Relevance Symposium

With too little information, an information processing system (*e.g.*, a reasoning or learning program) 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".

The AAAI Fall Symposium on "Relevance" was inspired by the recent flurry of interest in this topic from a range of communities, including (1) 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; (2) Machine Learning: removing irrelevant attributes and/or irrelevant training examples, to make feasible induction from very large databases; (3) Uncertainty Management: simplifying Bayesian nets (both topology and values) to permit real-time reasoning in applications; and (4) Neural Nets: identifying good topologies by eliminating irrelevant nodes. The symposium's goal was to obtain a better understanding of the various senses of the term "relevance", with an eye towards 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.

We were fortunate to receive a wealth of excellent submissions from researchers in several fields, including statistics, constraint satisfaction, information retrieval, cognitive science and philosophy, in addition to the fields mentioned above. These articles were tremendously diverse, in terms of 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). They helped to articulate and refine several important senses of relevance, and to identify *real* applications that can benefit from relevance reasoning.

To interrelate the various topics discussed, and thereby make the symposium more coherent, we enlisted recognized experts to place each presented paper in a larger context, by connecting it with other computational paradigms represented at the symposium and by familiarizing the audience with related work. We found this scheme worked very well in facilitating interchanges between the participants; *e.g.*, at least two joint projects between members of different communities were initiated at the symposium. The symposium also included two poster sessions, as well as a formal panel with experts, from several fields, presenting their thoughts on a "Research Agenda for Relevance Reasoning". The symposium concluded with a lively discussion on the variety of forms of relevance reasoning, and a debate on whether task-independent notions of relevance exist.

The workshop was very successful in bringing together researchers from various disciplines and for sharing perspectives, models and algorithms. The participants collectively felt that they benefited from learning about notions of relevance being pursued in other areas and about task-specific implementations of relevance reasoning. We are currently planning to publish an archival collection of "relevant" papers.