

A Formal Knowledge Management Ontology: Conduct, Activities, Resources, and Influences

C.W. Holsapple

School of Management, Carol M. Gattton College of Business and Economics, University of Kentucky, Lexington, KY 40506-0034. E-mail: cwhols@uky.edu

K.D. Joshi

School of Accounting, Information Systems and Business Law, College of Business and Economics, P.O. Box 644750, Washington State University, Pullman, WA 991164-4750. E-mail: joshi@wsu.edu

This article describes a collaboratively engineered general-purpose knowledge management (KM) ontology that can be used by practitioners, researchers, and educators. The ontology is formally characterized in terms of nearly one hundred definitions and axioms that evolved from a Delphi-like process involving a diverse panel of over 30 KM practitioners and researchers. The ontology identifies and relates knowledge manipulation activities that an entity (e.g., an organization) can perform to operate on knowledge resources. It introduces a taxonomy for these resources, which indicates classes of knowledge that may be stored, embedded, and/or represented in an entity. It recognizes factors that influence the conduct of KM both within and across KM episodes. The Delphi panelists judge the ontology favorably overall: its ability to unify KM concepts, its comprehensiveness, and utility. Moreover, various implications of the ontology for the KM field are examined as indicators of its utility for practitioners, educators, and researchers.

Introduction

Ontologies provide a simplified and explicit specification of a phenomenon that we desire to represent (Gruber, 1995). Ontologies are useful because they explicate components that define a phenomenon and, thus, can help in systematically understanding or modeling that phenomenon. This article advances a general-purpose knowledge management (KM) ontology that can be used by practitioners, researchers, and educators. The ontology is characterized in terms of formal definitions and axioms that have evolved from a

collaborative ontology design process. The ontology posited here may be extended, refined, modified, or even replaced, but in its current form it provides a foundation for systematic KM research, study, and practice. As such, it also provides a basis for designing and analyzing technological approaches to KM.

Over the past decade, the KM field has received considerable attention from researchers and practitioners. Despite this attention and consequent progress, KM researchers have not provided a well-integrated framework to the community that would help unify this discipline. This sentiment is appropriately summarized by Spender (2003) when he states, "But as we look at the KM literature it is immediately clear that it is neither homogeneous nor well integrated. There is no single set of terms or even theoretical constructs. . . ." In order to facilitate maturation and progress of KM as a credible research discipline, the field needs a formal ontology that not only offers a comprehensive understanding of KM phenomena, but also operates as an organizer for past research and a generative mechanism for future research directions.

As a step in this direction, we introduce a formal characterization of a KM ontology collaboratively developed with an international panel of KM practitioners and researchers. Prior papers have informally detailed various portions of this ontology and described panelists' piecemeal evaluations of them (Holsapple & Joshi, 2000, 2001, 2002c). These form a basis for the unified, formal specification of the entire ontology advanced here. We begin with a brief description of the ontology design process and report on the panelists' evaluation of the overall resultant ontology in terms of satisfaction, unification, completeness, and utility. The formal KM ontology is then presented as a system of more than 90 definitions and axioms organized into four components. Implications of the ontology for research, practice, and education are subsequently discussed to illus-

Received June 2, 2003; revised November 14, 2003; accepted November 18, 2003

© 2004 Wiley Periodicals, Inc. • Published online 25 February 2004 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/asi.20007

trate unfolding contributions of this ontology to the KM field.

Method and Evaluation

Here, we briefly outline the ontological engineering method used for this research. A more in-depth treatment of this method appears in Holsapple & Joshi (2002a). At the heart of this four-phase process, the KM ontology evolved through collaboration with a diverse panel of KM researchers and practitioners. We also present the panelists' evaluations of the overall ontology that resulted from this method.

The Ontology Design Process

The ontology design process is partitioned into four parts: the preparatory, anchoring, collaborative, and application phases.

The preparatory phase sets boundary conditions for the ontology, establishes a set of standards for guiding the ontology development, and defines design criteria for ontology evaluation. Three boundary conditions were used: we focused on KM in business settings (although the results may well be applicable to nonbusiness situations), we restricted the intent to describing KM phenomena (rather than prescribing or speculating on KM practices), and we sought to capture KM concepts at two or more levels of detail (while realizing that the ontology may be amenable to additional levels of detail). A collection of KM frameworks, KM case studies, KM surveys, and concepts evident in KM articles formed a set of standards used in the next two phases. Because the aim was to produce a unifying, relatively comprehensive KM ontology, we sought to ensure that all were accommodated in the ontology as it was developed. That is, they served as guidelines for what to be sure to include in the ontology as it evolved. The criteria chosen to evaluate the overall ontology were degree of satisfaction, degree of helpfulness, degree of comprehensiveness, and degree of unification.

In the second phase, we developed an anchor ontology by consolidating, synthesizing, organizing, and integrating concepts found in the set of standards. The anchor ontology was developed through several iterations. At the end of each iteration, the ontology was assessed relative to its coverage of the standards, and a decision about the need for another iteration was made. The final version of the anchor ontology was carried forward into the collaborative phase.

During the third phase, we collaborated with a panel of 31 KM practitioners and researchers to further develop the ontology. A list of potential collaborators was compiled, consisting of contributors to the KM literature, presenters at KM conferences, and faculty who designate KM as a major research area. Because their perspectives on KM can differ, care was taken to include both researchers and practitioners in the list. The list was then reduced to include only those persons for whom a mailing address could be readily deter-

mined, yielding a total of 122 candidates (70 academicians and 52 practitioners) for participation on the KM panel. This number seems to be reasonable for the purpose of capturing a variety of viewpoints covering the current state of KM thought (by comparison, the Bacon and Fitzgerald [1996] study involving a similar methodology, but different a subject, targeted 113 candidates). The result was a diverse set of candidates, each having an active interest and track record in KM practice and/or research.

All 122 candidates were invited to participate and 31 (25.4%) chose to do so. Demographics gathered from the panelists confirmed substantial diversity in type of industry, geographic region, and perspective. These participants had demonstrated experience in the area of KM, were evenly divided between practitioners and researchers, approached KM from diverse vantage points (e.g., sociology, human resources, philosophy, strategy, communications, technologies), and had work activities spanning four continents. A list of these contributors to the KM ontology appears in the Acknowledgements.

The anchor ontology and a questionnaire eliciting views on it were pilot tested for clarity, revised accordingly, and then sent to each panelist. The questionnaire was designed for a structured elicitation of critiques of the framework in terms of the evaluation criteria (comprehensiveness, correctness, clarity, utility, and conciseness) developed in the preparatory phase. Three sections of the instrument correspond to the three components of the initial framework (resources component, knowledge manipulation activity [KMA] component, and KM influences component). In each of these sections, participants' suggestions, concerns, and comments on the component's completeness, accuracy, clarity, and conciseness are captured using seven-point Likert-scale items, as well as via written critiques. Respondents are probed to elaborate on why they were not satisfied with a certain aspect of the framework. They are asked to suggest ways to further improve the framework. A fourth section gathered assessments of the framework's general utility, overall comprehensiveness, overall unification, and limitations.

We organized responses to the questionnaire and analyzed them to determine how to modify/refine the anchor ontology in order to address participants' critiques and suggestions, while still accommodating the set of standards. This yielded a revised ontology that, along with a summary of first-round responses and another questionnaire, was sent to panelists for further critique and comment. Delphi rounds continued in this way until the ontology had evolved to a point where participants expressed no major reservations. At this point, panelists gave an evaluation of the overall ontology, as described in the next section.

In the last phase, which is still unfolding, the utility of the developed ontology has been demonstrated in various ways, such as showing that it provides a unifying view of KM phenomena, using it to characterize KM technologies, using it to structure KM case studies, and adapting it to develop a KM model for competitive analysis. Such uses of




FIG. 1. Degree of success in providing a unified view.

the ontology are discussed near the end of this article, following the formal characterization of the full KM ontology.

Overall Evaluation of the Ontology

Participants were asked to give a final evaluation of the overall ontology on four criteria: their degree of satisfaction with it, the extent of unification it provided, its level of comprehensiveness, and the degree of usefulness offered to researchers and practitioners. Frequency distributions of participants' responses for these criteria are presented in Figures 1 through 4.

This descriptive technique for data analysis readily conveys a sense of the panel's positive view of the ontology. Although hypotheses could have been constructed and tested (e.g., the mean response for success in providing a unified view is significantly greater than the midpoint of moderately successful), the choice of a basis for comparison (e.g., midpoint) would be somewhat arbitrary. Moreover, such testing would not be quite as informative as seeing the underlying distributions, leaving the choice of a comparative basis to the observer.

Overall, 94% of the respondents were at least moderately satisfied with the resultant ontology. The mode for this item was moderate satisfaction, with 38% being more than moderately satisfied. With respect to furnishing a unified and

comprehensive view, at least 81% of the respondents gauge the ontology as at least moderately successful. The modes for unification and comprehensiveness are moderately successful and successful, respectively. For unification, 44% perceive more than moderate success. For comprehensiveness, 56% perceive more than moderate success.

Perhaps most important is the finding that a majority of respondents perceive that the ontology has utility for researchers and practitioners. Nearly 60% of the participants evaluate this ontology as being in the range "helpful" to "extremely helpful" for researchers. All but about 5% of the panelists regard it as at least "moderately helpful" for researchers. At about 30%, the mode of ontology utility for practitioners is "very helpful" and more than 70% of panelists indicate that the ontology is at least "moderately helpful" for practitioners. Several panelists indicated that practitioners would especially like to have a prescriptive KM framework, but this is outside the descriptive boundary established for ontology development in the preparatory phase. Nevertheless, the ontology offers a language for considering, discussing, and specifying KM prescriptions.

Formal Specification of the Knowledge Management Ontology

The formal ontology is presented in terms of *primitives*, *definitions*, and *axioms*. Primitives are terms that are con-




FIG. 2. Overall framework satisfaction.




FIG. 3. Degree of success in providing a comprehensive view.

sidered commonly understood and thus do not need to be characterized. A few examples of primitives used in this article include entity, organization, behaviors, environment, and technology. For the sake of simplicity, primitives are not enumerated; rather, all terms not defined in the ontology are assumed to be primitive.

Definitions form the ontology's backbone, defining its conceptual structure. Each definition is assigned a name using the following convention: The name starts with the letter "D," which is followed by a component acronym and a number. Here, an axiom is a rule or principle that is accepted as true as a result of the ontology development process. An axiom is specified using defined terms and primitives. Axioms are assigned names with the following convention: The name starts with a letter "A," which is followed by a component acronym and a number. The four component acronyms used in naming definitions and axioms correspond to the four components of the posited ontology: knowledge management conduct (KMC), knowledge manipulation activities (KMAs), knowledge resources (KRs), and knowledge management influences (KMIs).

Knowledge Management Conduct Component

We begin with the ontology's definitions and axioms that describe, at a broad level, the nature of knowledge manage-

ment phenomena. This component provides a foundation for appreciating the ontology's remaining three components.

DKMC1: Knowledge Management—An entity's systematic and deliberate efforts to expand, cultivate, and apply available knowledge in ways that add value to the entity, in the sense of positive results in accomplishing its objectives or fulfilling its purpose.

Knowledge management can be approached or studied on any of several levels, which vary in scope according to the nature of the entity. Much of the focus on KM has been at an organizational level, investigating knowledge management within the boundaries of a firm, enterprise, agency, or other organizational entity. However, it can be examined within a narrower scope, where the entity performing KM is an individual person rather than an organization. More broadly, the scope can be enlarged to transorganizational or national levels.

DKMC2: Personal Knowledge Management—Knowledge management conducted by an individual.

DKMC3: Organizational Knowledge Management—Knowledge management conducted by an organization.

DKMC4: Transorganizational Knowledge Manage-




FIG. 4. Degree of helpfulness.




FIG. 5. Knowledge for processor P in context C comprised of usable representations.

ment—Knowledge management conducted by multiple collaborating organizations.

DKMC5: National Knowledge Management—Knowledge management conducted by a nation.

Other scopes that can be used in considering the conduct of KM include interpersonal (involving multiple individuals not belonging to a specific organization), regional (within or across nations), and global. Construction of the ontology via the Delphi process emphasized the organizational scope, but the resultant constructs outlined below may be largely applicable to other kinds of entities as well.

Regardless of the boundary conditions selected for a particular KM study or initiative, the entity involved will have some knowledge and some skills for processing that knowledge. Many definitions of knowledge can be found. These involve considerable diversity and sometimes endeavor to define knowledge in terms of one or another of its attributes. In the Delphi process, no single definition of knowledge was developed. Rather, the intent was to develop a characterization of KM that accommodates different perspectives on the nature of knowledge. In this spirit, Newell's (1982) relatively inclusive view on knowledge is adopted here:

DKMC6: Knowledge—That which is conveyed in usable representations.

There are two key aspects to this simple, yet powerful, definition. First, the representations may be of many kinds including symbolic, audio, visual, behavioral, mental, and digital patterns (Holsapple, 1995, 2003a). These patterns range from static (e.g., object representations) to dynamic (e.g., process representations). Second, a representation

must be usable to some processor in the sense of having sufficient validity (e.g., accuracy, certainty, consistency) and utility (e.g., clarity, meaning, relevance, importance) for sense making (Holsapple & Whinston, 1996). For a given representation, there are, of course, degrees of usability that can vary from one processor to another, from one situation to another, and from one time to another.

Figure 5 illustrates the ontology's conception of knowledge, and the following axioms flesh out the formal characterization:

AKMC1: Knowledge representations include mental, behavioral, symbolic, digital, visual, audio, and other sensory patterns that may occur in various object and process formats.

AKMC2: The usability of a knowledge representation is a function of the validity and utility of knowledge it conveys for a particular processor in a particular context.

AKMC3: Knowledge has a variety of attributes including mode (tacit vs. explicit), type (descriptive vs. procedural vs. reasoning), orientation (domain vs. relational vs. self), applicability (local vs. global), accessibility (public vs. private), immediacy (latent vs. currently actionable), perishability (shelf life), and so forth.

Some of the attribute dimensions of knowledge are depicted in Figure 6. More complete and detailed listings of attribute dimensions for characterizing knowledge have been discussed, but are beyond the scope of this article (Holsapple & Joshi, 2001; Holsapple, 2003a). In a related vein, various knowledge taxonomies have evolved over the years (Nonaka, 1994; Alavi & Leidner, 2001; Marshall & Brady, 2001; Randall, Hughes, O'Brien, Rouncefield, & Tolmie, 2001; Sutton, 2001), but are not incorporated into




FIG. 6. A web of knowledge attributes.

the ontology. In the interest of being generic, the ontology is neutral on these differential views. A user of the ontology is free to adopt whatever specialized characterization of knowledge or attribute focus he/she desires to fit his or her KM context.

The classes of organizational resources traditionally studied in business school curricula are financial, human, and material assets. However, knowledge has been recognized as another important class of organizational resources (Drucker, 1993) and is increasingly being regarded as a basis for organizational competitiveness (Holsapple & Singh, 2001). KM is also beginning to be included in business curricula (Ruth, Shaw, & Frizell, 2003).

DKMC7: Resource—A source of value, revenue, wealth, or rent.

DKMC8: Knowledge Resource—Knowledge that an entity has available to manipulate in ways that yield value.

AKMC4: Four major classes of organizational resources are financial resources, human resources, material resources (including land, facilities, machinery, computer systems, inventories), and knowledge resources (KRs).

AKMC5: Over time, an organization's mix of resources is subject to change via acquisition, production, and elimination.

AKMC6: An organization's KRs can be manipulated by human resources and/or material resources (i.e., computer systems).

An entity's resources can be exchanged to acquire new resources from other entities. For instance, a firm may exchange some of its financial resources in order to acquire additional KRs (or vice versa). An entity's KR is comprised of knowledge that belongs to that entity, in the sense that it is available for manipulation (i.e., processing) by the entity's processors. The Delphi process yielded a taxonomy of KRs as one component of the knowledge management ontology. This is formally described later.

The knowledge processors can be human and/or technological. That is, there are usable representations for both

human- and computer-based systems, although the characteristics of their processing abilities/capabilities vary and the natures of the representations that are usable to them also vary.

DKMC9: Processor—A possessor of certain skills that allow it to implement some range of actions.

AKMC7: Some processors are more effective than others in implementing a particular type of action in a given situation.

AKMC8: A processor may be more effective in implementing one type of action than it is in implementing other types of actions.

AKMC9: Effectiveness of a processor's action can be impacted by the context within which that action is implemented.

DKMC10: Knowledge Processor—A part of (i.e., a participant in) an entity that possesses skills allowing it to implement some range of knowledge manipulations activities with varying degrees of effectiveness.

AKMC10: Knowledge processors can be human participants or computer-based parts in an entity.

AKMC11: A knowledge processor may be individual or collective (i.e., distributed).

Regardless of the processors involved and their respective skills (more or less developed), an entity's knowledge management work involves the manipulation of its KRs by those processors. This often results in a change in the state of the entity's own knowledge and/or knowledge flows that changes the state of KRs belonging to other entities with which it interacts. See Figure 7.

DKMC11: Knowledge Manipulation—The processing of usable representations.

DKMC12: Knowledge Manipulation Activity—A kind of knowledge processing that can be recognized and characterized independent of the nature of the knowledge representations being processed.

Alternative views have been advanced on what the major KMAs are. These views tend to use differing terminology, overlap in various ways, and deal with manipulation at different levels (Holsapple & Joshi, 2002b,c). As part of the




FIG. 7. A knowledge processor using its skills to yield knowledge flows.




FIG. 8. Architecture of a KM episode.

ontology, the Delphi process aimed to develop a relatively comprehensive, unifying, elemental characterization of the major KMAs that occur in an entity's conduct of knowledge management. An entity may perform these activities in various ways, at varying times, in varying sequences or configurations, via varying processors, and operating on varying KRs. The ontology's KMAs are formally described later.

An entity's conduct of knowledge management is shaped by, constrained by, and guided by various types of influences.

DKMC13: Influence—A factor that can affect resources, processors, and processes.

DKMC14: Knowledge Management Influence—A factor that determines how an entity's manipulation of knowledge unfolds in the course of knowledge management.

Specific knowledge management influences have been identified via the Delphi method and are formally described later.

Finally, an entity's knowledge management work can be seen as being organized into episodes (Holsapple et al., 1996). These episodes may be independent or interdependent, serial or parallel, of long or brief duration, simple or complex.

DKMC15: Episode—A distinctive process that can be recognized as separate from (but maybe related to) other processes.

DKMC16: Knowledge Management Episode—An entity's execution of some configuration of KMAs by some collection of processors, triggered by its intent to satisfy a knowledge need or opportunity, operating on available KRs, subject to knowledge management influences, and yielding learning and/or projections.

As illustrated in Figure 8, each knowledge management episode (KME) is triggered by the recognition of a knowledge need or opportunity; it culminates when that need/opportunity is satisfied or abandoned. Outcomes of a successful knowledge management episode are learning and/or projections. At a micro-level, KM influences affect what KRs are employed, how KMAs are configured, and what outcomes are yielded within a particular episode. At a macro-level, they affect the patterns of episodes that unfold in an entity's conduct of knowledge management.

It can be instructive to classify types of KM episodes as a basis for studying them (e.g., to investigate what methods or technologies work well for a particular episode class). There are various ways to classify episodes. One of these is given in the following axiom:

AKMC12: Knowledge management episodes can be categorized based on the nature of the intent, such as decision making, designing, researching, negotiating, problem solving, and brainstorming episodes.

The foregoing list of episode classes is suggestive rather than exhaustive, indicating that knowledge management episodes are knowledge-intensive processes that can vary in purpose and form.

Knowledge achieves direct returns along two dimensions of organizational performance: learning and projection. By affecting these, it can also indirectly affect other performance dimensions such as financial returns and, ultimately, competitiveness.

DKMC17: Learning—A process whereby KRs are modified; an outcome of a KME involving change in the state of an entity's knowledge.

Learning can be functional (a positive change in the state of an organization's KR) or dysfunctional (a negative change of state). These can be changes in amount, degree, quality, orientation, mode, structure, and so forth. A KME that results in learning may be regarded as a learning episode.

Expressions or manifestations of an entity's KME that are emitted into its environment are called projections. These can include not only knowledge, but financial, material, and human projections as well. In every case, knowledge management is either a prerequisite or co-requisite of projection.

DKMC18: Projection—A process whereby resources are emitted into the environment; an outcome of a KME involving an impact on the state of the entity's environment.

Examples of projections from a firm's KMEs include a commitment of funds to make a purchase (stemming from a decisional KME), a tangible product for a customer (stemming from some combination of design, production, and decisional KMEs), a personal service for a consumer (stemming from some combination of design, operations, decisional KMEs), and knowledge for a client (stemming from problem solving, research, or design KMEs). Fundamentally, a projection is packaged knowledge. Projections can be functional (i.e., positive outcomes for an entity or its environment) and/or dysfunctional (i.e., negative impacts for an entity or its environment). A KME that results in projection may be regarded as a projection episode.

AKMC13: Some knowledge management episodes are learning episodes, others are projection episodes, and still others may involve both learning and projection.

The ontology's episodic view is useful for studying, understanding, and shaping the conduct of knowledge management by an entity of interest.

DKMC19: Conduct—Behaviors that unfold during performance of action.

DKMC20: Conduct of Knowledge Management—An entity's ongoing execution of various knowledge management episodes, often configured in interrelated patterns and governed by knowledge management influences.

The ontology's characterization of knowledge management is more fully developed in the next three sections, beginning with KMAs and proceeding to the resource and influences components.

Knowledge Manipulation Activities Component

The Delphi process led to the ontology's identification of basic KMAs, their subactivities, and their interrelationships. We follow this sequence in presenting them formally as definitions and axioms. Recall that a KMA is an elemental type of knowledge processing that can be characterized

independently of the nature of knowledge representation being processed.

AKMA1: There are five types of KMAs that can occur in the conduct of knowledge management: knowledge *acquisition*, *selection*, *generation*, *assimilation*, and *emission*.

Each type of KMA constitutes a class comprised of many possible instances of that activity. For example, many instances of knowledge acquisition can occur in an entity's conduct of KM. These instances differ in terms of which processor(s) performs the acquisition, the kind of representation(s) used, the mechanism(s) employed in performing the acquisition, the subject matter of the acquired knowledge, and so forth. In the Delphi study, the activities of assimilation and emission were originally termed internalization and externalization, respectively. To avoid confusion with Nonaka's (1994) more specialized meanings of these latter terms, we do not use the original terminology here.

AKMA2: A specific instance of a KMA in an episode can be jointly performed by more than one of an entity's knowledge processors, or it can be performed by an individual processor.

AKMA3: A knowledge processor may have the skill/ability to perform more than one kind of KMA within a single knowledge management episode or across multiple episodes.

The knowledge resulting from a processor performing a KMA may be transferred for processing within other instances of KMAs. That is, there are flows of knowledge that emanate from KMAs. A knowledge flow may be initiated by a processor performing an instance of an activity, or it may be in response to a request made by a processor performing some other activity instance. So, beyond the main knowledge flows among processors performing KMAs, there are ancillary messages that can pass among activity instances. In addition to requests for a knowledge flow, these include feedback (e.g., about the suitability of a knowledge flow), clarification (e.g., about the meaning of a knowledge flow), evaluation (e.g., of the quality of a knowledge flow), and so forth. As illustrated in Figure 9, an instance of a KMA mediating knowledge flows and ancillary messages can involve any one of the five manipulation classes.

AKMA4: Instances of a KMA can result in flows of knowledge.

DKMA1: Knowledge Flow—The transfer of knowledge from one instance of a KMA to another instance, possibly involving a transformation of the knowledge representation.

DKMA2: Ancillary Message—A message that an instance of a KMA sends to another activity instance in order to issue a request or provide feedback, clarification, or evaluation.




FIG. 9. Knowledge flows and ancillary messages for an instance of a knowledge manipulation activity.

The Delphi process yielded the following definitions for the ontology's five classes of KMA:

DKMA3: Knowledge Acquisition—A KMA comprised of identifying knowledge in the entity's environment and making it available in a suitable representation to an appropriate activity.

Examples of knowledge acquisition include conducting consumer satisfaction surveys, purchasing a patented process or data bank, and hiring an employee.

DKMA4: Knowledge Selection—A KMA comprised of identifying knowledge within an organization's existing base of KRs and providing it in an appropriate representation to an activity that needs it.

Examples of this activity include selecting qualified employees to lend their knowledge to some task, selecting information from a database, and determining an appropriate policy or procedure from among those known to the entity. Knowledge selection differs from knowledge acquisition in that it focuses on knowledge already possessed by an entity rather than knowledge held by the environment. The difference is important because the nature, the skills, and the cost of processing may vary depending upon whether the resource is internal or external to the entity.

DKMA5: Knowledge Assimilation—A KMA that alters an entity's KR, resulting in learning.

Examples of this activity include creating/modifying an organization memory system, documenting best practices,

authoring/publishing a policy manual, and sharing knowledge among the entity's processors.

DKMA6: Knowledge Generation—A KMA whereby an entity derives or discovers knowledge in the context of existing knowledge.

Examples of this activity include deriving forecasts, discovering patterns, developing strategies, devising methods to improve customer satisfaction, and designing products and services. The knowledge generated may not necessarily be "new" to the entity; it may currently exist or may have previously existed in the entity. For instance, it may be more cost effective to generate new knowledge than to select it, the processor doing the generation may be unaware of its existence in the entity, or the knowledge was not properly assimilated when previously acquired or generated.

DKMA7: Knowledge Emission—A KMA that uses existing knowledge to produce projections for release into the environment.

Examples of this activity include manufacturing a product, developing an advertisement, producing a report, furnishing consultation advice, and providing customer service. Knowledge emission results in projections for external consumption, in contrast to assimilation, which results in learning involving KR retention.

DKMA8: Knowledge Use—The activity of applying existing knowledge to generate new knowledge and/or accomplish knowledge emission.

This term is used for convenience to refer to knowledge generation and/or emission.

As specified in the next five axioms, the ontology sees all of the KMAs as comprised of subactivities.

AKMA5: Knowledge acquisition is accomplished through a set of subactivities that includes *identification* of appropriate knowledge from the external sources, *capturing* the identified knowledge, *organizing* captured knowledge, and *transferring* the organized knowledge to an appropriate activity.

AKMA6: Knowledge selection involves a collection of subactivities that includes *identifying* appropriate knowledge within the entity's existing resources, *capturing* identified knowledge, *organizing* captured knowledge, and *transferring* organized knowledge to an appropriate activity.

AKMA7: Knowledge assimilation is realized through *assessing* and *valuing* knowledge to be assimilated, targeting KRs where knowledge would be assimilated, *structuring* knowledge into forms appropriate for the targets, and *transferring* the knowledge representations as targeted.

AKMA8: The generation of knowledge involves *monitoring* the entity's KRs and the external environment and obtaining required knowledge (via selection or acquisition); *evaluating* the obtained knowledge in terms of its utility and

TABLE 1. The ontology's knowledge manipulation activity component.

Knowledge manipulation activity	Subactivities within this activity	Knowledge flows into this activity from	Knowledge flows from this activity to
Acquisition	Identification, capturing, organizing, transferring	Entity's environment	Assimilation, generation, emission
Selection	Identification, capturing, organizing, transferring	Entity's knowledge resources	Acquisition, assimilation, generation, emission
Assimilation	Assessing/valuing, targeting, structuring, transferring	Acquisition, selection, generation	Entity's knowledge resources
Generation	Monitoring, evaluating, producing, transferring	Acquisition, selection	Assimilation, emission
Emission	Targeting, producing, transferring	Acquisition, selection, generation	Entity's environment

validity for the production of knowledge; *producing* knowledge by creating, synthesizing, analyzing, and constructing knowledge from a base of existing knowledge; and *transferring* the produced knowledge to an appropriate activity.

AKMA9: Emission is accomplished by *targeting* elements of the environment to determine what projections need to be produced, *producing* projections for the target by applying, embodying, controlling, and leveraging existing knowledge, and *transferring* the projections to targets which involves packaging and delivery.

Externalization is only partially a KMA because it can involve physical activities such as the act of producing a product through transformation of raw materials.

A processor engaged in an instance of a KMA communicates with other processors performing other activity instances through knowledge flows and ancillary messages as defined previously. The next five axioms characterize possible incoming and outgoing knowledge flows for each of the KMAs. These flows are summarized in Table 1, along with a summary of subactivities.

AKMA10: A knowledge acquisition activity receives knowledge flows from an entity's environment and delivers the acquired knowledge to an activity that immediately uses the knowledge and/or to one that assimilates it within the entity for subsequent use.

AKMA11: A knowledge selection activity receives knowledge flows from an entity's KRs and delivers the selected knowledge to the acquisition, use, and/or assimilation activities.

AKMA12: A knowledge assimilation activity receives knowledge flows from knowledge acquisition, selection, or generation activities and produces knowledge flows that are transferred/embedded into the entity's KRs.

AKMA13: A knowledge generation activity receives knowledge flows from knowledge selection or acquisition activities and delivers the generated knowledge to assimilation and/or emission activities.

AKMA14: A knowledge emission activity receives knowledge flows from knowledge selection, acquisition, and/or generation activities and delivers the packaged knowledge (i.e., projections) to targets in the environment.

Finally, there are axioms that relate the KMAs to other aspects of the ontology.

AKMA15: Assimilating knowledge is a culminating activity in entity learning.

AKMA16: Emitting knowledge is a culminating activity whereby an entity projects knowledge into its environment, adding value to that environment and possibly receiving resources in return.

Both learning and projection can add value to an entity. Learning does so in the sense of enhancing what an organization can potentially accomplish. Projection can add value to an entity and its environment. The value added can be in such forms as improved profits, image, customer loyalty, and visibility. Once emission occurs, its effects on the environment (e.g., in the forms of sales, profits, market impressions, and customer/supplier complaints/compliments) are candidates for knowledge acquisition.

The KM ontology is not prescriptive. It does not advocate any particular approach for coordinating KMAs. Rather, it recognizes that many arrangements are possible. The effectiveness of a particular configuration of KMAs relative to a particular set of participating knowledge processors is moderated by knowledge management influences.

AKMA18: The configuration of KMAs that occur in a knowledge management episode and the collection of knowledge processors that perform those activities are subject to variation.

The ontology's KMAs are applicable to any knowledge domain (e.g., marketing, manufacturing, consulting), knowledge mode (tacit or explicit), or knowledge type (e.g., descriptive, procedural, or reasoning knowledge). For instance, the knowledge selection activity may involve a selection of explicit knowledge by extracting records of how a problem was previously handled, or a selection of tacit knowledge by observing behaviors (e.g., seeking and distilling knowledge from an organization's culture).

AKMA19: A KMA can operate on different types of KRs.

Knowledge Resource Component

Recall from definitions DKMC6 and DKMC8 that a KR is conveyed by usable representations that are subject to

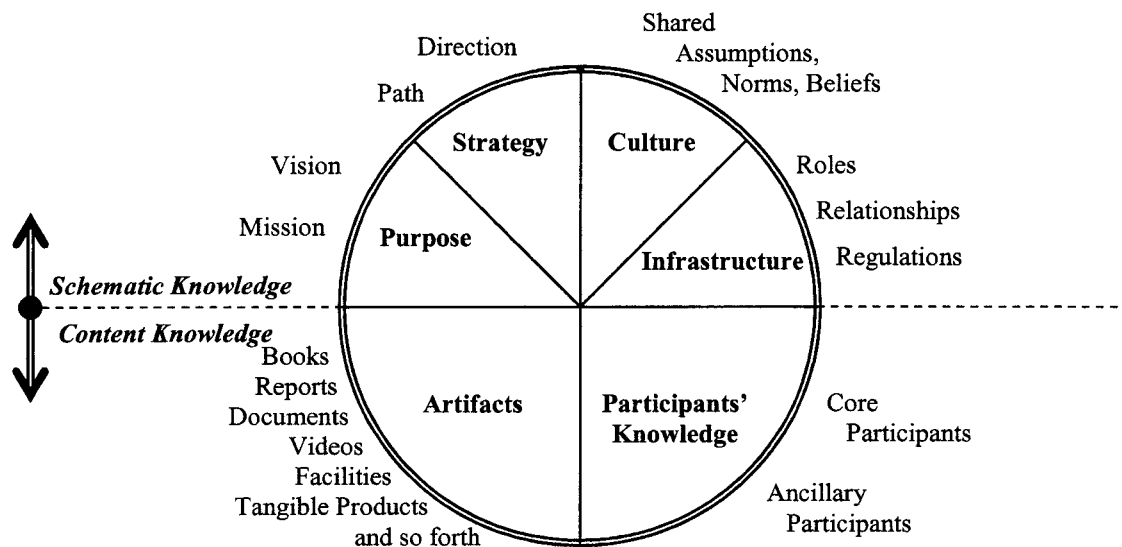


FIG. 10. Classes of organizational knowledge.

manipulation in ways that yield value. The ontology further develops this KM resource component via the following definitions and axioms. These are presented from the standpoint of entities that are organizations, but may well be applicable to other kinds of KM. Figure 10 furnishes an overview of the resource component.

AKR1: An organization has two classes of KRs: schematic and content KRs.

The basic premise underlying this categorization is that some KRs exist independently of an entity (e.g., organization), while the others depend on the entity for their existence.

DKR1: Schematic Knowledge Resource—a KR whose existence depends on the existence of the organization.

DKR2: Content Knowledge Resource—a KR that exists independently of an organization to which it belongs.

Schematic knowledge is represented or conveyed in the working of an organization. It manifests in the organization's behaviors. Schematic KRs establish an organization's ongoing identity. They are the basis for attracting, organizing, and deploying content resources.

AKR2: There are four kinds of schematic resources: culture, infrastructure, strategy, and purpose.

Each is a source of revenue or wealth for an organization. Each denotes a kind of organizational KR whose existence depends on the organization's existence. Each may change over time, but is invariably present as a KR.

An organization's values, principles, norms, traditions, unwritten rules, and informal procedures comprise its cultural KR.

DKR3: Culture—Basic assumptions and beliefs that are shared by members of an organization, that operate unconsciously, and that define in a basic taken-for-granted fashion an organization's view of itself and its environment (Schein, 1985).

An organization's infrastructure can be viewed as a formal counterpart to its cultural KR. It arranges an organization's participants in terms of the roles that have been defined for participants to fill, the relationships among those roles, and regulations that govern the use of roles and relationships (Holsapple & Luo, 1996).

DKR4: Infrastructure—The knowledge that defines an organization's roles, their interrelationships, and the regulations that govern the use of those roles and relationships.

An organization's strategy is comprised of plans for using an organization's infrastructure, culture, knowledge artifacts, and participants' knowledge (as well as other organizational resources). For instance, these can be plans for promoting a product or achieving effective resource allocation.

DKR6: Strategy—The knowledge that defines what to do in order to achieve organizational purpose in an effective manner.

An organization's purpose is directional knowledge with which its strategy, infrastructure, and culture need to be aligned.

DKR5: Purpose—The knowledge that defines an organization's reason for existence in terms of mission, vision, objectives, and goals.

The content KRs that exist at a given time qualify, condition, and color an organization's identity.

AKR3: The existence and use of content resources in the conduct of KM are both enabled and constrained by the schematic KRs.

AKR4: Content KRs are of two types: participants' knowledge and knowledge conveyed in/by artifacts.

A participant's knowledge or the knowledge represented in an artifact can have an existence apart from the organization that happens to host it at any given time. Each is not only subject to change over time, but is also subject to elimination.

DKR 6: Knowledge Artifact—An object that has no innate knowledge processing skills, but which is (or holds) a representation(s) of knowledge that may be usable to at least one knowledge processor in the organization.

Common examples of knowledge artifacts are video training tapes, books, memos, printed business plans, manuals, reports, patent documents, filing cabinet contents, facilities, layouts, and products (e.g., knowledge embedded in a manufactured vehicle).

DKR7: Participants' Knowledge—Knowledge possessed by a knowledge processor that participates in an organization.

Organizational participants include its employees, customers, suppliers, partners, consultants, and computer systems. The latter can function as substitutes for human participants, doing knowledge processing that may otherwise be performed by human participants. In so doing, they can overcome some of the cognitive, economic, and temporal limits confronting human knowledge processors. The Delphi panel was insistent that participants' knowledge includes the knowledge of not only employees, but also customers' knowledge and that of other ancillary participants in the organization.

DKR8: Core Participants—Possessors and manipulators of knowledge that belong to an organization's base of human and/or material resources.

See DKMI9 and DKMI10 below for definitions of human and material resources.

DKR9: Ancillary Participants—Possessors and manipulators of knowledge that do not belong to an organization's base of human and/or material resources, but whose knowledge is readily available for manipulation without having to be acquired.

Regardless of whether we are considering a core participant (e.g., employee) or an ancillary participant (e.g., a key customer), there are interesting questions such as: What

portion of the participant's knowledge can the organization regard as part of its own KR (vs. that portion of a participant's knowledge that is not made available to the organization)? What steps can be taken to maximize this portion, at least for the knowledge that would be helpful to the organization? How can this knowledge be preserved (if desired), even when the participant ceases participation?

The primary distinction between participants' knowledge and artifacts lies in the presence or absence of knowledge processing abilities. Participants have knowledge manipulation skills that allow them to process their own repositories of knowledge; artifacts have no such skills. A participant's knowledge is made available to an organization by means of that participant's knowledge manipulation skills. In contrast, an artifact is not accompanied by a processor and does not depend on a participant for its existence.

Existence of knowledge acquisition as a KMA implies the existence of knowledge in an organization's environment. The knowledge emission activity has a similar implication, in that projections can be viewed as embodiments of knowledge. The environment's KRs are a crucial source for replenishing and augmenting an organization's KRs.

DKR10: Environmental KRs—Knowledge that exists in an organization's environment that is potentially accessible/available for acquisition.

Perceptions of schematic knowledge can be captured and embedded in artifacts or participants' memories. For instance, we may represent culture, infrastructure, purpose, or strategy in an artifact (e.g., documentation), but its existence does not depend on the creation of an artifact.

AKR5: Schematic KRs exist independently of any one participant or artifact.

Various interplays exist among the different classes of KRs. For instance, strategy is distinct from purpose (i.e., alternative strategies are possible for a given purpose), yet strategy should conform to purpose; culture is distinct from infrastructure, yet culture can constrain infrastructure and vice versa; each schematic resource is distinct from content resources, yet a rendition of it may be represented in a participant's knowledge or as an artifact.

AKR6: The six types of KRs (participant's knowledge, knowledge artifacts, culture, infrastructure, purpose, and strategy) are both distinct and interrelated.

Knowledge Management Influences Component

Recall from definitions DKMC14 and DKMC16 that knowledge management influences are factors that determine how the manipulation of knowledge unfolds within and across knowledge management episodes. An entity engaged in a KM episode has discretion over some of these influences, while others function as constraints on the epi-

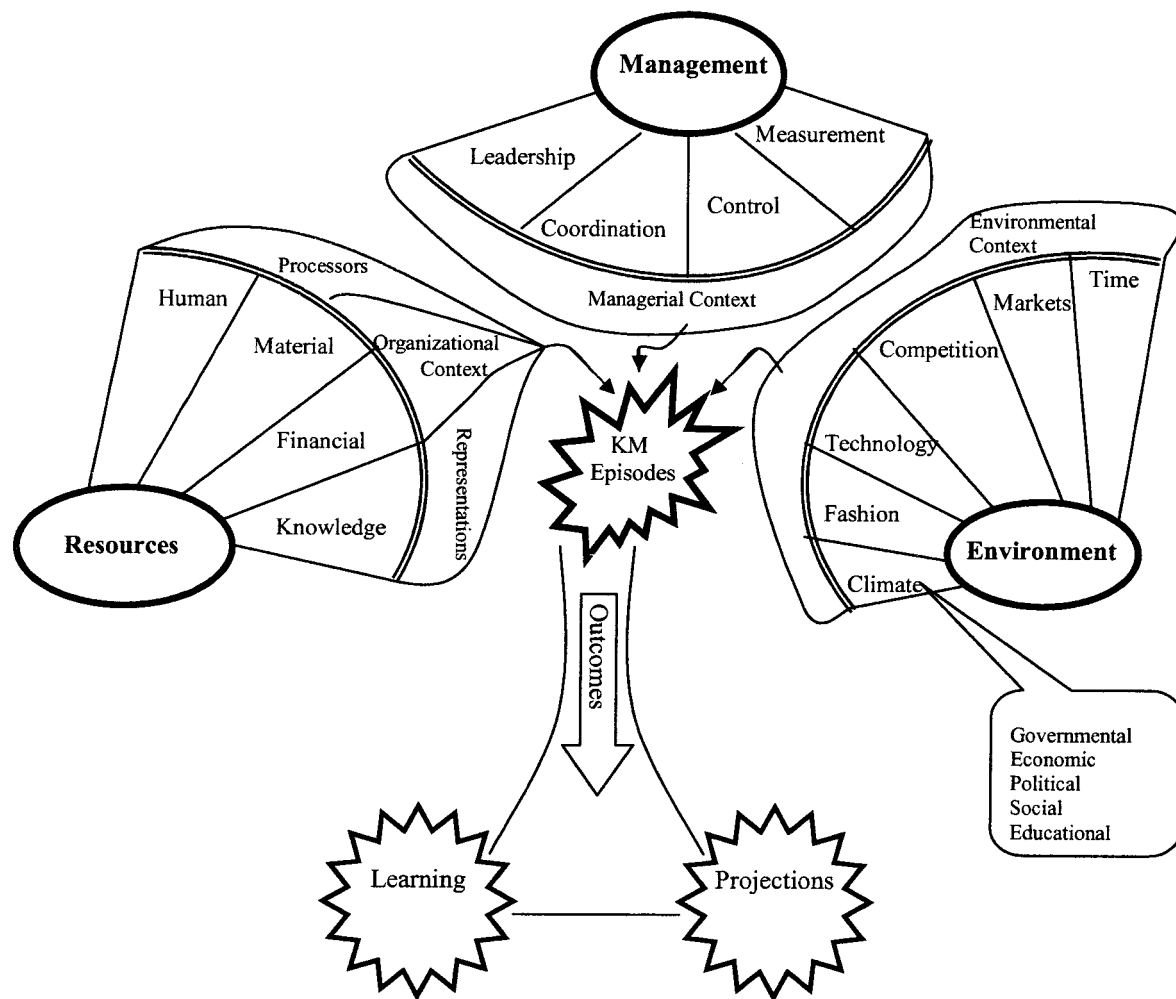


FIG. 11. Influences on the conduct of knowledge management.

sode and its outcomes. The ontological engineering process yielded a taxonomy of KM influences comprised of three major categories, as depicted in Figure 11.

AKMI1: There are three major classes of knowledge management influences: managerial influences, resources influences, and environmental influences.

DKMI1: Managerial Influences—Administrative efforts undertaken by an entity that affect its conduct of knowledge management.

DKMI2: Resource Influences—An entity's resources that are deployed to execute and affect its conduct of knowledge management.

DKMI3: Environmental Influences—Factors external to an entity (i.e., in its environment) that affect its conduct of knowledge management.

An entity may possess the best KRs and knowledge manipulation skills, but they are of relatively little use unless effectively harnessed in the conduct of KM. In the case of an organization, managerial influences emanate from those organizational participants responsible for directing its KM initiatives (such as a chief knowledge offi-

cer). The Delphi study yielded four main classes of managerial influences in the KM ontology.

AKMI2: The main classes of managerial influences are knowledge leadership, knowledge coordination, knowledge control, and knowledge measurement.

The notions of leadership, coordination, control, and measurement are not unique to the conduct of knowledge management. However, how to accomplish them with respect to KMC and the natures of their specific impacts on KMC are not definitively known. Their execution with respect to KMC may require special techniques, methodologies, and technologies.

Of the four managerial influences, leadership is primary. It establishes enabling conditions for fruitful KMC.

DKMI4: Knowledge Leadership—An entity's administrative efforts to create circumstances whereby knowledge processors can most effectively do the entity's knowledge work.

Coordination, control, and measurement are contributors to establishing these conditions, but there is an additional

aspect to fulfilling the leadership mission. This distinguishing characteristic of leadership is that of being a catalyst through such practices as inspiring, mentoring, setting examples, engendering trust and respect, instilling a cohesive and creative culture, listening, learning, teaching (e.g., through story-telling), and knowledge sharing. Core competencies for effective leaders of knowledge-intensive organizations include being a catalyst, being a coordinator, exercising appropriate control, and properly gauging resources, processors, processes, and influences. For discussions of key aspects of knowledge leadership, refer to Amidon and Macnamara (2003) and Bennet and Neilson (2003).

A generic characterization of coordination sees it as concerned with managing dependencies (Malone & Crowston, 1994). In the conduct of KM, dependencies that need to be managed include those among KRs, those among KMAs, those among processors performing these activities, those between KRs and other resources, and those between resources and KMAs. Coordination happens on both an intraepisode and interepisode basis.

DKMI5: Knowledge Coordination—Managing dependencies among KMAs, KRs, knowledge processors, knowledge management processes, and knowledge management episodes.

Coordination approaches suggested and used to manage dependencies in a knowledge-based organization include linking reward structures to knowledge sharing, establishing communications (e.g., communities) for knowledge sharing, constructing programs to encourage innovation and learning, and providing incentives to encourage fruitful KM behaviors on the part of knowledge processors.

Excellent coordination is to little avail if the entity's knowledge or knowledge processors are inadequate. In the conduct of KM, management can take steps to ensure that the entity has sufficient knowledge of sufficient quality (i.e., validity, utility) and sufficient processors sufficiently skilled in needed knowledge processing. Moreover, management needs to protect against harmful exposure/loss of its KRs and processors.

DKMI6: Knowledge Control—An entity's efforts to ensure that needed KRs and processors are available in sufficient quality and quantity, subject to required security.

It is important that the degree of control be appropriate—neither stifling nor cavalier. Guidelines for providing control in organizations are advanced by Jamieson and Handzic (2003). Assurance of knowledge control depends on having a clear picture of the entity's strengths and limitations. Measurement is a valuable basis for appreciating how well an entity's knowledge work is progressing. Such evaluation, in turn, can lead to changes in control, coordination, or leadership approaches, as well as directly influence the way in which knowledge management is done.

DKMI7: Knowledge Measurement—The entity's efforts at gauging and evaluating KRs, knowledge processors, KMAs, managerial influences, knowledge management episodes, and overall conduct of knowledge management.

Knowledge measurement can not only legitimize the KM initiatives within an organization, but can also allow for identification and recognition of value-adding activities and resources. Like the other managerial influences, knowledge measurement is challenging in practice (Stone & Warsono, 2003). Practical guidelines for measuring value of efforts in KM initiatives have been advanced by Hanley and Malafsky (2003).

Aside from managerial influences, the ontology recognizes an entity's resources as influences on the conduct of KM. This includes not just its KRs, but other more traditional resources as well. For instance, varying degrees of financial resources are allocated for KM episodes, human resources with varying knowledge manipulation skills are used in KM episodes, and computer systems are assigned to carry out various instances of KMAs.

AKMIK3: Resource influences on the conduct of knowledge management comprise four major categories: financial resources, human resources, material resources, and KRs.

DKMI8: Financial resources—An entity's financial assets.

DKMI9: Human resources—Skills possessed by an entity's human participants.

DKMI10: Material resources—Capabilities of the entity's material assets, including skills possessed by computer participants.

Recall that the KR definition has already been given in DKMC8 as knowledge conveyed by representations that are subject to manipulation in ways that yield value.

The ontology's third category of influences is concerned with an entity's environment. Unlike managerial influences (and, to a considerable extent, resource influences), environmental influences are factors over which an entity typically has limited (or no) control. These factors may operate as constraints, impediments, or facilitators of the entity's knowledge management efforts. The Delphi process yielded six major classes of environmental influences.

AKMI4: Major types of environmental influences on an entity's conduct of knowledge management include competition, fashion, markets, technology, time, and the GEPSE (governmental, economic, political, social, and educational) climate.

Competition refers to the competitive position in which an entity finds itself. Defending or improving this position may necessitate new knowledge management initiatives or adoption of particular approaches to knowledge management. The fashion factor refers to pressures that an entity experiences to align itself with trends that arise in its environment. Such pressures may inhibit or foster particular

TABLE 2. Some topics implied for a sample of KM issues.

Aspects of KM conduct	Ethics	Outsourcing knowledge	Reusing knowledge	Downsizing	Sharing knowledge
Knowledge manipulation activities	Are knowledge manipulation activities executed ethically?		How can assimilation and selection be executed to foster reuse of knowledge?	How does it affect an entity's ability to acquire and generate knowledge?	How can the assimilation activity be executed to enhance knowledge sharing?
Resources	Are knowledge resources handled ethically?			Does downsizing decrease the skill and knowledge base of an entity?	How can knowledge sharing behaviors become ingrained in culture?
Projection	Are the projections ethically sound?	Does it restrict projection?	Does reuse expedite projection?		Does knowledge sharing enhance the scope of knowledge application?
Learning		Does it have adverse effects on learning?		How does it affect learning?	How does knowledge sharing impact entity learning?
KM influences	Are the ethical values advocated by management, environment, and individuals in harmony?	How can coordination and control be performed?			How can a robust incentive structure that promotes knowledge sharing be instilled?

approaches that an entity employs in its knowledge work. The market factor encompasses both markets for resources that an entity can acquire (e.g., knowledge, knowledge processors) and markets for the entity's projections. The former can make for bottlenecks that restrict ingredients available for conducting KM episodes. The latter can drive those episodes toward particular kinds of knowledge emissions. The technology factor refers to the present state of technology available for an entity (and its competitors) to adopt in efforts to improve its KM efforts. Like the other classes of environmental influences, this factor is dynamic; technological advances demand continuing attention by those charged with designing and executing the managerial influences.

The time factor refers to pressures that an environment exerts on an entity to accomplish specific knowledge work before a deadline has passed. This can constrain the way in which, and quality with which, the knowledge work is accomplished. The time factor is operative within specific instances of KMAs, on the expeditious completion of a knowledge management episode, and on completing inter-related clusters of KM episodes. The final environmental influence is a complex construct comprised of the governmental/economic/political/social/educational climate in which an entity finds itself. For instance, the thrust and approaches of an entity's knowledge work may be restricted by governmental regulations (e.g., privacy laws), economic conditions (e.g., recession vs. expansion), political pressures (e.g., to terminate controversial programs, to tow the line), social climate (e.g., open vs. closed), and educational levels/availability (e.g., university training in KM).

AKMI5: The three categories of knowledge management influences are both distinct and interrelated.

The ontology does not consider the nature or complexities of these inter-relationships.

Implications for Research and Practice

The collaboratively developed KM ontology has several implications for research and practice. It furnishes a relatively unified and formalized view of KM phenomena. It provides researchers with a relatively comprehensive, organized foundation and common language for studying KM. It gives practitioners a frame of reference for assessing KM practices and recognizing KM opportunities. It points toward a structure and content for developing a formal KM curriculum. These contributions are briefly discussed below.

A Unified View of KM Phenomena

The ontology advanced here presents a formal, unified picture of KM by identifying and describing its principal components and their interrelationships in terms of definitions and axioms. In doing so, the ontology helps address the critical question: What is KM? Evidence of the ontology's unifying value exists not only in the Delphi contributors' evaluation shown in Figure 1, but can also be seen in a demonstration of its use as a framework for comparative analysis of other characterizations of KM phenomena (Holsapple & Joshi, 2002b). It gives a common way of viewing each, so that their overlaps become clear, their varying emphases are evident, and directions are suggested for extending each in order to deal with ontological elements not yet covered.

Guidance for KM Researchers

As an aid for conducting research, an ontology's completeness and utility are important, as is satisfaction with the

TABLE 3. Examples of ethical issues suggested by the KM ontology.

Aspects of KM conduct	Some ethical issues in managing knowledge
Knowledge acquisition	Issues: What can be acquired? What should (or should not) be acquired?
Knowledge selection	Issues: Ensuring that privacy requirements are known and enforced. Avoid inappropriate and inaccurate selection due to bias by a participant assigned to the selection activity.
Knowledge generation	Issues: Assessing the potential effects of new knowledge (e.g., is discovering how to clone humans ethical? Is research into creating biological and chemical weapons ethical?).
Knowledge emission	Issues: Abiding by law while projecting an output (e.g., warning customers about possible side effects); using environmentally safe packaging; targeting and transferring to the “right” audience (e.g., an output that is only for adult use).
Knowledge assimilation	Issues: Does selection of employees’ knowledge for assimilation into a computer system devalue humans or is it a means of knowledge sharing across generations that allows an organization to grow? Preventing assimilation of values into organizational resources that are detrimental to community at large.
Knowledge resources	Issues: Knowledge about ethics can be represented in any of the six knowledge resources. It can be in participants’ knowledge storehouses, documented in manuals, ingrained in the culture, appear in the form of infrastructure regulations, and reflected in purpose and strategy. A challenge for a CKO is to ensure that there is minimal conflict among the ethics knowledge embedded in the various knowledge resources.
KM influences	Issues: Ethics can guide how KM conduct is administered within an entity. Personal ethical values of human resources may affect the working of an organization’s ethical system. How do laws affect the conduct of KM? Certain types of laws may inhibit conduct of KM (e.g., they may hamper projection and learning).

ontology. Delphi participants’ favorable views of the ontology’s comprehensiveness, satisfaction, and utility are reflected in Figures 2, 3, and 4, respectively. Figure 4 shows specifically how participants gauge the ontology’s utility for researchers. A majority regards it as being in the helpful to extremely helpful range, and over 90% see it as at least moderately helpful. The ontology allows researchers to frame and generate research issues in a systematic manner and gauge the current state of KM research in an organized fashion. It provides a relatively comprehensive set of KM elements and interrelationships that can be used for generating and testing research models to investigate and discuss KM issues. For instance, Massey, Montoya-Weiss, & O’Driscoll (2002) applied the influence component of this ontology to examine the success of a process-oriented KM strategy at Nortel Networks. In doing so, they validate and suggest extensions for the KM influence component of this ontology.

As an illustration on how the ontology can be applied, a researcher can use it to systematically generate, study, and

discuss KM-related issues via the construction of exploratory matrices. Examples of such matrices are presented in Tables 2 and 3. Such matrices provide a mechanism to identify/generate/organize KM investigations.

Some crucial KM-related issues that are facing researchers and practitioners include ethical issues in managing knowledge, issues of outsourcing knowledge development versus developing it in house, downsizing and its impact on KM conduct, reusing versus regenerating (reinventing) knowledge, and knowledge hoarding versus knowledge sharing. When we consider these sample issues in terms of major components of the ontology, topics such as those shown in Table 2 emerge. The table’s rows list primary elements within the conduct of KM (i.e., the KMAs, KRAs, knowledge influences, learning, and projection). The table’s columns list a sample of KM-related issues. Each cell, at the intersection of each row and column, explores the issues that emerge from the relationship between the row and the column constructs. This type of matrix can help researchers identify research questions that need to be investigated.

TABLE 4. A sample CKO checklist for KM initiatives.

Aspects of KM conduct	Factors to consider in order to create and sustain organizational competitiveness through KM initiatives
Knowledge manipulation activities	<ul style="list-style-type: none"> <input type="checkbox"/> What kinds of mechanisms and techniques are necessary to effectively execute knowledge manipulation activities? <input type="checkbox"/> How to design knowledge manipulation strategies to facilitate effective transfer and reuse of knowledge to appropriate knowledge workers without creating information over load on knowledge workers.
Resources	<ul style="list-style-type: none"> <input type="checkbox"/> What type of programs needs to be created to attract, retain, and develop critical knowledge resources, such as organizational core and ancillary participants? <input type="checkbox"/> How to preserve the expertise of organizations’ core participants when they cease to be the members of the organization (e.g., due to turnover or retirement).
KM influences	<ul style="list-style-type: none"> <input type="checkbox"/> What type of incentives and reward systems needs to be introduced to cultivate and nurture an organizational culture where organizational participants create and share knowledge to continually improve and innovate organizational products and services? <input type="checkbox"/> How to value and account for knowledge resources in a formal fashion such that its impact on organizational growth indicators can be effectively measured.

TABLE 5. Applying the KM ontology to summarize exemplars of KM best practices.

Aspects of KM conduct	KM Best Practices			
	Chaparral Steel	Dow Chemicals	Skandia Inc.	Buckman Labs
Influences				
Managerial	Developed reward and incentive systems that encourage knowledge building	Created ways of measuring knowledge assets	Created ways of measuring and accounting for knowledge assets	Use of incentives and technology for effective coordination of knowledge manipulation activities
Resource	Creating culture conducive for knowledge building			
Activities	Instilling knowledge building activities into manufacturing processes			
Resources		Developed methods for extracting value from knowledge assets	Developed methods for accessing knowledge assets	

An exploratory matrix, such as the one shown in Table 2, can be further developed by selecting one issue and examining it in greater depth using more detailed ontology elements. This is illustrated in Table 3 by extending and elaborating on ethics issues using ontology subcomponents. This table can be read in a fashion similar to that used for Table 2. In this table, each cell lists a set of ethical issues in a more comprehensive manner using the subcomponents on the ontology.

Guidance for KM Practitioners

The ontology provides a common vocabulary and frame of reference that can enhance the communication and sharing of ideas among practitioners. It also provides a checklist of considerations to ponder in the course of planning and conducting KM initiatives. Figure 4 shows the relative frequency distribution of Delphi participants' views about the ontology's utility to practitioners. Notice that the mode is "very helpful," and more than three fourths think the ontology is at least moderately helpful to practitioners.

Exploratory matrices derived from the ontology can provide KM practitioners with a framework for designing, discussing, and evaluating KM initiatives. For instance, as illustrated in Table 4, a chief knowledge officer interested in investigating the impact of KM conduct on organizational growth can devise an exploration matrix to highlight relationships between each of the ontology's facets and various aspects organizational growth (e.g., in terms of learning and projections). This matrix is laid out with factors that a chief knowledge officer should consider in order to create and sustain competitive advantage (in light of indicators of growth) as columns and ontology elements as rows. Each cell in this matrix is a focal point for exploring relationships between the two intersecting constructs. Such exploration proceeds with the intent of identifying, discovering, or creating techniques, tools, or methods for KM that improve organizational growth.

In a related vein, the ontology has been used to devise the knowledge chain model, comprised of primary and secondary KM activities that are focal points for improving an entity's competitiveness in the directions of productivity, agility, innovation, and reputation enhancements (Holsapple & Singh, 2000). The knowledge chain model has been shown to be consistent with cases reported in the literature (Holsapple & Singh, 2001), supported by experiences of surveyed leaders of KM initiatives (Singh, 2000), and capable of more detailed development (Holsapple & Jones, 2003).

A Basis for KM Curriculum Development

Knowledge management is just beginning to find its way into university coursework (Ruth et al., 2003). This ranges from offering individual KM courses, to embedding KM as a theme within traditional business courses, to orienting traditional degree programs (e.g., MBA) around a KM theme, to furnishing a degree program or concentration in KM. The ontology provides a structure that systematically identifies candidate topics for examination in a KM course or curriculum.

Exploratory matrices can help in developing instructional content. For instance, issues identified in Tables 2 and 3 are candidates for discussion points and research paper assignments. The ontology's elements can be juxtaposed with such issue areas as technology (as illustrated in Tables 5 and 6), best practices (as illustrated in Table 7), outcomes (e.g., to explore impacts of KM technologies and practices on an entity's productivity, agility, innovation, reputation, and financial performance), cases (e.g., to explore specific, real-world examples of KM in action for any of the ontology's four components or their elements), lessons learned (e.g., to explore experience-based prescriptions for the conduct of KM), frameworks (e.g., to explore wide-ranging perspectives on KM in a comparative fashion by relating them to the ontology), traditional disciplines (e.g., to ex-

TABLE 6. Examples of knowledge management technologies.

Aspects of KM conduct	Related KM technologies
Manipulation activities	Search and retrieval technologies, intelligent agents; knowledge discovery (e.g., data or text mining), data warehousing, decision support systems (e.g., expert systems, executive information systems), document management technologies.
Resources	Computer participants, storage devices, electronic knowledge bases.
Influences	Groupware, intranets, extranets, messaging (e.g., e-mail), and collaborative technologies for coordination and control.
Projection	Use of technology to advertise (e.g., use of multimedia), provide services (e.g., online-services), design products, and manufacture products.
Learning	Technologies supporting computer-based training.

plore the implications of KM and the treatment of its elements in the contexts of financial management, production management, marketing, accounting, economics, strategy, human resource management, business law, business computing, international business, healthcare administration, and so forth), and entity levels (to explore commonalities, contrasts, and relationships among the personal through national and transorganizational levels of KM).

A matrix such as the one illustrate in Table 5 can be used in the classrooms to compare, contrast, and summarize KM best practices. The table's rows list primary elements within the conduct of KM, and its columns list the companies under evaluation. Each cell lists a best practice example that relates to its row's construct and its column's company.

Tables 6 and 7 can be used in the classrooms to illustrate how the ontology can be employed to systematically arrange an exploration of available technologies. Students can then study and report on the indicated technology clusters. Moreover, they can be used by practitioners seeking to identify/study clusters of KM technologies that have the potential to suit their needs, KM technology providers to evaluate their products and find functionalities that can be added to enhance them, or by researchers to devise taxonomies for KM technologies to guide investigations.

Examples of individual courses include Principles of KM for Management Students (a KM overview organized on the foundation furnished by the generic ontology), KM Practices (covering descriptive and prescriptive views of how KM is or should be practiced), and KM Technology (covering emerging and conventional technologies that can enable or facilitate KM conduct in an organization). For the latter, the ontology suggests a way to organize coverage of technologies based on their applications (e.g., technologies

related to managerial influences, technologies that perform KMAs, and technologies that can help enhance participants' skills in performing those activities). Another example of an individual course is Measuring and Assuring Knowledge Assets, which would be especially relevant for accounting, economics, or finance majors. The ontology's resource component, coupled with the measurement and control influences, would give structure to such a course. Similarly, courses focusing on Knowledge Coordination and Knowledge Leadership, integrated with traditional coverage of coordination and leadership, would be beneficial for a well-rounded appreciation of KM.

An example of a KM degree program would be one that focuses on developing future leaders of organizational KM initiatives. This program can have a series of courses designed to help students understand how knowledge-based organizations work, as well as equipping them with methods and tools for managing such organizations. The generic ontology can be used to help ensure completeness, unity, and parsimony in developing such a program. The program could be designed around the four ontology components. The influences component can be used to develop a module on KM influences (e.g., courses on KM administration dealing with leading, controlling, coordinating, and measuring conduct of KM in an organization; a course dealing with environmental enablers and constraints on an entity's knowledge work). The resource component can be used to develop a module examining KRs (intellectual capital, organizational memory, and the nature of knowledge assets) and knowledge processors—both human (individual and social) and technological (e.g., decision support systems). The knowledge manipulation component can be used to develop a module on understanding and exercising skills

TABLE 7. Examples of candidate technologies that support aspects of KM.

Aspects of knowledge manipulation activities	KM technologies for knowledge manipulation
Acquisition/selection	Knowledge identification tools (e.g., search technologies, intelligent agents), knowledge capturing tools (e.g., retrieval technologies), knowledge organization tools (e.g., knowledge visualization tools).
Assimilation	Tools that help in knowledge storage (by creating knowledge warehouses), capture organizational expertise (e.g., in the form of expert systems or case bases), aid in the assimilation process (through computer-based training).
Generation	Data mining, decision support systems; expert systems, executive information systems, modeling.
Knowledge flows	Messaging systems (e.g., e-mail), knowledge exchange (chat-rooms, electronic bulletin boards, video conferencing).

needed to operate on KRs (e.g., coursework on performing the various manipulation activities). The KM conduct component could lead to an introductory KM principles course and/or a capstone course that tie together the other modules, examine outcomes of the conduct of KM (e.g., organizational learning, knowledge commercialization via projections, competitive impacts advanced by the knowledge chain model), and delve into KM strategy (e.g., its alignment with business strategy and organizational vision/mision).

In a related vein, the ontology and consequent knowledge chain model has been used to conceive and structure the contents of a major KM reference book, the two-volume *Handbook on Knowledge Management* (Holsapple, 2003b,c). The contents of this book have been used to expand the Knowledge Management Professional Society's CKM Program for certification and are being woven into its eCKM curriculum as well (Weidner, 2003).

Conclusions

It is crucial to lay a strong foundation on which future KM research, practice, and education can develop. This research contributes to that foundation by formally advancing a generic, descriptive ontology that forms a cradle for KM research, study, and practice. The collaboratively developed ontology identifies and characterizes major elements of KM in a unified, relatively comprehensive manner. It describes organizational KRs where knowledge may be stored, embedded, and/or represented. The ontology identifies and relates KMAs that operate on those resources. It recognizes factors that influence the conduct of KM in an organization.

This ontology is intended to stimulate further conceptual development in the KM field. The ontology can further evolve through added breadth and depth. It can be extended in a normative direction by adding elements that prescribe methods and technologies for the conduct of KM. Future research and practice will more fully determine the extent of this ontology's utility and applicability.

Acknowledgments

This research was supported in part by the Kentucky Initiative for Knowledge Management, established in 1988 at the University of Kentucky. We are indebted to the following persons for their participation as Delphi panelists (those who participated in every round are denoted by an asterisk): Debra Amidon,* ENTOVATION International, Ltd., USA; Sulin Ba, University of Southern California, USA; Thomas J. Beckman,* George Washington University & IRS, USA; Kesper Deboer, Andersen Consulting, USA; Marc Demarest, The Sales Consultancy, USA; Alain Godbout,* Godbout Martin Godbout & Associates, Canada; Valerie Cliff, ICL Enterprise Consultancy, UK; Ming Ivory,* James Madison University, USA; Linda Johnson,* Western Kentucky University, USA; Mark A. Jones,

Andersen Consulting, USA; Sam Khoury, The Dow Chemical Company, USA; Kai Larsen, Center for Technology in Government, USA; Dirk Mahling, University of Pittsburgh, USA; Eunicka Mercier-Laurent,* EML Conseil-Knowledge Management, France; Philip C. Murray, Knowledge Management Associates, USA; Brian Newman, The New- man Group & The KM Forum, USA; David Paradice,* Texas A&M University, USA; Gordon Pettrash,* The Dow Chemical Company, USA; Dave Pollard,* Ernst & Young, Canada; Larry Prusak,* IBM Corporation, USA; David Skyrme,* David Skyme Associates Limited, UK; Charles Snyder,* Auburn University, USA; Kathy Stewart,* Georgia State University, USA; Karl Sveiby, Sveiby Knowledge Management, Australia; Robert Taylor,* KPMG Management Consulting, UK; Karl Wiig, Knowledge Research Institute, Inc., USA; Andrew Whinston,* University of Texas, Austin, USA; Fons Wijnhoven,* University of Twente, The Netherlands; Dennis Yablonsky, Carnegie Group, Inc., USA; Michael Zack, Northeastern University, USA. One participant prefers to remain anonymous.

References

- Alavi, M., & Leidner, D. (2001). Knowledge management and knowledge management systems: Conceptual foundations and research issues. *MIS Quarterly*, 25(1), 107-136.
- Amidon, D., & Macnamara, D. (2003). The 7 C's of knowledge leadership: Innovating our future. In C. Holsapple (Ed.), *Handbook on knowledge management—knowledge matters* (pp. 539-552). Berlin: Springer-Verlag.
- Bacon, J., & Fitzgerald, B. (1996). The field of IST: A name, a framework and a central focus. *Executive Systems Research Center Working Paper Series*, 96(5), 1-35.
- Bennet, A., & Neilsen, R. (2003). The leaders of knowledge initiatives: Qualifications, roles, and responsibilities. In C. Holsapple (Ed.), *Handbook on knowledge management—knowledge matters* (pp. 523-538). Berlin: Springer-Verlag.
- Drucker, P. (1993). *Post-capitalist society*. New York: HarperCollins.
- Gruber, T.R.. (1995). Toward principles for the design of ontologies used for knowledge sharing. *International Journal of Human and Computer Studies*, 43(5/6), 907-928.
- Hanley, S., & Malafsky, G. (2003). A guide for measuring the value of KM investments. In C. Holsapple (Ed.), *Handbook on knowledge management—knowledge directions* (pp. 369-390). Berlin: Springer-Verlag.
- Holsapple, C. (1995). Knowledge management in decision making and decision support. *Knowledge and Policy*, 8(1), 5-22.
- Holsapple, C. (2003a). Knowledge and its attributes. In C. Holsapple (Ed.), *Handbook on knowledge management—knowledge matters* (pp. 165-188). Berlin: Springer-Verlag.
- Holsapple, C. (Ed.). (2003b). *Handbook on knowledge management—knowledge matters*. Berlin: Springer-Verlag.
- Holsapple, C. (Ed.). (2003c). *Handbook on knowledge management—knowledge directions*. Berlin: Springer-Verlag.
- Holsapple, C., Johnson, L., & Waldron, V. (1996). A formal model for the study of communication support systems. *Human Communication Research*, 22(3), 421-446.
- Holsapple, C.W., & Jones, K. (2003). Toward an elaboration of the knowledge chain model. *Proceedings of the Americas Conference on Information Systems*, Tampa, FL, August 4-5.
- Holsapple, C.W., & Joshi, K.D. (2000). An investigation of factors that influence the management of knowledge in organizations. *Journal of Strategic Information Systems*, 9(2-3), 235-261.
- Holsapple, C.W., & Joshi, K.D. (2001). Organizational knowledge resources. *Decision Support Systems*, 31(4), 39-54.

- Holsapple, C.W., & Joshi, K.D. (2002a). A collaborative approach to ontology design. *Communications of the ACM*, 44(2), 42–47.
- Holsapple, C.W., & Joshi, K.D. (2002b). The evolution of knowledge management frameworks. In S. Barnes (Ed.), *Knowledge management systems: Theory and practice* (pp. 222–242). London: International Thomson Business Press.
- Holsapple, C.W., & Joshi, K.D. (2002c). Knowledge manipulation activities: Results of a Delphi study. *Information and Management*, 39(6), 477–490.
- Holsapple, C., & Luo, W. (1996). A framework for studying computer support of organizational infrastructure. *Information and Management*, 31(1), 13–24.
- Holsapple, C., & Singh, M. (2000). The knowledge chain. *Proceedings of the Annual Conference of the Southern Association on Information Systems*, Atlanta, GA, March 31–April 2.
- Holsapple, C., & Singh, M. (2001). The knowledge chain model: Activities for competitiveness. *Expert Systems with Applications*, 20(1), 77–98.
- Holsapple, C., & Whinston, A. (1996). *Decision support systems: A knowledge-based approach*. St. Paul, MN: West Publishing.
- Jamieson, R., & Handzic, M. (2003). A framework for security, control, and assurance of knowledge management systems. In C. Holsapple (Ed.), *Handbook on knowledge management—knowledge matters* (pp. 477–506). Berlin: Springer-Verlag.
- Malone, T., & Crowston, K. (1994). The interdisciplinary study of coordination. *ACM Computing Surveys*, 26(1), 87–119.
- Marshall, N., & Brady, T. (2001). Knowledge management and the politics of knowledge: Illustrations from complex products and systems. *European Journal of Information Systems*, 10(2), 99–112.
- Massey, A.P., Montoya-Weiss, M.M., & O'Driscoll, T.M. (2002). Knowledge management in pursuit of performance: Insights from Nortel Networks. *MIS Quarterly*, 26(3), 269–289.
- Newell, A. (1982). The knowledge level. *Artificial Intelligence*, 18(1), 87–127.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- Randall, D., Hughes, J., O'Brien, J., Rouncefield, M., & Tolmie, P. (2001). Memories are made of this: Explicating organisational knowledge and memory. *European Journal of Information Systems*, 10(2), 113–121.
- Ruth, S., Shaw, N., & Frizell, V. (2003). Knowledge management education: An overview of programs and instruction. In C. Holsapple (Ed.), *Handbook on knowledge management—knowledge directions* (pp. 581–604). Berlin: Springer-Verlag.
- Schein, E. (1985). *Organizational culture and leadership*. Washington, DC: Jossey-Bass Publishers.
- Singh, M. (2000). *Toward a knowledge management view of electronic business: Introduction and investigation of the knowledge chain model for competitive advantage*. Doctoral dissertation, University of Kentucky, Lexington.
- Spender, J.-C. (2003). Knowledge fields: Some post 9/11 thoughts about the knowledge-based theory of firm. In C. Holsapple (Ed.), *Handbook on knowledge management—knowledge matters* (pp. 59–71). Berlin: Springer-Verlag.
- Stone, D., & Warsono, S. (2003). Does accounting account for knowledge? In C. Holsapple (Ed.), *Handbook on knowledge management—knowledge matters* (pp. 253–270). Berlin: Springer-Verlag.
- Sutton, D.C. (2001). What is knowledge and can it be managed? *European Journal of Information Systems*, 10(2), 80–88.
- Weidner, D. (2003). *In thought & practice* [Book review; online]. Retrieved October 1, 2003, from http://www.kmpro.org/journal/Book_Review/Handbook_on_Knowledge_Management.cfm