Incremental Tradeoff Resolution in Qualitative Probabilistic Networks

Chao-Lin Liu and Michael P. Wellman
University of Michigan AI Laboratory
Ann Arbor, MI 48109-2110
{chaolin, wellman}@umich.edu

Abstract

Qualitative probabilistic reasoning in a Bayesian network often reveals tradeoffs: relationships that are ambiguous due to competing qualitative influences. We present two techniques that combine qualitative and numeric probabilistic reasoning to resolve such tradeoffs, inferring the qualitative relationship between nodes in a Bayesian network. The first approach incrementally marginalizes nodes in network, and the second incrementally refines the state spaces of random variables. Both provide systematic methods for tradeoff resolution at potentially lower computational cost than application of purely numeric methods.

Introduction

Researchers in uncertain reasoning regularly observe that to reach a desired conclusion (e.g., a decision), full precision in probabilistic relationships is rarely required, and that in many cases purely qualitative information (for some conception of "qualitative") is sufficient (Goldszmidt 1994). In consequence, the literature has admitted numerous schemes attempting to capture various forms of qualitative relationships (Wellman 1994), useful for various uncertain reasoning tasks. Unfortunately, we generally lack a robust mapping from tasks to the levels of precision required, and indeed, necessary precision is inevitably variable across problem instances. As long as some potential problem might require precision not captured in the qualitative scheme, the scheme is potentially inadequate for the associated task. Advocates of qualitative uncertain reasoning typically acknowledge this, and sometimes suggest that one can always revert to full numeric precision when necessary. But specifying a numerically precise probabilistic model as a fallback preempts any potential model-specification benefit of the qualitative scheme, and so it seems that one may as well use the precise model for everything.[1] This is perhaps the primary reason that qualitative methods have not seen much use in practical applications of uncertain reasoning to date.

The case for qualitative reasoning in contexts where numerically precise models are available must appeal to benefits other than specification, such as computation. Cases where qualitative properties justify computational shortcuts are of course commonplace (e.g., independence), though we do not usually consider this to be qualitative reasoning unless some inference is required to establish the qualitative property itself in order to exploit it. Since pure qualitative inference can often be substantially more efficient than its numeric counterpart (e.g., in methods based on infinitesimal probabilities (Goldszmidt & Pearl 1992) or ordinal relationships (Druzdzel & Henrion 1993)), it is worth exploring any opportunities to exploit qualitative methods even where some numeric information is required.

We have begun to investigate this possibility for the task of deriving the qualitative relationship (i.e., the sign of the probabilistic association, defined below) between a pair of variables in a Bayesian network. From an abstracted version of the network, where all local relationships are described qualitatively, we can derive the entailed sign between the variables of interest efficiently using propagation techniques. However, since the abstraction process discards information, the result may be qualitatively ambiguous even if the actual relationship entailed by the precise model is not.

In this paper, we report on two approaches that use qualitative reasoning to derive these relationships without necessarily resorting to solution of the complete problem at full precision, even in cases where purely qualitative reasoning would be ambiguous. Both approaches are incremental, in that they apply numeric reasoning to either subproblems or simplified versions of the original, to produce an intermediate model more likely to be qualitatively unambiguous.

The next section reviews the concepts of qualitative influences and tradeoffs in a network model. The third section explains the incremental marginalization approach, followed by the experimental results. We then discuss the state-space abstraction approach, and conclude with a brief comparison of our approaches with some others.

[1] If the qualitative formalism is a strict abstraction, then any conclusions produced by the precise model will agree at the qualitative level. Even in such cases, qualitative models may have benefits for explanation or justification (Henrion & Druzdzel 1991), as they can indicate something about the robustness of the conclusions (put another way, they can concisely convey broad classes of conclusions).

Qualitative Probabilistic Networks

Qualitative Influences

Qualitative probabilistic networks (QPNs) (Wellman 1990) are abstractions of Bayesian networks, with conditional probability tables summarized by the signs of qualitative relationships between variables. Each arc in the network is marked with a sign—positive (+), negative (−), or ambiguous (?)—denoting the sign of the qualitative probabilistic relationship between its terminal nodes.

The interpretation of such qualitative influences is based on *first-order stochastic dominance* (FSD) (Fishburn & Vickson 1978). Let $F(x)$ and $F'(x)$ denote two cumulative distribution functions (CDFs) of a random variable X. Then $F(x)\ FSD\ F'(x)$ holds if and only if (iff)

$$F(x) \leq F'(x) \text{ for all } x.$$

We say that one node positively influences another iff the latter's conditional distribution is increasing in the former, all else equal, in the sense of FSD.

Definition 1 ((Wellman 1990)) *Let $F(z|x_i\mathbf{y})$ be the cumulative distribution function of Z given $X = x_i$ and the rest of Z's predecessors $\mathbf{Y} = \mathbf{y}$. We say that node X positively influences node Z, denoted $S^+(X, Z)$, iff*

$$\forall x_i, x_j, \mathbf{y}. \quad x_i \leq x_j \Rightarrow F(z|x_j\mathbf{y})\ FSD\ F(z|x_i\mathbf{y}).$$

Analogously, we say that node X *negatively influences* node Z, denoted $S^-(X, Z)$, when we reverse the direction of the dominance relationship in Definition 1. The arc from X to Z in that case carries a negative sign. When the dominance relationship holds for both directions, we denote the situation by $S^0(X, Z)$. However, this entails conditional independence, and so we typically do not have a direct arc from X to Z in this case. When none of the preceding relationships between the two CDFs hold, we put a question mark on the arc, and denote such situations as $S^?(X, Z)$. We may apply the preceding definitions to binary nodes under the convention that **true** > **false**.

Inference and Tradeoffs

Given a QPN, we may infer the effects of the change in the value of one variable on the values of other variables of interest. The inference can be carried via graph reduction (Wellman 1990), or qualitative propagation techniques (Druzdzel & Henrion 1993).

If we are fortunate, we may acquire decisive answers from the inference algorithms. Often, however, the results of such qualitative reasoning are ambiguous. This might be because the relationship in question actually is ambiguous (i.e., nonmonotone or context-dependent), or due to loss of information in the abstraction process.

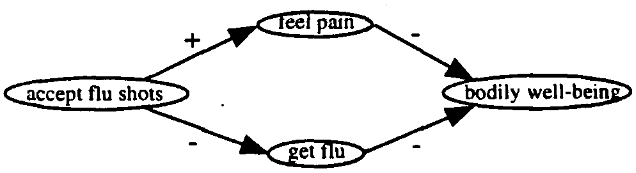

Figure 1: A simple case of qualitative ambiguity.

This can happen, for instance, when there are competing influential paths from the source node—whose value is tentatively modified—to the target node—whose change in value is of interest. While accept flu shots may decrease the probability of get flu, it also increases the probability and degree of feel pain. On the other hand, increasing either get flu or feel pain decreases overall bodily well-being, all else equal. As a result, qualitative reasoning about the problem of whether we should accept flu shots will yield only an ambiguous answer. The situation is illustrated by the QPN in Figure 1, where there is one positive path and one negative path from accept flu shots to bodily well-being. The combination of the two paths is qualitatively ambiguous. Worse, the ambiguity of this relationship would propagate within any network for which this pattern forms a subnetwork. For example, if this issue plays a role in a decision whether to go to a doctor, the result would be ambiguous regardless of the other variables involved.

Had we applied more precise probabilistic knowledge, such as a numerically specified Bayesian network, the result may have been decisive. Indeed, if accept flu shots and bodily well-being are binary, then a fully precise model is by necessity qualitatively unambiguous. However, performing all inference at the most precise level might squander some advantages of the qualitative approach. In the developments below, we consider some ways to apply numeric inference incrementally, to the point where qualitative reasoning can produce a decisive result.

Incremental Marginalization

Node Reduction

The idea of incremental marginalization is to reduce the network node-by-node until the result is qualitatively unambiguous. The basic step is Shachter's arc reversal operation.

Theorem 1 ((Shachter 1988)) *If there is an arc from node X to node Y in the given Bayesian network, and no other directed paths from X to Y, then we may transform the network to one with an arc from Y to X instead. In the new network, X and Y inherit each other's predecessors.*

Let $\mathbf{P}_X$, $\mathbf{P}_Y$, and $\mathbf{P}_{XY}$ respectively denote X's own predecessors, Y's own predecessors, and X and Y's common predecessors in the original network, and let

$\boldsymbol{P}_{Y'} = \boldsymbol{P}_Y - \{X\}$. *The new conditional probability distribution of Y and X are determined by the following:*

$$\Pr^{new}(y|\boldsymbol{P}_X\boldsymbol{P}_{Y'}\boldsymbol{P}_{XY}) = \sum_X \Pr^{old}(y|\boldsymbol{P}_Y\boldsymbol{P}_{XY})\Pr^{old}(x|\boldsymbol{P}_X\boldsymbol{P}_{XY}) \quad (1)$$

$$\Pr^{new}(x|y\boldsymbol{P}_X\boldsymbol{P}_{Y'}\boldsymbol{P}_{XY}) = \frac{\Pr^{old}(y|\boldsymbol{P}_Y\boldsymbol{P}_{XY})\Pr^{old}(x|\boldsymbol{P}_X\boldsymbol{P}_{XY})}{\Pr^{new}(y|\boldsymbol{P}_X\boldsymbol{P}_{Y'}\boldsymbol{P}_{XY})} \quad (2)$$

On reversing all the outgoing arcs from node X, the node becomes barren and can be removed from the network. The net effect of reversing arcs and removing barren nodes as described is equivalent to marginalizing node X from the network.

Marginalization and Qualitative Tradeoffs

Consider the QPN shown on the left-hand side of Figure 2. Since there exist both a positive path (through X) and a negative path (direct arc) from W to Z, the qualitative influence of W on Z is ambiguous. This local "?" would propagate throughout the network, necessarily ambiguating the relationship of any predecessor of W to any successor of Z.

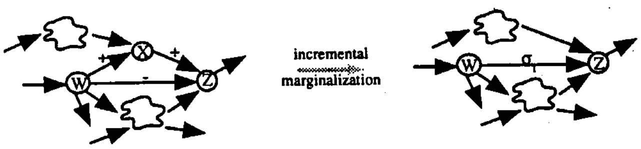

Figure 2: Marginalizing X potentially resolves the qualitative influence of W on Z.

Once we have detected the source of such a local ambiguity, we may attempt to resolve it by marginalizing node X. The new sign on the direct arc from W to Z can be determined by inspecting the new conditional probability table of Z, given by Equation (1). If we are fortunate, the qualitative sign will turn out to be decisive, in which case we have resolved the tradeoff.

This example illustrates the main idea of the incremental marginalization approach to resolving tradeoffs in QPNs. If we get an unambiguous answer from the reduced network after marginalizing a selected node, then there is no need to do further computation. If the answer is still ambiguous, we may select other nodes to marginalize. The iteration continues until a decisive answer is uncovered. We present the skeleton of the Incremental TradeOff Resolution algorithm below. The algorithm is designed to answer queries about the qualitative influence of a decision node on some target node, using some strategy for selecting the next node to reduce.

Algorithm 1
ITOR(decision,target,strategy)

1. *Remove nodes that are irrelevant to the query about decision's influence on target (Shachter 1988).*
2. *Attempt to answer the query via qualitative inference (Druzdzel & Henrion 1993).*
3. *If the answer to the query is decisive, exit; otherwise continue.*
4. *Select a node to reduce according to strategy, perform the node reduction, and calculate the qualitative abstractions of the transformed relationships. Return to Step 2.*

In general, we expect the incremental approach to improve performance over purely numeric inference. Since qualitative inference is quadratic whereas exact inference in Bayesian networks is exponential in the worst case, the qualitative inference steps do not add appreciably to computation time. On the other hand, when the intermediate results suffice to resolve the tradeoff, we save numeric computation over whatever part of the network is remaining.

Prioritizing Node Reduction Operations

For the evaluation of Bayesian networks, the objective of a node reduction strategy is to minimize computational cost to complete the evaluation. For our purpose, the aim is to minimize computational cost to the point that the qualitative tradeoff is resolved. The optimal strategies for the respective tasks will differ, in general. For example, a node that is very expensive to reduce at certain stage of the evaluation might be the best prospect for resolving the tradeoff.

We exploit intermediate information provided in qualitative belief propagation (Druzdzel & Henrion 1993) in determining which node to reduce next. If we can propagate the decisive qualitative influence of the decision node D all the way to the target node T, we will be able to answer the query. Otherwise, there must be a node X that has an indecisive relationship from D. Recall that we have pruned nodes irrelevant to the query, so any nodes that have indecisive relationship with D will eventually make the relationship between D and T indecisive. We have identified several conceivable strategies based on this observation, and have tried two of them thus far.

The first strategy is to reduce node X, as long as X is not the target node T. When X is actually T, we choose to reduce the node Y that passed the message to X changing its qualitative sign from a decisive one to "?". However, this Y cannot be D itself. If it is, then either (1) there are only two nodes remaining in the network, and there is no decisive answer to the query, or (2) there are other nodes, and we randomly pick among those adjacent to D or T.

The second strategy is similar to the first, except that we exchange the priority of reducing X and Y. We handle the situations where X and/or Y correspond to D and/or T in the same manner as in the first strategy.

These strategies have the advantage that finding the next node to reduce does not impose extra overhead in the ITOR algorithm. The selection is a by-product of the qualitative inference algorithm. However neither of these strategies (nor any that we know) is guaranteed to minimize the cost of resolving the tradeoff.

Experimental Results

We have developed an environment for testing the effectiveness of the algorithm using randomly generated network instances. The experiments are designed to examine how the connectivity of the network, the sizes of state spaces of the network, and the strategy for scheduling node reduction affect the performance of the algorithm.

Generating Random Networks

To carry out an experiment, we need two related networks: a QPN and its corresponding Bayesian network. The conditional probability distributions in the Bayesian network and the qualitative signs on the arcs in the QPN must agree with each other.

To create a random QPN with n nodes and l arcs, we first create a complete directed acyclic graph (DAG) with n nodes. We then remove arcs until the DAG contains only l arcs. Each arc is equally likely to be removed, under the constraint that the graph remains connected. After creating the network structure, we randomly assign qualitative signs (positive or negative) to the arcs.

We then build a Bayesian network that corresponds to the generated QPN, that is, respects its structure and qualitative signs. We select the cardinality of each node by sampling from a uniform distribution over the range $[2, MC]$, where MC denotes the maximum state-space cardinality. For nodes without predecessors, we assign prior probabilities by selecting parameters from a uniform distribution and then normalizing.

For a node X with predecessors P_X, the qualitative signs in the QPN dictate a partial ordering of the conditional probability distributions for various values of X, where the distributions are ordered based on the FSD relationship. To enforce this ordering, we identify the P_{X_i} that requires us to make the distribution $\Pr(X|P_{X_i})$ dominate distributions $\Pr(X|P_{X_j})$ for all other P_{X_j}. We assign the parameters $\Pr(X|P_{X_i})$ (as for priors) by sampling from a uniform distribution. We then assign the remaining distributions in stages, at each stage setting only those distributions dominated by the previously assigned distributions. We make these assignments using the same random procedure, but under the constraint that the resulting distribution must respect the qualitative signs given the previous assignments.

Results

In each experiment, we specify the number of nodes, the number of arcs, and maximum cardinality of state spaces for the randomly generated networks. In all experiments, we create networks with 10 nodes before pruning. We query the qualitative influence from the node1 to node10, and disregard the instances in which the influence of node1 on node10 is ambiguous after exact evaluation of the network.

Since the first step of the ITOR algorithm prunes nodes irrelevant to the query, the network actually used in inference is usually simpler than the original network. In Table 1, we record the *average* number of nodes and links after pruning, for each experiment. MC denotes maximum cardinality. All experiments reported used the first node selection strategy; results from the second strategy were virtually identical.

nodes	links	MC	R_1	R_2
8.0	14.2	2	0.697	0.722
8.0	14.4	3	0.730	0.754
9.2	26.1	2	0.846	0.869
9.4	26.8	3	0.855	0.874

Table 1: Experimental Results. Each experiment runs ITOR over 10000 random networks with decisive influence from node1 to node10.

We measure the performance of ITOR with two metrics. The first metric, R_1, is the ratio of the number of reduced nodes when the decisive answer is found to the number of nodes that would be reduced in exact numerical evaluation. The second metric, R_2, is the ratio of number of arc reversal operations already done when the solution is found to the number of arc reversal operations that would be carried out for exact numerical evaluation. The latter figure is based on an arbitrary node selection strategy (for reducing the remaining network after the tradeoff is resolved), however, and so would tend to be an overestimate. Table 1 reports averages for each metric. The savings due to incremental tradeoff resolution are $1 - R_1$ and $1 - R_2$, respectively, and so lower values of the metrics indicate better performance.

The results in Table 1 confirm the intuition that ITOR offers greater performance for sparsely connected networks and smaller state spaces. Further experimentation may lead us to more precise characterization of the expected savings achievable through incremental marginalization.

State-space Abstraction

Approximate evalu ation of Bayesian networks is a common strategy for time-critical problems. For qualitative inference, approximated distributions can be particularly useful if the qualitative relationships between nodes are preserved in the approximations.

In previous research (Wellman & Liu 1994), we have proposed iterative state-space abstraction (ISSA) as a technique for approximate evaluation of Bayesian networks. ISSA iteratively refines the state spaces of the nodes whose states are aggregated at the initial step of the algorithm. Approximated distributions get closer and closer to the true distributions in this process.

In the remainder of this section, we consider how we might determine the qualitative relationship of interest in a particular iteration of ISSA.

Controlled Approximations for Qualitative Inference

Consider the task of finding the qualitative influence from W to Z in the QPN shown in Figure 3, where the curly arcs represent paths between nodes. We assume that the overall influence is ambiguous, that is, $\sigma_1 \otimes \sigma_2 \neq \sigma_3$.

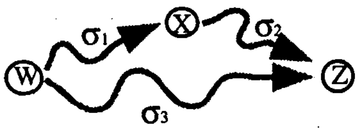

Figure 3: A QPN with ambiguous qualitative influence from W to Z.

In this situation, we may fruitfully abstract only the intermediate node X. To illustrate this, assume that we abstract W into two states, and that we find that the conditional, cumulative distribution function of Z given W being in the first abstract state dominates that of Z given W being in the second. With this observation, we still cannot infer with certainty that the unabstracted W negatively influences Z, although the approximated results suggest so. For a similar reason, results obtained by abstracting node Z cannot lead to definitive conclusions about the qualitative relationship in the original network.

Although aggregating the states of a node requires modification of conditional probability tables (CPTs) of the abstracted node whose states are aggregated and its successors, the effects of state-space abstraction can be analyzed as though only the successors' CPTs are modified, due to the way the CPT of the abstracted node are modified.

The qualitative influence from W to Z is determined by the signs of $F(z|w_i) - F(z|w_{i+1})$, for all i. Applying state-space abstraction, we compute approximations of these cumulative distribution functions (CDFs), that is, $F'(z|w_i)$ and $F'(z|w_{i+1})$. Let $F'(z|wx) = F(z|wx) + \delta(z,w,x)$ denote the new CDF of Z given W and X after we abstract X. The unconditional distribution, $F'(z|w_i)$, is given by

$$\begin{aligned} F'(z|w_i) &= \int_X F'(z|w_i x)\, dF(x|w_i) && (3) \\ &= \int_X (F(z|w_i x) + \delta(z,w_i,x))\, dF(x|w_i) \\ &= F(z|w_i) + \Delta_i, \end{aligned}$$

where $\Delta_i = \int_X \delta(z,w_i,x)\, dF(x|w_i)$.

Therefore, the difference $F'(z|w_i) - F'(z|w_{i+1})$ is equal to $(F(z|w_i) - F(z|w_{i+1})) + (\Delta_i - \Delta_{i+1})$. Notice that we cannot determine the sign of $(F(z|w_i) - F(z|w_{i+1}))$ purely based on the sign of $F'(z|w_i) - F'(z|w_{i+1})$. We need to know the sign of $\Delta_i - \Delta_{i+1}$ too. If we modify the CPT of Z such that $\Delta_i - \Delta_{i+1}$ is negative (or positive), then we may infer that $F(z|w_i) - F(z|w_{i+1})$ must be positive (or negative) when we find that $F'(z|w_i) - F'(z|w_{i+1})$ is positive (or negative).

The previous derivation reveals a way to apply state-space abstraction methods as well as other approximation methods to qualitative inference. Exact control of the monotonicity of Δ_i is essential. Without the control of the monotonicity, we may not infer the qualitative influence of interest based on the approximated cumulative distribution functions.

CPT Reassignment Policy

It is the CPT reassignment policy in the state-space abstraction that controls the monotonicity of Δ_i. The *average policy* applied in our previous work (Wellman & Liu 1994) does not ensure monotonicity of Δ_i, and thus we need to devise a new policy. The new policy is somewhat more complicated, and so would probably offset, to some extent, any computational savings from incremental approximation.

The task of controlling the monotonicity of Δ_i breaks into the tasks of controlling the monotonicity of $\delta(z,w,x)$ in multiple dimensions. For example, if σ_1="+" in Figure 3, then we want to make $\delta(z,w,x)$ an increasing function of x for all w and z, and an increasing function of w for all x and z. Given these, Δ_i will be an increasing function of w_i, as shown by Theorem 2 below. The first inequality in (4) follows from the fact that $\delta(z,w,x)$ is an increasing function of w, and we have the second inequality by application of Theorem 2 with (1) $F(x|w_{i+1})$ FSD $F(x|w_i)$, since σ_1="+", and (2) $\delta(z,w_{i+1},x)$ an increasing function of x.

$$\begin{aligned} \Delta_i &= \int_X \delta(z,w_i,x)\, dF(x|w_i) && (4) \\ &\leq \int_X \delta(z,w_{i+1},x) dF(x|w_i) \\ &\leq \int_X \delta(z,w_{i+1},x) dF(x|w_{i+1}) \\ &= \Delta_{i+1} \end{aligned}$$

Theorem 2 ((Fishburn & Vickson 1978)) *Let $F(x)$ and $F'(x)$ denote two cumulative distribution functions of a random variable X, and $g(x)$ a monotonically increasing function of X. Then, $F(x)$ FSD $F'(x)$ iff $\int g(x)dF(x) \geq \int g(x)dF'(x)$.*

Recall that $\delta(z,w,x)$ is the difference between the new and the old conditional probability of Z given W and X, i.e., $F'(z|wx) - F(z|wx)$. The difference is introduced when we reassign the conditional probability distributions associated with the abstracted nodes and their successors. Thus we need to know the exact conditional probability distribution $F(z|wx)$ for controlling the monotonicity of $\delta(z,w,x)$. When this distribution is available, we can apply the state-space abstraction technique. For instance, we may apply the idea to find the σ_1 on the direct arc from W to Z in the network shown in Figure 2, where $F(z|wx)$ is in the CPT of node Z.

Discussion

We have discussed the application of incremental marginalization and state-space abstraction methods to the qualitative tradeoff resolution task. The incremental marginalization approach iteratively reduces a node in the network. The state-space abstraction approach attempts to achieve the same goal by approximate evaluation of the Bayesian networks. Initial experiments with incremental marginalization suggest that nontrivial savings are possible, but definitive evaluation of both methods awaits further empirical and theoretical investigation.

The incremental marginalization approach bears some similarity to symbolic probabilistic inference, as in the *variable elimination* (VE) algorithm (Zhang & Poole 1996), in that we sum out one node from the Bayesian network at a time. The ITOR algorithm differs from VE in the strategy for elimination ordering, and of course in the stopping criterion.

Parsons and Dohnal discuss a semiqualitative approach for inference using Bayesian networks (Parsons & Dohnal 1993). The basic idea is similar to state-space abstraction. However, the center of their work is to design calculus for computing the the probability intervals of variables. Their methods may work even when the conditional probabilities in the Bayesian networks are not completely specified, but their methods cannot be applied to the qualitative tradeoff resolution task.

There are other approaches that make use of numerically specified knowledge in qualitative inference. Some of them are different from ours in that they do not assume the complete availability of the numerical information (Kuipers & Berleant 1988).

The incremental approaches we propose in this paper provide systematic ways to resolving qualitative tradeoffs at potentially lower computational cost than fully precise methods. Empirical results suggest that incremental tradeoff resolution can provide savings for some networks. How to use qualitative information to guide the scheduling of node reduction or state-space abstraction to achieve the best performance possible remain as open problems for future work.

References

Druzdzel, M. J., and Henrion, M. 1993. Efficient reasoning in qualitative probabilistic networks. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, 548–553. Washington, DC: Morgan Kaufmann.

Fishburn, P. C., and Vickson, R. G. 1978. Theroretical foundations of stochastic dominance. In Whitmore, G. A., and Findlay, M. C., eds., *Stochastic Dominance: An Approach to Decision making Under Risk.* Lexingtom, MA: D. C. Heath and Company.

Goldszmidt, M., and Pearl, J. 1992. Reasoning with qualitative probabilities can be tractable. In *Proceedings of the Eighth Conference on Uncertainty in Artificial Intelligence*, 112–120.

Goldszmidt, M. 1994. Research issues in qualitative and abstract probability. *AI Magazine* 15(4):63–65.

Henrion, M., and Druzdzel, M. J. 1991. Qualitative propagation and scenario-based approaches to explanation of probabilistic reasoning. In Bonissone, P.; Henrion, M.; Kanal, L.; and Lemmer, J., eds., *Uncertainty in Artificial Intelligence 6*, 17–32. North Holland: Elsevier.

Kuipers, B., and Berleant, D. 1988. Using incomplete quantitative knowledge in qualitative reasoning. In *Proceedings of the National Conference on Artificial Intelligence*, 324–329. AAAI.

Parsons, S., and Dohnal, M. 1993. A semiqualitative approach to reasoning in probabilistic networks. *Applied Artificial Intelligence* 7:223–235.

Shachter, R. D. 1988. Probabilistic inference and influence diagrams. *Operations Research* 36(4):589–604.

Wellman, M. P., and Liu, C.-L. 1994. State-space abstraction for anytime evaluation of probabilistic networks. In *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence*, 567–574. Seattle, WA: Morgan Kaufmann.

Wellman, M. P. 1990. *Fourmulation of Tradeoffs in Planning Under Uncertainty.* Pitman.

Wellman, M. P. 1994. Some varieties of qualitative probability. In *Proceedings of the Fifth International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, 437–442.

Zhang, N. L., and Poole, D. 1996. Exploiting causal independence in Bayesian network inference. *Journal of Artificial Inelligence Research* 5:301–328.