

Efficient Reasoning in Qualitative Probabilistic Networks*

Marek J. Druzdzel

Carnegie Mellon University
Department of Engineering
and Public Policy
Pittsburgh, PA 15213
marek+@cmu.edu

Max Henrion

Rockwell International Science Center
Palo Alto Laboratory
444 High Street, Suite 400
Palo Alto, CA 94301
henrion@camis.stanford.edu

Abstract

Qualitative Probabilistic Networks (QPNs) are an abstraction of Bayesian belief networks replacing numerical relations by qualitative influences and synergies [Wellman, 1990b]. To reason in a QPN is to find the effect of new evidence on each node in terms of the sign of the change in belief (increase or decrease). We introduce a polynomial time algorithm for reasoning in QPNs, based on local sign propagation. It extends our previous scheme from singly connected to general multiply connected networks. Unlike existing graph-reduction algorithms, it preserves the network structure and determines the effect of evidence on all nodes in the network. This aids meta-level reasoning about the model and automatic generation of intuitive explanations of probabilistic reasoning.

Introduction

A formal representation should not use more specificity than needed to support the reasoning required of it. The appropriate degree of specificity or numerical precision will vary depending on what kind of knowledge is available and what questions users want it to address. Qualitative Probabilistic Networks (QPNs) can replace or supplement quantitative Bayesian belief networks where numerical probabilities are either not available or not necessary for the questions of interest. QPNs have been found valuable for such tasks as planning under uncertainty [Wellman, 1990b] and for explanation of probabilistic reasoning [Henrion and Druzdzel, 1991]. Like other qualitative schemes, QPNs are weaker than their quantitative counterparts, but they can provide more robust results with much less effort.

QPNs are in essence a qualitative abstraction of Bayesian belief networks and influence diagrams. A QPN requires specification of the graphical belief network, expressing probabilistic dependence and independence relations. In addition, it requires specifica-

tion of the signs of influences and synergies among variables. A proposition A has a positive influence on a proposition B , if observing A to be true makes B more probable. Variable A is positively *synergistic* with variable B with respect to a third variable C , if the joint effect of A and B on the probability of C is greater than the sum of their individual effects. QPNs generalize straightforwardly to multivalued and continuous variables. An expert may express his or her uncertain knowledge of a domain directly in the form of a QPN. This requires significantly less effort than a full numerical specification of a belief network. Alternatively, if we already possess a numerical belief network, then it is straightforward to identify the qualitative relations inherent in it.

In previous work [Henrion and Druzdzel, 1991] we introduced an approach called *qualitative belief propagation*, analogous to message-passing algorithms for quantitative belief networks (e.g., [Kim and Pearl, 1983]), which traces the effect of an observation e on successive variables through a belief network to the target t . Every node on the path from e to t is given a label that characterizes the sign of impact. This was further developed, with particular emphasis on inter-causal reasoning in [Wellman and Henrion, 1991]. This approach differs from the graph reduction-based approach [Wellman, 1990b] in that it preserves the original structure of the network. The graph-reduction scheme performs inference by successively reducing the network to obtain the qualitative relation directly between e and t . There are usually several node reduction and arc reversal sequences possible at any step of the algorithm. As some of these sequences may lead to ambiguous signs, the algorithm needs to determine which sequences are optimal with respect to maximum specificity of the result. The computational complexity of this task is unknown [Wellman, 1990c]. The reason why different sequences of operators lead to different specificity of the results, is that, although the operation of arc reversal preserves the numerical properties of the network [Shachter, 1986], it leads to loss of the explicit qualitative graphical information about conditional independencies.

*This research was supported in part by the Rockwell International Science Center.

Earlier work on qualitative belief propagation applied only to a restricted class of singly connected belief networks (polytrees). The main contribution of this paper is to extend qualitative belief propagation to arbitrary networks and to present a complete belief propagation algorithm for QPNs.

All random variables that we deal with in this paper are multiple-valued, discrete variables, such as those represented by nodes of a Bayesian belief network. We make this assumption for convenience in mathematical derivations and proofs. Lower case letters (e.g., x) will stand for random variables, indexed lower-case letters (e.g., x_i) will denote their outcomes. In case of binary random variables, the two outcomes will be denoted by upper case (e.g., the two outcomes of a variable c will be denoted by C and \bar{C}). Outcomes of random variables are ordered from the highest to the lowest value. And so, for a random variable a , $\forall i < j$ [$a_i \geq a_j$]. For binary variables $C > \bar{C}$, or **true** > **false**.

Qualitative Probabilistic Networks

Formally, a QPN is a pair $G = (V, Q)$, where V is a set of variables or nodes in the graph and Q is a set of qualitative relations among the variables [Wellman, 1990b]. There are two types of qualitative relations in Q : *qualitative influences* and *additive synergies*. We reproduce their definitions from [Wellman and Henrion, 1991]. The qualitative influences define the sign of direct influence between two variables and correspond to an arc in a belief network.

Definition 1 (qualitative influence) We say that a positively influences c , written $S^+(a, c)$, iff for all values $a_1 > a_2$, c_0 , and x , which is the set of all of c 's predecessors other than a ,

$$\Pr(c \geq c_0 | a_1 x) \geq \Pr(c \geq c_0 | a_2 x).$$

This definition expresses the fact that increasing the value of a , makes higher values of c more probable. *Negative qualitative influence*, S^- , and *zero qualitative influence*, S^0 , are defined analogously by substituting \geq by \leq and $=$ respectively.

Definition 2 (additive synergy) Variables a and b exhibit positive additive synergy with respect to variable c , written $Y^+(\{a, b\}, c)$, if for all $a_1 > a_2$, $b_1 > b_2$, c_0 , and x , which is the set of all of c 's predecessors other than a and b ,

$$\begin{aligned} \Pr(c \geq c_0 | a_1 b_1 x) + \Pr(c \geq c_0 | a_2 b_2 x) \\ \geq \Pr(c \geq c_0 | a_1 b_2 x) + \Pr(c \geq c_0 | a_2 b_1 x). \end{aligned}$$

The additive synergy is used with respect to two causes and a common effect. It captures the property that the joint influence of the two causes is greater than sum of individual effects. *Negative additive synergy*, Y^- , and *zero additive synergy*, Y^0 , are defined analogously by substituting \geq by \leq and $=$ respectively.

If a qualitative property is not '+', '-', or '0', it is by default '?' ($S^?$ and $Y^?$ respectively). As all the definitions are not-strict, both '+' and '-' are consistent with '0'; for the same reason '?' is consistent with '0', '+', and '-'. Any qualitative property that can be described by a '0' can be also described by '+', '-', or '?'. Obviously, when specifying a network and doing any kind of reasoning, one prefers stronger conclusions to weaker ones and this is captured by the canonical order of signs: '0' is preferred to '+' and '-', and all three are preferred to '?' [Wellman, 1990b].

Qualitative properties can be elicited directly from a domain expert along with the graphical network using their common-sense interpretation or, alternatively, extracted from the numerical specification of a quantitative belief network using the definitions given above. It is worth noting that most popular probabilistic interactions exhibit unambiguous qualitative properties. It can be easily proven, for example, that bi-valued noisy-OR gates have always positive influences (S^+) and negative additive synergies (Y^-). Linear (Gaussian) models yield well defined qualitative influences (i.e., non-?) and zero additive synergies (Y^0).

Figure 1 shows an example of a QPN. This network is a small fragment of a larger belief network proposed for modeling an Orbital Maneuvering System (OMS) propulsion engine of the Space Shuttle [Horvitz *et al.*, 1992]. The OMS engine's fragment captured by the




Figure 1: An example of a qualitative probabilistic network.

network consists of two liquid gas tanks: an oxidizer tank and a helium tank. Helium is used to pressurize the oxidizer, necessary for expelling the oxidizer into the combustion subsystem. A potential temperature problem in the neighborhood of the two tanks (*HeOx Temp*) can be discovered by a probe (*HeOx Temp Probe*) built into the valves between the tanks. An increased temperature in the neighborhood of the two tanks can increase the temperature in the oxidizer tank (*High Ox Temp*) and this in turn can cause a leak in the oxidizer tank (*Ox Tank Leak*). A leak may lead

to a decreased pressure in the tank. A problem with the valve between the two tanks (*HeOx Valve Problem*) can also be a cause of a decreased pressure in the oxidizer tank. The pressure in the oxidizer tank is measured by a pressure gauge (*Ox Pressure Probe*). Of all the variables in this network, only the values of the two probes (*HeOx Temp Probe* and *Ox Pressure Probe*) are directly observable. The others must be inferred.

Links in a QPN are labeled by signs of the qualitative influences S^{δ} , each pair of links coming into a node is described by the signs of the synergy between them. Note that all these relations are uncertain. An increased *HeOx Temp* will usually lead to an increased reading from the *HeOx Temp Probe*, but not always — the probe may fail. But the fact that increased *HeOx Temp* makes an increased *HeOx Temp Probe* more probable is denoted by a positive influence S^+ .

Qualitative Intercausal Reasoning

In earlier work [Henrion and Druzdzel, 1991] we proposed a third qualitative property, called *product synergy*, which was further studied by Wellman and Henrion [1993]. Product synergy captures the sign of conditional dependence between immediate predecessors of a node that has been observed or has evidential support. The most common pattern of reasoning captured by product synergy is known as *explaining away*. For example, suppose my observed sneezing could be caused by an incipient cold or by a cat allergy. Subsequently observing a cat would explain away the sneezing, and so reduce my fear that I was getting a cold. This is a consequence of the negative product synergy between cold and allergy on sneezing.

A key desired feature of any qualitative property between two variables in a network is that this is invariant to the probability distribution of other neighboring nodes. This invariance allows for drawing conclusions that are valid regardless of the numerical values of probability distributions of the neighboring variables. Previous work on intercausal reasoning concentrated on situations where all irrelevant ancestors of the common effect were assumed to be instantiated. To be able to perform intercausal reasoning in arbitrary belief networks, we extended the definition of product synergy to accommodate this case. We reproduce here only the most important results. The complete derivations and proofs are reported in [Druzdzel and Henrion, 1993].

Definition 3 (half positive semi-definiteness) A square $n \times n$ matrix \mathbf{M} is called half positive semi-definite (half negative semi-definite) if for any non-negative vector \mathbf{x} consisting of n elements $\mathbf{x}^T \mathbf{M} \mathbf{x} \geq 0$ ($\mathbf{x}^T \mathbf{M} \mathbf{x} \leq 0$).

The following theorem states the sufficient condition for a matrix to be half positive semi-definite. Necessity of this condition in general remains a conjecture, although we have shown that it is necessary for 2×2 and 3×3 matrices.

Theorem 1 (half positive semi-definiteness) A sufficient condition for half positive semi-definiteness of a matrix is that it is a sum of a positive semi-definite and a non-negative matrix.

Definition 4 (product synergy) Let a , b , and x be predecessors of c in a QPN. Let n_x denote the number of possible values of x . Variables a and b exhibit negative product synergy with respect to a particular value c_0 of c , regardless of the distribution of x , written $X^-(\{a, b\}, c_0)$, if for all $a_1 > a_2$ and for all $b_1 > b_2$, a square $n_x \times n_x$ matrix \mathbf{D} with elements

$$D_{ij} = Pr(c_0|a_1 b_1 x_i) Pr(c_0|a_2 b_2 x_j) - Pr(c_0|a_2 b_1 x_i) Pr(c_0|a_1 b_2 x_j).$$

is half negative semi-definite. If \mathbf{D} is half positive semi-definite, a and b exhibit positive product synergy written as $X^+(\{a, b\}, c_0)$. If \mathbf{D} is a zero matrix, a and b exhibit zero product synergy written as $X^0(\{a, b\}, c_0)$.

Note that product synergy is defined with respect to each outcome of the common effect c . There are, therefore, as many product synergies as there are outcomes in c . For a binary variable c , there are two product synergies, one for C and one for \bar{C} .

Although the definition of product synergy may seem rather unintuitive, it simply captures formally the sign of conditional dependence between pairs of direct ancestors of a node, given that the node has been observed. This sign can be easily elicited directly from the expert. It is worth noting that most popular probabilistic interactions, the bi-valued noisy-OR gates exhibit negative product synergy (X^-) for the common effect observed to be present and zero product synergy (X^0) for the common effect observed to be absent, for all pairs of their direct ancestors.

Intercausal reasoning is an important component of the qualitative belief propagation allowing for sign propagation in cases where some of the network variables are instantiated. The following theorem describes the sign of intercausal reasoning for direct evidential support for the common effect node (see Figure 2). We prove an analogue theorem for indirect evidential support in [Druzdzel and Henrion, 1993].




Figure 2: Intercausal reasoning between a and b with c observed.

Theorem 2 (intercausal reasoning) Let a , b , and x be direct predecessors of c such that a and b are conditionally independent (see Figure 2). A sufficient and

necessary condition for $S^-(a, b)$ on observation of c_0 is negative product synergy, $X^-(\{a, b\}, c_0)$.

Qualitative Belief Propagation

In singly connected networks, evidence flows from the observed variables outwards to all remaining nodes of the network, and never in the opposite direction. In the presence of multiple connections, this paradigm becomes problematic, as the evidence coming into a node can arrive from multiple directions. For any link that is part of a clique of nodes, it becomes impossible to determine in which direction the evidence flows. Numerical belief propagation through multiply connected graphs, encounters the problem of a possibly infinite sequence of local belief propagation and an unstable equilibrium that does not necessarily correspond to the new probabilistic state of the network [Pearl, 1988, pages 195–223]. Algorithms adapting the belief propagation paradigm to multiply connected belief networks treat loops in the underlying graph separately and essentially reduce the graph to a singly connected one.

It turns out that the qualitative properties of the QPNs allow for an interesting view of qualitative belief propagation. The qualitative influences and synergies are defined in such a way that they are independent of any other nodes interacting with the nodes that they describe. This allows the propagation of belief from a node e to a node n to disregard all such nodes and effectively decompose the flow of evidence from e to n into distinct trails from e to n . On each of these trails, belief flows in only one direction, from e to n , and never in the opposite direction, exactly as it does in singly connected networks.

The belief propagation approach requires that qualitative changes be propagated in both directions. Product synergy is symmetric with respect to the predecessor nodes, so $X^\delta(\{a, b\}, c_0)$ implies $X^\delta(\{b, a\}, c_0)$. The following theorem shows that a qualitative influence between any two nodes in a network is also symmetric.

Theorem 3 (symmetry) $S^\delta(a, b)$ implies $S^\delta(b, a)$.

The theorem follows from the monotone likelihood property [Milgrom, 1981]. It shows merely that the sign of influence is symmetric. The magnitude of the influence of a variable a on a variable b can be arbitrarily different from the magnitude of the influence of b on a .

Of the definitions below, trail, head-to-head node, and active trail are based on [Geiger *et al.*, 1990].

Definition 5 (trail in undirected graph) A trail in an undirected graph is an alternating sequence of nodes and links of the graph such that every link joins the nodes immediately preceding it and following it.

Definition 6 (trail) A trail in a directed acyclic graph is an alternating sequence of links and nodes of the graph that form a trail in the underlying undirected graph.

Definition 7 (head-to-head node) A node c is called a head-to-head node with respect to a trail t if there are two consecutive edges $a \rightarrow c$ and $c \leftarrow b$ on t .

Definition 8 (minimal trail) A trail connecting a and b in which no node appears more than once is called a minimal trail between a and b .

Definition 9 (active trail) A trail t connecting nodes a and b is said to be active given a set of nodes L if (1) every head-to-head node with respect to t either is in L or has a descendant in L and (2) every other node on t is outside L .

Definition 10 (evidential trail) A minimal active trail between an evidence node e and a node n is called an evidential trail from e to n .

Definition 11 (intercausal link) Let a and b be direct ancestors of a head-to-head node t . An intercausal link exists between a and b , if t is in or has a descendant in the set of evidence nodes. The sign of the intercausal link is the sign of the intercausal influence between a and b determined by the product synergy.

Qualitative signs combine by means of sign multiplication and sign addition operators, defined in Table 1.

\otimes	+	-	0	?	\oplus	+	-	0	?
+	+	-	0	?	+	+	?	+	?
-	-	+	0	?	-	?	-	-	?
0	0	0	0	0	0	+	-	0	?
?	?	?	0	?	?	?	?	?	?

Table 1: Sign multiplication (\otimes) and sign addition (\oplus) operators [Wellman, 1990b]

Definition 12 (sign of a trail) The sign of a trail t is the sign product of signs of all direct and intercausal links on t .

Theorem 4 (evidential trails) The qualitative influence of a node e on a node n is equal to the sign sum of the signs of all evidential trails from e to n .

Proof: (outline) We demonstrate for each of the three qualitative properties that they are insensitive to the probability distribution of neighboring nodes. The presence of another, parallel trail through which evidence might flow changes only the probability distribution of the neighboring nodes, and this does not impact the qualitative properties of other nodes and paths.

This implies that none of the straightforward propagation rules for the singly connected networks will be invalidated by the presence of multiple trails. Qualitative change in belief in a node n given a single evidence node e can be viewed as a sum of changes through individual evidential trails from e to n . It will be well determined only if the signs of these paths are consistent (i.e., the sign sum is not '?'). \square

The algorithm for qualitative belief propagation (Figure 3) is based on local message passing. The goal is to determine a sign for each node denoting the direction of change in belief for that node given new evidence for an observed node. Initially each node is set to '0, except the observed node which is set to the specified sign. A message is sent to each neighbor. The sign of each message becomes the sign product of its previous sign and the sign of the link it traverses. Each message keeps a list of the nodes it has visited and its origin, so it can avoid visiting any node more than once. Each message travels on one evidential trail. Each node, on receiving a message, updates its own sign with the sign sum of itself and the sign of the message. Then it passes a copy of the message to all unvisited neighbors that need to update their signs.

```

Given: A qualitative probabilistic network,
          an evidence node  $e$ .
Output: Sign of the influence of  $e$ 
          on each node in the network
Data structures:
{ In each of the nodes }
  sign  $ch$ ;   { sign of change }
  sign  $evs$ ;  { sign of evidential support }
Main Program:
  for each node  $n$  in the network do  $ch := '0$ ;
  Propagate-Sign( $\emptyset, e, e, '+'$ );
Recursive procedure for sign propagation:
{  $trail$    visited nodes,
   $from$    sender of the message,
   $to$      recipient of the message,
   $sign$    sign of influence from  $from$  to  $to$  }
Propagate-Sign( $trail, from, to, sign$ )
begin
  if  $to \uparrow ch = sign \oplus to \uparrow ch$  then exit;
  { exit if already made the update }
   $to \uparrow ch := sign \oplus to \uparrow ch$ ;
  { update the sign of  $to$  }
   $trail := trail \cup to$ ;
  { add  $to$  to the set of visited nodes }
  for each  $n$  in the Markov blanket of  $to$  do
  begin
     $s :=$  sign of the link; { direct or intercausal }
     $sn := n \uparrow ch$ ; { current sign of  $n$  }
    if the link to  $n$  is active
      and  $n \notin trail$ 
      and  $sn \neq to \uparrow ch \otimes s$  then
        Propagate-Sign( $trail, to, n, to \uparrow ch \otimes s$ );
  end
end

```

Figure 3: The algorithm for qualitative sign propagation.

The character of the sign addition operator implies that each node can change its sign at most twice — first from '0 to '+, '-, or '?' and then, if at all, only to '?', which can never change to any other sign. Hence each node receives a request for updating its sign at most twice, and the total number of messages for the

network to reach stability is less than twice the number of nodes. Each message carries a list of visited nodes, which contains at most the total number of nodes in the graph. Hence, the algorithm is quadratic in the size of the network. Unfortunately, this propagation algorithm does not generalize straightforwardly to quantitative belief networks.

The sign propagation will reach each node in the network that is not d -separated from the evidence node e . It is possible to change the focus of the algorithm to determining the sign of some target node n and all nodes that are located on active trails from e to n . This requires a small amount of preprocessing, consisting of removal of irrelevant barren and ancestor nodes. The methods to do that are summarized in [Druzdzel, 1993].




Figure 4: Algorithm for qualitative belief propagation: An example.

Figure 4 shows an example of how the algorithm works in practice. Suppose that we want to know the effect of observing a high reading of the *HeOx Temp Probe* on other variables in the model. We set the signs of each of the nodes to '0 and start by sending a positive sign to *HeOx Temp Probe*, which is our evidence node. *HeOx Temp Probe* determines that its parent, node *HeOx Temp*, needs updating, as the sign product of '+' and the sign of the link '+' is '+' and is different from the current value of the node '0. After receiving this message, *HeOx Temp* sends messages to its direct descendants *High Ox Temp* and *Ox Tank Leak*, who also need updating. As the sign of *Ox Tank Leak* is already '+, *High Ox Temp* does not send any further messages. Seeing that *Ox Pressure Probe* needs updating, *Ox Tank Leak* will send it a message. The sign of this message is '-', because the sign of the qualitative influence between *Ox Tank Leak* and *Ox Pressure Probe* is '-. *Ox Pressure Probe* will not send any further messages and the algorithm will terminate leaving *HeOx Valve Problem* unaffected. The final sign in each node expresses how the probability of this node is impacted by observing a high reading of *HeOx Temp Probe*.

Conclusions and Applications

This paper has described an extension of belief propagation in qualitative probabilistic networks to multiply connected networks. Qualitative belief propagation can be performed in polynomial time. The type of reasoning addressed by QPNs and by the algorithm that we propose, namely determining the sign of evidential impact, is one of the few queries that can be answered in polynomial time in general networks, even when all nodes are of some restricted type, for example noisy-OR or continuous linear (Gaussian). Belief propagation is more powerful than graph reduction approach for two reasons: (1) it uses product synergy, which is a new qualitative property of probabilistic interactions, and (2) it offers a reasoning scheme, whose operators do not lead to loss of qualitative information and whose final results do not depend on the order of their application. Although examples of problems that can be resolved by belief propagation and not by graph reduction can be easily found, it is unfair to compare the strength of the two methods, as belief propagation uses an additional qualitative property, namely product synergy.

Wellman [1990a] describes several possible applications of QPNs, such as support for heuristic planning and identification of dominant decisions in a decision problem. The belief propagation approach proposed in this paper supports these applications, and has the additional advantage over the graph-reduction approach in that it preserves the underlying graph and determines the sign of the node of interest along with the signs of all intermediate nodes. This supports directly two new applications of QPNs. Firstly, it allows a computer program, in case of sign-ambiguity, to reflect about the model at a meta level and find the reason for ambiguity, for example, which paths are in conflict. Hence, it can suggest ways in which the least additional specificity could resolve the ambiguity. This may turn out to be a desirable property, as many multiply-connected networks that we used for testing the algorithm led in some queries, especially those involving intercausal reasoning, to ambiguous results. One reason for that is that the most common value of product synergy appears to be negative, which in loops often leads to conflicts with usually positive signs of links. A second application involves using the resulting signs for generation of intuitive qualitative explanations of how the observed evidence is relevant to the node of interest. The individual signs, along with the signs of influences, can be translated into natural language sentences describing paths of change from the evidence to the variable in question. A method for generation of verbal explanations of reasoning based on belief propagation is outlined in [Druzdzel, 1993].

Acknowledgments

We are most grateful to Michael Wellman and two anonymous reviewers for insightful comments.

References

- Druzdzel, M.J., and Henrion, M. 1993. Intercausal reasoning with uninstantiated ancestor nodes. Forthcoming.
- Druzdzel, M.J. 1993. *Probabilistic Reasoning in Decision Support Systems: From Computation to Common Sense*. Ph.D. diss., Department of Engineering and Public Policy, Carnegie Mellon University.
- Geiger, D.; Verma, T.S.; and Pearl, J. 1990. *d*-Separation: From theorems to algorithms. In Henrion, M.; Shachter, R.D.; Kanal, L.N.; and Lemmer, J.F., eds. 1990, *Uncertainty in Artificial Intelligence 5*. 139–148. North Holland: Elsevier.
- Henrion, M., and Druzdzel, M.J. 1991. Qualitative propagation and scenario-based approaches to explanation of probabilistic reasoning. In Bonissone, P.P.; Henrion, M.; Kanal, L.N.; and Lemmer, J.F., eds. 1991, *Uncertainty in Artificial Intelligence 6*. 17–32. North Holland: Elsevier.
- Horvitz, E.J.; Ruokangas, C.; Srinivas, S.; and Barry, M. 1992. A decision-theoretic approach to the display of information for time-critical decisions: The Vista project. In *Proceedings of SOAR-92*, NASA/Johnson Space Center, Houston, Texas.
- Kim, J.H., and Pearl, J. 1983. A computational model for causal and diagnostic reasoning in inference systems. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence, IJCAI-83*, Los Angeles, Calif. 190–193.
- Milgrom, Paul R. 1981. Good news and bad news: Representation theorems and applications. *Bell Journal of Economics* 12(2):380–391.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, Calif.: Morgan Kaufmann
- Shachter, R.D. 1986. Evaluating influence diagrams. *Operations Research* 34(6):871–882.
- Wellman, M.P. and Henrion, M. 1991. Qualitative intercausal relations, or explaining “explaining away”. In *KR-91, Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, Cambridge, MA. 535–546.
- Wellman, M.P., and Henrion, M. 1993. Explaining “explaining away”. *IEEE Trans. PAMI* 15(3):287–291.
- Wellman, M.P. 1990a. *Formulation of Tradeoffs in Planning Under Uncertainty*. London: Pitman.
- Wellman, M.P. 1990b. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence* 44(3):257–303.
- Wellman, M.P. 1990c. Graphical inference in qualitative probabilistic networks. *Networks* 20(5):687–701.