.

A relevant point of reference for elucidating Polanyi's (1966) core conception of tacit knowledge can be found in the work of an earlier philosopher, Charles Sanders Peirce (1931). Peirce's theory of abductive inference as a form of reasoning provides a logical framework into which Polanyi's (1966) conception of personal knowledge in science may fit well. As Larson (2021, pp. 160-173) argues, Peirce developed the concept of abduction to account for the formation of hypotheses. Before an inductive inference can be made, a plausible, relevant covering rule has to be selected from countless many. The inference to such a rule starts from a particular set of facts that need to be explained. Only after a plausible candidate hypothesis has been identified can statistical confirmation proceed.

As a form of reasoning, abduction is distinct from both deductive and inductive reasoning (Larson, 2021, p. 171). It involves the selection of a hypothesis out of a potentially infinite hypothesis space. As Larson notes, deduction is truth-preserving, while induction generalizes from a set of facts. Abduction, by contrast, reasons from a particular fact or event to a cause or interpretation. Both deduction and induction have distinct logical structures. Abduction, if formalized, appears as a fallacious deductive schema, namely, affirming the consequent ( $p \rightarrow q, q \mid p$ ) (pp. 170-172). In reality, the conditional is understood not in truth-functional terms but in causal or other explanatory terms. This requires background knowledge of the world, whether common knowledge or specialized theoretical knowledge. It also requires basic cognitive skills that allow us to recognize causal relations and determine semantic relevance. No doubt, our conceptual frameworks depend upon a linguistic community, but they also depend on cognitive abilities that developed as part of our evolution as social animals.

For our purposes, the main takeaway from these observations is that tacit knowledge, as Polanyi (1966) originally conceived it, cannot be understood independently of basic conceptions of human intelligence. As characterized by Landgrebe and Smith (2023) (building from Max Scheler), human intelligence is a combination of primary intelligence (shared with animals) and objective intelligence. Primary intelligence consists of the ability to respond to new situations without needing to learn how. It is organically bound and, in animals, limited to their specific environment. Objective intelligence, by contrast, involves the ability to abstract from specific environments and engage conceptually and imaginatively. These capabilities are fully integrated in humans and are required for and developed by language culture (Landgrebe & Smith, 2023, pp. 41-48). Deep tacit knowledge then needs to be understood in terms of human intelligence at the level of individual brains and collective culture. This point brings us back to the question of artificial intelligence and its relation to KM.

### *TACIT KNOWLEDGE, AI, AND KM*

In their paper, “Artificial intelligence and knowledge management: questioning the tacit dimension,” Sanzogni et al. (2017) use Collins’ (2010) threefold classification to answer the question of how far AI can go in automating tacit knowledge. They conclude that much RTK and some STK can be automated, but CTK cannot.

Some, but not all, relational tacit knowledge can be explicated and therefore automated using, for example, expert systems and/or neural networks to extract corporate social responsibility values from company documents and match these values to the financial outcomes of a company. Somatic tacit knowledge can also be explicated and mechanized using neural networks (Hidayati et al., 2016) and robotics through, for example, conversational agents who provide interactive advice to patients on a range of issues ... Collective tacit knowledge, however, cannot be codified, modelled, or mechanized. Machines cannot socialise or be meaningfully embedded in a social milieu since humans and machines are different in kind and materially. (Sanzogni et al., 2017, p. 43)

Their argument does not add much to Collins’ (2010) conclusion and misses some critical points. RTK can be converted to explicit knowledge in part but not in its entirety. So converted, it consists of what Collins (2010, p. 81) calls ‘strings’ in a natural language. These can be explicated in one of four meanings of the term Collins’ assigns. The first of these meanings is ‘explicable by elaboration,’ according to which a string S can be given its full meaning through a longer string, S. They can be explicated by transformation insofar that they can be converted into computer code. So serialized, they still need to be processed. Expert systems provide one model of processing, and I have discussed their limitations throughout the paper. Neural networks provide another model, though only some types will deal with explicit knowledge (EK). LLM’s are a prime example. Whatever the paradigm, expert systems, or LLMs, the explicated (codified) strings have to fit deductive or inductive inference structures. RTK, therefore, will have to go through multiple transformations before it is useful for AI automation, and the specific type of AI automation may or may not be even a partial specification of the relevant tacit knowledge, in this case, RTK. This is because AI algorithms may not model the human cognitive processes at all, as is the case with LLMs.

Turning to STK, two other senses of explicability may apply. For basic tasks such as painting a chair, a representation of the task may be created for purposes of automation, as is typical with stationary robots on assembly lines. This is explicable through Collins’ (2010) third form of explication, mechanization. Robotic process automation (RPA) is another example. RPA essentially consists of recording the clicking and typing motions of users. Where mechanization is at issue, the tacit knowledge of the individual is not extracted, but it is copied. So, a parasitic (though not extractive) relationship does, in fact, exist.

More skilled activity may not be imitable in this way, but it may be given an explanation in physics, physiology, or another relevant discipline. This is explication by explanation. In the case of explication by explanation, the inarticulate STK is not extracted. Rather, an external theory accounts for it. But, as Collins (2010, p. 79) emphasizes, the theoretical explanation will be of little use to the practitioner. A rugby player will not likely benefit from an explanation of the equations in calculus describing a ball’s trajectory, though tips from a coach can be quite helpful. These nuances are missing from Sanzogni et al.’s (2017) account, and as a result, they leave unclear the ways in which AI can be expected to automate tacit knowledge in a manner relevant to KM.

This is important because, if the research project of AI is successful, it will provide partial explications of cognitive processes in the fourth sense of the term, the explanatory sense. However, these explanations will be similar to neuro-physical explanations of bodily movement and skill, which are inarticulate to individuals. For this reason, I coined the term ‘NTK’ to refer to the inarticulate nature of our own thinking processes. These explanations, however, will not fall within the paradigm of extracting/converting tacit knowledge from individuals. Rather, they will constitute a partial scientific

explanation of intelligence and cognition. This is the goal of AI as a scientific endeavor as opposed to a set of engineering areas. It is not the project of KM.

The upshot of the above reflections on the role of tacit knowledge in procedural and strategic knowledge, and their relation to AI and KM, is to show that KM's original view of tacit knowledge as a resource for conversion into explicit knowledge is oversimplified and, as such, lends itself to extractive visions. Conversion of tacit knowledge to explicit knowledge still has a place in KM strategy, though its greatest benefits may simply be in the creation of documents that will be read (and internalized) by other humans. Extraction for automation purposes will be limited. AI/ML thrives on the use of data that is not extracted from knowledge workers' heads but collected from transactions and data generation. As KM endeavors to leverage AI, it needs to reappraise its understanding of tacit knowledge in order to fully understand how knowledge work can be enhanced by AI, not automated. A return to Polanyi's (1966) conception of tacit knowledge, which was distorted during its absorption into KM, and continuing research of tacit knowledge promise to provide a more fruitful discussion of the relation between human and artificial knowledge in the era of AI. It promises to support a human-enabling vision of AI in KM, which will lead to a more robust and independent engagement with AI technologies than one that is merely supportive of AI automation strategies.

## CONCLUSION

---

This paper re-examined an argument concerning knowledge management by Steve Fuller, who argued that KM is an exploitative management-technical strategy that aims to extract knowledge from knowledge workers in order to create intellectual property for the organization while alienating workers from their human capital. Fuller's argument was based on the potential capabilities of expert systems and knowledge engineering. This paper re-examined Fuller's argument in light of current developments in artificial intelligence as they relate to knowledge management technology. The paper then argued that current AI technologies do pose a threat to knowledge work as replacement technologies but that knowledge management has the conceptual resources to provide an alternative, human (knowledge) vision of knowledge of work. The resources lie in a reconceptualization of tacit knowledge, which is central to knowledge management strategy. This reconceptualization revives Polanyi's (1966) analysis of tacit knowledge as standing in a structural relationship with explicit knowledge in a way that supports the view that tacit knowledge is an essential element of knowledge work and a barrier to automation/replacement strategies. In addition, Collins' (2010) threefold distinction between RTK, STK, and CTK further articulates the concept of tacit knowledge in such a way as to make clearer the limitations of automation strategies. Finally, the paper argues that Peirce's (1931) conception of abductive reasoning elucidates and complements Polanyi's (1966) view of the role of tacit knowledge in scientific discovery and explains further why innovative thinking and insight are not reducible to formalizable reasoning procedures. Thus reconceptualized, the paper shows that tacit knowledge can play a critical role in KM research and strategy that makes human knowledge central to future attempts to integrate AI capabilities into KM systems and practices, and it can do so in a way that stands in stark contrast to Fuller's extractive view of KM.

## *FUTURE DIRECTIONS*

KM emerged in conjunction with the development and widespread use of information technologies such as databases, document repositories, and the Internet. A central concern was how to elicit tacit knowledge to fully leverage these technologies as drivers of knowledge in organizations. The emergence of modern AI is sparking a re-engagement with information technologies and requires a re-thinking of KM's epistemological assumptions. This paper has explored the role of tacit knowledge in KM in relation to AI. Its conclusions suggest a number of related research directions:

As noted in the article, AI presents challenges to KM that require a deeper engagement with cognitive science, epistemology, and knowledge studies. Tacit knowledge and its underlying cognitive mechanisms need to be further investigated within the context of KM and its concerns. A taxonomy

of tacit knowledge and supporting cognitive structures needs to be investigated and further articulated.

The role of different types of tacit knowledge emerging from further research should be mapped to different types of knowledge work and knowledge processes, and their role in supporting distinct types of knowledge work use cases should be explored. Further conceptual and supporting empirical research is needed.

The relation to KM and automation should be investigated historically and in relation to current technical developments. KM is historically related to a set of information and communication technologies. How automation technologies relate to KM's objectives presents a broad and important topic of research.

Do some forms of automation fit better within KM paradigms, e.g., knowledge graphs and ontologies, while others fall outside of the main research interests of KM, e.g., RPA?

How does AI augment knowledge work and enable knowledge workers? Do certain forms of AI enrich human knowledge work and help develop tacit knowledge?

Does KM provide a framework for researching and implementing augmentative AI as opposed to replacement automation?

## REFERENCES

---

- AIIM. (2019). *How to fit artificial intelligence into your information management strategy: 5 key use cases for AI in information management*. <https://info.aiim.org/how-to-fit-artificial-intelligence-into-your-information-management-strategy>
- Bercerra-Fernandez, I., Sabherwal, R., & Kumi, R. (2024). *Knowledge management: Systems and processes in the AI era*. Routledge. <https://doi.org/10.4324/9781003364375>
- Bilalić, M. (2017). *The neuroscience of expertise*. Cambridge University Press. <https://doi.org/10.1017/9781316026847>
- Brynjolfsson, E., Li, D., & Raymond, L. (2023). *Generative AI at work*. arXiv. <https://doi.org/10.48550/arXiv.2304.11771>
- Cawsey, A. (1998). *The essence of artificial intelligence*. Prentice Hall.
- Clark, R. C. (2008). *Developing technical training: A structured approach for developing classroom and computer-based instructional materials* (3rd ed). Pfeiffer/Wiley.
- Collins, H. (2010). *Tacit and explicit knowledge*. University of Chicago Press. <https://doi.org/10.7208/chicago/9780226113821.001.0001>
- Collins, H., & Evans, R. (2007). *Rethinking expertise*. University of Chicago Press. <https://doi.org/10.7208/chicago/9780226113623.001.0001>
- Drucker, P. (1988). The coming of the new organization. *Harvard Business Review*, 66, 45–53.
- El Naqa, I., & Murphy, M. J. (Eds.). (2022). *Machine and deep learning in oncology, medical physics and radiology* (2nd ed.). Springer. <https://doi.org/10.1007/978-3-030-83047-2>
- Feldman Barret, L. (2017). *How emotions are made: The secret life of the brain*. Harper Collins.
- Fuller, S. (2002). *Knowledge management foundations*. Butterworth-Heinemann.
- Gartner. (n.d.). *Hyperautomation*. <https://www.gartner.com/en/information-technology/glossary/hyperautomation>
- Goldman, A. (2018). Expertise. *Topoi*, 37, 3-10. <https://doi.org/10.1007/s11245-016-9410-3>

- Hidayati, R., Kanamori, K., Feng, L., & Ohwada, H. (2016). Combining feature selection with decision tree criteria and neural network for corporate value classification. In H. Ohwada, & K. Yoshida (Eds.), *Knowledge management and acquisition for intelligent systems* (pp.31–42). Springer. [https://doi.org/10.1007/978-3-319-42706-5\\_3](https://doi.org/10.1007/978-3-319-42706-5_3)
- Karim, M. R., Comet, L. M., Shajalal, M., Beyan, O., Rebholz-Schuhmann, D., & Decker, S. (2023). *From large language models to knowledge graphs for biomarker discovery in cancer*. arXiv. <https://doi.org/10.48550/arXiv.2310.08365>
- LaFayette, B., Curtis, W. C., Bedford, D. A. D., & Iyer, S. (2019). *Knowledge economies and knowledge work* (1st ed.). Emerald Publishing Ltd. <https://doi.org/10.1108/9781789737752>
- Landgrebe, J., & Smith, B. (2023). *Why machines will never rule the world: Artificial intelligence without fear*. Routledge. <https://doi.org/10.4324/9781003310105>
- Larson, E. J. (2021). *The myth of artificial intelligence why computers can't think the way we do*. Belknap Press. <https://doi.org/10.56315/PSCF12-21Larson>
- Lytras, M. D., Russ, M., Maier, R., & Naeve, A. (Eds.). (2008). *Knowledge management strategies: A handbook of applied technologies*. IGI Global. <https://doi.org/10.4018/978-1-59904-603-7>
- Majumder, S., & Dey, N. (2022). *AI-empowered knowledge management*. Springer. <https://doi.org/10.1007/978-981-19-0316-8>
- Mayo, R. C., & Leung, J. (2018). Artificial intelligence and deep learning – Radiology’s next frontier? *Clinical Imaging*, 49, 87–88. <https://doi.org/10.1016/j.clinimag.2017.11.007>
- McKendrick, J. (2023, April 18). *The state of knowledge management in 2023: Untapped potential for business value*. KMWorld. <https://www.kmworld.com/Articles/Editorial/Features/The-State-of-Knowledge-Management-in-2023-Untapped-Potential-for-Business-Value-158163.aspx>
- Mooradian, N. (2005). Tacit knowledge: Philosophic roots and role in KM. *Journal of Knowledge Management*, 9(6), 104–113. <https://doi.org/10.1108/13673270510629990>
- Mooradian, N. (2007). Computerizing knowledge: Ethical and philosophical dimensions of knowledge management. *Proceedings of the 7th International Conference of Computer Ethics*, 104–113. [https://www.researchgate.net/publication/351333697\\_Computerizing\\_Knowledge\\_Ethical\\_and\\_Philosophical\\_Dimensions\\_of\\_Knowledge\\_Management](https://www.researchgate.net/publication/351333697_Computerizing_Knowledge_Ethical_and_Philosophical_Dimensions_of_Knowledge_Management)
- Mooradian, N. (2021). Ethics. In P. C. Franks (Ed.), *The handbook of archival practice* (pp. 11–15). Rowman & Littlefield.
- Müller, M., Terziev, G., Metternich, J., & Landmann, N. (2020). Knowledge management on the shop floor through recommender engines. *Procedia Manufacturing*, 52, 344–349. <https://doi.org/10.1016/j.promfg.2020.11.057>
- Nissen, M. (2003). Inducing enterprise knowledge flows. In J. N. D. Gupta, & S. Sharma (Eds.), *Creating knowledge-based organizations* (pp. 185–202). IGI Global. <https://doi.org/10.4018/978-1-59140-162-9.ch009>
- Nonaka, I., & Takeuchi, H. (1995). *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation* (1st ed.). Oxford University Press. <https://doi.org/10.1093/oso/9780195092691.001.0001>
- Patton, S. (2022, October 12). *Microsoft Syntex: Content AI integrated in the flow of work*. Microsoft. <https://www.microsoft.com/en-us/microsoft-365/blog/2022/10/12/welcome-to-microsoft-syntex-content-ai-integrated-in-the-flow-of-work/>
- Peirce, C. S. (1931). *Collected papers of Charles Sanders Peirce* (C. Hartshorne, P. Weiss, & A. W. Burks, Eds.). The Belknap Press of Harvard University Press.
- Polanyi, M. (1958). *Personal knowledge: Towards a post-critical philosophy*. University of Chicago Press. <https://doi.org/10.1017/S0031819100066110>
- Polanyi, M. (1966). *The tacit dimension* (1st ed.). Doubleday.
- Ribiere, V., Gong, C., & Yang, K. (2021). Knowledge management from a technology perspective. In J. Liebowitz (Eds.), *A research agenda for knowledge management and analytics* (pp. 43–66). Edward Elgar Publishing.

- Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach*. (4th ed.). Pearson.
- Sanzogni, L., Guzman, G., & Bush, P. (2017). Artificial intelligence and knowledge management: Questioning the tacit dimension. *Prometheus*, 35(1), 37–56. <https://doi.org/10.1080/08109028.2017.1364547>
- Singhal, K., Azizi, S., Tu, T., Sara, M. S., Wei, J., Chung, H. W., Scales, N., Tanwani, A., Cole-Lewis, H., Pfohl, S., Payne, P., Seneviratne, M., Gamble, P., Kelly, C., Scharli, N., Chowdhery, A., Mansfield, P., Aguera y Arcas, B., Webster, D. A., & Corrado, G. S. (2023). Large language models encode clinical knowledge. *Nature*, 60, E19. <https://doi.org/10.1038/s41586-023-06455-0>
- Szulanski, G. (2003). *Sticky knowledge: Barriers to knowing in the firm*. Sage Publications. <https://doi.org/10.4135/9781446218761>
- Winch, C. (2012). *Dimensions of expertise: A conceptual exploration of vocational knowledge*. Bloomsbury Publishing.

## AUTHOR

---



**Norman Mooradian** has had a multifaceted career working as an information professional and an academic focusing on digital ethics and knowledge studies. He received his Ph.D. and M.A. in Philosophy from the Ohio State University. Mooradian has published articles in areas such as information ethics, business ethics, information privacy, knowledge management, enterprise content management, records management, and virtual reality. He is the author of *Ethics for Records and Information Management* (2018, ALA). His research and teaching interests connect concepts from digital ethics and philosophy of information to the information sciences. A central focus of his research is the intersection of ethics and knowledge in organizations and the knowledge society.

Prior to joining the faculty in the School of Information at SJSU, he was the Customer Education Manager at Konica Minolta in the Intelligent Information Management Division and worked in the IIM/ECM field for over 20 years in a variety of roles. Norman has been active in the information field, serving on standards committees for trusted systems, and has presented on numerous occasions to professional association groups such as ARMA, AIIM, ASIS&T, and ALA.