Using defeasible logic for a window on a probabilistic database:
Some preliminary notes

James Cussens[1] and Anthony Hunter[2]
[1]Centre for Logic and Probability in IT, King's College, Manresa Road, London, SW3 6LX
[2]Department of Computing, Imperial College, 180 Queen's Gate, London, SW7 2BZ

Introduction

Traditionally there has been a dichotomy between the probabilistic and logical views on uncertainty reasoning in AI. However, there does seem to be an intuitive overlap of the two views. Furthermore, given the relatively unclear understanding of the logical view, there is an argument for trying to establish the nature of the overlap between the probabilistic and logical views. Indeed there is some interesting relevant work that examines this overlap (including Pearl 1988, Neufeld 1990, Baccus 1989).

There is also a technological motivation for examining the inter-relationship between the probabilistic and logical views on uncertainity reasoning. In particular, in machine learning there seem to be some situations where a probabilistic representation has advantages over the logical one and others where a logical representation is superior (Gabbay 1991).

The particular research question we wish to focus on here is: Can we take a probabilistic database and view useful features of that database in a defeasible logic? In addressing this question, we develop an inference system for a probabilistic database that is based on the 'Principle of Total Evidence' and show how this is equivalent to a defeasible logic inference system with rule selection based on the specificity of the conditions of the rule. We compare the resulting inference systems with the 'extreme' probabilities system (Pearl 1988).

We develop our logics as Labelled Deductive Systems. This is part of a larger framework for logics in nonmonotonic reasoning (Gabbay 1989). An interesting subsidiary question in relating probabilistic and logical inference systems is: Does the relationship we establish provide a justification for defeasible logics with explicit representation of preference?

A brief review of defeasible logics

Within artificial intelligence there is a requirement for a "natural" representation of defeasible knowledge such as [1] and [2].

[1] "A match lights if struck"
[2] "A match doesn't light if the match is wet"

Such statements constitute "general" or default knowledge. However representing this kind of knowledge is not straightforward in classical logics. A number of approaches have been proposed (for reviews see Besnard 1989, Etherington 1988, Bell 1990, and Donini 1990). Though inter-relationships have been identified between approaches, there is as yet, no clear consensus. Here we focus on defeasible logics.

A number of formalisms for capturing aspects of defeasible reasoning have been proposed (Nute 1988, Pollock 1987, Hunter 1990, Laenens 1990, Loui 1986, Poole 1988), and are motivated by attempting to develop a lucid representation that does not involve listing all exceptions as conditions to rules, and in particular by supporting reasoning with competing arguments, and reasoning with different perspectives.

However, in developing formalisms for defeasible reasoning, if for example the following assertions, [3] and [4], are represented in the knowledge-base with [1] and [2], we do not want to infer a contradiction.

[3] "A match is struck"
[4] "A match is wet"

Some formalizations of defeasible logic, such as LDR (Nute 1988), resolve such a conflict be selecting the clause with the most specific antecedent. Other formalizations, (such as Loui 1986), base the resolution of such arguments on a number of criteria such as a preference for clauses that have more of the antecedent literals satisfied by monotonic reasoning.

Alternatively, there are formalizations with mechanisms that select the most appropriate clause based on a specified explicit ordering, such as KAL (Hunter 1990), and Ordered Logic (Vermier 1989, Laenens 1990). For example, we make non-monotonic inferences from an ordered set of clauses by only using the highest ordered formulae, with different defeasible logics defined to order formulae on different criteria. Such an approach avoids the unnatural representation of listing all exceptions to each clause as extra conditions in the antecedent, and furthermore provides an illuminating view on the structure and dynamics of a defeasible database. Under certain circumstances these approaches seem to be directly related to that of partially-ordered default theories (Brewka 1989).

The formalism LDR also captures the notion of a "defeater", which is a rule such that if the antecedent is satisfied, the rule defeats other rules lower in the ordering, but for which the consequent can not be detached. Defeater rules constitute arguments 'against' an inference. As with defeasible rules in LDR, defeater rules are selected on the basis of specificity.

Principle of Total Evidence

The idea here is that when assessing the probability of a given event α, we should calculate its probability value conditional on all the available evidence. This means selecting the most specific reference class within which to situate the event in question.

Let S = 'Match is struck', L = 'Match lights', and W= 'Match is wet'. Suppose that $p(L \mid S) = 0.8$ and $p(L \mid S \wedge W) = 0.1$. If we wish to assess the probability that a given match lights and all we know (all the evidence) is that it is struck, then using the principle of total evidence we infer that it has probability 0.8 of lighting. If we now learn that the match is struck and is wet then we employ the more specific conditional probability which will inform us that the match has a very low chance of lighting i.e. 0.1. In both cases, we use all the available evidence.

If we view an inference system as deriving propositions from a database, then by using the Principle of Total Evidence as the basis of an inference system, we can infer propositions from a probabilistic database. This contrasts with the usual view of a probabilistic database where from probability values one may only derive further values (or intervals of values).

It is clear that the probabilty-theoretic Principle of Total Evidence is analogous to the notion of specificity as used in defeasible logic. By using the Principle of Total Evidence in a probabilistic inference system, we hope to clarify this relationship. In the next section we develop an inference system based on the Principle of Total Evidence, and after that develop an equivalent logic-based inference system.

An inference system for a probabilistic database

The folllowing is a definition of a probabilistic database:

1. Let $K = \{ k_1, \ldots, k_n \}$ be a set of **attributes**. We can consider it be the set of all attributes of an event that we wish to consider in our database. If $q \in K$, then q and $\neg q$ are both termed **descriptors**, and term q as the complementary descriptor of $\neg q$, and $\neg q$ as the complementary descriptor of q. For any descriptor α, we denote $\alpha^\wedge$ as the complementary descriptor of α.

2. Let e be a **simple event** in the space generated by the set of attributes K, ie $e = q_1 \wedge q_2 \wedge .. \wedge q_n$ where for all i, s.t. $1 \leq i \leq n$, $q_i = k_i$ or $q_i = \neg k_i$. A **total description of e**, denoted [e], is defined as follows: if $e = q_1 \wedge \ldots \wedge q_n$, then $[e] = \{ q_1, \ldots, q_n \}$.

3. A **scenario** $\sigma(e)$ is a partial description of an event e, in other words $\sigma(e) \subseteq [e]$ holds. If there is a descriptor α s.t. $\alpha \in [e]$ and $\alpha \notin \sigma(e)$, then α is **open** by $\sigma(e)$.

4. A **probabilistic database**, denoted Δ, is **complete** w.r.t. a descriptor α if for any scenario $\sigma(e)$ where α is open by $\sigma(e)$, there is a conditional probability statement of the following form:

$$\text{'}p(\alpha | \bigwedge_i \beta_i) = \zeta\text{'} \quad \in \quad \Delta$$

where descriptors($\bigwedge_i \beta_i$) = $\sigma(e)$ and the function 'descriptors' is defined as follows:

descriptors($\delta \wedge \gamma$) = { δ } $\cup$ descriptors(γ) if δ is atomic
descriptors(δ) = { δ } if δ is atomic

Restricting our inference systems to using complete databases is, we feel, justified in certain applications. In particular, we are developing formalisms for applications in machine learning where obtaining complete databases is often not a problem.

Given a probabilistic database Δ that is complete w.r.t. α, and a scenario $\sigma(e)$, we wish to ascertain a probability value for α if α is open by $\sigma(e)$. If we use the Principle of Total Evidence, then there is only one conditional probability statement 'p($\alpha | \bigwedge_i \beta_i$) = ζ' $\in \Delta$ that can be used.

For reasoning about properties, if we compute probability values using Δ and $\sigma(e)$, we can make inferences about the ground properties: if we use some threshold θ, say 1/2 or 3/4, we can infer α from Δ with $\sigma(e)$, if the probability of α is greater than the threshold θ. Note, we must have θ at least 1/2 for consistency.

Furthermore we can reason with a probabilistic database qualitatively. For this we introduce the following definitions:

[D1] p(x | y) is high IFF 'p(x | y) = ζ' $\in$ Δ and $\zeta > \theta$

[D2] p(x | y) is low IFF 'p(x | y) = ζ' $\in$ Δ and $\zeta < (1 - \theta)$

[D3] p(x | y) is medium IFF 'p(x | y) = ζ' $\in$ Δ and $(1 - \theta) \leq \zeta \leq \theta$

Hence we can use the following qualitative relations between our conditional probability values:

[R1] p(x | y) is high IFF p($\neg$x | y) is low

[R2] p(x | y) is medium IFF p($\neg$x | y) is medium

We reason with a probabilistic database Δ, and a scenario $\sigma(e)$, as a pair (Δ, $\sigma(e)$). We formalize an inference from this pair, as follows, where θ is the probabilistic threshold value, and $\models$ is the inference relation:

(Δ, $\sigma(e)$) $\models$ α if $\alpha \in \sigma(e)$

(Δ, $\sigma(e)$) $\models$ α if 'p($\alpha | \bigwedge_i \beta_i$) = ζ' $\in \Delta$ and $\zeta > \theta$
and p($\alpha | \bigwedge_i \beta_i$) is the preferred conditional probability for α in Δ w.r.t $\sigma(e)$

The notion of preferred conditional probability for a descriptor α is intended to capture the intuition of the Principle of Total Evidence, and is defined as follows:

p($\alpha | \bigwedge_i \beta_i$) is the **preferred conditional probability** for α in Δ w.r.t $\sigma(e)$
if 'p($\alpha | \bigwedge_i \beta_i$) = ζ' $\in \Delta$
and for all i $\beta_i \in \sigma(e)$
and for all other 'p($\delta | \bigwedge_j \gamma_j$) = ψ' $\in \Delta$ s.t. head(α) = head(δ),
if for all j $\gamma_j \in \sigma(e)$ then descriptors($\bigwedge_j \gamma_j$) $\subseteq$ descriptors($\bigwedge_i \beta_i$)

where we define the function 'head' as head(α) = α, and head($\neg\alpha$) = α. Using these definitions, we can for example infer E follows from a scenario { A }, and a probabilistic database Δ, using the conditional statement 'p(E | A) = ζ' $\in \Delta$, where p(E | A) is high.

Forming a defeasible logic window

We now consider what a window on a probabilistic database is, and how we can form it. We start by considering the following example. The following conditional probabilities, (1), (2) (3) and (4) are in Δ, where Δ is complete with respect to the descriptor E.

(1) $p(A \mid D) = 1$
(2) $p(E \mid A)$ is high
(3) $p(\neg E \mid A \wedge D)$ is high
(4) $p(\neg E \mid D)$ is high

where if 'A' denotes 'adult', 'D' denotes 'dropout', and 'E' denotes 'employed', then (1) represents 'all dropouts are adults', (2) represents 'most adults are employed', (3) represents 'most dropouts, who are adults, are not employed' and (4) represents 'most dropouts are not employed'.

Since Δ is a complete probabilistic database w.r.t descriptor E, we can consider forming a defeasible logic window on Δ. We consider here a window based on just one descriptor, namely E. However, in general we could form a window based on more than one descriptor. The form of the defeasible rules in the defeasible logic representation is as follows, where a(i) is an ordering label, α is the consequent, β is the antecedent, and $\rightarrow$ is a conditional implication symbol:

$$a(i) : \quad \beta \rightarrow \alpha$$

The consequent is a descriptor, and the antecedent is a conjunction of descriptors (Our terminology is such that 'descriptor' is equivalent to the logical use of 'literal'). The ordering label represents a notion of preference for the rule over other rules with the same consequent, or its negation. A database of defeasible logic rules, denoted Γ, form a partial order. We define below a proof theory to support such databases, where if $\beta \in \sigma(e)$ then $c{:}\ \beta \in \Gamma$:

[A0] $\quad \Gamma \mid\!\sim \alpha$ if $a(i){:}\ \beta_1 \wedge .. \wedge \beta_n \rightarrow \alpha \in \Gamma$, and $c{:}\ \beta_1 \in \Gamma, \ldots,$ and $c{:}\ \beta_n \in \Gamma$

[A1] $\quad$ IF $c{:} \alpha \in \Gamma$ THEN $\Gamma \mid\!\sim \alpha$

Essentially, by [A0] we allow the inference α from a database Γ, if there is a rule in Γ, s.t. (1) the antecedent is satisfied, and (2) there is no rule in Γ that has the antecedent satisfied and is preferred by the ordering relation. We represent the logical rules formed from (1), (2), (3) and (4), as (1'), (2'), (3') and (4') respectively:

(1') $\quad a(i_1) : \; D \rightarrow A$
(2') $\quad a(i_2) : \; A \rightarrow E$
(3') $\quad a(i_3) : \; A \wedge D \rightarrow \neg E$
(4') $\quad a(i_4) : \; D \rightarrow \neg E$

where the ordering on the rules is $i_3 > i_2$ and $i_3 > i_4$ and $i_2 \parallel i_4$. In forming these rules, we have assumed that the conditional probability 'is high' means that the probability value is greater than our threshold value θ. Using the defeasible logic, we can make inferences from (2'),(3'), and (4'). For example, if we have a scenario where A and D hold, then we can infer $\neg E$ using the proof rules.

Essentially, if we use the 'Principle of Total Evidence' in our probabilisitic reasoning system, in order to make inferences from our probabilistic database, we can mimic the inference procedure in our defeasible logic system. The notion of 'Principle of Total Evidence' which is used to select the most appropriate conditional probability, is reflected in the logical system by selecting the highest rule in the partial order which has the antecdent satisfied.

Suppose we extend our probabilistic database with the following conditionals, where Y denotes 'youth', and then (5) represents 'some youths are employed', (6) represents 'some adult youths are employed', (7) represents 'most dropout youths are not employed, and (8) represents 'most dropout adult youths are not employed':

(5) $p(E \mid Y)$ is medium
(6) $p(E \mid Y \wedge A)$ is medium
(7) $p(\neg E \mid Y \wedge D)$ is high
(8) $p(\neg E \mid Y \wedge A \wedge D)$ is high

We now need to consider how to represent the medium conditional probabilities. Since by [R2] we also have the following:

(5a) $p(\neg E \mid Y)$ is medium
(6a) $p(\neg E \mid Y \wedge A)$ is medium

If we reason with the Principle of Total Evidence together with the threshold value high, then we can infer neither E nor ¬E given a scenario { Y }. However, if we are to form a defeasible logic window, how do we support corresponding reasoning? A solution is to represent defeater rules in our logic window. Such rules, if selected by the specificity ordering would over-ride any defeasible inferences lower in the ordering, but not allow their own consequent to be detached. In other words the defeater rules block inferences from the window, but do not add directly to the inferences. The form of the defeasible rules in the defeasible logic representation is as follows, where b(i) is an ordering label, α is the consequent, β is the antecedent, and $\rightarrow$ is a conditional implication symbol:

$$b(i): \quad \beta \rightarrow \alpha$$

The ordering label represents a notion of preference for the rule over other rules with the same consequent, or its negation. However, we need to amend the proof theory by replacing the rule [A0] with [A2] as presented below:

[A2] $\Gamma \mid\!\sim \alpha$ if $a(i): \beta_1 \wedge .. \wedge \beta_n \rightarrow \alpha \in \Gamma$, and $c: \beta_1 \in \Gamma, .. ,$ and $c: \beta_n \in \Gamma$
and NOT($a(j): \delta_1 \wedge .. \wedge \delta_m \rightarrow \alpha^\wedge \in \Gamma$, and $c: \delta_1 \in \Gamma, ..,$ and $c: \delta_n \in \Gamma$, and $j > i$)
and NOT($b(j): \delta_1 \wedge .. \wedge \delta_m \rightarrow \alpha \in \Gamma$, and $c: \delta_1 \in \Gamma, ..,$ and $c: \delta_n \in \Gamma$, and $j > i$)

The definition [A2] extends the definition [A0] by adding the extra condition that for a consequent of a rule to be detached, there is no defeater rule in the database that also has its antecedent satisfied and is also higher in the ordering. In the next section we clarify the relationship between the probabilistic representation, and the defeasible logic window.

Formalizing the relationship between the database and window

Given a complete probabilistic database Δ w.r.t a descriptor α, we denote a window function on Δ, w.r.t. to α, as $w(\Delta, \alpha)$, and define it as follows, where θ is the probabilistic threshold value, and $\mu, \lambda \in \{ a, b \}$:

[F1] IF '$p(\alpha \mid \delta) = \zeta$' $\in \Delta$ and $p(\alpha \mid \delta)$ is high
THEN $a(i): \delta \rightarrow \alpha \in w(\Delta, \alpha)$

[F2] IF '$p(\alpha \mid \delta) = \zeta$' $\in \Delta$ and $p(\alpha \mid \delta)$ is medium
THEN $b(i): \delta \rightarrow \alpha \in w(\Delta, \alpha)$ AND $b(i): \delta \rightarrow \alpha^\wedge \in w(\Delta, \alpha)$

[F3] IF '$p(\alpha \mid \delta) = \zeta$' $\in \Delta$ and '$p(\varphi \mid \gamma) = \xi$' $\in \Delta$
and head(α) = head(φ) and descriptors(δ) $\subseteq$ descriptors(γ)
and $\mu(i): \delta \rightarrow \alpha \in w(\Delta, \alpha)$ and $\lambda(j): \gamma \rightarrow \varphi \in w(\Delta, \alpha)$
THEN $(j, i) \in {}_{\alpha}\!\succ$

For a scenario $\sigma(e)$, we define the rewrite of the scenario, denoted $C(\sigma(e))$, as follows:

[F4] IF $\alpha \in \sigma(e)$ THEN $c: \alpha \in C(\sigma(e))$

From the window function applied to one, or more, predicates, together with rewrite of the scenario, we can form the defeasible logic database as follows:

$$\Gamma = w(\Delta, \alpha_1) \cup \ldots \cup w(\Delta, \alpha_n) \cup C(\sigma(e))$$

where the ordering relation, denoted $\succ$, is defined as $\succ =_{def} {}_{\alpha_1}\!\succ \cup \ldots \cup {}_{\alpha n}\!\succ$.

Essentially [F1] rewrites all the high conditional probabilities as defeasible rules, and [F2] rewrites all the medium conditional probabilities as defeater rules. [F3] rewrites the implicit ordering in the inference system based on the Principle of Total Evidence to an explicit ordering, and [F4] takes the assertions in the scenario and labels all of them with the symbol 'c' within the defeasible logic database. We can show the following equivalence between the probabilistic reasoning system and defeasible logic reasoning system:

[Result] For any descriptor, α s.t. $\alpha \in \{ \alpha_1, \ldots, \alpha_n \}$, then the following equivalence holds, where |~ is the defeasible logic consequence relation as defined by [A1] and [A2], and |≡ is the probabilistic consequence relation, where $\Gamma = w(\Delta, \alpha_1) \cup \ldots \cup w(\Delta, \alpha_n) \cup \Xi(\sigma(e))$:

$$(\Delta, \sigma(e)) \models \alpha \text{ iff } \Gamma \mid\sim \alpha$$

The ordering relation we have defined for our defeasible relation makes explicit the notion of preference based on specificity. Hence, for any rules on our logic database of the form:

$$a(i) : \alpha \rightarrow \beta$$
$$a(j) : \gamma \rightarrow \delta$$

if head(β) = head(δ), and $i \geq j$, then descriptors(γ) $\subseteq$ descriptors(α).

An Alternative View: Pearl's 'Extreme' Probabilities

Pearl claims that an approach such as ours can only be used if we use extreme probabilities. That is, we only adopt the defeasible rule X -> Y if $P(Y|X) = 1 - \varepsilon$ where ε is infinitesimally small. This contrasts with our approach were a probability has only to be 'high' (i.e. > β) rather than 'extremely high' for there to be a corresponding rule in the defeasible logic window.

Three axioms, which Pearl takes to be intuitive, are sound in his interpretation. They are

1.	Triangularity	$a \rightarrow b,\ a \rightarrow c$	$\Rightarrow$	$(a \wedge b) \rightarrow c$
2.	Bayes	$a \rightarrow b,\ (a \wedge b) \rightarrow c$	$\Rightarrow$	$a \rightarrow c$
3.	Disjunction	$a \rightarrow c,\ b \rightarrow c$	$\Rightarrow$	$(a \vee b) \rightarrow c$

These are not sound under our approach. To take the first example, suppose $\sigma(e) = \{a, b\}$ and Δ = {'p(c | a) > θ', 'p(b | a) > θ', 'p($\neg$c | a $\wedge$ b) > θ'} then we will have $a \rightarrow b$, $a \rightarrow c$ and $(a \wedge b) \rightarrow \neg c \in w(\Delta,c)$, contradicting triangularity. Note that such a Δ is indeed possible in our system, though it would not be if θ were infinitesimally close to one, as used in Pearl's system.

It is clear that Pearl's interpretation allows for nice intuitive axioms and Pearl gives examples of reasoning with it. But there are problems: is this a reasonable interpretation? Do people use extreme probabilities? We feel that the answer to both questions is no.

Pearl argues that any logic based on non-extreme logic will be 'extremely complicated'. He argues that the reasoning behind actual default rules is based on the fairly crude 'almost all' logic rather than what he calls the 'logic of the majority'. However we feel that (1) Pearl's 'infinitesimal' approach is considerably less intuitive than our own and (2) ours is an adequately simple method.

Discussion

Once the Principle of Total Evidence has been used to build an inference system for a probabilistic database, it is straightforward to construct an equivalent defeasible logic system. This highlights the similarity between the Principle of Total Evidence and the use of specificity orderings in defeasible logics.

Advantages that accrue from the equivalence of the two approaches includes a mechanism by which probabilistic data may be to used to derive logical databases. This may be enhanced by the possibility of merging such logical databases with background knowledge that is also in a logical form.

We have made use of labels to establish an equivalence between a probabilistic database and a defeasible logic window onto it. Currently, we are studying probabilisitc databases that include representation of the level of certainty we have in our probability values. We are also examining non-probabilistic belief functions, eg. Dempster-Shafer belief functions (Shafer 1976). In order to represent such features in a logic window we use an augmented labelling system. Recent work on Labelled Deductive Systems (Gabbay 1989) provides a framework for further work in this direction.

For translation from the probabilitic representation, Baccus (1989) adopts an analogous threshold notion (i.e. if p(a|b) > θ holds in the probabilistic representation, then $a \rightarrow b$ holds in the logic representation). However, our logical system also adopts a notion of defeater rules which enriches our knowledge representation capability. Furthermore, this notion is clearly supported by its probabilistic interpretation.

Interesting results from this work include the failure of the defeasible logic, as defined here, to meet the minimal properties (as presented in Gabbay 1985) expected of a non-monotonic consequence relation. Within the proof-theoretic framework of Labelled Deductive Systems we are considering strengthening the defeasible logic.

Another interesting direction is to consider the axioms such as 1 - 3 above as defeasible meta-level axioms. Then if we use incomplete probabilistic databases, we can use the axioms to non-monotonically complete the database. This direction could form part of a larger initiative to use defeasible meta-rules for making assumptions about the nature and usage of a probabilistic database. For example, in Grosof (Grosof 1988) non-monotonic meta-assumptions, such as independence, are made about the nature of the database.

Acknowledgements

This work is currently being funded by UK SERC grants GR/G 29861 and GR/G 29854 for the Rule-based Systems project. Special thanks are due to Dov Gabbay, Donald Gillies, Ng, and Barry Richards.

References

Baccus F (1989) A modest, but semantically well-founded, inheritance reasoner, IJCAI-89, Morgan Kaufmann
Bell J (1990) Non-monotonic reasoning, non-monotonic logics, and reasoning about change, Artificial Intelligence Review, 4, 79 - 108
Besnard P (1989) An Introduction to Default Logic, Springer
Brewka G (1989) Preferred sub-theories: An extended logical framework for default reasoning, IJCAI-89, Morgan Kaufmann
Donini F (1990) Non-monotonic reasoning, Artificial Intelligence Review, 4, 163 - 210
Etherington D (1988) Reasoning with incomplete information, RNAI, Pitman
Gabbay D (1985) Theoretical foundations for non-monotonic reasoning in expert systems, in Apt K, Proc. NATO Advanced Study Institute on Logics and Models of Concurrent Systems, Springer, 439 - 457
Gabbay D (1989) Labelled Deductive Systems, Technical Report, Department of Computing, Imperial College, London, also to be published by Oxford University Press
Gabbay D, Gillies D, Hunter A, Muggleton S, Ng Y, & Richards B (1991) The rule based systems project: using confirmation theory and non-monotonic logics for incremental learning, Proc. of the First Workshop on Inductive Logic Programming, Academic Press
Grosof B (1988) Non-monotonicity in probabilistic reasoning, Uncertainty in Artificial Intelligence, 3
Hunter A (1990) KAL: A linear ordered logic for non-monotonic reasoning, Draft Paper, Department of Computing, Imperial College, London
Laenens E & Vermier D (1990) A fixpoint semantics for ordered logic, J Logic & Computation, 1,
Loui R (1986) Defeat amongst arguments: a system of defeasible inference, Technical Report No 130, Department of Computer Science, University of Rochester, New York
Neufeld E, Poole D, and Aleliunas R (1990) Uncertainty in Artificial Intelligence, 4, 121 - 132
Nute D (1988) Defeasible reasoning and decision support systems, Decision Support Systems, 4, 97-110
Pearl J (1988) Probabilistic Reasoning in Intelligent Systems, Morgan Kaufmann
Poole D (1988) Logical framework for default reasoning, Artificial Intelligence, 36, 27 - 48
Pollock J (1987) Defeasible reasoning, Cognitive Science, 11, 481 - 518
Shafer G (1976) A Mathematical Theory of Evidence, Princeton University Press
Vermier D, Nute D & Geerts P (1989) A logic for defeasible perspectives, Proc. of the 1988 Tubingen Workshop on Semantic Networks and Non-monotonic Reasoning, vol 1, 1 - 27, SNS-Bericht 89-48