Proceedings of the Workshop on

# Models of Complex Human Learning

Cornell University
Ithaca, NY

June 27–28, 1989

Sponsored by:

Computer Science Dept.
Cornell University

Cognitive Science Branch
Office of Naval Research

Computer Science Dept.
University of Illinois

**Proceedings of the Workshop on**

# Models of Complex Human Learning

*Held in conjuction with the Sixth International Machine Learning Worshop*

**Rooms 700–702, Clark Hall**
**Cornell University**
**Ithaca, NY**

**June 27–28, 1989**

**Edited by David C. Wilkins**

**Sponsored by:**

**Computer Science Department, Cornell University**
**Cognitive Science Branch, Office of Naval Research**
**Computer Science Department, University of Illinois**

**To obtain copies of this document, request Report UIUCDCSR-89-1520 from:**

**Erna Amerman, Reports Librarian**
**Department of Computer Science**
**University of Illinois**
**1304 West Springfield Ave**
**Urbana, IL 61801**

# Foreword

Historically, much Artificial Ihtelligence (AI) research has been motivated by a desire to better understand human cognition. The ***Workshop*** on ***Models of Complex*** Human Learning is the first **AI** learning workshop to focus exclusively on learning research that gives attention to human data and has implications for understanding human cognition. Of particular interest at this workshop is learning research that models complex problem-solving tasks at the symbol level, such as acquisition of high-level programming skills, acquisition of problem-solving expertise, and learning via analogical reasoning between different problem domains.

Our desire to facilitate in-depth communication of research results in a multi-disciplinary gathering led to a decision to have long presentations and limit the workshop to forty-five participants. This decision precluded the acceptance of many excellent submissions; only 19% of the presentation submissions were accepted. A poster session allows an additional 11 submissions to receive a formal exposure. The proceedings is composed of abstracts, rather than full papers, as this facilitates presentation of very recent results. However, the abstracts contain publication references to recent full papers.

This ONR-sponsored workshop is an expansion of the annual **ONR** review of research, to include participants from the larger learning community; less than half of the abstracts describe research sponsored by **ONR**. Bringing together so many leading members of the cognitively-oriented machine learning community - with a rough balance between cognitive psychologists and AI computer scientists - provides a rare opportunity to review the entire enterprise. Here are four suggestions for issues to pursue formally and informally during the workshop.

First, what changes should occur in research in cognitively-oriented machine learning as the machine learning field enters an era of intense specialization? What opportunities are provided by new machine learning subfields such as computational learning theory and explanation based learning?

Second, as the machine learning field increases its emphasis on empirical validation and comparative analysis of results, learning testbeds will be established and become widely available. What form should these take, and how will they be of value to cognitive psychologists?

Third, how can machine learning research give better attention to data acquired **from** humans? The major methods appear to be the use of protocol analysis and fidelity to psychological effects, such as the log-log law of practice. Allen Newell estimates that there are 3000 psychological effects that relate to language acquisition, motor control, perceptual phenomena such as attention, decision making, and expert-novice differences. Would a catalog of these effects be of benefit to cognitively-oriented machine learning?

Lastly, how can computer scientists be made more aware of the methodological requirements in the proper use of human data? How can psychologists be made more aware of the challenges and pitfalls in the use of computational process models?

I would like to express my deep gratitude to Alberto Segre of Cornell University and Sharon Collins, Laura Mohlenkamp, and Beth **Shirk** of the University of **Illinois,** for their substantial efforts in organizing this workshop **and** in preparing the proceedings. The workshop would not have been possible without the intellectual and financial support of Susan Chipman.

David ***C.*** Wilkins
Program Chair

# Abstracts by Subject Area

## Acquisition of Programming Skills



## Cognitive Architectures and Learning



## Knowledge Acquisition



## Learning of Natural Concepts and Categories



## Mathematical Skill Acquisition

## Memory Organization and Learning



## Natural Language Acquisition



## Reasoning and Learning by Analogy



## Theory Revision

# Workshop Program Schedule

## Tuesday, June 27, 1989

| | |
|---|---|
| **8:00–10:00** | *Registration, Coffee and Tea,* Clark Hall |
| **9:10–9:30** | *Welcome,* David Wilkins and Susan Chipman |
| | **Reasoning and Learning by Analogy** |
| **9:30–10:20** | Dedre Gentner and Ken Forbus, University of Illinois<br>Psychological Studies of Access and Inference |
| | **Learning of Natural Concepts and Categories** |
| **10:20–11:10** | Douglas Fisher, Vanderbilt University<br>Model of Natural Category Structure and its Paradigmatic Implications |
| **11:10–11:40** | *Break, Refreshments* |
| **11:40–12:50** | Poster Presentations |
| **12:50–2:10** | *Lunch* |
| **2:10–3:00** | Raymond Mooney, University of Texas, Austin<br>Integrated Learning of Explanatory and Nonexplanatory Information |
| **3:00–3:50** | Michael Pazzani, University of California, Irvine<br>Computational Model of Influence of Prior Knowledge on Concept Acquisition |
| **3:50–4:20** | *Break, Refreshments* |
| | **Knowledge Acquisition** |
| **4:20–5:10** | David Wilkins, University of Illinois<br>Automated Knowledge Acquisition Using Apprenticeship Learning Techniques |
| **5:10–6:00** | David Kieras, University of Michigan<br>Explanation-Based Knowledge Acquisition of Electronics |

# Workshop Program Schedule

## Poster Presentations
### Tuesday, June 27, 1989
### 11:40–12:50

### Acquisition of Programming Skills

Robert Campbell, IBM T.J. Watson Research Center
Expertise in Programming: A Developmental Approach

Peter Pirolli and Margaret Recker, University of California, Berkeley
The Explanation **of** Programming Examples

### Learning of Natural Concepts and Categories

Doug Medin and Edward Wisniewski, University of Illinois
Category Learning: The Effect **of** Prior Knowledge and Different Conceptual Roles

Jeff Shrager, Xerox PARC
Reinterpretation **of** the Perceptual Microstructure **of** Conceptual Knowledge

### Natural Language Acquisition

Shyam Kapur, Cornell University
Conservative Language Acquisition from Positive Examples

Sheldon Nicholl, University of Illinois
A New Computational Model **of** Language Acquisition

### Reasoning and Learning by Analogy

Beth Adelson, Tufts University
Learning by Analogy: Find Causal Explanations

Tom Eskridge, New Mexico State University
Continuous Analogical Reasoning

Brian Falkenhainer, Xerox PARC
Explanation Through Physical Analogies

### Theory Revision

Paul O'Rorke, University of California, Irvine
Abduction and Learning Involving Physical and Psychological Explanations

Ronald **Yaeger** and Kenneth Ford, West Florida University
**A** Constructivist Model **of** Human Learning

# Workshop Program Schedule

## Wednesday, June 28, 1989

### Cognitive Architectures and Learning

**9:00–9:50** Kurt VanLehn, Carnegie Mellon University
Learning Events in Three Skills

**9:50–10:40** Pat Langley, University of California, Irvine
Icarus: An Integrated Cognitive Architecture

**10:40–11:10** **Break, Refreshments**

**11:10–12:00** Ryszard Michalski and Deborah Boehm-Davis, George Mason University
A Theory of Human Plausible Reasoning: Efforts on an Experimental Validation

**12:00–2:00** **Lunch, Andrew *Dickinson* White House**

### Memory Organization and Learning

**2:00–2:50** Kris Hammond, Univ. of Chicago and Colleen Seifert, Univ. of Michigan
Opportunistic Memory

**2:50–3:40** Thomas Bever, University of Rochester
The Induction of Abstract Representations During the Use of Symbolic Systems

**3:40–4:10** **Break, *Refreshments***

### Mathematical Skill Acquisition

**4:10–5:00** Sttllan Ohlsson, **LRDC**
The Function of Conceptual Understanding in the Learning of Arithmetic Procedures

**5:00–5:50** Sandra Marshall, San Diego State University
Acquisition of **Schema** Knowledge

# Abstracts

# Learning by Analogy: Finding Causal Explanations Across Domains

**Beth Adelson**
**Department** of **Computer Science**
Tufts **University**
**Medford, MA 02155**

The theory we present was developed within a problem-solving context, reflecting the purpose of analogical reasoning and, as a result, providing insights into the phenomenon. This has allowed us to move towards increasing the theory's specificity and valididty. The theory was developed using protocol data and has been implemented as a computer model.

Recent research suggests a class of theories of analogy which rest on the processes of retrieval, mapping, evaluation, debugging and generalization. Our work extends existing theories by specifying mapping, evaluation and debugging as simulation-based reasoning processes, constrained by the problem-solving context and dependent on knowledge about the relationship between function and structure. Below we discuss each of these processes in turn.

**1. Mapping:** Our mapper reflects the way in which students limit their focus of attention to the aspects of a domain that are relevant to the type of problem they are learning to solve (Burstein & Adelson, 1987). Our mapping mechanism selects one model among a set of models describing various ***operations*** in the base domain in various ways (e.g., in terms of behavior or mechanism). It then maps this model over to the target domain. After debugging the first model, additional models that provide explanations for the first model are also mapped into the target. This incremental mapping process allows the system to evaluate and debug one model at a time. It also allows the mapper to become 'tuned' during learning. That is, on later mappings the mapper leaves behind pieces of a base model that explain phenomena that have previously been found irrelevant to the target domain. The result is that evaluation and debugging are further simplified.

**2. Evaluation:** The base domain model provides an imperfect model of the target domain. Our system's evaluation mechanism decides whether elements in a model mapped over from the base remain appropriate in the target. In doing so, the system's evaluation mechanism uses knowledge about: 1. The class that the elements in the model belong to. 2. The class of objects to which an attribute can be applied. 3. The ***classes*** of objects and predicates appropriate to the domain being learned. 4. General knowledge about the way in which relations apply differently across analogous domains.

**3. Debugging:** When a functional aspect of a newly mapped model has been identified as appropriate, but the mechanism mapped from the base is inappropriate the system runs simulations of the model in the base and target domains in order to find analogous, domain-appropriate replacements. In choosing and evaluating these simulations, the system uses knowledge about the relationship among actions, the mechanisms they effect and the results they produce. This process allows the system to maintain functional aspects of analogical examples, while adding target-appropriate causal explanations.

Adelson, Beth. Cognitive modeling: Uncovering how designers design. ***The* Journal *of Engineering Design..*** **Vol 1,1. 1989.**

Adelson, Beth. A common-sense analogical reasoner. ***Cognitive Science.*** In preparation.

Adelson, Beth. When novices surpass experts: How the difficulty of a task may increase with expertise. ***Journal of Ezperimental Psychology: Learning, Memory and Cognition,*** July **1984.**

Burstein, M. and Adelson, B. Analogical Reasoning for Learning. in ***Applications of Artificial Intelligence to Educational Testing.*** R. Freedle (Ed.) Erlbaum: Hillsdale, NJ. In press.

# Implementation of a Rational Analysis in an ACT Architecure

**John R. Anderson**
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 37235

Since 1976 (Anderson, 1976) we have been working on developing the ACT framework for developing an architecture of human cognition. A rather complete proposal within the ACT framework, called ACT*, was developed in 1983 (Anderson, 1983). More recently we have been engaged in an effort to develop a rational analysis of human cognition within the ACT framework (Anderson, in press). This takes the form of deriving some prescriptions for optimality within the general framework. As it seems these prescriptions are also descriptively accurate of human cognition, we are now turning to the issue of how these prescriptions might be implemented to create a new specific architecture to replace ACT*.

This talk will briefly review the empirical status of the ACT* theory and the motivations for a rational analysis and overview the rational analysis of four key domains of human cognition: memory, categorization, causal inference, and problem solving. We will explore the implications of these analysis for an ACT architecture. This will involve making modifications to the activation computation and conflict resolution principles in ACT* to a form which is actually computationally more tractable. Another major modification is to move the inductive learning component from the procedural (production) memory to the declarative memory.

Anderson, J. R. Language, Memory, and Thought. Hillsdale, N.J.: Erlbaum, 1976.

Anderson, J. R. The Architecture of Cognition. Cambridge, Mass: Harvard, 1983.

Anderson, J. R. (in press) The Adaptive Charactive of Thought, Hillsdale, N.J.: Erlbaum.

# The Induction of Abstract Representations During the Use of Symbolic Systems

**Tom G. Bever**
University of Rochester
Department of Psychology
Rochester, NY

Mental maps and linguistic grammars have similar formal properties: they both represent what is true of their respective worlds (geometric, language) regardless of how the structure is being used at any given time. It has been alleged that map making in rats (and probably humans) and grammar making (in humans) are innate capacities. But, such allegations are unhelpful, in that they do not tell us what the learning circumstances are which evoke the innate capacity. We have been exploring several paradigms for map and language learning, as case studies in the relation between learning behaviors and learning an abstract structure which might underly them. We are testing the hypothesis that the formation of an abstract representations is elicted as a solution to a representational conflict at a more superficial level of representation.

We get around in the world either by following landmarks, or using mentalmaps. In a simple study with rats and humans, we have them learn to negotiate a digital figure 8 maze in which during initial training the center cross alley is never used though it can be explored that is, subjects learn to run from the top to the bottom of the maze using only one side or the other of the periphery (with rats, we use a standard wood alley maze, about 8 feet by four feet; with humans, we use the basement of our psychology building which happens to be in the form of a figure 8, about 40 by 25 feet) the human subjects also wear pinhole goggles to degrade their visual input). "One way" subjects always learn the maze from one end to the other; "two way" subjects learn the maze starting at both the top and the bottom of the maze, on different trials. We then test speed of learning to use the center cross alley, running the maze in a '5' or 's'. "Both way" subjects learn to use the cross alley faster than one way subjects.

Language is used for speaking and listening, with behavioral systems which seem to be at least partially independent of each other. We have used the acquisition of an artificial language to study the relationship between learning to map instances of a world and a symbol system, and learning the abstract grammatical structure inherent to that system: "perception" subjects learn to map grammatically structured symbol sequences onto an array of geometric forms; "production" subjects learn to map from geometric arrays to symbol sequences; "bidirectional" subjects learn to map in both directions. All subjects are periodically tested for their ability to make grammaticality judgments. The results indicate that when training goes in both directions, the two types of training interact via the formation of an abstract representation of the symbol system) i.e., subjects develop a grammar. The results also suggest that "production" is more intimately tied to the ability to access grammatical knowledge than is perception.

Bever, T.G., The aesthetic constraint on cognitive structures. In M. Brand and R. Harnish (eds.), The representation of knowledge and belief. Tuscon Arieona; University of Arizona Press. 1987.

Bever, T.G. and Hansen, R. E. The induction of mental structures while learning to use symbolic systems. In Proceedings of the tenth annual conference of the cognitive science society. LEA.

# A Theory of Human Plausible Reasoning: Efforts on an Experimental Validation

**Deborah Boehm-Davis**
**Department of Psychology**
**George Mason University**
**Fairfax, VA 22030**

**Ryszard Michalski**
**Department of Computer Science**
**George Mason University**
**Fairfax, VA 22030**

The paper will present a brief review of the main assumptions and components of the Collins and Michalski theory of human plausible reasoning. The components include assumptions about the hierarchical structure of human knowledge, the representation of knowledge as *traces* through the nodes of different hierarchies, and plausible reasoning as certain perturbations of those traces. The traces are annotated by a variety of parameters influencing the degree of belief associated with the corresponding piece of knowledge. The theory identifies a number of different types of reasoning based on specialisation, generalization, similarity, and dissimilarity among the different components of statements.

This work presents the results of studies conducted with human subjects who were asked to answer questions requiring them to conduct reasoning. Their answers were analyzed in terms of the concepts and inference rules developed in the theory. The purpose of this analysis was to validate the theory and to determine what enhancements or extensions were needed to account for the data. This analysis was restricted to the structural properties of the model and the types of inferences involved in reasoning. Future studies will examine the processes associated with assigning certainty to the conclusions. The analyses confirmed that people follow several lines of reasoning in reaching given conclusions. It also suggests some rules that were not captured in the original model.

Collins, A. and Michalski, R. S. (1989) The Logic of Plausible Reasoning: A Core Theory, ***Cognitive* Science,** June 1989

Dontas, K. and Zemankova, M.(1988): 'APPLAUS: An Implementation of the Collins-Michalski Theory of Plausible Reasoning.' The Third International Symposium on Methodologies for Intelligent Systems.

Kelly, J. (1988): PRS: A system of Plausiblle Reasoning. MS Thesis. Dept. of Computer Science. University of Illinois, Urbana.

# Expertise in Programming: A Developmental Approach

**Robert L. Campbell, Norman R. Brown and Lia A. DiBello**
User Interface Institute
IBM T. J. Watson Center, P.O. Box 704
Yorktown Heights NY 10508

One aim of research in human-computer interaction is to understand how people become expert programmers and how we might make it easier for them to do so. However, HCI research has not seriously examined the acquisition of expertise. Novices and experts are typically defined in terms of years of experience, not in terms of characteristic skills or accomplishments, and nothing is said about the process by which expertise is acquired.

Standard cognitive science approaches to expertise (those of Chi, Larkin, diSessa, Carey, and others) have not filled this gap in our understanding. Cognitive science accounts make binary comparisons between experts and novices without considering the process of acquisition. They do not explain how novices become experts. Neither information-processing models (e.g., self-modifying production systems) nor philosophical accounts of theory change in science offer developmental processes of sufficient power.

To understand the acquisition of expertise in programming, we look to constructivist accounts of development, such as those of Piaget, Bickhard, Vygotsky, and Feldman. Such accounts describe processes of learning and reflective abstraction that are powerful enough to construct new knowledge and reorganize old knowledge. Constructivist developmental theories treat learning as a variation and selection process that must move through relatively stable intermediate states. These accounts imply that development from novice to expert passes through a sequence of levels.

In our initial investigation, we interviewed 7 experts in various areas of programming, focusing on their accounts of their own development, and on their criteria for distinguishing the work of experts from the work of less advanced programmers. The programmers that we interviewed often thought of themselves as going through major qualitative changes in understanding, including changes in the meaning of key concepts in the language that they were mastering. They cited a number of areas of difference between programmers at different levels of expertise, such as the use of variables by novice, intermediate, and expert C programmers, and the increasing ability of more advanced programmers to consider alternatives to the design that they actually produced and defend their choice of that design.

A second study focused on three professional programmers and one non-programmer learning the Smalltalk/V language and environment. This was a short-term longitudinal study, lasting from 2 weeks to 2 months. Participants in the study kept diaries on audiotape as they worked through the Smalltalk/V tutorial handbook. The diary data disclose early misapprehensions of important Smalltalk constructs (e.g., the distinction between class and instance). They suggest an ordering of issues that arise in learning Smalltalk/V, from learning the language's syntax and precedence rules, to finding classes and methods in the hierarchy, to understanding the class/instance distinction, to understanding Model-View-Controller, to mastery of object-oriented design. They suggest an important difference in strategy between programmers and nonprogrammers; all three programmers took up the "programmer's burden." They took a depth-first approach when they encountered new methods and classes. Given the complexity of the Smalltalk/V environment, this depth-first strategy may be counterproductive, at least in the short term.

Campbell, R. L., & Bickhard, M. H. Knowing levels and developmental stages. Basel: S. Karger, 1986.

Campbell, R. L., Carroll, J. M., & DiBello, L. A. Expertise in human-computer interaction: The case for a developmental approach. Paper presented at the **Jean Piaget Society meeting,** Philadelphia, June, 1989.

Carroll, J. M., & Campbell, R. L. (in press). Artifacts as psychological theories: The case of human-computer interaction. **Behaviour and Information Technology.**

e

# Continuous Analogical Reasoning

**Thomas Eskridge**
Computing Research Laboratory
New Mexico State University
Las Cruces, NM 88003-0001

This research is aimed at developing a psychologically plausible cognitive model of analogical reasoning. This has resulted in the Continuous Analogical Reasoning theory, and its computer implementation in the ASTRA system (Eskridge 1989a,b). Continuous analogical reasoning is markedly different from other approaches to analogical reasoning in that it allows interactions between what is commonly referred to as the three stages of analogical reasoning: selection, mapping, and evaluation. In continuous analogical reasoning all stages influence the processing of the other stages. By developing a theory that accounts for all stages of analogical reasoning and their interactions, further strides in understanding how humans use analogies to reason and learn can be made.

We have selected a set of interactions to study and to model in the ASTRA program. The selection stage of analogical reasoning is affected by the mapping and evaluation stages in a manner in which the preferred source analog is one that is structurally consistent with the target and relevant to the system goals at hand. The mapping stage is affected by the selection and evaluation stages in a manner such that the correspondences produced will extend the partial mapping created during selection and will be relevant to solving the current goals of the system. Selection and mapping effect the evaluation stage by pressuring the directipn of the search through the problem space. The pressure from selection comes from the retrieval of a source analog that may cause new goals to be set for the system. The conjectures produced by the mapping stage effect the evaluation process by introducing new knowledge that must be taken into account. This knowledge, too, may cause the evaluation stage to set new goals for the reasoner.

We are in the process of completing the implementation of ASTRA, a computer implementation of the continuous analogical reasoning theory. ASTRA represents knowledge in a highly interconnected, content addressable parallel semantic network, in which semantic, schematic, and episodic memories are stored. A key feature of this representation is that it allows the connectionist ideas of distributed representation, spreading activation, and activation thresholds to be used in the process by which the three stages can interact. The mechanism used to accomplish the interaction between stages is a constrained, parallel marker-passing/spreading activation process. Activation is spread from nodes of interest along certain links specified by each of the three stages, in inverse proportion to the number of links emanating from the node. The spreading of activation ceases once the activation falls below a .preset threshold. This procedure is effective because it allows a large number of structures for analog retrieval, mapping, and evaluation to be suggested, while only the few with relatively high activation are actively pursued.

Current work on the continuous analogical reasoning theory is proceeding on three fronts: 1) acquiring evidence of the interactions between stages and determining the roles they play in analogical reasoning, especially the interactions effecting the retrieval of a source analog, 2) completing the implementation of the continuous analogical reasoning in the ASTRA program, and 3) using the ASTRA system to text hypotheses concerning analogical reasoning and as a method of increasing the flexibility of automated problem solvers and planners in unstructured task environments.

Eskridge, T.C., Principles of Continuous Analogical Reasoning, to appear in **Journal of *Theoretical* and Ezperimental Artificial Intelligence, *1*,3,** 1989a.

Eskridge, T.C. Continuous Analogical Reasoning: A Summary of Current Research, in **DARPA Workshop on Case-Based Reasoning,** pp. 253-257, 1989b.

# Explanation Through Physical Analogies

**Brian Falkenhainer**
Xerox **Palo** Alto Research Center
**3333** Coyote **Hill Road**
**Palo Alto, CA 94306**

Explanation, interpretation, and diagnosis are typically decoupled from theory formation and discovery in AI. Integration of these intimately related **tasks** into **a** unified view of explanation offers the potential for graceful degradation in the presence of an imperfect domain theory: provide **a** deductive explanation if possible and produce novel hypotheses where and when necessary.

This work presents an analogical approach to the problem for the task of constructing qualitative explanations of observed physical phenomena. It portrays explanation **as a** process of forming **physical *analogies*** – viewing the situation and its behavior as similar to familiar phenomena, conjecturing that they share analogous underlying causes, and using the plausible interpretation **as a** foot-hold to further understanding, analysis, and hypothesis refinement. Importantly, it suggests that distinctions and procedural separations between explanation and analogy are superfluous (Falkenhainer, to appear) Rather, a single analogical mechanism is used which provides smoother adaptability to unanticipated or underspecified phenomena by enabling **a** guessing facility capable of drawing upon past experience and knowledge of other domains. Distinctions between deductive, abductive, and analogical explanations emerge from the evaluation of candidate explanations instead of from individual procedures.

Primary emphasis is placed on two central questions. First, how are analogies elaborated to sanction new inferences about **a** novel situation? This problem is addressed by **conteztual *structure-mapping*** (Falkenhainer, 1988; Falkenhainer, 1989), **a** knowledge-intensive adaptation of Gentner's structuromapping theory. It presents analogy elaboration **as a map and *analyze*** cycle, in which two situations are placed in correspondence, followed by problem solving and inference production focused on correspondence inadequacies. Second, how is the quality of **a** proposed analogy evaluated and used for some performance task? **A** theory of ***verification-*baaed analogical learning** (Falkenhainer, 1986; Falkenhainer, 1988) is presented which addresses the tenuous nature of analogically inferred concepts and describes procedures that can be used to increase confidence in the inferred knowledge. Specifically, it relies on analogical inference to hypothesize new theories and simulation of those theories to analyze their utility and validity. It represents **a** view **of** analogy **as** an iterative process **of** hypothesis formation, testing, and revision.

These ideas are illustrated **via** PHINEAS (Falkenhainer, 1988), **a** program which uses analogical similarity to posit qualitative explanations for time-varying descriptions of physical behaviors. With knowledge of liquid flow, PHINEAS is able to produce explanations about observations of osmosis, heat flow, and liquid flow through an unknown object.

Falkenhainer, B. (1987). An Examination of the Third Stage in the Analogy Process: Verification-Based Analogical Learning, ***IJCAI-87.*** (Also UIUCDCS-R-86-1302, Computer Science, University of Illinois, October 1986)

Falkenhainer, **B.** (1988). Learning from Physical Analogies: **A** Study in Analogy and the Explanation Process, PhD thesis (Technical Report UIUCDCS-R-88-1479), University **of** Illinois at Urbana-Champaign, December, 1988.

Falkenhainer, B. (to appear). A Unified Analogy Model of Explanation and Theory Formation, In **J.** Shrager & P. Langley (Eds.), **Computational *Models of* Scientific Discovery** (to appear).

Falkenhainer, B, Forbus, **K. D,** and Gentner, D. (1987). The Structure-Mapping Engine: Algorithm and Examples. **Artificial Intelligence,** in press. (Also UIUCDCS-R-87-1361, University of Illinois, July 1987)

# A Model of Natural Category Structure and its Paradigmatic Implications

**Douglas Fisher**
Department of Computer Science
Vanderbilt University
Nashville, TN 37235

Cognitive modeling fits computational mechanisms to the constraints of psychological data. The problem of determining a starting point for this process has been addressed by several authors. Anderson (in press) suggests a *rational* analysis, whereby a general class of behaviors (e.g., concept formation) are associated with a performance function to be optimized. The guiding assumption is that natural organisms are rational, albeit resource-bounded decision makers (Simon, 1969). A similar but less formal view is implicit in speculative analyses (Hall & Kibler, 1985), which posit high-level computational principles that constrain human processing (e.g., Kolodner, 1983).

This talk traces the development of the COBWEB concept formation system (Fisher, 1987) from rational and speculative analyses by Gluck and Corter (1985), Kolodner (1983), and Lebowitz (1982). COBWEB learns classification hierarchies over environmental observations in a manner that is guided by two principles. First, learning should be incremental: observations should be efficiently assimilated into an evolving classification scheme as they are encountered. Second, learning should benefit performance at some task; in this case, predictions about unknown properties of environmental observations. To realize these objectives, COBWEB, borrows a measure of concept quality developed by Gluck and Corter in their work on human basic *level* effects. Kolodner's CYRUS and Lebowitz's UNIMEM provide general strategies of efficient classification that we modify to exploit category utility as a guide for concept formation. This union yields a system that meets the computational objectives of efficient assimilation and accurate prediction.

Incremental learning and prediction accuracy are dimensions of computational interest, but they are also high-level constraints on much of human learning. Fisher (1988) and Silber and Fisher (1989) refine COBWEB to account for a number of psychological effects, notably basic level, *typicality* (Rosch & Mervis, 1975) and fan (Anderson, 1976) effects. In fact, the model unifies these phenomena and suggests heretofore unexplored interactions between them.

Beyond an account of basic level, typicality, and fan effects, our work speaks to three paradigmatic concerns. First, COBWEB does not strongly distinguish between *probabilistic* and *exemplar* representations (Smith & Medin, 1981). Second, COBWEB's account of basic level effects assumes a local representation (i.e., one category corresponds to one node of the hierarchy), while typicality and fan effects assume that category members may be *distributed* over many nodes of the hierarchy. Finally, principles of object concept formation are being adapted to episodic and explanation-based learning in planning and problem-solving.

Fisher, D. A computational account of basic level and typicality effects. *AAAI-88.* St. Paul, MN: Morgan Kaufmann, 1988.

Silber, J. & Fisher, D. A model of natural category structure and its behavioral implications. *Cognitive* Science Conference. Aan Arbor, MI: Morgan Kaufmann, 1989.

Yang, H. & Fisher, D. Conceptual clustering of means-ends plans. Machine Learning *Workshop.* Ithaca, NY Morgan Kaufmann, 1989.

Yoo, J. & Fisher, D. Conceptual clustering of explanations. Machine Learning *Workshop.* Ithaca, NY: Morgan Kaufmann, 1989.

# Psychological Studies of Access and Inference.


**Dedre Gentner**
Department of Psychology
University of Illinois
Urbana, IL 61801

**Ken Forbus**
Department of Computer Science
University of Illinois
Urbana, IL 61801


Analogical learning involves transferring knowledge from a familiar, well-understood situation to a new, less-understood problem. Our research seeks a computational account of analogical learning. We have found that two decompositions are necessary. First, the processes involved in analogy must be differentiated into subprocesses, since each has distinct psychological properties (Gentner, 1988). Second, similarity itself must be differentiated, according to whether surface characteristics, structural characteristics or both are shared (Gentner, 1983). Our results, and those of others, indicate that these subprocesses involve these distinct classes of similarity to different degrees. In mapping, for example, people appear to rely on structural similarity, while access appears to be influenced strongly by surface commonalities.

This talk describes experiments which explore the implications of these ideas. We begin with experiments on human subjects. Using other stimuli, (e.g., proverbs), we replicated the finding that memory retrieval is superior for matches with surface commonality, despite subjects preferring relational matches as both more similar and more sound. Furthermore, our results indicate that (1) the preference for surface commonalities over structural in access does not depend on competition between surface and structural matches and (2) the preference persists when subjects are told to look for relational matches. In experiments on mapping, our results indicate that systematicity acts as a selection constraint, predicting which predicates are matched and predicting what inferences are drawn based on the match.

We also describe two cognitive simulation efforts. The first uses SME (Falkenhainer, Forbus, & Gentner, 1986, in press) in sensitivity analyses which probe issues of structural evaluation. We describe several principles for psychologically plausible algorithms, and conjecture that naturalistic representations often include a preponderance of appearance and low-order information (the **Specificity Conjecture).** We demonstrate that these principles and conjecture strongly constrain how structural evaluations should be performed (Forbus & Gentner, 1989). Second, we describe a computational model of similarity-based reminding and inference, MAC/FAC, built using SHE. The HAC stage uses a computationally cheap, coarse filter, based heavily on local matches. The FAC stage filters the output of the HAC stage, using more expensive (but more accurate) structural similarity computations to obtain better estimates of the quality of the match and to propose new inferences. We describe our initial explorations of this model, and our plans for future experiments.


Falkenhainer, B., Forbus, K. D. & Gentner, D. (1986). "The structuromapping engine". **Proceedings *of the* Meeting *of* the American *Association* for Artificial Intelligence. 1986.** Philadelphia. Revised version to appear in **Artificial Intelligence,** in press.

Forbus, K., & Gentner, D. Structural evaluation of analogies: What counts? In **Proceedings of the Eleventh Annual Meeting *of* the Cognitive Science Society,** Ann Arbor, MI. 1989, August.

Gentner, D. Structuromapping: A theoretical framework for analogy. **Cognitive Science, 7(2), 155-170.** 1983.

Gentner, D. Analogical inference and analogical access. In A. Prieditis (Ed.), **Analogical. Los** Altos, CA: Morgan Kaufmann. 1988.

# Opportunistic Memory

**Kristian Hammond**
Department of Computer Science
University of Chicago
Chicago, IL 60637

**Colleen Seifert**
Department of Psychology
University of Michigan
Ann Arbor, MI 48104

In earlier work (Hammond 1989a), we studied **ezpectation failures** (Schaiik 1982) that corresponded to actual **plan failures.** In our current research, we are looking at expectation failures that are failures to anticipate **planning opportunities.** We argue that a planner has to respond to both types of failure by repairing its current plan and by repairing the knowledge base which allowed it to create the plan.

**Our** approach uses episodic memory to organize, recognize and exploit opportunities. Briefly the algorithm includes the following features:

- Goals that cannot be fit into a current ongoing plan are considered blocked and, as such, are suspended.
- Suspended goals are associated with elements of episodic memory that can be related to potential opportunities.
- These same memory structures are then used to "parse" the world so that the planner can make execution-time decisions.
- As elements of memory are activated by conditions in the world, the goals associated with them are also activated and integrated into the current processing queue.

Because the planner's recognition of opportunities depends on the nature of its episodic memory structures, we call the overall algorithm presented here **opportunistic memory.**

Our experimental goals are centered around discoving the strategies used by people to index unsatisfied gods in memory. As a first experimental investigation, a re-examination of the Zeigarnik effect seems appropriate. A replication would provide a baseline to examine differences in recalling pending goals based on features of the new circumstances. The next experiments examine our taxonomy of goal blockage descriptions to determine if the vocabulary adequately describes the features people seem to utilize in indexing and recalling pending gods. Finally, these experiments lead to the investigation of the property that particular features of the environment may be ignored or noticed based upon their value to pending goals.

Hammond, K., **Case-bared Planning: Viewing planning ar a memory task.** Academic Press, **1989a.**

Hammond, K., Opportunistic Memory. To appear in **Proceedings *of* the Eleventh *International* Joint Conference on Artificial Intelligence.** 1989b.

Hammond, K., Converse, T., and Marks, M., Learning from opportunities: Storing and re-using execution-time optimizations, In **Proceedings *of* the Seventh National Conference on, Artificial Intelligence.** 1988.

Seifert, C., Goals in Reminding. In **Proceedings *of* the DARPA Conference *on* Case-Based Reasoning,** Clearwater Florida. 1988.

Seifert C., McKoon, G., Abelson, R. and Ratcliff, R., Memory Connections Between Thematically Similar Episodes. In **Journal *of* Ezperimental Psychology: Learning, Memory and Cognition.** 1985.

# Conservative Language Acquisition From Positive Examples

**Shyam Kapur**
Department of Computer Science
Cornell University
Ithaca, New York 14853

Language learnability is investigated in the **Gold** paradigm of inductive inference from positive data. A Characterization of learnable families in this framework was given by Angluin (1980). This work has attracted considerable attention of linguists. It is generally assumed that natural language acquisition is a *conservative* process, in that the child does not change her or his internal model of the language, unless evidence inconsistent with it is acquired. It is also assumed that the child learns from ***positive evidence.*** While more empirical studies are needed to substantiate these assumptions beyond doubt, it is interesting to explore their consequences **from** a formal viewpoint. The implications for learnability of conservativeness and other constraints such **as** ***consistency*** (at any stage the learning process must guess a language that contains **all** the evidence seen), and ***responsiveness*** (a guess must be made for any piece of evidence) are investigated.

It is easy to show that a conservative learner must always guess a least upper bound language corresponding to any input data. Berwick **(1982)** referred to this prescription **as** the ***subset principle,*** and argued that this principle can account for several linguistic phenomena. Since then the subset principle has been the subject of extensive discussion in linguistics. However, its formal basis and significance have often been misunderstood. Contrary to what has often been assumed, the subset principle is not necessarily **a** successful learner. **A** rather general condition on a family is shown to be sufficient for learnability by the subset principle.

Learnable families are characterized for (a) consistent, responsive and conservative learners. The consequences of dropping, individually or together, the constraints of responsiveness and consistency are explored. Learnable families are characterired for a learner which is: (b) conservative and responsive, (c) conservative and consistent, **and** (d) conservative. It is shown **that** the class **of** learnable **families** strictly incresses going from (a) to (b) and from (b) to (c), while it stays the same going from (c) to (d). It is also considered how the learnability conditions can be simplified for special families of languages such **as** those that admit an effective topological sorting, and those that are linearly ordered.

**Kapur, S.,** Lust, **B.,** Harbert, **W.** & **Martohardjono,** G. On relating Universal Grammar and Learnability theory: the case of Binding Domains and the Subset Principle. ***Proceedings of the conference on Knowledge and Language,*** Groningen, Netherlands. **1989.**

# Explanation-Based Knowledge Acquisition of Electronics

**David Kieras**
Technical Communication Program
University of Michigan
Ann Arbor, MI 481091

**John Mayer**
EE and Computer Science
University of Michigan
Ann Arbor, MI 481091

In the world of electronics and other technical domains, the knowledge of the domain is often presented in training materials and in equipment documentation in the form of explanatory text. If training is to be made more effective, we need to know how complex concepts and knowledge can be acquired from text. In addition, such knowledge might lead to methods for the automatic acquisition of knowledge from text, which would greatly assist in the development of expert systems and other knowledge-based technology.

The goals of this project are to determine how the knowledge of a complex domain can be assembled from explanations presented in textual form along with the usual accompanying graphics. This work will attempt to determine (1) the properties of good explanations, which can then be used to improve training materials, and (2) the mechanisms involved, which would constitute both a cognitive theory for how such knowledge is acquired, and also the basis for automated knowledge acquisition. The project uses a cognitive science approach, involving both AI and experimental cognitive psychology activities.

John Mayer has been developing an explanation-based learning system which forms new schemas for electronic circuits from explanations, and uses these to more efficiently understand later explanations. Mayer's system is able to form schemas for basic DC vacuum-tube circuits, and then can make use of this information to understand more complex circuits such as high-voltage DC voltage regulators. Mayer's target for his dissertation work is to devise a system which can understand explanations for basic radio frequency circuits such as oscillators and simple radio receivers and forming the appropriate schemas.

The psychological research, just starting, will first determine whether the effects of having acquired previous schemas demonstrated by Mayer's system are actually manifested in human learners, and how sensitive they are to having the schematic structure of the circuits pointed out,

# ICARUS: An Integrated Cognitive Architecture

**Pat Langley**
Department of Information **and** Computer Science
University of California
**Irvine, CA 02717 USA**

**A** cognitive architecture posits invariant properties of the human information-processing system, including aspects of representation, performance, and learning. This talk focuses on *carus*, a new architecture that differs from previous theories (e.g., Anderson, 1983; Laird, Newell, & Rosenbloom, 1986) along a number of dimensions. For instance, most research in this area has focused on cognition to the exclusion of perception and execution; in contrast, *carus* explicitly models object recognition and motor control. Moreover, the architecture employs symbolic representations, but grounds its symbols in sensori-motor descriptions. In addition, *carus* differs from most cognitive architectures by emphasizing the organization and indexing of knowledge in long-term memory. Most important, the theory diverges from earlier architectures in its basic representation and processes, which borrow heavily from Fisher's (1987) work on COBWEB. Specifically, long-term memory is represented as a probabilistic ***concept hierarchy,*** the underlying performance mechanism is a form of ***heuristic classification,*** and a single learning mechanism – incremental ***concept formation*** – is the basis for all performance improvement.

Since *carus* is still in the design stage, the talk will focus on its overall structure and its three implemented components. The first of these, LABYRINTH (Thompson & Langley, 1989), addresses the recognition and acquisition of composite concepts that involve multiple parts. The second component, DAEDALUS (Allen & Langley, 1989), generates plans using an augmented version of means-ends analysis and acquires plan expertise from successful solution traces. The MAGGIE component (Iba & Langley, 1987) deals with the execution, acquisition, and refinement of motor skills. All three modules employ another algorithm, CLASSIT (Gennari, Langley, & Fisher, in press), as a subroutine for retrieving and acquiring probabilistic concepts. Future work will focus on improving and integrating these components, incorporating drives for the generation of new gods, and developing a mechanism for attention. In addition, we will examine the architecture's. ability to account **for** results from experimental psychology.

Allen, J. A., & Langley, P. Using concept hierarchies to organiee plan knowledge. ***Proceedings of the Sizth International Workshop on Machine Learning.*** Ithaca, NY Morgan Kaufmann. 1983.

Gennari, J. H., Langley, P., & Fisher, D. E. (in press). Models of incremental concept formation. ***Artificial Intelligence.***

Iba, W., & Langley, P. A computational theory of motor learning. ***Computational Intelligence,*** 3, 338–350. 1987.

Thompson, K., & Langley, P. Incremental concept formation with composite objects. ***Proceedings of the Sixth International Workshop on Machine Learning.*** Ithaca, NY Morgan Kaufmann. 1989.

# Acquisition of Schema Knowledge

**Sandra P. Marshall**
Department of Psychology
San Diego State University
San Diego, CA 02182

This presentation focuses on the acquisition of schema knowledge by studedts interacting with a computer-based instructional program. The instructional program is SPS (Story Problem Solver), a system designed to instruct students about semantic relations embedded in arithmetic story problems. The presentation of instruction was designed to facilitate students development of and enrichment of specific schemas for solving arithmetic problems.

Two central issues will be addressed in the presentation. The first is a description of the theory of schema acquisition underlying SPS and governing the model of student learning. The second is a detailed report about students' acquisition of knowledge over five instructional sessions (spanning about 2-3 weeks). The data were gathered from students' responses while interacting with SPS, from responses by the students in brief clinical interviews following each session, and from experimental tasks posed to the students at the end of the interviews.

# Integrated Learning of Explanatory and Nonexplanatory Information

**Raymond J.** Mooney
Department of Computer Sciences
University of Texas
**Taylor Hall 2.124**
Austin, **TX** 78712

Woo-Kyoung Ahn
Department of Psychology
University of Illinois
603. E. Daniel St.
Champaign, **IL 61820**

Similarity-based learning (SBL) and explanation-based learning (EBL) constitute major research areas in machine learning. There is a relatively long history of SBL research in psychology and recent interest in the effect explanatory knowledge has on learning. Our initial experiments found that with sufficient background knowledge, subjects could learn a plan schema from a single instance, like an EBL system. However, even when subjects did not have sufficient knowledge to completely learn a concept from one instance, they still attempted to use existing knowledge. For example, when given a single example of a ceremony from an alien culture, features which subjects could somehow "explain" were immediately assumed to be relevant.

Many natural concepts have some aspects which can be explained using causal or intentional theories as well as others which cannot be explained. For example, many artifacts, like a cup, have some features with a clear functional purpose as well as conventional or aesthetic attributes, such as those which allow one to distinguish a wine glass from a beer mug. We believe that the acquisition of such concepts requires the integration of EBL and SBL.

To study this problem from a psychological perspective, we explored the learning of the Potlatch ceremony conducted by Indians of the American Northwest. If one has cultural knowledge about Northwest Indians, many of the components of this ceremony can be understood as a plan to increase the social status of the host. However, many other components are ritualistic and cannot easily be explained. We hypothesized that explanatory aspects are learned after one example and that multiple examples are required to learn nonexplanatory aspects. Experimental data support this conclusion and show that subjects have more confidence in explanatory aspects.

To study the problem from a machine learning perspective, we designed and implemented a learning method called Induction Over the Unexplained (IOU). IOU uses standard explanation-based learning techniques to learn the explainable aspects of a concept from a single example. The unexplainable features of the first and all subsequent examples are passed on to a standard inductive learning system which gradually acquires the nonexplanatory aspects of the concept as more and more examples are encountered.

Ahn, **W.**, Mooney, R.J., Brewer, **W.F.**, DeJong, **G.F.**, "Schema Acquisition **from** One Example: Psychological Evidence for Explanation-Based Learning," ***Proceedings of the Ninth Annual Conference of the Cognitive Science Society,*** Seattle, WA, **July 1987.**

Ahn, **W.** and Brewer, **W.**, "Similarity-Based and Explanation-Based Learning of Explanatory and Nonexplanatory Information," ***Proceedings of the Tenth Annual Conference of the Cognitive Science Society,*** Montreal, Quebec, August **1988.**

Mooney, R.J. and Ourston, D., "Induction Over the Unexplained: Integrated Learning of Concepts with Both Explainable and Conventional Aspects," ***Proceedings of the Sizth Internotional Workshop on Machine Learning,*** Ithaca, N.Y., June **1989.**

# A New Computational Model of Language Acquisition

**Sheldon Nicholl**
Department of Computer Science
University of Illinois
Urbana, IL 61801

This work describes a computational model of language acquisition. The topic is of great influence and importance: no grammatical theory can be viable if it is unlearnable, no matter how great its other attributes might be. It should come as no surprise, then, that the literature on language acquisition is vast. But the literature on computer models of language acquisition is almost indiscernibly minute: the major references are Anderson (1977), Berwick (1985), Langley (1982), and Selfridge (1986). There are at least two reasons for this. First, a computer model is necessarily a performance model, not a competence model, and second, no computer has access to the rich perceptual world known to every child. Workers using computers are therefore put under a double burden not shared by others.

In contrast to previous research, my work postulates that language learning requires two stages of processing and these stages use representations of different expressive power: propositional logic and context free grammars. I assume that the output of perception is fed into a "propositional learning algorithm" adapted to natural language. Within machine learning, a propositional or attribute-based learning algorithm is in general a procedure for forming rules expressed in propositional logic; see Quinlan (1986), and for another view see Rendell (1986).

The principal thesis of my work is that language learning requires two stages of processing and that all morphological and inflectional processes in language are governed by rules which are created by the propositional learning algorithm. The empirical evidence for my central thesis is a computer program that can learn certain grammatical agreement rules, like Subject-Verb agreement, by using a propositional learning algorithm. My program is new because previous programs have not addressed agreement, a major syntactic phenomenon in its own right, and because this is the first time that propositional learning has been integrated into a language acquisition system. This work shows for the first time that the study of language acquisition has properties in common with the rest of the field of machine learning at large.

In addition to the empirical evidence, my central thesis is supported by the following theoretical argument. Morphological and inflectional processes are all closed systems allowing only a finite number of possible combinations. Since the number of combinations is finite, propositional logic is sufficient; i.e., more powerful languages for expressing rules, like predicate logic, are unnecessary. This is desirable, because as a language increases in expressive power, it becomes much more difficult to learn. This does not mean that propositional rules can cover everything: when, and only when, the propositional learning algoritlim fails, the perceptual data are passed on to other more powerful rule-forming systems. Hence the two stages. Syntax, for example, is an open system that apparently allows an unbounded if not infinite number of combinations. Propositional learning must therefore fail for syntactic phenomena. These will then get passed on to other rule-learning systems, including one for creating context-free rules.

Nicholl, S. (1989) Computer acquisition of grammatical categories: a review. Draft manuscript.

Nicholl, S. (1988) Language Acquisition by a Computer Program based on Cognitive Universals. Proposal for PhD Preliminary Examination, unpublished. 41 pages.

Nicholl, S. (1988) Language Acquisition by a Computer Program based on First-Order Logic. MS Thesis, Department of Computer Science. UIUC Report No. UIUCDCS-R-88-1423. 137 pages.

Nicholl, S. and D. Wilkins (1989) SLANG: Semantics-directed LANGauge acquisition with induction rules. Draft manuscript.

# The Function of Conceptual Understanding In the Learning of Arithmetic Procedures

Stellhan Ohlsson and Edward Rees
Center for the Study of Learning
Learning Research and Development Center
Pittsburg, PA 15260

School children learn arithmetic procedures by rote, rather than by constructing them on the basis of their understanding of numbers. Rote learning produces lack of flexibility, nonsensical errors, and other difficulties in learning. Mathematics educators have proposed that if arithmetic procedures were constructed under the influence of conceptual understanding of the principles of arithmetic, then procedure acquisition would not suffer from these difficulties. However, little effort has been investigated in conceptual analysis of this hypothesis, or in proving its viability. We propose a theory of conceptual understanding and its role in the learning and execution of arithmetic procedures. The basic hypothesis of the theory is that principles constrain the possible states of affairs, and thereby enable the learner to monitor his/her own performance and to correct his/her errors. We propose a new knowledge representation, the **state constraint,** which captures this view of principled knowledge. The state constraint theory has been implemented in the Heuristic Searcher **(HS),** a computer model that learns arithmetic procedures on the basis of general principles encoded as constraints on search states. We have simulated (a) the discovery of a correct and general counting procedure in the absence of either instruction or solved examples, (b) flexible adaptation of an already learned counting procedure in response to changes in the task demands, and (c) the correction of errors in multi-column subtraction in the absence of external feedback. The state constraint theory provides novel answers to several questions with respect to conceptual understanding in arithmetic, generates counter-intuitive but testable predictions about human behavior, deals successfully with technical issues that cause difficulties for other explanations of the function of knowledge in learning, and fares well on evaluation criteria such as generality and parsimony. The state constraint theory is incomplete; it does not explain how procedure acquisition proceeds in the absence of conceptual understanding, or how learners overcome errors that can not be described as violations of principles. Future work will focus on the question of how knowledge and experience interact in procedural learning.

**Ohlsson, S.** and Rees, E. (1988). An Information Processing Analysis of the Function of Conceptual Understanding in the Learning of Arithmetic Procedures, Technical Report No. KUL-88-03, **LRDC,** August.

# Abduction and Learning Involving Physical and Psychological Explanations

**Paul O'Rorke**
Department of Information and Computer Science
University of California, Irvine
Irvine, CA 92717

C. S. Peirce (1839-1914) coined the term *abduction* for a certain kind of explanatory hypothesis generation. AI researchers use the term loosely so as to include evaluation as well as construction of explanations. At the University of California, Irvine, several different projects sharing a common frainework of domain-independent methods for integrating abduction and learning (O'Rorke 1988) are underway in specific domains. One project involves physical explanations, another focuses on psychological explanations. Information about human performance plays a role in each project, and we believe our work will have implications for understanding the relationships between explanations' and learning in human beings.

Our work on the role of physical explanations in learning involves case studies of scientific revolutions. O'Rorke, Morris, and Schulenburg (1989) suggests that abduction is a key to "world model revisions" — dramatic changes in systems of beliefs such as occur in children's cognitive development and in scientific revolutions. The paper describes a model of belief revision based upon hypothesis formation by abduction. It argues that when a contradiction between an observation and an existing model or theory about the physical world is encountered, the best course is often simply to suppress parts of the original theory thrown into question by the contradiction and to derive an explanation of the anomalous observation based on relatively solid, basic principles. This process of looking for explanations of unexpected new phenomena can lead by abductive inference to new hypotheses that can form crucial parts of a revised theory. As an illustration, the paper shows how one of Lavoisier's key insights during the Chemical Revolution can be viewed as an example of hypothesis formation by abduction.

Our work with psychological explanations focuses on learning to recognire plans involving affect (O'Rorke, Cain, and Ortony 1989). Motivation analysis and plan recognition, the task of understanding the mental states underlying observed actions, requires knowledge about the causal relationships between emotions and actions. Emotions serve to focus the recognition process on specific actions or events when people select particular plans of action based upon their emotional state, e.g. when someone runs away in fear or strikes someone in anger. We have built a system that uses knowledge of plans and emotions to construct explanations of actions. The system learns new recognition rules based on these explanations.

O'Rorke, P. Automated abduction and machine learning. In DeJong, G. (Ed.) *Proceedings of the AAAI Spring Symposium on Ezplanation-Based Learning.* 1988.

O'Rorke, P., Morris, S. & Schulenburg, D. Abduction and world model revision. In Ohlson, G. and Smith E. (eds), *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society,* Hillsdale, New Jersey: Lawrence Erlbaum Associates. 1989.

O'Rorke, P., Cain, T., & Ortony, A. Learning to recognize plans involving affect. In Segre, A. (Ed.) *Proceedingr of the Sizth International Conference on Machine Learning.* Los Altos, CA Morgan-Kaufmann. 1989.

# A Computational Model of the Influence of Prior Knowledge on the Ease of Concept Acquisition

**Michael J. Pazzani**
Department of Information and Computer Science
University of California, Irvine
Irvine CA 92717

We report on a series of experiments with human subjects that investigate the role of prior knowledge in learning. These experiments demonstrate two influences of prior knowledge:

- Subjects converge on an accurate hypothesis more quickly when a concept is consistent with prior knowledge.
- When several hypotheses are consistent with the set of training examples, subjects prefer a hypothesis consistent with prior knowledge.

Most early work on the ease of concept acquisition focused on the syntactic form of concept description. For example, it has been reliably found that conjunctive concepts are easier for human subjects to learn than disjunctive concepts. Here, we report some conditions under which this finding may not be true. In particular, we demonstrate that the prior causal knowledge of subjects can influence the rate of concept learning. Subjects were required to learn when a balloon could be inflated. The training examples are photographs of people with balloons. The photographs differ according to the size and color of the balloon, the age of the actor, and the action being performed (stretching the balloon or dipping the balloon in water.

We report on an experiment that indicates that disjunctive concepts that are consistent with prior knowledge (e.g., a balloon can be inflated if it is stretched **OR** if the actor is an adult) take fewer trials to learn than conjunctive concepts that are not consistent with prior knowledge (e.g., a balloon can be inflated if it is small **AND** red). A second experiment finds no difference between consistent conjunctive (e.g., a balloon can be inflated if it is stretched **AND** if the actor is an adult) and consistent disjunctive concepts. Finally, a third experiment demonstrates that when several hypotheses consistent with the data are of equal complexity, subjects prefer the hypothesis that is consistent with prior knowledge.

Purely empirical learning algorithms cannot account for the influence of prior knowledge in these experiments. Current explanation-based methods also cannot account for this learning task because they require that the domain knowledge be complete and consistent. Clearly, if domain knowledge facilitates the learning of both a conjunction and a disjunction of relevant factors, the domain knowledge cannot be complete and consistent. We postulate that subjects make use of a weaker form of domain knowledge that can be described as a set of influences of relevant factors. However, the domain knowledge does not indicate how these factors interact. We describe a program called PostHoc that makes use of such a domain theory and demonstrate how PostHoc can account for the experimental findings. PostHoc uses its domain theory to generate a plausible hypothesis for an observed outcome. The domain knowledge is also used to revise a hypothesis that fails to make accurate predictions of further observations. When an accurate hypothesis cannot be formed by combinations of known influences, additional factors are considered. In this cases, additional observations are used to revise or extend the hypothesis proposed by the domain knowledge.

Pazzani, M. & Schulenburg, D. The influence of prior theories on the ease of concept acquisition. ***Proceedings of the Eleventh Annual Conference of the Cognitive Science Society,*** **1989.**

# The Explanation of Programming Examples

**Peter Pirolli and Margaret Recker**
Graduate School of Education
4633 Tolman Hall
University of California, Berkeley
Berkeley, CA 94702

The ultimate goal of our project is to develop a model of knowledge acquisition and transfer that occurs across a fairly typical lessons on programming. Our current studies focus on learning in a lesson on recursive functions which takes place in a longer sequence of instruction on LISP programming. A typical programming lesson involves reading a text or listening to an instructor on some novel topic and then working through a set of relevant exercise problems. Typically the text or instructor will discuss somt illustrative examples to facilitate learning. In this learning situation, the learner actively constructs representations of texts and examples based on prior knowledge. This produces a set of example encodings and other relevant facts and principles that are stored as declarative knowledge in the learner's memory. Upon encountering a partially novel problem, the learner will use as much of her existing domain-specific skill as possible. At problem-solving impasses, in which no previously acquired skills are applicable, the learner resorts to weak-method problem solving. These methods operate on the declarative knowledge acquired from texts and examples. Knowledge compilation mechanisms summarize each novel problem-solving experience into new domain-specific skills.

We have constructed production system models of the cognitive skills acquired by subjects in learning recursion and have derived simple mathematical models to capture transfer effects as subjects work their way through programming problems. Similarly, production system models of analogy from presented examples have been developed and used to capture how examples modulate the acquisition and learning rates on cognitive skills. Our current work focuses on how subjects explain examples to themselves, and how these explanations interact with subsequent learning.

Protocol analyses reveal that skill acquisition in a programming lesson is correlated with the quantity and kinds of elaborations made by subjects when they initially try to comprehend instructional examples. We are currently working on a model of example explanation in the SOAR architecture. Example explanation is taken to be a process of search in a problem space in which the goal is to generate an explanation stiucture that satisfactorily interconnects the example to its intended purpose, to already acquired domain knowledge, and to new concepts, facts, principles, etc. that have just been introduced in a lesson. Our presentation will focus on analyses of subjects' explanation protocols and our first approximations of this data in SOAR.

Pirolli, P. and Bielacsyc, K. (in press). Empirical analyses of self-explanation and transfer in learning to program. Proceedings of the Annual Conference of the Cognitive Science Society.

# Reinterpretation and the Perceptual Microstructure of Conceptual Knowledge

**Jeff Shrager**
Xerox **Palo** Alto Research Center
**3333** Coyote **Hill Road**
**Palo** Alto, CA

We previously proposed that theory change involves a cognitive mechanism called "View Application" whose role is to **reinterpret one's knowledge in terms *of* newly uncovered abstractions** (i.e., "views"). Implementing View Application in a symbolic representational framework leads to two problems: **The Paradoz *of* Recognition:** How can views containing novel terms and relations be recognized **as** applicable to the current domain, if some of those terms and relations are not **already** available in the learner's current theory? **The Framework Alignment Problem:** How can semantic contact be made between terms and relations in the learner's current theory and those in a novel view without common terms shared between theory and view, or rules **of** translation between terms in the theory and those in the view?

These difficulties stem from thinking of views and theories in the form of models composed of categorical terms and relations. In this research we propose a new theory of "grounded representation" which resolves these problems. This theory rests upon the the claim that **perception and perceptual ezperience *form* the basis *of* conceptual knowledge.** More specifically, we replace symbolic representation in frames, views, scripts, etc, with a set of **"synchronization routines"** that mediate between traces in one modality and traces in another **(or** the same) modality. "Knowledge" thus consists of ***skills*** of ***identifying*** (and often **naming)** relevant features and concepts, and more importantly, **skills** for **acting** (i.e., executing appropriate actions) with respect to these entities. The basic approach to the framework alignment problem and the paradox of selection provided by grounded representation is that knowledge that is carried in different representational frameworks can be compared by understanding how they differentially interpret the experiences that compose their grounding. A central cognitive role is given to experiences themselves (or to quasi-perceptual traces of experiences themselves).

We are presently developing paradigms which will both help to reveal the specific quasi-perceptual content of conceptual knowledge, in accord with the above theory, and to provide support the theory. **I** describe **a** study of learning about laser physics (quantum optics) which serves both goals. We have **also** implemented a "qualitative" simulation of laser processes which learns about how lasers work using approximately the same information – particularly the figural information – that our experimental subjects have, and which can reason about the lasing process. **This** model contains two "working memories" in different modalities: **an** iconic (bitmap) memory in which animations take place and a "symbolic" (quasi-linguistic) memory in which explicit (rule-based) inference takes place. These are synchronized by inter-modality (inter-memory) "grounding" routines. Learning takes place by introducing routines specific to the application at hand, which serve to **label** the contents **of** the iconic memory **(by** making appropriate changes in the symbolic memory), and conversely, to make appropriate changes in the iconic memory whenever inference (or any other change in the symbolic memory) takes place.

Shrager, **J.** (in **press)** Commonsense Perception and the Psychology of Theory Formation. in Shrager & Langley (eds.) Computational Models of Discovery and Theory Formation. To appear from Lawrence Erlbaum; Hillsdale, **NJ.**

Shrager, **J. (1987)** Theory Change via View Application in Instructionless Learning. Machine Learning, **2: 247-276.**

Shrager, **J. (1989)** Reinterpretation and the perceptual microstructure of conceptual knowledge: Cognition considered **as** a perceptual **skill.** Proceedings of the Cognitive Science Society Conference. Lawrence Erlbaum Associates, Hillsdale, **NJ.** (to appear)

Shrager, **J. &** Klahr, **D. (1986)** Instructionless learning about a complex device: the paradigm and observations. Int. **J. of** Man-Machine Studies, **25, 153-189.**

# Learning Events in Three Skills

**Kurt Van Lehn**
Depts. of Computer Science and Psychology
Carnegie Melllon University
Pittsburgh, PA 37235

According to current theories of cognitive skill acquisition, new problem solving rules are constructed by proceduralization, production compounding, chunking, syntactic generalization, and a variety of other mechanisms. All these mechanisms are assumed to run rather quickly, so a rule's acquisition should be a matter of a few seconds at most. Such "learning events" might be visible in protocol. data.

This talk discusses a method for locating the initial use of a rule in protocol data. The method is applied to protocols of subjects learning three tasks: a river crossing puzzle, the Tower of Hanoi, and a topic in college physics. Rules were discovered at the rate of about one every half hour. Most rules required several learning events before they were used consistently, which is not consistent with the one-trial learning predicted by explanation-based learning methods. Some observed patterns of learning events were consistent with a learning mechanism based on syntactic generalization of rules. Although most rules seem to have been acquired at impasses-occasions when the subject does not know what to do next–there were clear cases of rules being learned without visible signs of an impasse, which does not support the popular hypothesis that all learning occurs at impasses.

VanLehn, K., Learning events in the acquisition of three skills. In Ohlson, G. and Smith, E. (eds), *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum, in press.

VanLehn, K. Discovering problem solving strategies: What humans do and machines don't (yet). In Segre, A. (Ed) *Proceedings of the Sizth Internation Conference on Machine Learning.* Los Altos, CA: Morgan-Kaufman, in press.

VanLehn, K. Learning events in the discovery of problem solving strategies. Manuscript submitted for publication, 1989.

# Automated Knowledge Acquisition Using Apprenticeship Learning Techniques

**David C. Wilkins**
Department of Computer Science
University of Illinois
**Urbana, IL 61801**

Apprenticeship is the most effective means for human problem solvers to learn domain-specific problem-solving knowledge in knowledge-intensive domains. This observation provides motivation to give apprenticeship learning abilities to knowledge-based expert systems. The paradigmatic example of an apprenticeship period is medical training. Our research investigated apprenticeship in a medical domain.

The described research illustrates how an explicit representation of the strategy knowledge for a general problem class, such as diagnosis, provides a basis for learning the domain-level knowledge that is specific to a particular domain, such as medicine, in an apprenticeship setting. Our approach uses a given body of strategy knowledge that is assumed to be complete and correct, and the goal is to learn domain-specific knowledge. This contrasts with learning programs such as LEX and LP where the domain-specific knowledge (e.g., integration formulas) is completely given at the start, and the goal is to learn strategy knowledge (e.g., preconditions of operators) (Mitchell, **1983).** Two sources of power of the Odysseus approach are the method of completing failed explanations and the use of a confirmation theory to evaluate domain-knowledge changes.

Apprenticeship learning involves the construction of explanations, but is different from explanation based learning as formulated in EBG (Mitchell, **1988)** and EBL (DeJong, **1986);** it is also different from explanation based learning in LEAP (Mitchell, **1989),** even though LEAP also focuses on the problem of improving a knowledge-based expert system. In EBG, EBL, and LEAP, the domain theory is capable of explaining a training instance and learning occurs by generalizing an explanation of the training instance. In contrast, in our apprenticeship research, a learning opportunity occurs when the domain theory, which is the domain knowledge base, is incapable of producing an explanation of a training instance. The domain theory is incomplete or erroneous, and all learning occurs by making an improvement to this domain theory.

Our approach is also in contrast to the traditional empirical induction from examples method of refining a knowledge base for an expert system for heuristic classification problems. However, with respect to the learning of certain types of heuristic rule knowledge, empirical induction over examples plays a significant role in our work. In these cases, an apprenticeship approach *can* be viewed as a new method of biasing selection of which knowledge is learned by empirical induction.

An apprenticeship learning approach, such as described in this talk, is perhaps the best possible bias for automatic creation of large 'use-independent' knowledge bases for expert systems. We desire to create knowledge bases that will support the multifaceted dimensions of expertise exhibited by some human experts, dimensions such as diagnosis, design, teaching, learning, explanation, and critiquing the behavior of another expert.

Wilkins, D.C., Knowledge Base Refinement Using Apprenticeship Learning Techniques, **Proceedings of the Seventh National Conference on Artificial Intelligence,** August, **1988, 646-651.**

Wilkins, D. C., "Apprenticeship Learning Techniques for Knowledge Based Systems," Report STAN-CS-**88-1242** and **KSL-88-14,** Department of Computer Science, Stanford University, December **1988, 153** pp.

Wilkins, D. C., "Automated Knowledge Acquisition Using Apprenticeship Learning Techniques." in ***Ma*chine Learning: *An Artificial Intelligence Approach, Volume III,*** **Y.** Kadratoff and **R.** Michalski (eds.), Los Altos, CA: Morgan Kaufmann, **1989,** in press.

Buchanan, B. G. and Wilkins, D. C. (eds.), **Readings in Knowledge Acquisition,** Menlo Park, CA: Morgan Kaufmann, to appear.

# Category Learning: The Effect of Prior Knowledge and Different Conceptual Roles

**Edward Wisniewski and Douglas Medin**
Department of Psychology
University of Illinois
**Champaign,** Illinois **61820**

In many category learning tasks, the experimenter presents examples of two or more categories, and subjects learn concepts that allow them to discriminate members of one category from those of others. This paradigm is similar to the similarity-based learning paradigm in machine learning. In similarity-based learning, a program examines a number of examples of different categories and creates generalized descriptions (concepts) of those categories. The descriptions enable the program to identify new category members.

There are at least two problems with this approach however. First, it focuses people and programs on forming concepts that emphasize only one of many roles that concepts can have (i.e., classification or discrimination). The importance of the concept is for accurately classifying category members. However, concepts must represent information about a category other than that used to identify its members. Otherwise, why have concepts? Second, the approach focuses people and programs on forming concepts that are based only on information that is explicit in the training examples. As a result, it ignores the effect of background knowledge.

This paper presents two studies that varied the roles of concepts during a classification learning task. Specifically, one group of subjects (the discrimination group) was given standard instructions to learn about pairs of categories. A second group of subjects (the goal group) was given these instructions but also was informed about the functions of the categories. The results of studies suggest that the two groups formed different concepts, even though they saw the same examples of the categories. The concepts of the discrimination group were based on those features in the examples that had predictive value–features with high cue and category validity. In contrast, the concepts of the goal group were based on predictive features and features that were important to the function of the category (the core features). Relative to the discrimination group, the goal group placed less emphasis on predictiveneis. The results are discussed in terms of their implications for standard classification tasks in psychology and explanation-based and similarity-based approaches in machine learning.

Wisniewski, E. Learning from examples: the effect of different conceptual roles. In Ohlson, G. and Smith, E. (eds), ***Proceedings* of the Eleventh Annual Conference of the Cognitive Science Society,** Hillsdale, NJ: Erlbaum, in press

Matheus, C. J., Rendell, L. R., Medin, D. L., & Goldstone, R. L. Purpose and conceptual functions: a framework for concept representation and learning in humans and machines. **In The Seventh Annual Conference of the Society for the Study of *Artificial* Intelligence and *Simulation* of Behavior,** in press.

# A Constructivist Model Of Human Learning

Ronald R. Yager
Machine Learning Institute
Iona College
New Rochelle, NY 10801-1800

Kenneth M. Ford
Division of Computer Science
University of West Florida
Pensacola, FL 32514

We will discuss a formal model of human and machine learning called participatory learning (Yager, to appear). In this model, the learner's previous beliefs play an important role in the assimilation of further information. A central aspect of the theory is the degree of compatibility (i.e., fuzzy "goodness of fit") between observations and belief. In addition, the role of arousal or anxiety (which occurs when we are continuously confronted with data that conflicts with our beliefs) is discussed.

Personal construct theory, as formulated by Kelly (1955), assumes that people typically use cognitive dimensions termed 'constructs' to evaluate their experience. In Kelly's theory, a necessary condition for organized thought and action is some degree of overlap between constructs in terms of their respective ranges of convenience. It is this overlap (or intersection) between the constructs' ranges of convenience that enables an event to be anticipated. Kelly's model of the personal scientist implies that each of us seeks to predict and control events by forming relevant hypotheses, and then testing them against available evidence. These hypotheses are derived from the specific relationships among constructs that articulate the 'logical' structure of an individual's personal construct system. In other words, as "personal scientists," we humans frequently anticipate the occurrence or non-occurrence of future events based on our willingness to project observed uniformities into the future. Thus, we continually glide from the past into the future with our previous experience preceding us – illuminating and organizing the manner in which subsequent events will be manifest to us (Ford, 1989).

The name participatory learning highlights the fact that the learner's current knowledge of the subject participates intimately in the learning process. A prototypical example of participatory learning is that of trying to convince a scientist to discard an old theory for a new one. In this situation we must relate and explain the new theory in terms of the scientist's view of the old theory. Thus the old theory (as construed by the scientist) must participate in the learning of the new theory. Central to participatory learning is the idea that an exogenous observation has the greatest impact on learning (i.e., revision of belief) when the observation is largely compatible with our present belief system. In particular, observations in conflict with our current core constructs (i.e., strongly held beliefs) are discounted. These core constructs serve as hidden hand editors; they are robust in the face of all but the strongest anomalous or discrepant feedback. Such powerful implicit feed-forward mechanisms not only help to predetermine the knowledge we construct or "discover", but also aid in maintaining and defending it. From a Kellyan vista, salient and/or massed negative feedback is the source of individual anxiety.

The formal model (discussed in the complete paper) will reflect the notion that participatory learning is optimal in situations in which it is necessary to change only a small part of the learner's current belief system. Informally, we can say that an intelligent reasoner endowed with the capacity for participatory learning employs sympathetic observations to modify itself. Occasional, very unsympathetic observations are discounted as erroneous.

Ford, K. M. A constructivist view of the frame problem in artificial intelligence. **Canadian Journal of Psychology** April, 1989.

Yager, R. R. (to appear) A model of participatory learning. ***IEEE Transactions* on Systems, Man, and Cybernetics.**

# Participants

Dr. Philip Ackerman
Dept. of Psychology
University of Minnesota
76 East River Road
N128 Elliott Ball
Minneapolis, MN 56466
(612)625-9812
eqz6511@umnacvx.bitnet

Dr. Beth Adelson
Department of Computer Science
Tufts University
Medford, MA 02166
(617)381-3616
adelsoncD.cs.tufts.edu

Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213
(412)441-4404
anderaon@psy.cmu.edu

Dr. Tom Bever
Dept. of Psychology
Heliora Hall
University of Rochester
Rochester, NY 14627
beverQpsych.rochester.edu

Dr. Gianfranco Bilardi
Department of Computer Science
Cornell University
4130 Upson Hall
Ithaca, NY 14853
(607)256-9212
bilardicDgvax.cs.cornell.edu

Dr. Deborah Boehm-Davis
Department of Psychology
George Mason University
Fairfax, VA 22030-4444
(301)961-8052
dbdaviscDgmuvax.gmu.edu

Dr. Mark Burstein
BBN Systems and Technologies Corp.
10 Moulton Street
Cambridge, MA 02138
(617)873-3861
bursteinQBBI.COM

Dr. Robert L. Campbell
IBH Thomas J. Watson Research Center
P. 0. Box 704
Yorktown Heights, NY 10698
(914)863-7722
rlc@ibm.com

Dr. Susan Chipman
Personnel and Training Research Program
Office of Naval Research, Code 1142PT
Arlington, VA 22217-6000
(202)696-4318
chipmanQnprdc.navy.mil

Kejitan Dontas
George Mason University
Dept. of Computer Science
4400 University Drive
Fairfax, VA 22030
(703)764-6057
kdontascDgmuvax2.gmu.edu

Dr. Tom Eskridge
Lockheed Austin Division
6800 Burleson Road
Dept. T4-41, Bldg. 30F
Austin, TX 78744
(612)448-5029
eskridge@austin.lockheed.com

Dr. Brian Falkenhainer
Xerox PARC
3333 Coyote Hill Rd.
Palo Alto, CA 94304
(415)494-4708
falkenhainer.pacDxerox.com

Dr. Douglas Fisher
Department of Computer Science
Vanderbilt University
Box 67, Station B
Nashville, TN 37235
(616)343-4111
dfisher@vuse.vanderbilt.edu

Dr. Kenneth D. Forbus
Department of Computer Science
University of Illinois
405 N. Uathews
Urbana, IL 61801
(217)333-0193
forbusQ.cs.uiuc.edu

Dr. Kenneth U. Ford
Division of Computer Science
The University of West Florida

11000 University Parkway
Pensacola, FL 32514
(904)474-2551
kford@uwf.bitnet

Dr. Dedre Gentner
Department of Psychology
University of Illinois
603 E. Daniel
Champaign, IL 61820
(217)333-1629
gentnerQcs.uiuc.edu

Dr. Art Graesser
Department of Psychology
Memphis State University
Memphis, TH 38152
(901)678-2742
agraesser@utmemi

Dr. Chris Hammond
Department of Computer Science
University of Chicago
1100 E. 58th Street
Chicago, IL 60637
krisQtartanrs.uchicago.edu

Dr. Yayne Iba
Dept. of Information and CS
University of California, Irvine
Irvine, CA 92717
(714)856-7210
ibaQics.uci.edu

Dr. Randolph Jones
Information and Computer Science
University of California
Irvine, CA 92717
(714)856-4196
rjones0ics.uci.edu

Mr. Karl Kadie
Department of Computer Science
University of Illinois
406 N. Mathews
Urbana, IL 61801
(217)244-1620
kadieQcs.uiuc.edu

Mr. Shyam Kapur
Department of Computer Science
Cornell University
4130 Upson Hall
Ithaca, NY 14863
(607)257-0827
kapurQavax.cs.cornell.edu

Dr. Dennis Kibler
University of California
Information and Computer Science Department
Irvine, CA 92717
(714)856-5961
kiblerQics.uci.edu

Dr. David Kieras
Technical Communication Program
TIDAL Bldg. 2360 Bonisteel Blvd.
University of Michigan
Ann Arbor, MI 48109
(313)761-7796
David,Kieras@ub.cc.umich.edu

Dr. Pat Langley
Dept. of Information and Computer Science
University of California, Irvine
Irvine, CA 92717
(714)856-1595
langleyQlics.uci.edu

Ms. Debbie Leishman
Knowledge Science Institute
University of Calgary
Calgary, Alberta
CANADA T2N 1N4
(403)220-5901
leishman@cpsc.ucalgary.ca

Dr. Sandra P. Marshall
Department of Psychology
San Diego State University
San Diego, CA 92182
(619)694-2696
q300020@calstate.bitnet

Mr. Christopher Matheus
Department of Computer Science
University of Illinois
405 B. Mathews
Urbana, IL 61801
(217)244-1620
matheus@cs.uiuc.edu

Dr. Stan Matwin
University of Ottawa
CSI Department
34 George Glinski/Private
Ottawa Ontario K1N 6N6
CANADA
(613)564-5069
stan0ca.i2.uofo.edu

Mr. John Mayer
University of Michigan
703 Church Street
Ann Arbor, MI 48104
(313)662-2846
jhm@c2en.engin.umich.edu

Dr. Douglas L. Medin
Dept. of Psychology
603 East Daniel
University of Illinois
Urbana, IL 61801
(217)333-7762
d-medin@h.psych.uiuc.edu

Dr. Alan L. Meyrowitz
Office of Naval Research, Code 433
1800 N. Quincy Rd.
Arlington, VA 22217
(202)696-4312

Dr. Ryszard Michalski
Dept. of Computer Science
George Mason University
4400 University Drive
Fairfax, VA 22030
(703)764-6259
michalskQgmuvax2.gmu.edu

Dr. Raymond Mooney
Department of Computer Science
The University of Texas at Austin
Taylor Hall 2.124
Austin, TX 78712
(512)471-9568
mooneyQcs.utexas.edu

Mr. Ken Murray
Department of Computer Science
The University of Teras at Austin
AI Lab, Taylor Hall 2.124
Austin, TX 78712-1188
murrayQcs.utexas.edu

Mr. Sheldon Nicholl
Dept. of Computer Science
University of Illinois
Urbana, IL 61801
(217)244-6077
nichollBarisia.xerox.com

Dr. Stellan Ohlsson
Learning R and D Center
University of Pittsburgh
Pittsburgh, PA 15260
(412)624-7460

Dr. Paul O'Rorke
Dept. of Information and Computer Science
University of California, Irvine
Irvine, CA 92717
(714)856-5563
ororke@ics.uci.edu

Dr. Michael J. Pazzani
Dept. of Information and Computer Science
University of California, Irvine
Irvine, CA 92717
(714)856-5888
pazzani0ics.uci.edu

Professor Peter Pirolli
Graduate School of Education, EUST Division
4533 Tolman Hall
University of California, Berkeley
Berkeley, CA 94702
(415)642-4206
pirolliQcogsci.berkeley.edu

Ms. Margaret Recker
Graduate Group in Education, EMST Division
4533 Tolman Hall
University of California, Berkeley
Berkeley, CA 04702
(416)642-4206
mimi0soe.Berkeley.edu

Dr. Allen A. Rovick
Rush Medical College
1653 West Congress Parkway
Chicago, IL 60612-3864

Dr. Colleen Seifert
Dept. of Psychology
University of Michigan
350 Packard Rd.
Ann Arbor, MI 48104
(313)763-0210
aeifert0camil.umich.edu

Dr. Jeff Shrager
Xerox PARC
3333 Coyote Hill Rd.
Palo Alto, CA 94304
(415)494-4338
shrager0xerox.com

Dr. Kurt Van Lehn
Dept. of Psychology

Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213
vanlehn@psy.cmu.educ

Dr. David C. Wilkins
Dept. of Computer Science
University of Illinois
405 North Matheus Drive
Urbana, IL 61801
(217) 333-2822
uilkinsQcs.uiuc.edu

Dr. Kent E. Williams
Institute for Simulation and Training
The University of Central Florida
12424 Research Parkway, Suite 300
Orlando, FL 32820

Dr. Edward Wisniewski
Dept. of Psychology
603 East Daniel
university of Illinois
Urbana, IL 01801
wisnieushiBvmd.cso.uiuc.edu

Dr. Ronald R. Yager
Machine Intelligence Institute
Iona College
leu Rochelle, BY 10801
(212)249-2047

e