TECHNICAL REPORT
R-51
February 1986

RESEARCH NOTE

On Evidential Reasoning in a Hierarchy of Hypotheses

Judea Pearl
Cognitive Systems Laboratory, Computer Science Department, University of California, Los Angeles, CA 90024, U.S.A.

Recommended by Wolfgang Bibel

1. Introduction

In a recent publication, "A Method of Managing Evidential Reasoning in a Hierarchical Hypothesis Space" [1], Gordon and Shortliffe (G-S) study the application of Dempster–Shafer (D-S) theory to evidential reasoning in a tree-structured hierarchy of hypotheses. They conclude: "Because the D-S approach allows one to attribute belief to subsets, as well as to individual elements of the hypothesis set, we believe that it is similar to the evidence-gathering process observed when human beings reason at various levels of abstraction," and they further state, "we are unaware of another model that suggests how evidence concerning hierarchically-related hypotheses might be combined coherently and consistently to allow inexact reasoning at whatever level of abstraction is appropriate for the evidence that has been gathered." The purpose of this note is to supplement the G-S analysis with a description of how evidential reasoning can be conducted in the same hypothesis space using a Bayesian formalism. The should give the reader an opportunity to compare the two approaches and judge their relative merits on both conceptual and computational grounds.

2. The Domain

We deal with a finite set $H = \{h_1, h_2, \ldots\}$ of hypotheses known to be mutually exclusive and exhaustive. Certain subsets of H have semantic interest, and these form a strict hierarchy, i.e., each subset has a unique parent set that

0004-3702/86/$3.50 © 1986, Elsevier Science Publishers B.V. (North-Holland)

contains it. The subsets can be viewed as nodes in a tree, with H as the root and the individual hypotheses as the leaves, and each intermediate node stands for the disjunction of its immediate successors. Initially, each singleton hypothesis h_i is quantified with a measure of belief $\text{Bel}(h_i)$, reflecting the probability that h_i is true given all previous evidence. The belief in each intermediate-level hypothesis is the sum of the beliefs given to its constituents (by mutual exclusivity). At this point, a new piece of evidence e arrives, which *directly bears* upon one of the subsets, say S, but says nothing about S's constituents. The degree to which the evidence confirms or disconfirms S is provided either by direct assessment of a human expert or, if e is referenced explicitly in some knowledge base, by the activation of a preassessed rule of the form $e \Rightarrow S$, or $e \Rightarrow \neg S$. It is required to calculate the impact of e on the belief of every hypothesis in the hierarchy.

3. Evidence Aggregation

The Bayesian way of calculating the effect of evidence e consists of the following 3-step process:

Step 1. *Estimation.* An expert determines the hypothesis set S upon which the evidence bears directly, and estimates the degree λ_S to which the evidence confirms or disconfirms S. λ_S is the likelihood ratio

$$\lambda_S = \frac{P(e \mid S)}{P(e \mid \neg S)} . \tag{1}$$

Confirmation is expressed by $\lambda_S > 1$; disconfirmation by $\lambda_S < 1$.

Step 2. *Weight distribution.* Each singleton hypothesis $h_i \in S$ obtains the weight $W_i = \lambda_S$ while every hypothesis outside S receives a unity weight, $W_i = 1$.

Step 3. *Belief updating.* The belief in each singleton hypothesis h_i is updated from the initial value of $\text{Bel}(h_i)$ to:

$$\text{Bel}'(h_i) = P(h_i \mid e) = \alpha_S W_i \,\text{Bel}(h_i) , \tag{2}$$

where α_S is a normalizing factor:

$$\alpha_S = \Big[\sum_i W_i \,\text{Bel}'(h_i)\Big]^{-1} . \tag{3}$$

The belief of each intermediate-level hypothesis is computed from the sum of the beliefs of its singleton elements.

This 3-phase process may be conducted recursively, where the updated beliefs calculated for evidence e_k serve as prior beliefs for the next evidence e_{k+1}. The normalization phase can also be *postponed* until several pieces of evidence $e_1, e_2, \ldots, e_n$ exert their impacts on their corresponding hypotheses $S_1, S_2, \ldots, S_n$. The weights are combined multiplicatively via

$$W_i(e_1, \ldots, e_n) = W_i^1 W_i^2 \cdots W_i^n, \tag{4}$$

where

$$W_i^k = \begin{cases} \lambda_{S_k}, & \text{if } h_i \in S_k, \\ 1, & \text{if } h_i \in -S_k. \end{cases} \tag{5}$$

4. Propagation-based Updating

An alternative way of updating beliefs, which avoids the normalization step, would be to start with the impacted node S and propagate the desired revisions up and down the tree by passing messages between neighboring nodes, as depicted in Fig. 1. The message-passing process is defined by the following rules:

Step 1. The impacted hypothesis S, with current belief Bel(S), updates its belief to

$$\text{Bel}'(S) = \alpha_S \lambda_S \, \text{Bel}(S), \tag{6}$$

where

$$\alpha_S = [\lambda_S \, \text{Bel}(S) + 1 - \text{Bel}(S)]^{-1}, \tag{7}$$

and transmits the following messages to its neighbors:

(a) a single message, $m^- = \alpha_S \lambda_S$, to each of its successors, and

(b) a pair of messages to its father: $m_1^+ = \text{Bel}'(S)$, $m_2^+ = \alpha_S$.

Step 2. Any node Y which receives a message m^- from its father, passes m^- to all its successors and modifies its belief by a factor m^-, i.e.,

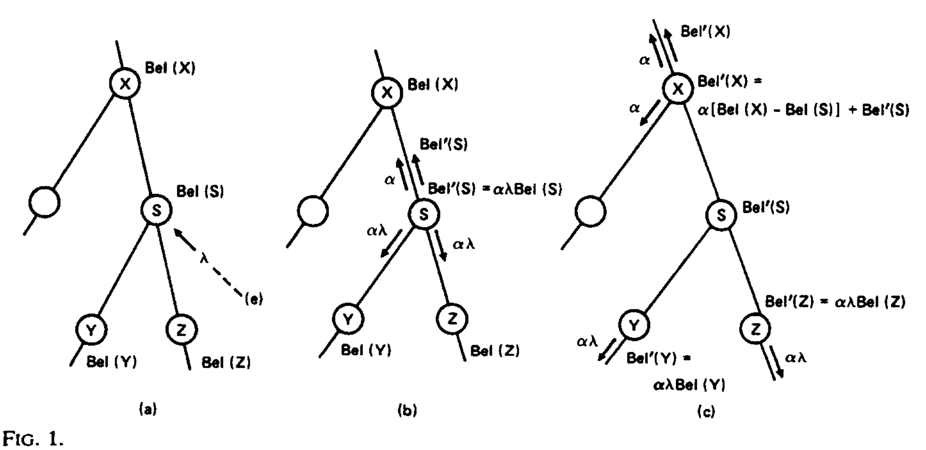

FIG. 1.

$$\text{Bel}'(Y) = m^- \text{Bel}(Y) . \tag{8}$$

Step 3. Any node X which receives the message pair (m_1^+, m_2^+) from one of its successors (say X_1), passes m_2^+ to all its other successors, updates its beliefs to

$$\text{Bel}'(X) = m_2^+ \text{Bel}(X) + m_1^+ - m_2^+ \text{Bel}(X_1) \tag{9}$$

and passes to its father the message pair:

$$(m_1^+)' = \text{Bel}'(X) , \qquad (m_2^+)' = m_2^+ .$$

The quantity $\text{Bel}(X_1)$ in (9) stands for the previous m_2^+ message that X received from X_1.

5. Probabilistic Justification

The operations described in the preceding sections follow from attributing a specific probabilistic interpretation to the statement: "evidence *e bears directly* on S but says nothing about the individual elements of S." We take this statement to mean that the mechanism which gave rise to the observation *e* is a unique property of the subset S, common to all its elements; therefore, contributes no information to make us prefer one element over another. This understanding is captured by the notion of conditional independence:

$$P(e \mid S, h_i) = P(e \mid S) , \quad h_i \in S , \tag{10}$$

stating that once we know that "S is true," the identity of h_i does not make e more or less likely. Simultaneously, we also assume

$$P(e \mid \neg S, h_i) = P(e \mid \neg S) , \quad h_i \in \neg S , \tag{11}$$

stating that the mechanism causing *e* in the absence of S could just as easily be present in each of the hypotheses outside S.

Equations (10) and (11) immediately give:

$$P(e \mid h_i) = \begin{cases} P(e \mid S) , & h_i \in S , \\ P(e \mid \neg S) , & h_i \in \neg S , \end{cases} \tag{12}$$

and, together with Bayes' rule, lead to

$$P(h_i \mid e) = \begin{cases} \alpha_S \lambda_S P(h_i) , & h_i \in S , \\ \alpha_S P(h_i) , & h_i \in \neg S , \end{cases} \tag{13}$$

where λ_S is given by (1) and α_S is a normalizing factor assuring that $P(h_i \mid e)$ sum to 1:

$$\alpha_S = [\lambda_S P(S) + 1 - P(S)]^{-1} . \tag{14}$$

Note that (13) holds not only when h_i is a singleton hypothesis but also for any subset of H, including S itself.

The correctness of equations (1)–(5) follows by identifying $\mathrm{Bel}(h_i)$ with $P(h_i)$, and $\mathrm{Bel}'(h_i)$ with $P(h_i \mid e)$. Note that an identical weight-distribution process, followed by normalization, would still be valid in a general graph hierarchy, not just trees.

The legitimacy of the propagation-based updating follows from the fact that the normalizing constant α_S can be computed directly at the impacted node S, giving (7) (or (14)). Moreover, each descendant of an impacted node S receives the same belief-modifying weight $m^- = \alpha_S\lambda_S$; these weights combine multiplicatively via (4), thus justifying the operation in (8). Similarly, since each node which is not a descendant or an ancestor of S should modify its belief by a constant factor α_S, the message $m_2^+ = \alpha_S$ is passed down to the siblings of S, to the siblings of its father and so on, as in Step 3. Finally, the belief revision appropriate for the ancestors of S is determined by the condition that each node should acquire a belief equal to the sum of the beliefs absorbed by its immediate successors. This gives rise to (9) via

$$\begin{aligned}\mathrm{Bel}'(\text{Siblings of } X_1) &= m_2^+ \, \mathrm{Bel}(\text{Siblings of } X_1) \\ &= m_2^+[\mathrm{Bel}(X) - \mathrm{Bel}(X_1)] .\end{aligned}$$

Although the belief of every singleton hypothesis can be determined without updating the beliefs of S's ancestors, the latter is a necessary calculation because subsequent evidence may directly impact any of these ancestors, and the magnitudes of the emerging messages depend upon the total belief that the impacted hypothesis merits just before the evidence arrives.

6. Conclusions

The updating scheme described exhibits the following characteristics:

(1) *Natural management of beliefs in a hierarchy of hypotheses.* Both evidence and beliefs are combined coherently and consistently at whatever level of abstraction is appropriate for the evidence that has been gathered. To specify the effect of e on the entire knowledge base, the expert need only quantify the relation between e and S (using λ_S), but, otherwise, is not required to apply any conscious effort whatsoever regarding other propositions in the system; that task is fully delegated to the background process of either normalization or propagation.

This feature, together with the ability to postpone the weight-distribution step until appropriate conditions develop, perfectly conforms to the popular metaphor of storing a quantity of uncommitted mass at S, and distributing it to its constituents only upon receiving further refining evidence. In the Bayesian formalism we retain the options of either holding that mass at S (in the form of undistributed λ_S), or distributing it *on a provisional basis* to S's constituents, to be retracted and rerouted later if additional evidence so warrants.

(2) *Clear distinction between (partial) confirmation and disconfirmation.* An evidence in favor of h could not possibly be construed as partially supporting the negation of h. (See G-S for apprehensions regarding this issue.) Confirmation is encoded by $\lambda > 1$, and disconfirmation by $\lambda < 1$. Working with the logarithms of λ would make the distinction even more pronounced by associating confirmation with a positive weight and disconfirmation with a negative weight. The logarithmic representation also provides a closer match to the mass-distribution methaphor, since masses combine additively, not multiplicatively.

(3) *Clearly stated assumptions.* The assumptions behind the updating procedure are stated in familiar meaningful terminologies. The expert can readily judge whether the conditions of equations (10) and (11) are satisfied in any given situation by answering basic, qualitative queries, of the type: "Does X influence Y given that we know Z?"

(4) *Meaningful parameters.* The expert is required to assess only one type of numerical parameter, the likelihood ratio λ_S. The epistemological meaning of this parameter is clearly understood (e.g. how much more likely would it be for e to occur under S as opposed to not-S) and, at least in principle, it can be derived from actual experiential data.

(5) *Transparency of inferences.* One of the main advantages of updating beliefs by message propagation, as opposed to global normalization, is that the former is far more *transparent* in the sense that the intermediate steps can be given intuitively meaningful interpretation. This is so because every computational step in a propagation process only obtains inputs from neighboring, semantically related hypotheses, and because the activation of these steps proceeds along semantically familiar pathways. As a result, it is possible to generate qualitative justifications mechanically by tracing the sequence of operations along the activated pathways and giving them causal or diagnostic interpretations using the appropriate linguistic expressions. This transparency can also be maintained in other hierarchies (e.g., causal models), where the assumption of mutual exclusiveness no longer holds [2, 3].

(6) *Ties to decision procedures.* The probabilistic interpretation underlying the updated beliefs provides a simple framework of converting these beliefs into meaningful decisions reflecting the following cost-benefit considerations: cost-benefit tradeoffs, uncertainties regarding the consequences of actions, and the utility of acquiring additional evidence.

ACKNOWLEDGMENT

I am indebted to Hector Geffner for urging me to expand on several subtle points and for pointing out the applicability of the updating scheme to general graph hierarchies.

Part of this work was supported by NSF Grant DCR 83-13875.

REFERENCES

1. Gordon, J. and Shortliffe, E.H., A method of managing evidential reasoning in a hierarchical hypothesis space, *Artificial Intelligence* **26** (1985) 323–357.
2. Kim, J. and Pearl, J., A computational model for combined causal and diagnostic reasoning in inference systems, in: *Proceedings Eighth International Joint Conference on Artificial Intelligence, Karlsruhe, F.R.G.* (1983) 190–193.
3. Pearl, J., Fusion, propagation, and structuring in Bayesian networks, Tech. Rept. CSD-850022, Cognitive Systems Laboratory, Computer Science Department, University of California, Los Angeles, June 1985.

Received September 1985; revised version received December 1985