Reuse and sharing of graphical belief network components.

Russell Almond *
StatSci division of MathSoft, Inc.

Jeffrey Bradshaw
Research and Technology, Boeing Computer Services

David Madigan
Dept. of Statistics, University of Washington

Abstract

A team of experts assemble a graphical belief network from many small pieces. This paper catalogs the types of knowledge that comprise a graphical belief network and proposes a way in which they can be stored in *libraries*. This promotes reuse of model components both within the team and between projects.

1 Introduction.

Graphical belief networks (Bayesian networks, influence diagrams, graphical belief models) have become a popular method for representing uncertain knowledge (Almond, 1990; Heckerman, 1991; Henrion, Breese, and Horvitz, 1991; Howard and Matheson, 1984; Pearl, 1988; Shachter, 1986; Shafer and Shenoy, 1988). Their attractiveness stems from the fact that they combine an easy to understand graphical notation with a rigorous computational model.

In our own work, we often encounter situations where modeling involves several people. Imagine, for example, that we are trying to model the system reliability of a complex machine. One engineer, the overall designer, might put together the overall structure of the model. A second engineer, a reliability expert, might determine what the failure states of the various components are and how they propagate. A third engineer, an expert in purchasing, might develop the models for individual component reliability, and so forth. The effectiveness of such a team will hinge on the quality of their communication. Of particular importance is the degree to which the components they are developing can be clearly described and easily shared among team members.

---

*This research was supported in part by the Graphical-Belief project, NASA SBIR Phase I Grant NAS 9-18669 and NIH SBIR Phase I Grant 1 R43 RR07749-01. Figures based on Graphical-Belief software

If the results of modeling efforts could be catalogued, others could build upon previous work, rather than starting from scratch. For example, the same component might appear many times in the completed system, and the model for its reliability can be reused. Or else the same combination of components (say a valve/actuator pair) may be used repeatedly, in each case replicating the same model fragment. Later, a different team may want to use the same components for a slightly different problem. To make this possible, they need a means to discover and exploit similarity among components in the graphical model.

The importance of effective sharing and reuse of model components has become even more apparent in light of recent developments in automated probabilistic and decision model construction systems (Bradshaw*et al.*, 1992a; Bradshaw *et al.*1991; Holtzman, 1989, Wellman, Breese and Goldman, 1991; Edgar, Puerta and Musen, 1992). In such systems, knowledge-based systems guide the configuration of situation-specific belief and decision models with components selected from an electronic library.

Although this paper concentrates on the problem of sharing and classifying knowledge *within* groups, there is an equally large an important problem of sharing knowledge between groups. Here the lack of standardization in terminology (just look at the number of synonyms for "graphical model") hinders the efforts of research groups from various schools to share examples and methods. Although this paper was derived from the design of a single system (GRAPHICAL-BELIEF, Almond, 1992b), a useful ontology for graphical belief networks will only be achieved by collaboration among many scientists involved in similar efforts.

To build and use libraries of resuable model components, we need at least three things: (1) a rigorous specification of the kinds of components such a library must contain—an *ontology* for graphical belief networks,—(2) a rich description of model components, and (3) a formalism to describe how components may be combined. Sections 2 and 3 briefly describe the first two elements of our approach.

2 An ontology for graphical belief networks.

A number of authors have argued the benefits for making conceptual commitments explicit in the form of ontologies (Bradshaw, *et al.*, 1992b; Gruber, 1991; Gruber, 1992b; Neches *et al.*, 1991; Skuce and Monarch, 1990). The term *ontology* is borrowed from the philosophical literature where it describes a theory of what exists. Such an account would typically include terms and definitions only for the very basic necessary categories of existence. However, the common usage of ontology in the knowledge sharing and reuse community is as a vocabulary of representational terms and their definitions at any level of generality. A knowledge-based system's *ontology* defines what exists for the program: in other words, what can be represented by it.

In this section, we discuss preliminary results in our efforts to define an ontology for graphical belief models. In Section 4 we describe some of the mechanisms we are exploring for exchanging these ontologies in computer interpretable form with other groups.

2.1 The central role of valuations.

The central theme of all graphical models is that the variables are represented by nodes connected by edges, which in some sense represent relationships between them. In Bayesian networks (Pearl, 1988), the edges are directed, and the relationships are conditional probability functions. In influence diagrams (Howard and Matheson, 1984), a value variable with utility to the decision maker is designated and compute so as to maximize is expected value. In graphical belief models (Almond, 1990, Dempster and Kong, 1988), the graph is a hypergraph and the "hyperedges" correspond directly to component belief functions.

Note that the graphical model provides a visual description of the structure of a mathematical model. Separation in the graph implies statistical independence (Pearl, 1988). The graph also implies a factorization of the model into distinct relationships (usually associated with edges). There is a close association between the independence conditions and the factorization, and under certain circumstances they are equivalent (Kong, 1988). As for the purposes of model construction the factorization is more important that the independence (indeed, for belief function models independence does not always imply factorization), we will leave aside issues of independence here.

To capture these diverse relationships in a single notation, Shenoy and Shafer(1990) introduce the term *valuation*. A valuation is defined over a set of variables called a *frame*. It maps sets of outcomes in that frame to values. We will define valuations and frames more formally below, but for now, we can consider them to be the generalization of the familiar probability relationship to include variations on the theme, such as utilities and belief functions.

Shenoy and Shafer define a set of operations, in particular, *combination* and *projections*, on valuations. Combination is the ability to combine two valuations defined over the same frame, for example by multiplying probability potentials together. Projection is the ability to change the frame (set of variables over which the valuation is defined) to a larger or smaller set; the most familiar example of this operation is the marginalization of probability distributions. Together with a theorem which allows limited commutativity of projection and combination, these operations can be used as the basis of local computation techniques.

Furthermore, the valuations define the graphical structure of the problem. Because for all graphical models there is one-to-one correspondence between the graphical structure and the factorization of the problem into valuations of appropriate types, the set of valuations must define the graphical model. Each valuation has an associated small fragment of graphical structure—the connection between the variables over which the valuation is defined. For example a valuation representing "If X and Y then (with probability ϕ) Z" would be represented by a directed edges from X and Y to Z or by undirected edges linking X, Y and Z. These fragments are assembled to form the full graphical model, although it is often more useful to run the association the other way, first defining the graphical structure and then defining the corresponding relationships (valuations).

Although it is easy to see how to assemble many valuations into graphs, one can also subdivide the valuations which describe the relationship structure into components. The frames of variables over which valuations are defined is worthy of further explanation, as are the variables themselves. Also, it is useful to partition the set of outcomes into groupings reflecting natural symmetry in the relationship. Finally, uncertainty about values can be represented by parameters with their own distributions. These substructures within a valuation are described below.

2.2 Variables and frames.

The place to start in defining a complex problem is with the *variables* or variables of the problem domain. Each variable has a set of *outcomes* which define the values it can take on. For example, a binary variable might be associated with the outcome set $\{0, 1\}$ or $\{\texttt{True}, \texttt{False}\}$.

There is a one-to-one correspondence between the variables of a problem domain and the nodes in the graphical model, although the node may have some additional information attached (such as the location, or the list of neighbors).

The domain of a valuation is a set of tuples over an ordered set of variables. We will refer to that domain as the *frame of discernment* or *frame*. There three different representations of the frame: (1) the *frame of variables* or list of variables, (2) the *frame vector* or tuple of outcome spaces associated with the variables, and (3) the *frame set* or set of possible tuples of outcomes, the cross product of the outcomes sets in the frame vector. For example, for three binary variables, the frame of variables might be (X, Y, Z), the frame vector would be $(\{0, 1\}, \{0, 1\}, \{0, 1\})$ and the frame set would be $\{(0, 0, 0), (0, 0, 1), \ldots, (1, 1, 1)\}$. The term *frame* is used when the distinction between the three views of the frame is unimportant.

If a valuation is defined over a given frame, one can think of marginalizing it to a smaller frame, or extending it to a larger frame. Extension can be defined very naturally for belief functions, but also can be done for probability potentials by replicating over the appropriate variables. Projecting a valuation onto a new frame is achieved by a combination of marginalization and extension.

2.3 Groups, groupings and partitions.

It is often useful to think of partitioning the frame set (the set of outcomes) into a number of groups. For example, one could partition a frame defined over two variables X and Y into those in which $X = Y$ and those in which that relationship does not hold. As another example, consider a system S with n identical components $C_1, \ldots, C_n$ in parallel. One might be interested in the probability of system failure, given that $0, 1, 2, \ldots, n$ of the components have failed. Thus there is a natural partition of the conditional part of the frame into sets representing k-out-of-n failures for $k = 0, \ldots, n$.

The *noisy-or* model (Pearl, 1988) is a simple form of this partitioning idea. Let $X_1, \ldots, X_n$ be a collection of binary input variables and Y be an output variable. Rather than specify a complete probability distribution for Y for each configuration of the n input variables, Pearl advocates assessing the conditional probability when two groups of input configurations: one in which at least one of the inputs has occurred and one in which none have occurred. Note that we do not need to restrict ourselves to "ors" and "ands"; any logical grouping of the attributes which we believe to be equivalent can be used.

We will formally define a *group* of outcomes as a set of outcomes about which we share a common pool of information. For example, the outcomes corresponding to the system with exactly k-out-of-n failures and where $X = Y$ both form a *group* of outcomes. Note that there is an implied frame associated with each group.

A set of groups over a particular frame is a *grouping*. Note that groups in a grouping need not be disjoint, nor need they span the entire frame set. Groupings are meant to reflect logical divisions of the valuation domain, and need not be true partitions.

An important subset of groupings is the *partition*. The groups in a *partition* must be pairwise disjoint and must span the entire outcome space. Partitions are particularly important, because probability valuations correspond to value assignments over partitions, as do simple utility valuations.

2.4 Formal definition of valuations.

We define a *valuation* as a mapping from a *grouping* over a particular *frame* to numeric *values*. In analogy with belief functions, the groups in the grouping are called *focal elements*. Perhaps the most obvious example of this is the mass function of a belief function. However, if the grouping over which the valuation is defined is a partition, then we can define a probability distribution as well. Recall the probability functions and belief functions are subclasses of valuations as they imply certain normalization constraints among the values.

An important subclass of valuations are those for which the grouping over which they are defined forms a partition. These valuations can be represented by an array of values, one element of that array corresponding to each tuple in the frame set. Such an array is called a *potential* which can be used to represent probabilities and utilities. The class of valuations which are defined over partitions and hence can be represented with potentials are called *simple valuations*. Valuations which are not simple must be represented by a more complex scheme, such as an association list between groups and values or an array indexed by subsets of the frame set (superpotentials).

Users of simple valuations, that is people who restrict their modelling effort to probabilities or probabilities and utilities, may not see the necessity of first defining a grouping, but may rather prefer to go directly to the array of values corresponding to the primitive tuples. This strategy, however, can quickly get out of hand for large problems, such a the system with many components. Pearl(1988) introduces the *noisy-or* model (see previous section) to address these situations. Identifying a grouping which reduces the effective domain of the valuation from the frame set to the set of groups could drastically reduce the number of values which must be specified. For example, in representing the knowledge "If X and Y hold then Z usually holds." we may only be concerned about assigning values to the two groupings which correspond to whether or not the rule holds. Furthermore, there is often uncertainty about the values (see Section 2.5), but usually not about the structure of the grouping.

Another important subclass of valuations are the conditional valuations. These valuations divide the frame variables into two groups, the conditions and the consequences. The value is thought to be a conditional value associated with the consequence group given the condition group, for example, a conditional probability of the consequence set given the condition set. Conditional valuations are usually represented by directed edges, where unconditional valuations are usually represented by undirected edges. Conditional valuations can come in both simple (maps to array) and complex varieties.

2.5 Parameters and laws.

Often there will be uncertainty about the numeric values of a valuation, that is an uncertainty about the strength of the relationship, but not the structure. Furthermore, the same numeric value, representing the same fragment of knowledge may appear in many valuations. In order to be able to trace and revise that knowledge, as well as express uncertainty about it, we must define an indirect pointer to the numeric values.

A *parameter* is just such a pointer to a numeric value. It is used as an alternative to the actual number in order to express uncertainty about the numeric value and promote re-use. As an example of both, the failure probability of a particular valve may be expressed as a parameter. Any place the valve is placed in the model, the same value for its failure probability should be used. If the valve is a new component, about which very little information is known, information about its failure rate may be uncertain or imprecise (or both). As test of the valve and other experience about it become available, the information will become more certain and precise, and the value of all parameters for that valve should be adjusted accordingly.

In order to express uncertainty about the value of parameters, parameters are allowed to have *laws*. These are probability distributions over the space of possible values for the parameters. Because some parameters are functionally linked (for example, the probability of A and not A), generally speaking the parameters will be dependent. In certain cases, it may be possible to make reasonable independence assumptions about some of the parameters.

Note that the term *law* is reserved for probability distributions over parameters, the term *valuation* is used for probability functions over variables. Parameters (in the statistical sense) of laws over parameters are called *hyperparameters*; this usage is consistent with the standard usage in Bayesian statistics. Distributions for hyperparameters are conceivable but hopefully unnecessary.

Spiegelhalter and Lauritzen(1990) use parameters to define a layered graphical model. The upper *quantitative* layer contains the distribution over the parameters, and the lower *qualitative* layer contains the graphical structure of the problem and the structure (groupings) within valuations. To answer questions in the *qualitative* layer, the best (average) values of the parameters are disseminated into the lower layer (in other words, each parameter is assigned a numeric value, the mean of its distribution). The now parameter free valuations are propagated through the graphical structure to answer questions. Finally, data from the consultation, can be used to update the distributions of the parameters in a Bayesian fashion. Almond(1990) uses a similar device, sampling from the distributions of the parameters to capture uncertainty about the parameters in the final estimation.

3 Model component libraries.

We now turn from the structure within valuations to the structure of many valuations. This is the *graph* of the graphical model. The strength of graphical modelling lies in the independence assumptions represented by separation in the graph, which in turn imply a factorization of the problem into component valuations. This in turn implies that an entire graphical model could be constructed by "dragging and dropping" a collection of valuations from a library into the model. This approach suggest how a design engineer might build a model from the work of a reliability engineer (the library designer); here the selection and placement is accomplished by a drag and drop interface. It is also a good model for how knowledge based model construction might work; here the selection and placement are accomplished by meta-rules which determine which knowledge is applicable when.

It is also possible to group the graphical structure into larger fragments. For example a subsystem might be a graph fragment which is repeated several times in the system. Modellers can obviously take advantage of such parallelism to reduce the modelling effort. Similarly, an intelligent program can take advantage of these symmetries to reduce computational cost.

Such a graph fragment, because of its portable nature, must be slightly different from a graph object. In particular, it will be necessary to duplicate the nodes (variables) in the graph fragment before adding it into the graph. Furthermore, there may be stub-nodes in the fragment which are meant to determine where the fragment will attach to other fragments already placed in the graph. Such stub-nodes will be resolved at model construction time.

Almond(1992a) has implemented a prototype library system that assists users in finding and reusing model components. Reusable model components in the library are packaged as *books*. A book consists of its *contents*—the associated model fragment;—it is labeled with a *title*—a brief description of the contents—and a set of *authors*—a list of contributors, allowing one to trace the sources of knowledge used in its construction—and wrapped in a *jacket*—a more thorough and detailed description of its function.

Users explore the contents of a library by means of two graphical interfaces: the *book editor* (Figure 1) and the *bookcase browser* (Figure 2). Figure 1 shows a *book editor* for a fragment of a graphical model. The display allows the user to examine the title, authors and jacket. The actual contents can be optionally examined for a more precise and detailed picture.

The *bookcase browser* presents a list of books by titles and allows the user to select or to open a *book editor* for any of them. If no appropriate book is found, a new one can be created from scratch or by editing an existing book. Ideally, the bookcase browser should be augmented by tools which would the list of books to be filtered via selection criteria (like electronic searching systems in libraries).

4 Conclusions and future directions.

The rate of progress in ontological issues will be largely determined by how well knowledge can be shared among those in the graphical belief networks community. Results of such analyses are currently shared very little, and where sharing takes place, it is usually either a) in the form of paper reports that take time to distribute and get outdated rapidly, or b) among scientists using some specific piece of not-widely-distributed or supported software. To facilitate development of ontologies it will be necessary for determine how

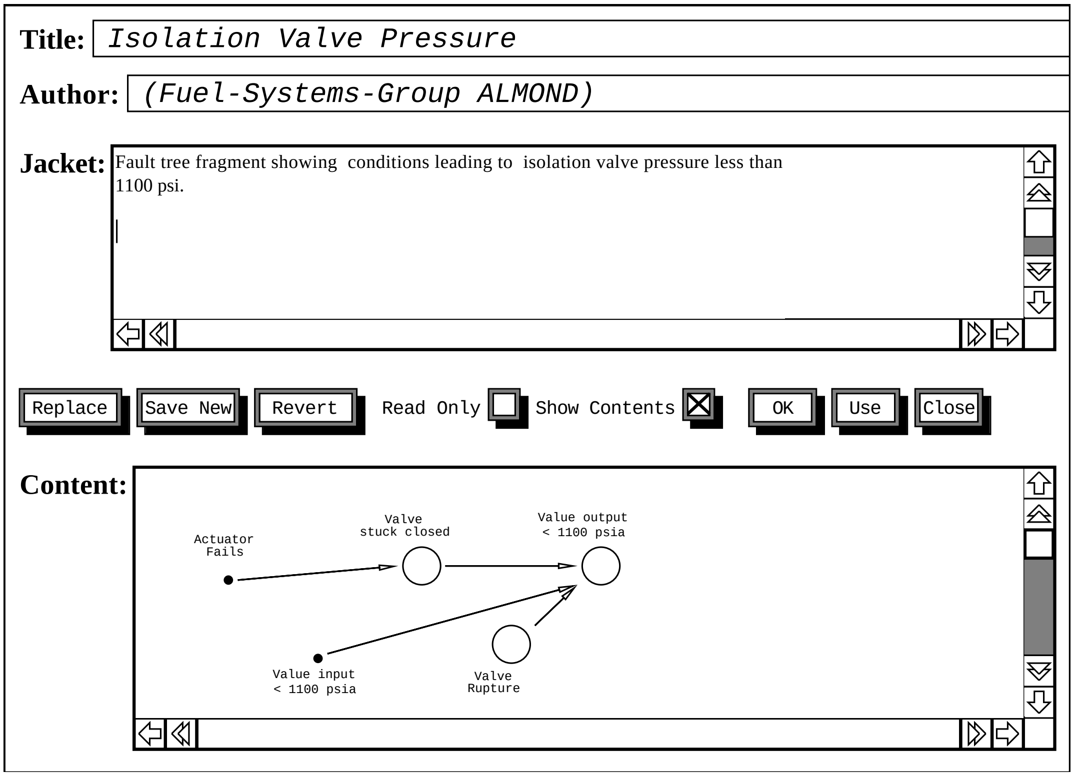

Figure 1: Book Editor for Graph Fragment
The book editor allows the user to inspect the title, author, jacket—detailed description of model fragment—and optionally the contents.

diverse software tools can exchange model fragments in computer-interpretable form.

Gruber's work on Ontolingua (Gruber, 1992a; Gruber, 1992b) currently provides the most promising mechanism for sharing ontologies between different tools and formalisms. Ontolingua extends the knowledge interchange format (KIF; Genesereth and Fikes, 1992) defined by the DARPA knowledge sharing effort with standard primitives for defining classes and relationships, and organizing knowledge in object-centered hierarchies with inheritance. Ontolingua facilitates the translation of KIF-level sentences to and from forms that can be used by various knowledge representation systems. Bradshaw *et al.*(1992b) and Lethbridge and Skuce (1992) describe the effort to integrate Ontolingua with other knowledge engineering tools.

References

[Almond 90] Almond, Russell G. (1990) *Fusion and Propagation in Graphical Belief Models: An Implementation and an Example.* Ph.D. dissertation and Harvard University, Department of Statistics, Technical Report S-130.

[Almond 92a] Almond, Russell G. (1992a) "Libraries, Books and Bookcases: A System for Organizing Re-Usable Knowledge." *Statistical Science Research Report 5,* StatSci, 1700 Westlake Ave, N, Seattle, WA.

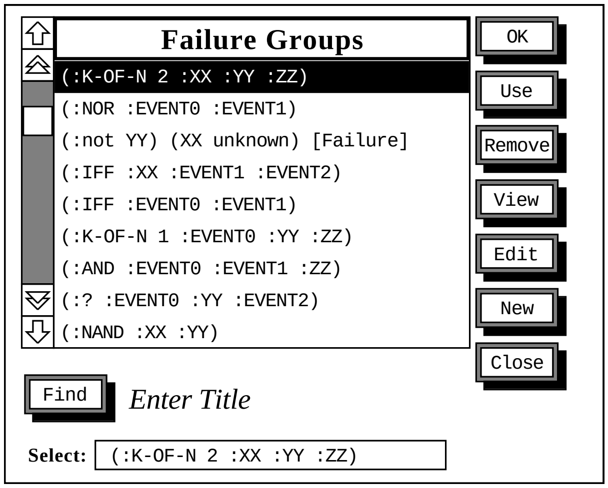

Figure 2: Bookcase Browser for Logical Groups

The bookcase browser allows the user to inspect a list of titles. Viewing the book shows an expanded description of the the contends and using the book employs the group in building a valuation.

[Almond 92b] Almond, Russell G.[1992b] "GRAPHICAL-BELIEF: Project Overview and Review." *Statistical Science Research Report 14* StatSci.

[Bradshaw 91] Bradshaw, Jeffery M., Covington, Stanley P., Russo, Peter J. and Boose, John H. (1991) "Knowledge acquisition techniques for decision analysis using `AXOTL` and `AQUINAS`." *Knowledge Acquisition* **3**, 49–77.

[Bradshaw 92a] Bradshaw, J.M., Chapman, C.R., Sullivan, K.M., Almond, R.G., Madigan, D., Zarley, D., Gavrin, J., Nims, J., and Bush, N. (1992a) "KS-3000: an application of DDUCKS to bone-marrow transplant patient support." Presented at the *Sixth Annual Florida AI Research Symposium (FLAIRS '93),* Ft. Lauderdale, FL.

[Bradshaw 92b] Bradshaw, J. M., Holm, P. D., Boose, J. H., Skuce, D., and Lethbridge, T. C. (1992b) "Sharable ontologies as a basis for communication and collaboration in conceptual modeling." *Proceedings of the Seventh Knowledge Acquisition for Knowledge-Based Systems Workshop*. Banff, Alberta, Canada.

[DempKong 88] Dempster, Arthur P. and Kong, Augustine (1988) "Uncertain Evidence and Artificial Analysis." *Journal of Statistical Planning and Inference,* **20**, 355-368.

[Edgar 92] Edgar, J.W., Puerta, A.R., Musen, M.A. (1992) "Graph-grammar assistance for modelling of decision." *Proceedings of the Seventh Banff Knowledge Acquisition for Knowledge-Based Systems Workshop,* pp 7:1–19. Banff, Alberta, Canada.

[Genesereth 92] Genesereth, M. R., and Fikes, R. (1992) "Knowledge Interchange Format Version 3.0 Reference Manual" *No. Logic Group Report, Logic-92-1*. Stanford University Department of Computer Science.

[Gruber 91] Gruber, T. R. (1991) "The role of common ontology in achieving sharable, reusable knowledge bases." In J. A. Allen, R. Fikes, and E. Sandewall (Eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference* (pp 601-602). Morgan Kaufmann, San Mateo, CA.

[Gruber 92a] Gruber, T. R. (1992a) "Ontolingua: A mechanism to support portable ontologies, Version 3.0." *No. Stanford Knowledge Systems Laboratory Technical Report KSL 91-66*. Stanford University Department of Computer Science.

[Gruber 92b] Gruber, T. R. (1992b) "A translation approach to portable ontology specifications." *Proceedings of the Seventh Knowledge Acquisition for Knowledge-Based Systems Workshop.* Banff, Alberta, Canada.

[Heckerman 91] Heckerman, D. (1991) *Probabilistic Similarity Networks,* ACM Press, New York.

[Henrion 91] Henrion, M., Breese, J.S., and Horvitz, E.J. (1991). "Decision analysis and expert systems." *AI Magazine,* 1991, 64-91.

[Holtzman 89] Holtzman, S. H. (1989) *Intelligent Decision Systems.* Addison-Wesley, Reading, Massachusetts.

[HowMath 81] Howard, R. A. and Matheson, J. E.[1981] "Influence diagrams." *Principles and Applications of Decision Analysis.* Strategic Decisions Group, Menlo Park, California.

[Kong 88] Kong, Augustine(1988) "A Belief Function Generalization of Gibbs Ensembles," *Joint Technical Report S-122* Harvard University and *No. 239* University of Chicago, Departments of Statistics.

[LethSkuce 92] Lethbridge, T. C., and Skuce, D. (1992) "Informality in knowledge exchange." Working Notes of the *AAAI-92 Knowledge Representation Aspects of Knowledge Acquisition Workshop*, 92.

[Neches 91] Neches, R., Fikes, R., Finin, T., Gruber, T., Patil, R., Senator, T., and Swartout, W. R. (1991) "Enabling technology for knowledge sharing." *AI Magazine,* pp 36-55.

[Pearl 88] Pearl, Judea [1988] *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, California.

[Shacter 86] Shachter, R. D.(1986) "Evaluating Influence Diagrams." *Operations Research,* **34**, 871–82.

[ShafShen 88] Shafer, Glenn and Shenoy, Prakash P.(1988) "Bayesian and Belief-Function Propagation." *School of Business Working Paper No. 192*. University of Kansas.

[ShenShaf 90] Shenoy, Prakash P. and Shafer, Glenn (1990) "Axioms for Probability and Belief-Function Propagation." in *Uncertainty in Artificial Intelligence,* **4**, 169-198.

[SkuceMon 90] Skuce, D., and Monarch, I. (1990) "Ontological issues in knowledge base design: Some problems and suggestions." *Proceedings of the Fifth Knowledge Acquisition for Knowledge-Based Systems Workshop,* Banff, Alberta, Canada.

[SpiegLaur 90] Spiegelhalter, David J. and Lauritzen, Steffen L. (1990) "Sequential Updating of Conditional Probabilities on Directed Graphical Structures." *Networks,* **20**, 579–605.

[Wellman 91] Wellman,M. Breese, J.S. and Goldman, R.P. (1991) “Working notes from the AAAI-91 Knowledge-Based Construction of Probabilistic and Decision Models.” Anahiem, CA.