

Planning to Learn

Lawrence Hunter
National Library of Medicine
Bldg 38A, MS-54
Bethesda, MD 20894
HUNTER@NLM.NIH.GOV

Abstract:

The thesis of this paper is that learning is planful, goal-directed activity – that acquiring knowledge is intentional action. I present evidence that learning from one's experiences requires making decisions about what is worth learning, regardless of the specific mechanisms underlying the learning or of the degree of consciousness or automaticity or level of effort of the learning. Decisions about what is worth learning are the expressions of *desires about knowledge*. I then sketch a theory of whence desires for knowledge arise, how they are represented, and how they are used. A taxonomy of learning actions is also proposed. This theory has been partially implemented in two computer models, which are briefly described.

Introduction

The central claim of this paper is that learning is planful, goal-directed activity – that acquiring knowledge is intentional action. If true, this thesis raises the relevance of both action psychology and AI planning to theories of learning. Action psychology (e.g. Tolman (1932) or Frese & Sabini (1985)) is based on the ideas that human behavior is directed towards the accomplishment of goals, that it is directed by plans, that those plans are hierarchically arranged, and that background knowledge and the environment interact in the creation and execution of plans for the guidance of action. As Frese & Sabini observe (p. xxiii) such a view is no doubt a good model of some behaviors and a poor model of others. Here I will try to demonstrate that this view plays an important role in understanding how people and machines learn from complex experiences.

Machine learning research has been dominated by the view that learning is a kind of search (see, e.g., Langley, Gennari & Iba, 1987). I believe that *planning* forms a better foundation than search for learning. In AI, planning is a set of techniques for selecting and combining actions to achieve explicit goals; see, e.g. Sacerdoti, (1971) or Charniak & McDermott (1985). The AI work originally focused on decomposing abstract goals into "primitive" actions expected to achieve the goal given specific resource limitations (such as time and energy). It has since evolved to consider issues such as uncertain environments, managing complex action over time, revising plans during their execution, taking advantage of unexpected opportunities, avoiding unexpected dangers (Birnbaum 1986), and "situating" the planner in the environment generally (Chapman 1985). The goals and actions considered in AI planning are grounded in the physical world. Domains tackled by AI planners thus far include stacking blocks, scheduling deliveries, creating recipes, suggesting battle tactics, and many others. Applying this research to the task of learning requires mapping work on planning in physical domains onto planning in mental domains. That is, a theory of planning to learn must describe the origins and nature of goals to learn, the actions that can be taken to learn, the mental and physical resources those actions contend for, the aspects of the environment that a learner must take into account when planning, and a process for selecting among and combining actions to accomplish the goals, given limited resources and the embedding of the learner in its environment. This paper sketches such a theory.

Learning Requires Decision-making

What is the evidence that learning is goal driven? There are several converging lines of argumentation. The simplest is the widespread use of goal/plan language by people discussing the acquisition of knowledge. "Inquiring minds *want to know*," "prerequisites for courses," and "research strategies" are references to goals, action preconditions and plans (respectively) for gaining knowledge. People can easily enumerate many different kinds of plans for learning, and the situations in which they are appropriate. Some general examples of common plans to learn include: ask an expert, look it up in a book, watch for it in the newspaper, try it yourself and see, wait until it happens again, or even run a scientific experiment. People are also capable of generating a variety of very specific knowledge acquisition plans for commonly occurring knowledge acquisition goal types. For example, when asked how to find Marvin Minsky's home phone number, one person suggested seven different potential plans: try Boston directory assistance, call MIT and ask in the CS department, ask some other senior AI

researcher, try the AAAI membership directory, send him email, ask the reporter who interviewed him for Science recently, or call the publisher of *Society of Mind* and ask them. He was also able to rate the likelihood of success of each of these plans, and could suggest general principles about how to get similar information in various circumstances. The widespread use of goal/plan language in describing learning, and the ability of people to easily generate plans when given statements of desired knowledge is, for now, merely anecdotal. A stronger line of argument comes from an analysis of the combinatorics of inductive learning from complex experience.

Inductive inference is limited in at least two ways. First, a recent proof by Deitterich (1989) demonstrated that learning algorithms, very broadly defined, can evaluate only a small proportion of the hypotheses compatible with the experiences they have. That is, there are far more hypotheses consistent with experience than can be distinguished among using that experience. A brief sketch of the proof is as follows: Consider inferential effort required to learn from experience, even in the drastically simplified situation assumed in the PAC-learnability literature (Valiant 1984). This simple task used for computational complexity analyses, involves a learner attempting to induce a subset h of the set of all Boolean n -tuples by processing m distinct examples, where an example is an n -tuple labelled as to its membership in h . There are $2^{(2^n - m)}$ hypotheses that are consistent with the m examples seen so far. That is, each of the $2^n - m$ possible n -tuples that are not examples could be in h or not. The size of that hypothesis space grows *double exponentially* in the complexity of the experiences. Information in the experiences grows only exponentially in their complexity. Learning by searching a space whose size is proportional to the space of hypotheses consistent with experience is therefore intractable.

The second problem with inductive algorithms stems from the fact that all known machine learning algorithms require time proportional to the number of features that can appear in their inputs, that is, the number of features they can "perceive." Learning systems that take time dependent on the number of perceptible features in the universe will be unable to account for human behavior, and are unlikely to be adequate for applying machine learning to desirable technological tasks (e.g. analyzing data from the human genome project). This is not to say that these algorithms may not prove to explain a *portion* of human learning, but alone they cannot form a sufficient theory.

To see how daunting the problem of learning algorithms that take time proportional to the complexity of their experiences is, consider the complexity of human experience. An order of magnitude estimate of the amount of afferent (incoming) information can be sketched easily: approximately 10^{10} – 10^{12} nerve cells in the body, conservatively 1% to 10% of them sensory nerves, each capable of carrying between 100 and 1000 bits per second. Combining to find a very conservative lower bound, humans routinely handle at least $10^{10} \times 10^{-2} \times 10^2 = 10^{10}$ bits per second at their sensory surfaces over their entire (waking) lifetimes, and quite possibly as much as 10^{13} or more bits per second. There is simply more information available in human-like experience than inductive algorithms can handle.

Machine learning theories of induction are not generally proposed as cognitive models of human learning generally (although they are sometimes advanced as models of category formation or categorical perception), and this is one of the reasons. The computational complexity of searching the hypothesis space generated by experiences anywhere near as complex as human experience would take computational power far exceeding even generous estimates of the computational power of the human brain. And since people use extensive background knowledge in learning, the complexity of the search space of is further increased by the interaction between experiences and all of memory. Schank, Collins and Hunter (1986) argued that inductive category formation approaches fails on other pragmatic grounds as well.

It is important to note here that parallelism does not provide a simple solution to these problems. The number of hypotheses consistent with a set of experiences grows *double exponentially* in the complexity of the experiences. Straightforward parallelization of the search through this space would require a number of processors exponential in the complexity of inputs to make this even an exponential time task. For comparison, current artificial neural net technology (e.g. Rummelhart, McClelland & the PDP group, 1986), typically uses a number of processors proportional to the complexity of the inputs; there are typically far fewer hidden nodes than input nodes in these networks, since too many hidden nodes reduces the usefulness of the learning these networks do.

Since the hypothesis space generated by human-like experience is far too large to be searched, even sublinearly, and all induction algorithms take time dependent on the complexity of their inputs, what

can people be doing when they learn from their experiences? The inescapable conclusion is that they must somehow drastically restrict the space of hypotheses that they consider during learning. How?

A key first step is the transformation of inputs to more compact representations of experience, capturing the “important” aspects and reducing the amount of “irrelevant” information. This process is hardwired perceptual processing and is likely to be automatic and fast (Fodor, 1985). However, even the transformed representational space will be quite large in systems capable of human-like behavior. It is simply the case that people are sensitive to a very large number of potentially relevant stimuli, and that this large number of “features” is overwhelming to known learning algorithms. So what can be done to restrict the size of *this* space to manageable proportions?

Existing machine learning methods have restricted the size of this space by applying *inductive biases*, e.g. Utgoff (1986), or by *a priori* limitations on the structure of the hypothesis space, through, for example, the use of decision trees or neural networks. These approaches can be considered *syntactic*, in that they constrain the *form* of the hypotheses considered, rather than their content.

I propose that the method of restricting potentially learnable hypotheses for both people and effective machine learning systems should be *content*-based. Explicit characterizations of *desirable knowledge* provide a principled method for restricting the realm of experience and background knowledge considered in learning, and thereby the size of the hypothesis space that must be considered. Having goals specifying what (kind of) knowledge is desirable provides a significant advantage for systems trying to learn from very complex experience.

Why does having explicit knowledge acquisition goals provide an advantage? The idea is to exert the broadest effective top-down constraint on the space of possible concepts to learn. Bidirectional inference, i.e. the ability to use top-down constraints (in this case, goals) as well as bottom-up information (here, processed perceptual data), is the most effective known technique for reducing the size of a space that has to be searched to find desired concepts (Birnbaum 1986).

This claim is consistent with a large body of psychological research on goal direction in selection of focus of attention, particularly from social psychology. Zukier’s (1986) review concludes: “Experimental studies have clearly demonstrated that a person will structure and process information quite differently, depending on the future use he or she intends to make of it. Information integration clearly is preceded by future-oriented decision-making processes, which guide data selection and the choice of an appropriate strategy or mode from among the several that are available,” (p. 495).

Hoffman, et al (1981) demonstrate that different goal orientations (e.g. “form an impression of a person in the following story” or “remember as much as you can from the following story.”) may influence not only to the use of different representations, but also the selection and use of different kinds of information processing. Although the goal orientations tested in that work are quite abstract, they significantly constrain the space of hypotheses consistent with the experimental materials. Srull & Wyer’s (1986) results, although divergent in important respects from those of Hoffman, et al, also provide evidence that different goal orientations have a strong effect on learning. These results bear an interesting relationship to the one of the implications that Deitterich (1989) draws from his proof about machine learning algorithms (p. 128):

[D]ifferent classes of learning problems may call for different learning algorithms. An important problem for future research is to attempt to identify relationships between types of learning problems... and types of hypothesis spaces....

That is, the combinatorics of learning require the selection of learning methods that are appropriate to particular kinds of problems, and goal orientation clearly effects the results of learning. This convergence of evidence from both psychological studies and from computational complexity analysis in machine learning suggests a hypothesis about the control of learning: *Goals about what would be desirable to learn are central to making required decisions about what and how to learn.*

Related Previous Research

Other cognitive theories have also included reference to desires for knowledge, although there are significant differences between those prior theories and the current claims. For example, consider the D-KNOW (delta-knowledge) class of goals, which are part of the conceptual dependency representation proposed by Schank & Abelson (1977). They are goals to “change knowledge state,” i.e.

to learn something. Examples of D-KNOW goals were to find out the location of food (in order to go to it and then eat it) or to find out the price of an item (in order to buy it). The generation of D-KNOW goals was always tied very specifically to a physical supergoal (e.g. satisfy hunger), and were not mentioned in the author's later theories of learning (e.g. Schank, 1982). Other theories, particularly from the animal learning psychology literature, have proposed general motivations to learn: a "will to perceive" (Thorpe), a "motivation for learning" (Thacker), and a "search by an information hungry organism" (Pribram – all reported in Livesey, 1986, p. 20-21), but these theorists did not propose any specific desires, just diffuse drives. Social psychologists have used various "goal orientations" as explanatory phenomena in theories of attention, recall and judgement, which are close in spirit to the goals to learn proposed here. However, social psychologist's goal orientations are generally specified at a very abstract level (e.g. "Form an impression," or "make predictions"), and as Zukier's (1986) review notes, "In general, however, little systematic research is available on goal orientation in inference, and no comprehensive taxonomies of 'middle-level' or concrete goals have emerged from these studies."

Also related to the current claims is the work of Horvitz, et al, (1989). They present a calculus for deciding when to do more inference (versus when to act) in medical decision making. Although based on highly idealized functions for estimating the expected value of additional inference (in their model, inference includes data gathering), it provides an attempt to model content-based decisions about *when* it is worthwhile to acquire knowledge. Although their model does not specify *what* is worth learning, it may be useful in deciding whether it is worth learning at all, potentially reducing the size of the potential hypothesis space to zero. Minton (1988) also proposes a model of judging whether it is worth learning, although his model involves computing the effect of learning on future performance *after the new concept is formed*, and is hence not useful for constraining the hypothesis space.

Both failure-driven (e.g. Schank 1982) and success-driven (e.g. DeJong & Mooney, 1986) computer models of learning posit very direct connections between experiences and (implicit) desires to learn. In these systems, the learning always takes place at the time of the failure (or success), and anything can be learned at that time is learned. A system that plans to learn may generate learning goals as a result of a success or failure, and may (or may not) be able to achieve those goals at that time. The role of failure (or success) in planning to learn systems is to identify knowledge that is worth pursuing, not (necessarily) to signal the time when knowledge can be acquired; they are failure (or success) *motivated*, not failure (or success) *driven*.

In the remainder of this paper, I sketch a theory of the origins and uses of explicit goals about what to learn. Some aspects of this theory of knowledge acquisition goals and knowledge acquisition planning are presented in greater detail in Hunter (1989, 1990a and 1990b).

Learning Goals

In order to make learning computationally feasible, learners must have goals specifying what they wish to learn, which are used to constrain the space of possibly inducible concepts. How are these goals represented? Where do they come from? How do they influence the learning process?

Desires about knowledge can be represented in at least two distinct ways. The first representational format is based on a description of the function that the desired knowledge will fulfill. These are generally stated as "desires to know how," such as the desire to know how to distinguish between mushrooms and toadstools, or how to recognize a potential good deal in the real estate ads. The other representational format is a description of the relationship of the desired knowledge to a set of existing knowledge; for example, the desire to know the capitols of all 50 states, or the names of your boss's children.

The relation-to-other-knowledge representation of learning goals is similar to Lehnert's (1978) work on representation of questions in natural language understanding. In her computer model, questions were represented by the same knowledge structures that held memories, but with some of the unfilled slots in those structures identified as the subject of a question. It is also possible to use her representational strategy for the internal representation of knowledge acquisition goals: goals can be represented as pointers to certain unfilled slots in memory structures. Ram (1989) presented a theory where relation-to-other-knowledge representations of questions were used to drive natural language understanding. Many of Ram's results apply to the design of knowledge acquisition planners generally.

Where do unfilled slots in memory structures come from? In general terms, they come from the incomplete instantiation of knowledge schemas. In order to generate relation-to-other-knowledge goal representations, a learner must have some knowledge of the structure of its knowledge. Consider a simple example: in order to represent the goal to find the capitols of all 50 states, a learner must know that states have capitols. That knowledge implies that the representation for each state will have a "Capitol" slot, and (presumably) the values for some of those slots are unfilled. Those unfilled slots can be the subject of a desire for knowledge. That is, the relation-to-other-knowledge representation of a goal to learn is the result of the application of some knowledge about the structure of knowledge to form a characterization of a gap, which is a representation of desired knowledge. Ram (1989) presents a much more detailed mechanism for generating these kinds of goals during the process of understanding natural language.

The other form of knowledge goal representation, based on the function of the desired knowledge, arise from inferences about knowledge useful for particular tasks. The knowledge underlying these inferences provide mappings from desired performance to desired knowledge. The resulting representations specify the processes in which the knowledge will be used, and the role that it will play in those processes. Another simple example: in order to do diagnosis, one must know (a) the kinds of things that can effect the behavior of a system and (b) methods for distinguishing among alternative potential causes of the to-be diagnosed behavior. When the need to diagnose, say, computer disk-drive failures arises, that high level knowledge can be used to generate goals to find out about the ways disk drives can fail and how to distinguish among them. The general knowledge must identify where in the diagnostic process the desired information will be used and for what, so that when it is found the information can be stored in the appropriate place for later use. See Hunter (1989) for detailed examples of the generation and representation of this kind of knowledge goal in a diagnosis domain.

Planning to Learn

The generation and representation of goals to learn is only the beginning of the learning process. The theoretical justification for generating them depends on their effectiveness at constraining combinatorics of learning from complex experience. I indicated briefly that learning should involve bidirectional inference: top-down, from learning goals *and* bottom-up, from experiential data. How can this be accomplished?

The idea is to use AI planning techniques for making decisions about which learning actions should be taken in what order to achieve the knowledge goals of an actor situated in the world. Generally speaking, these decisions are based on knowledge about available resources, knowledge about actions and knowledge about the current state of the world (including the actor's current knowledge state). Planning to learn is much like other kinds of planning, so here I will just try to describe the kinds of knowledge about resources, actions and states of the world that are necessary for planning about learning, rather than describe the process itself. The source of following characterizations are the computer models IVY and INVESTIGATOR. IVY was primarily an exploration in deciding what was worth learning, and INVESTIGATOR focuses on how to learn given a set of learning goals. They are described in detail in Hunter (1989) and (1990a), respectively.

Learning Actions

The actions that people take to acquire knowledge span a tremendous range, from looking up an answer in a reference book to designing and running scientific experiments. In order for a planner to select actions appropriate to goals, the actions must be annotated with the resources that they require, preconditions to executing the actions and expected outcomes of the actions (and perhaps information about possible alternative outcomes and relative probabilities of the alternatives). Here we will consider some of those actions and their representations.

In a system capable of taking a large number of possible actions, hierarchies of action classes can improve the combinatorics of the planning process. Classes of knowledge acquisition actions are, in effect, hypotheses about the component cognitive processes involved in learning. IVY and INVESTIGATOR, two computer models of planning to learn, use very different actions and learn from very different sorts of data, but their actions can nevertheless be grouped into four clearly defined classes:

- *Finding examples of specified phenomena.* This class of actions maps abstract characterizations (phenomena) to specific instances (examples). In IVY and INVESTIGATOR, these actions fall into two

subclasses: explicit data gathering and perceptual processing. INVESTIGATOR maps characterizations to instances by doing various kinds of database lookups. IVY works “perceptually,” checking for inputs that match a desired characterization while doing its main task of diagnosis. Both subclasses require as preconditions representations of the desired phenomena that can be used to acquire or recognize examples. In addition, the data gathering actions require access to the sources of data. Each particular action further specifies the general preconditions; e.g. to look up bibliographic records from Medlinetm, INVESTIGATOR must form a query in the Elhill retrieval language and be able to open a network connection to the Medlinetm server. The resources consumed by this class of actions (see below for a discussion of learning resources) are the time it takes to find the desired example, and the memory required to store the found examples. The expected time to find examples perceptually can be large (i.e. you do not know when you will find what you are looking for). The expected amount of memory required for some database searches can also be large.

- *Grouping examples.* The actions in this class create collections of related examples. Subclasses of these actions include finding similar examples (using various metrics), clustering examples into equivalence classes, and building hierarchical clusters. The precondition to this class of actions is a collection of examples. For example, given a genetic sequence (say, retrieved from a database) INVESTIGATOR can use sequence matching algorithms to find other genetic sequences it knows about that are similar to it. Very few resources are required for this action. INVESTIGATOR can also use Cheeseman’s (1989) Bayesian classification program Autoclass II to divide a collection of objects into clusters. That action requires significant amounts of time and CPU cycles. Some other grouping actions (e.g. hierarchical clustering) also require an applicable distance metric as a precondition.

- *Generating Abstract Characterizations of Groups.* This diverse class of learning actions includes many of the techniques traditionally associated with machine learning: concept formation, statistical analyses of collections of examples, and forming explanations of phenomena. This class of actions maps from a collection of examples and a collection of existing abstractions to a new abstract characterization of the collection of examples. INVESTIGATOR’s abstraction actions so far include an inductive category formation algorithm (which generates conjunctive definitions from groups of positive and negative examples) and ANOVA algorithms for determining the statistical features of collections of examples. These actions do not use existing characterizations: they map directly from a set of examples to an abstract characterization. Any learning method that uses domain knowledge uses both examples and existing abstractions (the domain knowledge) to form new abstractions (e.g. explanations of the examples). Although the actions in this class vary a great deal, their preconditions and expected results are similar enough so that it is possible to formulate useful planning knowledge that refers to this general class.

- *Mapping Abstract Characterizations from One Group to Another.* This class of learning actions transfers characterizations from one group to another. INVESTIGATOR currently has only one action in this class: a marker passing method for mapping a distinction in one hierarchy into another. The preconditions are two hierarchies, a distinction in one, and a mapping between the leaves of the hierarchies. This action was used to map a distinction in a taxonomy hierarchy (grouping organisms into classes) into a protein family hierarchy. The individual proteins were labelled with the organism that they came from, i.e. there was a map from the leaves of the protein hierarchy to leaves of the taxonomy hierarchy. Executing the action found protein families that were associated with specific taxa. Although not implemented in either program, this class also contains all of the learning actions involving analogy, as well as methods for mapping knowledge across dissimilar groups of examples (e.g. intersection search).

Individual learning actions are rarely able to satisfy knowledge acquisition goals; they must be assembled into sequences of actions – into plans to learn. In INVESTIGATOR, the generation of learning plans is done by top-down subgoal decomposition. Decomposition rules embody knowledge about what the various knowledge acquisition actions and classes of actions are good for. Knowledge acquisition goals are transformed into subgoals, and the subgoals are further decomposed until the process bottoms out in specific knowledge acquisition actions whose total resource consumption does not exceed preset limits, and all of whose preconditions can be satisfied. IVY did not do its own subgoal decomposition, but used programmer assembled stereotypical plans. However, INVESTIGATOR is strictly top-down, and cannot currently take advantage of unexpected opportunities, whereas IVY was able to select among potential knowledge acquisition plans based on opportunities detected during routine performance.

Work is currently underway to make INVESTIGATOR's planning more sensitive to its situation, creating a mechanism for exploring data and partial results in a more bottom-up, opportunistic fashion.

Learning Resources

With unlimited resources, planning is trivial. Unfortunately, there are always limits. Physical planners have to manage resources like energy, money and time. Planning to learn is similarly constrained, although the resources are different. In particular, learners have limitations on the amount of memory they have and on the amount of time they can spend on inference. Programs like INVESTIGATOR may also have limits on the amount of network traffic they generate. Traditional physical planning resources may also come into play, e.g. database access may cost money. Planners may have strict limits on resource consumption, or may merely try to avoid waste. INVESTIGATOR has estimates of the resources each of its actions will consume, and selects among alternative plans for accomplishing a goal by minimizing the resources consumed. It can also reject plans that exceed preset limits, e.g. would take thousands of CPU hours or gigabytes of storage.

For INVESTIGATOR, memory and CPU cycles are the constraining resources. Some of its knowledge acquisition actions are directly annotated with a formula for estimating the resources consumed. The resources consumed by others can be inferred from generalizations associated the class of which the action is a member. For example, grouping examples is assumed to take time and memory proportional to the number of examples. The Autoclass II grouping method overrides those defaults, specifying that it takes a large amount of time initially, plus time proportional to the number of examples times times the complexity of the examples times the number of expected classes). Because INVESTIGATOR tries to conserve resource consumption, Autoclass is not used unless the other grouping methods fail or unless its particular kind of output is a prerequisite for some other action.

The question of managing resources in learning raises the issue of learning over time. Existing machine learning research has focused on learning from a particular dataset. Conversely, human-like learning occurs over an entire lifetime. Learners need to decide not only whether and what to learn, but when to learn. IVY is able to keep "questions in the back of its mind," in the form of pending learning plans, which are executed as opportunities arise. A more sophisticated planner might manage a complex and interacting set of learning goals, making decisions about when to pursue a particular goal, based on its relationship to the program's other learning and performance goals and on on the current state of the world.

Conclusion: Decision-making in Learning

The space of possible lessons from experience is so large that it is combinatorically implausible to learn them all. A learner situated in a complex world must therefore make decisions about what is worth learning. The results of these decisions are explicit (although not necessarily conscious) goals about the knowledge a learner desires. Learning is not a passive process: learners act in order to learn. Their goals can be used to direct the selection of the actions taken.

Planning is decision-making based on expectations about the outcomes of actions. Effective learning decisions require knowledge about the kinds of actions that can be taken to acquire and transform knowledge, and the resources that those actions consume. Knowledge about learning actions used in planning includes information about the prerequisites for taking an action and about its expected results. Algorithms modeled on AI planners in physical domains can be used to select courses of action that can be expected to yield the desired results under resource constraints. Unlike physical planning domains, the limiting resources in learning are often inferential effort (CPU cycles for computer systems) and memory capacities.

The evidence for viewing learning as a planning process comes from both combinatoric arguments and empirical results in action psychology. This view raises a variety of issues not traditionally dealt with in the machine learning or cognitive psychology literature: How do learners come to have specific desires about knowledge? What kinds of desires to people have about knowledge? For example, can they fear specific kinds of knowledge? How do large numbers of goals to acquire knowledge interact? Can they interfere with each other the way physical goals can? How are they prioritized? The planning process raises questions of its own: How can learners recognize unexpected opportunities to learn? What are the actions that people take to learn? How are those actions organized and selected among? What do people know about the learning actions themselves, and can new actions be learned? Answers to these questions await future research.

References

- Birnbaum, L. (1986). Integrated Processing in Planning and Understanding. PhD. thesis, Yale University, New Haven, CT. (Technical Report YALEU/CSD/RR#489)
- Chapman, D. (1985). Planning for Conjunctive Goals. (Technical Report 802). Boston, MA: MIT AI Laboratory.
- Charniak, E., & McDermott, D. (1985). Introduction to Artificial Intelligence. Reading, MA: Addison-Wesley.
- Cheeseman, P., Kelly, J., Self, M., Stutz, J., Taylor, W. & Freeman, D. (1988). AutoClass: A Bayesian Classification System. in Proceedings of the Fifth International Conference on Machine Learning. University of Michigan, Ann Arbor. : Morgan Kaufman.
- Deitterich, T. (1989). Limitations on Inductive Learning. in Proceedings of the Sixth International Workshop on Machine Learning. Cornell University, Ithica NY. (pp. 125-128). San Mateo, CA: Morgan Kaufman
- DeJong, J., & Mooney, R. (1986). Explanation-based Learning: An Alternative View. Machine Learning, 1(2), pp. 145-176.
- Fodor, J. (1985). The Modularity of Mind. Boston, MA: MIT Press.
- Frese, M., & Sabini, J. (1985). Goal Directed Behavior: The Concept of Action in Psychology. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Hoffman, C., Mischel, W., & Mazze, K. (1981). The Role of Purpose in the Organization of Information About Behavior: Trait-based Versus Goal-based Categories in Person Cognition. Journal of Personality and Social Psychology, 39, 211-255.
- Horvitz, E., Cooper, G., & Heckerman, D. (1989). Reflection and Action Under Scarce Resources: Theoretical Principles and Empirical Study. (Knowledge Systems Laboratory Working Paper No. KSL-89-1). Stanford University, Stanford, CA
- Hunter, L. (1989). Knowledge Acquisition Planning: Gaining Expertise Through Experience. PhD. thesis, Yale University, New Haven, CT. (Technical Report YALEU/DCS/TR-678)
- Hunter, L. (1990a). Knowledge Acquisition Planning for Inference from Large Datasets. in Proceedings of the 23rd annual Hawaii International Conference on System Sciences, Software Track. Vol. 2 Kona, Hawaii. (pp. 35-44). Washington, DC: IEEE Press.
- Hunter, L. (1990b). Deciding What to Learn. Submitted to the Seventh International Conference on Machine Learning. Austin, TX.
- Langley, P., Gennari, J., & Iba, W. (1987). Hill Climbing Theories of Learning. Proceedings of the Fourth International Workshop on Machine Learning. Irvine, CA. (pp. 312-323). : Morgan-Kaufman.
- Lehnert, W. (1978). The Process of Question Answering. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Livesey, P. (1986). Learning and Emotion: A Biological Synthesis. Volume 1, Evolutionary Processes. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Ram, A. (1989). Question-driven Understanding: An Integrated Theory of Story Understanding, Memory and Learning. PhD. thesis, Yale University, New Haven, CT. (Tech report YALEU/CSD/RR#710)
- Rumelhart, D., McClelland, J., & the PDP Group (1986). Parallel Distributed Processing: Explorations in the Microstructure of Cognition (Volumes 1, 2 & 3). Cambridge, MA: MIT Press.
- Sacerdoti, E. (1971). A Structure for Plans and Behavior. New York, NY: American Elsevier.
- Schank, R. (1982). Dynamic Memory: A Theory of Reminding and Learning in Computers and People. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Schank, R., & Abelson, R. (1977). Scripts, Plans, Goals and Understanding. Hillsdale, NJ: Lawrence Erlbaum Associates
- Schank, R., Collins, G., & Hunter, L. (1986). Transcending Inductive Category Formation In Learning. Behavioral and Brain Sciences, 9(4), pp. 639-687.
- Srull, T., & Wyer, R. (1986). The Role of Chronic and Temporary Goals in Social Information Processing. in R. Sorrentino, & E. Higgins (eds.), Handbook of Motivation and Cognition: Foundations of Social Behavior. (pp. 503-549). Guilford, CT: The Guilford Press.
- Tolman, E. (1932). Purposive Behavior in Animals and Men. New York: Century Press.
- Utgoff, P. (1986). Shift of Bias for Inductive Concept Learning. in R. Michalski, J. Carbonell, & T. Mitchell (editors.), Machine Learning 2 Los Altos, CA: Morgan Kaufmann.
- Valient, L. (1984). A theory of the learnable. Communications of the ACM, 27(11), 1134-1142.
- Zukier, H. (1986). The Paradigmatic and Narrative Modes in Goal-Guided Inference. in R. Sorrentino, & E. Higgins (eds.), Handbook of Motivation and Cognition: Foundations of Social Behavior. (pp. 465-502). Guilford, CT: Guilford Press.

Medline is a trademark of the National Library of Medicine.