A Merged Future for Knowledge Management and Enterprise Modeling

H. T. Goranson, Kurt Kosanke, and Michael N. Huhns

Abstract Organizational knowledge typically comes from many independent sources, each with its own semantics. This paper describes a methodology by which information from large numbers of such sources can be associated, organized, and merged. The hypothesis is that a multiplicity of ontology fragments, representing the semantics of the independent sources, can be related to each other automatically *without* the use of a global ontology. That is, any pair of ontologies can be related indirectly through a *semantic bridge* consisting of many other previously unrelated ontologies, even when there is no way to determine a direct relationship between them. The relationships among the ontology fragments indicate the relationships among the sources, enabling the source information to be categorized and organized. A preliminary evaluation of the methodology has been conducted by relating 53 small, independently developed ontologies for a single domain. A nice feature of the methodology is that common parts of the ontologies reinforce each other, while unique parts are de-emphasized. The result is a *consensus* ontology.

1 INTRODUCTION

Corporate information searches can involve data and documents both internal and external to the organization. The research reported herein targets the following basic problem: a search will typically uncover a large number of independently developed information sources—some relevant and some irrelevant; the sources might be ranked, but they are otherwise unorganized, and there are too many for a user to investigate manually. The problem is familiar and many solutions have been proposed. The solutions range from requiring the user to be more precise in specifying search criteria, to constructing more intelligent search engines, or to requiring sources to be

more precise in describing their contents. A common theme for all of the approaches is the use of ontologies for describing both requirements and sources. Unfortunately, ontologies are not a panacea unless everyone adheres to the same one, and no one has yet constructed an ontology that is comprehensive enough (in spite of determined attempts to create one, such as the Cyc Project, underway since 1984). Moreover, even if one did exist, it probably would not be adhered to, considering the dynamic and eclectic nature of the Web and other information sources.

Enterprise modelling and knowledge management should be key contributors to decision making in an enterprise. Managers, engineers, and technicians all need knowledge and expertise in order to be most effective. Whether the necessary knowledge is internal or external to an enterprise, it needs to be located, reconciled, and focused on problems at the very moment when it can have the greatest benefit. From an idealistic viewpoint, the entire corporate expertise should be brought to bear on each problem or decision. For this to happen, the knowledge must be organized to be locatable and understandable: this can be provided by EM, with the result that the knowledge is isomorphic to the enterprise itself.

Both knowledge management and enterprise modeling are well established in enterprises today. In both one can find:

- A robust vendor and consulting community.
- Well-established university research groups, funded in different ways.
- An active press, targeting both managers and technicians.
- Enough promise — supported by case studies — to fuel continued investment and implementation.

However, there have been some spectacular failures and some vexing limits in successful implementations.

The workgroup examined the relationship between knowledge management (KM) and enterprise modeling (EM). The specific focus was on possible combined futures and the research roadmap these futures require. The workgroup concluded that a combination of techniques from KM and EM shows promise in addressing the limitations of each.

1.1 Enterprise Modeling and Integration: Background

Enterprise modeling is done for a purpose, and an important one is to support the optimization of operations, through what is termed Enterprise Integration (EI). This is a fundamental business need with direct and measurable benefit. For some time, there have been many techniques to model processes and other elements of the enterprise. Modeling in this context means creating an explicit representation, usually computable, for the purposes of understanding the basic mechanics involved. One often uses that understanding to measure, manage, and improve the process or element.

A basic problem is that there are many types of elements to be modeled in an enterprise, and many perspectives and contexts in which those models would be "viewed." Enterprise integration in this context combines models and their uses in such a way that the whole system can be seen in various coherent ways and for multiple purposes. EI provides a model framework in which components can be interrelated.

Some EM and EI systems are wholly computable. Enterprise Resource Planning (ERP) is one that focuses on specific tasks, delivering planning and control functions. It generally requires a constrained modeling approach and heavy use of generic models, thus restricting the processes for better or worse. The more general EI philosophy is framework based, such frameworks supporting:

- Levels of model genericity to enable model and best practice reuse.
- Relationships among different views (for instance views needed to see organizational linkages versus information flows).
- Relationships among different types of basic entities in the enterprise; for instance, activities need to be modeled differently than roles or resources.

CIMOSA is a strong example of such an integrating framework, major elements of which are standardized internationally as ISO 15704. EI frameworks are widely used, especially in the subset of ERP noted above and a similarly focused subset of Product Data Management (PDM), which supports activities centered on the evolution of product features as they are transformed by processes in the enterprise. Enterprises that use computer-aided design heavily implement EM in this fashion.

1.2 Enterprise Modeling: Problems and Limits

The major problems of EM are of two types. First, EM assumes that one knows what should be made or done, who will do it, and a precise notion (perhaps to change later) about how each element of work will be done. Because the primary leverage from the approach is the system view, some substantial part of the system must be included in the model. But those enterprises desiring a system view might wish to include strategic marketing and product design elements, if applicable. Such processes aren't as easily captured as process models however: they have "soft" elements like unknown futures, tacit knowledge, and poorly understood cultural and collaborative dynamics.

Second, EM usually deals with the normative, stable, deterministic case. In other words, managers expect their world to remain as it is because they are going to great lengths to engineer an operational enterprise. Dynamic environments, evolving processes, shifting partnerships and changing products are a way of life for many enterprises. So if EI is employed, it must

be more federated than unified. That means the EI system must ideally be cheap to assemble, must change the source models and process in little or no way, be responsive to change, even indicate change, and be to some extent self-organizing and adapting.

Adding KM techniques to the mix can mitigate these two problems, possibly in a revolutionary manner.

1.3 Knowledge Management: Background

Knowledge Management solutions address several needs that all share the underlying notion that enterprises depend heavily on individual and institutional knowledge, and the knowledge must be better understood and managed. KM is a set of philosophies, tools, and techniques to support various functions within this need. While both KM and EM address pressing business needs, EM originates from the industrial engineering and operational perspective and is technique-centric; KM originates from the management perspective and is needs-centric. The two communities have a poor history of deep collaboration, which may explain why such an apparent synergy has been hitherto unexploited.

The discrete problems addressed by the KM community are:

- A need to capture individual knowledge to make it "institutional" knowledge so that it can be reused in the enterprise, and persist when an expert leaves.
- A second order intent to use standardized knowledge elements and communication methods to develop and support corporate culture for competitive benefit.
- Support for the "learning organization:" education at the individual, team and enterprise level.
- The development of knowledge "metrics." Significant investment is wrapped up in knowledge, and there is currently no good way to quantify the value of the result. Metrics are needed by financial accountants to evaluate capital knowledge assets; and planners using simple cost-benefit analyses in decision-making.
- Auditability of intellectual property. Tracking the initiation of an idea and the various inputs can reveal who contributed what and when and prove it in court.
- Self-awareness. The better you "know" yourself and your relationship to the world, the better you can change and manage yourself. This notion is the very same driver as in EI, where it is focused on operations, but in the KM world it is more focused on strategic planning.
- KM is often invoked as the backbone around which diverse corporate cultures will be combined after a merger or acquisition.

Because the needs of KM are more diffuse, the tools and implementations are too. Many tools are simply ways of aiding collaboration by structuring the way information is stored, indexed, and shared. Also, many of the techniques are "soft" and merely philosophical, motivational or concerned with building awareness.

1.4 Knowledge Management: Problems and Limits

The general problems with KM systems are of two types:

- KM systems are "soft," almost by definition. They deal with intellectual property for which no good value metrics exist; they deal with collaborative contexts that are not well modeled; and they implicitly address the slippery reality of "tacit" knowledge. Many KM systems deal with strategic planning, which means they address uncertain futures, but without extrapolating from the current situation. The current situation is often described only by an EI or other operational system, whether or not formalized and automated.
- KM systems deal with both "know-what" and "know-how," but with little emphasis on the "how." In other words, the knowledge is not sufficiently bound to the work of the enterprise, or what that work might become. One part of this problem is the age-old lack of linkages between strategic planning and operational management — it is not just an impedance mismatch between functions, but between methods and basic representations as well. This mismatch frequently produces strategic decisions that make little sense.

Just from this brief overview, the reader may already be anticipating suggestions from the working group on how the strengths of one approach could strengthen the weaknesses of the other. KM needs formalisms (which might help with metrics) and anchoring in the enterprise's actual work; EM needs ways of dealing with knowledge about context and other soft elements, specifically including tacit knowledge.

2. NEAR TERM FUTURE: DEDUCTIVE TRUST AND PROCESS SITUATING

The workgroup recognized a few near term synergies between EM and KM.

"Knowledge" in the KM context is "justified true belief." Each of those three words conveys different dimensions of trust in the information. Usually that trust is "inductive;" the trust is based on (in ascending order of "closeness" to your own judgement):

- Authority: Someone in the enterprise represents that the knowledge is to be trusted. This person might be trusted by you, in which case you are trusting that person as a certifier of sorts; but usually you are delegating trust.
- Votes: The second case above involves a certifying agent that has the authority of the enterprise, which can be seen as a case of enough votes of the right kind. This type involves votes directly on the information itself. You might not have cause yourself to trust the information, but some group dynamic provides additional confidence, by aggregated authority or broadened depth. (There are likely several group mechanisms involved here, but the workgroup did not exhaustively explore them.)
- Experience: You have seen this case before with enough similarity and enough times to have confidence that it will turn out the same way the next time.

But there is a different basis on which one might base trust, a "deductive" basis that involves understanding the cause and effect mechanics behind the situation in sufficient detail to determine the outcome. For example, one may have experienced many sunrises so have inductive confidence that the sun will rise again tomorrow. Or that person may have deductive trust based on knowledge of the planetary mechanics that produce sunrises. Deductive trust produces a better foundation for justified true belief.

In the business enterprise, deductive trust is much preferred, because it is auditable: decision makers can — if so inclined — "audit" the trust behind the knowledge by zooming in on the underlying physics. Most knowledge in an enterprise is of the inductive type and this is reflected in current KM systems, whereas managers want most to be of the deductive type. Enterprise models capture cause-and-effect dynamics within the enterprise, so a marriage would seem manifest destiny. In this case, each element of knowledge in the KM system is linked to modeled processes (representing activities) in the EM.

Such a linkage can be made during the (already costly) modeling and knowledge capture processes without unduly extending the difficulty of either. The benefit to KM would be rather profound: some significant portion of the knowledge will be (or be expected to be) deductively auditable by linkage to actual processes. Another way of putting this is that knowledge in a KM system is know-how; current KM approaches focus on the "know," but not the "how." Linkage of KM to EM provides the how. And that "how" linkage provides a significant benefit: maintaining knowledge costs money — maintaining vitality in that knowledge base costs more.

Knowledge managers need to know which knowledge to "forget." If there is not a robust linkage to processes (current and future), the knowledge

has no apparent relevance to the business. That should prompt an examination with one of the following results:

- The EM is incomplete and needs to be extended. In this case, the existing knowledge indicates what processes need to be better modeled or added. Experience indicates that this can be a powerful technique for modeling processes that have "soft" mechanics, such as many marketing processes.
- The knowledge is determined to be not relevant, and can therefore be deliberately forgotten. The ability to know what is not relevant is an important step in a system's knowledge of itself, which in turn is a necessary condition for being a "learning organization." Knowledge should be deliberately forgotten because it is out of date; because of machine constraints on storage or search time; or because it can be more robustly handled by a collaborating agent.
- The knowledge is determined to be relevant, but poorly supported by processes in the existing enterprise. This would indicate modifying the enterprise. Often the solution in this instance is to develop business partnerships with entities that can support the knowledge process linkage either by supplementing the source enterprise, or maintaining that knowledge itself.
- The situation is the complement of the first case, where an EM is more complete than the knowledge base. This can be used as an indicator of knowing what you don't know within the universe of interest.

This rounds out the four likely conditions for full KM: knowing what you know in a trusted way; knowing what you can forget; knowing what you do not know; and knowing what you can delegate. Knowledge resides in the individual, but has value in the context of the enterprise. KM can be seen as the management of pieces of knowledge, while EM can be seen as the compositional framework for those pieces.

Another way of understanding the problem is to consider a breakdown of KM into four elements: revealing information; forming and managing facts; forming and managing relationships and contexts among facts; and applying that knowledge to effect. Today's KM systems do the first two well enough, but need help with the other two.

EM may help with understanding contexts. The basic idea behind EM is taking fragments of information within the enterprise and placing them in a larger context. EM provides a registration framework for the parts that relate one to another. But this framework relies on artifacts of the modeling process that capture local interdependencies. KM systems based on *ontologies* can allow global registration. Ontologies are formal descriptions of elements and behaviors, originally devised to help share knowledge between systems employing different representations.

A focus on ontologies should provide a bridge between EM and KM, but the leverage is likely to come more from the EM side, because enterprise models are based on the notion of activities and outcomes, which automatically captures a notion of local dependencies among information elements. This notion is what — at root — allows compositions into larger context and systems. The state of the art in process ontologies is the Process Specification Language, developed at the U. S. National Institute of Standards and Technology [PSL citation] and proposed as an international standard.

To provide a bridge between KM and EM, PSL is the likely starting point. In particular, the combination of a PSL-like ontology structure and CIMOSA-like composition strategies can be overlain on existing KM tools and theories to provide for system behavior and business context. Both PSL and CIMOSA (or substitutes) will have to be examined carefully for needed extensions. Neither was designed for this larger, more ambitious role.

The "effect" problem in KM is the problem of linking each piece of justified knowledge to a business role. The workgroup believes EM can help if there is a slight shift of emphasis from the normative notion of "task" in EM. EM is concerned with doing work, and processes that perform tasks are the logical currency. But knowledge is more naturally seen as being applied to solve problems. So a "problem-centric" notion of the basic unit is proposed as a bridging strategy. A problem is seen as a combination of a task (or set of tasks) together with an element (or elements) of knowledge.

At first glance, this seems an immediately implementable strategy to take short-term advantage of synergies between existing EM and KM tools and techniques. The workgroup proposes serious research focused on this likely "low hanging fruit."

There is a precedent for the sort of merger suggested here, and an example of how quickly the result can spread and become the normal way of doing things. Financial management is a matter of collecting many pieces of information and managing them in much the same way that KM intends to manage knowledge. In fact, financial knowledge is a simple case — the qualitative case — of general knowledge; so KM is a generalization of financial management.

About two decades ago, accounting reached a crisis very similar to the KM crisis today: (financial) knowledge was collected but not relevantly "situated." All of the problems noted above existed in some form. The response was Activity Based Costing (ABC), which simply uses a reduced form of enterprise model to ground individual costs and provide a way of intelligently assembling and relating them. ABC went from a proposal to standard practice in less than a decade; substantial benefits resulted. The near-term EM/KM proposal simply extends this logical evolution. As with the ABC revolution, a key strategy is to continue the same basic tools

already in place; in this case, that means to continue using the operational and business process modeling methods that are already part of the management toolkit.

In the KM context, most KM is non-formalized and non-managed so of course it is non-computable. Informal KM is a human-to-human phenomenon based on personal networks. So this end of the merged KM/EM system must leverage and ride on top of the human infrastructure.

3. MEDIUM TERM FUTURES: FACT BASED DECISION MAKING

EM is generally focused on tactical optimization and similar types of self-examination. But many enterprises have their most pressing needs in strategic planning in the context of uncertain futures. The more uncertain the future, the more significant the threats and opportunities, but the less valid are simple extrapolations from the past.

The importance of thinking about the future is paramount for many enterprises, and for these real resources must be committed for designing processes to be able to respond in an agile way. Decisions are weighty and should be deductive where possible. Often this is termed "fact based decision making," and it is frequently supported by iterative simulations of what-if situations.

The connection of this task with both EM and KM is straightforward and obvious. "Traditional" EM structures processes so that systems can be optimized. EM for simulation (though not recognized as such) does precisely this with the twist that the processes are executable representatives of the processes. Models in most conventional EI systems don't have this character, they are representatives used to understand, not control processes. But the extension to control is not so great in many cases, and indeed modern EI systems perform substantial but limited control. The further extension to simulatable elements is also not so great, generally involving substituting synthetic stimuli for real ones. So it seems quite logical and cost effective to speak of EM in the context of strategic simulation, especially when the basic unit is the problem as suggested above.

(It should be noted that the advantage does not flow the other way. Most built-from-scratch simulation systems use "models" that cheaply emulate the behavior of processes. This cheapness is usually achieved by not modeling the underlying "physics" of the system; also the granularity is not determined by the unit of work as seen at the level of the work, but at some coarser subsystem granularity. As a result, simulation-derived models cannot easily be adapted for wider purpose.)

The workgroup has three recommendations to make at this medium term horizon and in the context of strategic, fact-based simulation.

- The merger of EM and KM should be extended (and justified by) the use of the combined, structured knowledge/process base for simulation. The advantages are potentially profound because of the reuse of information, the running start in well-founded infrastructure that works, and the hard-won existing, practical binding to the way things are really done. The technical challenges seem to be in "packetizing" knowledge elements from the KM side and adding a few new expressions to modeling methods on the EM side.
- Notions of reuse should be better exploited. The advantages of this are seen as similarly profound. The basic problem is that KM systems are generally case-based, meaning that the knowledge and its representation are bound in specific cases containing details that are irrelevant artifacts of how the information appeared. It is hard work to wade through cases to find relevant insights, extrapolate what is needed, and apply it in a specific new context. The preferred alternative is to build analogy-based KM systems, which index and manage information at a more generic and reusable abstract level.

Analogy-based systems are hard to build, and certainly not expected in the near term. But the first step toward such systems may not be so far away. It concerns clear guidelines about what is generic and what is specific to a task, problem or application. As it happens, EI frameworks are nearly universal in dealing with this problem in some way. Unfortunately, the solution is a matter of art specific to the expert who is the source for the knowledge being modeled. It probably is the case that every practical determination of what is generic must be captured in this manner. In other words, it is a type of metaknowledge that is captured at the same time and using the same methods as the "base" knowledge. The format comes from the integrating framework.

The bottom line is that KM systems can take a large step toward identifying generic analogies by adopting EM methods when collecting knowledge from experts.

- The final medium term recommendation concerns knowledge feedback, or self-reinforcing truths. An example is when a prominent stock analyst predicts a stock will rise. It does in part because of her recommendation, which further reinforces confidence in her "analytical" ability. It turns out that many dynamics in an enterprise may be of this type. For example a quality metric may indicate quality because second order dynamics may have adjusted or grown up around it to promote quality results. For instance, a quality metric may be related to number of inspections, and the precision of those inspections adjusted to the fact that the system drives to many

inspections. In fact, the same quality could be achieved with fewer inspections, but only by breaking the cycle of driving toward many, promoted by the "truth" feedback.

Both EM and KM systems have this problem. Usually it is concealed in so-called "tacit" knowledge, which is the concern of many KM systems. But tacit knowledge is a famously black hole, not exhaustible. Good KM practices will help identify which tacit knowledge needs to be captured and why, and (sometimes) at what cost. But these truth feedback loops are best identified when they are deliberately broken as experiments, for instance actually trying to reduce the number of inspections while taking concurrent action elsewhere. One can practically do this only in simulated enterprises, which brings us back to the merger of EM, KM and strategic simulation.

The workgroup did not have time to make specific recommendations of steps and research issues toward solving this problem. But there is a general feel that opportunities are available when the problem is well stated and the more near term steps noted above are taken.

4. LONGER TERM FUTURES: SELF ORGANIZING ENTERPRISES

The workgroup considered the next generation of EI systems. These are likely to exhibit federating behavior and to do so using an agent system. They are also likely to cover much more of the enterprise. The new scope will include at minimum some strategic planning and product definition as one dimension of expansion, and some human, knowledge, and collaboration dynamics in the other dimension.

Agents in this context would likely be the result of evolution from first generation models that represent the superficial behavior of a process, and the second generation noted above where the models capture underlying physics and can be exercised in a simulation environment. Third generation models will be agents, small pieces of software code that include the model and have the ability to negotiate among themselves to optimize the system.

The result will be federated enterprise integration where the system self-integrates. Note that the integration is at the model level, not the enterprise proper. But since these models have the ability to control, the effect is much the same.

This vision of EI was already identified in the second ICEIMT when creating a capability model for integrated systems. A high level of integration was when a process had the ability to see itself, see its context in the system, and change itself to optimize the system — perhaps in collaboration with others — even when it would apparently "harm" the agent. Presumably, the risk-reward environment would be structured to

reward this behavior, and even reward an earnest but unsuccessful search for such optimization.

A higher level of integration is achieved when an agent has the ability to see into the system — following a relationship chain of some sort — discern a change in the system that would optimize the system, and effect that change. In this scenario, all of the agents involved would be rewarded in some way. For example, you may have a set of processes that do nothing but search and optimize for agility against a likely general change. If the enterprise were a virtual enterprise, this agent would be looking at processes involved in the work and others not currently engaged. All processes are in different formats, use only partially integrated applications, and cross business and cultural boundaries. Agents in these companies would be expected to enthusiastically support simulations that could eliminate them from the partnership. In fact, each company is expected to devise novel notions to support this process. This was considered an achievable goal.

In this case, distinctions among knowledge bases, operational process models, business processes, financial metrics and simulation agents will have all but disappeared. But there clearly are barriers. Perhaps the key barrier concerns realities of agent mechanics. As noted above, these agents need to know themselves and what they know, know what they do not, know where to get trusted information remotely, know what to forget, and know the system's goals and associated metrics. Perhaps it will collaboratively determine those goals.

Knowledge managed by these agents will include soft elements such as unknown futures, tacit knowledge and collaborative (cultural) dynamics. The system will integrate (in addition to factors currently handled by EI frameworks) product features, process features, and system features. (This latter incorporates the system optimization metrics.) The managing context will be through bounding constructs (for instance discretely supervised profit centers), practical constraints (such as physical transport of subassemblies) and financial and implementation motivations.

The good news is that lots of work by bright people is going into the general case. The business case provides a much simpler universe than "real life" because businesses (not necessarily their employees) are presumably motivated by financial rewards that are quantifiable. There are only complications about deferred rewards (market share, stock price, increased capability, new markets and the like). Moreover, the business application can justify significant investments in research and products — a repeatable improvement of only a few percent means hundreds of billions a year. Moreover, an agent-based system seems inevitable because it is the only scalable strategy for either knowledge or model management.

Agents are introduced to mitigate complexity, so agents themselves will be engineered for simplicity. One strategy will be to devise agents that all

behave the same. The reason is that each agent has to know how the others will behave; if they are all the same and the agent "knows itself" (or has recourse to examine itself), it can predict how others will behave.

There are likely to be many thorny research issues, but the workgroup focused on two related ones that are key. The first involves harmonizing the notion of uniform agents with the wild variety of models likely to be involved. Recall that at this level of federated integration, diversity of methods is expected, even encouraged. Obviously, some sort of agent wrapper must be devised. The work indicated in the near and midterm agenda sketched above indicates that this wrapper structure will almost certainly be designed at the ontology level, built on extensions to PSL. This work will begin on a firm basis because the first extensions to the existing PSL base will be known agent needs. The most prevalent approach would be to use "speech acts" which have several formal advantages and the elegant property of being intuitively related to processes as they are currently modeled.

The second challenge indicated for attention by the workgroup is the so-called multilevel agent problem. This problem has an analog in the real world: not all processes need or want the same level of freedom. Some collection of processes or organizational elements will be bound more tightly within the enterprise. For instance, several processes will typically be collected in a partner company. The processes act as agents, but the company does too, and one is not a simple sum of its constituents. Similar aggregations may occur by functions and many aggregations may overlap.

The research challenge is to design the wrapper so that it can both support the aggregation process and accommodate the agency of these higher-level agents. Clearly, this strategy will be framework-based, by methods extended from today's EI frameworks.

5. EXTRA CONSIDERATIONS

In addition to the ambitious agenda noted above, the workgroup raised three issues to be considered by the EI and KM communities.

The first is a common suggestion that needs to be underscored. EI and KM are generally thought of as something that large firms do to preserve their way of doing things, which is maintaining centralized control. The agenda above adds the clear alternative of smaller companies or profit centers opportunistically aggregating to act as large enterprises. That means that a future merged strategy must be devised with sensitivities to small and medium enterprises. Flexibility and tailorability must increase and complexity and cost must decrease from current practice.

The second is the complement. Implementing a new infrastructure with the level of cleverness outlined will change some fundamentals of how business is done. Some optimization must be considered at a higher level than the larger enterprise, beyond to national and societal interest. This is especially cogent as the initiating research will likely be funded by government agencies.

The final concern extends that notion in a structural way. Some technologies seem inherently abusable, while others seem self-correcting by design. For example, the Internet will likely be an inherently democratizing force despite the best efforts of large companies to "own" it or repressive governments to co-opt it. The workgroup recommends a project to study how to ensure that this new direction for merged EI/KM is inherently "good" and designed in a way that prevents capturing by inevitable corporate attempts to bend it one way or another for selfish purposes that compromise other elements of society.