

Chapter 4

BELIEF UPDATING BY NETWORK PROPAGATION

*Oh, would that my mind could let fall its dead ideas,
as the tree does its withered leaves.*

— *André Gide*

This chapter deals with fusing and propagating the impact of new evidence and beliefs through Bayesian networks so that each proposition eventually will be assigned a certainty measure consistent with the axioms of probability theory. We start in Section 4.1 by arguing that any viable model of human reasoning should be able to perform this task with a self-activated propagation mechanism, i.e., with an array of simple and autonomous processors, communicating locally via the links provided by the network itself. The impact of each new piece of evidence is viewed as a perturbation that propagates through the network via message-passing between neighboring variables, with minimal external supervision. In Section 4.2 we show that these objectives can be fully realized with *causal trees*, i.e., Bayesian networks in which each node has at most one parent. In Section 4.3 we extend the result to *causal polytrees*, i.e., trees with arbitrary arrow orientation, and thus allow multiple roots and multiple parents. In both cases, we identify belief parameters, communication messages, and updating rules to guarantee that equilibrium will be reached in time proportional to the longest path in the network and that at equilibrium each proposition will be accorded a belief measure equal to its posterior probability, given all the available evidence. In Section 4.3.2, we illustrate this propagation method's evidence-pooling and credit-assignment policies with a canonical model, where multiple causes are assumed to interact disjunctively. Several approaches to achieving autonomous propagation in multiply connected networks are discussed in Section 4.4, including clustering, conditioning, and stochastic simulation. Finally, Section 4.5 extends the inferential repertoire of Bayesian networks to include answering Boolean queries under propositional constraints, with a special emphasis on conjunctive queries.

4.1 AUTONOMOUS PROPAGATION AS A COMPUTATIONAL PARADIGM

Since a fully specified Bayesian network constitutes a complete probabilistic model of the variables in a domain (i.e., it specifies a joint distribution over the variables), the network contains the information needed to answer all probabilistic queries about these variables. The queries might be requests to interpret specific input data or, if utility information is provided, requests to recommend the best course of action. Interpretation requires instantiating a set of variables corresponding to the input data, calculating their impact on the probabilities of variables designated as hypotheses, and finally, selecting the most likely combinations of these hypotheses.

In principle, once we have a joint distribution function P , the interpretation task can be performed mechanically using purely algebraic methods. For example, the impact of the observation $X_j = x_j$ on another variable X_i can be obtained from the conditional probabilities $P(X_i | x_j)$ for each value x_j in the domain of X_j . Using the textbook definition

$$P(X_i | x_j) = \frac{P(X_i, x_j)}{P(x_j)}$$

we compute $P(X_i, x_j)$ by summing the joint distribution $P(X_1, \dots, X_n)$ over all variables except X_i and X_j . The summation can be executed in any order, but exponential savings can sometimes be realized by selecting a judicious ordering.

Network representations provide a valuable guide for making this selection. Summing over a variable, say X_k , amounts to eliminating X_k from the network while maintaining the proper dependencies among remaining variables—adding links between those neighbors of X_k that were d -separated by X_k and calculating the conditional probabilities associated with the new links. Since the size of a link matrix increases exponentially with the number of arrows that converge on a given variable, it is important to eliminate variables in an order that minimizes the number of converging arrows created throughout the process. Techniques for finding a good elimination ordering have been developed in the operations research literature [Bertelé and Brioschi 1972] and have been used for manipulating influence diagrams [Shachter 1986] (see Chapter 6). Such techniques do not, of course, avoid the exponential worst-case complexity of the interpretation task (the problem is NP -hard [Rosenthal 1975; Cooper 1987]), but they exploit the structural properties of sparse networks.

However, node elimination has several shortcomings. First, the process requires full supervision by a monitor that must access all parts of the network and use external computational facilities to decide, at any given stage, which variable should be eliminated next. The use of such a global supervisor is foreign to human reasoning because it requires faculties beyond most humans, e.g., comprehending

the entire structure of the network and quickly diverting attention from one section of the network to another. Our limited short-term memory, narrow range of attention, and inability to shift rapidly between alternative lines of reasoning suggest that our reasoning process is fairly local, progressing incrementally along preestablished pathways. Second, elimination techniques literally cover their tracks; they provide the final impact of a piece of evidence on a single hypothesis, but do not calculate the impact of the evidence on the nodes eliminated in the process. In many applications, we wish to know the updated belief of every variable in the network. Third, and perhaps most important, the elementary steps in the process of node elimination often have no conceptual correlates. They create and calculate spurious dependencies among variables normally perceived to be independent, and the dependencies are hard to explain in qualitative terms. Finally, elimination techniques are basically sequential, and there is growing interest in reasoning models that permit unsupervised parallelism. The interest is motivated both by technological advances in parallel computation and by the need to develop viable models of human reasoning. The speed and ease with which people perform some low-level interpretive functions, such as recognizing scenes, reading text, and even understanding stories, strongly suggest that such processes involve a significant amount of parallelism and that most of the processing is done at the knowledge level itself [Shastri and Feldman 1984].

4.1.1 Constraint Propagation and Rule-based Computation

We can model such phenomena by viewing a belief network not merely as a passive code for storing factual knowledge but also as a computational architecture for reasoning about that knowledge. This means the links in the network should be treated as the only mechanisms that direct and propel the flow of data through the process of querying and updating beliefs. Accordingly, we assume that each node in the network is given a separate processor, which maintains the parameters of belief for the host variable and manages the communication links to and from neighboring, conceptually related variables. The communication lines are assumed to be open at all times, i.e., each processor may, at any time, interrogate its neighbors and compare their parameters to its own. If the compared quantities satisfy some local constraints, no activity takes place. However, if any constraint is violated, the responsible node is activated and the violating parameter corrected. This, of course, activates similar revisions at neighboring nodes and initiates a multidirectional propagation that will continue until equilibrium is reached.

EXAMPLE 1: To illustrate the process of constraint propagation, consider the graph coloring problem depicted by Figure 4.1. Each node in the graph must be assigned one of

three colors, {1, 2, 3}, in such a way that no two adjacent nodes will have identical colors. The constraints in this problem are local; no node can assume a color seen at any of its neighbors. A local approach would be to assign a processor to each node and have it compare its current color to the colors of its neighbors. If equality is discovered, the processor should choose a new color, different (if possible) from the neighbors' colors. We assume there is no synchronization, so all possible activation schedules could be realized.




Figure 4.1. Demonstrating constraint propagation in the graph coloring problem.

Figure 4.1a shows the initial state of the system: all nodes are colored 1. Figure 4.1b shows the configuration after the nodes are activated in the order A, B, C, D: a stable solution is established once C selects the color 2. If, instead, the activation schedule happens to be A, C, B, D, and if C selects the color 3, a deadlock occurs (Figure 4.1c): B cannot find a color different from all its neighbors. A way out of such deadlock is to instruct each processor to change its color arbitrarily if no better one can be found. Indeed, if B changes its color to 2 (temporarily conflicting with A, as in Figure 4.1d), it forces A to choose the color 3, thus realizing a global solution (Figure 4.1e).

Even if a global solution exists, there is no guarantee, in general, that it will be found by repeated local relaxations. However, if the escape from deadlock is totally random, the probability of reaching a solution within a given time t approaches 1 as t increases. Note also that if the node activation is both parallel and synchronous, the system can fall into an indefinite dynamic loop without reaching a solution. For example, starting with the state of Figure 4.1a, all processors may simultaneously choose the color 2, then 3, etc. Such pathological behavior will not occur in the networks that we shall deal with.

One of the main reasons for adopting this distributed computation paradigm in evidential reasoning tasks is that it automatically exploits the independencies embodied in the network, via subtask decomposition, to gain a substantial reduction in complexity. For example, if the Markov neighbors of some variable X have successfully computed their combined distribution function, X can compute its own distribution without interacting with any variable outside its neighborhood (see Eq. (3.12)). Moreover, if the network has a tree structure, then X can compute its distribution by consulting each of its neighbors separately. If any of X's neighbors undergo change, X updates its own distribution and reports the update to the other neighbors, and so on until, at the network's periphery, we meet evidential variables whose probabilities are predetermined and the process halts.

The second reason to adopt this distributed paradigm is that it leads to a "transparent" revision process, in which the intermediate steps can be given an intuitively meaningful interpretation. Because a distributed process allows each computational step to obtain input only from neighboring, semantically related variables, and because the activation of these steps proceeds along semantically familiar pathways, people find it easy to give meaningful interpretation to the individual steps and thus gain confidence in the final result. It is also possible to generate qualitative justifications mechanically by tracing the sequence of operations along the activated pathways and giving them causal or diagnostic interpretations using appropriate verbal expressions.

The ability to update beliefs by an autonomous propagation mechanism also has a profound effect on sequential implementations of evidential reasoning. Of course, when this architecture is simulated on sequential machines, the notion of autonomous processors working simultaneously is only a metaphor, but it signifies the complete separation of the stored knowledge from the control mechanism. This separation is the proclaimed, if rarely achieved, goal of rule-based systems. It guarantees the utmost flexibility for a sequential controller: the computations can be performed in any order, with no need to remember or verify which parts of the network have already been updated. Thus, belief updating may be activated by changes occurring in logically related propositions (*spreading activation*), by requests for evidence arriving from a central supervisor (*goal-directed activation*), by a predetermined schedule, or entirely at random. The communication and interaction among individual processors can be simulated using a blackboard architecture [Lesser and Erman 1977], where each proposition is designated specific areas of memory to access and modify.

Finally, the uniformity of this propagation scheme makes it easy to formulate in object-oriented languages: the nodes are objects of the same generic type, and the belief parameters are the messages by which interacting objects communicate. The programmer need only specify how a single object reacts to changes occurring at its neighbors; he need not provide timing information or say where to store partial results.

Constraint propagation mechanisms have a special appeal for AI researchers because they are similar in many respects to logical inference rules. We already have seen that an inference rule of the form "If premise A, then action B" constitutes a very attractive unit of computation for three reasons:

1. The triggering mechanism is local, i.e., it grants a license to initiate the action whenever the premise A is true in the knowledge base K, regardless of the other information that K contains.
2. The action is computationally simple and normally involves only a minor adjustment in the knowledge base.
3. Both the triggering mechanism and the action are conceptually meaningful and are therefore easy to program, modify, and explain.

Almost identical features hold in the constraint propagation formalism. Each processor receives a permanent license to act whenever a certain condition develops among its neighbors. The action is simple, since it involves only local perturbation of the processor's parameters. Both the activation and the action are meaningful because they engage semantically related propositions. Thus, whenever a problem can be solved by a constraint propagation mechanism, it is also easy to formulate in a pure production-rule formalism.

4.1.2 Why Probabilistic Reasoning Seems to Resist Propagation

While constraint propagation mechanisms have been essential to many AI applications, e.g., vision [Rosenfeld, Hummel, and Zucker 1976; Waltz 1972] and truth maintenance [McAllester 1980], their use in evidential reasoning, surprisingly, has been limited to non-Bayesian formalisms [Lowrance 1982; Quinlan 1983; Shastri and Feldman 1984]. There have been several reasons for this, all based on the essential difference between the probabilistic statement $P(A|B) = p$ and the logical rule $B \rightarrow A$ (see Sections 1.1.4, 1.3.1, and 2.3.1).

First, the conditional probabilities characterizing the links in the network do not seem to impose definitive constraints on the probabilities that can be assigned to the nodes. Consider a pair of nodes A and B linked by an arrow $B \rightarrow A$ and quantified by the conditional probabilities $P(A|B)$ and $P(A|\neg B)$. The quantifier $P(A|B)$ restricts the belief accorded to a only in a very special set of circumstances, namely, when b is known with absolute certainty to be true and when no other evidential data is available. Normally, all internal nodes in the network will be subject to some uncertainty. Thus, if processor A inspects its neighbor B and finds it in a state of uncertainty with $P(b) < 1$, it still cannot proceed with a definite action on $P(a)$. A natural recourse would be to compute the weighted average

$$P(a) = P(a|b)P(b) + P(a|\neg b)P(\neg b).$$

After the arrival of some evidence e , however, the posterior distributions A and B are no longer governed by $P(a|b)$ but rather by $P(a|b, e)$, via

$$P(a|e) = P(a|b, e)P(b|e) + P(a|\neg b, e)P(\neg b|e),$$

which may be a totally different relationship. (To take an extreme example, if e fully confirms or denies a , the overall probability of a becomes totally insensitive to $P(a|b)$.) The result is that any arbitrary assignment of beliefs to the propositions a and b can be consistent with the value of $P(a|b)$ that was initially

assigned to the link connecting them; consequently, among these parameters, no violation of a constraint can be detected locally.

Second, the disparity between $P(a|b, e)$ and $P(a|b)$ suggests that once a new piece of evidence is introduced, the original weights on the link no longer retain their intended meaning; hence, they should not remain fixed but should undergo constant adjustment as new evidence arrives. This requires enormous computational overhead and an external unit to perform the adjustment, so it defeats the whole purpose of local propagation.

Finally, the presence of both top-down (*predictive*) and bottom-up (*diagnostic*) inferences in evidential reasoning has caused apprehensions that pathological instability, deadlock, and circular reasoning will develop once we allow the propagation process to run its course unsupervised [Lowrance 1982]. Indeed, if a stronger belief in a given hypothesis means a greater expectation of the occurrence of its various manifestations, and if, in turn, a greater certainty in the occurrence of these manifestations adds further credence to the hypothesis, how can one avoid infinite updating loops when the processors responsible for these propositions begin to communicate with one another?

EXAMPLE 2: You spread a rumor about person X to your neighbor N_1 . A few days later, you hear the same rumor from N_1 . Should you increase your belief in the rumor now that N_1 acknowledges it, or should you determine first whether N_1 heard it from another source besides you? It is clear that if you were N_1 's only source of information, your belief should not change, but if N_1 , independently confirmed the validity of the rumor, you have good reason to increase your belief in it.

Similar considerations apply to communicating processors that represent interdependent propositions. Imagine that a processor F , representing the event *Fire*, communicates asynchronously with a second processor S , representing the event *Smoke*. At time t_1 , some evidence (e.g., the distant sound of a fire engine) gives a slight confirmation to F , thus causing the probability of *Fire* to increase from $P(F)$ to $P_1(F)$. At a later time, t_2 , processor S may decide to interrogate F ; upon finding $P_1(F)$, it revises the probability of *Smoke* from $P(S)$ to $P_2(S)$ in natural anticipation of smoke. Still later, at t_3 , processor F is activated, and upon finding an increased belief $P_2(S)$ in *Smoke*, it increases $P_1(F)$ to an even higher value, $P_3(F)$. This feedback process may continue indefinitely, the two processors drawing steady mutual reinforcement void of any empirical basis, until eventually the two propositions, *Fire* and *Smoke*, appear to be firmly believed.

Such dangers are not unique to probabilistic reasoning, but should be considered in any hierarchical model of cognition where mutual reinforcement takes place between lower and higher levels of processing, e.g., connectionist models of reading [Rumelhart and McClelland 1982] and language production [Dell 1985].

To prevent these phenomena, we need a mechanism to keep track of the *sources* of belief, so that evidence is not counted twice and so that the impact of one piece of evidence is not fed back to its source. Unfortunately, source identification requires an overview of the entire network, and the question arises whether it can be represented and adjusted locally as an integral part of the propagation process.

This chapter demonstrates that in a large class of networks, coherent and stable probabilistic reasoning *can* be accomplished by local propagation mechanisms, keeping the weights on the links constant throughout the process. This is done by characterizing the belief in a proposition by a list of parameters, each representing the degree of support the host proposition obtains from one of its neighbors. In the next two sections we show that maintaining such a record of the sources of belief facilitates local updating of beliefs, and that the network relaxes to a stable equilibrium, consistent with the axioms of probability theory, in time proportional to the network diameter. Such a record of parameters is also postulated as a mechanism that permits people to retrace rationales and assemble explanations for currently held beliefs.

4.2 BELIEF PROPAGATION IN CAUSAL TREES

4.2.1 Notation

We shall first consider tree-structured causal networks, i.e., those in which every node except the one called *root* has exactly one incoming link. We allow each node to represent a multi-valued variable, comprising a collection of mutually exclusive hypotheses (e.g., the identity of an organism: *Org1*, *Org2*, ...) or observations (e.g., a patient's temperature: *High*, *Medium*, *Low*). Let variables be labeled by capital letters (A, B, \dots, X, Y, Z) and their possible values by the corresponding lowercase letters (a, b, \dots, x, y, z). In dealing with propositional variables, the symbols $+$ and \neg will be used to denote the affirmation and denial, respectively, of propositions. For example, $+a$ stands for $A = TRUE$, and $\neg a$ stands for $A = FALSE$. Each directed link $X \rightarrow Y$ is quantified by a fixed conditional probability matrix M , in which the (x, y) entry is given by

$$M_{y|x} \triangleq P(y|x) \triangleq P(Y=y|X=x) \quad (4.1)$$

$$\rightarrow y$$

$$\downarrow x \begin{bmatrix} P(y_1|x_1) & P(y_2|x_1) & \dots & P(y_n|x_1) \\ P(y_1|x_2) & P(y_2|x_2) & \dots & P(y_n|x_2) \\ \vdots & \vdots & \ddots & \vdots \\ P(y_1|x_m) & P(y_2|x_m) & \dots & P(y_n|x_m) \end{bmatrix}$$

Normally, the directionality of the arrow designates X as the set of causal hypotheses and Y as the set of consequences or manifestations of these hypotheses.

EXAMPLE 3: In a certain trial there are three suspects, one of whom has definitely committed a murder. The murder weapon, showing some fingerprints, was later found by police. Let X identify the last user of the weapon, namely, the killer. Let Y identify the last holder of the weapon, i.e., the person whose fingerprints were left on the weapon, and let Z represent the possible readings that may be obtained in a fingerprint laboratory.

The relations between these three variables normally would be expressed by the chain $X \rightarrow Y \rightarrow Z$; X generates expectations about Y , and Y generates expectations about Z , but X has no influence on Z once we know the value of Y .

To represent the commonsense knowledge that the killer is normally the last person to hold the weapon, we use a 3×3 conditional probability matrix:

$$M_{y|x} = P(y|x) = \begin{cases} 0.80 & \text{if } x = y \quad x, y = 1, 2, 3 \\ 0.10 & \text{if } x \neq y \quad x, y = 1, 2, 3. \end{cases} \quad (4.2)$$

To represent the reliability of the laboratory test, we use a matrix $M_{z|y} = P(z|y)$, satisfying

$$\sum_z P(z|y) = 1 \quad \text{for } y = 1, 2, 3.$$

Each entry in this matrix represents an if-then rule of the following type: "If the fingerprint is of Suspect y , then expect a reading of type z , with certainty $P(z|y)$."

Note that this convention is at variance with that used in many expert systems (e.g., MYCIN), where rules point from evidence to hypothesis (e.g., If symptom, then disease), thus denoting a flow of mental inference. By contrast, the arrows in Bayesian networks point from causes to effects, or from conditions to consequences, thus denoting a flow of constraints attributed to the physical world. The reason for this choice is that people often prefer to encode experiential knowledge in causal schemata [Tversky and Kahneman 1977], and as a consequence, rules expressed in causal form are assessed more reliably.[†]

Incoming information may be of two types: *specific evidence* or *virtual evidence*. Specific evidence corresponds to direct observations that affect the belief in some variables in the network. Virtual evidence corresponds to judgments based on undisclosed observations that are outside the network but have

[†] It appears that frames used to index human memory by and large are organized to evoke *expectations* rather than *explanations*. The reason could be because expectation-evoking frames normally consist of more stable relationships. For example, $P(y|z)$ in Example 3 would vary drastically with the proportion of people who have type z fingerprints; $P(z|y)$, on the other hand, depends merely on the similarity between the type of fingerprint that Suspect y has and the readings observed in the lab: it is perceived to be a stable local property of the laboratory procedure, independent of other information regarding Suspect y .

bearing on variables in the network. Such evidence is modeled by dummy nodes representing the undisclosed observations, connected by unquantified (dummy) links to the variables affected by the observations. These links will carry information one way only, from the evidence to the variables affected by it.

For example, if it is impractical for the fingerprint laboratory to disclose all possible readings (in variable Z) or if the laboratory chose to base its finding on human judgment, Z will be represented by a dummy node, and the link $Y \rightarrow Z$ will specify the degree to which each suspect is likely to bear the fingerprint pattern examined. For example, the laboratory examiner may issue a report in the form of a list (0.80, 0.60, 0.50), stating that she is 80% sure that the fingerprint belongs to Suspect 1, 60% sure that it belongs to Suspect 2, and 50% sure that it belongs to Suspect 3. If the examiner was totally unbiased before the test, such a profile of belief can be established only if the likelihood ratio is

$$P(Z_{\text{observed}} | y = 1) : P(Z_{\text{observed}} | y = 2) : P(Z_{\text{observed}} | y = 3) = 8 : 6 : 5,$$

which will be our standard way to characterize the impact of virtual evidence (see Sections 2.2.2 and 2.3.3). Because these numbers need not sum to unity, each judgment can be formed independently of the others—each suspect's fingerprints can be compared separately with those found on the weapon.

All incoming evidence, both specific and virtual, will be denoted by e and will be regarded as emanating from a set E of *instantiated* variables, i.e., variables whose values are known. For example, if the laboratory examination is the only evidence available in Example 3, we shall write $E = \{Z\}$ and $e = \{Z = Z_{\text{observed}}\}$. If several facts become known, say, $A = \text{TRUE}$, $B = \text{FALSE}$ and $X = x$, then $E = \{A, B, X\}$ and $e = \{+a, -b, X = x\}$.

For the sake of clarity, we will distinguish between the fixed conditional probabilities that label the links, e.g., $P(y|x)$, and the dynamic values of the updated node probabilities, e.g., $P(x|e)$. The latter will be denoted by $BEL(x)$, which reflects the overall belief accorded to proposition $X = x$ by all evidence so far received. Thus,

$$BEL(x) \triangleq P(x|e),$$

where e is the value combination of all instantiated variables.

Since we will be dealing with discrete variables, functions such as $\lambda(x)$, $P(x)$, and $BEL(x)$ can be regarded as lists, or *vectors*, with each component corresponding to a different value of X . For example, if the domain of X is $D_X = \{\text{Hot}, \text{Medium}, \text{Cold}\}$, we can write

$$\begin{aligned} BEL(x) &= (BEL(X = \text{Hot}), BEL(X = \text{Medium}), BEL(X = \text{Cold})) \\ &= (0.1, 0.2, 0.7). \end{aligned}$$

The product $f(x) \cdot g(x)$ of two such vectors will stand for term-by-term multiplication, e.g.,

$$(1, 2, 3) (3, 2, 1) = (1 \times 3, 2 \times 2, 3 \times 1) = (3, 4, 3).$$

Inner products (or *dot products*) will be denoted by a dot \cdot , e.g.,

$$f(x) \cdot g(x) = (1, 2, 3) \cdot (3, 2, 1) = 1 \times 3 + 2 \times 2 + 3 \times 1 = 10.$$

The dot symbol \cdot will also be used to indicate matrix products, e.g.,

$$f(x) \cdot M_{y|x} \triangleq \sum_x f(x) M_{y|x}.$$

The summation will always be taken over the repeated index, thus eliminating the need for transposing matrices or distinguishing between row and column vectors.

We shall use the symbol α to denote a *normalizing* constant, e.g.,

$$\alpha(1, 1, 3) = (0.2, 0.2, 0.6),$$

and the symbol β to denote an *arbitrary* constant, e.g.,

$$K\beta f(x) = \beta f(x) \quad \text{and}$$

$$\alpha\beta f(x) = \alpha f(x).$$

A vector of 1s will be written $\mathbf{1}$; for example,

$$BEL(x) = \alpha \mathbf{1} = \alpha(1, 1, 1) = (0.25, 0.25, 0.25).$$

4.2.2 Propagation in Chains

Consider the simplest of all tree-structured networks, namely, a network consisting of two nodes and a single link, $X \rightarrow Y$. If evidence $e = \{Y=y\}$ is observed, then from Bayes' Rule, the belief distribution of X is given by

$$BEL(x) = P(x|e) = \alpha P(x) \lambda(x), \quad (4.3)$$

where $\alpha = [P(e)]^{-1}$, $P(x)$ is the prior probability of X , and $\lambda(x)$ is the likelihood vector

$$\lambda(x) = P(e|x) = P(Y = y|x) = M_{y|x}. \quad (4.4)$$

In short, $\lambda(x)$ is simply the y 's column of the link matrix M , as in Eq. (4.1). Since this matrix is stored at node Y , $\lambda(x)$ can be computed at Y and transmitted as a message to X , enabling X to compute its belief distribution $BEL(x)$.

If Y is not observed directly but is supported by indirect observation $e = \{Z=z\}$ of a descendant Z of Y , we have the chain $X \rightarrow Y \rightarrow Z$, and we can still write

$$BEL(x) \triangleq P(x|e) = \alpha P(x) \lambda(x).$$

The likelihood vector $\lambda(x)$ can no longer be directly obtained from the matrix $M_{y|x}$, however, but must reflect the matrix $M_{z|y}$ as well. Conditioning and summing on the values of Y , we can write

$$\begin{aligned} \lambda(x) &= P(e|x) = \sum_y P(e|y, x) P(y|x) = \sum_y P(e|y) P(y|x) \\ &= M_{y|x} \cdot \lambda(y), \end{aligned} \tag{4.5}$$

using the fact that Y separates X from Z . Thus, we have shown that node X can still calculate its likelihood vector $\lambda(x)$ if it somehow gains access to the vector $\lambda(y)$.

Generalizing to the chain of Figure 4.2, every node can calculate the correct current value of its λ vector if it learns the correct λ vector of its successor.




Figure 4.2. Each node in a causal chain can calculate its λ from the λ of its successor.

Since the chain ends with an observed variable whose value is determined externally, the λ vector of all variables can be determined recursively. If the chain ends with an unobserved variable Z and we set $\lambda(z) = 1$ for all z . Eqs. (4.3) and (4.5) are still valid, because every variable will obtain $\lambda = 1$ and all belief distributions will coincide with the prior distributions. Assuming that each node constantly inspects the λ of its child and updates its own λ accordingly, we are guaranteed that every variable along the chain will obtain its correct λ , properly reflecting any changes that might have occurred in e . This updating process is analogous to the way soldiers are counted.

EXAMPLE 4: A platoon of soldiers is marching at night in enemy territory. The leader wants to know how many soldiers remain under his command. He sends a "count" signal to the soldier behind him. This person, in turn, looks behind, and if someone is there, he passes on the "count" signal; if no one is left behind him, he returns the signal "1" to the soldier in front of him. The soldier in front receives the "1," adds 1 (for himself) and sends "2" to the soldier in front of him, and so on. The leader eventually receives the correct count. (See Figure 4.3.) In fact, the leader need not be at the head of the platoon. He can initiate a "count" command to both his front and his back, wait for responses from both sides, and add the values received (see Figure 4.4).




Figure 4.3. Distributed soldier-counting.




Figure 4.4. Distributed soldier-counting with leader in line.

This common procedure suffers because the leader may not be aware of any missing soldiers until he decides to count, and even then the count gets back to him only after instructing each soldier to constantly update and communicate the messages without waiting for the "count" signal, or more efficiently, to initiate communication as soon as a

change is seen in the immediate environment. Thus, no communication takes place under normal conditions, but when any soldier suddenly finds himself at either the front or the end of the line, he will initiate a message-passing process that propagates toward the leader and eventually terminates at the leader (who is now passive) with the correct count (see Figure 4.5).




Figure 4.5. Soldier-counting initiated by changes at endpoints.

To draw the analogy closer to belief updating, let's remove the leader and force every soldier to be constantly aware of the current total count. In such a system, the messages, instead of stopping at any particular individual, should continue to propagate toward the periphery; the forward-moving messages should propagate all the way to the front of the line, and the backward-moving messages should propagate to the end of the line. Every soldier follows the same rule: receive a count from the person behind, add 1, and transmit the result to the person in front; receive a count from the person in front, add 1, and transmit to the one behind (see Figure 4.6). Note that each soldier must maintain and communicate two separate parameters, the back count and the front count; the overall count (the sum of the two) is not a message that can sustain the propagation.




Figure 4.6. Leaderless soldier-counting using dual-parameter communication.

BIDIRECTIONAL PROPAGATION

The need for dual-parameter communication also exists in belief updating, where new evidence can emerge from both a descendant of a node and its ancestor. For example, in the chain of Figure 4.7, variables are instantiated at both the head and the tail of the chain, and we may wish to calculate $BEL(X)$ as a function of the values, e^+ and e^- , that these variables take.[†]




Figure 4.7. A causal chain with evidential data at its head (e^+) and tail (e^-).

In Eq. (4.3), $\lambda(x)$ was defined by $P(e|x)$, e being the total evidence available. We now find it more convenient to handle the impact of e^- and e^+ by two separate vectors,

$$\lambda(x) = P(e^-|x) \tag{4.6a}$$

and

$$\pi(x) = P(x|e^+). \tag{4.6b}$$

Expressing the total belief distribution $BEL(x)$ in terms of $\lambda(x)$ and $\pi(x)$, with X separating e^+ from e^- , we have

$$\begin{aligned} BEL(x) &\triangleq P(x|e^+, e^-) = \alpha P(e^-|x, e^+) P(x|e^+) = \alpha P(e^-|x) P(x|e^+) \\ &= \alpha \lambda(x) \pi(x), \end{aligned} \tag{4.6c}$$

which is identical to Eq. (4.3) with $\pi(x)$ replacing $P(x)$.

To illustrate how information about $\pi(x)$ propagates from e^+ down the chain, let us condition Eq. (4.6b) on the values of the parent variable U :

$$\pi(x) = P(x|e^+) = \sum_u P(x|u, e^+) P(u|e^+).$$

Since U separates X from e^+ , we obtain

$$\pi(x) = \sum_u P(x|u) \pi(u) = \pi(u) \cdot M_{x|u}. \tag{4.7}$$

[†] e^+ might represent the background knowledge one has about T , in which case the prior probability $P(x)$ will be identical to $P(x|e^+)$.

We see that the forward propagation of π 's parameters is similar to the backward propagation of λ 's parameters; both involve vector multiplication by the appropriate link matrix. Each node can now compute its own π and λ after obtaining the π of its parent and the λ of its child. (See Figure 4.8.)




Figure 4.8. Belief calculation using bidirected message passing in causal chains.

EXAMPLE 5: Referring to the trial story of Example 3, let $e^- = \{Z = z\}$ represent the experience of examining the fingerprints left on the murder weapon, and let e^+ stand for all other testimony heard in the trial. So, $\pi(x) = P(x|e^+)$ stands for our prior certainty that the x -th suspect is the killer, $\pi(y) = P(y|e^-)$ stands for our prior certainty (before examining the fingerprints) that the y -th suspect was the last person to hold the weapon, and $\lambda(y) = P(e^-|y)$ represents the report issued by the fingerprint laboratory. Taking $\pi(x) = (0.8, 0.1, 0.1)$ and using the matrix of Eq. (4.2), we get

$$\pi(y) = (0.8, 0.1, 0.1) \cdot \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0.1 & 0.8 & 0.1 \\ 0.1 & 0.1 & 0.8 \end{bmatrix} = (0.66, 0.17, 0.17).$$

Prior to inspection of the fingerprints, all λ s are unit vectors 1 and the message profile on the chain is as shown in Figure 4.9.




Figure 4.9. Initial beliefs and messages in Example 5.

Now assume that a laboratory report arrives, summarized by $\lambda(y) = \beta(0.8, 0.6, 0.5)$. Node Y updates its belief to read

$$\begin{aligned} BEL(y) &= \alpha \lambda(y) \pi(y) = \alpha(0.8, 0.6, 0.5) (0.66, 0.17, 0.17) \\ &= (0.738, 0.143, 0.119) \end{aligned}$$

and delivers to X its $\lambda(x)$ vector. Upon receiving this message, node X computes its new $\lambda(x)$ vector,

$$\lambda(x) = M_{y,x} \cdot \lambda(y) = \beta \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0.1 & 0.8 & 0.1 \\ 0.1 & 0.1 & 0.8 \end{bmatrix} \cdot \begin{bmatrix} 0.8 \\ 0.6 \\ 0.5 \end{bmatrix} = \beta \begin{bmatrix} 0.75 \\ 0.61 \\ 0.54 \end{bmatrix}.$$

and its new belief distribution becomes

$$\begin{aligned} BEL(x) &= \alpha \lambda(x) \pi(x) = \alpha(0.75, 0.61, 0.54) (0.8, 0.1, 0.1) \\ &= (0.840, 0.085, 0.076). \end{aligned}$$

The messages are distributed as in Figure 4.10.

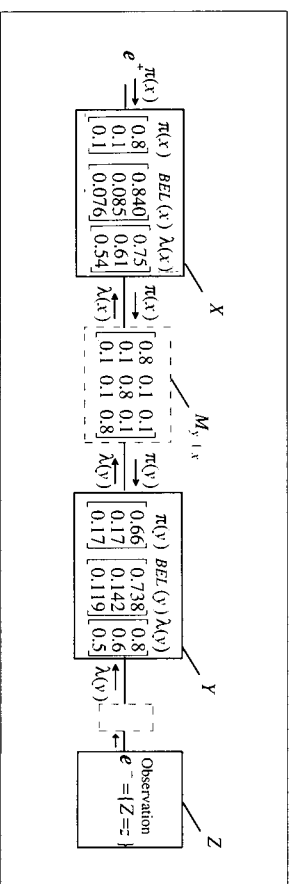


Figure 4.10. Beliefs and messages in Example 5, after obtaining the laboratory report $\lambda(y)$.

Now assume that Suspect 1 produces a very strong alibi, which (discounting fingerprint information) reduces the odds that he could have committed the crime from 0.80 to 0.28, thus yielding a revised prior probability of

$$\pi(x) = (0.28, 0.36, 0.36).$$

This change propagates toward Y and results in a revised $\pi(y)$:

$$\pi(y) = \pi(x) \cdot M_{y,x} = (0.28, 0.36, 0.36) \cdot \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0.1 & 0.8 & 0.1 \\ 0.1 & 0.1 & 0.8 \end{bmatrix} = (0.30, 0.35, 0.35).$$

Each processor can now compute the revised total belief of its variable, taking into account the evidential impact of the fingerprint findings:

$$\begin{aligned}
 BEL(x) &= \alpha\pi(x)\lambda(x) = \alpha(0.28, 0.36, 0.36)(0.75, 0.61, 0.54) \\
 &= \alpha(0.210, 0.220, 0.194) \\
 &= (0.337, 0.352, 0.311), \\
 BEL(y) &= \alpha\pi(y)\lambda(y) = \alpha(0.30, 0.35, 0.35)(0.80, 0.60, 0.50) \\
 &= \alpha(0.240, 0.210, 0.175) \\
 &= (0.384, 0.336, 0.280).
 \end{aligned}$$

Thus, Suspect 2 now becomes the strongest candidate for being the killer (with $P(X = 2) = 0.349$), though Suspect 1 is still most likely to be the owner of the fingerprint (with $P(Y = 1) = 0.384$). The final message distribution is shown in Figure 4.11.




Figure 4.11. Beliefs and messages in Example 5, incorporating alibi information $\pi(x)$.

Note how the separation between causal and evidential support (i.e., between the π s and the λ s) prevents feedback, circular reasoning, and indefinite relaxations of the type discussed in Section 4.1.1. Suspect 1's alibi renders him less likely (by a factor of 0.384/0.738) to be the owner of the incriminating fingerprints, but this reduction is not fed back to further influence his likelihood of being the killer (this would amount to counting the alibi information twice) because $\lambda(x)$ and $\lambda(y)$ are unaffected by changes occurring in $\pi(x)$. Keeping these two modes of support orthogonal to each other ensures stable updating in a single pass.

The local computations performed in Example 1, with two additional features: the relaxations of the type discussed in Example 1, with two additional features: the constraints are equalities, and they can always be satisfied locally. If the state of each processor is defined by its associated π and λ vectors, then the updating procedure can be written as a collection of local inference rules, identical in form and spirit to those used in constraint relaxation and logical deduction. For

example, assuming the content of π and λ is stored in two registers, called Π and Λ , the behavior of processor X in Figure 4.8 is specified by three inference rules:

$$\text{If } (X \rightarrow Y)_{M_{y|x}} \text{ and } \Lambda(Y) = \lambda(y) \text{ then } \Lambda(X) = \lambda(y) \cdot M_{y|x}, \quad (4.8)$$

$$\text{If } (U \rightarrow X)_{M_{x|u}} \text{ and } \Pi(U) = \pi(u) \text{ then } \Pi(X) = M_{x|u} \cdot \pi(u), \quad (4.9)$$

$$\text{If } \Lambda(X) = \lambda(x) \text{ and } \Pi(X) = \pi(x) \text{ then } BEL(x) = \alpha\lambda(x)\pi(x). \quad (4.10)$$

The first rule, for instance, reads:

If the rule "If $X = x$ then $Y = y$ " was assigned the certainty $M_{y|x}$ and the current content of $\Lambda(Y)$ is $\lambda(y)$, **then** put $\sum_y \lambda(y) M_{y|x}$ into $\Lambda(X)$.

Eqs. (4.8) through (4.10) are depicted in Figure 4.12. The reasons for formulating simple updating equations like Eqs. (4.5) and (4.7) as *inference rules* is to demonstrate their similarity to logical deductions. Like deductive rules of inference, they can be invoked at any time and in any order; postponing the activation of a rule or invoking it repeatedly may delay equilibrium but will not alter the final result. Like deductive rules of inference, the actions specified by Eqs. (4.8) through (4.10) are determined solely by the premisses, independent of the rest of the database. But these rules, unlike deductive rules, are non-monotonic in the sense that the conclusions (e.g., the belief measures $BEL(x)$) may undergo changes as new evidence arrives. Thus, Polya's aspirations of formulating patterns of plausible reasoning as rules of inference (see Section 2.3.1) are partially realized, and the pitfalls of his original patterns avoided, by distinguishing causal supports (π) from evidential supports (λ). Knowing that a neighboring proposition $Y = y$ has become "more credible" or "less credible" is insufficient to trigger an action in X ; we must first ascertain whether it is $\lambda(y)$ or $\pi(y)$ that has changed. This distinction is expressed in logical terms in Section 10.3.

The scheme described in Figures 4.9 through 4.12 requires that each processor gain access to matrices of both incoming and outgoing links. This may be inconvenient both for hardware implementation and rule-based programming. An alternative scheme, depicted in Figure 4.13, requires that each processor store just one matrix corresponding to the incoming link. Here, each processor receives as input the π of its parent and the λ of its own variable. Upon activation, each processor computes its own π (to be delivered to its child) and the λ of its parent (to be delivered to the parent). This convention will be used throughout the rest of this chapter because it closely reflects the basic construction of Bayesian networks, whereby each node is characterized by its relation to its parents.




Figure 4.12. Structure of individual processor, containing two link matrices.




Figure 4.13. Structure of individual processor, containing a single link matrix.

4.2.3 Propagation in Trees

We now examine a general tree-structured network where a node might have several children and one parent. The propagation scheme in trees is very similar to that of chains, with two distinctions: Each node must combine the impacts of λ -messages obtained from several children, and each node should distribute a separate π -message to each of its children.

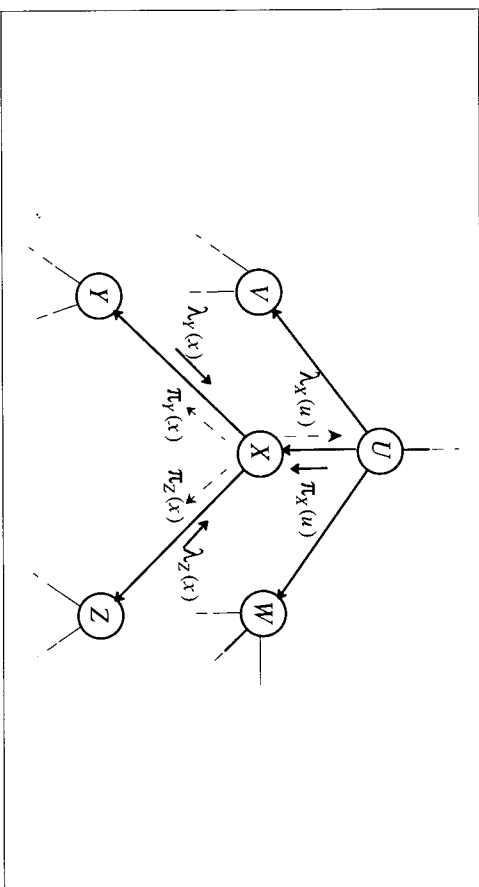


Figure 4.14. Fragment of causal tree, showing incoming (solid arrows) and outgoing (broken arrows) messages at node X.

Consider the tree fragment depicted in Figure 4.14. The belief in the various values of X depends on two distinct sets of evidence: evidence from the sub-tree rooted at X , and evidence from the rest of the tree. But the influence of the latter source of information on X is completely summarized by its effect on U , since U separates X from all variables except X 's descendants. More formally, letting $e_{\bar{x}}$ stand for the evidence contained in the tree rooted at X and letting $e_x^{\bar{x}}$ stand for the evidence contained in the rest of the network, we have

$$P(x|u, e_x^{\bar{x}}) = P(x|u). \tag{4.11}$$

This also leads to the usual conditional independence among siblings,

$$P(x, v|u) = P(x|u)P(v|u), \tag{4.12}$$

since the proposition $V = v$ is part of $e_x^{\bar{x}}$.

DATA FUSION

Assume we wish to find the belief induced on X by some evidence $e = e_{\bar{x}} \cup e_x^{\bar{x}}$. Bayes' Rule, together with Eq. (4.11), yields the product rule

$$\begin{aligned} BEL(x) &= P(x|e_x^{\bar{x}}, e_{\bar{x}}) = \alpha P(e_{\bar{x}}|e_x^{\bar{x}}, x)P(x|e_x^{\bar{x}}) \\ &= \alpha P(e_{\bar{x}}|x)P(x|e_x^{\bar{x}}), \end{aligned} \tag{4.13}$$

where $\alpha = [P(e_{\bar{x}}|e_x^{\bar{x}})]^{-1}$ is a normalizing constant.

Eq. (4.13) provides an interesting generalization of the celebrated Bayes product formula,

$$P(x|e) = \alpha P(e|x) P(x), \quad (4.14)$$

by identifying a surrogate— $P(x|e_x^+)$ —for the prior probability term $P(x)$ —with every intermediate node in the tree. In recursive Bayesian updating (see Section 2.1.4), the posterior probability can be used as a new prior, relative to the next item of evidence, only when the items of evidence are conditionally independent, given the updated variable X . Such recursive updating cannot be applied to networks because only variables that are separated from each other by X are conditionally independent. In general, it is not permissible to use the total posterior belief, updated by Eq. (4.13), as a new multiplicative prior for the calculation. Eq. (4.13) is significant because it shows that a product rule analogous to Eq. (4.14) can be applied recursively to any node in the tree, even when the observations are not conditionally independent, but the recursive, multiplicative role of the prior probability has been taken over by that portion of belief contributed by evidence from the sub-tree above the updated variable, excluding the data collected from its descendants. The root is the only node that requires a prior probability estimation, and since it has no network above it, e_{root}^+ should be interpreted as the background knowledge remaining unexplicated.

Eq. (4.13) suggests that the probability distribution of every variable in the tree can be computed if the node corresponding to that variable contains the vectors

$$\lambda(x) = P(e_x^- | x) \quad (4.15)$$

and

$$\pi(x) = P(x | e_x^+). \quad (4.16)$$

Here, $\pi(x)$ represents the causal or *predictive* support attributed to the assertion " $X = x$ " by all non-descendants of X , mediated by X 's parent, and $\lambda(x)$ represents the diagnostic or *retrospective* support that " $X = x$ " receives from X 's descendants. The total strength of belief in " $X = x$ " can be obtained by *fusing* these two supports via the product

$$BEL(x) = \alpha \lambda(x) \pi(x). \quad (4.17)$$

To see how information from several descendants fuses at node X , we partition the data set e_x^- in Eq. (4.15) into disjoint subsets, one for each child of X . Referring to Figure 4.14, for example, the tree rooted at X can be partitioned into the root, X ,

and two sub-trees, one rooted at Y and the other at Z . Thus, if X itself is not instantiated, we can write $e_x^- = e_y^- \cup e_z^-$, and since X separates its children, we have

$$\begin{aligned} \lambda(x) &= P(e_x^- | x) \\ &= P(e_y^-, e_z^- | x) \\ &= P(e_y^- | x) P(e_z^- | x). \end{aligned} \quad (4.18)$$

So $\lambda(x)$ can be formed as a product of terms such as $P(e_y^- | x)$, if these terms are delivered to X as messages from its children. Denoting these messages by subscripted λ 's,

$$\lambda_y(x) = P(e_y^- | x) \quad (4.19a)$$

and

$$\lambda_z(x) = P(e_z^- | x), \quad (4.19b)$$

we have the product rule:

$$\lambda(x) = \lambda_y(x) \lambda_z(x). \quad (4.20)$$

This product rule also applies when X itself is instantiated ($X = x$) if we model the new data by adding to X a dummy child D that delivers the message

$$\lambda_D(x) = \delta_{x,x'} = \begin{cases} 1 & \text{if } x = x' \\ 0 & \text{if } x \neq x'. \end{cases}$$

EXAMPLE 6: In the fingerprint story of Example 5, imagine that we receive reports from two independent laboratories, $\lambda_{Z_1}(y) = \beta$ (0.80, 0.60, 0.50) and $\lambda_{Z_2}(y) = \beta$ (0.30, 0.50, 0.90). The overall diagnostic support $\lambda(y)$ attributable to the three possible values of Y is

$$\lambda(y) = \beta(0.80, 0.60, 0.50) (0.30, 0.50, 0.90) = \beta(0.24, 0.30, 0.45),$$

and this, combined with the previous causal support $\pi(y) = (0.30, 0.35, 0.35)$, yields an overall belief of

$$\begin{aligned} BEL(y) &= \alpha(0.24, 0.30, 0.45) (0.30, 0.35, 0.35) \\ &= (0.215, 0.314, 0.471). \end{aligned}$$

What happens if Suspect 2 confesses, reliably, that he was the last weapon holder? We model this confession as a third report $\lambda_{Z_3}(y) = (0, 1, 0)$ which, by the product rule of Eq. (4.20), completely overrides the other two and yields $\lambda(y) = \beta(0, 1, 0)$ and $BEL(y) = (0, 1, 0)$.

Now we shall see if X can compute its $\pi(x)$ vector from information available at its parent U (see Figure 4.14).

Conditioning on the values of U we get

$$\begin{aligned}\pi(x) &= P(x | e_x^+) \\ &= \sum_u P(x | e_x^+, u) P(u | e_x^+) \\ &= \sum_u P(x | u) P(u | e_x^+).\end{aligned}$$

$P(x | u)$ is the matrix stored on the link $U \rightarrow X$, and $P(u | e_x^+)$ can be calculated by U and delivered to X as the message

$$\pi_X(u) = P(u | e_x^+), \quad (4.21)$$

yielding

$$\pi(x) = \sum_u P(x | u) \pi_X(u) = M_{x|u} \bullet \pi_X(u). \quad (4.22)$$

Substituting Eqs. (4.20) and (4.22) in Eq. (4.17) we have

$$BEL(x) = \alpha \lambda_Y(x) \lambda_Z(x) \sum_u P(x | u) \pi_X(u). \quad (4.23)$$

Thus, node X can calculate its own beliefs if it has received the messages $\lambda_Y(x)$ and $\lambda_Z(x)$ from its children Y and Z and the message $\pi_X(u)$ from its parent U .

PROPAGATION MECHANISM

Our next task is to determine how the influence of new information will spread through the network. In other words, we imagine that each node eventually receives from its neighbors the π - λ messages needed to calculate its own belief, as in Eq. (4.23), and we must determine how that node calculates the π - λ messages that its neighbors expect to receive from it. If the calculations can be accomplished by local computations, and if we let each node perform the calculations often enough, then the proper belief distributions are guaranteed to be reached, eventually, at every node.

Consider first the message $\lambda_X(u) \triangleq P(e_{\bar{x}} | u)$ that node X must send to its parent U (see Figure 4.14). If the value of X is known, say $X = x'$, then $\lambda_X(u)$ reduces to

$P(x' | u)$, which is column x' of the matrix $M_{x|u}$. If X is not known, we condition $P(e_{\bar{x}} | u)$ on $X = x$ and get

$$\begin{aligned}\lambda_X(u) &= \sum_x P(e_{\bar{x}} | u, x) P(x | u) \\ &= \sum_x P(e_{\bar{x}} | x) P(x | u) \\ &= \sum_x \lambda(x) P(x | u) \\ &= M_{x|u} \bullet \lambda(x).\end{aligned} \quad (4.24)$$

Thus, the message going to the parent U can be calculated from the messages received from the children and the matrix stored on the link from the parent. Note that Eq. (4.24) also holds if X itself is instantiated (say to $X = x'$) because in such a case $\lambda(x) = \delta_{x,x'}$, and Eq. (4.24) yields $\lambda_X(u) = P(x' | u)$ as required.

Now, consider the message that node X should send to one of its children, say Y :

$$\pi_Y(x) = P(x | e_y^+) = P(x | e_x^+, e_z^-).$$

Using Bayes' Rule, we get

$$\begin{aligned}\pi_Y(x) &= \alpha P(e_z^- | x, e_x^+) P(x | e_x^+) \\ &= \alpha P(e_z^- | x) P(x | e_x^+) \\ &= \alpha \lambda_Z(x) \pi(x) \\ &= \alpha \lambda_Z(x) \sum_u P(x | u) \pi_X(u).\end{aligned} \quad (4.25)$$

The second equality follows from the fact that X separates e_z^- from e_x^+ , the third equality follows from the definition of $\pi(x)$ (Eq. (4.16)), and the fourth follows from Eq. (4.22). Thus, the message sent from X to Y is calculated using the message it receives from its *other child* Z (in general, the messages it receives from all its children, except Y) and the message X receives from its parent U . This is precisely how double-counting of evidence is prevented.

Figure 4.15 summarizes the calculations for node X .

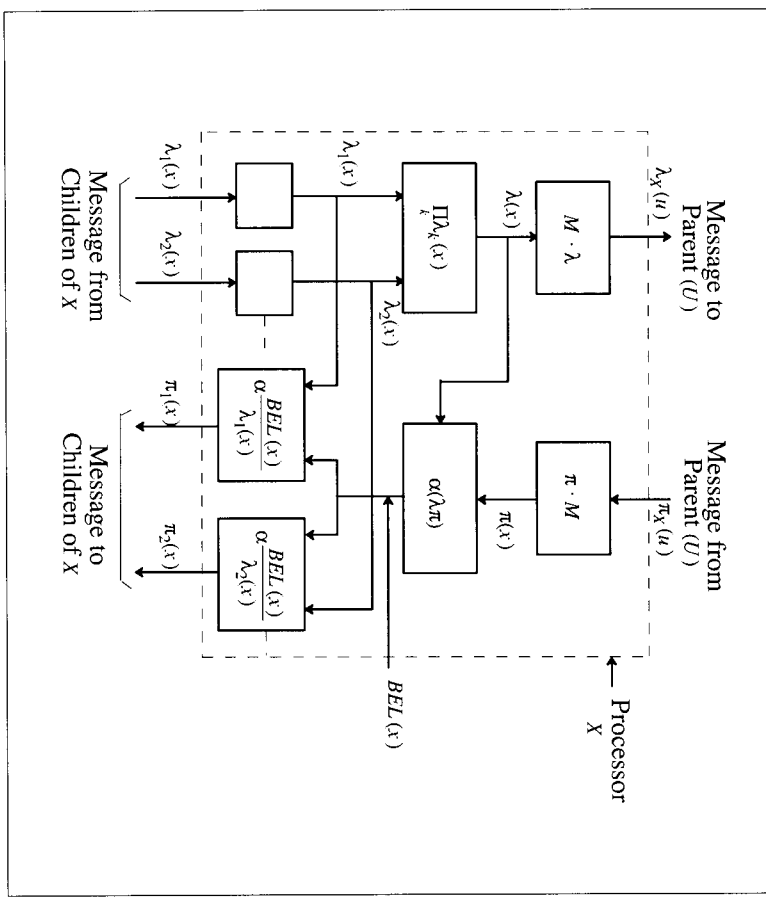


Figure 4.15. The internal structure of a single processor performing belief updating for the variable X .

By Eqs. (4.23) and (4.25),

$$\pi_Y(x) = \alpha \frac{BEL(x)}{\lambda_Y(x)}, \tag{4.26}$$

so it might be advantageous for node X , instead of sending each child a different message, to send all its children the value of its current belief, $BEL(x)$, and let each child recover its respective π message by dividing $BEL(x)$ by the value of the last message sent to X (caution should be exercised whenever $\lambda(x)$ is zero or is very close to zero). There is no need, of course, to normalize the π messages prior to transmission; only the $BEL(\cdot)$ expressions require normalization. The sole purpose of the normalization constant α in Eqs. (4.25) and (4.26) is to preserve the probabilistic meaning of these messages. It is a good engineering practice to encode the π and λ messages so that the smallest component of each will attain the value 1.

SUMMARY OF PROPAGATION RULES FOR TREES

We shall now summarize the steps involved in tree propagation by specifying the activities of a typical node X having m children, Y_1, Y_2, \dots, Y_m , and a parent U . The belief distribution of variable X can be computed if three types of parameters are made available:

1. The current strength of the causal support, $\pi_X(u)$, contributed by the parent of X ,

$$\pi_X(u) = P(u | e_X^+).$$

2. The current strength of the diagnostic support, $\lambda_{Y_j}(x)$, contributed by the j -th child of X ,

$$\lambda_{Y_j}(x) = P(e_{Y_j}^- | x).$$

3. The fixed conditional probability matrix $P(x | u)$ that relates the variable X to its immediate parent U .

Using these parameters, local belief updating can be accomplished in three steps, to be executed in any order.

Step 1—Belief updating: When node X is activated to update its parameters, it simultaneously inspects the $\pi_X(u)$ message communicated by its parent and the messages $\lambda_{Y_1}(x), \lambda_{Y_2}(x), \dots$ communicated by each of its children. Using this input, it updates its belief measure to

$$BEL(x) = \alpha \lambda(x) \pi(x), \tag{4.27a}$$

where

$$\lambda(x) = \prod_j \lambda_{Y_j}(x), \tag{4.27b}$$

$$\pi(x) = \sum_u P(x | u) \pi_X(u), \tag{4.27c}$$

and α is a normalizing constant rendering $\sum_x BEL(x) = 1$.

Step 2—Bottom-up propagation: Using the λ messages received, node X computes a new message, $\lambda_X(u)$, which is sent to its parent U :

$$\lambda_X(u) = \sum_x \lambda(x) P(x | u). \tag{4.28}$$

Step 3—Top-down propagation: X computes new π messages to be sent to each of its children. For example, the new $\pi_{Y_j}(x)$ message that X sends to its j -th child Y_j is computed by

$$\pi_{Y_j}(x) = \alpha \pi(x) \prod_{k \neq j} \lambda_{Y_k}(x). \quad (4.29)$$

The computations in Eqs. (4.27), (4.28), and (4.29) preserve the probabilistic meaning of the parameters. In particular,

$$\lambda_X(u) = P(e_{\bar{X}} | u), \quad (4.30)$$

$$\pi_Y(x) = P(x | e_Y^+), \quad (4.31)$$

$$BEL(x) = P(x | e). \quad (4.32)$$

Terminal and evidence nodes in the tree require special treatment. We must distinguish four cases:

1. *Anticipatory node*—a leaf node that has not been instantiated. For such variables, BEL should be equal to π , and we should therefore set $\lambda = (1, 1, \dots, 1)$.
2. *Evidence node*—a variable with instantiated value. Following Eq. (4.6a), if the j -th value of X is observed to be true, we set $\lambda(x) = (0, \dots, 0, 1, 0, \dots, 0)$ with 1 at the j -th position.
3. *Dummy node*—a node Y representing virtual or judgmental evidence bearing on X . We do not specify $\lambda(y)$ or $\pi(y)$ but instead post a $\lambda_Y(x)$ message to X , where $\lambda_Y(x) = \beta P(Observation | x)$, β being any convenient constant.
4. *Root node*—The boundary condition for the root node is established by setting $\pi(\text{root})$ equal to the prior probability of the root variable.

EXAMPLE 7: To illustrate these computations let us redo Example 5, using tree propagation on the network of Figure 4.16.

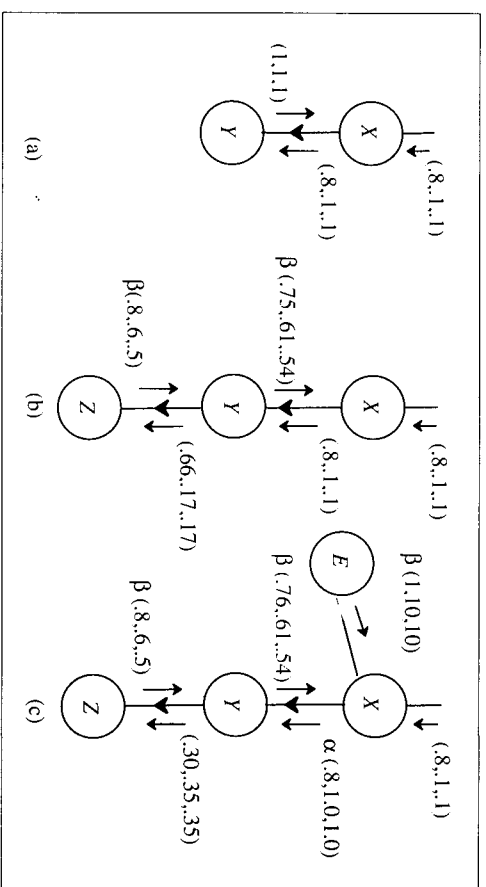


Figure 4.16. Belief updating in Example 7 using tree propagation. The alibi is modeled as a dummy node E generating a λ message $1:10:10$.

As before, let us assume that our belief in the identity of the killer, based on all testimony heard so far, amounts to $\pi(x) = (0.8, 0.1, 0.1)$. Before we obtain any fingerprint information, Figure 4.16a shows Y as an anticipatory node with $\lambda(y) = (1, 1, 1)$, which means $\lambda_Y(x) = \lambda(x) = (1, 1, 1)$ and $BEL(x) = \pi(x)$. $\pi(y)$ can be calculated from Eq. (4.22) (using $\pi_Y(x) = \pi(x)$), yielding

$$\begin{aligned} \pi(y) &= \pi_Y(x) \cdot M_{Y|X} = (0.8, 0.1, 0.1) \cdot \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0.1 & 0.8 & 0.1 \\ 0.1 & 0.1 & 0.8 \end{bmatrix} \\ &= (0.66, 0.17, 0.17) = BEL(y). \end{aligned}$$

Assume that a laboratory report arrives summarizing the test results (a piece of virtual evidence Z) by the message $\lambda_Z(y) = \lambda(y) = \beta(0.8, 0.6, 0.5)$, as in Figure 4.16b. Node Y updates its belief,

$$BEL(y) = \alpha \lambda(y) \pi(y) = \alpha(0.8, 0.6, 0.5)(0.66, 0.17, 0.17) = (0.738, 0.143, 0.119),$$

and based on Eq. (4.28), Y computes a new message, $\lambda_Y(x)$, for X :

$$\lambda_Y(x) = M_{Y|X} \cdot \lambda(y) = \beta \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0.1 & 0.8 & 0.1 \\ 0.1 & 0.1 & 0.8 \end{bmatrix} \cdot \begin{bmatrix} 0.8 \\ 0.6 \\ 0.5 \end{bmatrix} = \beta(0.75, 0.61, 0.54).$$

Upon receiving this message, node X sets $\lambda(x) = \lambda_Y(x)$ and recomputes its belief to

$$\begin{aligned} BEL(x) &= \alpha\lambda(x)\pi(x) = \alpha(0.75, 0.61, 0.54)(0.8, 0.1, 0.1) \\ &= (0.839, 0.085, 0.076). \end{aligned}$$

Now assume that Suspect 1 produces a very strong alibi, supporting his innocence ten times more than his guilt, i.e., $P(\text{Alibi} | X \neq 1) : P(\text{Alibi} | X = 1) = 10 : 1$. To fuse this information with all previous evidence, we link a new evidence node E directly to X and post the message $\lambda_E(x) = \beta(1, 10, 10)$ on the link (see Figure 4.16c). $\lambda_E(x)$ combines with $\lambda_Y(x)$ to yield

$$\begin{aligned} \lambda(x) &= \lambda_E(x) \lambda_Y(x) = \beta(0.75, 6.10, 5.40), \\ BEL(x) &= \alpha\lambda(x)\pi(x) = \alpha(0.75, 6.10, 5.40)(0.8, 0.1, 0.1) \\ &= (0.343, 0.349, 0.308), \end{aligned}$$

and generates the message $\pi_Y(x) = \alpha\lambda_E(x)\pi(x) = \alpha(0.8, 1.0, 1.0)$ for Y . Upon receiving $\pi_Y(x)$, processor Y updates its causal support $\pi(y)$ to (see Eq. 4.27)

$$\pi(y) = \pi_Y(x) \cdot M_{y,x} = \alpha(0.8, 1.0, 1.0) \cdot \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0.1 & 0.8 & 0.1 \\ 0.1 & 0.1 & 0.8 \end{bmatrix} = (0.30, 0.35, 0.35),$$

and $BEL(y)$ becomes

$$\begin{aligned} BEL(y) &= \alpha\lambda(y)\pi(y) = \alpha(0.8, 0.6, 0.5)(0.30, 0.35, 0.35) \\ &= (0.384, 0.336, 0.280). \end{aligned}$$

Finally, since Y has only one child— Z —Eq. (4.29) reduces to $\pi_Z(y) = \pi(y)$ (see also Eq. (4.26)). The purpose of propagating beliefs top-down to sensory nodes such as Z is twofold: to guide data-acquisition strategies toward the most informative sources (see Section 6.4) and to facilitate explanations for the system's choices.

Note that $BEL(x)$ cannot be taken as an updated prior of x for the purpose of calculating $BEL(y)$. In other words, it is wrong to update $BEL(y)$ via the textbook formula

$$BEL(y) = \sum_x P(y|x) BEL(x) \quad (4.33)$$

(see discussion of Jeffrey's Rule, Section 2.3.3), because $BEL(x)$ was affected by information transmitted from Y , and feeding this information back to Y would amount to counting the same evidence twice. Only the $\pi_Y(x)$ portion of $BEL(x)$ is fed back to Y ; it is based on evidence (E) not yet considered in $\lambda(x)$. Another way to view this is that once information is obtained from Z , the initial value of the link matrix $P(y|x)$ no longer represents the dependence between X and Y , so $P(y|x, z)$ should replace $P(y|x)$ in Eq. (4.33).

Note also that the activation steps need not be sequential but may be executed in parallel when several pieces of evidence arrive simultaneously. In the extreme case, we can

imagine that all processors are activated simultaneously by a common clock. For example, if the lab report and the alibi arrive together, then in the first clock cycle X and Y will simultaneously update their beliefs to

$$\begin{aligned} BEL(x) &= \alpha(0.08, 0.10, 0.10) = (0.286, 0.357, 0.357) \text{ and} \\ BEL(y) &= (0.738, 0.142, 0.111) \end{aligned}$$

and produce the messages

$$\pi_Y(x) = \alpha(0.08, 0.10, 0.10) \text{ and } \lambda_Y(x) = \beta(0.75, 0.61, 0.54),$$

respectively. In the second clock cycle, X and Y will simultaneously update their beliefs to

$$BEL(x) = (0.343, 0.349, 0.308) \text{ and } BEL(y) = (0.384, 0.336, 0.280)$$

and produce the same $\pi_Y(x)$ and $\lambda_Y(x)$ as before. From now on, the same beliefs and the same messages will be produced in every clock cycle unless additional evidence becomes available.

ILLUSTRATING THE FLOW OF BELIEF

Figure 4.17 shows six successive stages of belief propagation through a simple binary tree, assuming the updating is triggered by changes in the belief parameters of neighboring processors. Initially, the tree is in equilibrium, and all terminal nodes are anticipatory (Figure 4.17a). As soon as two data nodes are activated, white tokens are placed on the links from the nodes to their parents (Figure 4.17b). In the next phase, the parents, activated by these tokens, absorb them and manufacture enough tokens for their neighbors: white tokens for their parents and black ones for their children (Figure 4.17c). (The links from which the absorbed tokens originated do not receive new tokens because a π -message is not affected by a λ -message crossing the same link.) The root node now receives two white tokens, one from each of its descendants, triggering the production of two black tokens for top-down delivery (Figure 4.17d). The process continues in this fashion for six cycles, at which point all tokens are absorbed and the network reaches a new equilibrium. As soon as a leaf node posts a token for its parent, the leaf is ready to receive new data. If the new data arrives before the token was observed by the parent, a new token replaces the old one. In this fashion the inference network can track a changing environment and provide coherent interpretation of signals emanating simultaneously from multiple sources.




Figure 4.17. The impact of new data propagates through a tree by a message-passing process.

This updating scheme has the following properties:

1. The local computations it requires are efficient in both storage space and time. For a tree with m children per parent and n values per node, each processor should store $n^2 + mn + 2n$ real numbers and perform $2n^2 + mn + 2n$ multiplications per update.
2. The local computations and the final belief distribution are entirely independent of the control mechanism that activates the individual operations. These operations can be activated by either data-driven or goal-driven (e.g., requests for evidence) control strategies, by a central clock, or at random.
3. New information is diffused through the network in a single pass. Instabilities and indefinite relaxations are eliminated by maintaining a two-parameter system (π and λ) to decouple causal support from diagnostic support. The time required for completing the diffusion (in parallel) is proportional to the diameter of the network.

4.3 BELIEF PROPAGATION IN CAUSAL POLYTREES (SINGLY CONNECTED NETWORKS)

The tree structures treated in the preceding section require that exactly one variable be considered a cause of another given variable. This restriction simplifies computations, but its representational power is rather limited, since it forces us to form a single node from all causes sharing a common consequence. By contrast, when people see many potential causes for a given observation, they weigh one cause against another as independent variables, each pointing to a specialized area of knowledge. As an illustration, consider the situation discussed in Example 7 of Chapter 2:

Mr. Holmes receives a phone call at work from his neighbor notifying him that she heard a burglar alarm sound from the direction of his home. As he is preparing to rush home, Mr. Holmes recalls that the alarm recently was triggered by an earthquake. Driving home, he hears a radio newscast reporting an earthquake 200 miles away.

Mr. Holmes perceives two episodes as potential causes for the alarm sound—an attempted burglary and an earthquake. Even though burglaries can safely be assumed independent of earthquakes, the radio announcement still reduces the likelihood of a burglary, as it “explains away” the alarm sound. Moreover, the two causal events are perceived as individual variables pointing to separate frames of knowledge (crime-related information seldom evokes associations of earthquakes), so it would be unnatural to lump the two events together into a single node.

Treating $E = \text{Earthquake}$ and $B = \text{Burglary}$ as two separate entities (as in Figure 2.2) allows us to relate each of them to a separate set of evidence without consulting the other. For example, if $R = \text{The radio announcement}$ and $S = \text{The alarm sound}$, we can quantify the relation between E and R by the probabilities $P(R|E)$ without having to consider the irrelevant event of burglary, as would be required if the pair (E, B) were combined into one variable. Moreover, if R is confirmed, a natural way to update the beliefs of E and B would be in two separate steps, mediated by updating S . E and B are presumed to be independent unless evidence supporting S is obtained (e.g., the neighbor’s phone call); when this happens, E and B find themselves competing for a fixed amount of evidential support—information favoring one explanation (e.g., the radio report) would weaken the other explanation by undermining its connection with the mediator S .

This competitive interplay among multiple explanations is a prevailing feature of human reasoning and has been discussed in previous chapters (see Sections 1.2.2 and 2.2.4). When a physician discovers evidence in favor of one disease, it reduces the likelihood of other diseases that could explain the patient’s symptoms,

although the patient might well be suffering from two or more disorders simultaneously. When we find our driveway wet, the discovery that the sprinkler was on all night weakens the likelihood that it rained at night. The same maxim also governs the interplay of other frame-like (though not necessarily causal) explanations. For example, the sentence "Tweedy tasted wonderful" provides a clue that Tweedy, the celebrated non-flying bird from the AI literature (see Chapter 10), is not a penguin after all; a more likely explanation for Tweedy's reluctance to fly is that she is broiled.

This section extends our propagation scheme to Bayesian networks where a node may have multiple parents, thus permitting "sideways" interactions via common successors. The networks are, however, required to be *singly connected*, namely, no more than one path exists between any two nodes, as in Figure 4.18a. We call such networks *causal polytrees* because they can be viewed as a collection of several causal trees fused together at the nodes where arrows converge head to head. The absence of loops in the underlying network permits us to develop a local updating scheme similar to that used for causal trees. The derivation of the propagation rules, likewise, will correspond to the derivation of Eqs. (4.27) through (4.29). The impatient reader is advised to skip directly to the Summary, Eqs. (4.47) through (4.53).




Figure 4.18. (a) A fragment of a polytree and (b) the parents and children of a typical node X.

4.3.1 Propagation Rules

Consider a typical fragment of a singly connected network (Figure 4.18b), consisting of a node X, the set of all X's parents, $U = \{U_1, \dots, U_n\}$, and the set of all X's children, $Y = \{Y_1, \dots, Y_m\}$. As before, let e be the total evidence obtained, $e_{\bar{X}}$ be the evidence connected to X through its children (Y), and e_X^+ be the evidence connected to X through its parents (U), so that

$$\begin{aligned} BEL(x) &= \alpha P(e_{\bar{X}} | x) P(x | e_X^+) \\ &= \alpha \lambda(x) \pi(x). \end{aligned} \quad (4.34)$$

$e_{\bar{X}}$ and e_X^+ can be further decomposed into

$$e_{\bar{X}} = \left\{ e_{\bar{X}Y_1}, \dots, e_{\bar{X}Y_m} \right\}$$

and

$$e_X^+ = \left\{ e_{U_1X}^+, \dots, e_{U_nX}^+ \right\},$$

where $e_{\bar{X}Y_j}$ stands for evidence contained in the subnetwork on the *head* side of the link $X \rightarrow Y_j$, and $e_{U_jX}^+$ stands for evidence contained in the subnetwork on the *tail* side of the link $U_j \rightarrow X$.

To avoid cumbersome notation, we will treat all evidence as virtual, i.e., obtained from dummy children of variables whose values are known. Thus, all instantiated nodes in the networks are assumed to be leaf nodes. Now,

$$\begin{aligned} \lambda(x) &\triangleq P(e_{\bar{X}} | x) \\ &= P(e_{\bar{X}Y_1}, \dots, e_{\bar{X}Y_m} | x) \\ &= P(e_{\bar{X}Y_1} | x) \cdot P(e_{\bar{X}Y_2} | x) \cdots P(e_{\bar{X}Y_m} | x) \\ &= \prod_{j=1}^m \lambda_{Y_j}(x), \end{aligned} \quad (4.35)$$

where

$$\lambda_{Y_j}(x) = P(e_{\bar{X}Y_j} | x). \quad (4.36)$$

Also,

$$\begin{aligned} \pi(x) &\triangleq P(x|e_x^+) \\ &= P(x|e_{U_1, X}^+, \dots, e_{U_n, X}^+) \\ &= \sum_{u_1, \dots, u_n} P(x|u_1, \dots, u_n) P(u_1, \dots, u_n | e_{U_1, X}^+, \dots, e_{U_n, X}^+) \\ &= \sum_{u_1, \dots, u_n} P(x|u_1, \dots, u_n) P(u_1 | e_{U_1, X}^+) P(u_2 | e_{U_2, X}^+) \dots P(u_n | e_{U_n, X}^+), \end{aligned}$$

because each pair $\{U_i, e_{U_i, X}^+\}$ is independent of the other U_i 's and their evidence sets.

Let

$$\pi_X(u_i) = P(u_i | e_{U_i, X}^+). \quad (4.37)$$

Then

$$\begin{aligned} \pi(x) &= \sum_{u_1, \dots, u_n} P(x|u_1, \dots, u_n) \pi_X(u_1) \pi_X(u_2) \dots \pi_X(u_n) \\ &= \sum_u P(x|u) \prod_{i=1}^n \pi_X(u_i). \end{aligned} \quad (4.38)$$

Substituting Eq. (4.36) and Eq. (4.38) in Eq. (4.34), we get

$$BEL(x) = \alpha \left[\prod_{j=1}^m \lambda_{Y_j}(x) \right] \left[\sum_u P(x|u) \prod_{i=1}^n \pi_X(u_i) \right]. \quad (4.39)$$

Thus, node X can calculate its own beliefs if it receives the messages $\lambda_{Y_j}(x)$ from its children and $\pi_X(u_i)$ from its parents.

To prescribe how the influence of new information will spread through the network, we need to specify how a typical node, say X , will compute its outgoing messages $\lambda_X(u_i)$, $i = 1, \dots, n$, and $\pi_X(x)$, $j = 1, \dots, m$, from the incoming messages $\lambda_{Y_j}(x)$, $j = 1, \dots, m$ and $\pi_X(u_i)$, $i = 1, \dots, n$.

UPDATING λ

Consider the message $\lambda_X(u_i)$, which node X must send its parent U_i . By Eq. (4.36),

$$\lambda_X(u_i) = P(e_{U_i, X}^+ | u_i). \quad (4.40)$$

In deriving $\lambda_X(u_i)$ it is convenient temporarily to treat all parents (except U_i) as a single compound variable,

$$V = U - U_i = \{U_1, \dots, U_{i-1}, U_{i+1}, \dots, U_n\}, \quad (4.41)$$

connected to X via a single link $V \rightarrow X$, as in Figure 4.19.




Figure 4.19. Variables, messages, and evidence sets used in the derivation of $\lambda_X(u_i)$.

The evidence $e_{U_i, X}^+$ governing $\lambda_X(u_i)$ can now be decomposed into two components:

$$e_{U_i, X}^+ = \{e_{V, X}^+, e_{\bar{X}}^+\} \quad (4.42)$$

where

$$e_{V, X}^+ = \bigcup_{k \neq i} e_{U_k, X}^+. \quad (4.43)$$

Substituting Eq. (4.42) into Eq. (4.40) gives

$$\begin{aligned}
 \lambda_X(u_i) &= P(e_{VX}^+, e_{\bar{X}}^- | u_i) \\
 &= \sum_{x \ v} P(e_{VX}^+, e_{\bar{X}}^- | u_i, v, x) P(v, x | u_i) \\
 &\quad \text{(conditioning on } x \text{ and } v) \\
 &= \sum_{x \ v} P(e_{\bar{X}}^- | x) P(e_{VX}^+ | v) P(v, x | u_i) \\
 &\quad \text{(since } X \text{ separates } e_{VX}^+ \text{ from } e_{\bar{X}}^- \text{ and } V \text{ separates } e_{VX}^+ \text{ from } U_i) \\
 &= \beta \sum_{x \ v} P(e_{\bar{X}}^- | x) \frac{P(v | e_{VX}^+)}{P(v)} P(x | v, u_i) P(v | u_i) \\
 &\quad \text{(by Bayes' Rule)} \\
 &= \beta \sum_{x \ v} P(e_{\bar{X}}^- | x) P(v | e_{VX}^+) P(x | v, u_i) \\
 &\quad \text{(since } U_i \text{ and } V \text{ are marginally independent)}
 \end{aligned}$$

Restoring the meaning of V from Eqs. (4.41) and (4.43), we have

$$\begin{aligned}
 P(x | v, u_i) &= P(x | u), \\
 P(v | e_{VX}^+) &= \prod_{k \neq i} P(u_k | e_{VX}^+) = \prod_{k \neq i} \pi_X(u_k),
 \end{aligned}$$

and $\lambda_X(u_i)$ becomes

$$\lambda_X(u_i) = \beta \sum_x \lambda(x) \sum_{u_k: k \neq i} P(x | u) \prod_{k \neq i} \pi_X(u_k), \quad (4.44)$$

where $\lambda(x)$ is given in Eq. (4.35). As before, $\lambda_X(u_i)$ is not affected by $\pi_X(u_i)$; the two messages pass along the same arc (in the opposite direction) without interacting.

In the derivation above we assumed that X itself is not instantiated and therefore is not part of the evidence set $e_{U\bar{X}}$ in Eq. (4.42). This assumption does not affect the generality of Eq. (4.44) because the fact that X is an evidence node attaining the value x' can always be represented by instantiating a (dummy) child node Z , and thus delivering a message $\lambda_Z(x')$ to X , where

$$\lambda_Z(x) = \delta_{x,x'} = \begin{cases} 1 & x = x' \\ 0 & x \neq x' \end{cases}$$

UPDATING π

Consider the message $\pi_Y(x)$, which node X must send to its child Y . By Eq. (4.37),

$$\pi_Y(x) = P(x | e_{XY}^+).$$

e_{XY}^+ stands for the entire body of evidence e , excluding evidence found in the subnetwork on the head side of the link $X \rightarrow Y$,

$$e_{XY}^+ = e - e_{XY}^-.$$

Thus, $\pi_Y(x)$ is equal to $BEL(x)$ when the evidence e_{XY}^- is suppressed. Equivalently, the expression for $\pi_Y(x)$ can be obtained from $BEL(x)$, Eq. (4.39), setting $\lambda_Y(x) = 1$. This leads to

$$\pi_Y(x) = \alpha \prod_{k \neq j} \lambda_{Y_k}(x) \pi(x), \quad (4.45)$$

where $\pi(x)$ is given in Eq. (4.38). Alternatively, $\pi_Y(x)$ can be obtained from $BEL(x)$, writing

$$\pi_Y(x) = \alpha \frac{BEL(x)}{\lambda_Y(x)} \Big|_{\lambda_Y(x) = 1}. \quad (4.46)$$

BOUNDARY CONDITIONS

The boundary conditions are established as follows:

1. *Root nodes:* If X is a node with no parents, we set $\pi(x)$ equal to the prior probability $P(x)$.
2. *Anticipatory nodes:* If X is a childless node that has not been instantiated, we set $\lambda(x) = (1, 1, \dots, 1)$.
3. *Evidence nodes:* If evidence $X=x'$ is obtained (X being any node in the network, not necessarily a leaf node), we set $\lambda(x) = \delta_{x,x'} = (0, \dots, 0, 1, 0, \dots, 0)$ with 1 at the x' -th position. Alternatively, the evidence $X=x'$ can be simulated by adding to X the auxiliary child node Z , directing a message $\lambda_Z(x) = \delta_{x,x'}$ toward X .

Eqs. (4.44) and (4.45) demonstrate that the outgoing messages, $\pi_Y(x)$ and $\lambda_Y(u_i)$, are determined from the incoming messages available to X . They also demonstrate that perturbation of the causal parameter π does not affect the diagnostic parameter λ on the same link, and vice versa. The two are orthogonal to each other since

they depend on two disjoint sets of data. Therefore, any perturbation of beliefs in response to new evidence propagates through the network and is absorbed at peripheral nodes without reflection.[†] A new equilibrium state will be reached after a finite number of updates—if they are conducted in parallel, the number will be equal to the diameter of the network.

Eq. (4.44) also reveals that if no data are observed below X (i.e., all λ s pointing to X are unit vectors), then all λ s emanating from X are unit vectors, regardless of the incoming π messages. So, evidence gathered at a particular node does not influence any of its spouses until their common child gathers diagnostic support. This reflects the d -separation conditions established in Section 3.3.2 and matches our intuition regarding multiple causes. In Mr. Holmes' case, for example, prior to the neighbor's telephone call, seismic data indicating an earthquake would *not* have influenced the likelihood of a burglary.

SUMMARY OF PROPAGATION RULES FOR POLYTREES

The steps involved in polytree propagation are similar to those used with trees. We shall now summarize these steps by considering a typical node X having m children, Y_1, \dots, Y_m , and n parents, U_1, \dots, U_n , as in Figure 4.18*b*.

The belief distribution of variable X can be computed if three types of parameters are made available:

1. The current strength of the causal support π contributed by each incoming link $U_j \rightarrow X$:

$$\pi_X(u_j) = P(u_j | e_{U_j}^+ | X). \quad (4.47)$$

2. The current strength of the diagnostic support, λ , contributed by each outgoing link $X \rightarrow Y_j$:

$$\lambda_{Y_j}(X) = P(e_{Y_j}^- | X). \quad (4.48)$$

3. The fixed conditional-probability matrix $P(X | u_1, \dots, u_n)$ that relates the variable X to its immediate parents.

Using these parameters, local belief updating can be accomplished in three steps, to be executed in any order.

Step 1—Belief updating: When node X is activated, it simultaneously inspects the messages $\pi_X(u_i)$, $i = 1, \dots, n$ communicated by its parents and the messages

$\lambda_{Y_j}(X)$, $j = 1, \dots, m$ communicated by its children. Using this input, it updates its belief measure to

$$BEL(X) = \alpha \lambda(X) \pi(X), \quad (4.49)$$

where

$$\lambda(X) = \prod_j \lambda_{Y_j}(X), \quad (4.50)$$

$$\pi(X) = \sum_{u_1, \dots, u_n} P(X | u_1, \dots, u_n) \prod_i \pi_X(u_i), \quad (4.51)$$

and α is a normalizing constant rendering $\sum_X BEL(X) = 1$.

Step 2—Bottom-up propagation: Using the messages received, node X computes new λ messages to be sent to its parents. For example, the new message $\lambda_{Y_k}(u_i)$ that X sends to its parents U_i is computed by

$$\lambda_{Y_k}(u_i) = \beta \sum_X \lambda(X) \sum_{u_j: k \neq j} P(X | u_1, \dots, u_n) \prod_{k \neq i} \pi_X(u_k). \quad (4.52)$$

Step 3—Top-down propagation: Each node computes new π messages to be sent to its children. For example, the new $\pi_{Y_j}(X)$ message that X sends to its child Y_j is computed by

$$\begin{aligned} \pi_{Y_j}(X) &= \alpha \left[\prod_{k \neq j} \lambda_{Y_k}(X) \right] \sum_{u_1, \dots, u_m} P(X | u_1, \dots, u_n) \prod_i \pi_X(u_i) \\ &= \alpha \frac{BEL(X)}{\lambda_{Y_j}(X)}. \end{aligned} \quad (4.53)$$

Numerical examples illustrating these propagation rules in polytrees are given in Section 4.5.3 (see also Section 5.2.1).

The presence of multiple parents introduces an added dimension of complexity; although the computations in Eqs. (4.51) through (4.53) are still local, the summation ranges over all value combinations of the parent variables, which are exponential in n . If the number of parents is small, the summation can be performed by enumeration. However, if there are more than four or five parents, approximation techniques must be invoked that make use of the special structure of the link matrix $P(X | u_1, \dots, u_n)$.

In Chapter 2, we remarked that when n is large, we must use prototypical models of child-to-parents interaction, involving disjunctions and conjunctions (noisy OR- and AND-gates), in order to specify the link matrix. Here we see that

[†] A peripheral node is either a root with a single child or a leaf with a single parent. Every polytree must have at least two peripheral nodes.

considerations of computational complexity also dictate such usage. In the next section, we shall see that some of these prototypical models enable us to compute Eqs. (4.51) through (4.53) in closed form, thus reducing the computation of the outgoing messages to simple products of about n parameters.

4.3.2 Canonical Models of Multicausal Interactions

The original formulation of Bayesian networks (Section 3.4) required, for each variable X , that one assess the conditional probability $P(x|u)$, where U is a set of variables judged to be direct causes of X . We have also noticed that the individual elements of U often point to disparate frames of knowledge, making it difficult to assess $P(x|u)$. If U_1 and U_2 represent frames of knowledge truly foreign to one another (e.g., burglaries and earthquakes) except that they share X , as a common consequence (e.g., triggering the alarm), one cannot expect to find the matrix $P(x|u_1, u_2)$ prestored in memory. No reasoning system could spare the space required to permanently store the strength of connection between every conceivable event and every combination of conditions that might trigger that event. A more reasonable organization scheme would be to let each frame hold separately the weights associated with each of its likely consequences. Then, should a situation evoke more than one frame, the system can compute the weights of their common results "on the fly," using some *universal* combination rule. For example, in our burglary alarm scenario, it is reasonable to expect people to have some prestored idea of the likelihood that a burglary will trigger an alarm or that an earthquake will trigger it, but not of the likelihood that the combination, burglary and earthquake, will trigger it. A physician is expected to have prepackaged estimates of the chances that an individual disease will be accompanied by high fever, but when asked to estimate the likelihood of high fever given some rare combination of diseases, the physician cannot refer to prestored knowledge. Rather, she must resort to some canonical model of disease combinations in general, which most likely is also domain-independent, i.e., applicable to a wide class of interacting causes.

DISJUNCTIVE INTERACTION (THE NOISY OR-GATE)

One of the most common models of this type is *disjunctive interaction* (the "noisy OR-gate" of Chapter 2), which leads to a very convenient and widely applicable rule of combination. Disjunctive interaction occurs when any member of a set of conditions is likely to cause a certain event and this likelihood does not diminish when several of these conditions prevail simultaneously. For example, if each individual disease is likely to cause high fever, then a patient suffering from several of these diseases simultaneously would only be *more* likely to develop high

fever. Moreover, if the patient is also suffering from a disease that in isolation is not normally accompanied by high fever, this added information does not reduce the patient's likelihood of developing high fever from the other diseases.

Disjunctive interactions can be approximated by an elegant mathematical model based on two assumptions: accountability and exception independence.

Accountability requires that an event E be presumed false (i.e., $P(E) = 0$) if all conditions listed as causes of E are false. In the burglary alarm example, this assumption requires that we list explicitly the main conditions likely to trigger the alarm and lump together all those conditions that we prefer to keep implicit under the heading "All other causes."

Exception independence asserts that if an event E is a typical consequence of either one of two causal conditions c_1 and c_2 , then the mechanism that may inhibit the occurrence of E under c_1 is independent of the mechanism that may inhibit E under c_2 . Each exception to normal behavior acts as an independent variable. For example, the mechanism that inhibits the activation of the alarm during an earthquake could be low vertical acceleration, while the mechanism acting during a burglary could be the burglar's skill and sophistication. Since these two can safely be presumed independent of each other, exception independence holds. A power failure, on the other hand, would inhibit the alarm activation in both frames and thus would violate exception independence if it is a likely event.

These two assumptions are represented schematically in Figure 4.20. The event X represents a prediction or a consequence and is viewed as the output of a logical OR-gate. Each input to the OR-gate is the output of an AND-gate representing the conjunction of a causal explanation of X , U_i , and the negation of its specific inhibitory mechanism I_i . (Readers familiar with the IN and OUT justifiers in truth maintenance systems [Doyle 1979] or with the *abnormal* predicate in circumscription [McCarthy 1986] should note their resemblance to the U_i and I_i variables in Figure 4.20.)

The inputs $U = (U_1, U_2, \dots, U_n)$ are the parents of X in the Bayesian networks, and they normally represent explanations, hypotheses, conjectures, causal factors, or enabling conditions that may account for the occurrence of X . The inhibitors I_1, \dots, I_n represent exceptions or abnormalities that interfere with the normal relationship between U and X . These are normally not represented by nodes in Bayesian networks but are summarized implicitly by the link matrix $P(x|u_1, \dots, u_n)$. Here we explicate their structure and display them as root nodes in order to justify an especially useful form of the link matrix.

CONSTRUCTING THE LINK MATRIX

Denote by q_k the probability that the k -th inhibitor is active. If U_i is the only parent that is *TRUE*, X will be *TRUE* iff the inhibitor associated with U_i remains inactive. Hence, we have

$$P(X = \text{TRUE} | U_i = \text{TRUE}, U_k = \text{FALSE } k \neq i) = 1 - q_i.$$

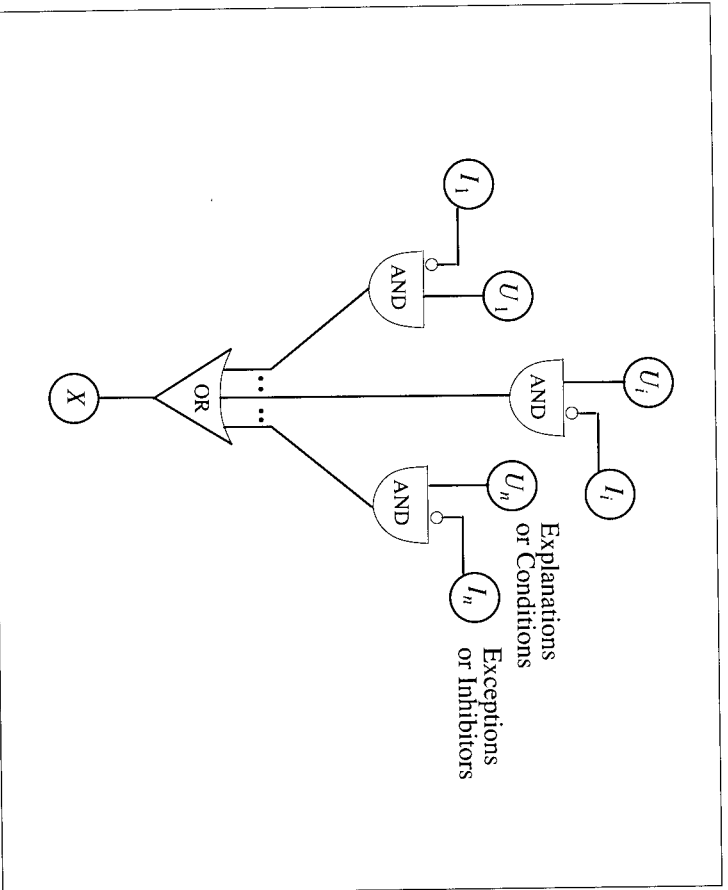


Figure 4.20. The noisy OR-gate. A canonical model of disjunctive interactions among multiple causes U_1, \dots, U_n predicting the same effect X .

Thus, the parameter

$$c_i = 1 - q_i$$

represents the degree to which an isolated explanation $U_i = TRUE$ can endorse the consequent event $X = TRUE$.

Let

$$\mathbf{u} = (u_1, u_2, \dots, u_n) \quad u_i \in \{0, 1\}$$

represent any assignment of truth values to the parent set U . The assumption of exception independence permits a closed-form calculation of the probability distribution of X given any assignment state \mathbf{u} . If T_u represents the subset of parents that are $TRUE$,

$$T_u = \{i: U_i = TRUE\},$$

then X is $FALSE$ iff all inhibitors associated with T_u are active. Thus, we can write

$$P(X = FALSE | \mathbf{u}) = \prod_{i \in T_u} q_i,$$

or

$$P(x | \mathbf{u}) = \begin{cases} \prod_{i \in T_u} q_i & \text{if } x = 0 \\ 1 - \prod_{i \in T_u} q_i & \text{if } x = 1, \end{cases} \tag{4.54}$$

where $x = 0$ and $x = 1$ represent the events $X = FALSE$ and $X = TRUE$, respectively. Eq. (4.54) constitutes a full specification for the link matrix. It also represents a general scheme for approximating a link matrix $P(x | \mathbf{u})$ from individual parent-child relationships, in cases with negligible cross-interaction.

COMPUTING $BEL(x)$

Following Eqs. (4.49) through Eq. (4.54), we write

$$BEL(x) = \alpha \lambda(x) \sum_{\mathbf{u}} P(x | \mathbf{u}) \prod_i \pi_X(u_i) = \begin{cases} \alpha \lambda_0 \sum_{\mathbf{u}} \left(\prod_{i \in T_u} q_i \right) \prod_k \pi_X(u_k) & \text{if } x = 0 \\ \alpha \lambda_1 \sum_{\mathbf{u}} \left(1 - \prod_{i \in T_u} q_i \right) \prod_k \pi_X(u_k) & \text{if } x = 1, \end{cases} \tag{4.55}$$

where $\lambda(x) = (\lambda_0, \lambda_1)$ represents the combined evidential support contributed by X 's children and $\pi_X(u_i)$ represents the message X receives from its i -th parent (i.e., the probability distribution of the i -th parent given all evidence from the subnetwork at the tail of the link $U_i \rightarrow X$).

The summation in Eq. (4.55) can be obtained in closed form by recursive summation over the individual U 's. Denoting by $+u_i$ and $-u_i$ the propositions $U_i = TRUE$ and $U_i = FALSE$, and letting

$$\pi_{X^+} = \pi_X(+u_i) = 1 - \pi_X(-u_i),$$

we define

$$\Pi'_u \triangleq \sum_n \left(\prod_{i \in T_n} q_i \right) \prod_k \pi_{kX}(u_k) = \sum_n \left[\prod_{i \in T_n} q_i \pi_{iX} \right] \left[\prod_{i \notin T_n} (1 - \pi_{iX}) \right],$$

and summing over the two states, $+u_j$ and $-u_j$, of one of the U 's, say U_j , we obtain

$$\begin{aligned} \Pi'_u &= \left[q_j \pi_{jX}(+u_j) + \pi_{jX}(-u_j) \right] \sum_{n-u_j} \prod_{i \in T_n} q_i \prod_{k \neq j} \pi_{kX}(u_k) \\ &= \left[q_j \pi_{jX} + 1 - \pi_{jX} \right] \Pi'_{u-u_j} = (1 - c_j \pi_{jX}) \Pi'_{u-u_j}. \end{aligned}$$

Thus, summing over u_j has the effect of pulling out a factor $(1 - c_j \pi_{jX})$ and leaving behind the same form with $\mathbf{u} - u_j$ instead of \mathbf{u} , where

$$\mathbf{u} - u_j = \{u_1, \dots, u_{j-1}, u_{j+1}, \dots, u_n\}.$$

Applying this recursively† gives

$$\Pi'_u = \prod_i (q_i \pi_{iX} + 1 - \pi_{iX}). \quad (4.56)$$

Hence,

$$BEL(x) = \begin{cases} \alpha \lambda_0 \prod_i (1 - c_i \pi_{iX}) & \text{if } x = 0 \\ \alpha \lambda_1 [1 - \prod_i (1 - c_i \pi_{iX})] & \text{if } x = 1. \end{cases} \quad (4.57)$$

Thus, we see that the overall belief distribution of a variable X can be formed by a multiplication over the individual contributions of X 's parents (and X 's children, via $\lambda(x) = (\lambda_0, \lambda_1)$). The product $c_i \pi_{iX}$ can be interpreted as the degree of *predictive endorsement* that the parent U_i lends to the proposition $X = TRUE$, hence the product $\prod_i (1 - c_i \pi_{iX})$ represents the overall *endorsement withheld* by all

parents. When each $c_i \pi_{iX}$ endorsement is small, Eq. (4.57) can be approximated by

$$\frac{BEL(x=1)}{BEL(x=0)} \approx \frac{\lambda_1}{\lambda_0} \sum_i c_i \pi_{iX}, \quad (4.58)$$

demonstrating that the combined endorsement by all parents is additive over the individual endorsements.

COMPUTING THE λ 's

Using Eq. (4.44) we write

$$\lambda_X(u_i) = \beta \sum_x \lambda(x) \sum_{n-u_i} P(x|\mathbf{u}) \prod_{k \neq i} \pi_{kX}(u_k). \quad (4.59)$$

The second summation in Eq. (4.59) is identical in form to that of Eq. (4.55), except that u_i is now excluded from the summation. Accordingly, using the formula Eq. (4.56) for Π'_u , we get

$$\lambda_X(u_i) = \beta \sum_x \lambda(x) F(x, u_i),$$

where

$$F(x, +u_i) = \begin{cases} q_i \prod_{k \neq i} (1 - c_k \pi_{kX}) & \text{if } x = 0 \\ 1 - q_i \prod_{k \neq i} (1 - c_k \pi_{kX}) & \text{if } x = 1 \end{cases}$$

and

$$F(x, -u_i) = F(x, +u_i) \Big|_{q_i=1}.$$

Inserting this in Eq. (4.59) yields

$$\lambda_X(u_i) = \begin{cases} \beta [\lambda_0 q_i \Pi'_i + \lambda_1 (1 - q_i \Pi'_i)] & \text{if } u_i = +u_i \\ \beta [\lambda_0 \Pi'_i + \lambda_1 (1 - \Pi'_i)] & \text{if } u_i = -u_i, \end{cases} \quad (4.60)$$

where

$$\Pi'_i = \prod_{k \neq i} (1 - c_k \pi_{kX}) \quad (4.61)$$

† This derivation is due to Charles Kalme. Another way of verifying Eq. (4.56) is to use the formula $\prod_i (a_i + b_i) = \sum_n \left(\prod_{i \in T_n} a_i \right) \left(\prod_{i \notin T_n} b_i \right)$, substituting $a_i = q_i \pi_{iX}$ and $b_i = 1 - \pi_{iX}$.

represents the overall endorsement withheld by all parents except U_i . A more concise formula for encoding $\lambda_X(u_i)$ is

$$\lambda_X(u_i) = \beta \left[\lambda_1 - q_i^H (\lambda_1 - \lambda_0) \Pi_i' \right] \quad u_i = 0, 1.$$

INTUITIVE INTERPRETATION AND CREDIT ASSIGNMENT POLICY

Since only the ratios

$$L_{XU_i} = \frac{\lambda_X(+u_i)}{\lambda_X(-u_i)} \quad \text{and} \quad L_X = \frac{\lambda_1}{\lambda_0}$$

enter the updating rules, we can write Eq. (4.60) as

$$L_{XU_i} = \frac{q_i \Pi_i' + (1 - q_i \Pi_i') L_X}{\Pi_i' + (1 - \Pi_i') L_X} = 1 + \frac{\Pi_i' c_i (L_X - 1)}{1 + (1 - \Pi_i') (L_X - 1)} \quad (4.62)$$

Note that $L_X = 1$ implies $L_{XU_i} = 1$, so if X gathers no evidential support, U_i will neither receive evidential support from X nor be affected in any way by the other parents. This phenomenon is true in general (see Eq. (4.44)), and it reaffirms our intuition that causal frames should remain uncoupled until their common slots receive indication of empirical confirmation or denial. However, the noisy-OR-gate model has one more condition under which causal frames remain uncoupled, namely, when X is completely denied. This can be seen by substituting $L_X = 0$ in Eq. (4.62), yielding

$$L_{XU_i} = 1 - c_i = q_i,$$

independent of c_k or π_{kX} , $k \neq i$. Thus, the denial of a common effect X results in each causal factor receiving a constant support $q_i \leq 1$, regardless of the existence of other factors (q_i is usually very small, so $L_{XU_i} = q_i$ means the withholding of support from U_i). From a network analysis viewpoint, this phenomenon means that any negatively instantiated variable X acts as an absorption barrier, similar to an anticipatory variable. The messages R_{kX} converging onto such a variable are totally absorbed and do not generate new λ messages. Negatively instantiated variables can even turn multiply connected networks into polytrees (see the example in Section 4.5.3, Figure 4.37).

There is another condition under which U_i receives no evidential support from X , namely, when at least one of the other parents, say U_k , extends a full endorsement to X by having $c_k \pi_{kX} = 1$. Under this condition $\Pi_i' = 0$, and again,

$\lambda_X(+u_i) / \lambda_X(-u_i) = 1$ regardless of c_i , meaning the connection between X and U_i is totally disrupted by U_k . This fits intuition; once we find a satisfactory explanation for a suspected symptom, that symptom no longer imparts confirmation to other explanations, no matter how sure we are of the symptom. For example, once I learn that the ignition wires in my car are disconnected, I no longer hypothesize a faulty battery no matter how sure I am that my car will not start. Explanations act as logical valves shutting off each other's flow of evidential support when they become confirmed. When two or more explanations achieve $c_k \pi_{kX} = 1$, none of them can receive evidential support from X because the Π_i' of each one, and of any other explanation of X , becomes zero. In fact, achieving $c_k \pi_{kX} = 1$ means the k -th explanations no longer need X 's support; it is firmly established by other sources of information, and from that position its sole effect would be to undermine the flow of support from X to every other explanation. In Chapter 10, we shall formulate this undermining effect in the framework of default logic.




Figure 4.21. The support that a partially confirmed event leads to its i -th explanation, U_i , as a function of the evidential weight accrued at X , and the amount of credit Π_i' left unclaimed by alternative explanations.

Figure 4.21 depicts the relationship between the likelihood ratios L_{XU_i} and L_X for several values of Π'_i . It shows that for a given value of Π'_i , there is a maximum amount of evidential support that can be transferred from X to U_i . That amount (see Eq. (4.62)) is equal to

$$L_{XU_i} = 1 + \frac{c_i \Pi'_i}{1 - \Pi'_i}, \quad (4.63)$$

and it is reached when $L_X = \infty$, namely, when $X = TRUE$ is confirmed with absolute certainty. Thus, Π'_i can be thought of as the maximum amount of *residual credit* left unclaimed by all alternative explanations combined.

An interesting situation develops when the a priori probabilities π_{kX} are extremely small. Under this condition, each cause individually provides only a small endorsement to the event $X = TRUE$. Yet when the truth of X is confirmed, the relative likelihoods of the various causes are determined by the relative magnitudes of their $c_k \pi_{kX}$ products. To see this, we write

$$\frac{BEL(+u_i)}{BEL(+u_j)} = \frac{\pi_{iX} L_{XU_i}}{\pi_{jX} L_{XU_j}}$$

Since $\pi_{kX} \ll 1$ renders $\Pi'_i \approx 1$ in Eq. (4.61), Eq. (4.63) yields

$$\frac{BEL(+u_i)}{BEL(+u_j)} \approx \frac{c_i \pi_{iX}}{c_j \pi_{jX}}$$

If we make the further assumption that only one of the causes can be true (*single cause assumption*), we get the formula

$$BEL(+u_i) = \frac{c_i \pi_{iX}}{\sum_k c_k \pi_{kX}}$$

Thus, the amount of credit deserved by explanation U_i is given by the product $c_i \pi_{iX}$, i.e., the endorsement offered by the explanation weighted by the probability that it is in fact *TRUE*.

To further demonstrate how parents compete for the evidential support provided by X let us assume that X is not confirmed with certainty, i.e., L_X is finite.

Making the rare-endorsement assumption

$$c_k \pi_{kX} \ll 1 \quad \text{for } k \neq i$$

and expanding Eq. (4.62) in Taylor's series, we obtain

$$L_{XU_i} - 1 \approx [L_X - 1] c_i \left[1 - \sum_{k \neq i} c_k \pi_{kX} \right]. \quad (4.64)$$

The term $L_X - 1$ can be identified with the overall *evidential weight* received by the proposition $X = TRUE$. Accordingly, Eq. (4.64) demonstrates that the evidential weight received from X by the i -th parent depends on three factors:

1. $L_X - 1$ = the total evidential weight accrued at X .
2. c_i = the degree to which U_i endorses X (if U_i is *TRUE*).
3. $1 - \sum_{k \neq i} c_k \pi_{kX}$ = the amount of *residual credit* left unclaimed by all other parents.

Thus, the struggle for a share of X 's support is settled by the following policy: he who risks the strongest endorsement in a prediction earns the greatest credit when that prediction materializes. Note that while parents compete for credit from their child's success, the same is not true for children; the more a child endorses a parent, the greater the benefit to its brethren. This distinction is further examined in Section 10.3.

OTHER CANONICAL MODELS

The basic noisy-OR-gate model of the preceding subsection is too restrictive in some applications, and more elaborate interactions among the causal conditions are sometimes needed. For example, conditions might interact conjunctively rather than disjunctively. The basic structure of such models remains the same, however: Boolean combinations of explicated conditions. The unexplicated conditions are summarized probabilistically under the assumption of exception independence.

The assumption of exception independence can easily be relaxed to represent global abnormality conditions, i.e., conditions that would inhibit the response event X even when several causal factors are triggered. For example, a power failure would inhibit the alarm from sounding under simultaneous burglary and earthquake conditions. To incorporate global inhibition in the model of Figure 4.20 we simply add another AND-gate between the OR-gate and the response

variable X , as in Figure 4.22a. When the global inhibitor I_0 is ON, X will be OFF, regardless of other causal factors.

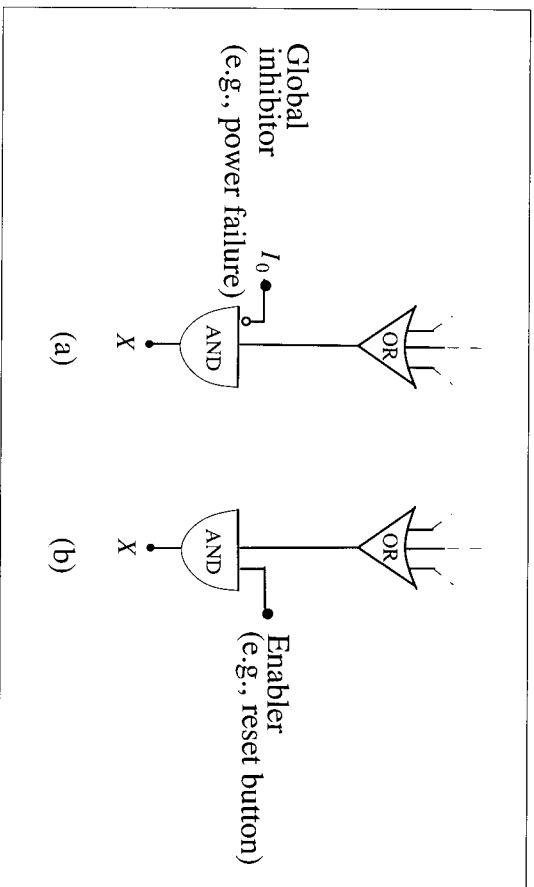


Figure 4.22. Canonical models of global inhibition (a) and enabling mechanisms (b).

In this fashion, we can also model various enabling mechanisms, i.e., conditions that have no influence of their own (on X) except to enable other influences to take effect. For example, if the alarm system has a reset button which Mr. Holmes occasionally forgets to push, setting this button is an enabling condition, as in Figure 4.22b.

SUMMARY

Canonical models can be thought of as default strategies for completing the specification of a Bayesian network whenever detailed interactions among causes are unavailable, too numerous to elicit, or too complex to be treated precisely. In particular, the disjunctive model of interacting causes has several advantages: it requires the specification of only n parameters (for a family of n parents), it executes the propagation routine in only n steps (for each node with n parents), and it leads to conclusions that match our intuition about how credit should be assigned among competing explanations. Having explicated the assumptions behind disjunctive interaction, i.e., accountability and exception independence, we can scrutinize the model's range of adequacy at a very basic level, and once this test is passed, the adequacy of the credit assignment policy is guaranteed.