Relevance Symposium

Russell Greiner
Siemens

Devika Subramanian
Cornell

Programme Committee

Nick Littlestone	NEC
David McAllester	MIT
Judea Pearl	UCLA
Bart Selman	AT&T

Answering a Query

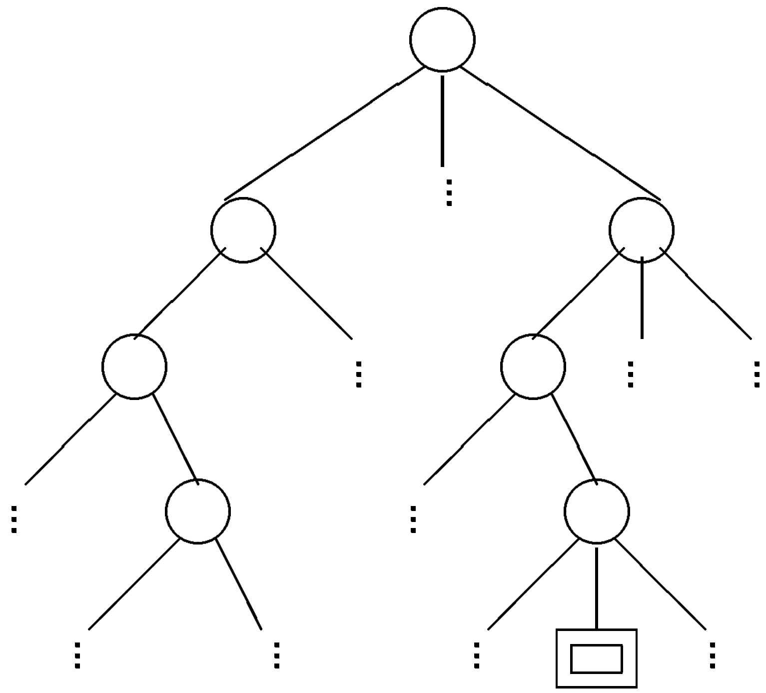

Learning a Classifier

$$
\begin{array}{lcccccccccc}
I_1: & \langle & + & + & + & + & - & + & - \cdots & \rangle & + \\
I_2: & \langle & + & - & + & - & - & + & + \cdots & \rangle & - \\
I_3: & \langle & - & - & - & - & + & - & - \cdots & \rangle & - \\
I_4: & \langle & - & - & - & - & - & + & - \cdots & \rangle & + \\
I_5: & \langle & + & - & + & + & - & + & - \cdots & \rangle & - \\
 & & & & & \vdots & & & & & \\
 & & & & & \vdots & & & & & \\
I_{f(n)}: & \langle & - & - & + & + & - & + & - \cdots & \rangle & -
\end{array}
$$

$$f(n) = O(\ldots n \ldots)$$

$$
\begin{array}{lcccccccccc}
I_1: & \langle & + & & + & + & & - & & - & \rangle & + \\
I_2: & \langle & - & & - & & - & & + & \rangle & - \\
I_3: & \langle & - & & - & & + & & - & \rangle & - \\
I_4: & \langle & - & & - & & - & & - & \rangle & + \\
I_5: & \langle & - & & + & & - & & - & \rangle & - \\
 & & & & & \vdots & & & & & \\
I_{g(n,r)}: & \langle & - & + & + & & - & + & - & \rangle & -
\end{array}
$$

$$g(n,r) = O(\ldots r \ln(n) \ldots)$$

Outline

- Introduction
  - Illustrative Examples of *Relevance*
  - Symposium Goals
- Framework (Types of Relevance)
- Symposium Organization & Logistics
- Questions

Idea of "Relevance"

- Systems require some information to succeed.

  But too much information is also bad...

Improve performance by *ignoring* some *information*

Task: Find irrelevant information, & ignore it.

but... Many types of *information*

∗ Background knowledge; Perceptual History; Input; ...

Many ways of *ignoring* information

∗ Remove it permanently; Remove it temporarily; De-emphasize it; ...

Many different *performance improvement*

∗ Improved accuracy; Increased efficiency; ...

Flurry of Recent Interest

Communities:
- Cognitive science
- Constraint satisfaction
- Control theory
- Information retrieval
- Knowledge representation:
  Default logic, Planning, Reasoning, ...
- Learnability: PAC, Inductive Inference, ...
- Machine learning
- Operations Research
- Philosophy
- Statistics
- Uncertainty Management

Methodologies:
- Theoretical analyses
- Empirical studies
- Psychological investigations

Applications:
- Data mining
- Game playing
- Object recognition
- Plan recognition
- Software agents
- Speech processing
- Text retrieval

Goals of Symposium

- Identify/Articulate common understanding of *RELEVANCE*

  ...towards exchanging useful formalisms, frameworks, implementations...

- Answers to ...
  - Are there useful, abstract notions of relevance that can be shared among different subfields?
  - What *real* applications can benefit from relevance reasoning?

Need common framework, vocabulary, task-descriptions, ...

Framework

- "Relevance" itself is *NOT* performance objective

$$\text{Relevance does NOT} \left\{ \begin{array}{l} \text{solve physics problem} \\ \text{acquire diagnosis fn of disease} \\ \text{operate steel factory} \end{array} \right\}$$

But, **Relevance Reasoning** can help

produce solution more efficiently
require fewer training instances
decide on fewer controls

So... "Performance System" using Relevance Module.

- Model of *Performance System*

Input-State-Output Agent
performing Performance Task
(eg, answering questions, or learning a fn)

...uses *Relevance Reasoning* Module

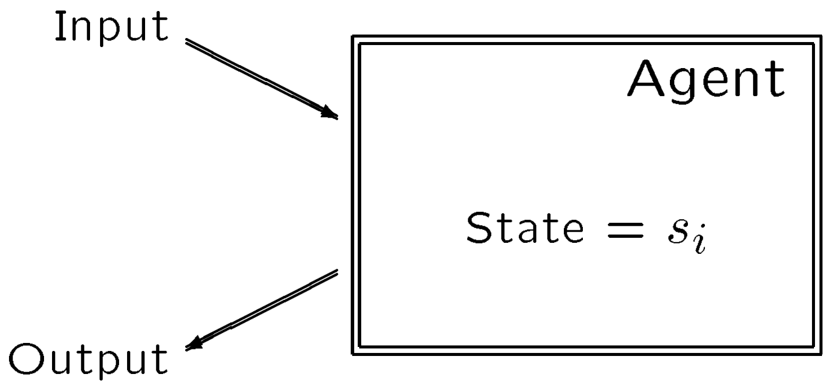

- Agent has internal "State"

  (...encoding knowledge of world, prior experiences, ...)

Agent takes input from world

(...queries, assertions, senses, instances, ...)

- ...changes state

  (...revised knowledge of world...)

- ...produces output

  (... "Yes", push-button, or ...)

Examples of I-S-O Agents

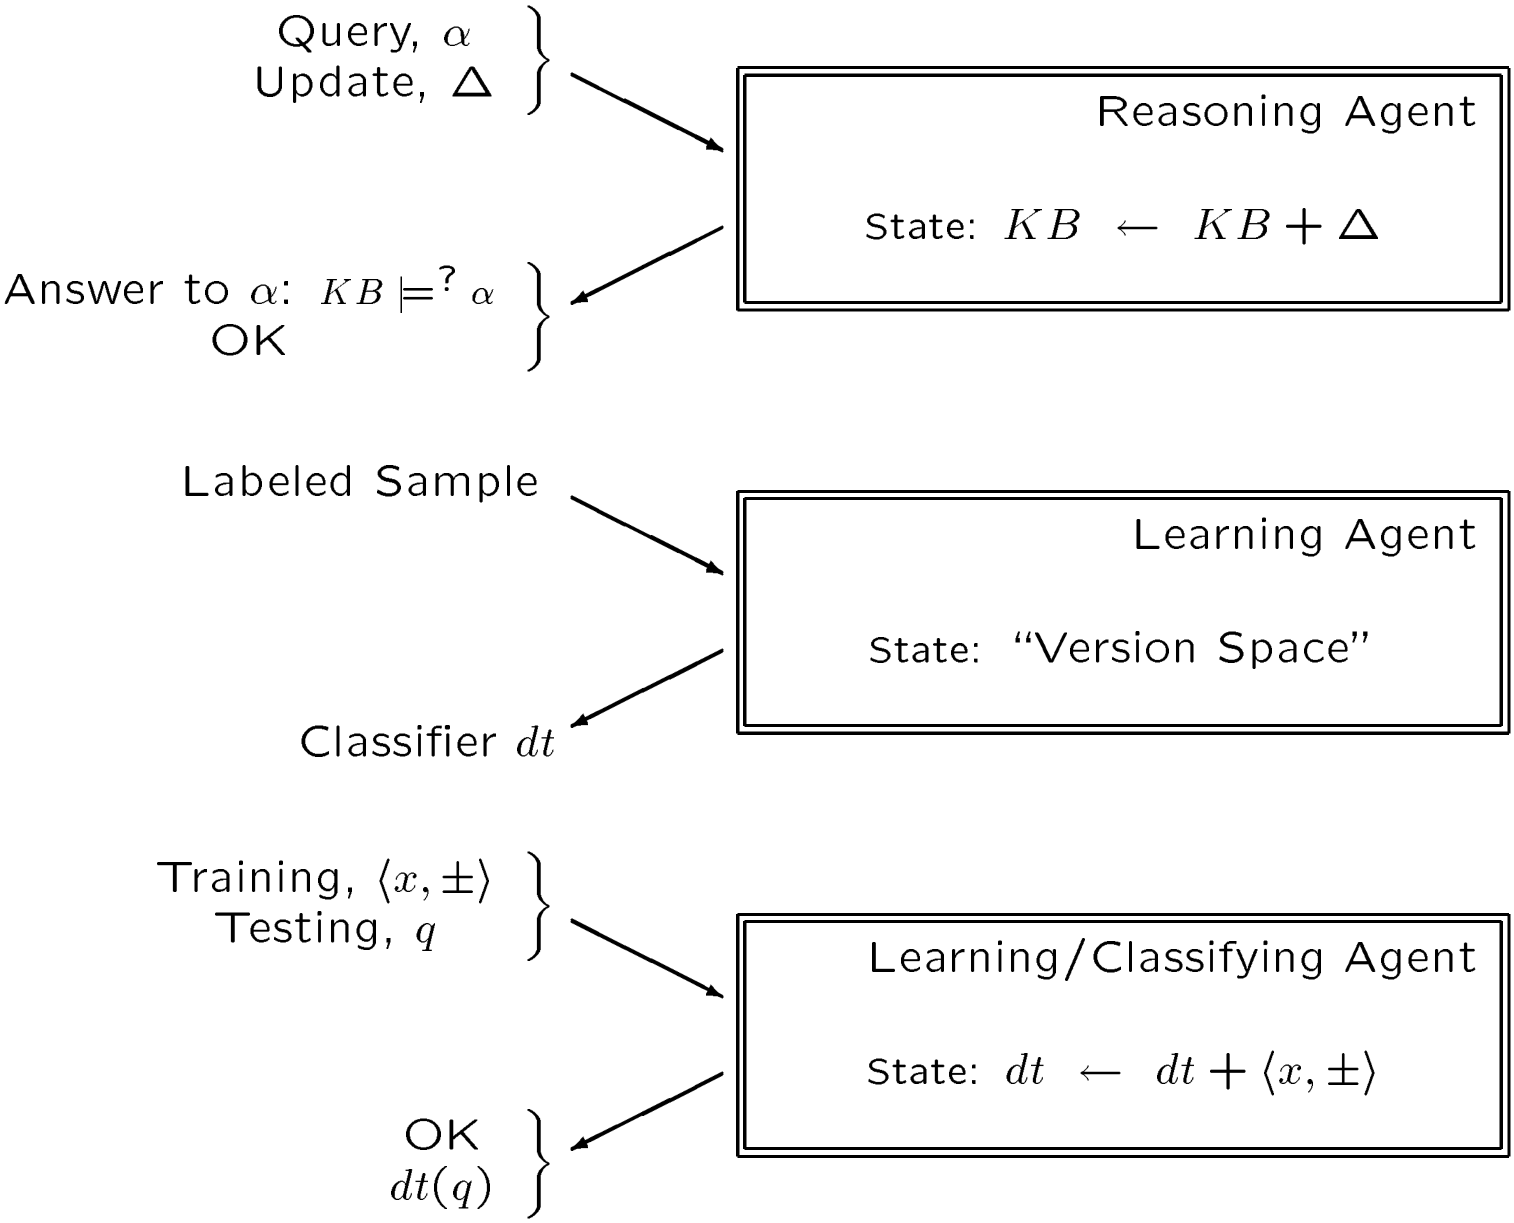

Relevance wrt Input

- Agent focuses on "Relevant" part of input (environment)

Eg1: Learning Agent:
- Subset of training instances
- Certain parts of each training instance
- ...

Eg2: Reasoning Agent:
- Portion of query, update
- ...

Eg3: Object-Recognizing Agent:
- Most "useful" part of feature-space
- ...

- Issues:
  - Are these parts appropriate for task?
  - Accuracy/Efficiency Tradeoff?
  - Ordering vs. Pruning

Relevance wrt State

- Agent focuses on "Relevant" portion of its "state"

- State $=^?$
  - World-model [Rule-Base, Bayesian net, ...]
  - ...rep'n of (subset of) Perceptual History
  - Neural net
  - ...

Eg during derivation,
some rules lead to known facts (& answer)
while others lead to dead-end
Better to IGNORE dead-end rules

- Issues:
  - Are these parts appropriate for task?
  - Accuracy/Efficiency Tradeoff?
  - Ordering vs. Pruning

Relevance wrt Output

- Agent produces only “Relevant” output

  ...ignore distinctions that don’t matter

Eg, # of bits in numerical computations
Only “Left” vs “Right”, rather than 360 angles
Only certain locations (not everywhere)

- Issues:
  - Determining which doesn’t matter
  - ...

Comments

Goal: Seek Commonalities:

Not KR-relevance vs ML-relevance, but
different types of Input-Relevance...

- Evaluation: "Did Relevance help?"
  wrt Distribution of Tasks
  Accuracy, Efficiency, combinations, ...

- Does **agent** reason explicitly re Relevance
  vs Is relevance implicitly encoded in agent?

- I/S/O Trichotomy is Fuzzy
  If State has no term for feature f,
  then no point in receiving $f(\cdot)$ as input.
  (f = color, or f = smell...)

Outline

◦ Introduction

◦ Framework (Types of Relevance)

• Symposium Organization

  – Discussants + Summarizers

  – Panels

  – Posters (+ Introduction)

  – "Definition List"

◦ Questions

Role of Discussants

- Many fields interested in RELEVANCE

Goal: finding commonalities in different senses of RELEVANCE

but... different disciplines

- use different terms for the same idea,
- work on different problems,
- apply different methodologies,
- ...

$\Rightarrow$ Cross-talk, not communication...

So... "Discussant" to provide framework for paper

Help people in other fields $\left\{ \begin{array}{c} \text{understand} \\ \text{appreciate} \end{array} \right\}$ context & contributions

Discussant Issues

- Why do people (in this research area) 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 it related to ideas in other fields?
- What are the major limitations of the work presented?

  ... Are extensions needed before it can be used effectively? ...

\+ Four "Area Summarizers"

Avrim Blum	Learnability
Pat Langley	Machine Learning
Judea Pearl	Uncertainty Management
Devika Subramanian	Knowledge Representation

Panels

1. “Research Agenda for Relevance Reasoning”

Don Kraft	Information Retrieval
Manfred Warmuth	Learnability
Pat Langley	Machine Learning
Peter Hammer	Operations Research
Adnan Darwiche	Uncertainty Management

Questions

- What real-world applications can benefit from the use of relevance reasoning? ...how ...?
- Which of these possible systems have been implemented? ...which have been deployed?
- Which have not? Why?
  (Ie, what are the important problems that have not yet been solved?)
- What are appropriate steps to take, now?
- What is the appropriate agenda for next 5 years? ...10 years?

2. “Tying it all together” Open Forum
Russ Greiner, Devika Subramanian,
Bart Selman, Nick Littlestone, Judea Pearl, ...

Other Topics

- Posters:
  - Embarassment of riches
  - Excellent material

    (just did not fit in as tightly)

  Short synopsis to help focus attention

  "2-sentence statement of PROBLEM your research is addressing, & method you use to address it"

- "Definition List"

  To encourage people to focus on various senses of relevance,

  seeking commonalities and distinctions

Task: Connect DEFINITION with Presentation

Outline

- ◦ Introduction
- ◦ Framework (Types of Relevance)
- ◦ Symposium Organization
- • Questions