An Approximate Nonmyopic Computation for Value of Information

David Heckerman, Eric Horvitz, and Blackford Middleton

An Approximate Nonmyopic Computation for Value of Information

David Heckerman
Department of Computer Science
University of California
Boelter Hall, Room 3531
405 Hilgard Avenue
Los Angeles, CA 90024-1596
<heckerman@cs.ucla.edu>

Eric Horvitz
Palo Alto Laboratory
Rockwell International Science Center
444 High Street
Palo Alto, California 94301
and
Section on Medical Informatics
Stanford University Medical Center
Medical School Office Building, Room 215
Stanford, CA 94305-5479
<horvitz@camis.stanford.edu>

Blackford Middleton
Medical Director, Clinical Information Systems
Clinical Assistant Professor of Medicine
Stanford University Medical Center
Medical School Office Building, Room 214
Stanford, CA 94305-5479
<bfm@camis.stanford.edu>

Abstract

Value-of-information analyses provide a means for selecting the next best observation to make, and for determining whether it is better to gather additional information or to act immediately. Determining the next best test to perform, given uncertainty about the state of the world, requires a consideration of the value of making all possible sequences of observations. In practice, decision analysts and expert-system designers have avoided the intractability of exact computation of the value of information by relying on a *myopic* assumption that only one additional test will be performed, even when there is an opportunity to make a large number of observations. We present an alternative to the myopic analysis. In particular, we present an approximate method for computing the value of information of a *set* of tests, which exploits the statistical properties of large samples. The approximation is linear in the number of tests, in contrast to the exact computation, which is exponential in the number of tests. The approach is not as general as is a complete nonmyopic analysis, in which all possible sequences of observations are considered. Also, the approximation is limited to specific classes of dependencies among evidence and to binary hypothesis and decision variables. Nonetheless, as we demonstrate with a simple application, the approach can offer an improvement over the myopic analysis.

Keywords: Probability, belief networks, decision theory, value of information, nonmyopic.

1 Introduction

When performing diagnosis, a person usually has the opportunity to gather additional information about the state of the world before making a final diagnosis. Such information gathering typically is associated with costs and benefits. These costs and benefits can be balanced with decision-theoretic techniques—in particular, procedures for computing *value of information.* These techniques form an integral part of many decision-theoretic expert systems for diagnosis, such as Gorry and Barnett's program for the diagnosis of congestive heart failure [**?**].

In most diagnosis contexts, a decision maker has the option to perform several tests, and can decide which test to perform after seeing the results of all previous tests. Thus, a person or expert system should consider the value of all possible sequences of tests. Such an analysis is intractable, because the number of sequences grows exponentially with the number of tests. Builders of expert systems have avoided the intractability of exact value-of-information computations by implementing *myopic* or *greedy* value-of-information analyses. In such analyses, a system determines the next best test by computing the value of information based on the assumption that the decision maker will act immediately after seeing the results of the single test [**?**].

The work presented in this article is motivated by Pathfinder, a decision-theoretic expert system that assists physicians with the diagnosis of lymph-node diseases [**?**, **?**, **?**]. The Pathfinder project began in 1983 as a joint project among researchers (David Heckerman, Eric Horvitz, Jaap Suermondt, Mark Fischinger, and Larry Fagan) in the Medical Computer Science Group at Stanford University and researchers at the University of Southern

California (Bharat Nathwani—the primary pathology expert—and Keung-Chi Ng). Currently, a commercial derivative of Pathfinder, called Intellipath, is being used by several hundred practicing pathologists and by pathologists in training as an educational tool [**?**]. The program reasons about over 60 diseases (25 benign diseases, 9 Hodgkin's lymphomas, 18 non-Hodgkin's lymphomas, and 10 metastatic diseases) and over 140 features of disease, including clinical, microscopic, laboratory, immunologic, and molecular biological findings.

In some instances of Pathfinder's use, a myopic value-of-information analysis is inappropriate. For example, suppose that a patient's primary physician has clinical information suggesting that the patient may have a serious lymph-node disease. At this point, one alternative available to the patient is a tissue biopsy: the surgical removal of one or more lymph nodes. If the biopsy is performed, a surgical pathologist examines the tissue using a microscope, and provides additional evidence for or against each possible disease. The tissue biopsy can provide a large amount of information, but is costly and subjects the patient to the risks of general anesthesia.

Pathfinder can assist the patient and physician with the decision of whether or not to perform a biopsy. Because the program uses a myopic value-of-information analysis, however, it can balance the cost of the biopsy with the value of only one of approximately 100 microscopic features. Thus, when a biopsy is cost effective, Pathfinder will not likely recommend one.

In this article, we present a tractable real-time solution to this problem. In particular, we develop an approach that takes advantage of the statistical properties of large samples to compute approximately the value of information for *sets* of tests. The approximation is linear in the number of tests, in contrast to the exact computation, which is exponential in the number of tests. The approach is not as general as is a complete nonmyopic analysis, in which all possible sequences of observations are considered. Also, the approximation is limited to specific classes of dependencies among evidence and to binary hypothesis and decision variables. Nonetheless, as we demonstrate with the biopsy example, the approach can be an improvement over the myopic analysis.

2 A Decision-Theoretic Model for Diagnosis

The diagnostic model for Pathfinder as well as other decision-theoretic expert systems is represented by the influence diagram in Figure 1. In this model, the chance node H represents a mutually exclusive and exhaustive set of possible hypotheses, and the decision node D represents a mutually exclusive and exhaustive set of possible actions or alternatives. The value node U represents the utility of the decision maker, which depends on the outcome of H and the decision D. The chance nodes $E_1, \ldots, E_n$ are observable pieces of evidence or tests about the true state of H. Pieces of evidence in Pathfinder are called features.

Insert Figure 1 about here.

In the first part of this article, we make several simplifying assumptions. First, we assume that H is a binary chance variable and D is a binary decision variable. We use H and $\neg H$ to denote the two instances of H, and D and $\neg D$ to denote the two alternatives associated with D. For definiteness, we assume that the decision maker chooses D (as opposed to $\neg D$), when

H occurs. Second, we assume that each piece of evidence, $E_1, \ldots, E_n$, is binary. Finally, we assume that each piece of evidence is conditionally independent of all other evidence, given H and $\neg H$. In Section 5, we relax several of these assumptions.

Using Bayes' theorem and the assumption of conditional independence of evidence, we can calculate the ratio of the posterior probability of H to that of $\neg H$:

$$\frac{p(H|E_i, \ldots, E_m)}{P(\neg H|E_i, \ldots, E_m)} = \frac{p(E_1|H)}{p(E_1|\neg H)} \cdots \frac{p(E_m|H)}{p(E_m|\neg H)} \frac{p(H)}{p(\neg H)}$$

We can write this equation more compactly in odds-likelihood form as

$$O(H|E_i, \ldots, E_m) = O(H) \prod_{i=1}^{m} \lambda_i \tag{1}$$

where $O(H|E_i, \ldots, E_m)$ is the posterior odds of H, λ_i is the likelihood ratio $\frac{p(E_i|H)}{p(E_i|\neg H)}$, and $O(H)$ is the prior odds of H.

Because D and H are binary, it follows from the axioms of decision theory that there exists a threshold probability p^*, such that we should take action D if and only if the probability of H exceeds p^*. This threshold is the probability of H at which the decision maker is indifferent between acting and not acting. That is, p^* is the point where acting and not acting have equal utility, or

$$p^*U(H, D) + (1 - p^*)U(\neg H, D) = p^*U(H, \neg D) + (1 - p^*)U(\neg H, \neg D) \tag{2}$$

In Equation 2, $U(H, D)$ is the decision maker's utility for the situation where H occurs and action D is taken, $U(H, \neg D)$ is the utility when H occurs and action D is not taken, and so on. Solving Equation 2 for p^*, we obtain

$$p^* = \frac{C}{C + B} \tag{3}$$

where C is the *cost* of the decision

$$C \equiv U(\neg H, \neg D) - U(\neg H, D) \tag{4}$$

and B is the *benefit* of the decision

$$B \equiv U(H, D) - U(H, \neg D) \tag{5}$$

If the decision maker has observed pieces of evidence $E_1, \ldots, E_m$, then the decision maker should choose action D if and only if

$$p(H|E_1, \ldots, E_m) > p^* \tag{6}$$

In terms of the odds formulation, Equation 6 becomes

$$O(H|E_1, \ldots, E_m) > \frac{p^*}{1 - p^*} \tag{7}$$

Equations 1 and 7 imply

$$\prod_{i=1}^{m} \lambda_i \quad > \quad \frac{p^*}{1-p^*} \; / \; O(H) \tag{8}$$

Taking the logarithm of both sides of Equation 8, we see that the decision maker should choose action D if and only if

$$\sum_{i=1}^{m} w_i \quad > \quad \ln \frac{p^*}{1-p^*} - \ln O(H) \tag{9}$$

where $w_i \equiv \ln \lambda_i$ is called the *weight of evidence* E_i for H. With the definitions

$$W \equiv \textstyle\sum_{i=1}^{m} w_i \qquad\qquad W^* \equiv \ln \frac{p^*}{1-p^*} - \ln O(H) \tag{10}$$

we have the simple prescription that the decision maker should choose action D if and only if

$$W > W^* \tag{11}$$

3 Myopic Analysis

Let us assume that the user of a diagnostic system has instantiated zero or more pieces of evidence in the influence diagram shown in Figure 1. We can propagate the effects of these instantiations to the uninstantiated nodes, and remove the instantiated nodes from the influence diagram. This removal leaves an influence diagram of the same form as that shown in Figure 1. To simplify our notation, we continue to refer to the remaining pieces of evidence as $E_1, \ldots, E_n$. Also, we use $p(H)$ to refer to the probability of the hypothesis H, given the instantiated evidence.

The decision maker now considers whether he should observe another piece of evidence before acting. A myopic procedure for identifying such evidence computes, for each piece of evidence, the expected utility of the decision maker under the assumption that *the decision maker will act after observing only that piece of evidence.* In addition, the procedure computes his expected utility if he does not observe any evidence before making his decision. If, for each piece of evidence, the expected utility given that evidence is less than the expected utility given no evidence, then the decision maker acts immediately in accordance with Equation 11. Otherwise, the decision maker observes the piece of evidence with the highest expected utility. Then, the myopic procedure *repeats this computation to identify additional evidence for observation.* Because the myopic procedure allows for the gathering of additional evidence, the procedure is inconsistent with its own assumptions. We return to this observation in the next section.

In the remainder of this section, we examine the computation of expected utilities and introduce notation. Let $EU(E, C_E)$ denote the expected utility of the decision maker who will observe E at cost C_E, and then act. Let $CE(E, C_E)$ be the certain equivalent of this situation. That is,

$$U(CE(E, C_E)) \equiv EU(E, C_E) \tag{12}$$

or

$$CE(E, C_E) = U^{-1}(EU(E, C_E)) \tag{13}$$

where $U(\cdot)$ is the decision maker's *utility function*: a monotonic increasing function that maps the value of an outcome (e.g., in dollars) to the decision maker's utility for that outcome. Similarly, let $EU(\emptyset, 0)$ denote the expected utility of the decision maker if he acts immediately, and let $CE(\emptyset, 0)$ denote the certain equivalent of this situation. Thus, in the myopic procedure, a decision maker should observe the piece of evidence E for which the quantity

$$CE(E, C_E) - CE(\emptyset, 0) \tag{14}$$

is maximum, provided it is greater than 0.

To simplify the discussion, we assume that the delta property holds.[1] The *delta property* states that an increase in value of all outcomes in a lottery by an amount $\triangle$ increases the certain equivalent of that lottery by $\triangle$ [**?**]. Under this assumption, we obtain

$$CE(E, C_E) = CE(E, 0) - C_E \tag{15}$$

where $CE(E, 0)$ is the certain equivalent of observing E *at no cost.* Therefore, we have

$$CE(E, C_E) - CE(\emptyset, 0) = VI(E) - C_E \tag{16}$$

where

$$VI(E) \equiv CE(E, 0) - CE(\emptyset, 0) \tag{17}$$

is the *value of information* of observing E.[2] The quantity $VI(E)$ represents the largest amount that the decision maker would be willing to pay to observe E. When we compare Expression 14 with Equation 16, we see that, in the myopic procedure, a decision maker should observe the piece of evidence E (if any) for which the quantity

$$VI(E) - C_E \equiv NVI(E) \tag{18}$$

is maximum and positive. We call $NVI(E)$ the *net value of information* of observing E.

The decision maker usually specifies directly the cost of observing evidence. In contrast, we can compute $VI(E)$ from the decision maker's utilities and probabilities. Specifically, from Equations 13 and 17, we have

$$VI(E) = U^{-1}(EU(E, 0)) - U^{-1}(EU(\emptyset, 0))$$

To simplify notation, we use the abbreviations

$$EU(E, 0) \equiv EU(E) \quad \text{and} \quad EU(\emptyset, 0) \equiv EU(\emptyset)$$

Thus, we obtain

$$VI(E) = U^{-1}(EU(E)) - U^{-1}(EU(\emptyset)) \tag{19}$$

[1] The primary result of this research—that we can use the central-limit theorem to make tractable an approximate nonmyopic analysis—is unaffected by this assumption.

[2] Other names for $VI(E)$ include the value of perfect information of E and the value of clairvoyance on E.

The computation of $EU(\emptyset)$ is straightforward. We have

$$EU(\emptyset) = \begin{cases} p(H)\ U(H, \neg D) + p(\neg H)\ U(\neg H, \neg D), & p(H) \leq p^* \\ p(H)\ U(H, D) + p(\neg H)\ U(\neg H, D), & p(H) > p^* \end{cases} \tag{20}$$

by definition of p^*.

To compute $EU(E)$, let us assume that E is defined such that observing E to be true increases the probability that H is true. If $p(H|E) > p^*$ and $p(H|\neg E) > p^*$, then $VI(E) = 0$, because the decision maker will not change his decision if he observes E. Similarly, if $p(H|E) < p^*$ and $p(H|\neg E) < p^*$, then $VI(E) = 0$. Thus, we need only to consider the case where $p(H|E) > p^*$ and $p(H|\neg E) < p^*$. Let us consider separately the cases H and $\neg H$. We have

$$EU(E|H) = p(E|H)\ U(H, D) + p(\neg E|H)\ U(H, \neg D) \tag{21}$$

and

$$EU(E|\neg H) = p(E|\neg H)\ U(\neg H, D) + p(\neg E|\neg H)\ U(\neg H, \neg D) \tag{22}$$

where $EU(E|H)$ and $EU(E|\neg H)$ are the expected utilities of observing E, given H and $\neg H$, respectively. To obtain the expected utility of observing E, we average these two quantities over H:

$$EU(E) = p(H)\ EU(E|H) + p(\neg H)\ EU(E|\neg H) \tag{23}$$

To compute $VI(E)$, we combine Equations 19, 20, and 23.

4 A Special-Case Nonmyopic Analysis

As we mentioned in the previous section, the myopic procedure for identifying cost-effective observations includes the incorrect assumption that the decision maker will act after observing only one piece of evidence. This myopic assumption can deleteriously affect the performance of an expert system, as described in the introduction.

In a correct identification of cost-effective evidence, an expert system should take into account the fact that a person can observe more than one piece of evidence before acting. In its most general form, this computation should consider all possible observation strategies. An example of an observation strategy is

> Observe E_3. If E_3 is present, then observe E_2; otherwise, make no further observations and make the diagnosis. If E_3 and E_2 are present, then observe E_7, and make the diagnosis. If E_3 is present and E_2 is absent, then make the diagnosis.

In this article, we consider a special-case nonmyopic analysis that considers only two observation strategies: (1) perform a *set* of tests, and then make the diagnosis, and (2) make the diagnosis immediately (the trivial observation strategy). The general nonmyopic analysis reduces to this special case when there is a specific dependency among the costs of performing tests. Namely, the general nonmyopic analysis reduces to this special case when there are a set of tests such that the cost of performing any test in the set is high, and once any test in the set has been performed, the cost of performing additional tests in the

set is significantly reduced. This special-case analysis is appropriate for the biopsy example discussed in the introduction.

Let us suppose that the decision maker has the option to observe a particular subset of evidence $\{E_1, \ldots, E_n\}$ before acting. We assume that the costs of observing the pieces of evidence in this set are dependent as described in the previous paragraph, and that the decision maker can specify directly the initial cost of observing a piece of evidence in this set. There are 2^n possible instantiations of the evidence in this set, corresponding to the observation of E_i or $\neg E_i$ for every i. Let $\mathcal{E}$ denote an arbitrary instantiation; let $\mathcal{E}_D$ and $\mathcal{E}_{\neg D}$ denote the set of instantiations $\mathcal{E}$ such that the optimal decision is D and $\neg D$, respectively.

The computation of the value of information for the observation of the set $\{E_1, \ldots, E_n\}$ parallels the myopic computation. In particular, we have

$$EU(E_1, \ldots, E_n) = p(H)\ EU(E_1, \ldots, E_n|H) + p(\neg H)\ EU(E_1, \ldots, E_n|\neg H) \tag{24}$$

where

$$EU(E_1, \ldots, E_n|H) = \left[\sum_{\mathcal{E} \in \mathcal{E}_D} p(\mathcal{E}|H)\right] U(H, D) + \left[\sum_{\mathcal{E} \in \mathcal{E}_{\neg D}} p(\mathcal{E}|H)\right] U(H, \neg D) \tag{25}$$

and

$$EU(E_1, \ldots, E_n|\neg H) = \left[\sum_{\mathcal{E} \in \mathcal{E}_D} p(\mathcal{E}|\neg H)\right] U(\neg H, D) + \left[\sum_{\mathcal{E} \in \mathcal{E}_{\neg D}} p(\mathcal{E}|\neg H)\right] U(\neg H, \neg D) \tag{26}$$

To obtain $VI(E)$, we combine Equations 19, 20, and 24.

When n is small, we can compute directly the sums in Equations 25 and 26. When n is large, we can compute these sums using an approximation that involves the central limit theorem as follows. First we express the sums in terms of weights of evidence. We have

$$\sum_{\mathcal{E} \in \mathcal{E}_D} p(\mathcal{E}|H) = p(W > W^*|H) \tag{27}$$

$$\sum_{\mathcal{E} \in \mathcal{E}_D} p(\mathcal{E}|\neg H) = p(W > W^*|\neg H) \tag{28}$$

$$\sum_{\mathcal{E} \in \mathcal{E}_{\neg D}} p(\mathcal{E}|H)) = 1 - p(W > W^*|H) \tag{29}$$

$$\sum_{\mathcal{E} \in \mathcal{E}_{\neg D}} p(\mathcal{E}|\neg H)) = 1 - p(W > W^*|\neg H) \tag{30}$$

where W and W^* are defined in Equation 10. The term $p(W > W^*|H)$, for example, is the probability that the sum of the weight of evidence from the observation of $E_1, \ldots, E_n$ exceeds W^*. That is, $p(W > W^*|H)$ is the probability that the decision maker will take action D after observing the evidence, given that H is true.

Next, let us consider the weight of evidence for one piece of evidence. We have

w_i	$p(w_i\|H)$	$p(w_i\|\neg H)$
$\ln \frac{p(E_i\|H)}{p(E_i\|\neg H)}$	$p(E_i\|H)$	$p(E_i\|\neg H)$
$\ln \frac{p(\neg E_i\|H)}{p(\neg E_i\|\neg H)}$	$p(\neg E_i\|H)$	$p(\neg E_i\|\neg H)$

To simplify notation, we let $p(E_i|H) = \alpha$ and $p(E_i|\neg H) = \beta$. The expectation and variance of w, given H and $\neg H$, are then

$$EV(w|H) = \alpha \ln \frac{\alpha}{\beta} + (1-\alpha) \ln \frac{(1-\alpha)}{(1-\beta)} \tag{31}$$

$$Var(w|H) = \alpha(1-\alpha)\ln^2 \frac{\alpha(1-\beta)}{\beta(1-\alpha)} \tag{32}$$

$$EV(w|\neg H) = \beta \ln \frac{\alpha}{\beta} + (1-\beta) \ln \frac{(1-\alpha)}{(1-\beta)} \tag{33}$$

$$Var(w|\neg H) = \beta(1-\beta)\ln^2 \frac{\alpha(1-\beta)}{\beta(1-\alpha)} \tag{34}$$

Now, we take advantage of the additive property of weights of evidence. The central-limit theorem states that the sum of independent random variables approaches a normal distribution when the number of variables becomes large. Furthermore, the expectation and variance of the sum is just the sum of the expectations and variances of the individual random variables, respectively. Because we have assumed that evidence variables are independent, given H or $\neg H$, the expected value of the sum of the weights of evidence for $E_1, \ldots, E_n$ is

$$EV(W|H) = \sum_{i=1}^{m} EV(w_i|H) \tag{35}$$

The variance of the sum of the weights is

$$Var(W|H) = \sum_{i=1}^{m} Var(w_i|H) \tag{36}$$

Thus, $p(W|H)$, the probability distribution over W, is given by

$$p(W|H) \sim N(\sum_{i=1}^{m} EV(w_i|H), \sum_{i=1}^{m} Var(w_i|H)) \tag{37}$$

The expression for $\neg H$ is similar.

Finally, given the distributions for H and $\neg H$, we evaluate Equations 27 through 30 using an estimate or table of the cumulative normal distribution. We have

$$p(W > W^*|H) = \frac{1}{\sigma\sqrt{2\pi}} \int_{W^*}^{\infty} e^{\frac{-(t-\mu)^2}{2\sigma}} dt \tag{38}$$

where $\mu = EV(W|H)$ and $\sigma = Var(W|H)$. The probability that the weight will exceed W^* corresponds to the shaded area in Figure 2. Again, the expression for $\neg H$ is similar. In this analysis, we assume that no probability ($p(E_i|H)$ or $p(E_i|\neg H)$) is equal to 0 or 1. Thus, all expected values and variances are finite. We relax this assumption in the next section.

Insert Figure 2 about here.

5 Relaxation of the Assumptions

We can relax the assumption that evidence is two-valued with little effort. In particular, we can extend easily the odds-likelihood inference rule, Equation 1, and its logarithmic transform, to include multiple-valued evidential variables. In addition, the computation of means and variances for multiple-valued evidential variables (see Equations 31 through 34) is straightforward.

In addition, we can relax the assumption that no probability is equal to 0 or 1. For example, let us suppose that

$$0 < p(E_j|H) = \alpha < 1 \qquad p(E_j|\neg H) = \beta = 1$$

and, for all $i \neq j$,

$$0 < p(E_i|H) < 1 \qquad 0 < p(E_i|\neg H) < 1$$

Using Equations 31 through 34, we obtain

$$\begin{array}{rclcrcl} EV(w_j|H) & = & +\infty & & Var(w_j|H) & = & +\infty \\ EV(w_j|\neg H) & < & 0 & & Var(w_j|\neg H) & = & 0 \end{array}$$

Therefore, although the computation of $p(W > W^*|\neg H)$ is straightforward, we cannot compute $p(W > W^*|H)$ as described in the previous section. Instead, we compute $p(W > W^*|H)$, by considering separately the cases E_j and $\neg E_j$. We have

$$p(W > W^*|H) = p(E_j|H)\ p(W > W^*|H, E_j) + p(\neg E_j|H)\ p(W > W^*|H, \neg E_j) \qquad (39)$$

If $\neg E_j$ is observed, $W = +\infty$, and $p(W > W^*|H, \neg E_j) = 1$. Consequently, Equation 39 becomes

$$p(W > W^*|H) = p(E_j|H)\ p(W > W^*|H, E_j) + p(\neg E_j|H)$$

We compute $p(W > W^*|H, E_j)$ as described in Equations 35 through 38, replacing $EV(w_j|H)$ with w_j in the summation of Equation 35, and $Var(w_j|H)$ with 0 in the summation of Equation 36. The other terms in the summations remain the same, because we have assumed that evidence variables are independent, given H or $\neg H$. This approach generalizes easily to multiple-valued evidence variables and to cases where more than one probability is equal to 0 or 1.

We can extend our analysis to special cases of conditional dependence among evidence variables. For example, Figure 3 shows a schematic of the belief network for Pathfinder. In this model, there are groups of dependent evidence, where each group is conditionally independent of all other groups. We can apply our analysis to this model by using a clustering technique described by [**?**][pages 197–204]. As in the previous section, suppose we want to compute the value of information for the set of evidence $S = \{E_1, \ldots, E_n\}$. For each group of dependent features G^k, we cluster those variables in the intersection of S and G^k into a single variable. Then, we average out all variables in the belief network that are not in S. We obtain clusters of variables each of which are conditionally independent, given H and $\neg H$. We can now apply our analysis—generalized to multiple-valued variables—to this model.

Insert Figure 3 about here.

There are special classes of dependent distributions for which the central-limit theorem is valid. We can use this fact to extend our analysis to other cases of dependent evidence. For example, the central-limit theorem applies to distributions that form a Markov chain, provided the transition probabilities in the chain are not correlated [**?**]. Thus, we can extend our analysis to belief networks of the form shown in Figure 4. We can generalize the value-of-information analysis even further, if we use the Markov extension in combination with the clustering approach described in the previous paragraph.

Insert Figure 4 about here.

It is difficult for us to extend the analysis to include multiple-valued hypotheses and decisions. The mathematics becomes more complex, because the simple p^* model for action no longer applies. There is, however, the opportunity for applying our technique to more complex problems. In particular, we can abstract a given decision problem into one involving a binary hypothesis and decision variable. For example, we can abstract the problem of determining which of n diseases is present in a patient into one of determining whether the disease is malignant (i.e., cancer) or benign. In doing so, we ignore details of the decision maker's preferences, and we introduce dependencies among evidence variables. Nonetheless, the benefits of a nonmyopic analysis may outweigh these drawbacks in some domains.

6 A Simple Application

Let us return to the situation described in the introduction: A patient's primary care physician believes, based on clinical evidence, that the patient may have a malignant lymph-node disease. The patient may receive a lymph-node biopsy, at high cost, before a treatment decision is made. If the biopsy is performed, a pathologist can inspect the tissue microscopically, thereby providing a large number of observations that are clues about the patient's disease.

As described in the previous section, we abstract the diagnostic problem to that of determining whether or not the patient has a malignant or benign disease. In addition, we assume that there are only two treatment alternatives: (1) treat the patient as if he had a malignant disease—that is, treat the patient with chemotherapy, surgery, radiation therapy, or some combination of these procedures—or (2) do not treat the patient, but merely watch his progress carefully.

To simplify the discussion, we consider only a fraction of clues made available by the pathologist. In particular, we consider only those features that describe follicles—spherical aggregates of multiplying white cells—in a lymph-node section. Also, we assume that the clinical and microscopic observations are conditionally independent, given the patient's disease. Consequently, we do not have to consider interactions among the two information sets.

The influence diagram for the pathologist's diagnostic task is shown in Figure 5. The hypothesis node contains the two disease instances: malignant (H) and benign ($\neg H$). The decision node contains two alternatives: treat (D) and watch ($\neg D$). The node U represents the patient's utility for the four possible outcomes: (Malignant, Treat), (Malignant, Watch), (Benign, Treat), and (Benign, Watch). The evidence variables represent microscopic observations about the follicles that provide clues about the disease state of the patient. For

example, the feature "Area" represents the percent area of the lymph-node section occupied by follicles; and the feature "Polarity" represents whether one or more follicles have a uniform appearance or exhibit different distributions of cell types at opposite poles. The influence diagram was constructed from data (48 patients) using the K2 algorithm [**?**].[3]

Insert Figure 5 about here.

To simplify the discussion further, we express the utilities of the four possible outcomes in dollars. The values we use are

$$\begin{aligned} U(\text{Malignant}, \text{Treat}) &= -\$300K & \qquad U(\text{Malignant}, \text{Watch}) &= -\$800K \\ U(\text{Benign}, \text{Treat}) &= -\$100K & \qquad U(\text{Benign}, \text{Watch}) &= \$0 \end{aligned}$$

In addition, we assume that the decision maker is an expected-value decision maker. That is, we assume $U(X) = X$, so that expected value and expected utility are the same quantity, and so that the delta property holds. Finally, for the cost of the biopsy, we use

$$C_{\text{Biopsy}} = \$30K$$

This utility model is inappropriate for most medical decisions, including this one. Utility models appropriate for medicine can be found in [**?**], [**?**], and [**?**].

Let us assume that, given the clinical information available to the patient's primary care physician, $p(\text{Malignant}) = 0.1$. From Equations 4 and 5, we have

$$C = \$0 - (-\$100K) = \$100K \qquad B = -\$300K - (-\$800K) = \$500K$$

where C and B are the cost and benefit of treating the patient, respectively. Thus, from Equation 3, we obtain

$$p^* = \frac{\$100K}{\$100K + \$500K} = \frac{1}{6}$$

where p^* is the probability above which the patient should be treated. Consequently, from Equation 10, we have

$$W^* = \ln \frac{1/6}{5/6} - \ln \frac{0.1}{0.9} = 0.588$$

The patient should be treated if and only if W—the weight of evidence that the patient has a malignant disease—exceeds this value of W^*.

Figure 6 is a plot of $p(W > W^*|\text{Malignant})$ and $p(W > W^*|\text{Benign})$ as a function of W^*, assuming that all of the features in Figure 5 are observed. The curves labeled "exact" show the exact values; the curves labeled "approx" show the values obtained from the central-limit-theorem approximation with the generalizations for nonbinary and dependent features described in Section 5. Note the goodness of the approximation with only eight observed features. With $W^* = 0.588$, the approximate values for $p(W > W^*|\text{Malignant})$ and $p(W > W^*|\text{Benign})$ obtained from the approximation are

$$p(W > W^*|\text{Malignant}) = 0.923 \qquad p(W > W^*|\text{Benign}) = 0.028$$

The inequality $p(W > W^*|\text{Malignant}) > p(W > W^*|\text{Benign})$ states that it is more likely for the evidence to suggest a malignancy when the patient has a malignancy then when the patient has a benign disease—a reasonable result.

[3]The full specification of the influence diagram, including probabilities, is available from the first author.

Insert Figure 6 about here.

We can now compute the net value of information for a biopsy that permits the observation of all features in Figure 6. From Equations 25, 27, and 29, we have

$$EU(\text{Biopsy}|\text{Malignant}) = (0.923)(-\$300K) + (0.077)(-\$800K) = -\$338K$$

where $EU(\text{Biopsy}|\text{Malignant})$ is the expected utility of obtaining the biopsy, given that the patient has a malignant disease. Similarly, from Equations 26, 28, and 30, we obtain

$$EU(\text{Biopsy}|\text{Benign}) = (0.028)(-\$100K) + (0.972)(\$0) = -\$3K$$

where $EU(\text{Biopsy}|\text{Benign})$ is the expected utility of obtaining the biopsy, given that the patient has a benign disease. Thus, from Equation 24, we have

$$EU(\text{Biopsy}) = (0.1)(-\$338K) + (0.9)(-\$3K) = -\$36K$$

To obtain the patient's expected utility without a biopsy, $EU(\emptyset)$, we apply Equation 20, with $p < p^*$.

$$EU(\emptyset) = (0.1)(-\$800K) + (0.9)(\$0) = -\$80K$$

Consequently, from Equation 19, the value of information of the biopsy, $VI(\text{Biopsy})$, is given by

$$VI(\text{Biopsy}) = -\$36K - (-\$80K) = \$44K$$

Finally, from Equation 18, we have

$$NVI(\text{Biopsy}) = \$44K - \$30K = \$14K$$

for the net value of information of the biopsy. Because this value is greater than 0, the biopsy should be performed. We obtain the same recommendation using the exact values for $p(W > W^*|\text{Malignant})$ and $p(W > W^*|\text{Benign})$ ($NVI(\text{Biopsy}) = \$12K$).

In a myopic analysis of value of information, a biopsy would not be recommended. In particular, of all the features, "Polarity" has the greatest value of information—$VI(\text{Polarity}) = \$25K$—which is less than the cost of the biopsy.

7 More General Nonmyopic Analyses

The nonmyopic analysis described in this article is unlikely to be useful unless the dependencies among observation costs fit the model described in Section 4. Nonetheless, we can use the techniques developed in the article for more general nonmyopic analyses.

For example, suppose that n pieces of evidence are available for observation, and that the myopic analysis determines that no single piece of evidence has a positive net value of information. We may be able to identify evidence whose observation is cost effective by (1) enumerating sets of evidence whose observation are likely to be cost effective, and (2) applying our approximate analysis to each such set.

One heuristic for identifying sets of evidence whose observation are likely to be cost effective is as follows. First, arrange the pieces of evidence in descending order of their

net values of information. Specifically, label the pieces of evidence $E_1, \ldots, E_n$, such that $NVI(E_i) \geq NVI(E_j)$, if $i < j$. Then, consider subsequences of $E_1, \ldots, E_n$ that begin with E_1. That is, identify for consideration the sets $\{E_1, \ldots, E_m\}$, $m = 2, \ldots, n$.

Empirical studies are needed to determine whether this or other generalizations provide significant improvements over a myopic analysis.

8 Summary

We have described an approach using the central-limit theorem to compute the value of information for a set of tests. Our procedure provides a nonmyopic, yet tractable alternative to the traditional myopic analysis for determining the next best piece of evidence to observe. Our approach is limited to information-acquisition decisions for problems involving specific classes of dependencies among evidence variables, and binary hypothesis and action variables. Nonetheless, as we have demonstrated, the approach can offer an improvement over the myopic analysis.

Acknowledgments

Edward Herskovits constructed the influence diagram in Figure 5 using data generated by Bharat Nathwani and David Heckerman. This work has been supported by the National Cancer Institute under Grant RO1CA51729-01A1, and by the Agency for Health Care Policy and Research under Grant T2HS00028.

Figures

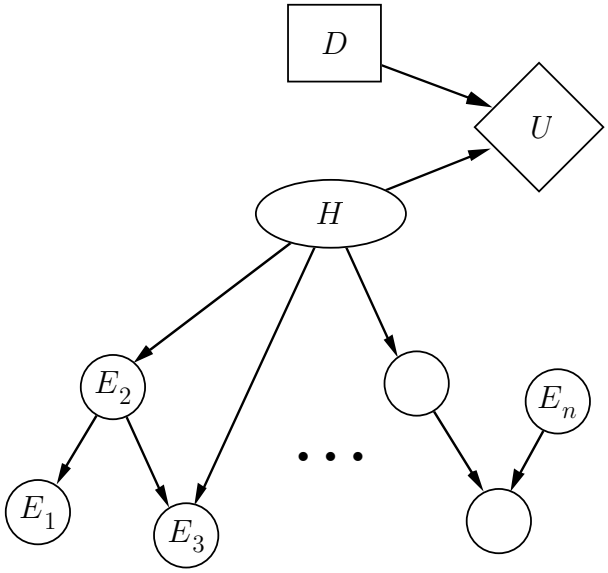

Figure 1: The Pathfinder influence-diagram for diagnosis. The decision-maker's utility (diamond node, U) depends on a hypothesis (oval node, H) and a decision (square node, D). The variables E_i are pieces of evidence or tests about the true state of H.

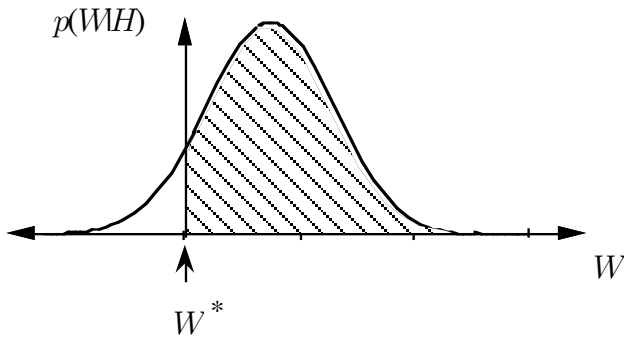

Figure 2: The probability that the total weight of evidence will exceed the threshold weight is the area under the normal curve above the threshold weight W^* (shaded region).

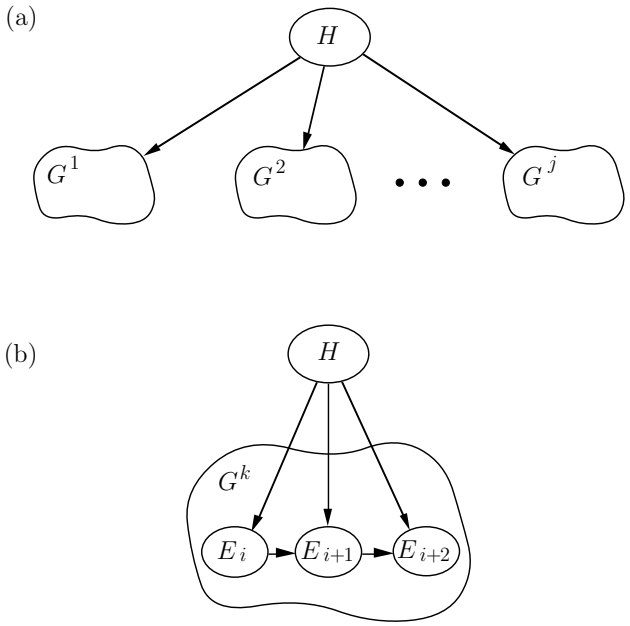

Figure 3: A schematic belief network for Pathfinder. (a) The features in Pathfinder can be arranged into groups of evidence variables $G^1, G^2, \ldots, G^j$. The variables within each group are dependent, but the groups are conditionally independent, given the disease variable H. (b) A detailed view of the evidence variables E_i, E_{i+1}, and E_{i+2} within group G^k.

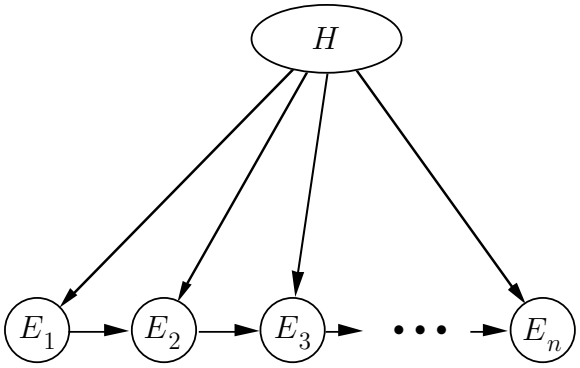

Figure 4: A conditional Markov chain. The evidence variables form a Markov chain conditioned on the variable H. We can extend our analysis involving the central-limit theorem to this case.

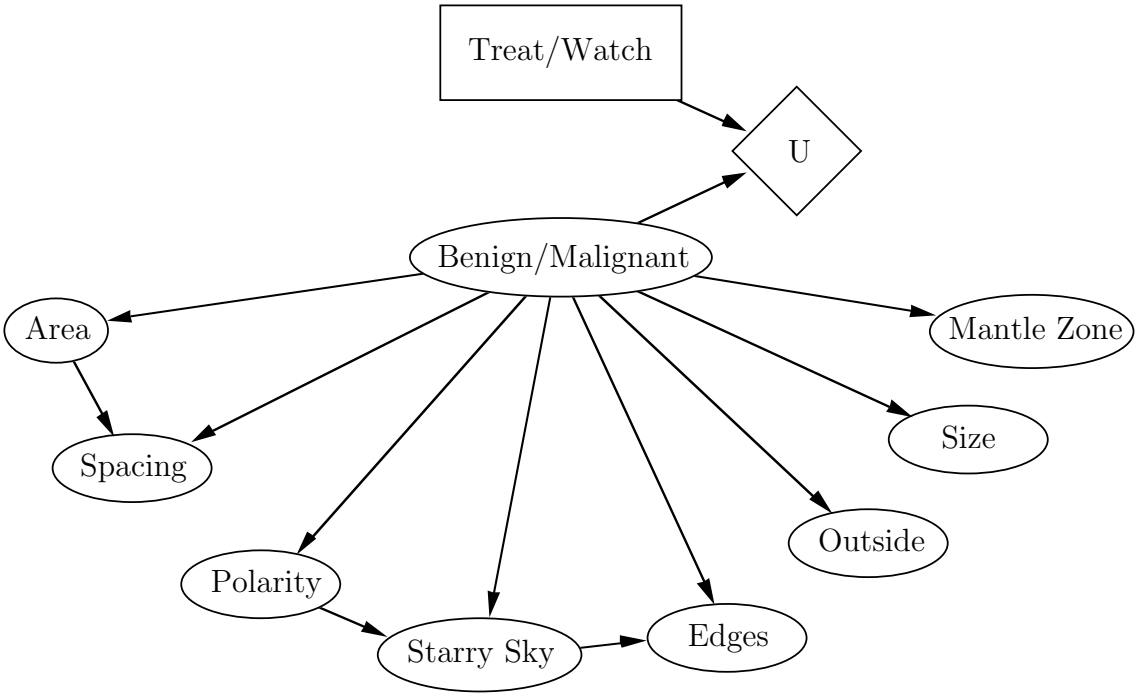

Figure 5: An influence diagram for a subset of lymph-node diagnosis. The hypothesis node represents whether the patient has a malignant or benign disease. The decision node represents the two alternatives: treat and watch. The node U represents the patient's utility for the four possible outcomes. The evidence variables represent follicular features that are clues about the disease state of the patient.

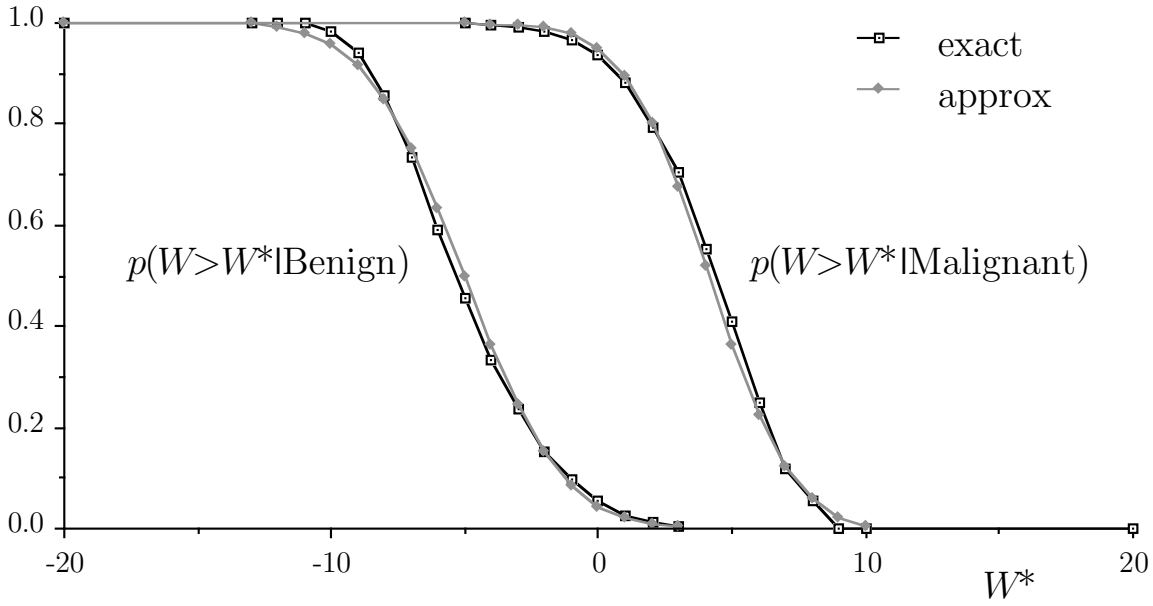

Figure 6: A plot of $p(W > W^*|\text{Benign})$ and $p(W > W^*|\text{Malignant})$ as a function of W^*, showing both the the exact and approximate values.