

①

ARI Research Note 87-77

DTIC FILE COPY

AD-A191 237

HILL-CLIMBING THEORIES OF LEARNING

Pat Langley, John H. Gennari, and Wayne Iba
University of California at Irvine

for

Contracting Officer's Representative
Judith Orasanu

BASIC RESEARCH LABORATORY
Michael Kaplan, Director

DTIC
SELECTED
JAN 28 1988
S D E



U. S. Army

Research Institute for the Behavioral and Social Sciences

December 1987

Approved for public release; distribution unlimited.

88 1 25 041

U. S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency under the Jurisdiction of the
Deputy Chief of Staff for Personnel

EDGAR M. JOHNSON
Technical Director

WM. DARRYL HENDERSON
COL, IN
Commanding


Research accomplished under contract
for the Department of the Army

University of California at Irvine

Technical review by

Dan Ragland

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



This report, as submitted by the contractor, has been cleared for release to Defense Technical Information Center (DTIC) to comply with regulatory requirements. It has been given no primary distribution other than to DTIC and will be available only through DTIC or other reference services such as the National Technical Information Service (NTIS). The views, opinions, and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other official documentation.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ARI Research Note 87-77	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Hill-Climbing Theories of Learning		5. TYPE OF REPORT & PERIOD COVERED Interim Report July 86 - July 87
7. AUTHOR(s) Pat Langley, John H. Gennari, and Wayne Iba		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of Information and Computer Science University of California at Irvine Irvine, CA 92717		8. CONTRACT OR GRANT NUMBER(s) MDA903-85-C-0324
11. CONTROLLING OFFICE NAME AND ADDRESS U.S. Army Research Institute for the Behavioral and Social Sciences, 5001 Eisenhower Avenue, Alexandria, VA 22333-5600		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 2Q161102B74F
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) --		12. REPORT DATE December 1987
		13. NUMBER OF PAGES 13
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE n/a
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) --		
18. SUPPLEMENTARY NOTES Judith Orasanu, contracting officer's representative		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Artificial Intelligence Computer Intelligence World Modeler CLASSIT MAGGIE		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This research note proposes "hill climbing" as a metaphor for much of human learning, and considers a number of computer systems that learn in this manner. The paper focuses on CLASSIT, a model of concept formation that incrementally acquires a conceptual hierarchy, and MAGGIE, a model of skill improvement that alters motor schemas in response to errors.		

Hill-Climbing Theories of Learning

PAT LANGLEY

JOHN H. GENNARI

WAYNE IBA

(LANGLEY@CIP.UCI.EDU)

(GENNARI@CIP.UCI.EDU)

(IBA@CIP.UCI.EDU)

Irvine Computational Intelligence Project, Department of Information & Computer Science, University of California, Irvine, CA 92717 U.S.A.

Abstract

Much human learning appears to be gradual and unconscious, suggesting a very limited form of search through the space of hypotheses. We propose hill climbing as a framework for such learning and consider a number of systems that learn in this manner. We focus on CLASSIT, a model of concept formation that incrementally acquires a conceptual hierarchy, and MAGGIE, a model of skill improvement that alters motor schemas in response to errors. Both models integrate the processes of learning and performance.

1. Introduction

Search has proved to be a powerful metaphor for understanding the nature of learning (Mitchell, 1982; Langley & Carbonell, in press). Describing a learning system in terms of its states, operators, and evaluation criteria has led to insights into learning tasks themselves and into relations between different learning methods. However, much of the search-based work on empirical (inductive) learning methods has relied on methods like depth-first search, breadth-first search, and beam search. Although these may be useful for applied learning systems, they seem implausible as models of human learning.

1.1 Hill Climbing as a Metaphor for Learning

In many domains, human learning seems to occur in a gradual, unconscious fashion. Obvious examples of this mode include concept formation, grammar acquisition, and motor learning. But even complex belief structures – such as those occurring in scientific theories – may gradually evolve in this manner. We will argue that psychological theories of such learning should be constrained along three dimensions:

- learning must be incremental; there should be no extensive reprocessing of previously encountered instances;
- the learner can entertain only one 'hypothesis' at a time; i.e., competing alternatives are not retained;
- The learner has no memory of previous hypotheses that it has held; thus, there can be no direct backtracking.

Taken together, these constraints rule out nearly all forms of search. However, there is one very weak search framework – *hill climbing* – with the requisite characteristics.

In this paradigm, one begins with some initial structure in memory, often degenerate in form (e.g., an empty decision tree). Given a new instance, the learner can modify the current structure in a variety of ways, and each choice constitutes a step through the space of structures. In order to select between the alternatives, the learner invokes an evaluation function, selecting that structure with the best score. The previous state of memory is forgotten, along with the alternative structures that were not selected.¹ This process continues as long as new instances are encountered. In some cases, a constrained state generator replaces the evaluation function, producing the new state deterministically from the current state and the instance.

This algorithm differs from standard hill-climbing methods in that steps are taken only as instances arrive, but the basic structure is the same. Thus it is subject to the same limitations, such as the tendency to halt at a local optimum. However, one does not require optimal behavior in models of human learning; one only requires them to mimic human behavior. Simon (1969) has argued that in complex domains, humans tend to *satisfice*. In this light, the limits of hill-climbing methods may be an asset.

1.2 Earlier Work in the Hill-Climbing Paradigm

Until the resurgence of machine learning research in the late 1970's, hill-climbing approaches to learning were reasonably common. For instance, the 'parameter tuning' method used in Samuel's (1963) checker player employed a form of hill climbing, and the incremental learning schemes used in neural networks can also be viewed in this light. Both classes of algorithm step through a space of numeric parameters, with the direction and amount of motion controlled by the most recent instance. There is no memory for alternative or previous states, but the states themselves are quite complex, consisting of many terms/links and their associated weights.

Winston's (1975) early work on learning from examples provides another instance of the hill-climbing paradigm. In this case, each state consisted of a complex structural definition of the goal concept, with operators for specializing and generalizing this structure. As with Samuel's system and the work on neural nets, there was no explicit evaluation function, but given a new instance the system selected a single response. Again, there was no memory for previous concept descriptions, so no explicit backtracking could occur. However, the presence of inverse operators (generalization could undo specialization and vice versa) could produce a backtrack-like effect in certain cases.

Research on grammar acquisition has also employed the hill-climbing metaphor. The best example is Wolff's (1982) SNPR system, which induced a phrase-structure grammar from sample strings. This program included operators for defining both chunks (words and phrases) and clusters (word classes). SNPR incorporated an evaluation function that measured the tradeoff between a grammar's simplicity and its 'compression' of the data. At each stage in its processing, the system defined the

¹Note that we have placed no restrictions on the complexity of the memory structures, the sophistication of the evaluation function, the power of the state generator, or whether instances are stored. The only limits involve memory for alternative states and the manner in which instances are used.

chunk or cluster that led to the best value on this criterion. Wolff's system was only semi-incremental, processing a number of strings to compute the scores for competing grammars. These grammars could become quite complex but, like the other systems we have examined, it stored only one such structure in memory at a time.

We can contrast hill-climbing theories of learning with methods that incorporate more memory-intensive search schemes. For instance, Mitchell (1982) describes a depth-first search algorithm for learning from examples that remembers both instances and previous states. He also describes the version space algorithm, which carries out a breadth-first search through the space of concept descriptions by maintaining a frontier of hypotheses. Michalski (1983) describes another algorithm that employs a beam search; this uses an evaluation function, but it differs from hill-climbing methods in maintaining the N best states at each level of the search. Some strength-based methods, such as those proposed by Holland (1986) and Langley (1987), come closer to the hill-climbing metaphor, but these retain competing hypotheses in memory.

In the remainder of the paper, we will present two models of learning based on the hill-climbing analogy, both drawn from the UCI branch of the World Modelers Project (Carbonell & Hood, 1986; Langley, 1986). The first involves the task of incremental concept formation, in which the learner must construct a concept hierarchy for objects it encounters in the environment. The second addresses the task of improving motor skills with practice. We close with some other instances of hill-climbing systems that operate in more symbolic domains.

2. A Model of Incremental Concept Formation

Much of the AI research on concept learning has occurred within the 'learning from examples' framework, in which a tutor presents positive and negative instances of goal concepts at a single level of abstraction. Yet we know that a human can acquire concepts in the absence of a tutor, and human memory appears to have a complex hierarchical organization. In recent years, research in conceptual clustering (Michalski & Stepp, 1983; Fisher & Langley, 1985) has responded to both these issues. However, most of this work has assumed that learning is nonincremental and that concepts are represented as necessary and sufficient conditions, neither of which hold for human concept formation. In this section we present CLASSIT, a model that acquires hierarchies of 'fuzzy' concepts using an incremental algorithm.

In the following pages we describe the system in terms of its representation of data and concepts, its mechanisms for classification and learning, and the evaluation function it employs to direct search through the space of concept hierarchies. The model borrows from earlier concept formation systems, including Feigenbaum's EPAM (1963), Lebowitz's UNIMEM (1986), and especially Fisher's COBWEB (in press).² Like its three predecessors, CLASSIT can be viewed as a hill-climbing learning system.

²We should note that CLASSIT's learning algorithm is identical to that used in Fisher's COBWEB, and that the two systems differ only in their representations and evaluation functions. Many of our ideas on hill-climbing approaches to learning emerged from discussions with Doug Fisher about COBWEB.

2.1 Representing Objects and Object Concepts

The World Modelers Project is concerned with learning in a reactive, physical environment. Thus, CLASSIT accepts input consisting of descriptions for three-dimensional physical objects. Each instance is specified as a set of cylinders having a length, radius, location, and orientation. Marr (1982) has argued that such descriptions constitute plausible output from the human vision system. This representation is heavily numeric and differs considerably from the more abstract semantic network and predicate calculus representations used by Winston (1975) and others.

For instance, our model represents a particular animal (say a cat) as a set of eight cylinders - representing the head, neck, torso, tail, and four legs. The size, shape, and orientation of a given animal are represented by 72 real-valued attribute-value pairs, with nine attributes for each cylinder. The *concept* for a cat (as distinct from a particular cat) is represented using the same attributes, but specifying the mean and variance for each attribute instead of a particular value. Some attributes will vary considerably, while others will be nearly constant; the latter can be viewed as more central (or criterial) to the concept than the former. Thus, both instances and concept descriptions are closely linked to the the sensory level.

2.2 Classification and Learning

In CLASSIT, the processes of classification and learning are intertwined; one cannot occur without the other. Concepts are organized into a *concept hierarchy*, with more general concepts on top and their more specific children below. Each time the system encounters a new instance, it sorts that instance down the concept hierarchy. At each level, it decides whether to place the instance into an existing class or whether to create an entirely new (disjunctive) class. In the former case, the attribute-values of the new instance are 'averaged into' the existing means and variances; this changes the 'definition' of the class. The instance is then compared to the children of this class and the process is applied recursively. If a new class is created, the values of the instance become the initial means of that class. Such a decision actually changes the structure of the concept hierarchy.

The model also includes operators for merging and splitting classes; these provide some ability to recover from poor hierarchies that may result from non-representative experiences early in the learning process. This gives a backtracking-like effect without the memory overhead of that mechanism. In summary, the system is incremental; it retains only one 'hypothesis' at each point in its evolution, and it has no memory of its earlier stages. The states themselves are quite complex, consisting of an entire hierarchy of complex concept descriptions. This complexity makes the notion of retaining multiple states seem implausible, and thus lends plausibility to the hill-climbing approach we have taken in CLASSIT.

2.3 CLASSIT's Evaluation Function

Most clustering systems attempt to maximize some tradeoff between within-class similarity and between-class differences. In a similar spirit, CLASSIT computes both the within-class variance W and the between-class variance B for each attribute in a potential class. These terms can be stated:

$$W = \frac{\sum_{j=1}^J n_j \sum_{i=1}^{n_j} (x_{ij} - \bar{x}_j)^2}{N - J + 1} \quad \text{and} \quad B = \frac{\sum_{j=1}^J n_j (\bar{x}_j - \bar{x})^2}{J}$$

where J is the number of classes, N is the number of instances, n_j is the number of instances in class j , \bar{x}_j is the average value of the attribute for class j , and \bar{x} is the average over all classes in the partition (a set of disjoint classes). The first measure corresponds to an attribute's *predictability* (how well it is predicted by membership in the class), whereas the second measure corresponds to an attribute's *predictiveness* (how well the attribute predicts membership in a class).

CLASSIT's evaluation function - which we call *category quality* - takes both of these terms into account, summing over all K attributes:³

$$\text{category quality} = \sum_{k=1}^K \frac{B_k}{W_k}$$

This measure lets CLASSIT find clusterings of instances that maximize within-class similarities and that minimize between-class differences. Note that the variance W for a class incorporates the number of instances in that group. Retaining this number lets the model incrementally update its means and variances (and thus category quality) as it observes new instances.

CLASSIT uses the *category quality* metric to determine which action to take at each level in the hierarchy. The system considers placing the new instance in each of the existing classes and computes the resulting score. Next it compares the best of these values to the score that would result from creating a new class containing only that instance. The program then forms that partition with the best score, generating a new 'state.' CLASSIT also uses this measure to determine when to combine and decompose concepts; Fisher (in press) provides the details of this process.

2.4 Experimental Results

We have evaluated CLASSIT's behavior under a variety of conditions. Figure 1 summarizes an experiment in which we 'defined' four classes - cats, dogs, horse, and giraffes - with different amounts of variation. The column labeled 'exact' represents runs in which all members of a class were identical, giving zero within-class variation.

³If a class has only one member, then its variance is zero and division by W is undefined. To avoid this problem, we use a minimum variance for each attribute. This parameter corresponds to the notion of a 'just noticeable difference' in psychophysics.

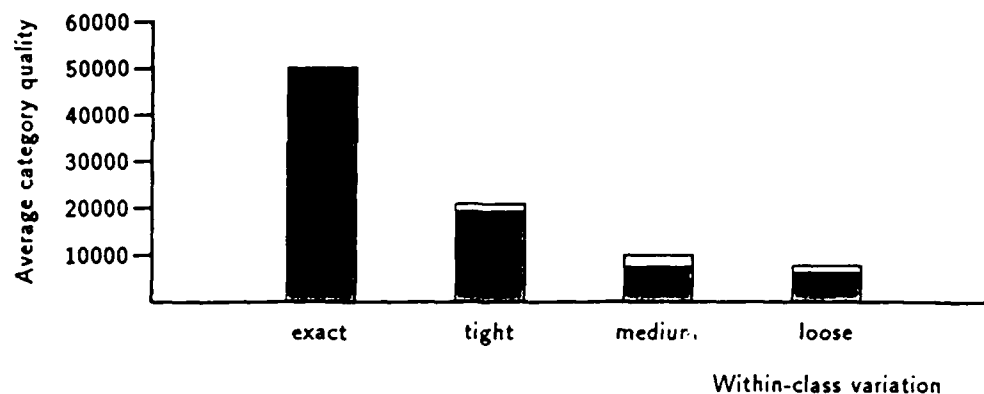


Figure 1. Category qualities resulting from different within-class variation.

The 'tight', 'medium' and 'loose' conditions introduced successively more variance. The scores for each column were averaged over ten executions, each with 30 instances that were randomly generated from means and variances for each category. The heights of the light bars indicate the average category quality of the final hierarchy after CLASSIT processed the 30th instance.⁴

The graph shows that as the regularity within each category decreases, the category quality also decreases. It is difficult to determine whether CLASSIT is actually finding the optimal clustering in each case, since this would require an exhaustive search of the clustering space. However, we have noted that as variation increases, the system's hierarchies tend to diverge from the 'desired' hierarchies used to generate the data, typically producing more than the four 'ideal' top-level categories. This suggests that CLASSIT's behavior degrades as the within-class variation increases and between-class variance decreases, as one would expect. The system performs well in orderly environments, but its ability falls off as more variation occurs.

This version of CLASSIT retains all instances it has ever encountered, storing them as terminal nodes in the concept hierarchy. Although this does not conflict with our hill-climbing philosophy, it does clash with our intuitions about human long-term memory. Thus, we have also tested a 'memory-limited' version that constrains the depth of the concept hierarchy. Naturally, this variant loses information that the unlimited-memory version retained, and this limits the extent to which the program can simulate 'backtracking' by combining and decomposing existing categories. This in turn makes the program more sensitive to the order in which it encounters instances. The heights of the dark bars in Figure 1 show the scores that result when CLASSIT retains only one level of categories. Except for the 'exact' condition, the system's behavior clearly degrades as memory limitations are introduced. However, its behavior still serves as a reasonable approximation of the original, while considerably reducing the load on memory.

⁴Since instances were created with a random number generator, different runs within a condition could produce quite different hierarchies. This was our reason for averaging across ten executions.

3. A Model of Motor Skill Improvement

Although many researchers have examined procedural learning, there has been little AI work on the improvement of motor skills. Our concern with reactive environments led us to implement a simulated jointed arm and to use this in modeling motor behavior. Below we describe MAGGIE, our model of motor skills and their acquisition. As with the work on concept formation, we have tried to remain consistent with knowledge of human behavior. Again, we begin with representational issues and then turn to problems of performance and learning. Naturally, the latter incorporates a hill-climbing approach.

3.1 Two Representations for Motor Schemas

Following Schmidt (1982), we will use the term *motor schema* to refer to some stored description of a motor skill. More precisely, we represent a schema as a temporal sequence of points (X_1, X_2, \dots, X_n) , where each point describes the location and velocity for the joints involved in the schema. Within this framework, two natural representations suggest themselves, each based on a different coordinate system.

The first scheme uses Cartesian three-space with the origin at the base (the first joint) of the arm. We will call this a *viewer-centered* representation. It corresponds to the view an agent receives as it carries out the skill. We assume that such information is available from the sensory system during execution of the motor schema. Thus, this framework can be used for recognition and monitoring purposes.

An alternative representation involves *joint-centered* descriptions, in which each joint has its own spherical coordinate system. The coordinate system for a particular joint is defined in relation to the joint to which it is connected. For instance, the coordinates for an elbow would be described in the reference frame of its associated shoulder joint. Thus, each joint has a coordinate system in which location and velocity are represented using distance from the origin, an angle of rotation about the x-axis, and an angle of rotation about the y-axis. We assume this form of information is available as proprioceptive feedback during execution; this representation can also be used to actually generate motor behavior.

3.2 Generating Motor Programs

We will assume MAGGIE has somehow acquired a viewer-centered schema that describes some desired behavior. The first step in carrying out this skill involves translating the viewer-centered description into a joint-centered representation that can be directly executed. We will not consider the details of this transformation process, but we will assume that it is serial in nature, and therefore costly. Transformations must be done for each joint in a serial manner, starting with the base joint and considering each successive joint in turn.

However, the joint-centered representation specifies only selected points involved in the skill; to actually generate behavior, one must have the desired locations and

velocities for every joint at every point in time. We will use the term *motor program* to refer to such an interpolated schema. Motor programs are not stored in memory; they are generated in real time as the skill is executed. In our model, the agent interpolates the points making up a motor program by generating a spline for each joint, connecting the sparser points in the joint-centered schema. There is evidence that humans can 'set' their limbs in desired positions even in the absence of feedback. Thus, we have not attempted to model the low level mechanisms by which an arm actually moves; it simply follows the specified motor program.

The interpolation process leads to smooth curves that cross the specified points at the desired velocities. However, the interpolated locations and velocities may be quite different from those that would result from interpolating the viewer-centered scheme. For instance, a schema for moving the hand in a straight line can be specified in viewer-centered coordinates using a few points, and splining these points would in fact produce straight line behavior. However, when MAGGIE translates this schema into joint-centered coordinates and uses splining to generate a motor program, a sequence of arcs result, with the end of each arc corresponding to a point in the motor schema.

3.3 Recovering from Errors

In other words, translating from the initial representation to an executable one can introduce errors. This means the performance system must be able to monitor its own behavior and to correct errors as they occur. In our model, this is done by generating a 'pseudo-motor program' by splining points in the viewer-centered representation and comparing these to the actual points generated as the motor program runs. MAGGIE cannot execute the pseudo-program, but it does specify the desired position at each time during execution. When the monitoring process notices a significant difference (i.e., exceeds a threshold), it invokes the error correction process.

This mechanism applies a 'burst of force' in a direction that will reduce the size of the error. The correction function has an inverted U shape, starting with minor alterations, increasing to a peak, and then decreasing to zero after a time. If the error does not increase or decrease, the path of the limb will return to the desired path after the correction process has ended. However, whether this occurs will depend on the nature of the movement. If the error had been increasing when it was detected, then undercompensation will occur. If it had been decreasing, then overcompensation will cause the arm to overshoot the mark. In such cases, the agent must reinvoke the error recovery mechanism a number of times.

3.4 Improving Motor Schemas

Although monitoring and error correction give immediate aid in carrying out desired behaviors, learning provides a longer-term solution. Although the viewer-centered and joint-centered representations lead to different interpolated behavior, one scheme can be made to approximate the other by adding selected points to the schema. For instance, one can simulate straight-line behavior with a joint-centered

schema by connecting a sequence of very small arcs. Although other forms of learning are possible, in this paper we will focus on learning by the addition of points to the joint-centered description.

We have seen that error detection invokes the error recovery process, but it also serves as the trigger for learning. Whenever the path of a joint diverges noticeably from the desired path, MAGGIE attempts to add another point to its joint-centered schema. Learning occurs only after the execution has been completed, with the location and velocity of the added point being based on the largest error that was detected during the run. Thus larger errors are reduced before smaller ones, giving a learning curve roughly similar to the power laws observed in human skill acquisition.

In this manner, MAGGIE gradually transforms its initial, sparsely defined motor schema – containing only a few points – into a more detailed schema containing many points. This incremental process continues until the monitoring can no longer detect any differences or until the addition of new points fails to improve performance. Of course, some behaviors require more learning than others; since the joint-centered representation describes arc-like motions quite well, skills involving such motions require the insertion of many fewer points.

The details of this model differ radically from our theory of concept formation, but note that the overall idea is the same. MAGGIE's schemas begin as relatively simple structures, and details are added as it gains experience in a domain. Our model of motor learning retains no memory of instances or previous schemas, nor does it maintain competing alternatives in memory. Although MAGGIE uses an intelligent generator in place of an evaluation function, it meets all the criteria set forth at the outset and constitutes another instance of a hill-climbing theory of learning.

3.5 Experimental Studies

Our model is independent of a limb's dimensions and rotational constraints, but we have tested the system using a two-jointed arm with roughly human characteristics. This includes an upper arm and a forearm, the first rotating at a shoulder joint and the second at an elbow joint. We have run a number of experimental studies with MAGGIE, all in two dimensions. For instance, we have shown that, as in humans, there is a tradeoff between the speed at which a motor skill is executed and its accuracy. We have also studied the relation between speed of execution and overcompensation effects. However, these involve the performance of the system, and our focus here lies with learning.

Naturally, we would expect that as the system detects errors and adds new points to its joint-centered schema, its errors will decrease on later executions. Figure 2 shows the results of a series of eight runs with the 'straight line' schema, indicating that the model's performance gradually improves with practice. Figure 3 presents another result that makes intuitive sense. As the skill level improves, the tradeoff between speed and accuracy becomes less evident. As more points are added to the schema, its behavior comes to approximate the desired behavior even without monitoring. This

means that MAGGIE can execute the schema at a higher speed – even though there are fewer chances for monitoring – without seriously diverging from the target path. This simulates the gradual transition of motor skills from closed-loop processing to open-loop mode, in which feedback is unnecessary.

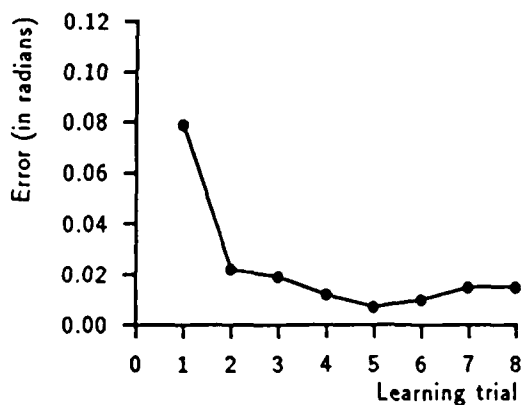


Figure 2. Error as a function of practice.

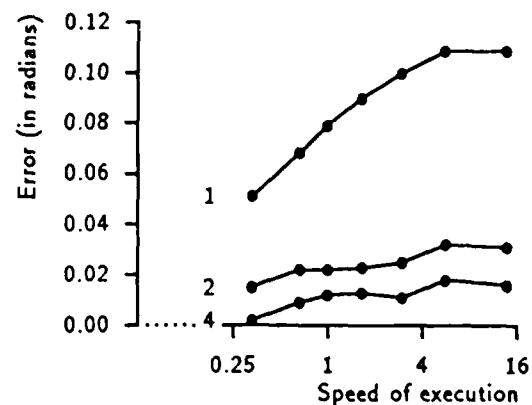


Figure 3. Speed vs. accuracy after 1, 2, and 4 learning trials.

4. Generality of the Hill-Climbing Metaphor

Both CLASSIT and MAGGIE employ low-level sensory-motor representations, since we feel such representations play an important role in learning about complex physical environments. However, the hill-climbing approach is not limited to such representations. For instance, Rose and Langley (1986) have described STAHLp, a model of scientific reasoning that incorporates techniques from belief revision and truth maintenance. The system operates in the domain of chemistry, accepting chemical reactions as input and generating componential models of various substances as output. At each point in its search, the system holds a set of beliefs that cover the known data.

Upon finding an inconsistency in its belief structures, STAHLp invokes an assumption-based reasoning technique that identifies the problematic premises and suggests changes that would eliminate the inconsistency. The program then evaluates each modification in terms of its impact on the belief system, selecting the revision that causes the least overall change. Despite its complex reasoning processes, STAHLp can be viewed as a hill-climbing learner in the sense we have defined the term. At each point, the system maintains a single 'state' in memory – its entire belief system – and when change is required, it selects a single successor state from a set of alternatives. Once the new belief system has emerged, the program has no memory for previous states or for competing belief systems.

It seems natural to associate the hill-climbing metaphor with empirical learning systems like CLASSIT and MAGGIE, but the approach can also be used within an analytic or explanation-based framework. Given a positive instance of some concept

or operator application, an analytic learning system constructs some explanation for why that instance satisfies the goal concept. Using the proof tree from the explanation, it then formulates a general rule that can be used in future cases.

Although the second step in this process (from explanation to rule) is algorithmic, the first step (constructing an explanation) can involve considerable search and can invoke heuristic techniques to evaluate the quality of competing explanations. If the learner selects only one explanation (or even a few) to transform into rules, then we have another case of hill-climbing learning. At each step, only one state exists in memory – the set of rules that constitute the compiled proofs of previous explanations. There is no memory for previous states to support backtracking, nor is there any memory of explanations that were abandoned in favor of better ones.

In summary, the hill-climbing framework extends across the traditional boundaries of machine learning. It can be applied to symbolic or sub-symbolic representations, and it can be used in conjunction with weak (empirical) learning methods or knowledge-intensive (analytic) methods. We believe that many aspects of human learning operate in this mode, and we have presented evidence – through CLASSIT and MAGGIE – that viable and interesting learning can occur in this fashion. We encourage other researchers to explicitly adopt the hill-climbing metaphor, and to explore the characteristics of this constrained but promising approach to learning.

Acknowledgements

This work was supported by Contract MDA 903-85-C-0324 from the Army Research Institute and by a gift from the Hughes Aircraft Company. Many of our ideas emerged from conversations with members of the UCI machine learning group; Doug Fisher played an especially important role in developing the notion of hill-climbing approaches to learning.

References

- Carbonell, J. G., & Hood, G. (1985). The world modelers project: Objectives and simulator architecture. In *Proceedings of the Third International Machine Learning Workshop* (pp. 14-16). Skytop, PA.
- Feigenbaum, E. A. (1963). The simulation of verbal learning behavior. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. New York: McGraw-Hill.
- Fisher, D. (in press). Knowledge acquisition via incremental conceptual clustering. *Machine Learning*.
- Fisher, D. & Langley, P. (1985). Approaches to conceptual clustering. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence* (pp. 691-697). Los Angeles, CA: Morgan Kaufmann.

- Holland, J. H. (1986). Escaping brittleness: The possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2). Los Altos, CA: Morgan Kaufmann.
- Langley, P. (1986). *A cognitive architecture for learning in reactive environments* (Technical Report 86-21). Irvine: University of California, Department of Information & Computer Science.
- Langley, P. (1987). A general theory of discrimination learning. In D. Klahr, P. Langley, & R. Neches (Eds.), *Production system models of learning and development*. Cambridge, MA: MIT Press.
- Langley, P., & Carbonell, J. G. (in press). Machine learning. In S. Shapiro (Ed.), *Encyclopedia of artificial intelligence*. New York: John Wiley & Sons.
- Lebowitz, M. (1986). Concept learning in a rich input domain: Generalization-based memory. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2). Los Altos, CA: Morgan Kaufmann.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W. H. Freeman.
- Michalski, R. S. (1983). A theory and methodology of inductive learning. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*. Los Altos, CA: Morgan Kaufmann.
- Michalski, R. S., & Stepp, R. (1983). Learning from observation: Conceptual clustering. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*. Los Altos, CA: Morgan Kaufmann.
- Mitchell, T. M. (1982). Generalization as search. *Artificial Intelligence*, 18, 203-226.
- Rose, D., & Langley, P. (1986). Chemical discovery as belief revision. *Machine Learning*, 1, 423-451.
- Samuel, A. L. (1963). Some studies in machine learning using the game of checkers. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. New York: McGraw-Hill.
- Schmidt, R. A. (1982). The schema concept. In J. A. S. Kelso (Ed.), *Human motor behavior: An introduction*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Simon, H. A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- Winston, P. H. (1975). Learning structural descriptions from examples. In P. H. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill.
- Wolff, J. G. (1982). Language acquisition, data compression, and generalization. *Language and Communication*, 2, 57-89.