The tyranny of tacit knowledge: What artificial intelligence tells us about knowledge representation

Kurt D. Fenstermacher
MIS Department, University of Arizona
KurtF@Eller.Arizona.edu

Abstract

Polanyi's tacit knowledge captures the idea "we can know more than we can tell." Many researchers in the knowledge management community have used the idea of tacit knowledge to draw a distinction between that which cannot be formally represented (tacit knowledge) and knowledge which can be so represented (explicit knowledge). I argue that the deference that knowledge management researchers give to tacit knowledge hinders potentially fruitful work for two important reasons. First, the inability to explicate knowledge does not imply that the knowledge cannot be formally represented. Second, assuming the inability to formalize tacit knowledge as it exists in the minds of people does not exclude the possibility that computer systems might perform the same tasks using alternative representations. By reviewing work from artificial intelligence, I will argue that a richer model of cognition and knowledge representation is needed to study and build knowledge management systems.

1. Introduction

If we are to build and study knowledge management systems, we must understand something of the nature of knowledge. While philosophers might restrict their enquiries into the nature of knowledge to the mind, those who study how organizations collect, generate and share knowledge must concern themselves with questions beyond the mind. For example, if we wish to develop systems that capture the decision-making processes of experts to enable their insights to more easily share their expertise, we must understand something both of how experts think and how to best capture that thinking in information systems. This dichotomy — understanding how people think and how best to capture that thinking in machines — naturally suggests two different tacks. One, we could distinguish between *human knowledge* and *machine knowledge* and then develop theories that bridge the two theories of knowledge. Or, two, we could develop a single theory that explains *knowledge* in both people and machines. Moreover, given that people hold much of the relevant knowledge we seek, we must determine how to represent that knowledge in a computer.

Because many researchers in the knowledge management (KM) field begin by appealing to philosophical work in epistemology, I begin with a review of the most common philosophical theories cited in the KM literature. Next, I review the knowledge model presented in several papers from the KM literature. Following that analysis of the KM literature, I discuss some of the shortcomings inherent in applying the common philosophical concepts to KM problems. I then offer a survey of relevant work in artificial intelligence (AI). Finally, I conclude by appealing for a more nuanced notion of knowledge and an expanded role for computers in managing knowledge.

2. Historical definitions of knowledge

The philosophical field of epistemology has long struggled with defining the term *knowledge* and so many authors in knowledge management have looked to epistemology for guidance in describing knowledge. Unfortunately, despite debating the topic for millennia, philosophers have yet to agree on a definition of knowledge themselves; several perspectives have proven useful to researchers discussing knowledge management. A common starting point in epistemology is the Socratic dialog *Theaetetus* [1], in which Socrates, Theodorus, and Theaetetus discuss the nature of knowledge and conclude that knowledge is "justified, true belief."

What then is justified true belief? Working backwards, I will start with "belief". Belief (also called *opinion* in some translations of the dialog) can be anything. You can believe that I am the King of England; in this case, your belief would be false, but it is a belief nonetheless. A true belief is a belief that you hold and is an accurate representation of the world. I can believe that the Los Angeles Lakers won the 2004 Western Conference Finals and because they did so, I hold a true belief. Socrates adds justification to ensure that "knowing" cannot occur by chance. For example, I might flip a fair coin and say that, "I know it will come up 'heads'." Even if the coin does indeed land 'heads', Socrates (and many later philosophers) would say that I did not *know* that it would

0-7695-2268-8/05/$20.00 (C) 2005 IEEE

land 'heads' because I could not have justified my claim. I might believe my claim; my claim even turned out to be true. However, true belief alone is not enough. I must also be able to explain logically why I believed a statement to be true.

Leaping ahead more than two thousand years brings us to Gilbert Ryle, who in his 1949 work, *The Concept of Mind* [2], distinguished between *knowing-that* and *knowing-how*. In Ryle's terms, knowing-that means storing and recalling facts, e.g., "Ronald Reagan was the fortieth President of the United States." Knowing-how is practical knowledge, e.g., how to tell a joke or diagnose an illness. For Ryle, this was an important distinction in his argument against Cartesian dualism [3], which claims that the mental and physical worlds are fundamentally different. In the Cartesian view, an expert's mind would dictate to the body needed physical actions and the body would send perceptions to the mind. In Ryle's view, experts practicing their craft demonstrate know-how and do so without conscious reflection. Ryle's distinction between know-how and know-that carries through to Michael Polanyi's theory of personal knowledge.

Undoubtedly, the most popular philosophical perspective in knowledge management is Polanyi's theory of personal knowledge. While the idea that "we know more than we can tell" [4] has been around for centuries, Polanyi's term for this "tacit knowledge" has since been popularized by Nonaka's study of knowledge-creating companies in Japan [5, 6]. In his 1997 work, Polanyi offers the example of face recognition: "We know a person's face, and can recognize it among a thousand, indeed among a million. Yet we usually cannot tell how we recognize a face we know. So most of this knowledge cannot be put into words." Polanyi continues by stating the police (who often need to solicit facial descriptions from witnesses) have worked to address this difficulty:

> [T]he police have recently introduced a method by which we can communicate much of this knowledge. They have made a large collection of pictures showing a variety of noses, mouths, and other features. From these the witness selects the particulars of the face he knows, and the pieces can then be put together to form a reasonably good likeness of the face. This may suggest that we can communicate, after all, our knowledge of a physiognomy, provided we are given adequate means for expressing ourselves. [4]

Another example commonly given is our knowledge of how to ride a bicycle. Riding a bicycle is clearly a skill that requires knowing how (to use Ryle's terminology). Because I can engage in the skill, I have the knowledge (it is "specifiably known") in Polanyi's words. However, I would be hard-pressed to write down detailed instructions that you could follow to ride a bicycle (and hence it is "tacit").

In Polanyi's view, there is also explicit knowledge, which is similar to Ryle's knowing-that. Explicit knowledge consists of theoretical knowledge, facts and other elements that you are consciously aware of as you think. To continue with the bicycle example, I could write down a list of the rules that are relevant to riding a bicycle, for example, "The faster you pedal, the faster you'll go," and "It's easier to stay upright if you have a little speed." Such rules are examples of both Ryle's know-that and Polanyi's explicit knowledge (although these two are not synonymous [7]).

Polanyi makes the case for tacit knowledge by studying human perception in the Gestalt tradition. His baseline examples of tacit knowledge stem from the observation and analysis of perceptual actions. The ability to recognize a face in its totality but the inability to describe how the face's components lead to such recognition is a classic illustration of Gestalt phenomena. In addition to face recognition, Polanyi also describes an experiment in which subjects received an electric shock whenever the experimenters presented certain nonsense syllables. The subjects came to anticipate the shocks and reacted upon seeing the syllables, but the subjects could not name the syllables when questioned. In Polanyi's view, the subjects had developed tacit knowledge of the linkage between the syllables and the shocks.

More specifically, Polanyi argues that tacit knowledge is linked to the act of attending (in the sense of attention) to aspects of perception. He distinguishes between proximal and distal aspects (by analogy to the adjectives proximal and distal from anatomy) of attention by stating that the proximal is that which we associate with ourselves, while the distal is the world beyond ourselves. Polanyi offers the example of the probe a blind person uses to explore the world. When the person first uses a probe, he has only distal knowledge of it; the probe is simply a stick held in the hand. Over time, as he learns to use the probe, it becomes an extension of him and his knowledge shifts from the distal to the proximal. Together, the proximal and distal (each a kind of knowing) compose tacit knowledge.

In the syllable-shock experiment, the subjects had proximal knowledge of the shock syllables and the shock associations and had distal knowledge of the shock itself. In Polanyi's terms, the subjects *attend to* the distal term (the shock) while *attending from* the proximal term (the shock syllables and the association). The subjects "know" the shock syllables in the sense that they react differently to them than non-shock syllables, but they do not consciously recognize them. Instead, they have knowledge of the syllables in one context (the experimental setting) and for one purpose (anticipating the shock).

What do electrical shocks have to do with, say, capturing the expertise of seasoned consultants? While Polanyi's theory of personal knowledge has its roots in the analysis of perception, he claims a much broader sweep. Polanyi contends that "by elucidating the way our bodily processes participate in our perceptions we will throw light on the bodily roots of all thought, including man's highest creative powers." [4] While Polanyi's argument for this connection is complex, the essence of it is that to know something we must make it a part of us in the same way that we must integrate our perception of the world, a process that Polanyi labels *interiorization*.

Suppose that you have heard other researchers discuss the Technology Acceptance Model (TAM), but do not know the model yourself. A brief search would reveal Davis's 1989 *MISQ* article [8] as the seminal work on the topic and so you might begin by reading his article. At this point, you have acquired explicit knowledge of the topic. You could list the two key elements in technology adoption and define each in your won words. You could then extend your understanding by reading articles in which other researchers adopted the TAM, tested its predictions, or criticized its assumptions. If you wish to adopt the TAM for your own work, however, you must learn to see the world though the lens of the TAM. Once you have begun to see the world in this light, but without consciously thinking of the TAM and its principles, you have interiorized the theory — you no longer see the theory itself, but only the world through the lens of the theory.

My use of the words "see" and "lens" is intentional in the previous paragraph. Polanyi uses the analogy to perception to justify the extension of tacit knowledge from a theory or perceptual knowledge to a broader theory of knowing. An expert in TAM is one who has interiorized TAM. Prior to this interiorization, you had to rely on general-purpose reasoning skills and the individual elements of the theory, puzzling out its application in each new instance — the hallmark of a novice in Polanyi's conception.

While the foregoing is a very highly condensed introduction to epistemology, the above conceptions of knowledge (i.e., "justified true belief", knowing-that vs. knowing-how and tacit knowledge) are the most common philosophical starting points in the knowledge management literature. In Section 3, I will describe the knowledge models commonly presented and in the following section discuss some of the benefits and shortcomings of these models.

3. Current models of knowledge

Before considering how to research, design, and develop knowledge management systems, we should consider the elusive question "What is knowledge?" Because knowledge is central to human endeavor, every field (and likely every person) has contemplated the question at some point and so much has been said and written already. In this paper, I will focus primarily on two related disciplines that have considered the nature of knowledge and its relation to information technology: management information systems and information science. In the following review, I consider several representative papers by first summarizing the knowledge model employed and then stating the implications of such a model on the design of a knowledge management system (KMS). The current section addresses conceptual models of knowledge and KMSs, while the following section reviews work that both describes implemented KMSs and posits reasons for their success or failure.

3.1. What's in a model?

As used here, a knowledge model should address several issues that bear on managing knowledge, which support the broader question of "What is knowledge?" First, a knowledge model should offer a working definition of knowledge, for example, Plato's definition in *Theaetetus* [1], "justified, true belief". Ideally, the definition should explain the structure of knowledge. Second, while I am not convinced of the usefulness of distinguishing data, information and knowledge, most authors in the field explain how they distinguish these three concepts. By studying each author's definitions, we can learn something of her definition of knowledge. Third, authors often draw a distinction between a person's capacity for knowledge and a computer's. Fourth, the knowledge model should offer insight into the nature of knowledge management and KMSs. At a minimum, the model should include a definition of knowledge management that in turn draws on the author's definition of knowledge.

The review of knowledge models that follows is not exhaustive, but is instead representative of authors with differing perspectives and backgrounds, publishing in different venues. In some of the earliest work in the field of knowledge management, Alavi and Leidner [9] considered the nature of knowledge management broadly, outlining a knowledge model and surveying common practices in knowledge and conclude by making suggestions for the direction of research in the knowledge management field. In a brief survey of definitions of knowledge from the knowledge management literature, Spiegler [10] argues that earlier authors have offered definitions of knowledge that fail to distinguish it from information. Moreover, they fail to establish a clear vision for the field of knowledge management that differentiates it from previous concepts, including management information systems (MIS), data management, and information management. In a more recent article, Blair [11] assesses the state of KM research and practice.

In their 1999 *CAIS* article "Knowledge management systems: Issues, challenges, and benefits" [9], Alavi and Leidner begin by noting that organizations have worked to capture and disseminate knowledge for many years, but information technology can expand the means for sharing knowledge. They begin with a definition of knowledge that draws on earlier work by Huber and Nonaka as "a justified personal belief that increases an individual's capacity to take action."

The two key elements of this definition are that knowledge is personal and that knowledge enables action. The idea that knowledge is personal leads directly to Alavi and Leidner's claim that knowledge exists solely in the mind of a person. The second element, that knowledge enables action, should be read broadly as action in this context means physical action, thinking, and a combination of the two (for example, driving a car). Under this definition, everything is an action and so the more interesting claim is that of the subjective nature of knowledge.

While many authors view data, information, and knowledge as varying along a continuum, Alavi and Leidner see a sharp break between information and knowledge. Knowledge is the result of one person's processing of information and so has a subjective, personal quality that information does not. What appears in books, is stored in computer memories, or explained in lectures is not knowledge but information. For example, I know how to use a biscuit joiner (a woodworking tool) to join two pieces of wood. I learned to do this by reading books and manuals, which are information, and experimenting with wood and tools. According to Alavi and Leidner, I cannot share this knowledge with you directly through the written word or even by standing at your side as you work. Instead, I can only render my knowledge as information and you must then internalize that information and thereby create your own knowledge.

Alavi and Leidner see knowledge management as a process, defined as "a systemic and organizationally specified process for acquiring, organizing and communicating both tacit and explicit knowledge of employees so that other employees may make use of it to be more effective and productive in their work." Under this definition, knowledge management is much broader than a computer system (and rightly so); knowledge management could include guidelines for documenting work, regular cross-departmental lunches, and a weekly departmental newsletter.

The authors then define a KMS specifically as "information systems designed specifically to facilitate the codification, collection, integration, and dissemination of organizational knowledge". While Alavi and Leidner do not discuss the interplay between knowledge management and technology in detail, they do cite Vandenbosch and Ginzberg's study of Lotus Notes deployment [12] to support the claim that "in the absence of an explicit strategy to better create and integrate knowledge in the organization, computer systems which facilitate communication and information sharing have only a random effect at best."

Alavi and Leidner's contention that knowledge exists only in the minds of people indeed has important implications for the design of KMSs. Because knowledge cannot be stored in computers by definition, the focus of the technological component is on linking people together. Thus, a KMS should incorporate descriptions of individual's knowledge and should support searching for experts in a desired area. (The common term for such systems is "yellow pages", by analogy to American business directory listings by industry or skill, which are usually printed on yellow paper.) Alavi and Leidner would further add that it is important that the descriptions of experts' knowledge be accessible to likely searchers. For example, an electrical engineer who served as an expert on the semiconductor industry at a consulting firm would need a description that non-engineers would understand — using terms like "chip" and "semiconductor" rather than "wafer" and "molecular beam epitaxy".

A year after Alavi and Leidner's article, Israel Spiegler analyzes knowledge management in his *CAIS* article "Knowledge management: A new idea or a recycled concept?" [10]. After listing many definitions of knowledge, Spiegler offers a list of knowledge "components" drawn from Dreyfus [13] and Wittgenstein [14]: context, experience, basic truths, best practices, common sense, judgment, rules of thumb, values and belief, needs, emotions, desires, socializing into a culture. Spiegler then defines knowledge as, "[T]he process of knowing, a reflexive process that takes data and information, in a social context, together with the [components of knowledge], and generates new data, information, and/or knowledge."

While Spiegler surveys many definitions of data and information, he does propose a definition for either himself. His survey concludes that information is data with value and knowledge is information coupled with insight. Without a clear definition of data, value, or insight, however, it is impossible to analyze Spiegler's differentiation of data, information, and knowledge.

Spiegler does discuss the differing nature of knowledge with respect to computers and people. He begins by stating that technology is not a substitute for knowledge, which is surely the case. While knowing is a process, technology is a conduit for data and information. Spiegler claims unequivocally that technology can neither replace nor generate human knowledge. In fact, Spiegler appeals to the literature in artificial intelligence (AI) claiming, "Attributing knowledge to humans rather than to machines is a frequent discussion in AI…. Humans deal with and possess knowledge whereas machines handle the representations of knowledge, at least one step lower in the abstraction of reality. This level is really data or information." This comment suggests that a common view in AI holds that a representation of knowledge is not knowledge, but simply data or information. Because Spiegler does not give any references to support his characterization of the AI literature, it is difficult to investigate his support for this claim. I would argue, however, that Spiegler misrepresents the prevailing view in AI; indeed, the perspective he offers is that of AI's critics, not its practitioners. Many AI researchers have argued for decades that the brain itself stores a representation of knowledge [15-17], albeit biochemically rather than in the silicon of digital computers. I will discuss the interplay between AI and knowledge in more detail in a later section.

While Spiegler never offers a definition of knowledge management, his criticism of others' definitions gives some insight into how he would define the field. Drawing on his definition of knowledge and his commentary, we can infer some properties of Spiegler's perspective on knowledge management. First, because knowledge is the *process* of knowing, one cannot maintain a store of knowledge because knowledge is dynamic, not static. This view supports the use of knowledge maps [18], which are similar to the yellow pages mentioned by Alavi and Leidner [9] — because knowledge is not static it cannot be captured in a store, thus the best strategy is to point at the people and organizations who process the knowledge.

Second, knowledge is reflexive. While Spiegler does not explain the notion of reflexivity, he appears to use it as a synonym for Schön's notion of reflection [19] — people can use knowledge to consider their own knowledge in a recursive process. This suggests that organizations that wish to manage knowledge must encourage their members to pause and reflect.

Third, the process of knowing can act not only on itself, but also to synthesize data and information. The challenge that many database and information retrieval researchers confront is finding the right data or information to present to users. The process of knowing occurs in a social context, which suggests that knowledge changes as it moves from one context to another. To promote the sharing of knowledge then, organizations should work to capture the social context in which knowledge is generated and encourage sensitivity throughout the organization to the contextual nature of knowledge. Lastly, because the act of knowing generates new data, information, or knowledge, organizations should work to capture the linkage as well as the newly generated data, information, or knowledge.

David Blair asks "Knowledge management: Hype, hope, or help?" in his 2002 *JASIST* article [11] of the same name. Blair builds his notion of knowledge by starting with definitions of data and information. Taking his cue from Wittgenstein's observation that we should study language by its use [14], Blair offers examples of the differing usage of the terms data and information, which people use similarly, and knowledge, whose use is of a different character. People often speak of having and losing data and information, but only speak of losing knowledge when they refer to specific facts, which is not really knowledge at all, but data or information. Unfortunately, Blair never offers a clear definition of knowledge and his linguistic analysis is problematic because there are counter-examples to each of his examples. For example, many people say, "I know her phone number," but Blair's analysis suggests that a single phone number doesn't constitute knowledge. Moreover, people will often say, "I used to know the capital of Delaware, but I don't anymore."

Blair's argument suggests that knowledge is of two kinds: knowing-how and knowing-that. In Blair's conception, only knowing-how is true knowledge, while knowing-that is the same as having data or information. He notes that these two often go hand-in-hand: applying know-how requires knowing-that. Nevertheless, knowing-how is a critical component. If I lack know-how, but I have all the needed information, I cannot practice the skill. For example, I may have a copy of my car's auto repair manual, but that does not make me a mechanic. (Conversely, a mechanic who did not have my car's repair manual might be at a loss, despite her skill, given the unique features of today's modern cars.)

While he seems data and information differently from each other, they are both qualitatively different from knowledge. Data "are simply 'facts' and 'figures' that are meaningful in some way." Blair adopts Drucker's definition of information as "data that has been organized for a particular purpose" [20]. The distinction between the two is still fuzzy, however, since it is not clear how much organization is sufficient or how narrow a use is intended. For example, Blair does not consider a corporate phone book to be information, even though it is organized (perhaps in more than one way, for example, both alphabetically by last name and by department) to make it easy to find specific employees (admittedly a fairly broad use). However, a list of active customers by region, their contact information, and past purchases, prepared for the sales force to make follow-up calls would be information.

In distinguishing between people and computers, Blair states that both computers and people can have data and information, but "only a person can be knowledgeable, that is, only a person can have and exercise knowledge" [11].

Blair also criticizes those who believe that any human skill (including cognitive skills) can be described by a set of rules or procedures, saying that "there is evidence that in most areas of expertise, it does not hold" [11]. Moreover, Blair references Magee's [21] quote of philosopher of science Hilary Putnam, "It is a mistake to think that merely because one practices an activity one can give a theory of it." Blair offers the example of the wine connoisseur who can separate good from not-so-good wines, but can't say why. Finally, Blair cites work by Hubert and Stuart Dreyfus's book, *Mind over Machine* [22], as evidence that most expertise has a significant inexpressible component. According to Hubert and Stuart Dreyfus, the rules that *can* be expressed in written form support at best competence in a particular skill and not an expert level of that skill. Blair contends that it is very important for the KM field to distinguish between expert knowledge that has not been expressed, but could be, and expert knowledge that cannot be expressed.

Blair's split of knowledge into know-how and know-that suggests to him that knowledge management should address both of these by managing data and information (know-that) and the management of expertise (know-how). Blair is sanguine about KM's future success because, unlike expert systems (which attempt to replace human expertise) and decision support systems (which work to supplement human expertise), KM systems aim only to facilitate the use of human expertise.

While the above articles are not the whole of work on the nature of knowledge and its role in knowledge management, they offer a snapshot over several years of research on the question of knowledge. While each of the above papers has its own perspective, several common themes run through them and the broader literature. First, knowledge is difficult to define. While each of the above purports to offer a definition of knowledge, some never do offer a clear statement of what knowledge is. Second, there is general agreement that data, information, and knowledge are related, although perhaps not simply as a matter of degree. Third, knowledge is of two different kinds, know-how and know-that. Fourth, the knowledge of experts that makes them expert is tacit and cannot be captured in computers. Fifth, because the most valuable knowledge is restricted to people, knowledge management systems should provide pointers to expertise (e.g., knowledge maps [18] and "yellow pages"). While Blair appeals to Dreyfus (and Dreyfus), most KM authors draw on Polanyi's theory of personal knowledge, and his distinction between tacit and explicit knowledge, to support the argument that knowledge is a distinctly human phenomenon.

4. Problems in applying philosophical conceptions to knowledge management

Philosophers have debated the meaning of knowledge and other epistemological questions for millennia, and while appealing to their work is a good strategy, there are differences among philosophers. More importantly, the findings of epistemology may not be directly applicable to problems face by KM researchers and practitioners. For example, most businesspeople would take issue with the notion of knowledge as "justified, true belief" on closer examination. If I were to say, "Third-quarter profits in the Western region are up 15% over last quarter" this could meet the Socratic definition of knowledge. If indeed profits increased by the stated percentage, then I would hold a true belief. When pressed on how I know this to be true, I could answer, "Each store in the region reports its revenue and expenses, and the region's profit is simply the sum of the individual stores' differences between revenue and expense." Thus, I could justify my belief as true. Most people in business, however, would say I have simply stated a fact — a single datum. Why were profits up? Did we expect the increase to be larger? Does this portend a fourth-quarter rise?

Many authors studying knowledge in organizations expect that knowledge is true and that it can be justified, but there is also a notion that knowledge must add value. This is of course a slippery slope because it is difficult to say beforehand what will add value or what the value threshold is. If you, as the general manager of a large petrochemical plant, need to consult with a plant engineer regarding an impending disaster, then having contact "information" for the engineer at that moment could be very valuable. If the contact information could be shown to be accurate, then you would have a justified true belief that added value and suddenly the engineer's mobile phone number would be knowledge. Of course, few people (other than a Socratic epistemologist) would claim that a phone number is knowledge.

4.1. Know-how vs. know-that

One way out of this difficulty is to focus on the process of knowing. In most cases, it is the process of knowing — not knowing the cell phone number of an engineer, but knowing how to find her immediately —that adds value for organizations. We want experts to help devise solutions to new problems, not to simply recall their solutions to past problems. Blair highlights this aspect by emphasizing the importance of inference in knowledge. Focusing on the process reminds us of Ryle's distinction between knowing-how (the process) and knowing-that (specific facts). As discussed earlier, Ryle's aim was to discredit Cartesian dualism; he wished to demonstrate that one could not transform know-how into know-that. (Under Cartesian dualism, know-that could be the sole type of knowledge and through dualist interaction, the mind would then direct the body.)

However, a more measured view sees know-how and know-that as intertwined. For example, I cannot write a computer program (which surely requires know-how) without knowing the purpose of the program, its inputs and outputs, and the rules of the chosen programming language (which are all instances of know-that). A. J. Ayers argues for a similar linkage of know-how and know-that in *The Problem of Knowledge* [23].

Incidentally, the view that know-how and know-that are distinct but are in fact interrelated is much older than Ryle, having been advanced by Aristotle in his *Nicomachean Ethics* as the difference between *technê* (craft or practical knowledge) and *epistêmê* (theoretical knowledge) [24].

4.2. Polanyi's tacit knowledge

Polanyi's notion of tacit knowledge is not synonymous with Ryle's know-how in part because Polanyi sees tacit and explicit knowledge to be more tightly linked than Ryle's notions of know-how and know-that. However, Polanyi's theory of personal knowledge has two significant problems. First, there are many layers of perception and interaction among them, demanding a much richer model than simply distinguishing between tacit and explicit knowledge. There is no doubt that we are not conscious of many perceptual actions. For example, human vision is foveated, meaning that only a small region of our visual field is in focus at a given moment. However, neural processing "pieces together" high-resolution images gathered through the fovea so that our entire field of view appears to be in focus. You can observe this for yourself by focusing on a pencil directly in front of you and slowly moving the pencil out of your field of view (while continuing to look directly ahead). As you move the pencil, it will become blurry in your vision as it passes out of foveal range and will eventually disappear as you move it out of your field of view altogether. As you conduct this experiment, you will become aware of the foveated nature of your vision in a way that you would not by simply reading about it. However, regardless of how intensely you consider this phenomenon, you will not become conscious of the neural integration need to produce a complete, high-resolution view of your surroundings. To use Schön's term, you cannot reflect on the experience of integrating foveal vision. Perception is a complex topic that requires deeper investigation and richer models than those put forth by Polanyi.

Second, and more problematic for the knowledge management community, Polanyi's extension of his theory from perception to all of knowledge is problematic. Because Polanyi does not offer a structured theory of perception, he is unable to carefully map a clear path from perceptual processing to higher-level thought. Experts can "think aloud" as they solve problems, explain their reasoning, and compare solution methods for a problem. While some might argue that this is not sufficient to enable the codification of their knowledge, it certainly demonstrates that such knowing is of a different character that many perceptual activities.

While I see significant problems in applying Polanyi's theory of personal knowledge to expertise in organizations, I am not claiming that we should leave Polanyi behind. Instead, I argue that we should first understand the origin of Polanyi's theory and the implications that stem from its roots as a theory of perception, not expertise. Moreover, while the notion that "we know more than we can tell" is important, it is not by itself an argument against attempts to codify expert knowledge. For the KM community, the most critical flaw is that Polanyi fails to explain how we come to know.

4.3. How do we come to know?

If we are to manage knowledge and share expertise across organizations, then we must have a basis for understanding how people develop expertise. However, the knowledge management literature eschews theories on the development of expertise in favor of taxonomies of knowledge. While understanding types of knowledge is important, the irony is that the enterprise of managing knowledge pays scant attention to how individuals manage knowledge. This lack of attention diminishes the usefulness of the taxonomies presented by enabling authors who make broad claims about knowledge to side step the difficult issue of explaining how humans learn under their taxonomy of knowledge.

Alavi and Leidner [9], for example, state that knowledge exists only in the minds of people and so cannot be codified in computers. Blair makes a similar claim stating "only a person can be knowledgeable, that is, only a person can have and exercise knowledge" [11]. The question left unanswered by these claims is "How do people come to know?" While it is certainly reasonable to argue that people can know things a way that computers can't, the only way to defend such a claim is to offer a model of human knowledge and explain why computers cannot mimic this same model. The philosopher Daniel Dennett commented [25]:

> It is rather as if philosophers were to proclaim themselves expert explainers of the methods of stage magicians, and then, when we ask how the magician does the sawing-the-lady-in-half trick, they explain that it is really quite obvious: the magician doesn't really saw in her in half; he just makes it appear that he does. "But how does he do *that*?" we ask. "Not our department," say the philosophers.

The stance that "people know, computers can't" is not limited to Alavi and Leidner, Spiegler and Blair, as it is common in the literature. Indeed, "people know, computers can't" is simply a slice of the broader claim that people are intelligent, but computers cannot be. In next section, I will discuss work on these very questions of acquiring and representing knowledge.

5. Man with machine: a division of knowledge

Many authors in the KM literature claim that having knowledge is a uniquely human quality, but do not justify the claim either by novel argument or by reference to others who can support the claim. Therefore, it is not possible to directly address the evidence underlying the claims. In his analysis of KM, however, Blair draws on work by Hubert and Stuart Dreyfus, who make similar claims about the limitations of "artificial reasoning". I will first address the objections raised by Dreyfus (and Dreyfus). Most authors (including Dreyfus and Dreyfus) imply that it is the special knowledge of experts that enables them to solve problems that we cannot codify in computers. As there is agreement that one can show knowledge of such a skill by performing it, I will then present examples of computers engaged in this process of knowing.

5.1. Can rules explain everything?

As an AI researcher, I am in complete agreement with Dreyfus's claim that expert knowledge cannot be captured in a static rule set. In particular, it is generally impossible to specify all the needed conditions — a problem well known in AI for decades as the *qualification problem* [26]. Of course, people cannot specify all the needed preconditions for an action either, and yet still function successfully. There are two resolutions to the problem: constrain the "world" so tightly as to enable an exhaustive listing of necessary conditions or adopt a different architecture that does not require an exhaustive listing of logical rules and preconditions. Researchers in AI often adopt the second approach by using learning techniques and probabilistic reasoning systems. Thus, Dreyfus's primary criticism (admittedly now somewhat old, having most recently published in 1992, but a view still held by Dreyfus) is an objection to an agent architecture, not the enterprise as a whole.

As quoted in Russell and Norvig's text on AI [15], Dreyfus claims "A mere chess master might need to figure out what to do, but a grandmaster just sees the board as demanding a certain a move…the right response just pops into his or her head." Again, the problem is that Dreyfus does not explain how this "pop" occurs, reminding of us Dennett's observations on philosophers and stage magicians. While it is certainly plausible that the grandmaster in question might have trouble explicating his reasoning, this does not mean that the grandmaster is not reasoning or that a computer could not engage in the same reasoning. Of course, if you are building a chess-playing program, you cannot select a grandmaster at random and ask, "How do you play chess?" and simply transcribe the answer into the computer to create a strong chess player.

In their 1986 work, Dreyfus and Dreyfus [22] offer a theory of skill acquisition from simple rule following through automatic expert skill. Moreover, they offer a model of the mechanism at work in humans to achieve this progression: a large neural network organized around cases. They then present an analysis of problems in implementing the proposed mechanism in a computer. While I will not discuss their objections in detail here, their objections include lack of an adequate learning mechanism and the inability to direct attention to gathering needed information. Both kinds of problems have been by addressed — the learning mechanisms by researchers in the field of machine learning and problems of attention in the field of active vision.

As with Polanyi, I do not wish to suggest that the objections raised by Hubert and Stuart Dreyfus are without merit. The problems they highlight are the very same problems that AI researchers have struggled with for years. The healthy debate between AI researchers and critics spurs new research. However, the objections we must consider the objections carefully. For those who claim that people have know-how that computers will never master, the most convincing counter argument should be the demonstration of that know-how by computers, which I discuss next.

5.2. What computers *can* do

Perhaps the best-known example of computers knowing something is that computers play chess [27]. Using Blair's adoption of Wittgenstein's linguistic analysis, we find it acceptable to say, "She knows how to play chess" and there's general agreement that playing chess well requires special skill. Computers perform many other complex tasks that require such skills. Russell and Norvig [15] offer an example from David Heckerman's 1991 work on probabilistic reasoning [28]:

> [A] leading expert on lymph-node pathology scoffs at a program's diagnosis of an especially difficult case. The creators of the program suggest that he ask the computer for an explanation of the diagnosis. The machine points out the major factors influencing its decision and explains the subtle interaction of several of the symptoms in this case. Eventually, the expert agrees with the program.

In the commercial realm, Russell and Norvig [15] offer XCON (known as R1 [29] in its research incarnation) as an early example of a commercially successful rule-based system that "contained several thousand rules for designing configurations of computer components for customers of the Digital Equipment Corporation." Computers can also schedule transportation, arrange logistics, drive cars and perform surgery. Many of these tasks are ones that experts perform without conscious reflection — the very same tasks that some philosophers and KM researchers claim require knowledge that cannot be stored in a computer.

But what if computers perform these skills differently than people do? What if, inside, computers adopt a different strategy and use different tools than a human expert would? The answer is that computers do use a different approach; today's computers are digital electronic devices and function differently than the brains of people. Even more abstractly, a chess-playing computer might "see" the board differently than a human grandmaster. The simplest reply is "So what?" Many in the KM literature agree that ultimately we wish to use knowledge to create value. If computers perform the same tasks as well (or better) and create the same (or greater) value, then why should the difference matter?

In some cases, simple computer techniques outperform human experts. Meehl's 1955 study on clinical (i.e., human) versus statistical analysis [30] reports on an 1928 study of Illinois parolees in which the researcher compared a simple statistical approach (combining 21 factors in an unweighted fashion) with the clinical judgment of three prison psychiatrists. The psychiatrists' expert judgment was slightly more accurate in the successful cases (parolees who were not recidivists) but much less accurate in the unsuccessful cases. Moreover, the psychiatrists offered no opinion in many of the cases, while the statistical approach made predictions in every case. Grove and Meehl found similar results in the more recent 1996 study [31], stating "Empirical comparisons of the accuracy of the [clinical and statistical] methods (136 studies over a wide range of predictands) show that the mechanical method is almost invariably equal to or superior to the clinical method." While many researchers would argue that a linear regression model is too shallow to be taken as knowledge, in many cases the models yield more accurate judgments than those of seasoned experts. If we are to say that such experts use knowledge in making their predictions, why do we not consider a model with more accurate predictions of the same phenomena to be knowledge?

While people may do a poor job where computers succeed, it is nevertheless important that computers can present their findings in a way understandable to humans. (This is the same point made more generally about knowledge sharing between people by Alavi and Leidner [9].) Because people are often more skeptical of a computer than another expert, it is doubly important that computers offer clear insight into why they make the recommendations they do.

5.3. What are the implications for managing knowledge?

Because codification is a prerequisite to computer programming, the existence of computer systems that perform as well as (or better than) experts disproves two common assumptions in the KM literature: the inexpressibility of skill knowledge (know-how) and the limitation that only people can possess knowledge. The denial of the first assumption suggests that we should broaden KM efforts to include codification of experts' knowledge. I am not suggesting that we ask experts to simply "write down" everything they know. While many would have previously equated computers with the written word and lumped AI systems together with phone books and both apart from people, maybe it would be more accurate to group people with computers, apart from books and papers. The flexibility of computers means that codification in computer form is qualitatively different from simply authoring a document. I also do not wish to suggest that codifying expert knowledge in a computer system is a simple task. It is not. However, we should not dismiss it as impossible.

From the examples given above, it is clear that computers can possess know-how; computers offer a potential store of knowledge. The continuous availability of computers and ease of replicating systems from one computer to the next offers an opportunity to augment the human network of expertise within every organization

6. Conclusions: Adding a new theme

If we are to study the effective management of knowledge in organizations, we must understand something of the nature of nature of knowledge. In analyzing current models of knowledge in the KM literature, I have argued that common perspectives of knowledge have little justification and are in some case based on questionable philosophical underpinnings. The tyranny of tacit knowledge is that many researchers classify expert knowledge as tacit knowledge, conclude that it cannot be expressed, and so cannot be captured by computers. Thus, they conclude, we should dismiss computers as knowing entities and instead focus on linking together the people in organizations, perhaps by keeping electronic phone books on computer for easy access. But computers can know. Thus, knowledge management researchers should consider the question of how to divide and share knowledge between people and computers.

I am not arguing that computers should displace people or even that KM researchers should discontinue a current line of inquiry. I ask for an expansion of the research agenda. We should study techniques for capturing expertise, the complementary roles that people and computers can play and use computers to help better our understanding of how people think.

7. Acknowledgments

I would like to David W. Aha of the Naval Research Lab for sponsoring me during the summer, enabling me to write this paper. I would also like to thank Jim Sheffield for his extensive comments on my initial proposal and several anonymous reviewers. Lastly, I'd like to thank my five-year old son, who's forever asking, "How do you know that?"

8. References

[1] Plato, "Theaetetus," Project Gutenberg. Last accessed: June 9, 2004. Available at: <http://www.gutenberg.net/etext99/thtus10.txt>. 1999.
[2] G. Ryle, *The Concept of Mind*. Chicago, IL: University of Chicago Press, 2000.
[3] M. Rozemond, *Descartes's Dualism*. Cambridge, MA: Harvard University Press, 1998.
[4] M. Polanyi, "The Tacit Dimension," in *Knowledge in Organizations*, L. Prusak, Ed. Boston, MA: Butterworth-Heinemann, 1997, pp. 135-146.
[5] I. Nonaka, "The Knowledge-Creating Company," *Harvard Business Review*, 1991.
[6] I. Nonaka and H. Takeuchi, *The Knowledge Creating Company*. New York, NY: Oxford University Press, 1995.
[7] D. Barbiero, "Tacit knowledge," in *Dictionary of Philosophy of Mind*, C. Eliasmith, Ed., May 11, 2004 ed. St. Louis, MO: Washington University in St. Louis, 2004.
[8] F. D. Davis, "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology," *MIS Quarterly*, vol. 13, no. 3, pp. 319-340, 1989.
[9] M. Alavi and D. E. Leidner, "Knowledge Management Systems: Issues, Challenges, and Benefits," *Communications of the Association for Information Systems*, vol. 1, no. Article 7, 1999.
[10] I. Spiegler, "Knowledge Management: A New Idea Or a Recycled Concept?," *Communications of the Association for Information Systems*, vol. 3, no. Article 14, 2000.
[11] D. C. Blair, "Knowledge management: Hype, hope, or help?," *Journal of the American Society for Information Science and Technology*, vol. 53, no. 12, pp. 1019-1028, 2002.
[12] B. Vandenbosch and M. J. Ginzberg, "Lotus Notes and Collaboration: Plus ça change," *Journal of Management Information Systems*, vol. 13, no. 3, pp. 65-82, 1996.
[13] H. L. Dreyfus, *What Computers Still Can't Do*. Cambridge, MA: MIT Press, 1997.
[14] L. Wittgenstein, *The Blue and Brown Books*. New York: Oxford University Press, 1960.
[15] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 2nd ed. Saddle Hill, NJ: Prentice Hall, 2002.
[16] A. Turing, "Computing machinery and intelligence," in *Computers and Thought*, E. A. Feigenbaum and J. Feldman, Eds. Menlo Park, CA: AAAI Press, 1951, pp. 11-37.
[17] H. A. Simon and A. Newell, "GPS, A program that simulates human thought," in *Computers and Thought*, E. A. Feigenbaum and J. Feldman, Eds. Menlo Park, CA: AAAI Press, 1963, pp. 279-295.
[18] E. F. Vail, III, "Knowledge Mapping: Getting Started With Knowledge Management," *Information Systems Management*, vol. 16, no. 4, pp. 16-23, 1999.
[19] D. A. Schön, *The Reflective Practitioner: How Professionals Think in Action*, Paperback ed: Arena, 1995.
[20] P. Drucker, "The coming of the new organization," in *Harvard Business Review on Knowledge Management*. Cambridge, MA: Harvard Business School Press, 1998.
[21] B. Magee, *Men of Ideas*. New York, NY: The Viking Press, 1997.
[22] H. Dreyfus and S. Dreyfus, *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. New York, NY: The Free Press, Macmillan, Inc., 1986.
[23] A. J. Ayer, *The Problem of Knowledge*. Harmondsworth, Middlesex, U.K.: Penguin Books, 1956.
[24] R. Parry, "Epistêmê and Technê," Stanford University. Last accessed: June 15, 2004. Available at: <http://plato.stanford.edu/archives/sum2003/entries/episteme-techne/>. 2003.
[25] D. C. Dennett, "Cognitive wheels: the frame problem of AI," in *Minds, Machines, and Evolution: Philosophical Studies*, C. Hookway, Ed. Cambridge, U.K.: Cambridge University Press, 1984, pp. 129-151.
[26] P. Hayes, "The Frame Problem and Related Problems in AI," in *Artificial and Human Thinking*, A. Elithorn and D. Jones, Eds.: Jossey-Bass, Inc. and Elsevier Scientific Publishing Company, 1973.
[27] P. Hoffman, "Who's best at chess? For now, it's neither man nor machine," *The New York Times*, February 8, 2003, Late Edition - Final, Section B1, Page 2.
[28] D. Heckerman, *Probabilistic Similarity Networks*. Cambridge, MA: MIT Press, 1991.
[29] J. McDermott, "R1: A rule-based configurer of computer systems," *Artificial Intelligence*, vol. 19, no. 1, pp. 39-88, 1982.
[30] P. Meehl, *Clinical vs. Statistical Prediction*. Minneapolis, MN: University of Minnesota Press, 1955.
[31] W. M. Grove and P. E. Meehl, "Comparative efficiency of informal (subjective, impressionistic) and formal (mechanical, algorithmic) prediction procedures: The clinical statistical controversy," *Psychology, Public Policy, and Law*, vol. 2, no. 2, pp. 293-323, 1996.