How to develop and use a Bayesian Belief Network

R Montironi, W F Whimster, Y Collan, P W Hamilton, D Thompson, P H Bartels

Institute of
Pathological Anatomy
and Histopathology,
University of Ancona,
Ancona, Italy
R Montironi

Department of
Histopathology, King's
College School of
Medicine, London
W F Whimster

Institute of Pathology,
University of Turku,
Turku, Finland
Y Collan

Department of
Pathology, The
Queen's University of
Belfast, Belfast
P W Hamilton

Optical Sciences
Center, University of
Arizona, Tucson,
Arizona, USA
D Thompson
P H Bartels

Correspondence to:
Professor Rodolfo
Montironi, Institute of
Pathological Anatomy and
Histopathology, University
of Ancona, Nuovo Ospedale
Regionale, Via Conca, I-
60020 Torrette di Ancona,
Italy.

Accepted for publication
18 October 1995

Introduction

Histopathologists are accustomed to assessing histological features expressed in linguistic or descriptive terms, words and concepts and integrating them in order to reach a diagnosis in their minds. If several histopathologists review a case there is seldom 100% agreement on the diagnosis. Sometimes they will accept a majority opinion; sometimes the opinion of the most respected, experienced or forceful histopathologist prevails. It has been difficult to examine the differences in histopathologists' analyses in any meaningful way, although there have been many attempts to weight particular features both for diagnostic and prognostic purposes, including that of Bloom and Richardson in breast cancer.[1] There have been attempts to give numerical gradations to verbal descriptions, which has the effect of increasing the thoroughness with which the slides are examined, and, one hopes, the certainty with which the diagnosis is reached.

Pathologists have been trying for some time to combine their human skills in histopathological diagnosis with the advantages offered by computer systems. Expert systems are computer programs designed to process knowledge and reach diagnostic decisions in the same way as human experts. Methods which allow diagnostic knowledge to be represented within a computer system include rules,[2] artificial neural networks[3] and inference networks[4]. Of particular interest to this group is a particular implementation of inference networks called Bayesian Belief Networks (BBNs).[4] BBNs can be visualised as "nodes" connected by "links" where the nodes represent chunks of knowledge and the links represent the relation among these bits of knowledge. In their simplest form they can be represented as a series of diagnostic clues, each contributing to a single diagnostic decision. BBNs are simple to construct and when used, can provide a numerical value of belief in the diagnostic alternatives. Abmayr *et al*,[5] Hamilton *et al*[6] and Montironi *et al*[7] have published results from using BBNs in the diagnosis of skin lesions, breast lesions and prostate lesions. These papers have largely dealt with how BBNs can be applied to specific problems and we thought that a more basic paper describing how the calculations could be performed was necessary for pathologists.

The ability to make diagnostic or prognostic decisions in the presence of uncertainty is also an important characteristic of an expert system. In spite of dealing constantly with uncertainty, histopathologists may be unaware that the study of uncertainty, its nature, sources and effects on decision sequences, and numerical assessment is a field of research in its own right.[8] The management of uncertainty has two major theoretical bases, namely, probability theory and possibility theory. Probability theory is used when the uncertainty is due to randomness. Probabilities are derived from observing the relative frequency with which an event occurs (for example, in a group of patients) and using this result to predict the occurrence of the event in future patients. The assumption is made that the patients are homogeneous in every other respect. This calls for rigorous matching, which may be impossible. While frequency counts represent the most rigorous method for obtaining probability measures they may not be available or even obtainable. Alternatively, probabilities can be assigned to events, based on the previous experience of "experts". These are by no means arbitrary and tend to capture knowledge in the form of human experience. These are termed subjective or personal probabilities.

Possibility theory is applied when the uncertainty is due to vagueness—for example, in the terms used to describe a histological feature. "Fuzzy logic"[9] was developed in the 1960s to assess situations in which mathematical probability was seen as inappropriate but in which possibility theory could provide consistent methodology. There are other points to consider—for example, probabilities must sum to unity over all possible outcomes; possibilities have no such constraint. The rules of combining evidence in a probabilistic approach and in possibility based system are different.

When uncertainty is assessed by a probabilistic measure, the combining of evidence may be based on Bayes theorem (Bayes, Reverend Thomas, 1702–61), which allows one to compute the probability of an event— for example, a diagnosis, if some other event has taken place—for example, the observation of a histological feature. How the theorem works is clearly explained by Bartels *et al*,[4] using the probability of a high-risk woman developing severe cervical intraepithelial neoplasia (CIN III) as an example, in which the known prior probability for all women to develop CIN III is combined with the conditional probability (probability of CIN III given the presence of a certain histological feature) to obtain a new probability that the woman falls into the CIN III group.

The assumption has to be made that the conditional probabilities are independent of each other, because otherwise the computational load becomes enormous—that is,

2^n for n clues (4 billion conditional probabilities for 32 clues!). But in histopathological diagnosis and prognosis the clues are almost never independent. To deal with this one can use the knowledge of experts to draw up a network that uses only the most effective and independent clues and then apply Bayes theorem to it: hence the Bayesian Belief Network, which has been extensively researched by Pearl.[10]

Construction of a BBN begins by defining the diagnostic problem. This might be the simple distinction between benign and malignant for a particular tissue, or the grading of a lesion (for example, grade 1, 2, 3) or a differential diagnosis (for example, atypical ductal hyperplasia/ductal carcinoma in situ). The possible diagnostic outcomes for a given problem are held in a **decision node**. Next the pathologist defines each of the microscopic and/or clinical features which are important in making the diagnostic decision in question. Each feature is held in a separate **evidence node** and includes the name of the feature and its possible outcomes (for example, tubular formation: well defined/moderately well defined/poorly defined). The relation between each feature (evidence node) and the diagnostic decision (decision node) is then defined. A clear advantage of this approach is that one can retain the terminology that pathologists use in the description of histological or cytological clues. This is potentially valuable because, if it lives up to expectations, its use will not only facilitate more uniform reporting of pathological material but provide insight into how diagnostic decisions are made.

An example

In the example below we show (A) how to design and create a BBN and (B) how to use it to reach a diagnostic decision. The calculations needed to do this are given. They are relatively simple and can be worked out using a pocket calculator, which the reader may also like to use to follow the process. The authors have found it necessary to understand the calculations in order to appreciate how BBNs work. In practice, however, software is available which facilitates the creation of a network and which carries out all the diagnostic calculations very rapidly.

In our example we develop and use a BBN for the grading of breast cancer according to the Bloom and Richardson scheme.[1]

(A) Creating a BBN

There are four steps in creating a BBN.

STEP A1. DEFINE THE DECISION NODE (FIG 1)
In our hypothetical case of breast cancer we want to decide the Bloom and Richardson grade. We therefore create the decision node which contains the grade, for which we have three possible outcomes, grades 1, 2 or 3. We now need to allocate a prior probability to each

of the three possible outcomes; we assume that the pathologist has no preconception as to whether a case will turn out to be grade 1, 2 or 3, so we give an equal Prior Probability to each of the three grades (table 1). Also stored at the decision node are two additional values for each diagnostic outcome. These are called the Relative Likelihood Ratio (abbreviated to Lambda) and the Belief. The belief is calculated by multiplying the prior probability by the lambda. Initially, the lambda is set to 1·0. At the start, therefore, the belief equals the prior probability. While these values are not important at this stage, they are essential and will be explained later when we use the network to make a decision.

STEP A2. DEFINE THE EVIDENCE NODES (FIG 1)
We must next define the features of importance in making the diagnostic decision. In our example three histological features are used as evidence in the grading of a breast cancer. These are listed with their possible outcomes as defined by us:

● tubule formation; possible outcomes: well defined, moderately well defined, poorly defined;
● mitotic figures; possible outcomes: rare mitoses, 2–3 mitoses/high power field (hpf), several mitoses/hpf;
● pleomorphism; possible outcomes: uniform nuclei, moderate pleomorphism, severe pleomorphism.

These are also illustrated in fig 1 as evidence feeding into the decision node. Each evidence node also requires prior probability, lambda and belief values. These are calculated in the initialisation step.

STEP A3. DEFINING THE RELATION BETWEEN THE EVIDENCE AND THE DECISION
The features described by the evidence nodes in Step A2 have a bearing on the diagnostic decision. This relation is defined by estimating the probability of finding the evidence node outcomes (for example, tubule formation: well defined) given the decision node outcomes (for example, grade 1). These are naturally called conditional probability matrices (CPM) and are established by the author of the BBN from the best information available, whether it is actual frequency data, published data or personal experience. In our example we judge from our own experience that the probability is that in a case of grade 1 carcinoma of the breast we will find tubular formation to be well defined in 86% of cases (a probability of 0·86), moderately well defined in 9% of cases (a probability of 0.09) and poorly defined in 5% of cases (a probability of 0·05) (fig 2). This accounts for 100% of cases. Our judgements of what we would expect tubular formation to show in grades 2 and 3, and the matrices established from what we would expect to find in each grade with regard to mitotic figures and nuclear pleomorphism are shown in table 2.
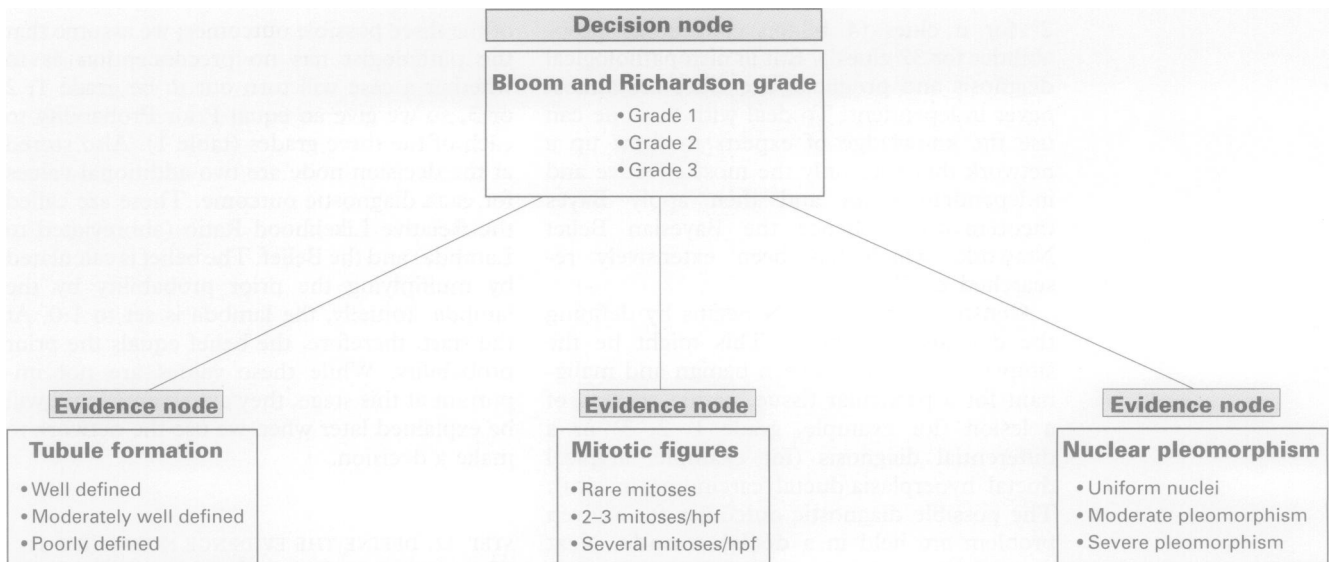
Figure 1   Shallow network with an open-tree hierarchic topology with a diagnostic decision node and three first-level descendant nodes.

Table 1   *The decision node showing the figures entered at the initialisation stages*

Grade	Prior probability	Lambda	Belief
1	0·33	1·0	0·33
2	0·33	1·0	0·33
3	0·34	1·0	0·34

Prior probability × lambda = belief.

STEP A4. INITIALISATION OF THE NETWORK

As stated previously, prior probabilities, lambdas and beliefs need to be calculated for the evidence nodes before the network can be used. Table 3 shows that this is done by combining the numbers in the decision node (table 1) with the figures in the CPM (table 2). For the tubules the prior probability of a case being grade 1 (0·33), grade 2 (0·33) or grade 3 (0·34) is multiplied by the conditional probability of well defined, moderately well defined, and poorly defined tubular formation being seen. These products are summed as shown in table 3A. The totals are entered as

the prior probability in the tubular formation evidence node (table 3B). The same procedure is completed for the mitotic figures and nuclear pleomorphism evidence nodes (fig 3) (table 3B). As in Step A1, the lambda values for the evidence node outcomes are initially set to 1·0 and multiplication of the prior probability with the lambda for each outcome gives the belief values. In the BBN creation stage the lambda and belief columns in table 3B are merely awaiting evidence.

(B) Using the BBN to make a decision

In order to use the BBN to reach a new belief in the diagnostic outcomes we have to enter evidence about a specific case into the network. The process of how this evidence modifies the diagnostic beliefs takes place in four steps.

STEP B1. ENTERING THE EVIDENCE INTO THE EVIDENCE NODES

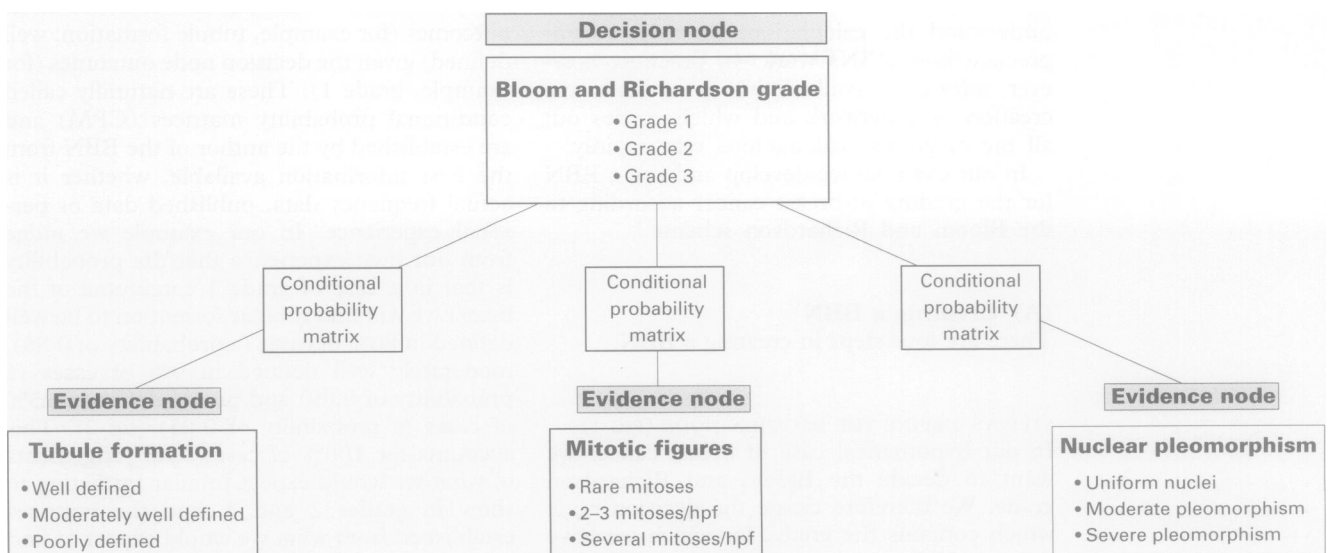The BBN operates interactively with the pathologist providing information on each piece of

Figure 2   Figure 1 showing conditional probability matrices in position.

Table 2 *Conditional probability matrices for the links between the decision node and the tubule formation, mitotic figures and nuclear pleomorphism evidence nodes. This numerical relation between evidence and decision nodes is defined by the "expert"*

| | Decision node | | |
| | Grade | | |
Evidence node	1	2	3
Tubule formation			
well defined	0·86	0·10	0·01
moderately well defined	0·09	0·70	0·09
poorly defined	0·05	0·20	0·90
Mitotic figures			
rare	0·80	0·10	0·01
2–3/hpf	0·15	0·45	0·29
several/hpf	0·05	0·45	0·70
Nuclear pleomorphism			
uniform	0·75	0·15	0·01
moderate	0·20	0·70	0·19
severe	0·05	0·15	0·80

Table 3A *Calculation of the prior probability numbers for the tubule formation evidence node*

Decision node		×	CPM (for well defined)
Grade 1	0·33		0·86 = 0·2838
Grade 2	0·33		0·10 = 0·0330
Grade 3	0·34		0·01 = 0·0034
			total = 0·3202

Decision node		×	CPM (for moderately well defined)
Grade 1	0·33		0·09 = 0·0297
Grade 2	0·33		0·70 = 0·2310
Grade 3	0·34		0·09 = 0·0306
			total = 0·2913

Decision node		×	CPM (for poorly defined)
Grade 1	0·33		0·05 = 0·0165
Grade 2	0·33		0·20 = 0·0660
Grade 3	0·34		0·90 = 0·3060
			total = 0·3885

Table 3B *Evidence node values for tubule formation after calculating prior probability*

	PP	Lambda	Belief
Well defined	0·32	1·0	0·32
Moderately well	0·29	1·0	0·29
Poorly defined	0·39	1·0	0·39

PP = prior probability.

evidence. This must be presented to the network in the formation of a relative likelihood ratio (lambda) vector for each feature. This vector is nothing more than a set of numbers which describes the relative likelihood that the outcomes for a specific feature are present in the current case that is being examined. These numbers can be estimated subjectively. For example, we may want to be discrete about our assessment of features and give the following relative likelihood ratio for tubule formation:

Tubule formation		
well defined	moderately well defined	poorly defined
10	1	1

This lambda vector of [10 1 1] indicates that we are certain that the current case shows well defined tubules (see Abmayr et al[5]). We can, however, allow ourselves more freedom in our allocation of a relative likelihood ratio for histological features. If we wish, we can set the scale between 0·0 and 1·0 and define the following likelihood ratio:

Tubule formation		
well defined	moderately well defined	poorly defined
0·95	0·20	0·01

This lambda vector of [0·95 0·20 0·01] expresses a strong likelihood that the case shows well defined tubules, but that it may also show moderately defined tubules (see Hamilton et al[6]). This allows us to express uncertainty, which is common in the assessment of histological features. Note that in both examples the values do not have to sum to unity, because ultimately the belief vector values are recalculated to sum to unity ("normalised").

Another promising way in which the relative likelihood ratio can be arrived at, without the pathologist having to convert what he sees into
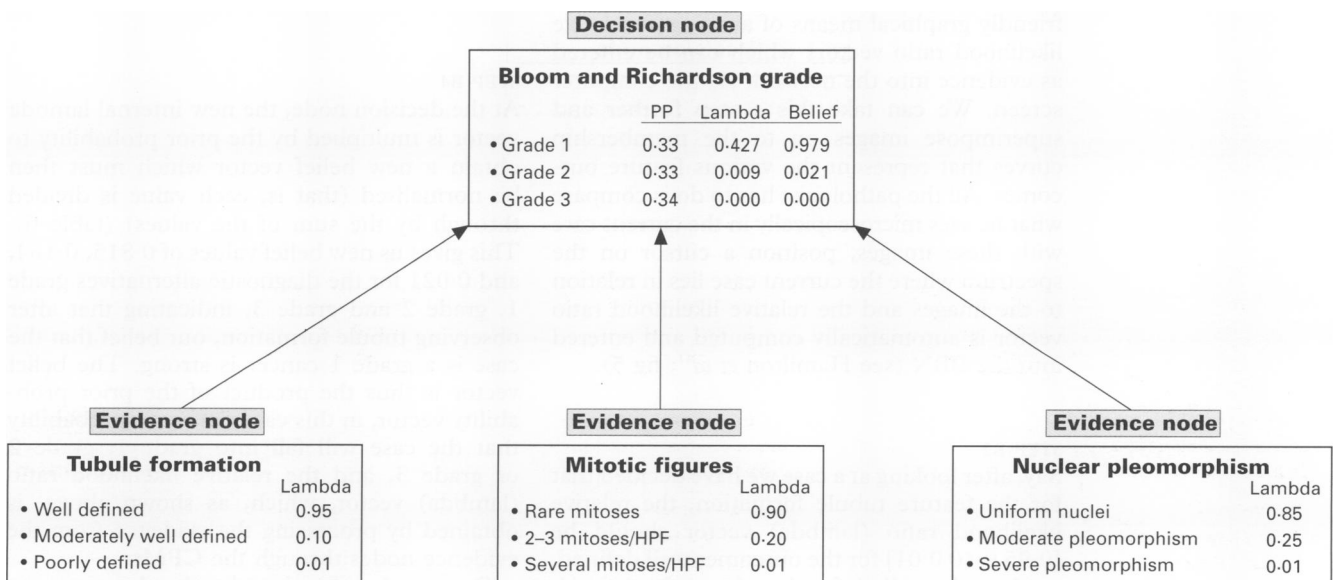

Figure 3 *Data resulting from initialising of the network and entry of data into the three evidence nodes.*

Poorly defined   Moderately defined   Well defined
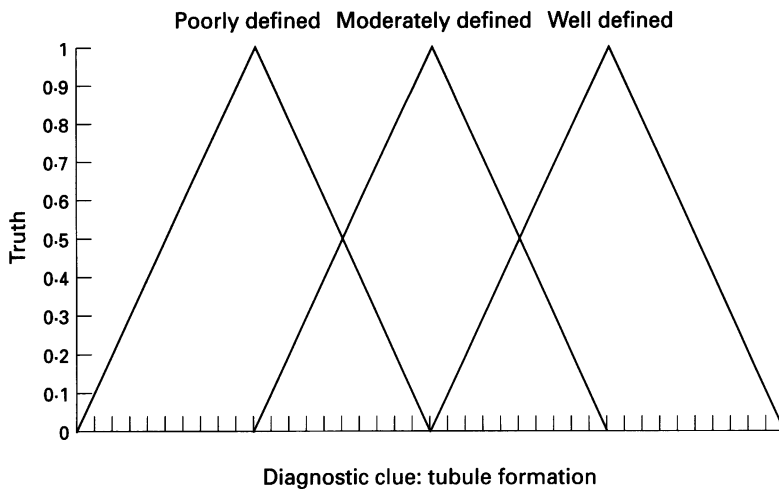

Diagnostic clue: tubule formation

Figure 4   Graphical representation of overlapping distribution of membership functions for the feature, tubule formation. The dense vertical line represents the position where a new case is thought to lie in relation to the feature outcomes. The likelihood vector for this position can be obtained by reading the values on the y-axis where this line cuts the membership functions. Here we have a likelihood vector of [0 0·65 0·35] for poorly, moderately and well defined tubule formation.
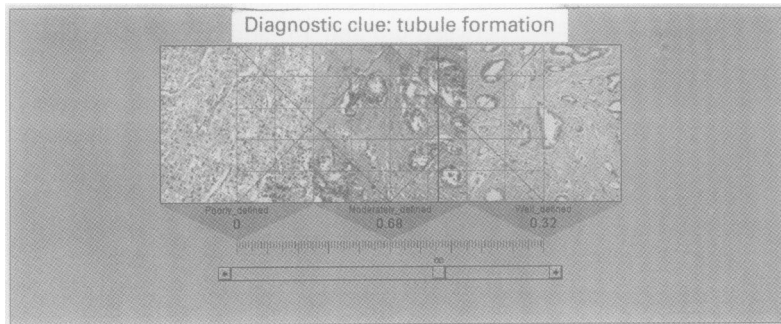

Figure 5   Images of the different feature outcomes for tubule formation showing the membership functions.

Table 4   The hypothetical evidence from the hypothetical example case for each feature is entered as a relative likelihood ratio (lambda) vector for each evidence node

Evidence node	Lambda
Tubule formation	
well defined	0·95
moderately	0·10
poorly defined	0·01
Mitoses	
rare	0·90
2–3/hpf	0·20
several/hpf	0·01
Pleomorphism	
uniform	0·85
moderate	0·25
severe	0·01

for mitoses and pleomorphism are shown in table 4.

We begin our calculations by entering the relative likelihood ratio vector for tubule formation into the BBN. The relative likelihood vector is first multiplied by the initial lambda values at the tubule formation node to give a new lambda vector. As the initial lambda vector values were originally set to 1·0 (see Step A4), multiplication gives a new lambda vector at that evidence node which simply equals the relative likelihood ratio vector.

We now calculate the product of the new lambda vector and the CPM which provides an "external" lambda vector for entry into the decision node. This involves simple matrix algebra and the calculations are shown in table 5. Here, the external lambda vector for the decision node is calculated to be [0·827 0·167 0·021].

numbers, is to draw a graph and represent the outcomes of each feature as an overlapping distribution or membership curve (fig 4). By marking the position on the x-axis where we feel the case in question falls, the relative likelihood ratio vector for that feature is expressed by the ordinate values of where that position cuts the membership curves. This is essentially a friendly graphical means of arriving at relative likelihood ratio vectors which can be entered as evidence into the network via the computer screen. We can take this a step further and superimpose images on to the membership curves that represent the various feature outcomes. All the pathologist has to do is compare what he sees microscopically in the current case with these images, position a cursor on the spectrum where the current case lies in relation to the images and the relative likelihood ratio vector is automatically computed and entered into the BBN (see Hamilton *et al*[11]; fig 5).

STEP B2

Say, after looking at a case we have decided that for the feature tubule formation, the relative likelihood ratio (lambda) vector should be [0·95 0·10 0·01] for the outcomes well defined, moderately well defined and poorly defined. This and the other relative likelihood vectors

STEP B3

The "external" lambda vector for the decision node is multiplied by the "internal" lambda vector at the decision node to give a revised internal lambda vector. Again as the initial internal lambda values at the decision node were originally set to 1·0 (see Step A1), the new internal lambda equals the external lambda (that is, [0·827 0·167 0·021]).

STEP B4

At the decision node, the new internal lambda vector is multiplied by the prior probability to obtain a new belief vector which must then be normalised (that is, each value is divided through by the sum of the values) (table 6). This gives us new belief values of 0·815, 0·164, and 0·021 for the diagnostic alternatives grade 1, grade 2 and grade 3, indicating that after observing tubule formation, our belief that the case is a grade 1 cancer is strong. The belief vector is thus the product of the prior probability vector, in this case the equal probability that the case will fall into grade 1, grade 2 or grade 3, and the relative likelihood ratio (lambda) vector, which, as shown above, is obtained by processing the evidence from the evidence nodes through the CPMs.

One can therefore imagine the above process as a message which begins by entering evidence

Table 5   Calculation of evidence node external lambda vector message for tubule formation to the decision node

Tubule formation	Lambda	×	CPM (for grade 1)	
well defined	0·95		0·86	=0·817
moderately well	0·10		0·09	=0·009
poorly defined	0·01		0·05	=0·0005
				total=0·8265
	Lambda	×	CPM (for grade 2)	
well defined	0·95		0·10	=0·095
moderately well	0·10		0·70	=0·07
poorly defined	0·01		0·20	=0·002
				total=0·167
	Lambda	×	CPM (for grade 3)	
well defined	0·95		0·01	=0·0095
moderately well	0·10		0·02	=0·002
poorly defined	0·01		0·90	=0·009
				total=0·0205

Lambda message for decision node from tubule formation = [0·827 0·167 0·021].

into the evidence node, the internal lambda is calculated, this message is passed via the CPM to form an external lambda for the decision node. This updates the internal lambda of the decision node which in turn updates the beliefs in the diagnostic alternatives. The message is carried by the lambda values within the BBN and so this propagation of change from evidence node to decision node is often called the lambda message.

What next happens is that another message (called the pi message) is passed down to the remaining evidence nodes. This updates their prior probabilities. This process has no effect on the decision in a single level BBN such as our current example so the calculations will not be described here. They have been given previously by Bartels et al.[4]

Steps B1 to B4 are repeated for the mitoses evidence. These calculations are summarised in box 1 but are carried out in an identical manner to the entry of evidence for tubule formation. The only important thing to remember is that the internal lambda values at the diagnostic node have now been changed by the entry of tubule formation and updating the evidence for the next evidence nodes must use these new values. After the entry of both tubule formation and mitoses, the beliefs in the diagnostic alternatives are as follows: grade 1, 0·949; grade 2, 0·046; and grade 3, 0·005.

Steps B1 to B4 are repeated for the pleomorphism evidence. These calculations are summarised in box 2 and are carried out as before. After the entry of all three pieces of evidence, the belief in the diagnostic alternatives are as follows: grade 1, 0·979; grade 2, 0·021; and grade 3, 0·000.

Discussion

The clear cut example shown above was designed to illustrate the calculations involved. The benefit really comes when we are less

certain about how defined the tubules are, how many mitoses there are or how uniform the nuclei really are across the lesion, as we can express the outcomes for each feature as a number ranging from 0·00 to 1·00 for each. This approach allows the combination of such "uncertain" evidence to provide an overall belief for the diagnostic outcomes. We maintain the use of descriptive terminology in our assessment of cases but obtain numerical data to support our decision. As can be seen, the collection of evidence is sequential and the belief in the diagnostic alternatives is based on the accumulated evidence to date—that is, a final value of belief will be given even if certain evidence is missing.

Working with the pocket calculator becomes laborious even with three features and only one level in the hierarchy of evidence nodes as above. Other diagnostic scenarios are much more complex. For example, in a BBN designed by Montironi et al,[13] a total of eight diagnostic clues was used to make the diagnostic distinction between benign, low grade prostate intraepithelial neoplasia (PIN), high grade PIN, cribriform carcinoma, and large acinar pattern carcinoma. Another BBN created to support the diagnosis of "benign" or "malignant" in fine needle aspirates of breast lesions[6] consisted of 10 features with 27 possibilities (outcomes). Here, computer programs designed to accept evidence and process these through a BBN to arrive at a final decision node belief are vital.[5] A suitable software program is now commercially available (DIAGSOFT, Kleselstrasse 14a, D-8099, Munich, Germany) for the development and investigation of BBNs in pathology. This requires only a cheap computer (386 processor, 4 Mb RAM, Windows 3·0 or higher). Other software packages are being developed with specific applications in mind.[11]

BBNs have been created and applied to microscopic (histological or cytological) diagnosis and/or prognostic grading of melanocytic lesions,[5] breast lesions[6 11] and prostatic lesions.[7 12 13] Others—for example, lung cancer biopsy typing, are in progress. Composite networks, in which several networks are combined and the most suitable one pursued for the diagnosis that emerges, are in prospect.

It is proving easy for the "experts" to define the diagnostic or prognostic outcomes for the

Table 6   Recalculation of the decision node belief after entering external lambda for tubule formation from table 5

Grade	PP		Lambda		Belief	Normalised belief
1	0·33	×	0·827	=	0·273	0·815
2	0·33	×	0·167	=	0·055	0·164
3	0·34	×	0·021	=	0·007	0·021
						1·000

PP = prior probability.

Box 1   Calculations for evidence of mitoses

(i) At the evidence node
Relative likelihood
vector (evidence)   Internal lambda   New internal lambda
$$[0.90\ 0.20\ 0.01] \times [1.0\ 1.0\ 1.0] = [0.90\ 0.20\ 0.01]$$

Internal lambda   CPM   External lambda
$$[0.90\ 0.20\ 0.01] \times \begin{bmatrix} 0.80 & 0.15 & 0.05 \\ 0.10 & 0.45 & 0.45 \\ 0.01 & 0.29 & 0.70 \end{bmatrix} = [0.750\ 0.185\ 0.074]$$

(ii) At the decision node
External lambda   Internal lambda   New internal lambda
$$[0.750\ 0.185\ 0.074] \times [0.827\ 0.167\ 0.021] = [0.620\ 0.031\ 0.002]$$

New internal lambda   Prior probability   Belief
$$[0.620\ 0.031\ 0.002] \times [0.33\ 0.33\ 0.34] = [0.205\ 0.010\ 0.001]$$

Normalised belief
$$[0.949\ 0.046\ 0.005]$$

Box 2   Calculations for evidence of pleomorphism

(i) At the evidence node
Relative likelihood
vector (evidence)   Internal lambda   New internal lambda
$$[0.85\ 0.25\ 0.01] \times [1.0\ 1.0\ 1.0] = [0.85\ 0.25\ 0.01]$$

New internal lambda   CPM   External lambda
$$[0.85\ 0.25\ 0.0] \times \begin{bmatrix} 0.75 & 0.20 & 0.05 \\ 0.15 & 0.70 & 0.15 \\ 0.01 & 0.19 & 0.80 \end{bmatrix} = [0.688\ 0.304\ 0.064]$$

(ii) At the decision node
External lambda   Internal lambda   New internal lambda
$$[0.688\ 0.304\ 0.064] \times [0.620\ 0.031\ 0.002] = [0.427\ 0.009\ 0.000]$$

New internal lambda   Prior probability   Belief
$$[0.427\ 0.009\ 0.000] \times [0.33\ 0.33\ 0.34] = [0.141\ 0.003\ 0.000]$$

Normalised belief
$$[0.979\ 0.021\ 0.000]$$

---

parent node and not too difficult to select the microscopic clues and outcomes for the evidence nodes; easy to obtain the prior probabilities for the parent node but more difficult to "estimate" the numbers for the CPM. It is proving more time-consuming for even an experienced pathologist to read the slides with attention to the microscopical clues and outcomes (nine in the example above; 27 in the study by Hamilton *et al*[6]) and to record them and read them into the computer program than for him to use his normal mental processing and descriptive procedure. Different observers do not necessarily come up with the same answers even with BBNs[14] and the reasons for that must be explored. It is likely that this is related to the translation of observations into numbers, a task with which most pathologists are unfamiliar. However, as mentioned above, Hamilton's group and others are already working on the use of stored computer images as guides to visual assessment and as a means of evidence entry into the BBN, thus combining the recording and entering step.[6 14] It is also possible to examine the sequence of data entry to see how the diagnostic outcomes were affected by the features chosen and to compare the sequences of different observers.[6 14] We expect the latter, together with the need for specific feature observation and the creation of sets of

reference images, to become very valuable in the training of microscopists, in audit, continuing education, and in routine diagnostic pathology.

1 Bloom HJG, Richardson WW. Histological grading and prognosis in breast cancer. *Br J Cancer* 1957;11:359–77.
2 Bartels PH, Hiessl H. Expert systems in histopathology. II. Knowledge representation and rule based systems. *Anal Quant Cytol Histol* 1989;11:147–53.
3 Cross SS, Stephenson TJ, Diez Y, Harrison RF, Underwood JCE, Downs J. Diagnosis of breast fine needle aspirates using human observations and a multilayer perception neural network (abstract). *J Pathol* 1995;175(Suppl): 109A.
4 Bartels PH, Thompson D, Weber J. Expert systems in histopathology. IV. The management of uncertainty. *Anal Quant Cytol Histol* 1992;14:1–13.
5 Abmayr W, Stolz W, Rehberg D, Thompson D, Bartels PH. Bayesian belief network for Windows. A tool for computer aided diagnosis of macroscopic melanocytic lesions investigated by dermatoscopy. In: Wied GL, Bartels PH, Rosenthal DL, Schenk U, eds. *Compendium on the computerised cytology and histology laboratory.* Chicago: Tutorials of Cytology, 1994:96–106.
6 Hamilton PW, Anderson N, Bartels PH, Thompson D. Expert system support in the diagnosis of breast fine-needle aspiration biopsy using Bayesian Belief Networks. *J Clin Pathol* 1994;47:329–36.
7 Montironi R, Bartels PH, Thompson D, Scarpelli M, Hamilton PW. Prostate intraepithelial neoplasia: development of Bayesian Belief Network for diagnosis and grading. *Anal Quant Cytol Histol* 1994;16:101–12.
8 Shafer G, Pearl J (eds). *Readings in uncertain reasoning.* San Mateo, California: Morgan Kaufman Publishing, 1990.

9 Zadeh LA. Fuzzy sets. *Information Control* 1965;8:338–53.
10 Pearl J. *Probabilistic reasoning in intelligent systems*. San Mateo, California: Morgan Kaufman Publishing, 1988.
11 Hamilton PW, Anderson N, Bartels PH, Diamond J, Thompson D, Gregg B, *et al*. An interactive decision support system for breast fine needle aspiration cytology. *Anal Quant Cytol Histol* (in press).
12 Hamilton PW, Montironi R, Abmayr W, Bibbo M, Anderson N, Thompson D, *et al*. Clinical applications of Bayesian belief networks in pathology. *Pathologica* 1995;87:237–45.
13 Montironi R, Bartels PH, Thompson D, Scarpelli M, Hamilton PW. Prostatic intraepithelial neoplasia (PIN). Performance of a Bayesian Belief Network for diagnosis and grading. *J Pathol* 1995;177:153–62.
14 Whimster WF, Anderson N, Humphreys S, Boyle M, Sundaresan M, Rainey A, *et al*. Reproducibility of Bayesian belief network (BBN) assessment of breast FNAs. *Anal Quant Cytol Histol* (in press).