RATIO-SCALE ELICITATION OF DEGREES OF BELIEF

Shimon Schocken
Department of Information Systems
NYU Stern School of Business
Room 9-80, 44 West 4th Street
New York, NY 10012
(212) 998-0841
(212) 995-4228 (fax)
sschocke@stern.nyu.edu

September 1988

Working Paper Series
STERN IS-88-95

Abstract

Most research on rule-based inference under uncertainty has focused on the normative validity and efficiency of various belief-update **algorithms**. In this paper we shift the attention to the **inputs** of these algorithms, namely, to the degrees of beliefs elicited from domain experts. Classical methods for eliciting continuous probability functions are of little use in a rule-based model, where propositions of interest are taken to be causally related and, typically, discrete, random variables. We take the position that the numerical encoding of degrees of belief in such propositions is somewhat analogous to the measurement of physical stimuli like brightness, weight, and distance. With that in mind, we base our elicitation techniques on statements regarding the **relative** likelihoods of various clues and hypotheses. We propose a formal procedure designed to (a) elicit such inputs in a credible manner, and, (b) transform them into the conditional probabilities and likelihood-ratios required by Bayesian inference systems.

1. Introduction

Most rule-based expert systems fall in the category of deductive, inexact, classification models: given a set of observable clues, $M' \subseteq M$, an inference engine algorithm attempts to discern a set of hypotheses, $H' \subseteq H$, which provides the best explanation to M'. M and H are sets of propositions related to each other through inexact inference rules. For example, consider the following reasoning chain: the act of smoking (a disposition) increases the likelihood of a heart disease (an hypothesis), which, in turn, is sometimes manifested through a chest ache radiating to the left arm (a manifestation). This line of reasoning is plausible, but not necessarily categorical; many smokers will not develop any heart problems; likewise, chest ache is not a unique manifestation of heart disease. Hence, although causal information is indeed useful, any inference drawn from it must be qualified by the impreciseness of the underlying rules.

During the past decade, a number of models were put forward to represent and carry out rule-based inference under uncertainty. Chief among these are rule-based belief-update algorithms (Duda and Shortliffe, 1977), influence diagrams (Howard and Matheson, 1981), and belief networks (Pearl, 1986). These models are closely related to each other at the elicitation level, requiring human experts to specify a coherent set of degrees of belief reflecting the uncertainty associated with inference rules. This elicitation task, which is normally delegated to a knowledge engineer, is the focus of this paper. Throughout the paper, we use the **belief network** paradigm as our working environment; at the same time, our results are equally applicable to any inference model involving inexact rules.

A belief network is an acyclic, directed graph, consisting of propositional nodes and causal arcs. The directed arc (x,y) emanating from node x to node y represents our belief that x causes y directly. The strength of this causal relationship is modeled through the conditional probability $P(y|x)$. If a node y has multiple causes, $\{x_1,...,x_n\}$, the degree of belief associated with this complex relationship is the conditional probability tensor $P(y|x_1,...,x_n)$.

A belief network can be encoded as a set of inexact inference rules. In the rule-based terminology, an arc (x,y) and its label $P(y|x)$ correspond to the rule IF x THEN y WITH DEGREE OF BELIEF $P(y|x)$. These rules are elicited from human experts. Automatic inference is carried out by following reasoning chains from the observable evidence, M', back to its possible explanations, H'. This backward reasoning process is accompanied by a belief-update algorithm designed to order the prospective hypotheses $h' \in H'$ in terms of the posterior beliefs $P(h'|M')$. Unfortunately, this belief-update procedure is exponential in the size of the network, and, in general, is NP-hard (Cooper, 1987). If, however, the topology of the network meets certain criteria, posterior beliefs may be computed efficiently (Pearl, 1986).

A belief network is constructed through an elaborate knowledge elicitation procedure involving domain experts and knowledge engineers. Now, humans are normally good at suggesting causal relationships between evidence and hypotheses. At the same time, humans have serious problems in estimating the uncertainty associated with such conjectures. For example, a physician can swiftly suggest that the cause of swollen ankles (m) might be a certain heart disease (h). However, this same expert might be at loss when asked to estimate the subjective probability associated with this rule, $P(m|h)$. When pressed to do so, the expert will probably produce a number, but the validity of this estimate is clearly questionable. We propose a global approach to elicitation which minimizes the guesswork and yields a credible set of degrees of belief.

2. The Problem

Unlike early rule-based architectures, the belief network model has a sound interpretation on deductive as well as on probabilistic grounds. From a deductive standpoint, the network's nodes are viewed as propositions, and the network's topology reflects causal relationships among these propositions. From a probabilistic standpoint, these nodes are viewed as random variables associated with an unknown joint probability distribution function (PDF). The network's topology is analogous to a set of independence assumptions imposed on the PDF. Generally speaking, if two nodes are not connected, the random variables that they represent are assumed independent. These assumptions simplify the PDF considerably; hence, they also place an explicit constraint on the family of PDF's that might be modeled as belief networks.

Consider, for example, the simple network depicted in Figure 1, taken from the domain of diagnosing heart diseases. This network has the following interpretation: the disease h may be caused by either c1, c2, or both. h, in turn, manifests itself through subsets of the symptoms {m1,m2,m3}. These relationships are taken to be inexact. For example, it is possible that both c1 and c2 obtain but h does not obtain. Likewise, it is possible (but not likely) that h obtains and none of the m's obtain.

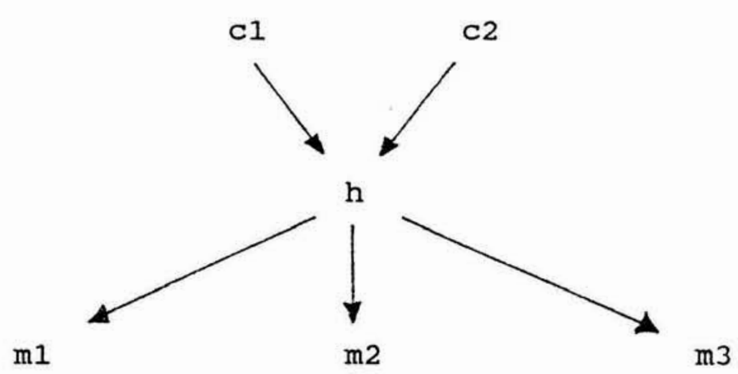

Figure 1

Hypothesis:

h: heart disease

Predispositions:

c1: high Cholesterol count
c2: smoking

Symptoms:

m1: chest ache radiating to the left arm
m2: swollen ankles
m3: shortness of breathing

How can we elicit and represent the PDF $P(c_1,c_2,h,m_1,m_2,m_3)$? one naive approach is to set up a global contingency table which specifies the joint frequencies of all possible combinations of propositions. Clearly, this is not a practical solution. Alternatively, if the propositions in question are arranged in a network like Figure 1, one can interpret the topology of the network as a graph-theoretic definition of the PDF:

$$P(c1,c2,h,m1,m2,m3) = P(m1|h)*p(m2|h)*P(m3|h)*P(d|c1,c2)*P(c1)*P(c2) \quad (1)$$

This derivation is based on a standard theorem in probability, the "chain rule," along with a set of marginal and conditional independence assumptions imposed on the PDF by the network's topology (Pearl, 1986). Pearl has shown that a PDF like (1) is suspectable to an efficient belief-update algorithm which is consistent with probability theory. This algorithm computes posterior beliefs in hypotheses given any subset of observable clues (terminal nodes in the network), in time linear to the network's size. The algorithm involves only local computations, and, in principle, can proceed in parallel.

We now turn to the elicitation problem: in order to fully specify the belief network associated with a PDF like (1), the knowledge engineer must elicit three types of probabilities:

(a) marginal probabilities associated with nodes with no parents, e.g. P(c1) and P(c2).

(b) conditional probability tensors associated with multi-cause relationships, e.g. P(h|c1,c2).

(c) elementary conditional probabilities, e.g. P(m1|h), P(m2|h), and P(m3|h).

The elicitation of conditional probability <u>tensors</u> is largely impractical, as the number of questions that one is required to ask grows exponentially with the number of propositions following the conditioning bar. There exist heuristic techniques, though, designed to approximate $P(h|c_1,...,c_n)$ from the set of elementary probabilities $P(h|c_1),...,P(h|c_n)$ (Kim and Pearl, 1987). Hence, we see that the elicitation problem is primarily one of assessing <u>elementary</u> conditional probabilities. Therefore, we'll restrict our attention for now to the bottom tier of Figure 1, focusing on the relationships among h and its 3 manifestations. Our analysis can be easily extended to any number n>3 of propositions, so we'll sometimes use n=3 for brevity.

3. Two Propositions on Elicitation

Our approach to elicitation is motivated by two key propositions regarding the direction and ratio-scale properties of implicit degrees of belief. These propositions are briefly discussed below.

Proposition 1: some questions are easier to answer if you turn them around.

Consider the rule <IF m THEN h WITH DEGREE OF BELIEF d(m,h)>. This is sometimes referred to as "abductive" or "backward" reasoning. That is, although the causal relationship between the hypothesis h and its manifestation m is h->m, we are typically faced with the problem of assessing the likelihood of the unknown h in light of the observable fact m. The interpretation of the belief function $d(.|.)$ depends on our choice of a belief language. Most of these languages, though, are "unidirectional." That is, they consist of either diagnostic or causal inference, but not of both. For example, the certainty factors calculus and the Dempster-Shafer model require the expert to specify **diagnostic** degrees of belief in terms of $CF(h|m)$ and $Bel(h|m)$, respectively. Conversely, the Bayesian belief network specifies the causal relationship "h causes m" directly, requiring the expert to estimate the **causal** degree of belief $P(m|h)$, P being a probability.

Indeed, there exist a growing body of literature suggesting that humans find it easier to "think forward in reverse," preferring causal on diagnostic explanations of evidential reasoning (Einhorn and Hogarth, 1987, Shachter and Heckerman, 1987). These findings clearly render cognitive justification to the Bayesian belief network formalism. At the same time, the insistence that in any given situation human reasoning is confined to proceed in only one direction seems to be overly restrictive. For example, consider the rule <IF x smokes THEN it is likely that x will develop a heart disease>. Denote this rule by S -> HD. Which subjective probability is more credibly available from a human expert, $P(HD|S)$ or $P(S|HD)$?

The answer seems to depend largely on the experience of the expert, and, in particular, on his or her ability to retrieve examples from the S and HD populations. If the expert knows relatively far more smokers than she knows people with heart diseases, it is probably safer to use the smokers population as a reference group (Figure 2-a) and go on to assess $P(HD|S)$. If, alternatively, we force this expert to specify $P(S|HD)$, she will have to resort, in her mind, to a small sample (Figure 2-b), yielding a highly unreliable estimate of $P(S|HD)$. The situation changes if the expert happens to be a heart disease specialist. This latter expert will probably find it easier to assess $P(S|HD)$, due to the large sample of people with heart disease that she can retrieve from her clinical work experience.

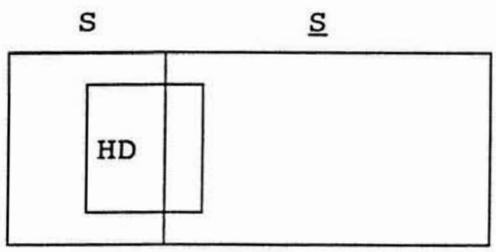

Figure 2-a

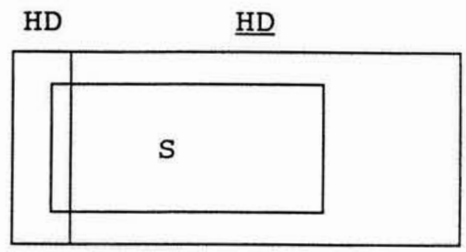

Figure 2-b

We see that the factors that influence implicit beliefs are both availability and representativeness (Tversky and Kahneman, 1974). The availability heuristic leads the expert to focus on the population which is more salient or vivid in his mind. Strictly speaking, this heuristic is beneficial only if it coincides with a larger population. If the selected background population is small, most people will still be willing to use it as a representative image of the overall population. This insensitivity to sample size, or the "law of small numbers," is a manifestation of the well-known representativeness bias. To debias this flaw, the expert should be encouraged to retrieve as many examples as possible from both populations. The larger sample should then be selected as the conditioning assumption. Hence, we wish to support experts who are willing to express their reasoning using diagnostic inference instead of or in tandem with causal inference.

Proposition 2: relative degrees of belief are more credibly available then absolute degrees of belief.

We take the position that the numerical encoding of degrees of belief is somewhat analogous to the measurement of such physical stimuli as brightness, weight, and distance. Can you specify the aerial distances between New York, Tokyo, Cairo, and Seattle with good confidence? probably not. It seems obvious that you would rather prefer a statement like: Tokyo is about twice as far away from New York as Cairo is. In general, people find it easier to express differences between physical quantities (as well as preferences) using relative, pair-wise judgement, rather than a cardinal scale of measurement (Stevens, 1957, Stevens and Galanter, 1964, Krantz, 1973). Going back to the previous question, let's take New York as our point of departure, and consider the following pair-wise distance comparisons:

New York	Tokyo	Cairo	Seattle
Tokyo	1	2	3
Cairo		1	2
Seattle			1

The entries in this table read as follows: from New York, Tokyo is 2 times as far as Cairo, and 3 times as far as Seattle. Also, Cairo is 2 times as far as Seattle. Note that these judgement are not only inaccurate but also inconsistent, as is normally the case with human inputs (the inconsistency may be resolved, for example, by setting the third entry in the second row to 1.5). How can we synthesize this set of n(n-1)/2 pair-wise comparisons into a vector of n weights reflecting the true distances implicit in the human's answers? Several such methods have been proposed, e.g. least squares (Cogger and Yu, 1983), and logarithmic least squares (De Graan, 1980). However, perhaps the only method capable of synthesizing **inconsistent** inputs in a credible manner is the Eigenvector method proposed by T.L. Saaty (1980). This method is illustrated in the next section.

3. One-Way Elicitation

To restate the elicitation problem, consider an hypothesis h and a series of related manifestations $\{m_1,...,m_n\}$. Our goal is to estimate the probabilities vector $P=<P(m_1|h),...,P(m_n|h)>$ using pair-wise comparisons, elicited from a human expert. That is, rather than asking the expert to specify the absolute probabilities $P(m_i|h)$ and $P(m_j|h)$, we ask her to estimate the extent to which mi is more likely than mj in light of h. These local judgments, which are likely to be inconsistent, will be further synthesized into a ratio-scale of probabilities.

To illustrate, consider the following set of questions:

> Assume that a person X has a heart disease (h).
> Now consider the following two observations:
> X suffers from a chest ache radiating to the left arm (m1)
> X suffers from swollen ankles (m2).
>
> In your opinion, which observation (m1,m2) is more likely in light of h? ___________ (assume the expert answered m1)

> To what extent is m1 more likely than m2? For example, if you think that m1 and m2 are equally likely, enter the number 1. If you think that m1 is twice as likely as m2, enter the number 2. Feel free to enter any number greater than 1 that you think describes the extent to which m1 is more likely than m2: ______

With n manifestations, we have to ask the expert n(n-1)/2 such questions. The expert's answers are recorded in an nxn likelihood matrix, A, in which $a_{ij}=P(m_i|h)/P(m_j|h)$. This approach is similar to Gupta and Wilson's (1987), who had experts express their opinions regarding the performance of competing forecast models. These inputs were recorded in a matrix A in which the element a_{ij} was the perceived odds that forecast model i will outperform forecast model j.

The cognitive complexity of the elicitation may be somewhat mitigated if we ask the expert to first rank-order the manifestations in terms of decreasing perceived likelihoods, yielding (without loss of generality) an ordered set $<m_1,...,m_n>$ with $P(m_i|h)>P(m_j|h)$ if $i<j$. This order sets up (an empty) likelihood matrix A in which (a) only the entries above the diagonal have to be specified, (b) all of these entries must be greater than or equal to 1, and, (c) the relation $a_{ij}<a_{ik}$ must obtain for all i and $j<k$.

The resulting matrix, A, is positive reciprocal, with $a_{ij}>0$ and $a_{ij}=1/a_{ji}$. We are thus in the familiar domain of Saaty's analytic hierarchy process, and we can proceed to compute the principal Eigenvector of A, denoted W. The crux of this approach is the notion that W is "ratio equivalent" to the unknown P, with $w_i/w_j = P(m_i|h)/P(m_j|h)$ for all pairs $<i,j> \in \{1,...,n\} x \{1,...,n\}$.

For example, suppose n=3 and the <u>true</u> probabilities vector $P=<P(m1|h),P(m2|h),P(m3|h)>$ is as follows:

$$P = <0.8,0.2,0.1> \quad (2)$$

Had we had access to an expert whose judgement are perfectly calibrated with reality, we would end up with the following matrix:

A =

h	m1	m2	m3
m1	1	4	8
m2	1/2	1	2
m3	1/8	1/2	1

Each entry a_{ij} in the matrix represents the degree to which the manifestation m_i is more likely than the manifestation m_j in light of h, in the expert's mind. For example,

$$a_{13} = P(m_1|h)/P(m_3|h) = 0.8/0.1 = 8$$

Since A in the above ideal example is consistent, its principal Eigenvector is given by any of its columns. Thus, focusing for example on the last column, we obtain the following weights vector:

$$W = \langle 8, 2, 1 \rangle \quad (3)$$

Due to A's consistency, the true probabilities vector P (2) and the weights vector W (3) are ratio-identical. In reality, though, the vector P is unknown, and the matrix A is inconsistent. That is, the set of answers to the n(n-1)/2 questions presented to the expert will yield intransitive responses, with

$$\frac{P(m_i|h)}{P(m_j|h)} * \frac{P(m_j|h)}{P(m_k|h)} \neq \frac{P(m_i|h)}{P(m_k|h)}$$

Hence, in the general case, we have to deal with an inconsistent matrix. To illustrate, refer to the A matrix, and assume that the expert's above-diagonal judgement a_{12}, a_{13}, and a_{23} were +25%, -25% and +50% off mark (with respect to A), respectively, yielding the following inconsistent matrix:

A' =

h	m1	m2	m3
m1	1	5	6
m2	1/5	1	3
m3	1/6	1/3	1

The normalized principal Eigenvector of this matrix is:

$W' = <0.717, 0.195, 0.088>$

Where do we proceed from here? the key assumption underlying our approach is that in spite of A's inconsistency, W' is still a good "ratio-estimate" of the unknown probabilities vector P. Thus, had we had a-priori knowledge, say, that $P(m_2|h)=0.2$, we could have used it as an anchor to compute the following estimate P' of P:

$$P' = \left\langle \frac{0.717}{0.195} * 0.2 \ , \ 0.2 \ , \ \frac{0.088}{0.195} * 0.2 \right\rangle$$

$$= <0.73, 0.2, 0.09>$$

The difference between this result and the true probabilities vector $P=<0.8,0.2,0.1>$ stems from the inconsistency of A' and the imperfect knowledge that it represents. However, considering the expert's biasdeness and inconsistency, this result is surprisingly good, illustrating the robustness of the normalized Eigenvector to data perturbations.

Of course, the ability to construct P' hinges on our a-priori knowledge of any one of its underlying probability elements. In some situations, we may anticipate this requirement ex-ante and augment the initial set of manifestation $<m_1,...,m_n>$ with an additional clue, m^*, whose conditional probability $P(m^*|h)$ is either known or can be credibly estimated. For example, in the above heart disease scenario, m^* might be the fact that the underlying patient is a male, the assumption being that the probability P(male|heart disease) is well-known.

There exist situations, however, in which we have no prior knowledge of any one of the underlying degrees of belief. In these cases, we can use an extension of the above technique in order to estimate the likelihood-ratio vector $<P(m_1|h)/P(m_1|\underline{h}),...,P(m_n|h)/P(m_n|\underline{h})>$, $\underline{h}$ being "not h." Methods to compute such vectors for dichotomous and multi-valued propositions are described in the next section. We conclude the present section with a brief comment regarding the scale of measure used throughout the paper.

Our approach to elicitation is based on the assumption that humans are capable of describing relative likelihoods using numbers. This controversial assumption was challenged by many, not the least of them is H.R. Haldeman, Chief of Staff of President Nixon. Describing Kissinger's persistent concern about a Russian attack on China, Haldeman recalls how "I used to tease him about his use of percentages. He would say there was a 60% chance of a Soviet strike on China, for example, and I would say: why 60, Henry? Couldn't it be 65% or 58%? (Kotz and Stroup, 1983). Clearly, Haldeman's point is well taken. Although there exist evidence that people find it easier to assess odds rather than probabilities, the credibility of **any** numeric measure of intuitive judgement is an open question.

It was this motivation which led to Saaty's (1980) 1 to 9 scale, which is often accompanied with a verbal interpretation. This scale can be slightly modified to suit our elicitation context, as follows:

> 1 - propositions p and q are **equally** likely
> 3 - p is **weakly more** likely than q
> 5 - p is **strongly more** likely than q
> 7 - p is **very strongly more** likely than q
> 9 - p is **absolutely more** likely than q
>
> The numbers 2,4,6, and 8 are used to express intermediate judgement between adjacent scale values.

This scale has been justified by Saaty on analytical, psychological, and experimental grounds. In a similar vein, Lichtenstein and Newman (1967) have shown empirically that verbal descriptions of uncertainty may be mapped on ranges of probabilities. These findings are especially relevant to the Eigenvector method, which is insensitive to the type of scale being used. As Harker and Vargas (1987) argue, "One scale may be appropriate for one application and may not be appropriate for another. In this situation, a different scale could and should be chosen for each application." For example, in situations where little is known about a particular set of hypotheses, a 1 to 3 scale might be used. Clearly, the freedom to modify the scale of measure or develop a totally new one adds significant flexibility to the elicitation task.

4. Two-way Elicitation

Most realistic inference problems involve multi-valued propositions. For example, the dichotomous "heart disease" hypothesis can be made finer by considering the following possibilities:

h_1: X is likely to suffer a fatal stroke
h_2: X is likely to suffer a mild stroke
h_3: X has no heart disease

From a logical perspective, to say that h assumes one of the values in $H=\{h_1,h_2,h_3\}$ is equivalent to assigning truth values to the three dichotomous propositions h_1, h_2, and h_3. We will assume henceforth that H enumerates all the possible values that h can attain and that these values are mutually exclusive.

Using the techniques described in the previous section, we can condition our set of causal elicitation questions on each value of h, yielding three matrices and three weight vectors W_{h1}, W_{h2}, and W_{h3}, where W_{hi} is conditioned by the background hypothesis h_i. Now, these three matrices and respective vectors are disjoint: one cannot use them to calculate likelihood-ratios <u>across</u> competing hypotheses. That is, although we can calculate the ratios, say, $P(m_i|h_1)/P(m_j|h_1)$ and $P(m_i|h_2)/P(m_j|h_2)$ <u>within</u> the vectors Wh1 and Wh2, there is no sufficient information to compute the more useful likelihood-ratios $P(m_i|h_1)/P(m_i|h_2)$ and $P(m_j|h_1)/P(m_j|h_2)$. These ratios were termed by Alan Turing the "weights of evidence" carried by mi and mj, respectively, to the statement "h_1 is preferred on h_2" (Good, 1950). These weights of evidence play an important role in many Bayesian belief-update algorithms.

Recall that the ultimate objective of any belief-update algorithm is to compute the posterior beliefs in the competing hypotheses H in light of $\{m_1, \ldots, m_n\}$. If these manifestations are ratio-independent with respect to H (Grosof, 1986), the posterior odds favoring h_i on h_j is given in terms of Bayes rule, as follows:

$$\frac{P(h_i|m_1,\ldots,m_n)}{P(h_j|m_1,\ldots,m_n)} = \frac{P(m_1|h_i)}{P(m_1|h_j)} * \ldots * \frac{P(m_n|h_i)}{P(m_n|h_j)} * \frac{P(h_i)}{P(h_j)} \qquad (4)$$

This formula has the following "mechanical" interpretation: in the absence of any relevant evidence, the posterior odds are set to the prior odds, $P(h_i)/P(h_j)$. When we know that a certain manifestation is present, say m_k, we multiply these odds by the likelihood-ratio $P(m_k|h_i)/P(m_k|h_j)$. As more manifestations become available, the posterior odds are updated in a similar fashion.

Hence, what we are after are likelihood-ratio vectors of the form

$$\left\langle \frac{P(m_1|h_i)}{P(m1|h_j)}, \ldots, \frac{P(m_n|h_i)}{P(m_n|h_j)} \right\rangle \qquad (5)$$

These vectors must be specified for all pairs of competing hypotheses $<h_i,h_j> \in HxH$.

How do we go about computing these vectors? we propose a solution, called two-way elicitation. The first stage of this procedure is identical to the one-way elicitation described in the previous section. Having completed this line of causal questions, we ask the expert to go through a second stage of diagnostic questions. The expert's answers are then synthesized into a set of n(n-1)/2 likelihood-ratio vectors like (5). For now, we assume that the second stage involves roughly the same number of questions as the first stage. As we'll see shortly, the second stage is far less demanding.

A typical diagnostic question has the following form:

> Assume that a person X suffers from swollen ankles (m2). Now consider the following two possibilities: X will suffer a mild heart attack (h2), X has no heart disease (h3).
>
> In your opinion, which possibility (h2,h3) is more likely in light of m2 ___________ (assume the expert answered h2)
>
> To what extent is h2 more likely than h3? For example, if you think that h2 and h3 are equally likely, enter the number 1. If you think that h2 is twice as likely as h3, enter the number 2. Feel free to enter any positive real number that you think describes the extent to which h2 is more likely than h3: ______

With the manifestation m2 fixed, three such pair-wise comparisons are required to construct a likelihood matrix and a diagnostic weight vector, denoted W_{m2}. W_{m2} is taken to be ratio-equivalent to the probabilities vector <P(h1|m2),P(h2|m2),P(h3|m2)>. Two-way elicitation thus consists of two independent elicitation stages: in the first stage (described in the previous section), the three causal weight vectors W_{h1}, W_{h2}, W_{h3} are computed. The second stage yields the three diagnostic vectors W_{m1}, W_{m2}, W_{m3}. The restriction of n=3 can be easily lifted, and the number of hypotheses and manifestations need not be equal. We now show how these disjoint sets of vectors can be synthesized into the n(n-1)/2 likelihood-ratio vectors (5).

The data gathered in the two-way elicitation procedure can be represented in two related, directed graphs (Figure 3). The nodes M_{ij} and H_{ij} represent the unknown probabilities $P(m_i|h_j)$ and $P(h_i|m_j)$, respectively. The directed arcs (M_{ij}, M_{kl}) and (H_{ij}, H_{kl}) are labeled M_{ijkl} and H_{ijkl}, respectively. Each triangle H_k in the causal graph represents the causal weights vector W_{hk}, k=1,2,3, with $M_{ikjk}=P(m_i|h_k)/P(m_j|h_k)$. As was argued earlier, these ratios are not particularly useful in their present form, and we are more interested in the likelihood-ratios $M_{ijik}=P(m_i|h_j)/P(m_i|h_k)$, measuring the degree to which the manifestation m_i serves to discriminate between the hypotheses h_j and h_k. In Figure 3, M_{ijik} is the label of the "external" arc connecting M_{ij} and M_{ik} across the two triangles Hj and Hk.

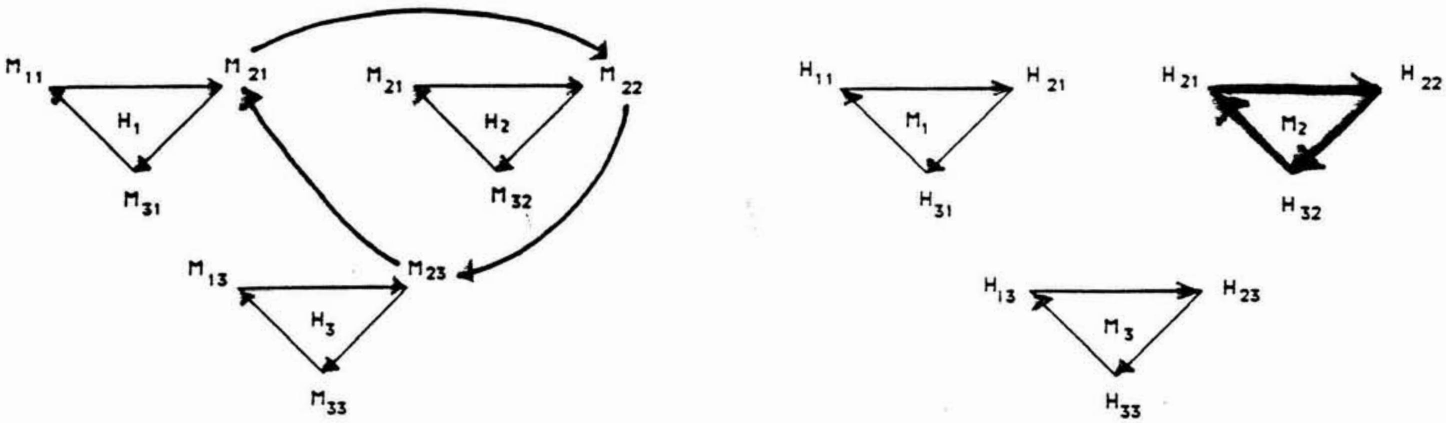

Causal graph Diagnostic graph

Figure 3

Let's focus on the following two sets of labels:

$M_{ijik}=P(m_i|h_j)/P(m_i|h_k)$ (external, causal arcs) (6)

$H_{jiki}=P(h_j|m_i)/P(h_k|m_i)$ (internal, diagnostic arcs) (7)

Also, denote the prior odds favoring hypothesis h_k on h_j by:

$$O_{kj} = P(h_k)/P(h_j) \tag{8}$$

Given this notation, (6), (7), and (8) are related through Bayes rule:

$$M_{ijik} = H_{jiki} * O_{kj} \tag{9}$$

Hence, we see that the two graphs mirror each other. That is, every cross-triangle arc M_{ijik} at the causal graph is proportionally related to a within-triangle arc H_{jiki} at the diagnostic graph. For example, the bold circuit (triangle M_2) of the diagnostic graph is the "dual image" of the external circuit connecting all M_{2k} nodes, k=1,2,3 across all triangles in the diagnostic graph. Hence, had we had access to a prior weights vector $W_0=\langle w_1,...,w_n \rangle$ which is ratio-equivalent to the prior probabilities vector $\langle P(h_1),...,P(h_3) \rangle$, we could have used (9) to compute the desired likelihood-ratio vectors $\langle P(m_1|h_i)/P(m_1|h_j),...,P(m_n|h_i)/P(m_n|h_j) \rangle$ for any pair of competing hypotheses $\langle h_i,h_j \rangle$ HxH.

It turns out that one half of the two-way elicitation procedure can be cut down considerably, due to some favorable graph properties. To illustrate, consider Figure 3. Once we "construct" any <u>one</u> triangle in the **diagnostic** graph, we can compute its dual external circuit in the **causal** graph, through (9). And, due to the topology of the causal graph, any such single circuit makes the entire causal graph connected, meaning that there is a path between any two given nodes. Let us define the "intensity" of a directed path as the product of all the labels along the path. We immediately get, from the label's definition, that this is a telescopic product. For example, the intensity of the path <(M31,M11),(M11,M21),(M21,M22),(M22,M32)> is:

$$\frac{P(m3|h1)}{P(\cancel{m1|h1})} * \frac{P(\cancel{m1|h1})}{P(\cancel{m2|h1})} * \frac{P(\cancel{m2|h1})}{P(\cancel{m2|h2})} * \frac{P(\cancel{m2|h2})}{P(m3|h2)} = \frac{P(m3|h1)}{P(m3|h2)}$$

Similarly, the intensity of a loop is 1. This implies that the intensities of all the paths connecting a given pair of nodes are the same. This is true because the paths are directed, and the only difference between two paths with the same end-points is that one might consist of one or more loops. However, the intensity of these loops is 1, so they have no effect on the overall path intensity.

Let us illustrate the relevance of this analysis to the elicitation problem. Suppose that we have gone through the first stage of the elicitation procedure, yielding the three triangles H1, H2, and H3. We now turn to the diagnostic, second stage of the elicitation. After reviewing the various manifestations, we find out that M2, say, is the most reliable background population (recall proposition 2). Therefore, we proceed to compute the diagnostic weight vector W_{m2}, which is represented by the bold triangle M2 in Figure 3. At that stage, if we have access to the prior weights vector, W_0, we can proceed to construct the circuit connecting all M_{2k} nodes, k=1,2,3, in the causal graph. By "construct" we mean that we can now compute the circuit's labels {M2122,M2223,M2321}, which, in turn, allow us to compute any likelihood-ratio pertaining to m_2. Moreover, we can now proceed to compute <u>any</u> other desired likelihood-ratio $M_{ikil}=P(m_i|h_k)/P(m_i|h_l)$, by simply taking the intensity of any path connecting the nodes M_{ik} and M_{il}. Since the causal graph is connected, at least one such path must exist. And due to the intensity uniqueness, this solution is also unique.

Estimating prior probability vectors: We now turn to discuss the elicitation of the prior weights vector $W_0=\langle w_1,\ldots,w_n\rangle$ which is assumed to be ratio-identical to the prior probability vector $\langle P(h_1),\ldots,P(h_n)\rangle$. In some situations, e.g. when the $\{h_i\}$ represent well-known diseases, this vector can be estimated from textbook information, field records, or other relevant background knowledge. In other situations, though, we have to elicit W_0 from a domain expert. In the latter case, we can simply use the one-way elicitation procedure described in Section 3. That is, the expert will be asked to compare the relative likelihoods of all pairs $\langle h_i,h_j\rangle \in HxH$, forming a likelihood matrix. The desired vector, W_0, will then be taken to be the Eigenvector of this matrix.

5. Conclusion

Our elicitation approach relies heavily on the Eigenvector method, the cornerstone of Saaty's Analytic Hierarchy Process (AHP). As Harker and Vargas (1987) put it, the AHP framework is designed to cope with intuitive, rational, and irrational judgement, with and without certainty. It is thus natural, in our opinion, to apply it to the problem of eliciting degrees of belief, where rational knowledge is often combined with intuitive guts feeling and, on occasion, with inconsistent judgement. It was argued by Fischhoff et al (1980) that it is inappropriate to think of a person's opinion about a set of events as existing within that person in a precise, fixed fashion, just waiting to be measured. And yet, asking experts to provide numeric degrees of belief and adding this information verbatim to a knowledge-base is a common practice among many knowledge engineers. We think that the elicitation problem deserves a more serious treatment. This paper is a step in this direction.

6. References

Cogger, K.O. and Yu, P.L., "Eigen Weight Vectors and Least Distance Approximation for Revealed Preference in Pair-Wise Weight Ratios," School of Business, University of Kansas, 1983.

Cooper, G.F., "Probabilistic Inference Using Belief Networks is NP-Hard," Tech. Report KSL-87-27, Medical Computer Science Group, Knowledge Systems Laboratory, Stanford University, Stanford, CA, May, 1987.

de Graan, J.G., "Extensions of the Multiple Criteria Analysis Method of T.L. Saaty," Paper presented at EURO IV, Cambridge, July 22-25, 1980.

Duda, R. O., and Shortliffe, E. H., "Expert Systems Research," Science, 220, 1981, pp. 261-268.

Einhorn, H. J. and Hogarth, R. M., "Decision Making: going Forward in Reverse," Harvard Business Review, January-February 1987, pp. 66-70.

Fischhoff, B., Slovic, P., and Lichtenstein, S., "Knowing What You Want: Measuring Labile Values," in T.S. Wallsten (ed.), Cognitive Processes in Choice and Decision Behavior, Erlbaum Associates, Hillsdale, New Jersey, 1980.

Good, I. J., "Probability and the Weighing of Evidence," New York: Hafner, 1950.

Grosof, B. N., "Evidential Information as Transformed Probability," in Lemmer, J. F., and Kanal, L. (eds.), Uncertainty in Artificial Intelligence. North Holland, 1986.

Gupta, S., and Wilton, P., "Combinations of Forecasts: an Extension," Management Science, Vol. 33, No. 3, March 1987, 356-372.

Harker, P.T. and Vargas, L.G., "The Theory of Ratio-Scale Estimation: Saaty's Analytic Hierarchy Process," Management Science, Vol. 33, No. 11, November 1987, pp. 1383-1403

Howard, R.A. and Matheson, J.E., "Influence Diagrams," in "Readings on the Principles and Applications of Decision Analysis" (R.A. Howard and J.E. Matheson, Eds.) Strategic Decision Group, Menlo Park, CA pp. 719-762, 1981.

Kim, J.H., and Pearl, J., "CONVINCE: A Conversational Inference Consolidation Engine," IEEE Trans. on Systems, Man, and Cybernetics, Vol. SMC-17, No. 2, March/April 1987, 120-131.

Kotz, S. and Stroup, D.F., "Educated Guessing," Marcel Dekker, 1983, p. 18.

Krantz, D.H., "A Theory of Magnitude Estimation and Cross-Modality Matching," J. Math. Psychology, 9, 2 (1972), 168-199.

Lichtenstein, S. and Newman, J.R., "Empirical Scaling of Common Verbal Phrases Associated with Numerical Probabilities," Psychonomic Science, 9, 1967, 563-564.

Pearl, J., "Fusion, Propagation, and Structuring in Belief Networks," Artificial Intelligence, September, 1986.

Saaty, T.L., "The Analytic Hierarchy Process," McGraw-Hill, New York, 1980.

Shachter, R.D., and Heckerman, D.E., "Thinking Backward for Knowledge Acquisition," AI Journal, Fall, 1987, pp. 55-61.

Stevens, S. S., "On the Psychological Laws," Psychological Review, 64, 1957, pp. 153-181.

Stevens, S.S. and Galanter, E., "Ratio-Scale and Category-Scale for a Dozen Perceptual Continua," J. Experimental Psychology, 54 (1964), 377-411.

Tversky, A. and Kahneman, D., "Judgment Under Uncertainty: Heuristics and Biases," Science, 185, September, 1974, pp. 1124-1131.