

REVEREND BAYES ON INFERENCE ENGINES: A DISTRIBUTED
HIERARCHICAL APPROACH^(*)^(**)

Judea Pearl
Cognitive Systems Laboratory
School of Engineering and Applied Science
University of California, Los Angeles
90024

ABSTRACT

This paper presents generalizations of Bayes likelihood-ratio updating rule which facilitate an asynchronous propagation of the impacts of new beliefs and/or new evidence in hierarchically organized inference structures with multi-hypotheses variables. The computational scheme proposed specifies a set of belief parameters, communication messages and updating rules which guarantee that the diffusion of updated beliefs is accomplished in a single pass and complies with the tenets of Bayes calculus.

Introduction

This paper addresses the issue of efficiently propagating the impact of new evidence and beliefs through a complex network of hierarchically organized inference rules. Such networks find wide applications in expert-systems [1], [2], [3], speech recognition [4], situation assessment [5], the modelling of reading comprehension [6] and judicial reasoning [7].

Many AI researchers have accepted the myth that a respectable computational model of inexact reasoning must distort, modify or ignore at least some principles of probability calculus. Consequently, most AI systems currently employ ad-hoc belief propagation rules which may hinder both the inferential power of these systems and their acceptance by their intended users. The primary purpose of this paper is to examine what computational procedures are dictated by traditional probabilistic doctrines and whether modern requirements of local asynchronous processing render these doctrines obsolete.

We shall assume that beliefs are expressed in probabilistic terms and that the propagation of beliefs is governed by the traditional Bayes transformations on the relation $P(D|H)$, which stands for the judgmental probability of data D (e.g., a combination of symptoms) given the hypothesis H (e.g., the existence of a certain disease). The unique

(*)The paper "An Essay Towards Solving a Problem in the Doctrine of Chances by the late Rev. Mr. Bayes", Phil. Trans. of Royal Soc., 1763, marks the beginning of the science of inductive reasoning.

(**) Supported in part by the National Science Foundation, Grant IST 80 19045.

feature of hierarchical inference systems is that the relation $P(D|H)$ is computable as a cascade of local, more elementary probability relations involving intervening variables. Intervening variables, (e.g., organisms causing a disease) may or may not be directly observable. Their computational role, however, is to provide a conceptual summarization for loosely coupled subsets of observational data so that the computation of $P(H|D)$ can be performed by local processes, each employing a relatively small number of data sources.

The belief maintenance architecture proposed in this paper is based on a distributed asynchronous interaction between cooperating knowledge sources without central supervision similar to that used in the HEARSAY system [4]. We assume that each variable (i.e., a set of hypotheses) is represented by a separate processor which both maintains the parameters of belief for the host variable and manages the communication links to and from the set of neighboring, logically related variables. The communication lines are assumed to be open at all times, i.e., each processor may at any time interrogate its message-board for revisions made by its neighbors, update its own belief parameters and post new messages on its neighbors' boards. In this fashion the impact of new evidence may propagate up and down the network until equilibrium is reached.

The asynchronous nature of this model requires a solution to an instability problem. If a stronger belief in a given hypothesis means a greater expectation for the occurrence of a certain supporting evidence and if, in turn, a greater certainty in the occurrence of that evidence adds further credence to the hypothesis, how can one avoid an infinite updating loop when the two processors begin to communicate with one another? Thus, a second objective of this paper is to present an appropriate set of belief parameters, communication messages and updating rules which guarantee that the diffusion of updated beliefs is accomplished in a single pass and complies with the tenets of Bayes calculus.

A third objective is to demonstrate that proper Bayes inference can be accomplished among multi-valued variables and that, contrary to the claims made by Pednault, Zucker and Muresan [8], this does not render conditional independence incompatible with the assumption of mutual exclusivity and exhaustivity.

