

Combining Argumentation and Bayesian Nets for Breast Cancer Prognosis

MATT WILLIAMS¹ and JON WILLIAMSON²

¹*Advanced Computation Laboratory, Cancer Research UK*

E-mail: matthew.williams@cancer.org.uk

²*Department of Philosophy, Logic and Scientific Method, London School of Economics*

E-mail: j.williamson@lse.ac.uk

Abstract. We present a new framework for combining logic with probability, and demonstrate the application of this framework to breast cancer prognosis. Background knowledge concerning breast cancer prognosis is represented using logical arguments. This background knowledge and a database are used to build a Bayesian net that captures the probabilistic relationships amongst the variables. Causal hypotheses gleaned from the Bayesian net in turn generate new arguments. The Bayesian net can be queried to help decide when one argument attacks another. The Bayesian net is used to perform the prognosis, while the argumentation framework is used to provide a qualitative explanation of the prognosis.

Key words: argumentation, logic, bayes theorem, bayesian networks

1. Introduction

There are two key ways in which one may offer a prognosis for a patient suffering from breast cancer. One can find arguments in favour of recurrence and arguments against recurrence and try to balance the former against the latter to reach a conclusion. Or one can construct a statistical model to estimate the probability of recurrence given the patient's symptoms. Under the former approach the reasons for the decision are clear, but the weighing up of arguments is difficult. The latter approach offers reliable conclusions but no obvious qualitative chain of inference to the conclusion.

Through a judicious combination of the two approaches one can reap the rewards of each while avoiding their disadvantages. A statistical model can be used for the prognosis itself, while arguments can be used to present the reasons for and against such a prognosis. The task of this paper is to present such a combination of the two approaches and to apply it to breast cancer prognosis.

Logic and probability can be combined in various ways. One strategy involves creating a new formalism that encapsulates both logic and probability and new techniques for operating within this formalism, i.e. for constructing and updating

theories or models, and inferring conclusions from them. A second strategy involves adopting standard logical and probabilistic formalisms, keeping the logical and probabilistic components separate, but letting them interact. We take this latter, modular approach: we adopt a standard argumentation formalism, a standard probabilistic network formalism, and the two interact (the argumentation component helps generate the probabilistic component which in turn is used to augment the argumentation component). This modular methodology has several advantages. First, it tends to be simpler: logical and probabilistic formalisms, while fairly simple on their own, soon become very complicated when amalgamated. Second one can directly employ well-understood techniques for constructing logical theories and probabilistic models and drawing inferences from them, rather than having to develop new methods. Third, one can more easily replace one or other component (or both) with one's preferred formalism, or tailor components to particular applications.

The plan is first to introduce the breast cancer problem domain in Section 2, as well as current methods for prognosis, argumentation and combining probabilistic methods with argumentation. Then, in Sections 3 and 4 respectively, we shall describe the logical and probabilistic formalisms in more detail. In Section 5 the two are combined in the context of the breast cancer problem. Finally in Section 6 we discuss extensions of the resulting methodology.

2. Background

2.1. BREAST CANCER

Breast Cancer is one of the commonest cancers in the Western World. It is the commonest non-skin cancer in women in the UK and US, and accounts for approximately 1/3 of cancers in women, with lifetime rates of 1 in 10. Some 36 000 cases are diagnosed each year in the UK, of whom about a 1/3rd will die from the disease (McPherson et al., 2000). Consequently there has been a considerable amount of research focused on breast cancer, and death rates have fallen over the last 10 years (Quinn and Allen, 1995).

The mainstay of treatment for breast cancer remains surgery and radiotherapy (Richards et al., 1994) with hormonal and chemotherapeutic agents often used to treat presumed micro-metastatic disease. One of the advantages of surgery is that, as well as removing any local disease, a sample can also be taken of the axillary lymph nodes. These are a common site of metastatic spread for the cancer, and their removal not only removes any spread that may have occurred, but also allows analysis of the nodes to describe the degree of spread. The two main aims of treatment are to provide local control of, and to prevent premature death from, disease.

Examination of the primary tumour and lymph nodes lets us define certain characteristics of the disease that make local recurrence and death more likely. These characteristics are primarily the grade of the tumour, (which represents the

degree of abnormality displayed by the cells, scored 1–3), the size of the tumour (as its maximum diameter, in mm) and the number of involved nodes (Richards et al., 1994). There are also newer tests for the presence or absence of certain proteins on the cell surface that may predict tumour behaviour or response to certain drugs (Veer et al., 2005; Cristofanilli et al., 2005).

The central aim of therapy planning is to match treatment with the risk of further disease. Thus those at high risk should be treated aggressively while those at low risk should be treated less aggressively. This allows more efficient use of resources, and restricts the (often considerable) side effects of intensive treatment to those patients who would benefit most.

2.2. CURRENT PROGNOSTIC TECHNIQUES

These prognostic characteristics are currently modelled using statistical techniques to provide an estimate of the probability of survival and local recurrence. There are two commonly used prognostic systems. The Nottingham Prognostic Index (NPI) (Galea et al., 1992), which uses data from large UK studies, and the SEER database, which contains results derived from the American Surveillance, Epidemiology and End Results survey (Ries et al., 2004). Both techniques rely on multivariate analyses of large volumes of data (based on over 3 million people for SEER) to calculate prognostic formulae.

These tools, and others like them, are effective at providing estimates of risk of death and local recurrence. In contrast, humans are often poor at manipulating explicit probabilities (Kahneman and Tversky, 1973; Borak and Veilleux, 1982), but there are two issues with such approaches. Firstly, the lack of human-readable explanations of the risk result in a situation where we may ‘know’ what the risk *is*, but not *why* it is so. Secondly, clinicians have the ability to consider other data (such as the presence and impact of other co-existing conditions) and adapt the estimates in the light of new information (such as the discovery of Her-2neu, a cell-surface protein that is a marker for more aggressive disease). This allows them to (partially) individualise risk assessments, be confident that such extra information has been incorporated into the decision-making process, and be able to *explain* why the risk estimate is what it is. Now, we seek to combine the strengths of these two approaches. Part of this requires that we present the information on prognosis to clinicians in a form that they can incorporate into their own reasoning. We therefore seek to provide explicit probabilistic estimates of risk whilst also allowing the reasoning behind this prognosis to be clear in order to help clinicians not just *know* what the risk is, but to *understand* why that risk is so, and to be *confident* that they know what factors have been considered by the system.

2.3. ARGUMENTATION

Clinicians, and more generally humans, use arguments to make decisions. The argumentative method goes back at least 2500 years, and extends beyond the

Graeco-Roman tradition (Gard, 1961). Arguments have the advantage that they can present not only a conclusion, but also its justification.

The idea of trying to base decision-making on arguments has a long history. The first clear example of an algorithm for doing so was described by Franklin (1887). Over the last 5–10 years, there has been an increasing amount of work in developing formal logics for argumentation (Fox and Parsons, 1997). The first major work was by Dung (1995) who described a very abstract method for dealing with a set of arguments that attack each other, and from this defined formal semantics to allow calculation of stable and grounded extensions. More recent work has drawn together developments in non-monotonic logic and insights from cognitive science to produce a number of different argumentation frameworks (Amgoud et al., 2004; Hunter and Besnard, 2001; Krause et al., 1995). There have also been some early attempts to produce implemented systems (Pollock, 1999; Sutton and Fox, 2003), some of which have been trialled in medical domains.

Although argumentation formalisms differ, they share some common features (Prakken and Sartor, 1996). They define:

- A logical language, \mathcal{L}
- A structure for an argument
- Conflicts between different arguments
- A notion of defeat of arguments
- A record of the status of arguments

Some are more abstract than others - for example, Dung's argumentation framework leaves the choice of logic open, and defines semantics in terms of attack relations. Irrespective of this, an 'argument' in an argumentation framework should correspond to a proof in the underlying logic. In this, they are the same as other non-monotonic logics. However, the ideas of attack, defeat and status mark them out as separate. Furthermore, since arguments that attack other arguments may, themselves, be defeated, we require some semantics to define the final status of the arguments. In this paper we use an abbreviated version of Dung's Argumentation framework, as presented in Dung (1995).

2.4. RELATIONSHIP TO OTHER WORK

Our approach builds on work in both argumentation and Bayesian reasoning. Work on using qualitative constraints in Bayesian networks has been discussed by Wittig and Jameson (2000), and previous work on fusing the two approaches has been discussed by Parsons (2003, 2004) and Poole (2002). Although both use Bayesian networks as a basis for developing arguments, neither use both Bayesian and argumentative techniques at the same time. Furthermore, both authors concentrate on theoretical approaches, whereas we present Bayesian networks and arguments learned from real data. Poole (2002) adopts a decision-theoretic framework and

uses logic programs with probabilities to represent knowledge; in contrast under our approach logic (the argumentation framework) and probability (the Bayesian net) are quite separate, but interact. Saha and Sen (2004) combine Bayesian nets and argumentation in a rather different context, *negotiation*, and Bayesian nets are used to model other agents' belief states rather than for argument generation. Similarly, McConachy et al. (1998) use a Bayesian net to model the user's belief state in order to determine the persuasiveness of an argument.

3. Logical Formalism

3.1. INTRODUCTION

We may informally define an argument as being a "chain of reasoning that leads to a conclusion via some inference mechanism." We accept that conclusions are defeasible, and may be overturned in the light of new evidence via some argumentation mechanism. Many different formalisms have been proposed for both argumentation and inference, and each has their strengths and weaknesses. Here, we shall use a fragment of a relatively basic formalism, argument-based logic programming,¹ Prakken and Sartor (1996), for inference, together with Dung's framework for argumentation. Although these do not have all the features we would wish to find in a framework for large-scale decision making, they have a simple syntax which makes them a good exemplar for our work. We use a fragment of Prakken & Sartor's rule logic, ignoring preference ordering on rules, and using Dung's semantics for argumentation rather than theirs. Below, we present an abbreviated description of our formalism. For a further discussion, see their original papers (Prakken and Sartor, 1996; Dung, 1995).

3.2. RULES

We define a binary connective, \Rightarrow , which forms rules from literals, and two forms of negation: Strong (or classical) negation, $\neg L$, and weak negation, or negation-as-failure $\sim L$. A strong literal is an atomic first-order formula, or such a formula preceded by strong negation. A weak literal is $\sim L$ where L is a strong literal.

We define two different sorts of rules, r : and s :. The first, r :, are defeasible rules. That is, rules with whom's conclusions we may disagree. The second, s :, are strict rules. These capture definitional rules, and as such are non-defeasible. Strict rules follow the same form and logic as defeasible rules; they are merely epistemologically different, and we use the \rightarrow operator in strict rules for the sake of intelligibility. Strict rules are intended to be used to express transformation

¹ We have based our rule logic on that developed by Prakken and Sartor and the argumentation on Dung's semantics. Readers familiar with these papers may wish to omit the next sections, noting that we have not used preference ordering in our rule logic, and we have not discussed grounded semantics for the argumentation.

between things that are, for our purposes, identical, while defeasible rules are for those about which we might wish to argue. To illustrate our definitions, we use the following example:

EXAMPLE 3.1.

- A: Radiotherapy reduces mortality
- B: Chemotherapy reduces mortality
- C: An aggressive tumour increases mortality
- D: Spread of the tumour to the lymph nodes increases mortality
- E: Low mortality is the opposite of high mortality

DEFINITION 3.1. A rule is an expression of the form either r : or s :

$$r \in \mathcal{D} : L_0 \wedge L_1 \wedge L_2 \wedge \cdots \wedge \sim L_k \wedge \overline{L_m} \Rightarrow L_n$$

$$s \in \mathcal{S} : L_0 \wedge L_1 \wedge L_2 \wedge \cdots \wedge \sim L_k \wedge \overline{L_m} \rightarrow L_n$$

where we denote the complement of L as \overline{L} , and r : and s : are the name of the rules and each L_i ($0 \leq i \leq n$) is a strong literal. The terms to the left of the arrow are the antecedent, and those to the right are the consequent. Together, both types of rules form the input to our system. \mathcal{S} is the set of all strict rules, which represents those facts that are ‘beyond dispute’, while \mathcal{D} is the set of defeasible rules.

A rule may have zero antecedents - such a rule is universally applicable. Variables will be denoted by x, y, z .²

EXAMPLE 3.1.

- A: *Radiotherapy*(x) \Rightarrow *Low Mortality*(x)
- B: *Chemotherapy*(x) \Rightarrow *Low Mortality*(x)
- C: *Aggressive tumour*(x) \Rightarrow *High Mortality*(x)
- D: *Tumour Spread*(x) \Rightarrow *High Mortality*(x)
- E: *Low Mortality*(x) $\rightarrow \neg$ *High Mortality*(x)

We now need to assemble our rules into arguments. We may regard an argument as being formed from a set of rules, such that the antecedent of one rule is satisfied by the consequent of the preceding rule:

DEFINITION 3.2. An *argument* is a finite sequence $A = [r_0 \dots r_n]$ of ground instances of rules such that:

² We will use the notation “ x has Low Mortality” and “*Low Mortality*(x)” interchangeably to aid readability.

1. For every $i(0 \leq i \leq n)$ and for every strong literal L_j in the antecedent of r_i there is a $k < i$ such that L_j is the consequent of r_k .
2. No two rules in the argument have the same consequent.

Condition 1 defines arguments as being formed by chaining rules together, while condition 2 prevents circular chains of rules within an argument. One consequence of (1) is that we may “ignore” weak literals, and use a rule for conjoning strong literals and a rule of ‘defeasible modus ponens’ of the form:

$$\frac{r : L_0 \wedge \dots \wedge L_j \wedge \sim L_k \wedge \dots \wedge \sim L_m \Rightarrow L_n}{L_0 \wedge \dots \wedge L_j} \quad L_n$$

An argument A may be said to be based on a theory $(\mathcal{S}, \mathcal{D})$ iff $A \in \mathcal{S} \cup \mathcal{D}$.

DEFINITION 3.3. For any argument, A :

- A is strict iff it does not contain any defeasible rule; it is defeasible otherwise.
- L is a conclusion of A iff L is the consequent of some rule in A .

Note that the above definition of conclusion is unusual in that we define the conclusion of *any* rule in an argument as a conclusion of the argument.

3.2.1. Conflicts between arguments

So far we have described a defeasible form of modus ponens. In order to allow argumentation, we must define how we infer an attack between two arguments (that is, define one argument as a counter-argument of another). There are three commonly defined forms of attack between arguments. Consider the following argument:

$$\frac{r : hasAggressiveTumour(x) \Rightarrow hasIncreasedMortality(x)}{hasAggressiveTumour(x)} \quad hasIncreasedMortality(x)$$

- A *premise* attack is to dispute that the tumour is aggressive; perhaps the pathologist was wrong about its features
- An *undercutting* attack is to dispute the link between the two; perhaps aggressive tumours do not lead to increased mortality
- A *rebutting* attack is to directly attack the conclusion; perhaps we show a reason why x would be expected to have low mortality

For our purposes, we need only the third form, *rebuttal*. It is worth noting that rebuttal-attacks are symmetrical (that is, if A attacks B , B attacks A) whereas the other two are not.

Let us now consider how to formalise our definition of attack. Using rules A and B from Example 3.1:

$$\begin{aligned} \textit{Chemotherapy}(x) &\Rightarrow \textit{Low Mortality}(x) \\ \textit{Aggressive tumour}(x) &\Rightarrow \textit{High Mortality}(x) \end{aligned}$$

We note that at the moment there is no explicit conflict between the two rules – as we have nothing to connect *High Mortality*(x) and *Low Mortality*(x). We therefore need to include rule E from our example in order to demonstrate the conflict:

$$\textit{Low Mortality}(x) \rightarrow \neg \textit{High Mortality}(x)$$

Having done this, we may now develop two arguments, which conflict:

$$\begin{aligned} A_1: \textit{Chemotherapy}(x) &\Rightarrow \textit{Low Mortality}(x), \\ &\textit{Low Mortality}(x) \rightarrow \neg \textit{High Mortality}(x) \\ A_2: \textit{Aggressive tumour}(x) &\Rightarrow \textit{High Mortality}(x) \end{aligned}$$

However, we need to be careful in our definition of conflict. If we merely take the final conclusion of A_1 and A_2 , then we seem to have conflict between the defeasible rule: *Aggressive tumour*(x) \Rightarrow *High Mortality*(x) and the strict rule: *Low Mortality*(x) \rightarrow \neg *High Mortality*(x). This is counter-intuitive, as we would normally understand the conflict as lying between the two defeasible rules: *Chemotherapy*(x) \Rightarrow *Low Mortality*(x) and *Aggressive tumour*(x) \Rightarrow *High Mortality*(x), rather than with the strict rule whose function is more to provide semantic unification than make new inferences. In order to capture this, we wish to define attack between arguments as being between defeasible rules *within* the arguments, rather than the strict rules that may extend them.

DEFINITION 3.4. We define an attack thus:

- If an argument, A, has a conclusion or assumption L we should consider A being attacked by all those arguments that can with strict rules only be extended to an argument with conclusion \bar{L} .
- Let A be an argument, and T be a sequence of rules; then $A + T$ is the concatenation of A and T.
- Let A_1 and A_2 be two arguments. Then A_1 attacks A_2 iff there are sequences S_1, S_2 of strict rules such that $A_1 + S_1$ is an argument with conclusion L and $A_2 + S_2$ with conclusion \bar{L} or A_2 is an argument with assumption \bar{L} .

3.3. ARGUMENTATION

We are now concerned with resolving arguments; that is, with trying to decide which of a set of arguments are justified. We use Dung's argumentation framework.

DEFINITION 3.5. We define an Argumentation Framework:

- An Argumentation Framework (AF) is a pair $AF = \langle AR, attacks \rangle$ where AR is a set of arguments and $attacks$ is a binary relationship on AR : $attacks \subseteq AR \times AR$.
- For two arguments A, B , $attacks(A,B)$ means that A attacks B .

Using Example 3.1:

$$AF_{Ex} = \langle \{A,B,C,D,E\}, \{(A,C),(A,D),(B,C),(B,D),(C,A),(C,B),(D,A),(D,B)\} \rangle$$

DEFINITION 3.6. We define the following extensions:

- A set S of arguments is said to be *conflict-free* if there are no arguments A, B in S such that A attacks B .
- An argument $A \in AR$ is *acceptable* with respect to S iff for each argument $B \in AR$, if B attacks A , then B is attacked by S .
- A conflict-free set S is *admissible* iff each argument in S is acceptable wrt S .
- A *preferred extension* of AF is a maximal admissible set of AF .
- A conflict-free set, S , is called a *stable* extension iff S attacks each argument which does not belong in S .
- An admissible set S is called a *complete* extension iff each argument which is acceptable wrt S belongs to S . Intuitively, this captures the notion of believing everything that can be defended while still having consistent beliefs.

EXAMPLE 3.2. We use the rules from Example 3.1, as in AF_{Ex} :

$$AF_{Ex} = \langle \{A,B,C,D,E\}, \{(A,C),(A,D),(B,C),(B,D),(C,A),(C,B),(D,A),(D,B)\} \rangle$$

We define two subsets of AR : $S_1 = \{A,B\}$, $S_2 = \{C,D\}$. We note the following properties:

- S_1, S_2 are both *conflict-free*.
- A and B are *acceptable* wrt S_1 ; C and D are *acceptable* wrt S_2 .
- All $x \in S_1, S_2$ are *acceptable* wrt S_1, S_2 ; S_1, S_2 are both *conflict-free*; therefore, S_1, S_2 are both *admissible*.
- S_1, S_2 are both *preferred* extensions of AR , as they are of equal size.
- S_1, S_2 are both *stable* and *complete* extensions.

We have dealt with the arguments by dividing them into two sets - one for increased mortality, and one for decreased. What Dung's semantics suggest is that we cannot rationally hold *both* sets of beliefs at the same time, but must commit to one (or neither). One of the slight problems with our example is that S_1 , S_2 are equally preferred, and so there is no basis for choosing between them. That this problem is, in part, an artifact of our example is shown in Section 5.

4. Probabilistic Formalism

Having described the logical framework that we adopt, we now turn to the probabilistic framework: *causally-interpreted Bayesian nets*. This formalism offers the ability to represent causal relationships amongst the domain variables as well as probabilistic relationships.

4.1. EPISTEMIC CAUSALITY

The *epistemic* view of causality is concerned with determining the causal beliefs that an agent should adopt on the basis of her background knowledge (Williamson, 2005, Chapter 9). In our application, background knowledge consists of a database of past patient observations and a set of qualitative logical rules that concerns the measured variables. We are interested in determining a set of causal relationships suggested by this background knowledge. (It should be emphasised that any causal relationships hypothesised on the basis of the background knowledge are merely tentative beliefs—the agent's best bet as to the causal relationships—and these beliefs would be expected to change as background knowledge improves.)

How then are causal beliefs to be gleaned from background knowledge? The patient database determines a probability distribution over the measured variables. Under this probability distribution there will be a number of probabilistic dependence relationships amongst the variables. Now a probabilistic dependency may be attributable to one of several different explanations: the dependency may be induced by a causal or a semantic or a logical relationship amongst the variables for instance.³ However, in applications such as ours where variables tend to be associated with spatio-temporal events, causal explanations predominate. Hence if background knowledge does not include a non-causal explanation of the dependency, it is reasonable to attribute it to a causal connection. Thus the agent's causal beliefs should account for all the dependencies in the database that are not already accounted for by the agent's other background knowledge.

More formally, we model an agent's causal beliefs by a directed acyclic graph whose nodes are the variables and whose arrows correspond to direct causal

³ Dependencies also arise if the variables in question are related by some mathematical equation or non-causal physical law or if a dependency is forced by boundary conditions or induced by time series. See (Williamson, 2005, § 4.2) for a more detailed discussion of this point.

relationships. This causal graph should contain an arrow from A to B if A and B are dependent when intervening to fix the value of A and when controlling for B 's other direct causes (i.e. if they are probabilistically dependent conditional on B 's other direct causes and conditional on a set of A 's non-effects that includes all its direct causes) and this dependency is not accounted for by other background knowledge. Moreover, this causal belief graph should be minimal: it should not contain any arrows that are not warranted by the agent's background knowledge (Williamson, 2005, Section 9.5).

In certain circumstances one can use readily available algorithms to isolate a causal belief graph that is rational on the basis of background knowledge. If, as in our application, there are no non-causal inducers of dependencies (e.g. no semantic or deductive logical relationships amongst the variables) then the causal belief graph must satisfy the *Causal Markov Condition*: each variable will be probabilistically independent of its non-effects conditional on its direct causes (Williamson, 2005, §9.6). Moreover, the causal belief graph will be a minimal graph out of all those that satisfy the Causal Markov Condition and any causal constraints in the background knowledge. Now minimal graphs satisfying the Causal Markov Condition have been investigated in some detail (Pearl, 2000; Spirtes et al., 1993) and there are several software packages for finding such graphs, many of which allow one to impose further causal constraints (Korb and Nicholson, 2003, Appendix B). These packages can then be used to find those causal belief graphs that are rational on the basis of background knowledge.

4.2. CAUSAL NETS

A *Bayesian net* consists of two components:

DAG a directed acyclic graph whose nodes are the domain variables $V = \{A_1, \dots, A_n\}$,

Probability Tables the probability distribution of each variable conditional on its parents in the graph,

together with an assumption:

Markov Condition each variable is probabilistically independent of its non-descendants in the graph, conditional on its parents.

A Bayesian net represents a joint probability distribution over the domain variables: under the Markov Condition, the probability of an assignment $a_1 \dots a_n$ of values to the variables is a product of conditional probabilities,

$$p(a_1 \dots a_n) = \prod_{i=1}^n p(a_i | \text{par}_i),$$

where par_i is the assignment of values to the parents of A_i . Bayesian nets have been widely researched in the literature and there are a whole host of algorithms for performing probabilistic inference using Bayesian nets.⁴

A *causally-interpreted Bayesian net* or *causal net* is a Bayesian net whose graph is interpreted causally: an arrow from A_i to A_j means that A_i is a direct cause of A_j . Note that under a causal interpretation the Markov Condition is just the Causal Markov Condition.

Since the causal belief graph satisfies the Causal Markov Condition one can construct a causally-interpreted Bayesian net by extracting from the database the probability distribution of each variable conditional on its parents. Thus the DAG in the Bayesian net is the causal belief graph, the probability tables are determined from the database, and the Markov Condition is guaranteed to hold by construction of the causal belief graph.

The resulting Bayesian net can then be used to perform the prognosis: when input a patient's symptoms the net can be queried to determine the probability of recurrence of breast cancer.

The causal net can also be used to generate new qualitative arguments for the argumentation system. The causal graph provides reasons in favour of the prognosis, e.g. if the prognosis is recurrence, the patient has breast cancer in both breasts, and breast cancer in both breasts is a cause of recurrence, then the latter two facts provide an argument in favour of the particular prognosis. Probabilities inferred from the Bayesian net can be used to determine the direction of arguments. If the symptom renders recurrence *more likely* then it is a *positive cause* of recurrence and the causal argument *supports* a prognosis of recurrence. On the other hand, if the symptom renders recurrence *less likely* then it is a *negative cause* or *preventative* of recurrence and the causal argument *attacks* any argument for recurrence. These arguments then provide a qualitative chain of reasoning to back up or qualify the prognosis.

5. Combining Formalisms

5.1. METHOD

We are interested in exploiting both argumentative and Bayesian techniques. Argumentation has the advantage of providing a justification (the 'chain of reasoning') for a conclusion, but cannot easily handle probabilities and is therefore often imprecise in evaluating risk. Bayesian techniques are precise, but are often impenetrable to lay users. A combination of the two therefore offers both accuracy and perspicuity.

In this section we shall show in more detail how a Bayesian network can be used as a basis for generating a prognostic probability for individual cases, and

⁴ See for example the proceedings of the conferences on Uncertainty in Artificial Intelligence, www.auai.org. Software packages are listed in (Korb and Nicholson, 2003, Appendix B).

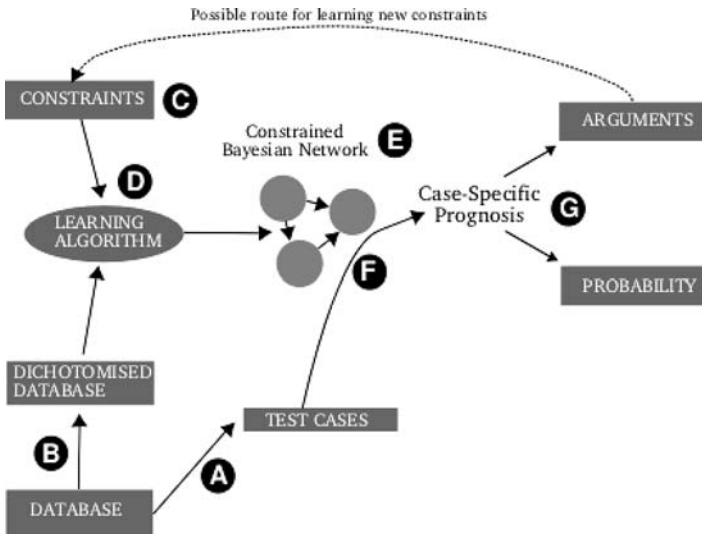


Figure 1. Diagram showing our methodology.

then provide arguments relevant to that case. Our approach is depicted in Figure 1 and consists of steps A-G:

- A: Separate the training and test sets.

We remove the test cases from our training data in order to avoid testing the system on cases it has learned over.

Logical arguments deal with two-valued literals rather than many-valued variables, so we need to generate binary variables from the measured variables:

- B: For each variable, dichotomise the range of values to give two equal populations.

Ideally we would dichotomise so as to divide the variable perfectly; however, in the situation where the data is already grouped to some degree, we cannot divide these groups; we therefore choose to dichotomise so as to minimise the difference in the two halves:

DEFINITION 5.1. If a variable A takes more than two possible values, a^1, \dots, a^k , and these values are ordered $a^1 < \dots < a^k$, then we *dichotomise* the variable to produce a two-valued variable A' such that $A' = a'^1$ iff $A \in \{a^1, \dots, a^j\}$ and $A' = a'^2$ iff $A \in \{a^{j+1}, \dots, a^k\}$, where j is chosen so that the number of data points satisfying $A' = a'^1$ and $A' = a'^2$ are as equal is possible, i.e. $|(N_1 + N_2 + \dots + N_j) - (N_{j+1} +$

$N_{j+2} + \dots + N_k)$ is minimised where N_i is the number of data entries satisfying $A = a^i$.

Next we start on the construction of the causally interpreted Bayesian net. The structure of the causal graph in the net is constrained by domain knowledge available in the logical component of our system.

- C: Define constraints on the network structure based on background knowledge.

These constraints are what may be considered “incontrovertible” knowledge – that which is true irrespective of the data in the data set. Typically, these constraints concern the *absence* of causal relationships between measured variables. Such constraints may be deduced from a variety of sources that may be relatively domain-specific and include temporal and spatial knowledge.

DEFINITION 5.2. We define an operator, \leftrightarrow , such that $A \leftrightarrow B$ means that A definitely does not cause B.

Next we use the database to generate a causal net that satisfies these constraints:

- D: Learn the constrained network.

We have a choice of algorithms for learning Bayesian networks, and we can use any of those which will accept external constraints in their learning process. The resulting network is ‘constrained’ in that we have removed the possibility of certain links.

- E: Isolate the node(s) of interest.

Datasets typically include several variables, of which we are usually only interested in a few. We therefore concentrate on these nodes, and develop probabilities and arguments for particular conclusions, rather than retrieving them for all possible nodes.

- F: Expose the test case(s) to the constrained network.

We use the constrained network that we developed in D to calculate the probability of our variables of interest.

- G: Develop arguments for the test case(s).

DEFINITION 5.3. We develop rules from probabilistic data thus:

- $r: y \Rightarrow x$ represents the fact that for $x,y \in \mathcal{L}$ iff $p(x | y, Z) \geq p(x | \neg y, Z)$ for all $Z \in \{z, \neg z\}$ for which there is a rule $r: z \Rightarrow x$ or $z \Rightarrow \neg x$
- $r: y \Rightarrow \neg x$ represents the fact that for $x,y \in \mathcal{L}$ iff $p(x | y, Z) \leq p(x | \neg y, Z)$ for all $Z \in \{z, \neg z\}$ for which there is a rule $r: z \Rightarrow x$ or $z \Rightarrow \neg x$
- $s: x_1 \rightarrow \neg x_2$ represents the fact that for $x,y \in \mathcal{L}$ iff $p(x_1) = 1 - p(x_2)$

We use the first two definitions to develop defeasible rules about the influence of y on x . The third states that if we have a variable, x , which takes two exclusive values (x_1, x_2) , then we may define an argument for x_1 as being an argument against x_2 .

A Simple Example. We have used the example of the effect of Age on Menopause to illustrate our approach. In this instance, we use only one test case, who is removed from the sample. We then dichotomise the variable Age. We generate two new values for Age, *under 49* and *over 49*, in order to divide the sample as closely as possible into two equal halves. We can then calculate the probabilities for these composite groups. (See Tables I and II, where n is the number of patients satisfying the classification.)

We then constrain the network, using the fact that we know that menopause does not cause aging. Thus we define the constraint:

$$\text{Menopause} \leftarrow \text{Age}$$

and then use the data on age and menopause to learn the constrained network, whose causal graph is shown in Figure 2.

We use the case of a female patient, x , who is 45 years old, and for whom we wish to know the probability of menopause. Using the Bayesian net we are able to

Table I. Probability Distribution for Age and Menopause

Age	20–29	30–39	40–49	50–59	60–69	70–79
Menopausal Status						
p(Pre)	1	0.97	0.91	0.36	0	0
p(Post)	0	0.03	0.09	0.64	1	1
n	1	34	87	90	54	5

Table II. Probability Distribution for dichotomised Age and Menopause

Age	≤ 49	>49
Menopausal Status		
p (Pre)	0.93	0.22
p (Post)	0.07	0.78
n	123	149




Figure 2. Age and Menopause.

estimate the probability of x being pre-menopausal. In this case, $p(\text{Pre-Menopausal}) = 0.93$. We are also interested in presenting pertinent arguments. As x is 45, she clearly falls into the first age group; We know that there are no other factors that affect the relationship between Age and Menopause.⁵ We are then able to develop the rule relating Age and Menopause that is relevant in her case:

$age_1: x \text{ has Age} \leq 49 \Rightarrow x \text{ is Pre-Menopausal}$
 $s : \text{Premenopausal}(x) \rightarrow \neg \text{Post Menopausal}(x).$

5.2. OUR DOMAIN

For our application we chose the problem of predicting recurrence in breast cancer, a well-characterised problem in breast cancer prognosis. We used the Ljubljana Breast Cancer Dataset (Zwitter and Soklic, 1988), a set of 286 instances of real patient data with a binary outcome (Recurrence/ No Recurrence) and 9 possible predictive attributes. This dataset has been used in the past for several machine learning projects (Michalski et al., 1986; Clark and Niblett, 1987). The dataset contains the following variables:

- Age: The age (in years at last birthday) of the patient at the time diagnosis.
- Menopause: Whether the patient is pre- or post- menopausal at time of diagnosis.
- Tumour Size: The greatest diameter (in mm) of the excised tumour.
- Degree of Malignancy: The histological grade (range 1-3) of the tumour. Tumours that are grade 1 predominantly consist of cells that, while neoplastic, retain many of their usual characteristics. Grade 3 tumours predominately consist of cells that are highly abnormal. Such abnormalities include marked variation in cell size and a high index of mitotic activity in the cells.
- Inv_nodes: The number (range 0 - 26) of axillary lymph nodes that contain metastatic breast cancer visible on histological examination. Since the axillary lymph nodes act as a primary site of drainage for the breast, they are a common site of early metastasis.
- Node Caps: If the cancer does metastasise to a lymph node, although outside the original site of the tumour it may remain “contained” by the capsule of the lymph node. However, over time, and with more aggressive disease, the tumour may replace the lymph node and then penetrate the capsule, allowing it to invade the surrounding tissues.

⁵ In this case, because there are no other factors that influence Menopause.

- Breast Quadrant: The breast may be divided into four quadrants, using the nipple as a central point. Breast cancer most commonly occurs in the upper-outer quadrant.
- Breast: Breast cancer may occur in either breast, although there is no difference in incidence between breasts. There are some possible implications of cancer laterality, however, as left-sided breast cancer, if treated by radiotherapy, would involve exposing the heart to a substantial radiation dose, which might increase the subsequent risk of cardiovascular disease.

We used an existing implementation of a machine learning algorithm (Hugin Lite), a feature reduced version of the commercially available package, A/S (1989), which constructs a Bayesian net from data and can also be used to perform probabilistic inference from the net. We excluded five cases from the data to act as test cases, and the rest of the data was used to construct the network (incomplete examples deleted, NPC algorithm, $p = 0.1$, $n = 272$). Variables were dichotomised and the following constraints were used:

Recurrence, Irradiation \leftrightarrow inv_nodes
 Recurrence, Irradiation \leftrightarrow age
 Recurrence, Irradiation \leftrightarrow menopause
 Recurrence, Irradiation \leftrightarrow tumour-size
 Recurrence, Irradiation \leftrightarrow degree malignancy
 Recurrence, Irradiation \leftrightarrow breast quadrant
 Recurrence, Irradiation \leftrightarrow breast
 Recurrence, Irradiation \leftrightarrow node-caps
 Recurrence \leftrightarrow Irradiation

These constraints were provided by existing background domain knowledge, and are based on temporal ordering of events (i.e. Recurrence or Irradiation cannot be a cause of other things in our example, as it occurs later than them).

After constructing the net, we then isolated the component of the graph that involves variables connected to the Recurrence variable. Other variables were ignored – they were unconnected to Recurrence and therefore had no bearing on it.

5.3. RESULTS

We removed five cases from the dataset to act as a test set. Their characteristics are shown below (Table III).

Using the constraints above and the remainder of the data (with parameters as specified above) led to the network shown in Figure 3.

The following trends were seen in $p(\text{Recurrence} | X)$ for different X . Since the options of recurrence/no-recurrence are mutually exclusive, we have excluded the probabilities of no-recurrence from the tables, and the data has been dichotomised in accordance with our method (Tables IV–VI).

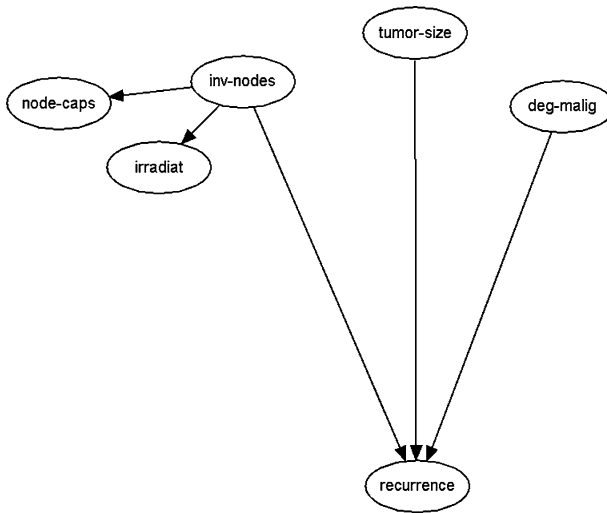


Figure 3. Factors relating to recurrence.

Following our method above, we exposed each of the five cases in turn to the network, and calculated the probability of their recurrence. We also developed a set of rules that were relevant to each case (Table VII). The rules developed were of the form var_n , where var is the variable of interest and $n = 1$ if in the lower end of the range and $n = 2$ if in the upper. These rules may be read as follows:

TS_1 : x has Tumour ≤ 29 mm $\Rightarrow x$ will have Non-Recurrence

TS_2 : x has Tumour > 29 mm $\Rightarrow x$ will have Recurrence

Table III. Details of test cases

Patient No	Age	Menopause	Tumour size	No. Nodes	Hist. Grade
1	60–69	Post	15–19	0–2	2
2	30–39	Pre	25–29	0–2	2
3	30–39	Pre	35–39	0–2	3
4	40–49	Post	20–24	0–2	3
5	50–59	Pre	25–29	3–5	3

Table IV. Tumour size and recurrence

Tumour size (in mm)	≤ 29	> 29
$p(\text{Recurrence})$	0.24	0.36
n	160	112
Rule	TS_1	TS_2

Table V. Number of Involved Nodes and Recurrence

Number of involved nodes	≤ 2	> 2
p (Recurrence)	0.20	0.55
n	205	67
Rule	NIN_1	NIN_2

Table VI. Histological Grade and Recurrence

Histological grade	≤ 2	3
p (Recurrence)	0.19	0.53
n	193	79
Rule	HG_1	HG_2

Table VII. Hugin predicted p(Recurrence) and rules instantiated, by patient

Patient no	p (Recurrence)	Rules instantiated
1	0.17	TS_1, NIN_1, HG_1, S_1
2	0.25	TS_1, NIN_1, HG_1, S_1
3	0.2	TS_2, NIN_1, HG_2, S_1
4	0.22	TS_1, NIN_1, HG_2, S_1
5	0.5	TS_1, NIN_2, HG_2, S_1

NIN_1 : x has Number of Involved Nodes $\leq 2 \Rightarrow x$ will have Non-Recurrence

NIN_2 : x has Number of Involved Nodes $> 2 \Rightarrow x$ will have Recurrence

HG_1 : x has Histological Grade $\leq 2 \Rightarrow x$ will have Non-Recurrence

HG_2 : x has Histological Grade $> 2 \Rightarrow x$ will have Recurrence

S_1 : x will have Recurrence $\rightarrow \neg(x$ will have Non-Recurrence)

These rules were then used to instantiate an individual argumentation framework for each patient (Table VII)

The argumentation frameworks developed for each case are shown below:

$$AF_1 = \langle \{TS_1, NIN_1, HG_1, S_1\}, \{\} \rangle$$

$$AF_2 = \langle \{TS_1, NIN_1, HG_1, S_1\}, \{\} \rangle$$

$$AF_3 = \langle \{TS_2, NIN_1, HG_2, S_1\}, \{(TS_2, NIN_1), (NIN_1, HG_2), (NIN_1, TS_2), (HG_2, NIN_1)\} \rangle$$

$$AF_4 = \langle \{TS_1, NIN_1, HG_2, S_1\}, \{(TS_1, HG_2), (NIN_1, HG_2), (HG_2, TS_1), (HG_2, NIN_1)\} \rangle$$

$$AF_5 = \langle \{TS_1, NIN_2, HG_2, S_1\}, \{(TS_1, NIN_2), (TS_1, HG_2), (NIN_2, TS_1), (HG_2, TS_1)\} \rangle$$

6. Discussion

AF_1 and AF_2 are conflict free; this is because for both, all data symptoms lead to arguments for non-recurrence. Despite this, there is a significant difference in the $p(\text{Recurrence})$; this shows that the probabilistic component of the system provides information that is complementary to that of the logical component. The other three AFs all contain symmetrical attack relations. This comes from our definition of attack as rebuttal, which will give rise to symmetrical AFs. In all three cases, there are arguments both for and against recurrence; in AF_3 and AF_5 there are two for recurrence and one against, while AF_4 has one for recurrence and two against. Despite this, $p(\text{Recurrence})$ is almost identical in cases 3 and 4, while it is substantially different in case 5.

6.1. EXTENSIONS IN $AF_1 - AF_5$

Cases 1 and 2 are clearly both unusual examples, as there is no conflict. Our definitions of semantics for the AF was based on the notion of a set of arguments, S , where $S \subseteq \text{AR}$. For AF_1 and AF_2 , $S = \text{AR}$, and S is clearly *conflict-free*, and the *preferred* extension of $AF_n = S$.

AF_3 has two *conflict-free* subsets: $S_1 = \{NIN_1\}$, $S_2 = \{T S_2, HG_2\}$. Because of the symmetrical nature of our attacks, these two subsets are then both also *admissible*, and both are *preferred extensions* of AR , as both are maximal wrt set inclusion. Both S_1 S_2 are also *stable extensions*, as they are able to attack any argument in AR that attacks them. We can define similar extensions for AF_4 and AF_5 .

6.2. LEARNING ALL THE ARGUMENTS FROM THE NETWORK

We have so far presented an approach that involves constructing a Bayesian network from some data, and then using individual patient characteristics to calculate a probability and arguments that pertain to that patient. However, there is an alternative approach. Instead of using individual patients to instantiate arguments, we could instead use the structure of the network to develop all the possible arguments. Such an approach would have the advantage of ‘capturing’ all of the network, even if we did not yet have examples that triggered all possible nodes in the network.

If we apply this approach to our network, we see that we develop the following rules:

$TS_1: x \text{ has Tumour} \leq 29 \text{ mm} \Rightarrow x \text{ will have Non-Recurrence}$

$TS_2: x \text{ has Tumour} > 29 \text{ mm} \Rightarrow x \text{ will have Recurrence}$

$NIN_1: x \text{ has Number of Involved Nodes} \leq 2 \Rightarrow x \text{ will have Non-Recurrence}$

$NIN_2: x \text{ has Number of Involved Nodes} > 2 \Rightarrow x \text{ will have Recurrence}$

$HG_1: x \text{ has Histological Grade} \leq 2 \Rightarrow x \text{ will have Non-Recurrence}$

$HG_2: x \text{ has Histological Grade} > 2 \Rightarrow x \text{ will have Recurrence}$

The obvious thing to note is that it gives us the same rules as our initial, individualised approach! However, this is an artefact of our method. Consider our results had we used only two test cases – the list of arguments that we developed would be far smaller, and would not cover the entire network. If we were to use a network of 200 nodes, it would be arduous to generate a set of test cases to extract all the arguments.

Under this alternative approach of extracting all rules from the net, one only needs to extract rules once. However, large networks will lead to very large systems of rules, many of which may not be needed to back up individual prognoses.⁶ Moreover, if new data induces modifications to the net, the set of rules will need to be reconstructed. Under the original approach we only extract rules as and when we need them (i.e. for each patient). However, a large number of patients will lead to some replication of rules. Clearly the application will determine which approach is most appropriate; in our current work the choice between the two approaches is small and more a matter of taste than substance.

6.3. IMPLICATIONS FOR ARGUMENTATION

Our work presents a technique for combining causally interpreted Bayesian networks and argumentation; since we have learned the Bayesian networks from (real-life) data, we have also demonstrated how Bayesian networks can be used as an intermediary stage in the production of argumentation frameworks from data. Taken as a whole, our results point to some important lessons on the real-world use of argumentation. Firstly, the number of arguments for and against a conclusion cannot be used in isolation to predict the likelihood of the outcome - some arguments count more than others, and the probabilistic component provides a way to measure this. Secondly, argumentative structures need to be richer than those used here - we need to expand our example, and our argumentation, to capture a richer semantics than simple rebuttal. Doing so will prevent our AFs being necessarily symmetrical, as the ones here are. Thirdly, our notion of argument is strongly dependent on the background on which it is used; the actual probability of recurrence in our sample was approximately 0.3; thus arguments ‘for’ and ‘against’ recurrence need to be taken in the light of this background rate. Failure to do so leaves argumentation open to the same biases as human beings (notably the base rate fallacy). Finally, our rule structure ensured that each case would generate an equal number of rules.⁷

⁶ Large networks are quite common in other applications, e.g. vision systems, natural language processing and bioinformatics (Friedman, 2004).

⁷ This is due to the coverage of the rules: Consider some case, C_1 , in a domain which has variables x , y and z . Since we dichotomise the variables, $(x_1, x_2, y_1, y_2, z_1, z_2)$ to produce our rules, whatever C_1 's particular attributes (e.g. $x = 0.4, y = 7, z = true$), she will instantiate 3 rules; if we now consider some other case, C_2 , with $x = 1.2, y = 3, z = false$, we see that she will also instantiate 3 rules – although they may not be the same rules as C_1 . Therefore each case will have an AF of the same size.

Such a property may seem desirable from a theoretical point of view, but from a practical aspect it seems unlikely that every patient characteristic could be so easily dichotomised and form the basis for one pair of disjoint rules. Therefore, in a more realistic system we would expect the size of the AF to vary between patients, instead of being uniform.

7. Conclusion

We have demonstrated a simple methodology for combining the explanatory power of argumentation frameworks and argumentative logics with the precision of Bayesian techniques. Our approach involves keeping the logical and the probabilistic components distinct and letting them interact in a productive way; our results show the practicality of this approach.

In addition, our work supports the development of richer semantics for argumentation, and a better understanding of how to weigh up competing arguments. Our examples have used relatively sparse Bayesian networks, leading to simple argumentation structures; future work will concentrate on using richer network structures and argumentation semantics, in particular the development of *chains* of arguments, and the integration of non-probabilistic arguments into our structure.

Acknowledgements

Our thanks to Rory Steele and John Fox for ideas and discussions that contributed to this paper.

References

- Amgoud, L., Cayrol, C., and Lagasque-Schiex, M.-C., 2004, "On Bipolarity in Argumentation frameworks," in: *10th International Workshop on Non-Monotonic Reasoning (NMR 2004)*, J.P. Delgrande and T. Schaub eds., *Whistler, Canada, June 6–8, 2004, Proceedings*. pp. 1–9.
- A/S, H.E.: 1989, 'Hugin'.
- Borak, J. and Veilleux, S., 1982, "Errors of intuitive logic among physicians," *Soc. Sci. Med.* **16**, 1939–1947.
- Clark, P. and Niblett, T., 1987, "Induction in Noisy Domains," in: *Proceedings of the 2nd European Working Session on Learning. Bled Yugoslavia*:. Sigma Press.
- Cristofanilli, M., Hayes, D., Budd, G., Ellis, M., Stopeck, A., Reuben, J., Doyle, G., Matera, J., Allard, W., Miller, M., Fritsche, H., Hortobagyi, G., and Terstappen, L., 2005, "Circulating tumor cells: A novel prognostic factor for newly diagnosed metastatic breast cancer," *J Clin Oncol* **23**, 1420–1430.
- Dung, P., 1995, "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games," *Artificial Intelligence* **77**, 321–357.
- Fox, J. and Parsons, S., 1997, "On Using Arguments For Reasoning About Actions And Values," in: *Proc AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, Stanford.
- Franklin, B., 1887, *Collected Letters*, Putnam New York.

- Friedman, 2004, "Inferring cellular networks using probabilistic graphical models," *Science* **303**, 799–805.
- Galea, M., Blamey, R., Elston, C.E., and Ellis, I., 1992, "The Nottingham Prognostic Index in primary breast cancer," *Breast Cancer Research and Treatment* **3**, 207–219.
- Gard, R., 1961, *Buddhism*, George Braziller Inc New York.
- Hunter, A. and Besnard, P., 2001, "A logic-based theory of deductive arguments," *Artificial Intelligence* **128**, 203–235.
- Kahneman, D. and Tversky, A., 1973, "On the psychology of prediction," *Psychol. Rev.* **80**, 237–251.
- Korb, K.B. and Nicholson, A.E., 2003, *Bayesian Artificial Intelligence*, London: Chapman and Hall / CRC Press.
- Krause, P., Amblar, S., Elvang-Goransson, M. and Fox, J., 1995, "A logic of argumentation for reasoning under uncertainty," *Computational Intelligence* **11**, 113–131.
- McConachy, R., Korb, K.B., and Zukerman, I., 1998, "A Bayesian approach to automating argumentation," in *Proceedings of New Methods in Language Processing & Computational Natural Language Learning (NeMLaP3/CoNLL98)* D.M.W. Powers ed., pp. 91–100.
- McPherson, K., Steel, C., and Dixon, J.C., 2000, "Breast cancer: Epidemiology. Risk factors and Genetics," *BMJ* **321**, 624–628.
- Michalski, R., Mozetic, I., Hong, J., and Lavrac, N., 1986, "The Multi-Purpose Incremental Learning System AQ15 and its Testing: Application to Three Medical Domains," in: *Proceedings of the Fifth National Conference on Artificial Intelligence*, Philadelphia PA, pp. 1041–1045.
- Parsons, S., 2003, "Order of magnitude reasoning and qualitative probability," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* **11**(3), 373–390.
- Parsons, S., 2004, "On precise and correct qualitative probabilistic reasoning," *International Journal of Approximate Reasoning* **35**, 111–135.
- Pearl, J., 2000, *Causality: Models, Reasoning, and Inference*, Cambridge: Cambridge University Press.
- Pollock, J.L., 1999, "Rational Cognition in OSCAR," in: *ATAL*, pp. 71–90.
- Poole, D., 2002, "Logical argumentation, abduction, and Bayesian decision theory: A Bayesian approach to logical arguments and its application to legal evidential reasoning," *Cardozo Law Review* **22**, 1733–1745.
- Prakken, H. and Sartor, G., 1996, "Argument-based extended logic programming with defeasible priorities," in *Working Notes of 3rd ModelAge Workshop: Formal Models of Agents*, P.-Y. Schobbens, ed., Sesimbra, Portugal.
- Quinn, M. and Allen, E., 1995, "Changes in incidence of and mortality from breast cancer in England and Wales since introduction of screening," *BMJ* **311**, 1391–1395.
- Richards, M., Smith, I., and Dixon, J., 1994, "Role of systemic treatment for primary operable breast cancer," *BMJ* **309**, 1263–1366.
- Ries, L., Eisner, M., Kosary, C., Hankey, B., Miller, B., Clegg, L., Mariotto, A., Feuer, E., and Edwards, B., 2004, *SEER Cancer Statistics Review 1975-2001*, National Cancer Institute.
- Saha, S. and Sen, S., 2004, "A Bayes Net approach to Argumentation," in: *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'04)*, Vol 3, 1436–1437.
- Spirtes, P., Glymour, C., and Scheines, R., 1993, *Causation, Prediction, and Search*, Cambridge MA: MIT Press, second (2000) edition.
- Sutton, D. and Fox, J., 2003, "The syntax and semantics of PROforma," *J Am Med Inform Assoc.* **10**(5), 433–443.
- Veer, L., Paik, S., and Hayes, D., 2005, "Gene expression profiling of breast cancer: A new tumor marker," *J Clin Oncol.* **23**, 1631–1635.
- Williamson, J., 2005, *Bayesian nets and Causality: Philosophical and Computational Foundations*, Oxford: Oxford University Press.

- Wittig, F. and Jameson, A., 2000, "Exploiting Qualitative Knowledge in the Learning of Conditional Probabilities of Bayesian Networks," In: C. Boutilier and M. Goldszmidt (eds.): *Uncertainty in Artificial Intelligence: Proceedings of the Sixteenth Conference*.
- Zwitter, M. and Soklic, M., 1988, "Breast Cancer Characteristics and Recurrence Data."