Poster Summaries

Session I

Foued Ameur: *Space-bounded Identification of Halfplanes in the Discrete Grid*

In the learning model of Angluin we address the following probelm: Is it possible to learn if we are allowed to forget the majority of the examples presented in the learning process? We give a positive answer to this question for the simple concept class of halfplanes over the discrete grid.

Avrim Blum & Ravi Kannan: *Finding Relevant Subspaces in Neural Network Learning*

PROBLEM: Neural networks with only a few hidden units are effectively selecting out some low-dimension "relevant subspace". How can we hone in on it?
METHOD: theoretical analysis.

Robin Cohen, Ken Schmidt & Peter van Beek: *Relevance in Plan Recognition for Advice-giving Systems*

The problem we address is to select parts of a plan library to ask a user about, in order to peform plan recognition in an advice-giving setting. Our method is theoretical analysis.

Scott Davis & Stuart Russell: *NP-Completeness of Searches for Smallest Possible Feature Sets*

James Delgrande & Jeff Pelletier: *A Formal to Relevance: Extended Abstract*

Given a formal theory which represents default truths as statements of conditionals ('birds fly' is represented as 'if something is a bird then typically/normally it flies'), how do we represent the notion that some property is or is not relevant to that conditional? The idea is that properties which make us retract the application of the conditional to an object which satisfies the antecedent (i.e., properties that a bird might have that make us wish to say it probably doesn't fly) are *relevant* to the original conditional. But we are interested in the *logic* of the notion of relevance in this setting: what complex formulas are relevant, given that some part or parts of that complex are relevant?

Aditi Dhagat & Lisa Hellerstein: *Notes on Learning with Irrelevant Attributes in the PAC Model*

In many learning situations, the learner is presented with data that includes many attributes, but the function being learned depends on only a small number of those attributes. We consider this situation from a theoretical perspective in Valiant's PAC model of learning, and develop learning algorithms (for particular classes of functions) that output hypotheses depending on only a small number of attributes.

Marek Druzdzel & Henri Suermondt: *Relevance in Probabilistic Models: Backyards in a Small World*

The paper addresses the problem of relevance in probabilistic models. It reviews various methods that can be used to identify those variables in a model that are needed for reasoning about the impact of some evidence variables on a set of target variables. It provides a comprehensive review of the available methods, developed either by the authors or others in the context of computation and explanation in Bayesian belief networks.

Renee Elio & Jeff Pelletier: *On Relevance in Non-monotonic Reasoning: Some Empirical Studies*
Renee Elio and Jeff Pelletier are concerned with identifying factors that influence people's default conclusions on simple default reasoning problems — the same problems, in fact, that are used as benchmark problems for AI default reasoning systems. Their position is that how people do default or non-monotonic reasoning defines the phenomenon that AI aims to mimic, so we need actual data, rather than intuitions, on what are appropriate default conclusions in different scenarios. They report the results of pychological studies on default reasoning that concern the relevance of factors such as object similarity and whether the objects in question come from naturally-occurring or artifical classes.

Rusins Freivalds, Efim Kinber & Carl Smith: *Quantifying the Amount of Relevant Information*
In order to accurately access space utilization of learning algorithms in a true complexity theoretic style a model of learning must satisfy the following three properties:
1. space utilization must be measured as a function of the amount of input received by the algorithm,
2. both a long term and a short term memory are necessary,
3. the long term memory must house the state information of the learning algorithm.
Our paper surveys dozens of attempts to produce a model of space utilization in learning algorithms. We present the first model that satisfies 1., 2. and 3. above. Several fundamental results, including a hierarchy theorem, are presented.

Ken Gemes: *Information Loss Versus Information Degradation*
Not every deductive consequence of a statement S counts as a relevant consequence of S; for instance, while 'p' is a relevant consequence of '$p\&q$', '$p \vee \neg q$' is not. In this paper we show how the class of consequences of an arbitrary wff of any first-order propositional or quantificational language may be partioned into the class of relevant consequences and the class of non-relevant consequences.

Robert Korsan: *The Value of Relevance*
This paper describes an algorithm to price any relevance arrow, or group of relevance arrows, in a given influence diagram. These prices extend sensitivity analysis and allow a analyst to materially simplify an influence diagram in some cases. Only marginal probability distributions need to be assessed for the algorithm.

Shusaku Tsumoto & Hiroshi Tanaka: *Characterization of Relevance & Irrelevance in Empirical Learning Methods Based on Rough Sets & Matroid Theory*
PROBLEM: What kind of algebraic structure characterizes relevance of attributes to classfication via symbolic learning methods?
METHOD: Without domain knowledge, relevant information can be viewed as an independent space of the matroid generated by given training examples and attributes.

Session II

Raj Bhatnagar: *Hypothesizing Relevant Situation Models*

Kathleen Dahlgren: *Relevance Reasoning in Test Retrieval*

This research concerns text retrieval to raw text in response to unrestricted English queries. It consists of a theoretical analysis of the detection of similarity between predications in the query and the text, and a computational implementation of the findings.

Marie desJardins: *The Use of Relevance to Evaluate Learning Biases*

Intelligent agents must filter out irrelevant and marginally relevant aspects of the environment in order to minimize the costs of learning. This must be done in a context-sensitive way (i.e., dependent on the goals and constraints of the learning problem). This paper presents a method for evaluating learning biases by formally analyzing the tradeoff between expected accuracy and complexity of alternative biases. Background knowledge about relevance is used to computer the expected accuracy.

Susan Epstein: *Identifying the Right Reasons: Learning to Filter decision Makers*

Given a set of heuristic, decision-making agents for use in a domain of related problem classes, this research addresses the identification of those agents that prove to be irrelevant, self-contradictory, or less trustworthy than others for a particular class of problems. A perceptron-like model is used to learn weights for the agents.

Mark Fasciano: *Task Selection in a Dynamic World*

Relevance is the measure of appropriateness of the many currently active tasks of a planner at work in a complex, dynamic world; a good relevance function provides a way to choose between the short and long term tasks the planner faces. MAYOR is a planning project which provides an architecture for exploring the arbitration between the many tasks that are active at any moment, according to such measures of relevance.

Leona Fass: *Learning from Relevant & Irrelevant Information*

Our research addresses problems of learnability and the selection and conveyance of "relevant" information that can lead a system to learn. Our approach to such problems involves a theoretical analysis, emphasizing techniques of inductive inference and the interactions of logic, language, and computation.

Igor Jurisica: *How to Retrieve Relevant Information?*

PROBLEM: Efficient and flexible locating of relevant information in information repositories.
METHOD: Similarity-based retrieval algorithm with flexible context specification, where context is a basis for relevancy measure-theoretical analysis.

Angela Kennedy: *Understanding Relevance Vis-a-Vis Internal Transfer*

Learning methods that can select and retrieve past experiences that are most relevant to the current goal are needed to overcome the expensive search encountered in problem solvers operating on complex problems in complicated domains. Here, we define a relevant experience as one that can inform the problem solver about which operators will lead it more directly to its current goal, thus, increasing its efficiency. In hopes of discovering an effective relevance metric, we ran a verbal protocol experiment to study subjects solving problems in the physical sciences via analogical reasoning, and the results of the study will be presented in the poster.

Don Kraft & Carol Barry: *Relevance in Textual Retreival*
The problem we address is the determination of the nature of what is relevant in terms of a collection of textual (natural language) records, including factors of topicality (subject content) and non-topicality (e.g., recency, readability, validity, accessibility). Moreover, we look at retrieval mechanisms, especially those based on relevance feedback, to modify queries, using such techniques as genetic algorithms, in order to improve the ability to provide users with relevant information.

Pat Langley & Stephanie Sage: *Pruning Irrelevant Features from Oblivious Decision Trees*
PROBLEM: We are concerned with developing induction methods that scale well to domains that involve many irrelevant features.
METHOD: We are carrying out experimental studies that compare the behavior of different induction algorithms in both natural and synthetic domains.

Eric Neufeld: *Discarding Irrelevant Parameters in hidden Markov model based part-of-speech taggers*
Knowledge is more relevant if it is more specific, and knowledge is more relevant if it is more precise. But it may be that one kind of relevance can only be had at the expense of the other.

Daniel Oblinger & Gerald DeJong: *Dynamic-Bias Induction*
We employ both theoretical and empirical analysis to the problem of inductive learning biased with a specific form of relevance knowledge.

Guiyou Qiu: *Nonmonotonic Logic for Analogical Reasoning*
In this paper, we provide a nonmonotonic logic for analogical reasoning and a graphical representation method to facilitate the calculation of analogies.

Jeff Rickel & Bruce Porter: *Automated Modeling for Answering Prediction Questions: Selecting Relevant Influences*
Our research addresses the task of automatically constructing a model of a complex physical system in order to answer a given prediction question. Our approach is empirical; we have developed criteria and methods for performing this task, implemented them in a program, and tested the program in the domain of plant physiology.

William Stubblefield & George Luger: *Improving Source Selection in Analogical Reasoning An Interactionist Approach*
Our work addresses the problem of selecting a relevent source for analogical reasoning. We have proposed a model of source selection and tested it in a computer implementation.