Introduction

With too little information, reasoning and learning systems cannot work effectively. Surprisingly, too much information can also cause the performance of these systems to degrade, in terms of both accuracy and efficiency. It is therefore important to determine what information must be preserved, *i.e.*, what information is "relevant".

This, the AAAI Fall Symposium on "Relevance", was inspired by the recent flurry of interest in this topic in a number of different communities, including

Knowledge Representation: reasoning about irrelevance of distinctions to speed up computation, whether in the context of standard logical derivation, a production system execution, or constraint satisfaction;

Machine Learning: removing irrelevant attributes and/or irrelevant training examples, to make feasible induction from very large databases;

Probabilistic reasoning: simplifying Bayesian nets (both topology and values) to permit real-time reasoning in applications; and

Neural Nets: identifying good topologies by eliminating irrelevant nodes,

to name but a few. Our goal is a better understanding of the various senses of the term "relevance", with a focus on techniques for improving a system's performance by ignoring or de-emphasizing irrelevant and superfluous information. These techniques will clearly be of increasing importance as knowledge bases, and learning systems, become more comprehensive to accommodate real-world applications.

With the goal of articulating and refining a common understanding of relevance, we invited submissions from the various research communities, seeking answers to both theoretical questions

> Are there useful, abstract (formalism-independent) notions of relevance that can be shared among these different subfields?

and practical concerns,

> What are the *real* applications that would benefit from relevance reasoning, especially in the context of exploring and understanding large information repositories?

We were fortunate to receive a wealth of great submissions, including many top-notch papers from excellent researchers in several fields, including knowledge representation, machine learning and learnability, statistics, constraint satisfaction, information retrieval, psychology and philosophy. These submissions represented tremendous diversity, in both computation models (*e.g.*, belief networks, default logics, decision trees), as well as specific applications (*e.g.*, text retrieval, game playing, real-time object recognition, plan recognition, speech processing, software agents, data mining). We accepted a diverse blend of papers — ranging from normative descriptions and theoretical analyses, to strong empirical studies, and even some psychological investigations. (In fact, a handful of the papers presented both intriguing theoretical insights as well as compelling empirical results on non-trivial application domains!)

We also arranged a panel, peopled with major researchers from several different fields, to help set the research agenda for relevance for the next five years.

Synopsis: In general, "relevance reasoning" is an (often implicit) component of an overall system that is performing a "performance task", such as answering questions or learning a function. Without loss of generality, we can view this system as one that receives inputs (percepts) from the world, has some internal state (which could include encoded knowledge about the world and its goals), and which affects the world through its actions (these can be recommendations for action (asking a human to push a button), or concrete actions (pushing a button)).[1] By specifying the system only by this input-state-output description, we are deliberately hiding the details of the system's inner working (for instance, whether it reasons explicitly with propositions, or not).

Using this model, we can perform:

- **relevance analysis of the percepts**. To what aspects of the environment does the system pay attention? Are these the appropriate aspects, for the system's task? For instance, which training examples (or alternatively, which parts of each training instance) does a learning agent use, or in the context of using a set of sub-agents to answer queries, which parts of each given query does each sub-agent attempt to answer? Such an analysis allows us to define the perception policy used by the agent, and characterize its behaviour as a function of the environment and its task.

- **relevance analysis of internal state**. An agent's internal state often corresponds to its model of the world (encoded perhaps as a propositional rule-base, Bayesian net, or a neural net), possibly augmented with aspects of the agent's perceptual history. We can then ask what aspects of its perceptual history should the agent remember? Are these the right choices for its environment and task? For instance, an agent traversing the World-Wide-Web looking for documents in a specific area may benefit by remembering landmark locations (Mosaic home pages) to aid in navigation, much as a physical robot benefits from remembering physical landmarks (the McGraw tower at Cornell).

- **relevance analysis of actions**. Which actions or aspects of actions are superfluous for a given task/environment pair? Can the system be directed to take more appropriate actions? For example, a classification agent that distinguishes among 100 different road conditions for the NavLab system (and emits a label from 1–100) could potentially be simplified into a system that produces one of only two labels (perhaps "slippery" and "dry"), if the underlying effector system has only two control programs for driving under these two conditions.

[1] *Cf.*, Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, Englewood Cliffs, NJ: Prentice Hall, 1995.

Two of the many obvious measures for assessing how well a system's relevance policies work are: (1) accuracy, or how appropriate are the responses generated by the system, and (2) efficiency, for example, how well does the system employ its resources in delivering responses of a given accuracy. Of course, we can consider various combination of these objectives, which correspond to particular trade-offs; and in general, when considering systems that reason under resource bounds (*e.g.*, time deadlines), these two performance measures are quite interdependent.

We can also distinguish between agents that *autonomously derive* appropriate policies for monitoring relevant aspects of the environment, versus agents that use *fixed* policies, hard-wired in by their designers for the agent's specific task and environment. This issue is especially prominent when considering the complexity of making relevance decisions online.

We invite readers to characterize the notion of "relevance" used in the included papers along these various dimensions: identifying whether the system's input, state, and/or output is affected, (*i.e.*, what specific policies are used to isolate the relevant percepts, internal state and/or actions), the purpose motivating these relevance-based modification (accuracy versus efficiency) and whether the system is using a fixed relevance policy versus a dynamic relevance policy.[2]

Discussants and Area Summarizers: The range of submissions show that people from many, very different fields are interested in using something they call "relevance". One of the symposium's goals is to identify what (if anything) these uses have in common, with an eye towards mutually-beneficial exchanges — of everything from useful formalisms and frameworks, to perhaps implementations. A major obstacle to such interactions, of course, is that people from different disciplines use different terms for the same idea, work on different problems, apply different methodologies, etc etc etc.

To interrelate the various topics discussed, and thereby make the symposium more coherent, we found first-rate researchers to provide the "framework" for each of the presented papers, and help tie this paper in with other pertinent ideas — especially those presented in the symposium. Each of these "discussants" was asked to place the discussed paper in a larger context, by familiarizing the audience with related work (including his/her own) and by connecting the discussed work with the other computational paradigms represented at the symposium, addressing questions like

- Why do people in this research area (Knowledge Representation, Learnability, Statistics, ...) care about this *problem*?
- Why should people care about this *answer*?
- Who else *has used* this idea?
- Who else *should have* used this idea?
- How is this related to other ideas within other fields?
- What are the major limitations of the work presented — *e.g.*, are extensions needed before it can be used effectively?

We also enlisted four exceptional speakers to provide summaries of each of the major areas represented: Avrim Blum (Learnability), Judea Pearl (Uncertainty Management), Pat

[2]There are, of course, other secondary dimensions as well, for example, whether the relevance criteria are used to *prune* or *order* some aspect, etc.

Langley (Machine Learning), and Devika Subramanian (Knowledge Representation). Each was asked to first discuss a particular paper, and then to provide a survey of their respective areas that was more comprehensive than the other discussants.

We were overwhelmed with the high quality of the resulting discussions and summaries; and heartily recommend reading them!

Thanks: Many people have contributed generously to the organization and planning of this symposium. We owe a great debt to all members of the program committee for working with heavy reviewing burdens under rigid time constraints: Nick Littlestone, Judea Pearl, David MacAllester and Bart Selman. We would especially like to thank Bart for helping make final decisions about both individual papers as well as the overall schedule. We thank each of the discussants and summarizers for commenting on their discussee's papers — a difficult task made even more challenging by the realization that they knew the authors would know the discussant's (aka critic's) identity! Finally, we would like to thank the AAAI, especially its Fall Symposium Committee, for organizing these symposia in general, and helping us arrange our symposium in particular.

We look forward to stimulating presentations and discussions at the Symposium. See you all in New Orleans this November!

Russ Greiner, Princeton, NJ
Devika Subramanian, Ithaca, NY
29 August 1994