Intelligence in Context

Douglas H. Fisher
Vanderbilt University

- Intelligence is subordinate to … ?
- Synthesis/abstraction and compartmentalization/specialization
- Classroom: teaching the context too
- Research: synthesis and metaphor
- A proposal for placing intelligent system research in context
- Final thoughts: balanced communities

Compartmentalization and synthesis in research

Where is the "basic level" for most researchers?

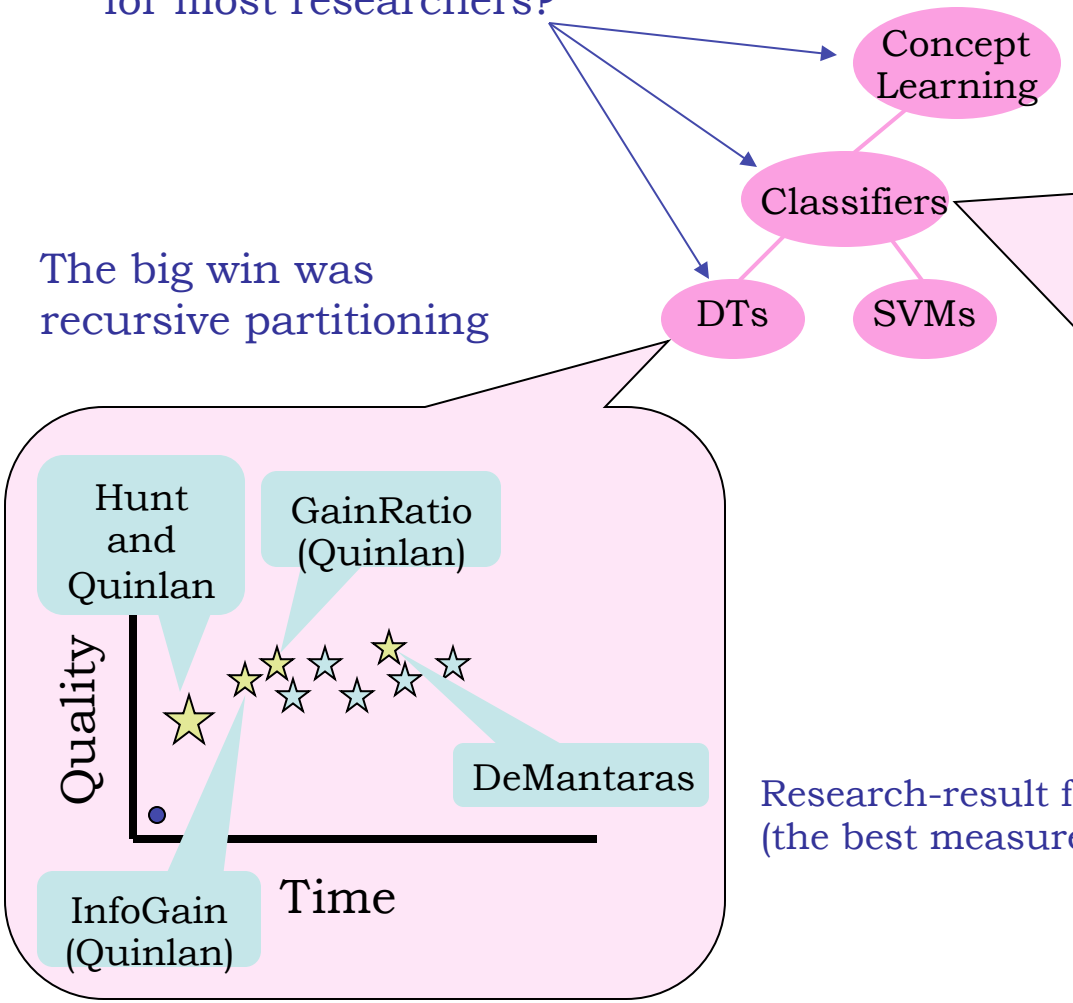

The big win was recursive partitioning

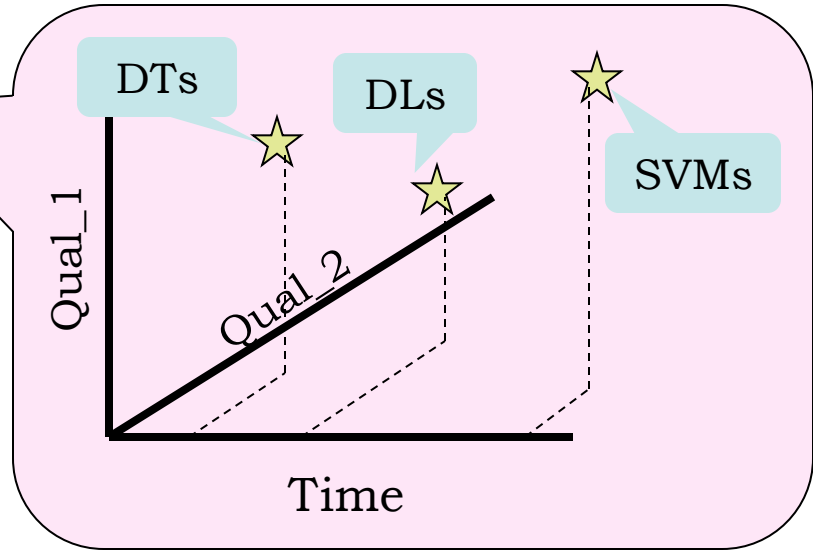

Research-result fan effect
(the best measure may get lost)

Richer Reward Systems

- Counting Publications
- Counting Citations
- Counting data sets defined, collected, released, and *used*
- Counting educational materials developed, released, and *used* *(next)*
- Tutorials, surveys, synthesis

A many dimensional community should have many rewards

Research and academic communities

- UCI
- ICMLs and IWMLs
- AIStats
- IDAs
- AAAIs and IJCAIs
- MLJ
- JMLR

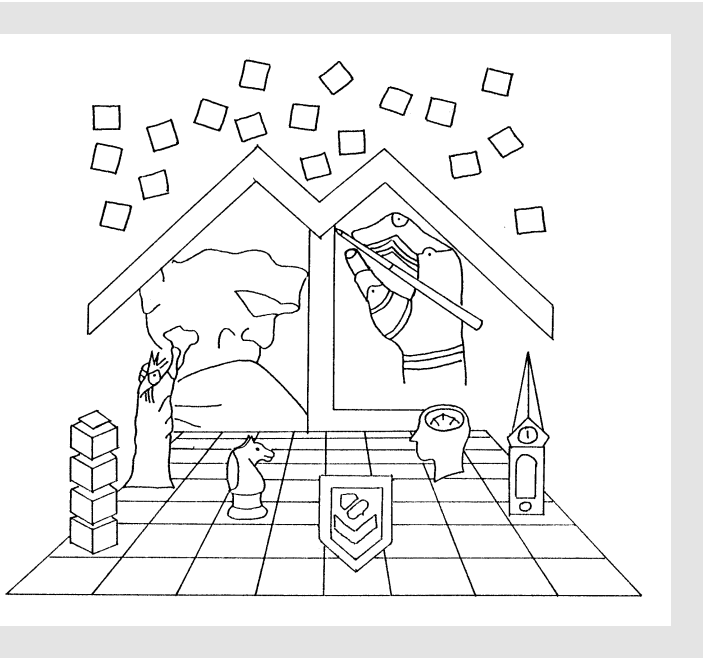

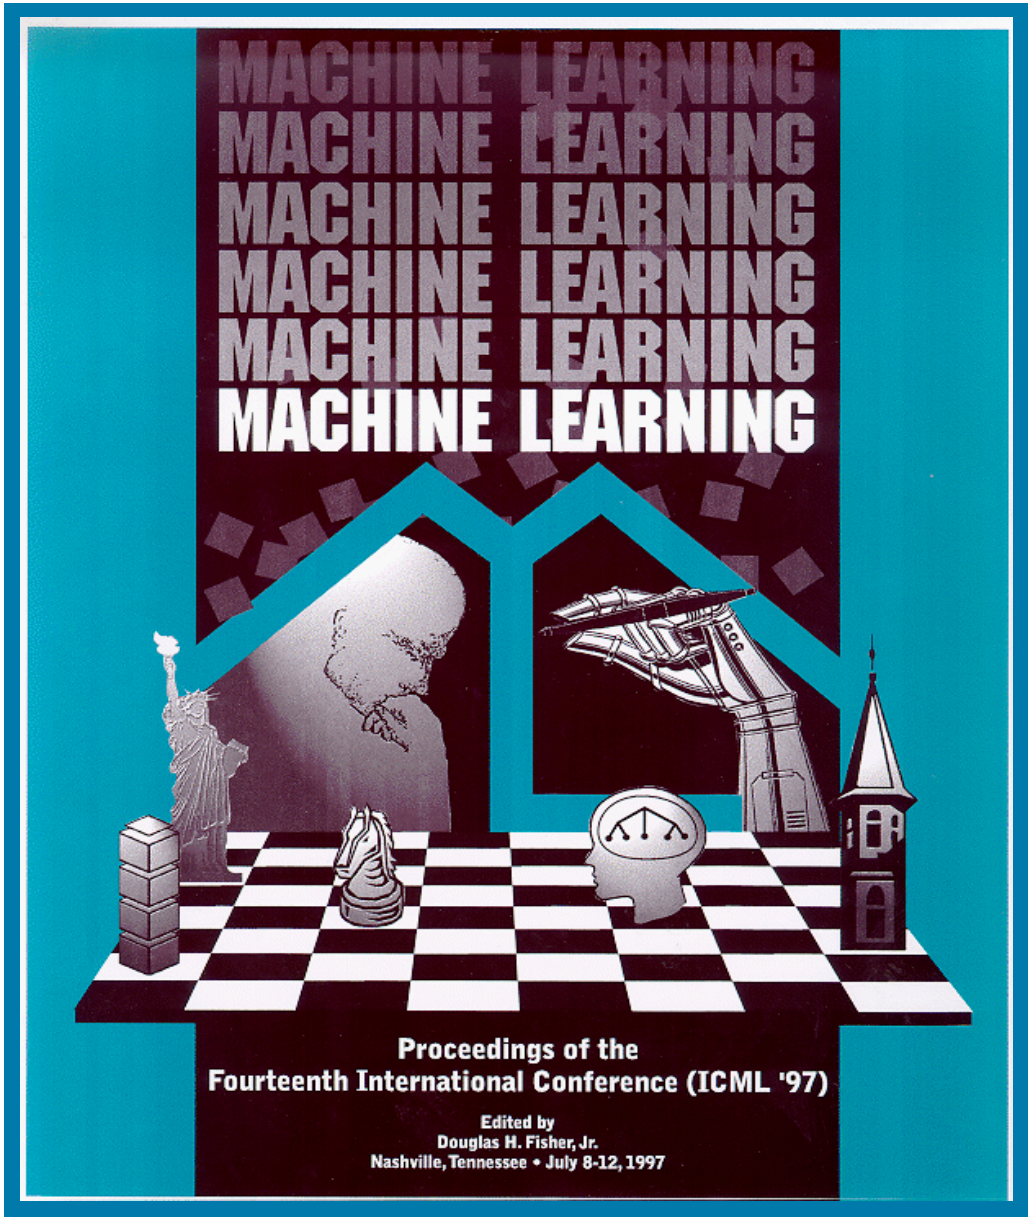

Classroom: teaching the context too

- Compartmentalization in the curriculum
- Informal analysis of textbooks
- ABET accreditation
- IEEE Code of Conduct
- NSF Ethics Education in Science and Engineering
- Two models for embedding ethics and contemporary issues

Embedding ethics and contemporary issues

- Synthesis at curriculum level with separate classes covering ethics and contemporary issues

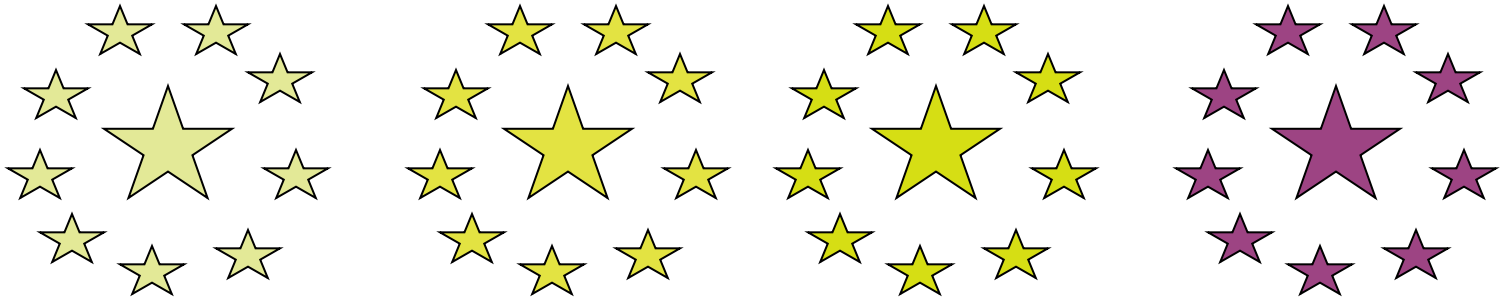

- Materials exist to support this level of synthesis, but they are incomplete

Embedding ethics and contemporary issues

- Synthesis at course level
- Materials rare to support this level of synthesis

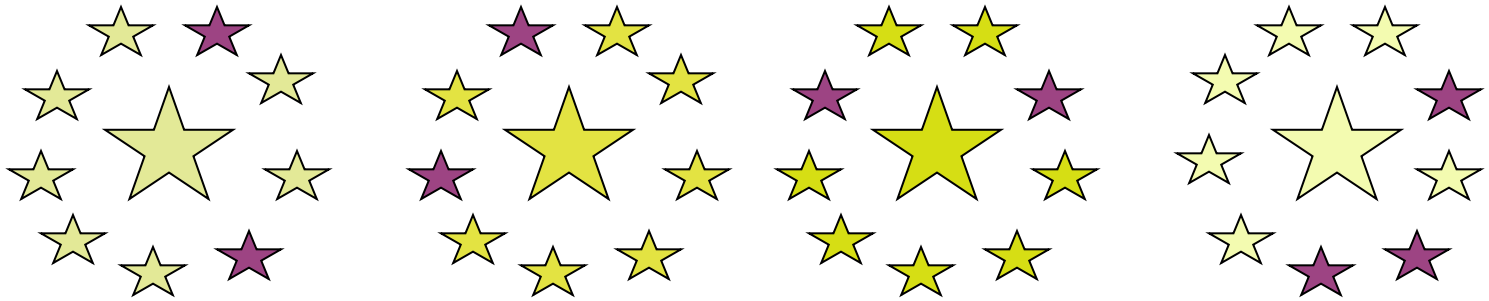

- Ultimately, we want synthesis in the mind of the student.
- Which organization best supports familiarity with possible ethical and societal consequences and opportunities?
- Appeal to whole person (e.g., data security, privacy, identity theft)

Evaluate the pedagogical models and if appropriate encourage and reward the development and deployment of materials that integrate ethics and contemporary issues *at the course level* for information science courses…

Research: synthesis and metaphor

- Cobweb (Fisher, 1987): a diverse ancestry
- Basic level, fan and typicality effects
  (Silber & Fisher, 1989; Fisher & Langley, 1990)
- Supervised and unsupervised learning
  (Fisher, 2001, 1996; 1987; Frey, Fisher, Aliferis, … 2003)
- Exor: learning concepts and learning to problem solve.
  (Yoo & Fisher, 1991; Fisher & Yoo, 1993)
- Induction and prior knowledge
  (Evans and Fisher, 2001; 1994; Fisher, Edgerton, et al 2003; 2006)

A proposal for placing intelligent system research in context:
A taxonomy of climate change issues and tasks

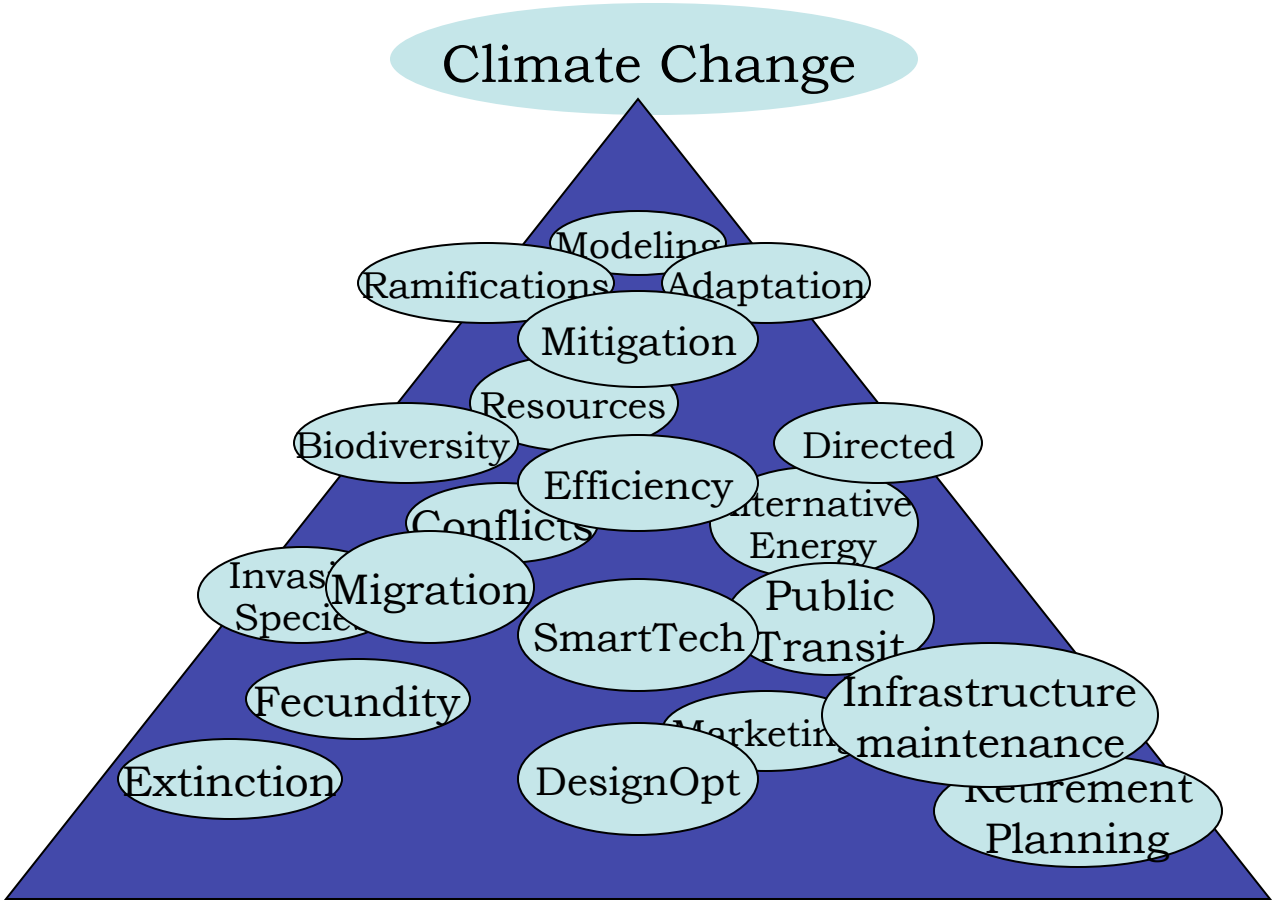

A survey of Climate Change (CC) and AI work

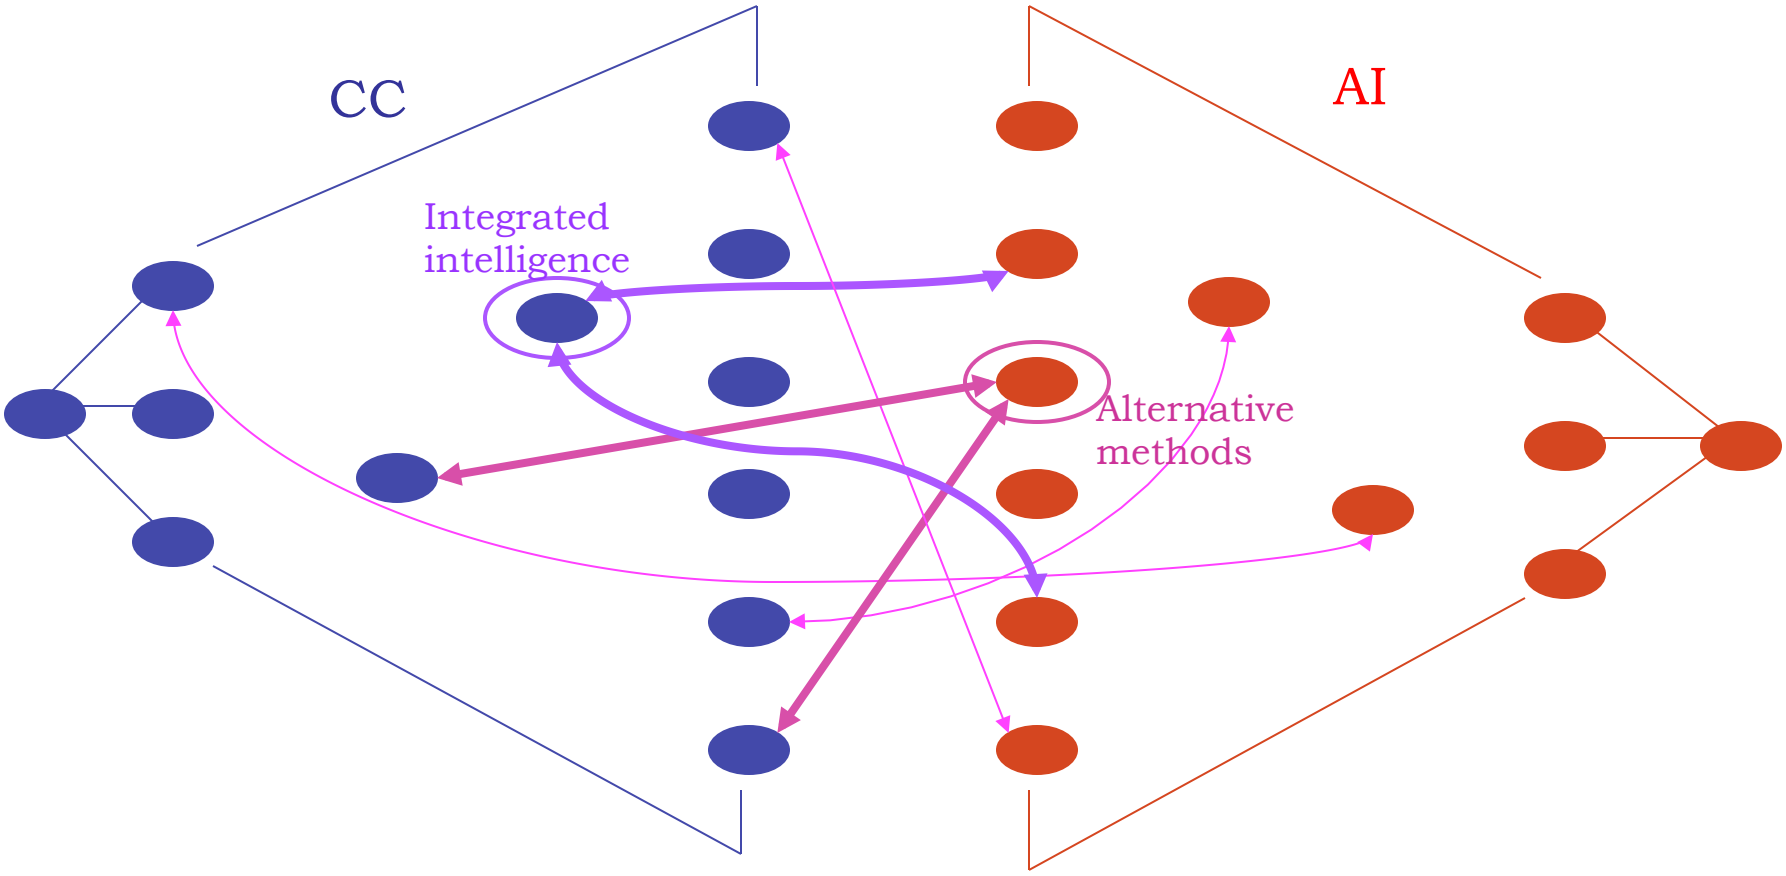

Loosely-coupled collaboration

Examples

CC/**Modeling**/…/ModelAdjustment/
"Here we assess the range of warming rates over the coming 50 years that are consistent with the observed near-surface temperature record as well as with the overall patterns of response predicted by several general circulation models."
Quantifying the uncertainty in forecasts of anthropogenic climate change
Allen, M. R., Stott, P. A., Mitchell, J. F. B., Schnur, R. & Delworth, T. L.
Nature **407**, 617–620 (2000).

"Regression analysis is used to estimate the scaling factor *a* that produces the best match between observations and the simulated climate-change signal."
Weaver, A. J. & Zwiers, F. W., Nature **407**, 571-572 (5 October 2000)

Suggestive of the use of other methods of combining models and/or experts:

- Each model's prediction becomes a feature that augments the data and can be used by inductive learning (e.g., SVMs, regression trees, ANNs) (Cox, 1999, Ortega & Fisher, 1995)

- Can be used for regional modeling

Examples

CC/**Modeling**/…/ModelAdjustment/
"One or more experts are used to define a Bayesian prior distribution to each of the selected attributes, and the interattribute links, of the system under study. Posterior probabilities are calculated interactively, indicating consistency of the assessment and allowing iterative analysis of the system. Illustration is given by 2 impact studies of surface waters. In addition to climatic change studies, the approach has been designed to be applicable to conventional EIA. Insufficient attention has thus far been devoted to the probabilistic nature of the assessment and potential inconsistencies in expert judgment."
BENE-EIA: A BAYESIAN APPROACH TO EXPERT JUDGMENT ELICITATION WITH CASE STUDIES ON CLIMATE CHANGE IMPACTS ON SURFACEWATERS,
VARIS, O. & KUIKKA, S. *Climatic Change* **37:** 539–563, 1997.

Additionally, AI search-based methods might be profitably applied in many circumstances associated with high uncertainty, looking for conditional outcomes with less conditional uncertainty.

Examples

CC/**Ramifications**/Biodiversity/…
"Here we forecast the potential distribution of zebra mussels in the United States by applying a machine-learning algorithm for nonparametric prediction of species distributions (genetic algorithm for rule-set production, or GARP) to data about the current distribution of zebra mussels in the United States and 11 environmental and geological covariates. Our results suggest that much of the American West will be uninhabitable for zebra mussels." *(p. 931). The Potential Distribution of Zebra Mussels in the United States*, DRAKE, J. M. & BOSSENBROEK, J. M. BioScience *Vol. 54 No. 10* 931-941

"Unification of predictive analyses across these two phenomena (invasions and climate change) is completely feasible, yielding predictions of opportunities for invasions in the face of global climate change. Integrating projections of invasions with other scenarios of change, such as human-induced changes in land use and land cover, is equally feasible. A limitation of these explorations, however, is the lack of appropriate baseline data sets to permit quantitative statistical validation of predictivity across multiple scenarios of change." (p. 429). *PREDICTING THE GEOGRAPHY OF SPECIES' INVASIONS VIA ECOLOGICAL NICHE MODELING Volume 78*, No. 4 December 2003 *THE QUARTERLY REVIEW OF BIOLOGY, 419-433.*

Examples

Some researchers may not be conscious that they are working on climate change problems

Ramifications and adaptation

"A Machine Learning (ML) System known as ROAMS (Ranker for Open-Auto Maintenance Scheduling) was developed to create failure-susceptibility rankings for almost one thousand 13.8kV-27kV energy distribution feeder cables that supply electricity to the boroughs of New York City." (p. 1) "We have a number of theories as to why performance was better during the summer. The first is that many of the input features to our machine learning algorithm were developed by Con Edison with a specific focus on modeling the electric distribution system during heat waves. The second is that distribution system failures may have more deterministic causes during heat waves, as the load and stress contribute directly to cable, joint, and transformer problems, while in the cooler months, failures tend to be more random and difficult to model." (p. 5) *Predicting Electricity Distribution Feeder Failures Using Machine Learning Susceptibility Analysis, Gross, P. et al. AAAI (2006)*

"The domain for our experimental investigation is a popular computer war strategy game called FreeCiv. FreeCiv is a multiple-player game in which a player competes either against several software agents that come with the game or against other human players. Each player controls a civilization that becomes increasingly modern as the game progresses. As the game progresses, each player explores the world, learns more about it, and encounters other players. Each player can make alliances with other players, attack the other players, and defend their own assets from them. In the course of a game (that can take a few hours to play) each player makes a large number of decisions for his civilization ranging from when and where to build cities on the playing field, to what sort of infrastructure to build within the cities and between the civilizations' cities, to how to defend the civilization. FreeCiv provides a highly complex, extremely large, non-deterministic, partially-observable domain in which the agent must operate." *Using Model-Based Reflection to Guide Reinforcement Learning,* Ulam, P., Goel, A., et al

"The need for decomposition in learning problems has been widely recognized. One approach to making learning in large state spaces tractable is to design a knowledge representation composed of small pieces, each of which concerns a more compact state space than the overall problem. Techniques that would be intractable for the problem as a whole can then be applied successfully to each of the learning subproblems induced by the set of components." *Knowledge Organization and Structural Credit Assignment,* Jones, J. & Goel, A. (IJCAI 05 Workshop)

A survey of Climate Change (CC) and AI work

- Survey
- Prototype development
- Publicly available
- Publicly updatable
- Integrate with other taxonomies
- Ascribe utilities for policy making
- Promote balanced community

Education, Research, Press, Public, *engage a whole community*

Research results, data sets, educational material, software, art

Other engineering, science, medicine, political, social

Final Thoughts: balanced communities

- A balanced community of researchers is not (necessarily or even probably) a community of individually-balanced researchers
- What types of scholars should be part of a balanced community?
  - Specialists
  - Synthesizers
  - Educators
  - Communicators
  - Ethicists, Artists, Theologians (e.g., MIT), …

Balanced research communities

- At what level of community do we want balance and participation of the various types?
  - Within individual?
  - Within research group?
  - Within Institution?
- Differences within small, medium, and large grant teams?
- How is community defined at NSF:
  - includes RI, IIS, CISE, NSF, other granting agencies?
- Cross-cutting programs

Auxiliary Slides

Intelligence is subordinate to … ?

- Why ask?
  - To manage complexity
- Descriptive, not prescriptive
- Benefits of thinking about AI in context
  - Lines blur and can be redrawn

Benefits of thinking about AI in context

- Collaborative Intelligence
  - Interactive Induction and data engineering
  - Asimo can dance, but can Asimo dance with a person?
  - Human-centered to Robust Intelligence and vice versa
- Embedded Intelligence
  - Understanding and exploiting domain constraints for specialized intelligence (e.g., monitoring for accidents in a parking garage)
- Embodied Intelligence
  - Embodiment on the Internet, a vehicle, a city
  - Emotion and intelligence are a function of the body; what are characteristics of non-human intelligences and emotions
- Expansive Intelligence
  - Intelligence isn't greedy search

An Informal Analysis of Textbooks

- AI Textbook(s)
  - Don't discuss whether intelligent decision support systems in medicine contribute to sloppier and/or more careful physicians?
  - Don't discuss whether intelligent buildings create environmentally stupider or smarter people?

- Database Textbook(s)
  - Don't discuss data *privacy* (illustrated by identity theft)?
  - Don't discuss ethics of decision support?

- Patterson and Hennessy's "Computer Organization and Design"
  - Computing and networking for the Third World, ecological monitoring, and grassroots news
  - Ethics of premature chip release

ABET Accreditation

Outcomes and Program Criteria Assessment

(a) an ability to apply knowledge of mathematics, science, and engineering
(b) an ability to design and conduct experiments, as well as to analyze and interpret data
(c) an ability to design a system, component, or process to meet desired needs within realistic constraints such as economic, environmental, social, political, ethical, health and safety, manufacturability, and sustainability
(d) an ability to function on multi-disciplinary teams
(e) an ability to identify, formulate, and solve engineering problems
(f) an understanding of professional and ethical responsibility
(g) an ability to communicate effectively
(h) the broad education necessary to understand the impact of engineering solutions in a global, economic, environmental, and societal context
(i) a recognition of the need for, and an ability to engage in life-long learning
(j) a knowledge of contemporary issues
(k) an ability to use the techniques, skills, and modern engineering tools necessary for engineering practice
…Knowledge of …. basic sciences, computer science, and engineering sciences necessary to analyze and design (a) complex electrical and electronic devices, (b) software, and (c) systems containing hardware and software components, as appropriate to program objectives.

IEEE Code of Conduct

1. to accept responsibility in making engineering decisions consistent with the safety, health and welfare of the public, and to disclose promptly factors that might endanger the public or the environment;
2. to avoid real or perceived conflicts of interest whenever possible, and to disclose them to affected parties when they do exist;
3. to be honest and realistic in stating claims or estimates based on available data;
4. to reject bribery in all its forms;
5. **to improve the understanding of technology, its appropriate application, and potential consequences;**
6. to maintain and **improve** our technical competence and to undertake technological tasks for others only if qualified by training or experience, or after full disclosure of pertinent limitations;
7. **to seek, accept, and offer honest criticism of technical work,** to acknowledge and correct errors, and to credit properly the contributions of others;
8. to treat fairly all persons regardless of such factors as race, religion, gender, disability, age, or national origin;
9. to avoid injuring others, their property, reputation, or employment by false or malicious action;
10. **to assist colleagues and co-workers in their professional development and support them in following this code of ethics.**

NSF Ethics Education in Science and Engineering

"… will develop learning units that focus on topics of such active public interest that there is ongoing, rapid change in the laws that provide context for discussions of the associated professional ethics issues. Examples include **whistle-blowing; reverse engineering; investigation of security vulnerabilities in running systems; and conflicts of interest and intellectual property rights** associated with university laboratories and faculty-owned businesses that commercialize university-developed research." Cem Kaner kaner@kaner.com(PI); Ephraim P. Glinert (PM) IIS

"… the main deliverable of the project being a book having four main sections: **ownership; privacy; access; and safety, reliability, and liability** … the PIs' book will address ethical and social issues in a manner that attends to the cognitive, social, and affective aspects of ethical development of human beings during early adulthood…" Melissa Dark dark@purdue.edu(PI) A collaborative proposal. Ephraim P. Glinert (PM) IIS

"The project will provide an intellectual contribution to the scientific community by **teaching scientists how their discoveries fit into broader social and humanistic contexts**, and to the philosophical community by creating new theoretical and practical tools for applied ethics …the project also promotes teaching through the production of new educational materials and a coordinated program of study…" N. Dane Scott dane.scott@mso.umt.edu (PI); Priscilla Regan (PM), Division of Social and Economic Sciences, Directorate for Social, Behavioral & Economic Sciences

Cobweb: an incremental system for hierarchical clustering (Fisher, 1987)

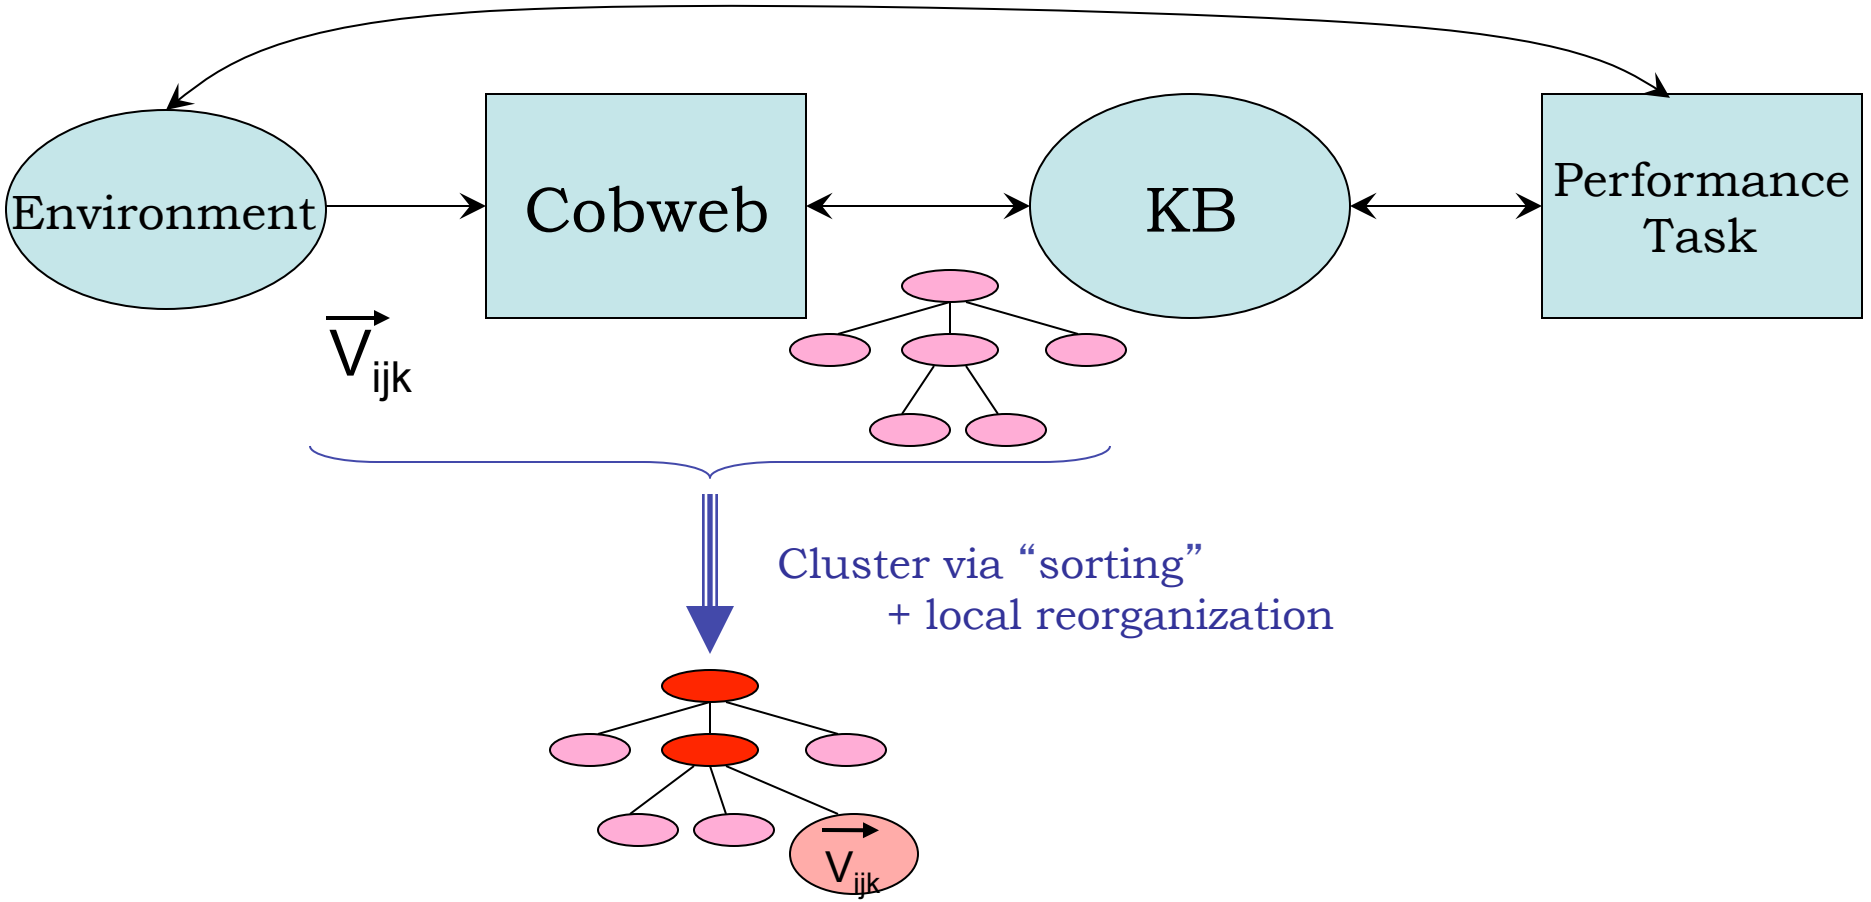

Cobweb resulted from a synthesis of ideas:

- Michalski (clustering as search)
- Lebowitz and Kolodner (sorting or a hill-climbing search, motivation for, prediction)
- Gluck and Corter (a measure for predicting basic levels as an evaluation function)

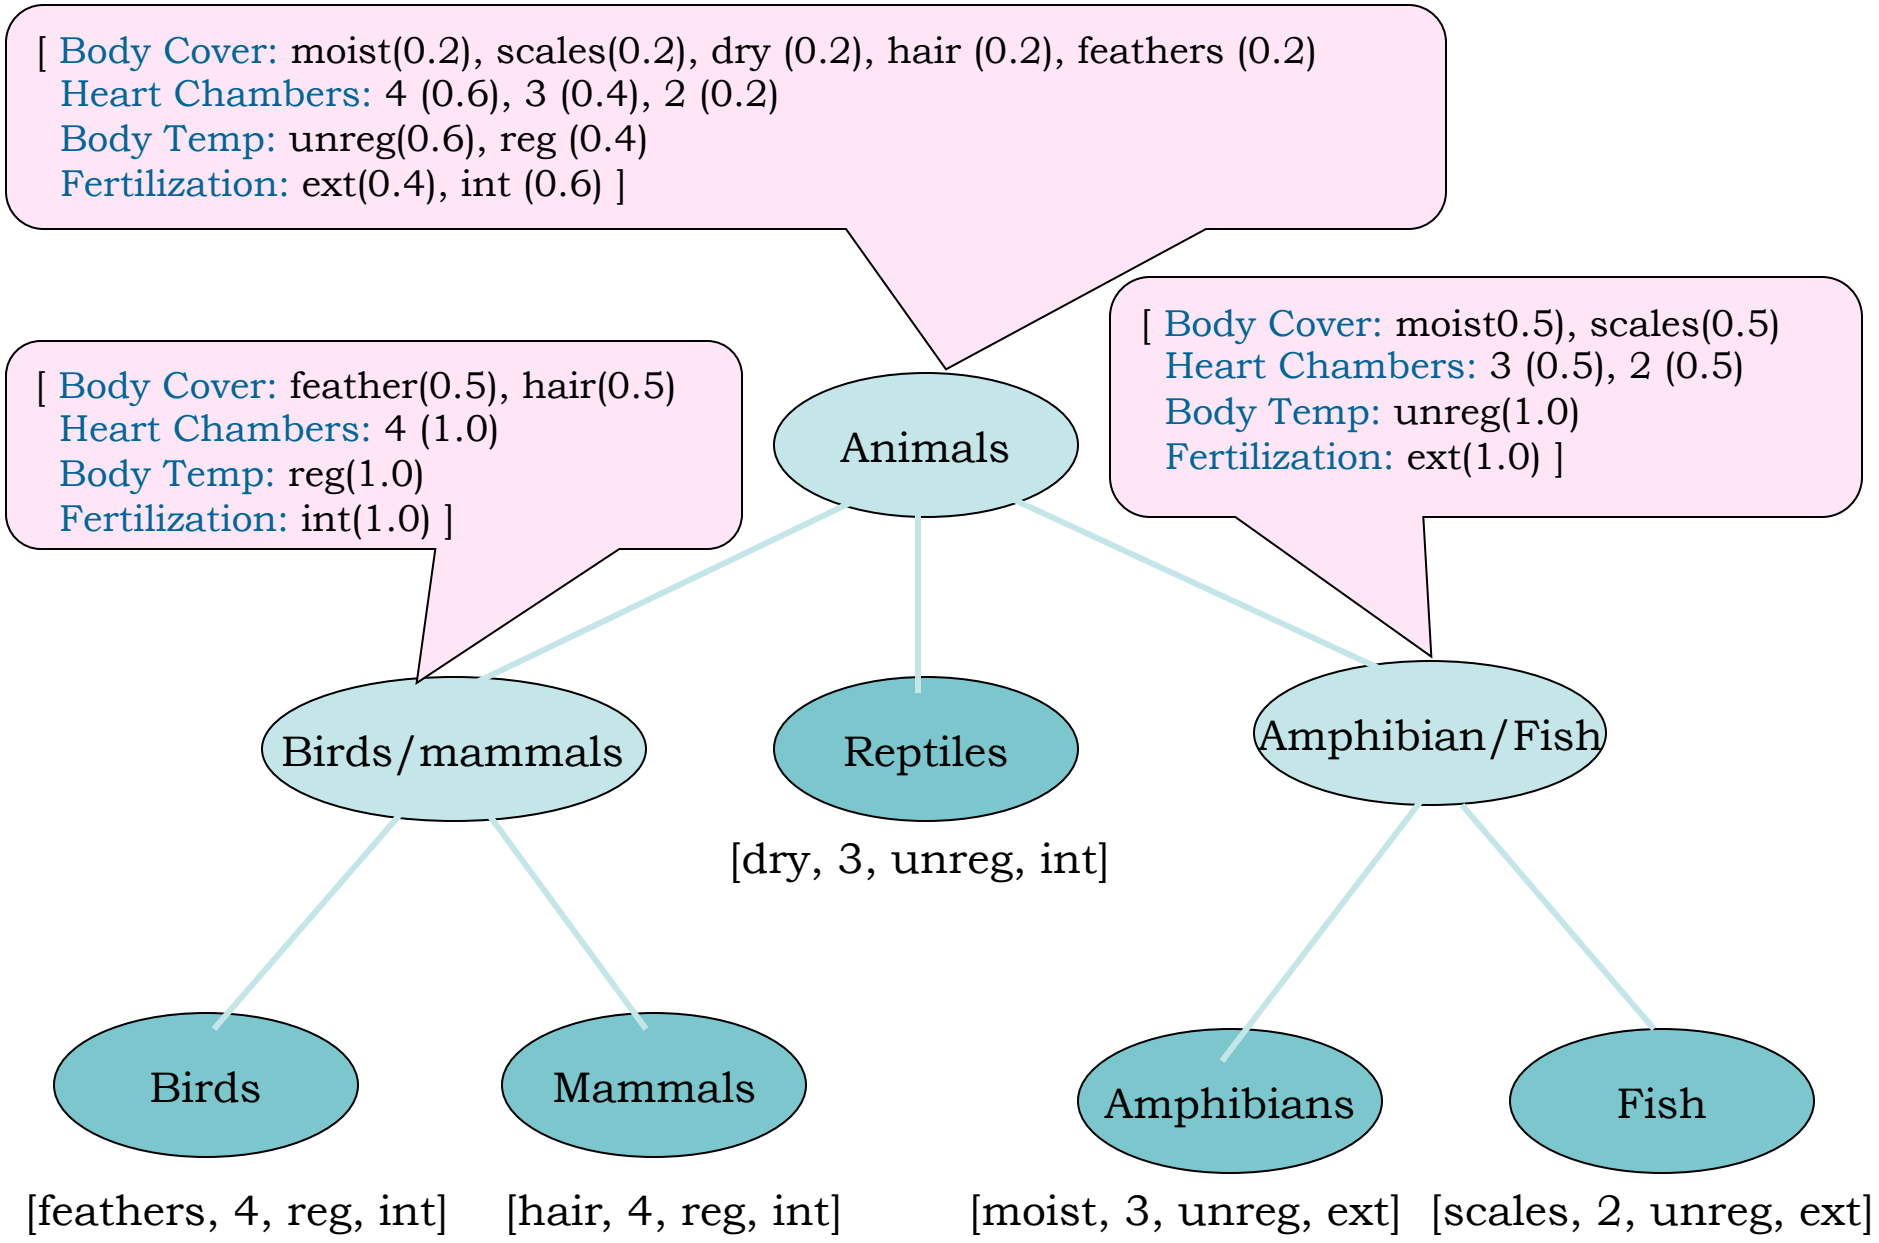

Example of data and hierarchy

Cobweb: pattern completion (Fisher, 1987; Fisher, 1996)

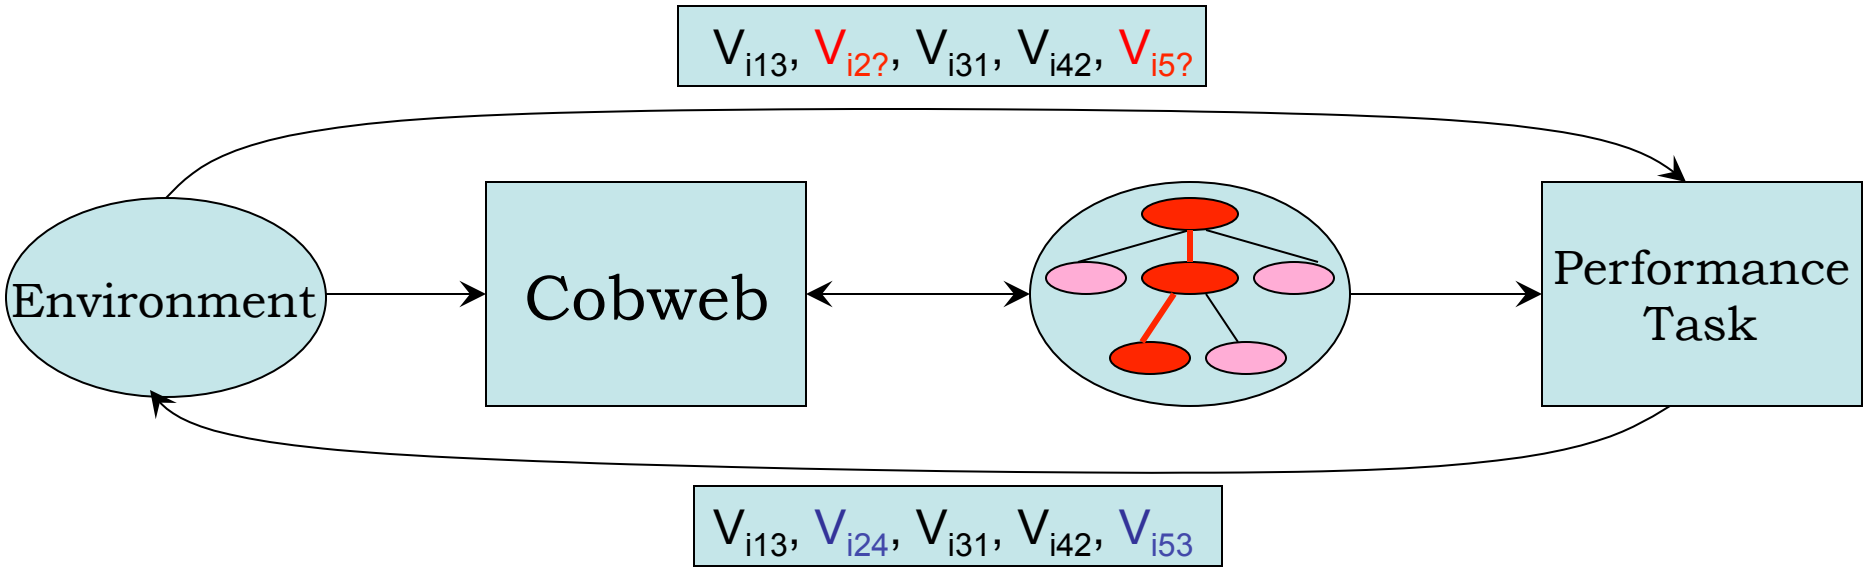

Other approaches to Unsupervised Learning (Fisher, 2001) that can be adapted to pattern completion:

- Learning Association Rule Sets
- Clustering
- Learning Bayesian Networks

Related learning paradigms:

- Multi-Task Learning (Caruana, 1997)
- Data mining clustering (Fisher, 1995, 1996)

The Basic Level: getting the most **bang** for the **buck**

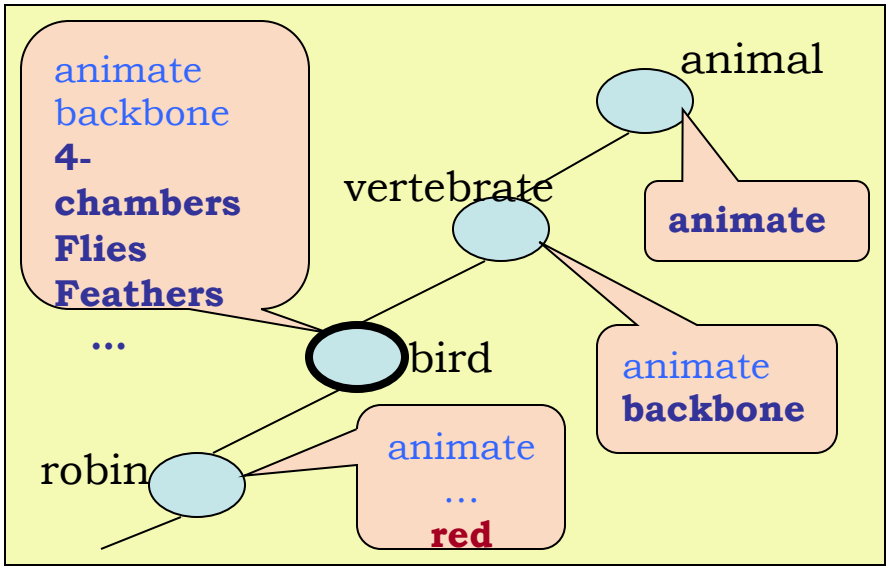

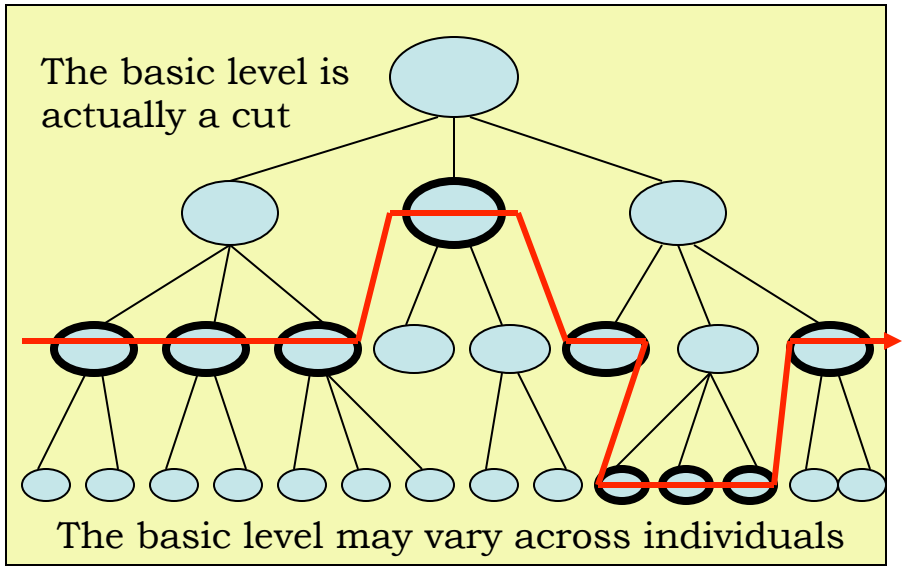

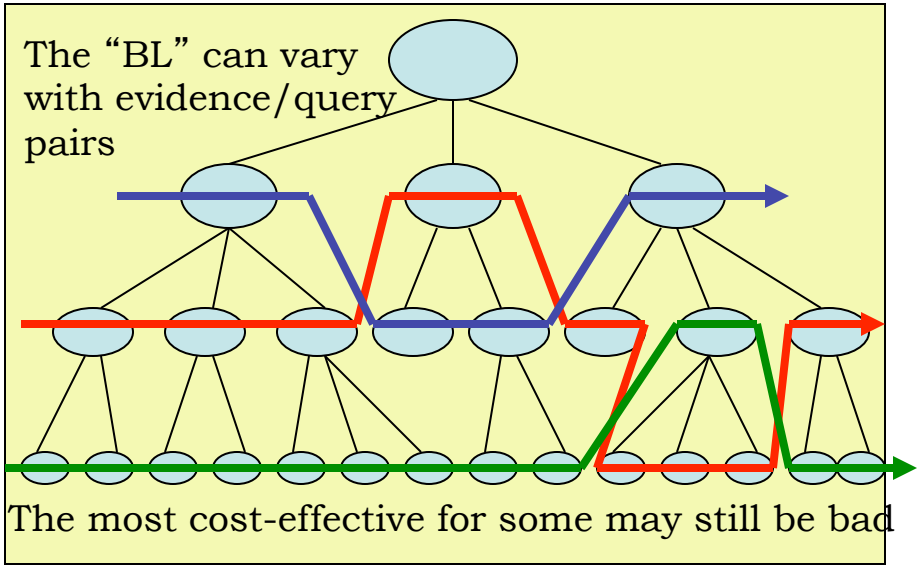

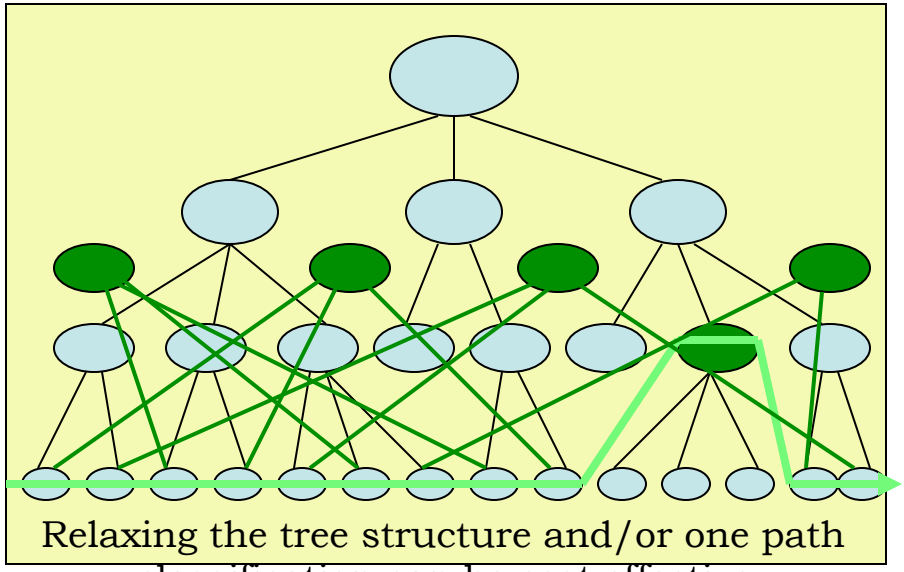

Exor: Casting learning to problem solve as concept learning (Yoo & Fisher, 1991; Fisher & Yoo, 1993)

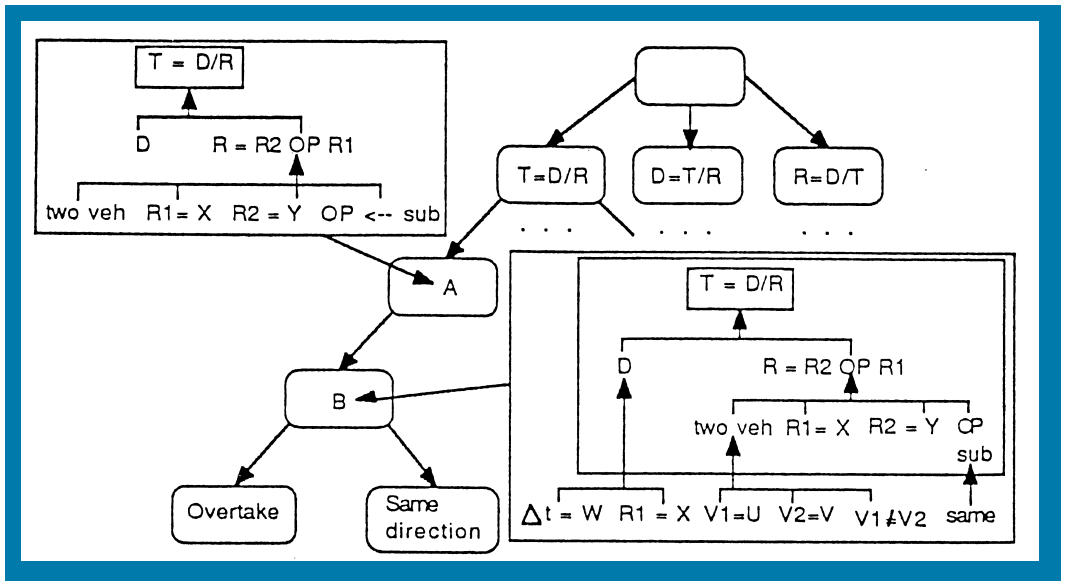

Synthesis of ideas from sources such as Expert vs. Novice problem solving (Chi et al, 1981), Learning operator preferences (Langley, 1985; Mitchell et al, 1986), Selective utilization (Minton, 1988; Mooney, 1989; Markovitch & Scott, 1989), Case-based problem solving (Callan, Fawcett, & Rissland, 1991)

Problem solutions arranged in an abstraction hierarchy

New problems solved via a combination of classification with respect to known solutions and domain theory search

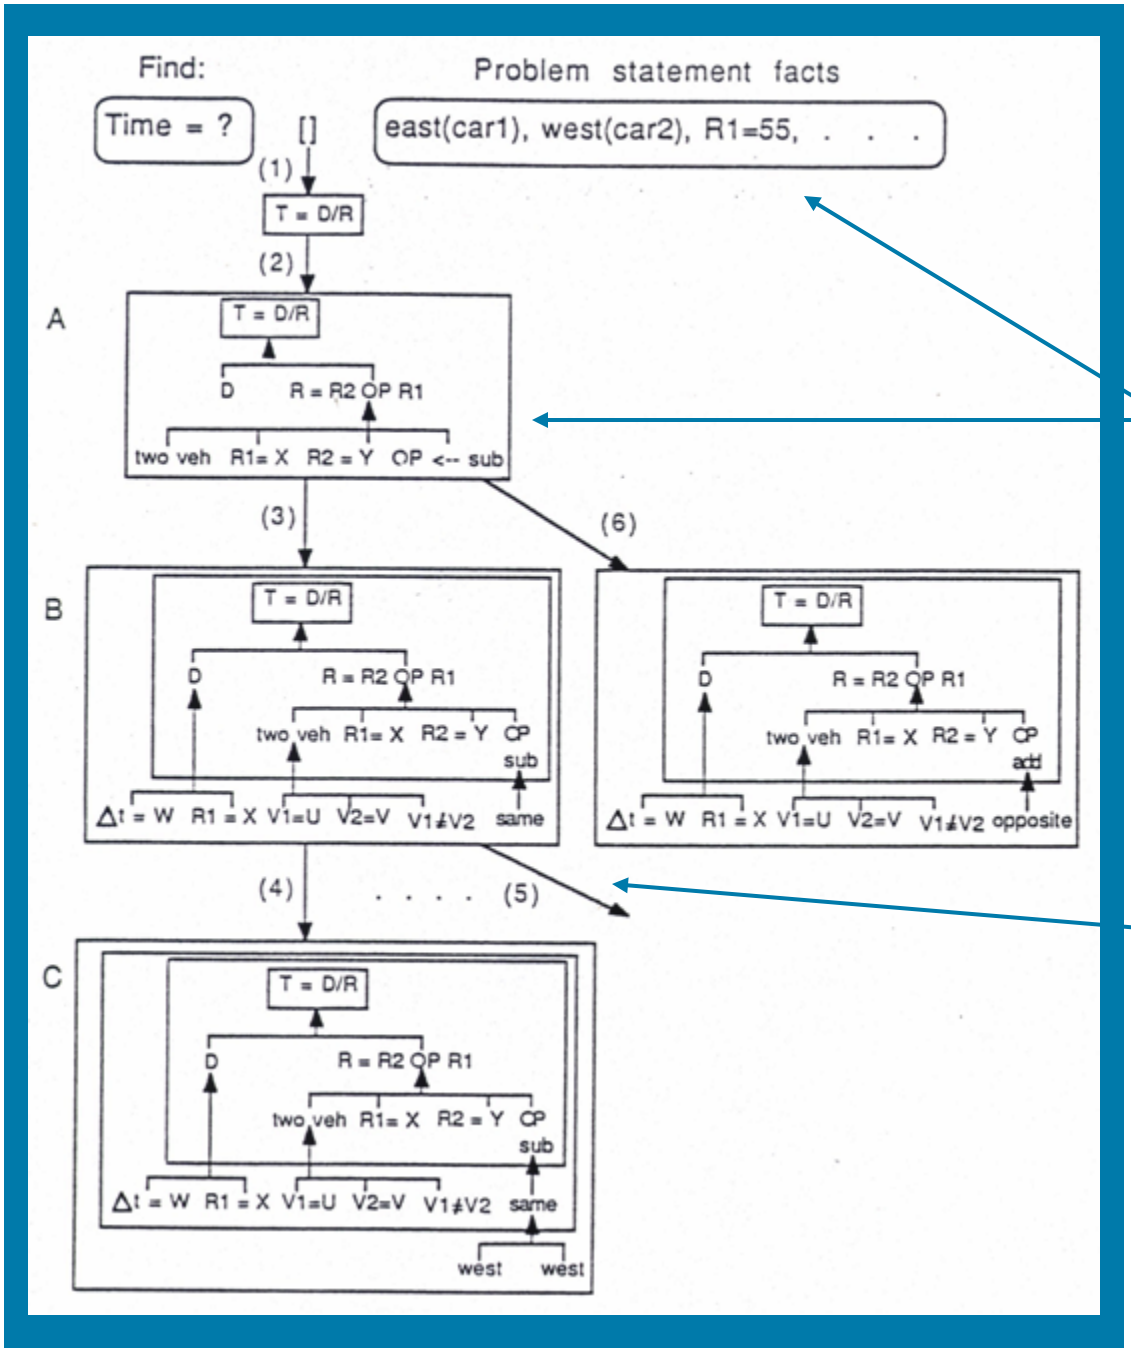

An illustration of Exor's classification-driven problem solving

- Problem's categorized by matching knowns (observed or inferred) against class/cluster distributions (the latter no shown).
- The cluster's partial solution is asserted as a partial solution to new problem.
- If a contradiction is found, backtracking occurs (and retraction of previously-asserted solutions) and a next-best choice is tried.

- Exor extended *boundary of operationality* to cost effective features, and also theorized about what deep features were used by experts:

$$EC(C_i) > P(F_k \mid C_i)\,[\,EC(C_i \mid F_k) + EC(\text{prove } F_k)\,] + [1 - P(F_k \mid C_i)\,]\,[EC(C_i \mid \neg F_k) + EC(\text{prove } \neg F_k)]$$

- Exor exploited other concept learning strategies such as *pruning* to improve problem solving performance

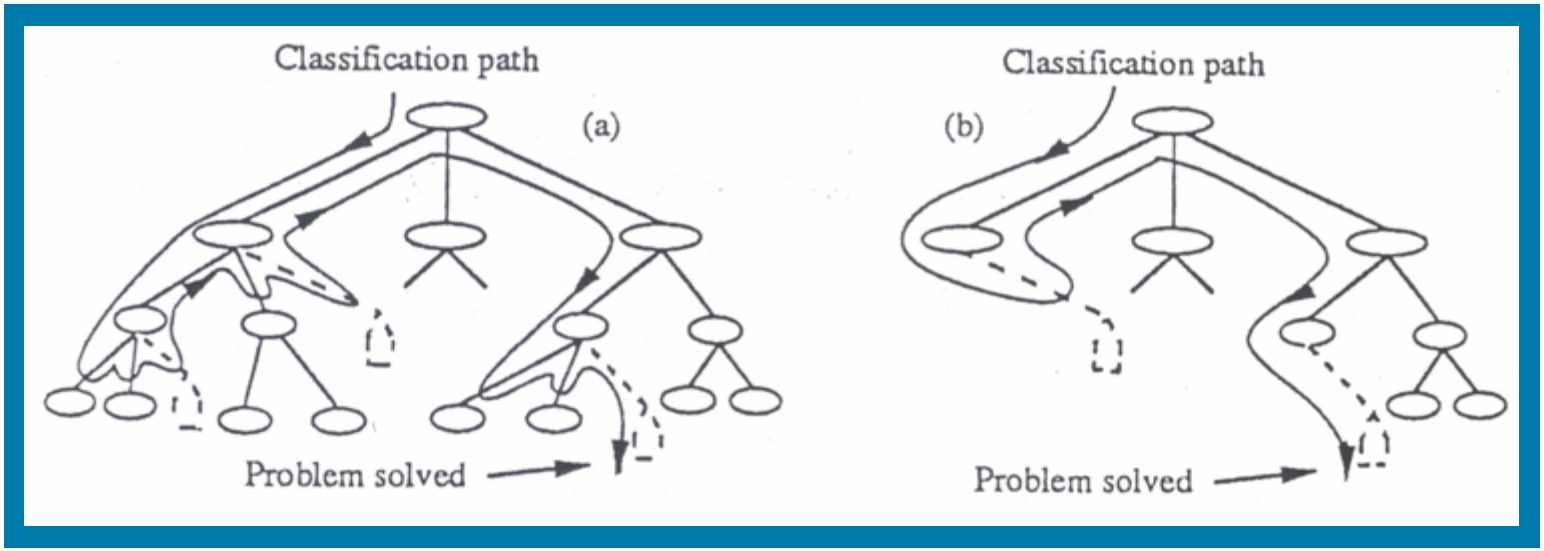

- Exor introduced the ideas of context-based utilization (in contrast to selective utilization) and context-based examination

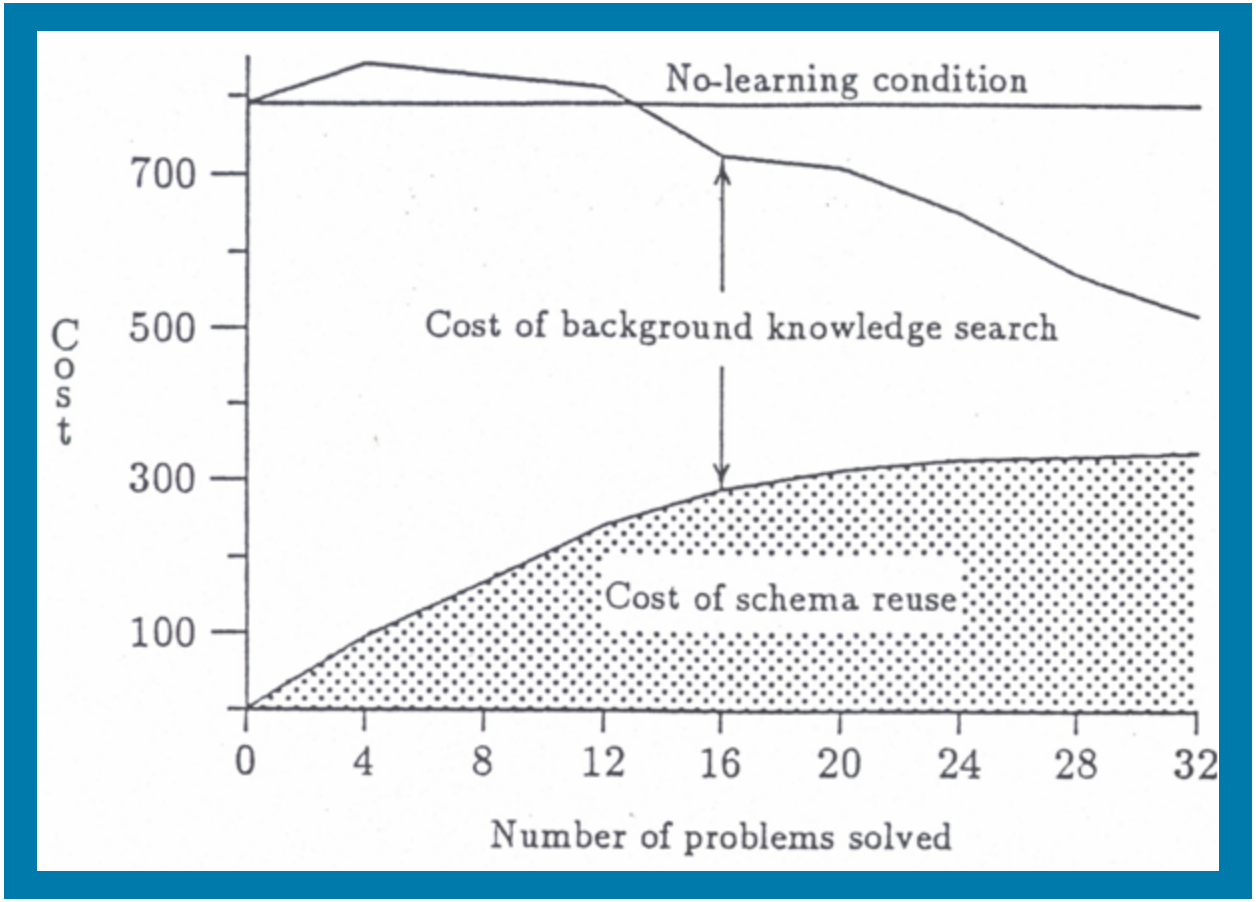

Related to

- Hierarchical Case-Based Reasoning (Smyth, Keane, & Cunningham. 2001)

Backward-Chaining Rule Induction (Fisher, Edgerton, et al 2005, 2006)

- Outcome-influenced "Association Rule Learning"

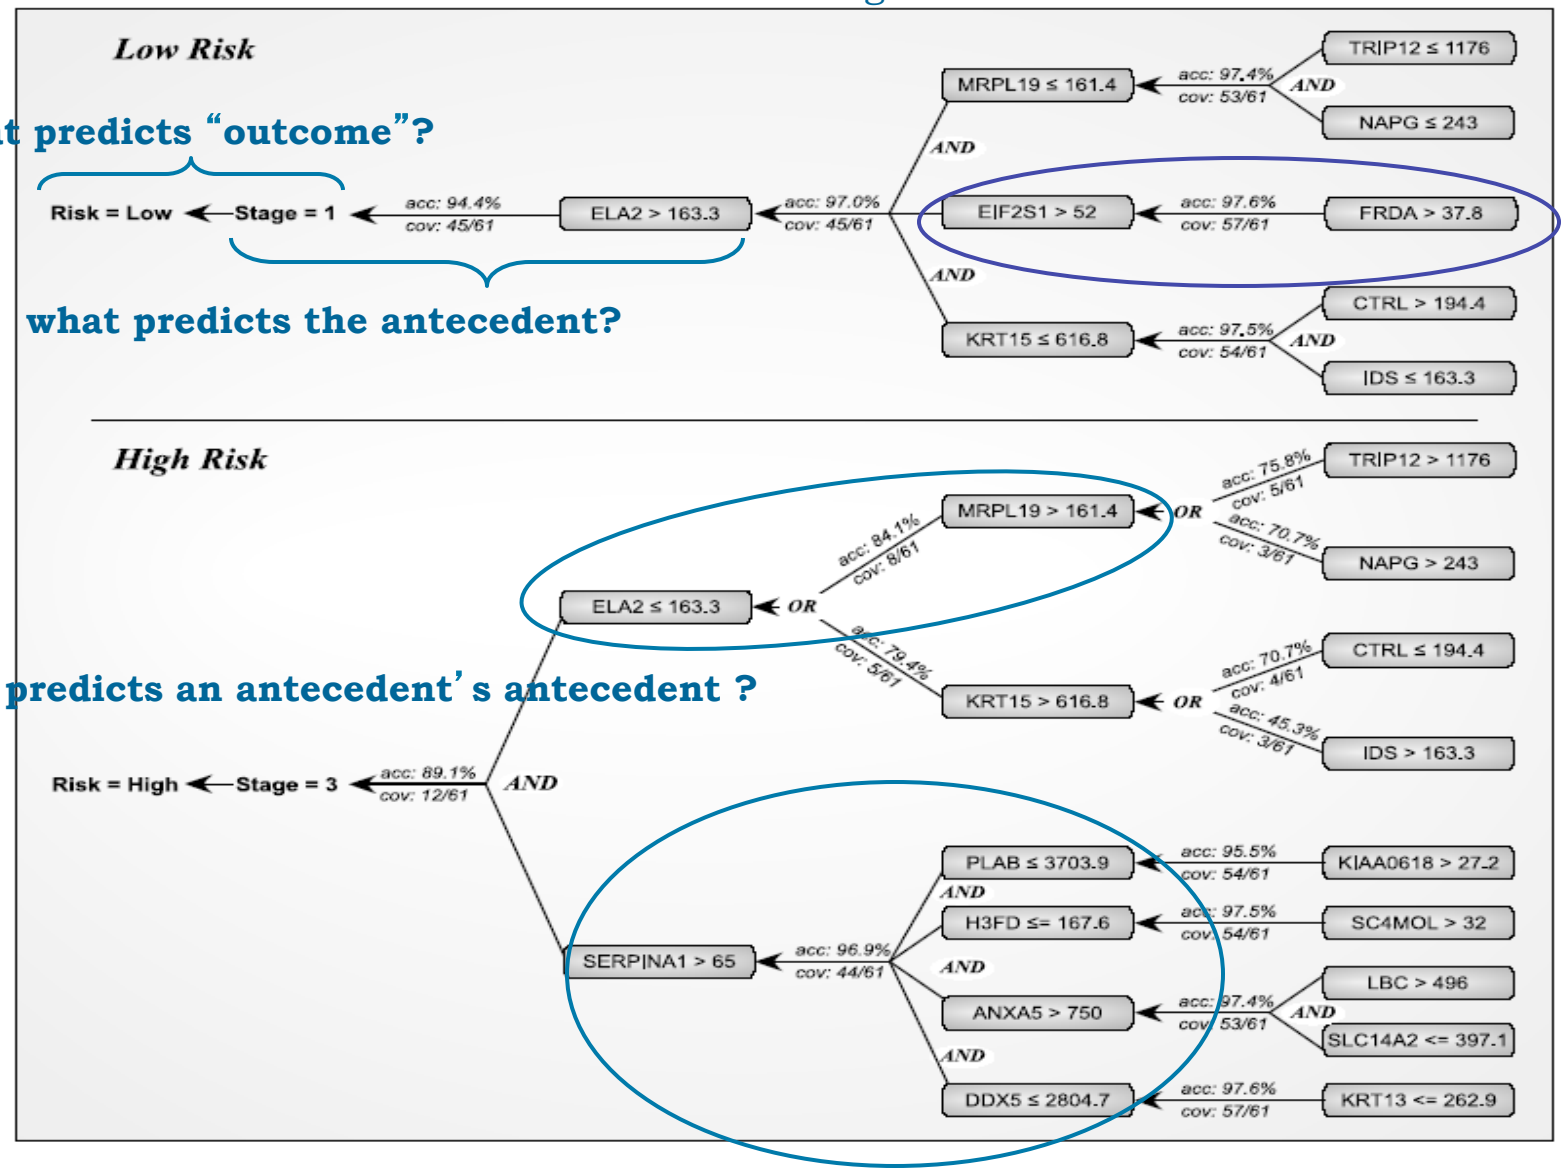

Example on next slide

- BCRI is a strategy for exploration and hypothesis generation, not classifier construction

Compartmentalization and synthesis in research

Promote synthesis through

- Climbing up the organizational hierarchy
- History and education
- Cognitive architectures
- Overriding applications (contextualize research and research results)
- Balanced community

Compartmentalization and synthesis in research

- What is the mean number of citations per pub of NSF-funded work?
- How do we evaluate citations beyond mean number (e.g., breadth or scope)
  - Authority
  - "Markov Blanket" or discounted measures
    (Authority + descendent and ancestor authority)
- How to reach backward, particularly past the "Web horizon",
  in an age of information loss and overload
- Adapt collaborative filtering to scholarship
- Adapt RL to looking back through citations chains