ON THE RATIONAL SCOPE OF PROBABILISTIC RULE-BASED INFERENCE SYSTEMS[1]

by

Shimon Schocken

Information Systems Area
Graduate School of Business Administration
New York University
90 Trinity Place
New York, N.Y. 10006

June 1987

Center for Research on Information Systems
Information Systems Area
Graduate School of Business Administration
New York University

Working Paper Series

STERN #IS-87-59

[1]This Paper will appear in Uncertainty In Artificial Intelligence, Volume II, L.N. Kanal and J.F. Lemmer [Eds.], Amsterdam: North-Holland, 1987

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

ON THE RATIONAL SCOPE OF PROBABILISTIC RULE-BASED INFERENCE SYSTEMS

Shimon Schocken*

Information Systems Area
Graduate School of Business Administration
New York University

Belief updating schemes in artificial intelligence may be viewed as three dimensional languages, consisting of a *syntax* (e.g. probabilities or certainty factors), a *calculus* (e.g. Bayesian or CF combination rules), and a *semantics* (i.e. cognitive interpretations of competing formalisms). This paper studies the rational scope of those languages on the syntax and calculus grounds. In particular, the paper presents an endomorphism theorem which highlights the limitations imposed by the conditional independence assumptions implicit in the CF calculus. Implications of the theorem to the relationship between the CF and the Bayesian languages and the Dempster-Shafer theory of evidence are presented. The paper concludes with a discussion of some implications on rule-based knowledge engineering in uncertain domains.

1. INTRODUCTION

In order for a computer program to be a plausible model of a (more or less) rational process of human expertise, the program should be capable of representing beliefs in a language that is (more or less) calibrated with a well-specified normative criterion, e.g. the axioms of subjective probability [1], the theory of confirmation [2], formal logic, etc. According to Shafer and Tversky, the building blocks of a probabilistic language are syntax, calculus, and semantics [3]. The syntax is a set of numbers, commonly referred to as *degrees of belief* (e.g. standard probabilities or certainty factors), used to parameterize uncertain facts, inexact rules, and competing hypotheses. Typically, a set of atomic degrees of belief is elicited directly from a human expert, while compound degrees of belief are computed through a set of operators collectively known as a belief calculus. The semantics of the language can be viewed as a mapping from a real-life domain of expertise onto the belief language. This mapping provides a cognitive interpretation as well as descriptive face-validity to both the syntax and the calculus dimensions of the language.

Given the critical role that a belief language plays in determining both the low-level mechanics and the high-level ordinal ranking of the recommendations generated by an expert system, it is clear that the implicit rationality of the language is directly related

---

*This Research was carried out while the author was at the Decision Sciences Department, the Wharton School of the University of Pennsylvania. The author's current address is 624 Tisch Hall, Washington Square, New York, NY 10003.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

to both the internal and external validities of computer-based expertise. By 'rationality' I refer here to the normative criteria of *consistency* and *completeness* [1] as well as to the psychometric criteria of *reliability* and *validity* [4]. It is argued that the performance of any expert, whether a human being or a computer program, should be evaluated and rated along those lines.

The two mainstream belief languages in rule-based inference systems are the normative *Bayesian* and the descriptive *certainty factors* (CF) languages, the latter being representative of a wide variety of ad-hoc calculi of uncertainty. It seems that the CF method is currently the most widely used belief language in applied expert systems, primarily due to the popularity of such CF-based shells as EMYCIN [5], M.1 [6], and Texas Instrument's Personal Consultant [7]. Bayesian inference has been traditionally much less popular, with the exception of some notable examples, e.g. PROSPECTOR [8], which uses a Bayesian syntax and an ad-hoc version of a Bayesian calculus. Recently, new techniques designed to cope with the computational complexity of a complete Bayesian design are emerging, giving rise to the concept of a *Bayesian inference net* [Pearl,9].

Notwithstanding the critical importance of exploring the practical scope of non-categorical rule-based inference systems, few studies have compared belief languages on rational as well as cognitive grounds. Furthermore, practitioners are often oblivious to the theoretical limitations inherent in the representation and synthesis of degrees of belief. This has led to a number of commonly held misconceptions regarding some properties of the CF and the Bayesian languages, such as the following two conjectures:

C1: Classical Bayesian methods are either too simplistic or too complex: in order for a Bayesian updating procedure to be computationally feasible, strict statistical independence must prevail. This requirement is rarely met in practice, where interaction effects among clues and hypotheses make the Bayesian solution unmanageable on combinatorial grounds. The CF calculus, on the other hand, does not make explicit assumptions of statistical independence; therefore, it can be used to model realistically complicated problems that defy a normative Bayesian interpretation.

C2: Both the Bayesian and the CF calculi are special cases of the general Dempster-Shafer theory of evidence [10]. Hence, they may be construed as two alternative and competing belief languages, each specialized to deal with a particular class of problems and probabilistic designs.

Toward the end of the paper a rather different interpretation of both C1 and C2 will be presented. The organization of the paper is as follows: Section 2 presents some necessary background and terminology. Section 3 provides a brief review of the CF language and its rational (Bayesian) interpretation. Section 4 presents three lemmas that are further integrated into an endomorphism theorem. This theorem shows that the CF language is a special case of the Bayesian language. Similar results have been proven in the past by Adams [11], Heckerman [12], and Grosof [13]. Section 5 presents some preliminary thoughts about the computational complexity of wholistic (not conditionally independent) inference problems. The paper concludes with some implications on knowledge engineering and future research directions

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

2. BACKGROUND AND NOMENCLATURE

Consider an n-dimensional propositional space S defined over a certain domain of expertise, e.g. medical diagnosis. Each dimension of S is interpreted as an attribute of the domain, e.g. chest pain, headache, allergy, etc. An *inference problem* in S is a tuple $<h,e_1,...,e_m>$, $0 \leq m \leq n$, where the attribute h is interpreted as a prospective hypothesis and the attributes $e_1,...,e_m$ are interpreted as pieces of evidence relevant to h. Let $F:S \rightarrow [0,1]$ be a (possibly transcendental) joint distribution function defined over S. Although we don't have direct access to F, we assume that there exists a domain expert who is capable of making judgments that can be further interpreted as a subjective function BEL which approximates F. In particular, given an inference problem $<h,e_1,...,e_m>$, the *posterior belief* $BEL(h|e_1,...,e_m)$ reflects the expert's belief that h is true in light of the evidence $e_1,...,e_m$.

In order to avoid the apologetic debate of whether or not the function F exists, we note that F is presented here primarily for the sake of clear exposition. In fact, the relationship between BEL and F is at the center of an intensified philosophical debate that has been going strong for more than 300 years. In short, under a Bayesian interpretation (e.g. Ramsey), $BEL = P \equiv F$, where P is the standard Savage/de Finetti subjective probability function. Objectivists (like Popper) argue that F stands aloof from P (or, for that matter, from any personal BEL), and, hence, in general, $P \not\equiv F$. Proponents of the logical school of probability model BEL through Carnap's "confirmation" function $C(\cdot)$ [2]. A similar approach is taken by the certainty factors formalism, which sets BEL = CF.

Finally, pragmatic Bayesians (like myself) feel that BEL = P is our best shot at F, a shot whose accuracy is directly related to the operational characteristics of the elicitation procedure designed to construct P. Since P is subject to an internal axiomatic system, I term |BEL-P| an 'internal bias' and |BEL-F| an 'external bias.' Attempts to reduce those biases are termed 'debiasing' or 'corrective procedures' in the cognitive psychology literature [e.g. 14].

Let V be the subset of all "interesting" (i.e. non-arbitrary) inference problems $<h,e_1,...,e_m>$ defined over S. The following set of definitions partitions V into three classes of problems that vary in terms of their computational (and cognitive) complexity. This partitioning is a reflection of the fact that some problems that require expertise may be simple 'open and shut' cases, while other problems may be complicated and vague. In what follows, I wish to provide a more precise definition of this taxonomy of problems, based on the underlying complexity of their diagnostic structures.

Diagnostic Structure: the diagnostic structure of a problem $q = <h,e_1,...,e_m>$ is the conditional distribution $F(e_1,...,e_m|h)$.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

Weakly Decomposable Problems: the set of weakly decomposable problems WD is defined as follows:

$$WD = \{q \mid q = \{h, e_1, \ldots, e_m\} \in V \text{ and } F(e_1, \ldots, e_m|h) = F(e_1|h) \cdot \ldots \cdot F(e_m|h)\}$$

Decomposable Problems: the set of decomposable problems D is defined as follows:

$$D = \{q \mid q = \{h, e_1, \ldots, e_m\} \text{ is weakly decomposable and } F(e_1, \ldots, e_m) = F(e_1) \cdot \ldots \cdot F(e_m)\}$$

Wholistic Problems: a problem q is wholistic if $q \in V - WD$

Corollary: $D \subset WD \subset V \subset S$

Note that "decomposability" is a weaker notion of statistical independence. The latter requires that events be independent in all subsets, e.g. $P(abc) = P(a)P(b)P(c)$, $P(ab) = P(a)P(b)$, $P(ac) = P(a)P(c)$, and $P(bc) = P(b)P(c)$. Decomposability requires only the first constraint, i.e. $P(abc) = P(a)P(b)P(c)$.

3. THE CF LANGUAGE AND ITS RATIONAL INTERPRETATION

This section provides a brief account of the definition and interpretation of the certainty factors language, as stated by Shortliffe and Buchanan in [10]. Given a problem $q = \langle h, e_1, \ldots, e_m \rangle \in V$, the CF syntax approximates the posterior belief associated with q through the difference between a measure of increased belief (MB) and a measure of increased disbelief (MD) in the hypothesis h in light of the clues $\langle e_1, \ldots, e_m \rangle$:

$$-1 \leq CF(h|e_1, \ldots, e_m) = MB(h|e_1, \ldots, e_m) - MD(h|e_1, \ldots, e_m) \leq 1$$

The CF calculus is a set of operators designed to combine atomic CF's into compound CF's (e.g. compute BEL(h|a,b) from BEL(h|a) and BEL(h|b)). This paper focuses only on a subset of this calculus, denoted hereafter (M1-M2):

$$MB(h|a,b) = \begin{cases} 0 & \text{if } MD(h|a,b) = 1 \\ MB(h|a) + MB(h|b) \cdot (1 - MB(h|a)) & \text{otherwise} \end{cases} \quad \text{(M1)}$$

$$MD(h|a,b) = \begin{cases} 0 & \text{if } MB(h|a,b) = 1 \\ MD(h|a) + MD(h|b) \cdot (1 - MD(h|a)) & \text{otherwise} \end{cases} \quad \text{(M2)}$$

Note that (M1-M2) appears to convey a certain descriptive appeal: if you open the parentheses of (M1) for example, you obtain the sum of MB(h|a) and MB(h|b) minus their multiplicative interaction effect. The resulting combination rule is both commutative and associative, as one would expect.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

Shortliffe and Buchanan have also suggested a syntactical mapping from Bayesian probabilities to certainty factors, defined as follows:

$$MB(h|a) = \begin{cases} 1 & \text{if } P(h)=1 \\ \frac{\max\{P(h|a),P(h)\}-P(h)}{1-P(h)} & \text{otherwise} \end{cases} \qquad \text{(R1)}$$

$$MD(h|a) = \begin{cases} 1 & \text{if } P(h)=0 \\ \frac{\max\{P(h|a),P(h)\}-P(h)}{-P(h)} & \text{otherwise} \end{cases} \qquad \text{(R2)}$$

I term the (R1-R2) mapping a *rational interpretation* for three reasons. First, the mapping is intended to convey a certain degree of descriptive face-validity to the CF syntax. For example, (R1) represents the measure of increased belief in h in light of the piece of evidence a as a normalized difference between the posterior P(h|a) and the prior P(h). Second, the mapping relates CF's to a subset of the real interval [0,1] which is consistent with the seven rational postulates of Savage and de Finetti (i.e. the axioms of subjective probability). Third, my notion of a rational interpretation is consistent with Shortliffe and Buchanan who suggest: *"Behavior is irrational if actions taken or decisions made contradict the result that would be obtained under a probabilistic analysis of the behavior"* [10, p. 251].

Note in passing that (M1-M2) is very different from (R1-R2). The former pair of combination rules is the nucleus of the CF calculus, designed to compute the compound strength of belief of two parallel pieces of evidence. The latter pair of definitions is a suggested ex-post Bayesian interpretation of certainty factors which is not necessarily unique.

4. THE CF LANGUAGE AS A SPECIAL CASE OF THE BAYESIAN LANGUAGE

Lemma 1 If the MB combination rule (M1) is used to approximate the posterior belief associated with a problem $\langle h,e_1,...,e_m\rangle$, then (M1) is mutually consistent with the rational MB interpretation (R1) if and only if $\langle \neg h,e_1,...,e_m\rangle$ is decomposable.

Lemma 2 If the MD combination rule (M2) is used to approximate the posterior belief associated with a problem $\langle h,e_1,...,e_m\rangle$, then (M2) is mutually consistent with the rational MD interpretation (R2) if and only if $\langle h,e_1,...,e_m\rangle$ is decomposable.

Lemma 3 If the CF calculus (M1-M2) is used to approximate the posterior belief associated with a problem $\langle h,e_1,...,e_m\rangle$, and if (M1-M2) is mutually and jointly consistent with the rational interpretation (R1-R2), then both $\langle h,e_1,...,e_m\rangle$ and $\langle \neg h,e_1,...,e_m\rangle$ are weakly decomposable.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

The Endomorphism Theorem Let CF be the set of all problems that have an approximate posterior solution derived by the CF calculus (M1-M2). Let V be the set of all problems that have an approximate posterior solution derived by a Bayesian calculus. Let T be the rational interpretation (R1-R2). Under these conditions, T is an endomorphic transformation $T: CF \rightarrow WD \subset V$, where WD is the subset of weakly decomposable problems in V.

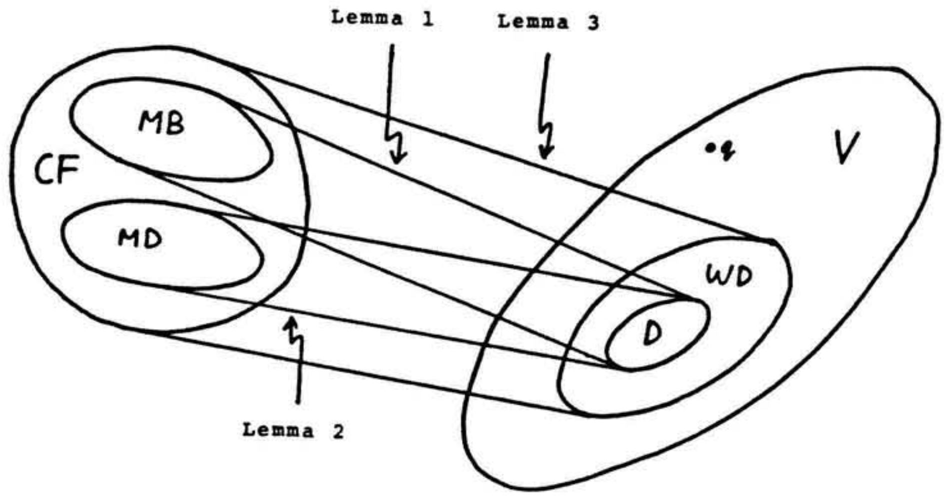

Figure 1

5. DISCUSSION

The endomorphism theorem says that the CF calculus (R1-R2) has a rational interpretation if and only if it is restricted to weakly decomposable problems. Under these conditions, the CF belief synthesis rule is equivalent to the likelihood ratio version of the Bayesian belief updating rule. As was mentioned at the beginning of the paper, the fact that the CF calculus makes implicit assumptions of conditional independence was proved elsewhere, e.g. by Adams [11], Heckerman [12], and Grosof [13]. The present theorem is useful in that it highlights the important implications of the CF/Bayesian relationship on the rational scope of CF-based inference systems.

In the pictorial illustration of the theorem, q is a wholistic problem that is outside the rational scope of the CF language (e.g. $q = \langle h,a,b \rangle$ has a "synergistic" diagnostic structure, i.e. $F(a,b|h) > F(a|h) \cdot F(b|h)$). At the same time, however, q does have a (complicated) Bayesian posterior belief, by virtue of its membership in V. This dichotomy means that one cannot trade rationality for efficiency, as is sometimes being done in AI. Furthermore, the likelihood that real-life problems exist in WD is very small, due to the underlying complexity of inference problems that require expertise [15].

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

The preceding paragraph implies that most CF-based expert systems (and, hence, most applied expert systems) are inconsistent with their rational interpretation (R1-R2). This finding is disturbing in view of the impressive decision-making performance of some CF-based systems [8]. There may be (at least) two potential explanations for the disparity between the narrow normative foundation and the de-facto face-validity of the CF language.

First, I argue that experienced knowledge engineers intuitively know that the endomorphism theorem is true, and, in fact, take advantage of it. In particular, designers of complicated expert systems often feel that the more granular the knowledge-base, the higher is the validity of the system [16]. This heuristic amounts to augmenting an evidence/hypotheses inference net with a multitude of sub hypotheses and intermediate states, designed to partition the knowledge-base and achieve a higher degree of granularity. This judicious decomposition is done in an attempt to explicitly account for interaction effects, and, thereby, induce more conditional independence on the evidence/hypotheses space, as was proposed by Charniak [15] and by Winter and Girse [17].

In the context of the present paper, we can describe this practice as follows: when a CF knowledge engineer faces a wholistic problem q which is outside the scope of a rational interpretation, he or she first modifies the diagnostic structure of the original problem, thus creating a transformation from q to $q' \in WD$, which is a rational CF territory. If there exists a problem q' whose diagnostic structure is indeed a plausible (weak) decomposition of q, a CF-based system applied to q' is likely to provide a (close) rational belief representation to q as well.

The second explanation of the CF descriptive/normative contrast may be that the original rational interpretation of certainty factors (R1-R2) is subject to doubt. In other words, it seems that the CF language is indeed a novel formalism that deserves a serious look, especially on practical and descriptive grounds. Indeed, the fact that the CF language has been going strong for more than a decade in spite of its unrealistically narrow rational interpretation suggests that the model is basically powerful although its normative foundation is weak. Hence, future research is needed to explore new interpretations to the CF language that will be more plausible on rational, cognitive, and philosophical grounds. An Example of such an undertaking may be found in Heckerman's work [12] on alternative probabilistic interpretations of certainty factors.

5.1. Conjecture C1 Revisited

We now turn to the casual conjecture C1, which attributes the impracticality of the Bayesian language vis a vis the CF language to the fact that real life domains of expertise are not statistically independent. The reader has perhaps realized by now that this statement is based on a semantic rather than a substantive argument. In particular, note that the phenomenon of statistical independence is not directly expressible in the CF formalism. This is consistent with Shafer and Tversky, who observe that some mathematical properties are not translatable from one belief language to another [3]. However, the fact that a particular characteristic of the world cannot be described in a certain language does not necessarily imply that this characteristic in nonexistent.

The statistical independence phenomenon is an attribute of nature which stands aloof from the CF/Bayesian debate. To clarify this distinction, we may use an analogy from physics. The presence or absence of statistical independence is a unique property of a domain of expertise just as the mass is a unique physical property of a brick. Notwithstanding the mass uniqueness, the weight of the brick varies with different scales (or on different planets). Thus an absolute unique property of nature may be mapped onto different manifestations under different circumstances. Similarly, the manifestation of the independence property may be explicit in some belief languages and vague or even null in others. The crispness of this expression should be construed as a property of the language, not a property of nature.

5.2. Conjecture C2 Revisited

The C2 conjecture suggests that both the CF and the Bayesian languages are special cases of the Dempster-Shafer theory of evidence. Although this premise is indeed correct, this truth is quite different from its popular interpretation. That the Bayesian design is a special case of the Dempster-Shafer model is a trivial corollary that can be found in [18]. Similarly, Gordon and Shortliffe gave a Shaferian belief interpretation to the CF calculus. They then proceeded to conclude that *"The Dempster-Shafer combination rule includes the Bayesian and the CF functions as special cases"* [19, p.273].

In my view, the popular interpretation of this correct argument is depicted in Figure 2a:

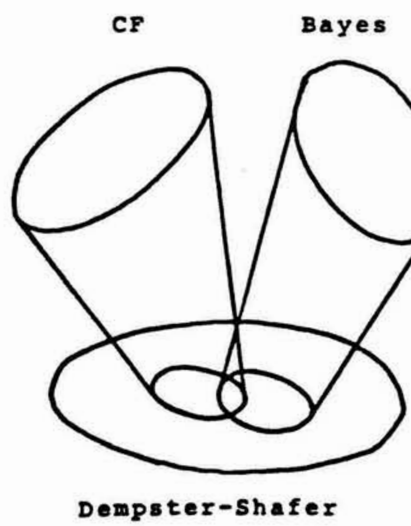

Figure 2a

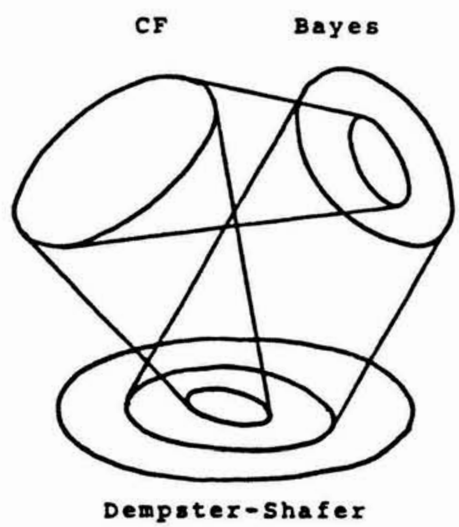

Figure 2b

However, in light of the endomorphism theorem, a more accurate description of the CF/Bayesian/Shafer relationships is as depicted in Figure 2b. The relationships depicted in the latter figure are consistent with Grosof's analyses of the CF/Bayesian/Shafer interplay [13]. Grosof has also provided the explicit transformations under which the CF and the Bayesian languages are special cases of the Dempster-Shafer language.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

1. Apply a rule-based algorithm as though the problem in not wholistic:

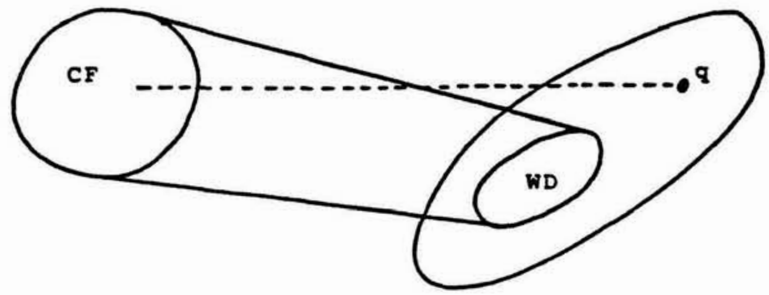

Figure 3a

2. Devise a new rule-based algorithm to wholistic problems:

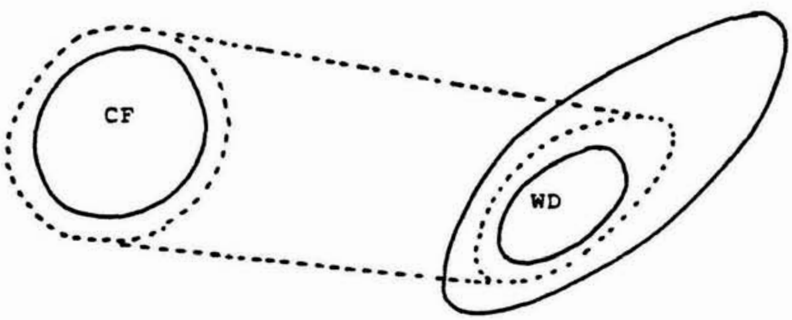

Figure 3b

3. Transform the wholistic problem q into a more complicated problem q' that nonetheless is (roughly) weakly decomposable. Then apply a rule-based algorithm to q':

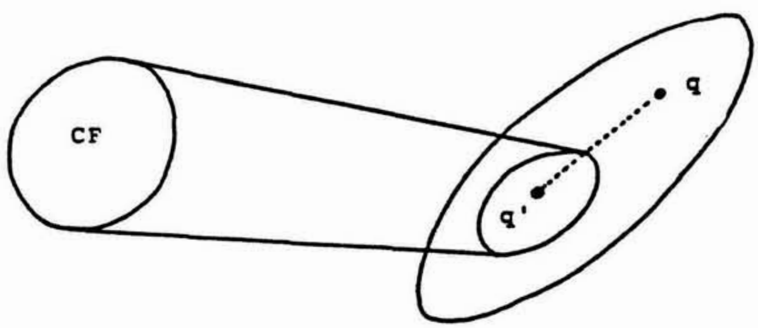

Figure 3c

Option 1, which basically amounts to fudging, is, in my opinion, the leading practice among practitioners who use rule-based expert system shells (this impression is not supported by any firm empirical evidence). Moreover, the endomorphism theorem shows that it doesn't really matter if you use a decomposable Bayesian or a CF approach; both fail to handle wholistic diagnostic structures, although the latter appears to be oblivious of this limitation.

Option 2 presents a very tough challenge. Basically, it requires the development of an

In short, the present discussion indicates that within the subset of weakly decomposable problems (WD), CF's are remarkably Bayesian after all. Outside the subset WD, the CF language may have a variety of free-form and appealing descriptive interpretations. At the same time, those ad-hoc interpretations will not be accountable or testable on rational (Bayesian) grounds. Of course, this restriction may be lifted if either (R1-R2) or (M1-M2) are modified or extended in order to cover a larger superset of WD.

6. IMPLICATIONS ON KNOWLEDGE ENGINEERING AND FUTURE RESEARCH

Several authors have stressed the fact that most real-life domains of expertise include problems that are *dependant, adjacent,* or *wholistic.* Clearly, all those definitions basically imply that the diagnostic structure of realistically complex problems is not weakly decomposable. This observation has far-reaching computational implications which may be summarized in the following proposition, which has not yet been proven:

Proposition: If a problem $<h,e_1,...,e_m>$ is wholistic, then the rule-based computation of the posterior belief $BEL(h|e_1,...,e_m)$ is *NP*-Complete.

Heuristic argument: This proposition is based on the intuition that an m-cities Travelling Salesman Problem (TSP) is reducible to the integer programming (IP) formulation of the rule-based comutation of $BEL(h|e_1,...,e_m)$, provided that such plausible formulation exists. The objective function of the IP model might be based on a logarithmic transformation of $BEL(h|e_1,...,e_m)$ which amounts to a linear combination of all the possible interaction effects within subsets of the $\{h,e_1,...,e_m\}$ space. Each individual interaction effect might be multiplied by an integer $0-1$ variable that determines whether or not the interaction obtains. Furthermore, it is felt that the typical TSP constraint designed to avoid a sub-tour among $k \leq m$ cities can be mapped onto the constraint that given that the underlying inference problem is wholistic, the computation of BEL cannot afford to disregard a dependency of degree k within the $\{h,e_1,...,e_k\}$ space. The objective here is to design these 2^m logical constraints in a way that will force the appropriate $0-1$ variables to be set to 0 or 1, thus determining whether or not the respective interaction effect should enter the computation of $BEL(h|e_1,...,e_m)$ in the objective function.

We now turn to the implications of the endomorphism theorem on knowledge engineering in light of the proposition just presented. In particular, we wish to focus on the key question that ought to be addressed, namely: how can a rule-based expert system compute the posterior belief associated with a wholistic problem. Basically, there seem to be three alternative options:

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

optimal belief updating solution to wholistic problems which is also rule-based, or polynomial in the size of the problem, in some sense. In light of the proposition regarding the computational complexity of wholistic problems, such algorithm will imply that $P = NP$, amounting to the most staggering finding in complexity theory, and a very unlikely one. This leaves us with more realistic lines of attack which are suboptimal, but, nonetheless, feasible. Examples of such efforts are Lemmer's work on incompletely specified distributions [20] and Cheeseman's maximum entropy algorithm [21].

It seems that proponents of the classical rule-based approach to inference are left with the pragmatic option 3, which, in my view, has not received a sufficient amount of research. If we manage to construct a plausible decomposition q' of q, then we can safely apply a rule-based algorithm to q'. Moreover, the goodness of this solution will be a function of the structural proximity of q' to q, which might be estimated by the knowledge engineer. Decomposition might be carried out syntactically, as in Pearl's technique of structuring causal trees [9], or semantically, as in Charniak's notion of 'intermediate states' [15].

If we view the optimal solution of q as a complete combinatorial Bayesian design, and the rule-based solution of q' as a heuristic solution of q, we may bring upon some very strong findings from the probabilistic analysis of the TSP. For example, there is a greedy algorithm that solves the TSP in polynomial time, giving a solution that may not be optimum but is guaranteed to be no worse than twice the optimum path. Furthermore, if some very plausible assumptions are made regarding the layout of the cities (viz, the topology of the diagnostic structure of the problem), this error can become as small as 5% [22].

7. CONCLUSION

The approach taken in this paper was to explicitly define the class of inference problem that are solvable in the CF language in a way which is consistent with the theory of probability. The resulting endomorphism theorem is yet another way to show that the CF language makes strong independence assumptions. In the final analyses, it is obvious that conditional dependencies is a phenomenon that we cannot afford to ignore, regardless of the belief language that we choose to adopt. Moreover, it seems that the conditional independence assumption is structurally inherent in any rule-based algorithm, when applied to probabilistic domains. This is unfortunate, since the rule-based architecture is a well established inference technique with some very appealing characteristics. With that in mind, it is argued that future research should concentrate on manipulating the <u>problem</u> space, rather than the <u>algorithm</u>, in order to make it more amenable to a rule-based solution which will also be valid on probabilistic grounds.

8. Appendix: Proofs

For the sake of brevity, the following proofs are limited to inference problems with two pieces of evidence. Due to the commutative and associative nature of both Bayes rule and (M1-M2), these results can be easily extended to any finite number of pieces of evidence.

8.1. Proof of Lemma 2

Applying (M2) to an inference problem q={h,a,b} amounts to

$$MD(h|a,b) = MD(h|a) + MD(h|b) \cdot (1-MD(h|a)) \qquad (M2)$$

It is easy to show (from the (R2) interpretation of MD) that

$$MD(h|e) = 1 - \frac{P(e|h)}{P(e)} \qquad (\text{if } P(h),P(e)>0)$$

Applying this to the r.h.s. of (M2) gives:

$$MD(h|a,b) = \left(1-\frac{P(a|h)}{P(a)}\right)+\left(1-\frac{P(b|h)}{P(b)}\right)\cdot \frac{P(a|h)}{P(a)} = 1-\frac{P(a|h)}{P(a)}+\frac{P(a|h)}{P(a)}-\frac{P(a|h)\cdot P(b|h)}{P(a)\cdot P(b)}$$

Hence, (M2) is equivalent to

$$MD(h|a,b) = 1-\frac{P(a|h)\cdot P(b|h)}{P(a)\cdot P(b)} \qquad (M2')$$

Alternatively, the rational interpretation (R2) of MD(h|a,b) gives:

$$MD(h|a,b) = \frac{P(h|a,b)-P(h)}{-P(h)} \qquad (R2')$$

To prove the IF direction, we begin with a decomposable {h,a,b}:

$$P(a,b|h) = P(a|h)\cdot P(b|h) \qquad (1)$$

$$P(a,b) = P(a)\cdot P(b) \qquad (2)$$

Now, plugging (1) and (2) in the r.h.s. of (M2') and using Bayes Rule yields:

$$1-\frac{P(a|h)\cdot P(b|h)}{P(a)\cdot P(b)} = 1-\frac{P(a,b|h)}{P(a,b)} = 1-\frac{P(h|a,b)\cdot P(a,b)}{P(h)\cdot P(a,b)} =$$

$$= \frac{P(h)-P(h|a,b)}{P(h)} = \frac{P(h|a,b)-P(h)}{-P(h)}$$

which is identical to (R2'). Alternatively, we could have applied Bayes rule to (R2'), plug (1)-(2) in the result, yielding an expression which is identical to (M2'). Since (M2') and (R2') are consistent with (M2) and (R2), respectively, we have proven that (M2) and (R2) are mutually consistent.

To prove the ONLY IF direction, we assume that (M2') and (R2') are mutually consistent, and equate them:

$$\frac{P(h|a,b)-P(h)}{-P(h)} = 1-\frac{P(a|h)\cdot P(b|h)}{P(a)\cdot P(b)}$$

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

Which, after some algebraic manipulations, gives

$$P(h|a,b) \quad = \quad -\cancel{P(h)} + \frac{P(a|h) \cdot P(b|h) \cdot P(h)}{P(a) \cdot P(b)} + \cancel{P(h)} \qquad (3)$$

Now, if {h,a,b} is subject to probabilistic interpretation, then Bayes rule dictates that:

$$P(h|a,b) \quad = \quad \frac{P(a,b|h) \cdot P(h)}{P(a,b)} \qquad (4)$$

Equating (3) and (4) recovers the implicit assumptions:

$$P(a,b|h) \quad = \quad P(a|h) \cdot P(b|h)$$

$$P(ab) \quad = \quad P(a) \cdot P(b)$$

which imply that {h,a,b} is a decomposable problem.

8.2. Proof of Lemma 1

We begin by applying the fact $MB(h|a) = MD(\neg h|a)$ to (R1), obtaining:

$$MD\,(\neg h|a,b) = MD\,(\neg h,a) + MD\,(\neg h,b) \cdot (1\text{-}MD\,(\neg h,a))$$

Applying Lemma 2 to this completes the proof.

8.3. Proof of Lemma 3

The proof follows from lemmas 1 and 2: The former says that (M1) is equivalent to the Bayesian computation of $P(\neg h|a,b)$ under the assumption that $\{\neg h,a,b\}$ is decomposable. The latter says that (M2) is equivalent to the Bayesian computation of $P(h|a,b)$ under the assumption that {h,a,b} is decomposable. Taken together, these two belief updating operations are equivalent to the odds-ratio form of Bayes rule:

$$\frac{P(h|a,b)}{P(\neg h|a,b)} \quad = \quad \frac{P(a|h)}{P(a|\neg h)} \cdot \frac{P(b|h)}{P(b|\neg h)} \cdot \frac{P(h)}{P(\neg h)}$$

Which is based on the assumption that {h,a,b} and $\{\neg h,a,b\}$ are weakly decomposable.

8.4. Proof of the Endomorphism Theorem:

Let the rational interpretation (R1-R2) be a transformation T, $T:CF \rightarrow V$. Let $WD \subseteq V$ be the subset of weakly decomposable problems in V. Since there exist many wholistic problems in V which are not weakly decomposable, we get $WD \subset V$. Let q be such a wholistic problem with $q \in V$ and $q \notin WD$. According to Lemma 3, $T(CF) = WD$. Hence, we have found a problem $q \in V$ which is outside the range of T. This implies that $T:CF \rightarrow V$ is an endomorphism.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59

REFERENCES

[1] Savage, L. J., The Foundations of Statistics, Wiley, 1954.

[2] Carnap, R., Logical Foundations of Probability, Chicago: University of Chicago Press, 1954.

[3] Shafer G., and Tversky, A., "Languages and Designs for Probability Judgment," Cognitive Science 9, 1985.

[4] Wallsten, T. A., and Budescu, D. V., "Encoding Subjective Probabilities: a Psychological and Psychometric Review," Management Science, Vol. 29, No. 2, Feb. 1983.

[5] Van Melle, W., Shortliffe< E. H., and Buchanan B. G., "EMYCIN - a Knowledge Engineer's Tool for constructing Rule-Based Expert Systems," in: Buchanan, B. G., and Shortliffe, E. H. (Eds.), Rule-Based Expert Systems, Addison-Wesley, 1984, pp. 302-313.

[6] Webster, R., "M.1 Makes a Direct Hit," PC Magazine, April 16, 1985, pp. 145-157.

[7] Personal Consultant Expert Systems Development Tools (User Guide), Texas Instruments, January, 1985.

[8] Duda, R. O., Hart, P. E, and Nilsson, N. J., "Subjective Bayesian Methods for Rule-Based Inference Systems," Proceedings of the 1976 National Computer Conference, AFIPS, Vol 45, 1976, pp. 1075-1082.

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

[10] Shortliffe, E. H., and Buchanan, B. G., "A Model of Inexact Reasoning in Medicine," in: Buchanan, B. G., and Shortliffe, E. H. (Eds.), Rule-Based Expert Systems, Addison-Wesley, 1984, pp. 233-262.

[11] Adams, J. B., "Probabilistic Reasoning and Certainty Factors' in Rule-Based Expert Systems," in: Buchanan, B. G. and Shortliffe, E. H. (Eds.), Rule-Based Expert Systems, Addison-Wesley, 1984, pp. 263-271.

[12] Heckerman, D. E., "Probabilistic Interpretation for MYCIN's Certainty Factors," in Lemmer, J. F., and Kanal, L. (eds.), Uncertainty in Artificial Intelligence. North Holland, 1986.

[13] Grosof, B. N., "Evidential Confirmation as Transposed Probability," in Lemmer, J. F., and Kanal, L. (eds.), Uncertainty in Artificial Intelligence. North Holland, 1986.

[14] Fischhoff, B., "Debiasing," in: Kahneman, D., Slovic, P., and Tversky A. (Eds.), Judgment under Uncertainty, Heuristics and Biases, Cambridge University Press, 1982.

[15] Charniak, E., "The Bayesian Basis of Common Sense Medical Diagnosis," Proceedings of the National Conference in Artificial Intelligence, 1983, 3, pp. 70-73.

[16] Barr, A., and Feigenbaum, E. A. (Eds.), "The Handbook of Artificial Intelligence." Los Altos, CA: William Kaufmann, 1981, p 195.

[17] Winter, C. L., and Girse, R. D., "Evolution and Modification of Probability in Knowledge Bases," The 2nd Conference on Artificial Intelligence Applications, Miami Beach, Florida, 1985.

[18] Shafer, G., A Mathematical Theory of Evidence, Princeton University Press, 1976.

[19] Gordon, J., and Shortliffe, E. H., "The Dempster Shafer Theory of Evidence." in: Buchanan, B. G., and Shortliffe, E. H. (Eds.), Rule-Based Expert Systems, Addison-Wesley, 1984, pp. 272-294.

[20] Lemmer, J. "Generalized Bayesian Updating of Incompletely Specified Distributions," Large Scale Systems 5, Elsvier Science Publishers, 1983.

[21] Cheeseman, P., "A Method of Computing Generalized Bayesian Probability Values for Expert Systems," Proceedings of the 8th International Joint Conference on Artificial Intelligence, Vol. 1, Karlsruhe, West Germany, 1983.

[22] Papadimitriou, C. H., and Steiglitz, K., Combinatorial Optimization: Algorithms and Complexity, Prentice-hall, 1982, pp.406-429.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-87-59