Int. J. Human–Computer Studies (1997) **46,** 327–337

Impediments in the use of explicit ontologies for KBS development

Daniel E. O'Leary

University of Southern California 3660 Trousdale Parkway Los Angeles, CA 90089-1421, USA, email: oleary@rcf.usc.edu

This paper explores some of the impediments in the development of libraries of reusable ontologies. First, ontologies generated in multiple agent environments derive from political processes. As a result, it is impossible to choose an ontology that maximizes the utility of all agents in the process and the group. An "ontology impossibility theorem" is formulated and discussed. Second, ontologies are not necessarily stationary over relevant time periods, providing an impediment to librarying ontologies. Third, scaling up ontologies is a difficult matter. Fourth, since each ontology is the result of a political process, it is difficult to interface multiple ontologies.

1. Introduction

In their seminal paper, van Heijst, Schreiber and Wielinga (section 3.0) argue that:

> "... there are two impediments that hinder the development of libraries of reusable ontologies: the *hugeness problem* and the *interaction problem*. The hugeness problem concerns the overwhelming amount of knowledge in the world. This makes the construction of reusable domain ontologies a daunting exercise."

The interaction problem, discussed by Bylander and Chandrasekaran (1988), argued

> "Representing knowledge for the purpose of solving some problem is strongly affected by the nature of the problem and the inference strategy to be applied to the problem."

This paper argues that there are additional important impediments, including: existence of multiple agents in design, development or as computational agents; ontology stationarity; scaling up individual ontologies; interfacing independently developed ontologies; and formality of ontologies. Case studies are used extensively throughout the paper in order to facilitate discussion of some of the issues.

2. Multiple agents as an impediment

The existence of either multiple computational agents or multiple agents from which expertise is being gathered or multiple development agents can require the generation of an ontology different than any one agent would find optimal. This is because different agents have different incentive systems, different resource investments, resource constraints, different organizational settings and a wide range of other differing concerns. As a result, in multiple agent settings, I argue that the choice of an ontology is a political process that requires "negotiation" and that it is impossible to find an ontology that is preferred by each agent in the process.

1071-5819/97/2/30327 + 11$25.00/0/hc960094

2.1. CASE STUDY: DEVELOPING MULTIPLE AGENT SYSTEM ONTOLOGIES REQUIRES "NEGOTIATION"

What kinds of processes are actually used to develop ontologies for multiple agents settings and what is the nature of those processes? Typically, development of an ontology to be used by multiple agents is a political process, requiring substantial negotiation, etc. Ontology issues might be discussed on a pairwise basis between concerned agents. One example of that approach was given with a discussion of the development of the Palo Alto Collaborative Testbed (PACT) system by Cutkosky *et al.* (1993) who noted:

> "... the agents involved in any transaction must agree on a common ontology, which defines a standard vocabulary for describing time-varying behavior under each view of time that is needed. What went on behind the scenes in PACT, and is not represented in computational form at all, was a careful negotiation among system developers to devise the specific pairwise ontology that enabled their systems to cooperate. The developers met and emulated how their respective systems might discuss, say, the ramifications of increasing motor size. In this fashion, they ascertained and agreed upon what information had to be exchanged, and how it would be represented."

2.2. DEVELOPMENT OF ONTOLOGIES AS A POLITICAL PROCESS

Decision making has been categorized as a political process (e.g. Zeleny, 1982: p. 480). Since ontology developments are for systems to assist in decision making, development of ontologies is also a decision making process. In addition, ontology choice is a political process for other reasons, including the following.

First, the incentive system of the user (developer) of the ontology may actually be more important in the choice of an ontology than the particular ontology that is chosen. As a result, the key to system success may be less dependent on the ontology than other factors. Thus, tradeoffs might be made in the development of ontologies in order to accommodate other variables.

> "... the acceptability and explicability of proposed solutions must be very important dimensions in the decision maker's objectives. They are actually more important than the optimality or objective rationality of the proposal. A suboptimal solution successfully implemented may often be more desirable than an optimal solution which is actively resisted, sabotaged during implementation, or rejected altogether." (Zeleny, 1982: p. 481–482).

Second, there is individual interpretation by developers and users. That interpretation is wed to the "circumstances" in which they find themselves. Ontologies ultimately are generated in developer and or user contexts, not independent of them.

> "... there is no 'objective reality' or 'rationality'; these categories are context dependent, observer dependent, and subject to individual interpretation under given circumstances. The subject is inseparable from the circumstances of its deliberation, judgment and decision making. Jose Ortega y Gasset condensed this simple, important and often forgotten truth into a single sentence 'I am myself plus my circumstance'." (Zeleny, 1982: p. 481).

Researchers may want to focus on the search for objective and rational ontologies, however, this ultimately ignores the setting in which ontology usage must reside.

Because development of ontologies is dependent on political processes and negotiation, the resulting ontologies are likely to capture the result of those political processes and negotiations within them. Accordingly, this can "damage" much of the potential of storage for *reusability*.

Further, political processes are likely to influence the choice of the *initial use* of ontologies from the library. In particular, no choice might be made from the library, because of a need for a "unique" competitive edge or "non-optimal" choices could be made, because of alterntive views to those under which the ontology was developed (e.g. surgeon vs. chemotherapist). These behaviors have been seen in many of the reengineering studies done lately with firms "reinventing" solutions previously discovered.

> ". . . many analysts want to be the professional servants of objective knowledge and rationality rather than of a decision maker who is a human being, whose policies are formulated and implemented in and around an organization of human beings. Complex social relationships of human systems involve power, influence, negotiation, persuasion, and a large dose of organizational politics. The rationality and desirability of most decisions are not self evident outside the circumstances of a particular individual and organization. An analysis of political feasibility must be a part of the analyst decision maker interaction." (Zeleny, 1982: p. 482).

2.3. CASE STUDY: DEVELOPMENT OF MULTIPLE AGENT-BASED ONTOLOGIES ELICITS MANY AGENT DIFFERENCES, MAKING CHOICE OF ANY ONE ONTOLOGY, MAXIMIZING EVERYONE'S PREFERENCES, IMPOSSIBLE

Developing multiple agent-based ontologies is difficult. Generation of ontologies in multiple agent situations generally substantiates the existence of many alternative views. The choice of any one view, abstraction or position provides an advantage to one developer of computational agents (e.g. one ontology may be more natural or more easily developed), generally at the expense of another developer. Or similarly, one computational agent may be more efficient or effective under one ontology as compared to another. For example, with the development of PACT, Cutkosky *et al.* (1993) note:

> "Designing an integrated system by committee is always hard, and PACT was no exception. The most difficult task, by far, was agreeing on the ontological commitments that enable knowledge level communication among the systems. Designing a shared ontology is difficult because it must bridge differences in abstractions and views."

The process of generating agreement in the development of ontologies is likely to be similar to generating agreement in other human actor-based structures. Thus, negotiation and other devices are used to come to an agreement. Unfortunately, this often means that actors may have to trade away positions of optimality in order to obtain agreement with other agents for overall ontology development.

2.4. IMPOSSIBILITY THEOREM FOR ONTOLOGIES (ARROW'S IMPOSSIBILITY THEOREM APPLIES TO ONTOLOGIES)

As a result, a major concern with ontologies is that in virtually every "real life" set of circumstances, multiple feasible alternative ontologies might be used. Those multiple alternatives are likely to better meet the needs of different decision makers or designers. Thus, we need to address the issue of which ontology should be chosen to best meet everyone's needs. For example, which ontologies do we capture for or from a library (e.g. all of them?) and which are used or not used because there is another that better meets the needs of the user? Arrow (1963) addressed this issue, in general, with his "impossibility theorem". This section applies that theorem and its discussion to the problem of choosing ontologies.

Suppose that we want to compare alterntive ontologies developed for or by different agents. It is assumed that the ontologies are substantially different in terms of their impact on resource allocation (or else decision makers would not care which ontology was used) and that individual preferences with respect to the ontologies may differ.

Let Ω be the set of feasible ontologies. If individual i regards ontology j at least as desirable as k, then the expected utility associated with the use of j is greater than or equal to k, i.e. $E(U_i \mid n_i^*(j), j) \geq E(U_i \mid n_i^*(k), k)$, where $n_i^*(j)$ is the most preferred action for individual, i, given use of ontology j. That preference is expressed as $j >_{(i)} k$, for j and k in Ω, for individual i. Further, by making the appropriate rationality assumptions we can assure that the binary relation $>_{(i)}$ for each individual i is complete and transitive over the set of ontologies, Ω. In a similar manner, if j is at least as desirable as k, at the society level S, $j >_{(S)} k$. The optimality question then takes the form of specifying the relationship between the individual preference relations $i, i = 1, 2, \ldots, n$ and the social preference relation (S). Accordingly, we are interested in finding a collective choice rule $f(.)$. $>_{(S)} = f(>_{(1)}, \ldots, >_{(i)}, \ldots, >_{(n)})$ Thus, the problem is to find an ontology that meets the preferences of everyone. There are a number of collective choice rules. However, in his well-known theorem, Arrow found that the imposition of a set of apparently desirable conditions reduces the acceptable set to the null set. Informally, there is no solution that satisfies the needs of all individuals and the group.

2.4.1. Arrow's impossibility theorem for the ontology problem

(a) *Universal domain.* All logically possible orders of Ω are admissible. The domain of $f(.)$ must include all combinations of complete and transitive orders of Ω. "Weird" people are not disallowed, as long as they are completely and transitively "weird."

(b) *Pareto optimality.* If for any pair j and k in Ω, all individuals strictly prefer j to k, $f(.)$ must guarantee social preference for j over k. This is a necessary property, since not requiring it would amount to constructing a theory of choice among ontologies based on systematically denying individuals their preferred option.

(c) *Independence of irrelevant alternatives.* For any subset of Ω, any pair of sets of individual orderings that rank identically the ontologies in the specified subsets must have identical social choices of the ontologies in the subset. This condition eliminates interpersonal utility comparisons. For example it requires that the choice between some pair j and k in Ω, be determined solely by the individual's preferences for j and k, other alternatives are not allowed to influence the social choice between j and k.

(d) *Dictatorship.* There is no individual i such that $(>_{(S)}) = (>_{(i)})$ regardless of the other individuals' preferences.

The set of collective choice rules $f(.)$ that provides complete and transitive ranking of Ω and also satisfies the universal domain, Pareto optimality, independence of irrelevant alternatives and non-dictatorship conditions is null. That is, the four conditions are mutually inconsistent. For proof and extensive discussion of this theorem and related issues see Arrow (1963), Luce and Raiffa (1957) and Demski and Feltham (1976).

There are at least two implications of this result for multiple agent-based ontology development. First, given Arrow's structure this result precludes viewing the ontology library problem as optimization of some aggregate level utility, cost–benefit or social welfare function. Second, since these choices must be made and the conditions imposed by Arrow are inconsistent, we must violate at least one of them regardless of the manner in which ontologies are chosen or archived.

In addition, the impossibility theorem has some implications for the types of negotiations that can take place between agents in the development of ontologies. As noted by Kleindorfer, Kunreuther and Schoemaker (1993: p. 265), so called "agenda setting" can influence the choice of alternatives that are ultimately adopted if pairwise comparison and adoption is done. Suppose that $j >_{(S)} k$, $k >_{(S)} x$ and $x >_{(S)} j$, that is, on a pairwise basis, j is preferred to k and k is preferred to x, and x is preferred to j. If the agenda is set (j vs. k) vs. x then pairwise choice leads to x. However, if the agenda is set as (k vs. x) vs. j then the choice leads to j. Real world settings that employ pairwise analysis are subject to issues such as "agenda setting".

2.5. ALTERNATIVE APPROACHES AND IMPLICATIONS

As noted by Arrow (1963) if all preferences have a single peak, majority voting will still work. However, do we want to generate libraries based on majority voting (consensus)? O'Leary (1994) investigated some of the mathematical structure associated with consensus in multiple agent systems. Unfortunately, one concern with consensus is the potential lack of stability in ontologies. In those situations where consensus approaches 50%, changing coalitions and system requirements can disrupt consensus solutions. Even in situations where there are 60%/40% coalitions can result in relatively unstable situations since only a small percentage is necessary to change the consensus ontology.

The lack of stability can be driven by the existence of multiple conflicting paradigms to solve the same problem, e.g. surgery vs. chemotherapy, or other factors. The impossibility theorem indicates that any ontology developed by or for multiple agent settings will have a number of compromises that will ultimately limit applicability to other situations. Further, compromises built into these ontologies suggest that these ontologies will not be optimal in other situations. Although developers might try to capture the extent of those compromises through documentation, such efforts will be limited by their own views and political constraints. There is no objective reality in which to anchor such descriptions.

3. Ontology stationarity

Another impediment is ontology stationarity. Ontologies are not static, they need to evolve. Although evolution has been examined for some computer-based systems (e.g. Chen, McLeod & O'Leary, 1995), there has been limited research in evolving ontologies.

Ontology evolution likely is more of an impediment in some disciplines than other disciplines. For example, in the area of technology the changes have been rapid and substantial. To illustrate this rapid change, consider the table of contents in subsequent years of the "Technology Forecast", issued by Price Waterhouse consulting firm. Although these documents do not provide a formal ontology for computer-based systems, they do illustrate lack of ontology stationarity. The table of contents in version 4, published August of 1993, featured seven sections, including, Components, Hardware Infrastructure, Software Infrastructure, Architectures, Systems Engineering, Applications and Management. The table of contents in version 5, published September of 1994 also had seven sections, and only one of the names was the same, Components and Transmission Media, Computing and Communications Platforms, Computing and Communications Operating Systems, Computing and Communications Architectures, Enabling Technologies, Applications and Application of Technology. Further, the index changed substantially over that same period. I reviewed the change in the "Q's" because of the feasibility of the size, but similar change has occurred throughout. In the "Q's" there are eight items in each year, however, only two were the same. Rapid ontology change would greatly inhibit the ability to reuse (or even develop).

On the other hand, some ontologies might be very stable for periods of time (e.g. "Greek History"), because our knowledge about them changes little. In such cases ontology reuse would be a more likely option.

4. Scaling up ontologies

Another critical issue in ontology design is that of "scaling up." Ontologies are by definition structured in a precise manner, since they form the definitions, on which dialogues are based. As a result, ontologies are particularly vulnerable to scaling up issues, because of the relationship between complexity and definitional precision. As noted by Zadeh (1973):

> "As the complexity of a system increases, our ability to make precise and yet significant statements about its behavior diminishes until a threshold is reached beyond which precision and significance can no longer exist."

4.1. CASE STUDY: SYSTEM DESIGN OF AGENT INTERACTION

It is difficult to anticipate many of the issues that can be important factors in system design as the system is scaled up. Factors that are important early in the development cycle can lose importance as the problem is scaled up. For example, in the development of SHADE, Kuokka, McGuire, Weber, Tenenbaum, Gruber and Olsen (1993) noted

> "In the current year, the functionality of a matchmaking agent has been sketched out in detail. This led to the enhancement of the KQML specification with several new performatives vital to matchmaking. In particular, the "advertise" performative was added, which allows an agent to describe the KQML messages it is capable of supporting. However, during the definition of matchmaking, it was discovered that it is easier to specify an interest (which looks like a knowledge base query) than it is to characterize an agent's capabilities (which needs to account for all possible interest requests)."

Feasibility does not guarantee system design, or that it will scale up. In the case of SHADE, advertising was feasible if there were only a few properties, however, as the number of properties to advertise increased, the feasibility of design decreased, ultimately requiring a change in system design.

4.2. EXISTENCE OF A WORKING ONTOLOGY DOES NOT MEAN IT CAN BE SCALED UP

Most descriptions of ontology generation describe it as an iterative process, that can include development of working systems or sub-systems. Usage finds gaps which leads to additions in the ontology, etc., gradually growing the ontology, possibly even finding use in selected systems.

However, given a library ontology there is no guarantee that ontology will scale up or will integrate with other related ontologies. As noted by Hayes (1985: p. 3), the problem of scaling up is a deceptive one.

> "It is perilously easy to conclude that because one has a program which works (in some sense), its representation must be more or less correct (in some sense). Now this is true in some sense. But representation may be adequate to support a limited kind of inference, and completely unable to be extended to support a slightly more general kind of behavior. It may be wholly limited by scale factors, and therefore tell us nothing about realistically complicated worlds."

This suggests that if ontologies are libraried then there needs to be correlated information about them that captures issues, such as scaling up problems or extent to which issues of scaling up were investigated. Again, we can anticipate that political issues can interfere with capturing information that would facilitate scaling up.

4.3. OTHER SCALING UP ISSUES

There are a number of "scaling up issues" in systems and in ontology development. Unfortunately, that area of ontologies has only received limited attention in the literature.

Ontology scaling up issues might include "robustness", "overlap" and other concerns. Robustness can be defined as the ability of the basic design to stay the same as the ontology increases in size. In the above case study, a design based on "interest" is more robust than one based on "agent capabilities." Overlap could be characterized as the similarity between components of sub-ontologies. In some cases highly diverse subject areas may employ the same ontology for very different issues. There are many potential examples of conflicts in different ontologies. Consider generating an ontology for business computing. Generally, in the business community, the term "CPA" is used to represent "Certified Public Accountant." However, in the "Style Guide of the IEEE Computer Society", "CPA" is used to refer to the "Computer Press Association, San Francisco". There can be many other such conflicts.

5. Interfaces between ontologies

Although ontologies might be developed independently, libraries of ontologies would likely result in ontologies that would be used in conjunction with each other. Unfortunately, independently developed ontologies might not effectively integrate with other ontologies, for a wide range of reasons, ranging from similarity of vocabulary to conflicting views of the world (e.g. surgeon vs. chemotherapist).

5.1. CASE STUDY

In their discussion of PACT, Cutkosky *et al.* (1993) found that rather than a single ontology for different concepts that there were multiple, integrated concepts that needed to be addressed.

> "Agreements must be reached about concepts in the natural world, such as position and time, shape and behavior, sensors and motors. For each concept, agreement is required on many levels, ranging from what it means to how it is represented. For instance, how should two agents exchange information about the voltage on a wire (what units, what granularity of time?): how should manipulator dynamics be modeled (as simultaneous equatons or functions? in what coordinate frame?). The four systems comprising PACT used various coordinate systems and several distinct representations of time (e.g. discete events, points in continuous time, intervals of continuous time, piecewise approximations). These representations were chosen for valid task and context-dependent reasons. They cannot simply be replaced by one standard product model (e.g. with a single representaton of time)."

It would have been difficult to take an independently developed component for any of the four systems and use it directly in their system. Accordingly, based on this description (and others), interfaces between ontologies provides another impediment to effectively generate a library of ontologies for effective reuse.

5.2. INTERFACING MULTIPLE ONTOLOGIES

As noted above in Section 2, ontologies are the result of political processes, whose contexts are, to a certain extent, in the eye of the beholder. Accordingly, the task of interfacing multiple ontologies is even greater than the task of generating one. Often as noted in the case study, critical issues must be decided in concert with each other, limiting the ability to interface libraries of ontologies. Decisions about basic design issues made in one ontology can be in direct conflict with alternative design ontologies or other libraried ontologies. As an example, the term "CPA" discussed above indicates some of the potential impediments of interfacing multiple ontologies.

6. Formality

Development of ontologies can employ either informal or more formal processes. Since it takes equivocality to remove equivocality, uncertain and highly complex ontology problems may require solution using informal, "out of the box" thinking. As a result, some complex ontologies are likely to be developed using informal and unstructured processes.

6.1. CASE STUDY

In real world settings, ontologies often seem to be developed using relatively informal processes. For example, Cutkosky *et al.* (1993) found that:

> "...what PACT actually demonstrates is a mechanism for distributing reasoning, not a mechanism for automatically building and sharing a design model. The model sharing in PACT, as in other efforts, is still implicit—not given in a formal specification enforced in software. The ontology for PACT was documented informally in email messages among developers of the interacting tools."

6.2. ASHBY'S LAW OF REQUISITE VARIETY

Ashby's (1956: pp. 206–207) law of requisite variety is that "only variety can destroy variety". As a result, *informal* development processes actually might be necessary for complex ontology development. For a unique and complex ontology, formal approaches can limit the view of the world captured in the ontology. Less structued approaches might be used to capture creative ontologies.

Ashby's law implies that much of the richness of the environment needs to be captured and used. If there were a library of ontologies, this suggests that the ontology carry descriptor information regarding the nature of the development process, along with its political constraints and concerns, personnel capabilities, time, organizational setting information and a variety of other sets of information. Unfortunately, the very political processes for which information is needed can limit the capture of these kinds of information. Further, as noted above, much context information is in the eye of the beholder. As a result, much context information is likely to not be captured. In addition, any information captured is likely to be biased, with substantial situation specific meaning.

7. What kinds of situations are best for reusability of ontologies?

There have been a number of arguments presented in this paper that are critical of the concept of reusable ontologies. However, those criticisms are representative of only one side of the discussion. They basically assume that the ontology being constructed is highly equivocal and nonstationary, subject to political concerns and controversy, with possible paradigm differences represented in the ontology.

However, with that said, there are likely to be situations where the problems are not as equivocal or controversial or are stationary at points in time. In particular, development of ontologies in so-called "well-structured" problems are likely to be less costly than other problems. In these well-structured problems there may be well-established ontologies in use, and no (or few) current conflicts over paradigms. In these situations, the "impediments" discussed in this paper become less influential, and the likelihood of ability to reuse increases.

Unfortunately, in these situations there may be what I call the "paradoxical lack of reuse." As the impediments become less influential and the ability to reuse increases, the ease of development is also likely to increase so that there is less need for reusable ontologies. In those situations, developers can opt for development to meet specific needs, rather than reuse.

8. Summary and extensions

8.1. SUMMARY

This paper has addressed the issue of "impediments" to ontologies. Five different impediments were discussed. First, multiple agents make the choice of an ontology a political process. An outcome of the nature of that process is the notion of "impossibility." The impossibility theorem, when applied to ontologies, asserts that no ontology can be maximum for all individuals and the group, i.e. some individuals or the group will lose when an ontology is adopted over some other ontology. The resulting ontology will provide more benefits to some agents. Second, ontologies are not necessarily stationary. Third, scaling up ontologies is not a straight forward process. Issues such as density and breadth influence salability. Fourth, the notion of libraries of ontologies suggests relative independence between different ontologies. The interface of those ontologies is also an impediment, particularly since each of those ontologies was developed in the context of a political process with various compromises, further limiting the ability to treat ontologies as generic books on the library shelves. Fifth, Ashby's law of requisite variety indicates that it takes variety to destroy variety. As a result, less formal approaches that feature creative thinking and environments may be more likely to generate ontologies for complex situations. Finally, these impediments are likely to become less influential in highly structured and noncontroversial situations.

8.2. EXTENSION

Two other concerns, not addressed here, include "What incentives are there to library an ontology that had been developed?" and "What would be the quality of libraried ontologies and how could quality be assured?"

It is easy to imagine both positive and negative incentives to library ontologies. Corporations anxious to send competitors down what they think is the wrong course of development might readily agree to library unsuccessful ontologies or small scale ontologies, while not labeling them as such. Alternatively, corporate developed ontologies placed in the library might be only partially developed, or with critical portions missing, in hopes that the corporation will be contacted for consulting. Neither one of these stories is particularly appealing to the potential librarying of ontologies.

Further, any library of ontologies would need to generate a means of assuring or measuring the "quality" of an ontology, however that would be defined—a difficult issue in and of itself. A library of low quality contributions would provide little benefit and receive little use. As a result, there would be a need to develop some assurance as to the quality of the ontologies and measures of quality.

The author would like to acknowledge the comments of the referees on an earlier version of this paper.

References

ARROW, K. (1963). *Social Choice and Individual Values, Cowles Foundation Monograph 12*. New York: Wiley.

ASHBY, W. (1956). *An Introduction to Cybernetic.* New York: John Wiley and Sons.

BYLANDER, T. & CHANDRASEKARAN, B. (1988). Generic tasks in knowledge level reasoning: the right level of abstraction for knowledge acquisition. In B. GAINES & J. BOOSE, Eds. *Knowledge Acquisition for Knowledge-based Systems,* Vol. 1, pp. 65–77. London: Academic Press.

CUTKOSKY, M., ENGELMORE, R., KIKES, R., GENESERETH, M., GRUBER, T., MARK, W., TENENBAUM, J. & WEBER, J. (1993). PACT: An experiment in integrating concurrent engineering systems. http://www.eit.com/creations/papers/pact

CHEN, J., MCLEOD, D. & O'LEARY, D. (1995). Domain knowledge guided schema evolution. *Expert Systems with Applications,* **9,** 491–502.

DEMSKI, J. & FELTHAM, G. (1976). *Cost Determination: A Conceptual Approach.* Ames, IA: The Iowa State Press.

HAYES, P. (1985). The second naive physics manifesto, In J. HOBBS & B. MOORE, Eds. *Formal Theories of the Commonsense World,* pp. 1–36. Sussex: Ablex Publishing.

KLEINDORFER, P. KUNREUTHER, H. & SCHOEMAKER, P. (1993). *Decision Sciences: An Integration Perspective.* Cambridge: Cambridge University Press.

KUOKKA, D. R., MCGUIRE, J. WEBER, J. C. TENENBAUM, J. M. GRUBER, T. R. & OLSEN, G. R. (1993). SHADE: knowledge-based technology for the re-engineering problem. http://www-ksl.stanford.edu/knowledge-sharing/papers/shade-overview.html

VAN HEIJST, G., SCHREIBER, A. TH. & WIELINGA, B. J. (1997). Using explicit ontologies in KBS development. *International Journal of Human–Computer Studies,* **46,** 183–292.

LUCE, R. & RAIFFA, H. (1957). *Games and Decisions.* New York: Wiley.

O'LEARY, D. E. (1994). Models of consensus for multiple agent systems. In R. LOPEZ DE MANTARAS & D. POOLE, Eds. *Uncertainty in Artificial Intelligence.* pp. 447–453. San Francisco, CA: Morgan Kaufmann Publishers.

ZADEH, L. (1973). Outline of a new approach to the analysis of complex decision processes. In J. COCHRANE & M. ZELANY, Eds. *Multiple Criteria Decision Making,* pp. 686–725. South Carolina: University of South Carolina Press.

ZELENY, M. (1982). *Multiple Criteria Decison Making.* New York: McGraw-Hill.