Getting Intelligence into the Minds of People

Robert W. Lawler

Abstract

In conversation, Seymour Papert asked me "What's the point of studying Artificial Intelligence if not to get intelligence into the minds of people?" His question inspires my juxtaposition of explorations of Natural Learning and Constructed Personal Knowledge. Since "you can't learn about learning without learning about learning something," [1] the analyses will proceed with two examples. The first, focused on strategy learning at tic-tac-toe, concludes that learning depends on specific relationships among the elements of the context in interaction with processes of incremental cognitive change. (Even if not news, the details may be interesting.) The second analysis, focused on mastering a solution for Rubik's Cube, argues 1.) the importance of reformulation of representations as a strategy for learning in more complex situations, and 2.) that the integration of multiple modalities of representation can be a key to "getting the intelligence into the minds of people."

Introduction

In **Sciences of the Artificial**, Herbert Simon introduced his discussion of "The Psychology of Thinking" with a novel perspective on agent - context interaction:

> "We watch an ant make his laborious way across a wind- and wave-moulded beach. He moves ahead, angles to the right to ease his climb up a steep dunelet, detours around a pebble.... Thus he makes his weaving, halting way back to his home.... Viewed as a geometric figure, the ant's path is irregular, complex, hard to describe. But it's complexity is really a complexity in the surface of the beach, not a complexity in the ant.... An ant, viewed as a behaving system, is quite simple. The apparent complexity of its behavior over time is largely a reflection of the complexity of the environment in which it finds itself...." [2]

Simon concludes that his aim was exploring a similar vision of behavior, but with "human being" substituted for "ant." Over many years he pursued that aim, especially with his collaborator Alan Newell. In a broad and detailed case study of my first daughter's learning, I captured a complete history of her play at tic-tac-toe. Although simpler than Chess and various puzzles that Newell and Simon explored,[3] tic-tac-toe touched upon some of the same issues that were grand themes in the early days of Artificial Intelligence research, such as representing knowledge, strategic play, and interactions of thinking and environment. With further inspiration from psychological observations and the epistemological focus of Piaget, and from the data-orientated ecological psychologists, Barker and Wright, I was able to exploit my case study material for the computational generation of paths of learning. Together, these enable a new idea of what makes discovery-learning possible and memorable (computer modeling turns this verbal "idea" into a worked example, with process detail).

---

[1] This is an echo, of course,of Papert's well known quip that "You can't think about thinking without thinking about thinking about something."

[2] Sciences of the Artificial, 1969, MIT Press.

[3] Human Problem Solving, Prentice-Hall, 1972. See especially section 4.

Natural Learning

Many people believed, years ago, that computer modeling of student knowledge would enable the development of computer tutors sufficiently intelligent that they could replace poor or non-existent teachers.[4] That engineering effort has largely succeeded and been absorbed into the technologies of our day. The **scientific** frontier of Artificial Intelligence, in the Minsky paradigm,[5] focused on how knowledge could function (intelligence) and change (learn) through failures and repairs in articulated (specified and organized) knowledge structures.

What was most wonderful in my human case study[6] is seeing how learning can occur through serendipitous successes -- events at least as important as failures -- because such learning exemplifies learning through side effects, which may be more common for people than learning through the analysis of process failures, as in debugging.
If recall of surprising successes can be reformulated as a new plan for future action, the process can be a prototype for Natural Learning, processes by which we all learn a bit more than we knew before about the business of everyday life; in Artificial Intelligence (AI), such knowledge is called "common sense." Exhibiting such processes in extensive detail was the agenda of the project in which SLIM (Strategy Learner: Interactive Model) was developed.[7] Our effort was to use information from the child learning case study for guiding our knowledge representation, for following the development of strategies, and for examining the interplay of such development and the possibilities for learning based on specific experiences of play. The final focus was on the question of what it was that made learning possible at all, given limited initial knowledge and weak learning mechanisms.

Learning Landmarks in the Human Case: (1)

At the start of the study, Miriam played solely in a tactical way. Her moves were focused on selection of cells by preference for the center cell if available, then any corner, with side cells moves as a last resort. Although she knew forming forks was desirable (we called that "getting two ways to win"), she did no show any notion of how to achieve a fork. This changed after one specific incident:

> "I visited the Boston Children's Museum with Miriam, Robby, and his friend John. One single incident made a tremendous impression on Miriam and markedly influenced all her subsequent play. John and Miriam both played tic-tac-toe against the Museum's computer. When they fell in the pitfall of a side response to the computer's invariable middle opening, they were regularly defeated. The computer - which could be accurately described as playing with the Newell and Simon rules, always responded in the corner to their middle openings, so the children's best game was no better than a draw. Eventually Robby got a turn (the others peering over his shoulder) and opened with a corner move. The game unfolded... and Robby won. John and Miriam were astounded. Robby clearly owned some powerful knowledge.

---

[4]See the section "Intelligent Tutoring Systems," in Artificial Intelligence and Education, Vol. 1.
[5] See in "Materials Referenced," Minsky (1979, 1986), and Sussman (1975).
[6] The following summary is reported in greater detail in primary sources, "The Articulation of Complementary Roles," and "Consider the Particular Case," both online now at NLCSA.net as well as in print sources. See Materials Referenced.
[7] Undertaken with the guidance of Oliver Selfridge and some programming help by Bud Frawley at GTE's Fundamental Research Laboratory.

> Since he was glad to share that knowledge with them, that one three-corner fork entered Miriam's repertoire and exerted a dominant influence in the development of her play. Tic-tac-toe became a game of frequent choice as The Intimate Study began.... This particular form of the game is my archetypical example of strategic play, quintessentially a game length plan for the single player -- as contrasted with the single cell move preferences, either choices or responses, of tactical play...." [8]

Miriam now had a pattern oriented goal but no plan to achieve that pattern. Play with Rob degenerated to negotiation (taking turns at winning), complaining (when blocked), and cheating (to circumvent blocks). I stopped their play against each other and became Miriam's primary opponent. Her play was "egocentric," in the sense of being focused on her goal, not because she was insensitive to her opponent's possible moves, but because she did not comprehend and had no way to deal with the opponent.

Learning Landmarks in the Human Case: (2)

The second major advance was the beginning of mental play. This grew out of a practice we called "Turning the tables," thus: When we played and Miriam lost a game. I urged her to replay the same game with our roles reversed, i.e. we would replicate a game, move for move, which she would win in the replay because she was playing as the opponent to whom she had lost. Mental play began this way

> "At one point, I was called away from play by our doorbell and asked the two children not to play any games until I returned.... Coming back, I found Miriam had been playing tic-tac-toe against herself, in her words, "making smart moves for me and the other guy."

Her multi-role play worked by setting herself the strategic goal and assigning tactical play to "the other guy." This split was eventually the foundation for the next landmark.

Learning Landmarks in the Human Case: (3)

At the end of the study, Miriam invented a strategy different from a variation of her archetypical three-corners fork.

> "Reviewing her play of the intervening years, Miriam stated that she played rarely and mostly against herself (30 games in two years was her estimate). As I introduced my objective, Miriam proposed one of her own: to apply a newly discovered strategy against me. When I asked her to discuss it with me, she responded, "No way! Then whenever I play you'll know my tactics."

Miriam's new strategy, the Middle and Corners Fork, is nonetheless clear from her play on the game and... in a final comment on the discovery of this new strategy, Miriam added that she had stopped the game and only days later noticed the fork created with her third move. (This is verified by the appearance of the original game: the first six markers were made in brown crayon and the last winning move in ball pen ink.)

Conclusions for Modeling

In relating case study details to computing models, one might use those details much as one uses boundary conditions to specify the particular form of a general solution to a differential equation. So we use the psychological study here, as a foundation for the

---

[8] from"The Articulation of complimentary Roles," at http://nlcsa.net/lc2b-tis-2/lc2b-analyses/lc2ba4/

representations used in the models, and as justification for focusing on key issues: the centrality of egocentricity in cognitive self-construction and the particularity of the naive agent's knowledge.

What I drew for representation from these naturalistic observations was:
1. a decision to represent separately a goal, as a pattern, and a plan as a list
2. a willingness to see new learning as recognition of a new strategy, after the fact, by reflection on the player's own actions.
3. a recognition that mental play is necessary for analysis because simulation of "the other guy" implies the player can understand the opponent's reasons for action, as compared to some "mysterious and incomprehensible other" with whom one can cope only by social interactions or wishful thinking.

With respect to focusing on the development of a strategy, I review here the one simplest strategy to invent, to exhibit the interaction of knowledge with opportunities within the intellectual "landscape." With respect to better grasping that "landscape" of learning in this mental environment, we used a computer-embodied model to generate the suite of all games that could be played from a single opening move, then to examine those games in which the opening player experienced an unplanned win. By cataloging the sets of games which could be won without a plan, based upon a given known plan in the model's repertoire, we construct genetic pathways of plan descent.

Using Simulations to Track Paths of Learning:

Learning through Interaction

One virtue of machine learning studies is that they allow us no "miracles." They can completely and unambiguously cover some examples of learning with mechanisms simple enough to be comprehensible. Further, the computer's aid in systematically generating sets of all possible conditions helps liberate our view of what possible experiences might serve as paths of learning. When we do generate **all** possible interactions through which learning might occur, we can more thoroughly explore alternate paths and the suite of relationships among elements of the ensemble.

Strategies for achieving specific forks are the knowledge structures of SLIM (Strategy Learner, Interactive Model). Each has three parts: a Goal, a sequence of Actions, and a set of Constraints on those actions (each triple is thereby a GAC). We simulated operation of such structures in a program where SLIM plays tic-tac-toe against variations of REO (a programmed Reasonably Expert Opponent. REO is "expert" in the sense of applying uniformly a set of cell preference rules for tactical play.[9]). Applying these strategies leads to moves that often result in winning or losing; that leads to the creation of new structures, by specific modifications of the current GACs. The modifications are controlled by a small set of rules, so that GACs are interrelated by the ways modifications can map from one to another.[10]

---

[9] REO's preferences are those common in tactical play: first, win if possible; next block at need; finally choose a free cell, preferring the center cell first, and then any corner cell and finally a side cell. See Human Problem Solving, Newell and Simon (1972).

[10] Subject to limitations based on rotational symmetry and the strategies adopted by the psychological subject, the generation of possible games was exhaustive. The analysis focused on won-games in which SLIM moved first in cell 1. See figure 1. The focus on corner opening play

To evaluate specific learning mechanisms in particular cases, one must examine and specify which forks are learned from which predecessors in which sequence and under which conditions of opponent cell preferences. The simulation avoided abstraction, in order to explore learning based on the modification of fully explicit strategies learned through particular experiences.[11] The results are first, a catalog of specific experiences through which learning occurs within this system and second, a description of networks of descent of specific strategies from one another. The catalog permits a specification of two desired results: first, which new forks may be learned when some predecessor is known; and second, which specific interaction gives rise to each fork learned. The results obviously also depend on the specific learning algorithm used by SLIM.

Consider how SLIM can learn the symmetrical variation to one particular fork (refer to the cell, plan, and game representations in Figure 1, below). Suppose that SLIM begins with the objective of developing a fork represented by the Goal pattern in the set{1 3 9} and will proceed with letter moves in the sequenced plan [1 9 3] (as in Figure 1). SLIM moves first to cell 1. REO (numbers) prefers the center cell (5), and moves there. SLIM moves in cell 9, according to plan. When REO's second move is to cell 3, SLIM's plan is blocked. The strategic goal {1 3 9} is given over -- but the game is not ended. SLIM, now playing tactically with the same set of rules as REO, moves into cell 7, the corner cell remaining. Unknowingly, SLIM has created a fork symmetrical to its fork- goal. SLIM can not recognize the fork. It has not the knowledge or information to do so. What happens ? REO blocks one of SLIM's two ways- to- win, choosing cell 4. SLIM, playing tactically, recognizes it can win and moves in cell 8. This is the key juncture.

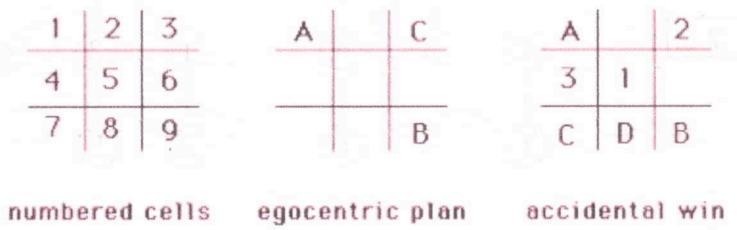

Figure 1

SLIM is programmed to recognize "winning without expecting to do so" as a special circumstance. Even more, SLIM assumes that it won through creating an unrecognized fork (otherwise REO would have blocked the win). SLIM takes the set of its first three moves as a pattern for a fork. That pattern {1 7 9} is made the goal of a new GAC. SLIM examines its known plans for creating a fork (there is one, [1 9 3]) with the list of its own moves, executed in sequence before the winning move was made [1 9 7]. The terminal step of the plan is the only difference between the two. SLIM modifies the prototype plan terminal step to create a new plan, [1 9 7]. SLIM now has two GACs, two forking strategies, for future play.[12]

---

is based on the behavior of the subject and the relative variety of strategies such openings permit.
[11]This focus of the models is precisely where the egocentricity of naive thought and cognitive self-construction of the psychological subject are embodied. The models are " egocentric" in the specific sense that no consideration of the opponent is taken unless and until the current plan is blocked. The psychological subject played this way. When a plan is blocked, SLIM drops from strategy driven play into tactical play based on preferences for cells valued by type (center, corner) and not by relation to others.

The complete set of results involves consideration of all paths of possible learning, even those deemed unlikely a priori, and concludes with the complete specification of all possible paths of learning every fork given any fork prototype. For corner opening play, the first six GACs form a central collection of strategies. Their interrelations can be represented as trees of derivation or descent (as shown in Figure 2).

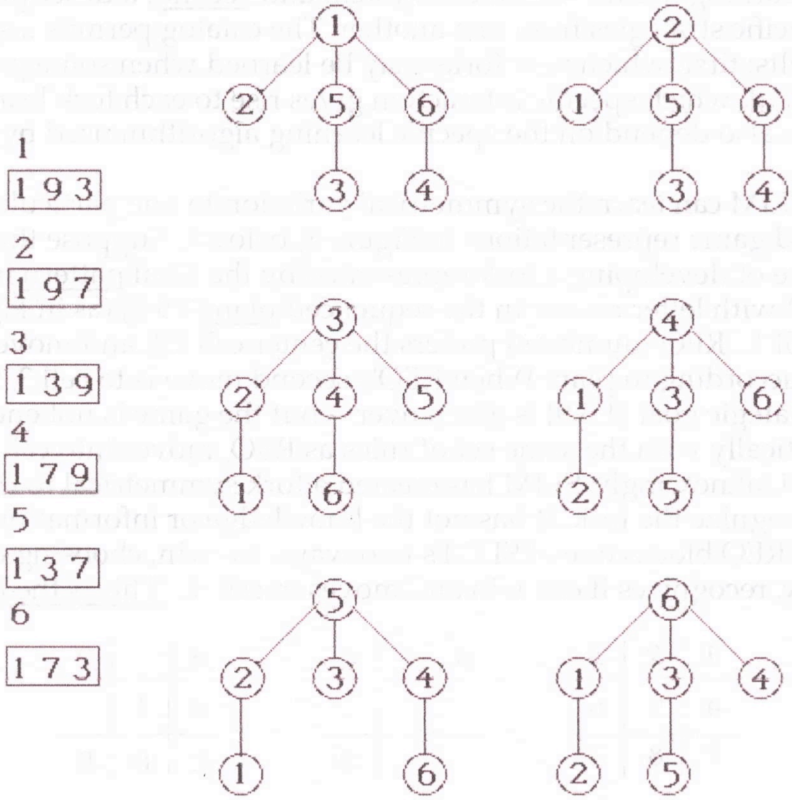

Figure 2
Plans Learnable from the Plan Specified as the Top Node of the Descent Trees

The tree with strategy three as top node (beginning with only GAC 3 known) may be taken as typical. Play in five specific games generates the other five central GACs, i.e. five central strategies. The specialness of the six central nodes is a consequence of their co-generatability. Some of those are directly generatable, can generate each other (such as GACs 1 and 2) ; they are <u>reciprocally</u> generatable. Some lead to each other through intermediaries (such as GAC 1 and 3); they are <u>cyclically</u> generatable. For these six central strategies, the trees of structure descent can fold together into a connected network of descent whose relations of co-generativity are shown in Figure 3.

[12] This form of learning by modifying the last term of a plan is one of two sorts; the other involves generating two possible plans based on deletion of the second term of a prototype plan. We know the forks achieved by plans [1 9 3] and [1 9 7] are symmetrical. SLIM has no knowledge of symmetry and no way of knowing that the forks are related other than through descent, i.e. the derivation of the second from the first. This issue is discussed in the longer versions of this text. See "On the Merits of the Particular Case."

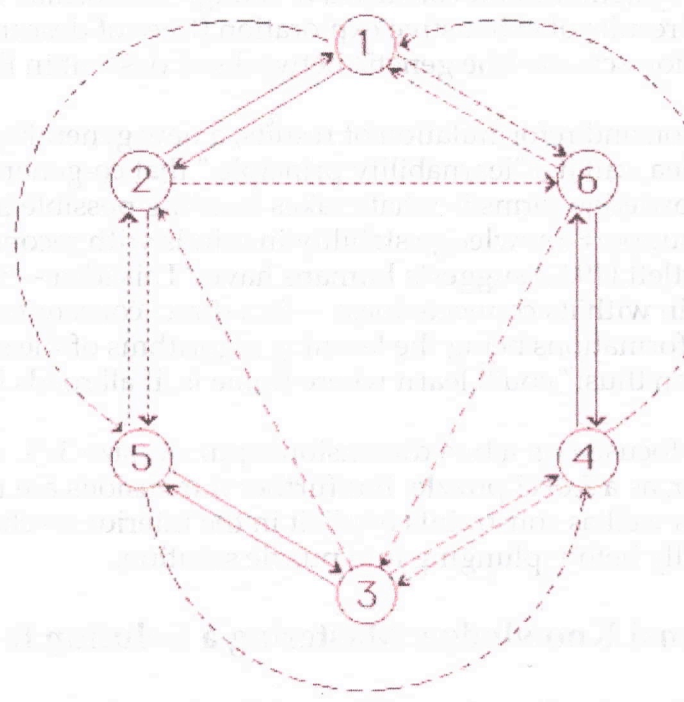

Figure 3

The form of these descent networks is related to symmetry among forking patterns. But they include more: they reflect the play of the opponent, the order in which the forks are learned, and the specific learning mechanisms permitted in the simulations. These descent networks are summaries of results.

Reflective Summary

Working with SLIM started with the general principle that learning happens through interaction. The model was constructed to represent the behavior of both the learner and the opponent in explicit detail with specification of representations and learning algorithms giving the notion a precise meaning. Through focusing on particular detail, we are exploring how knowledge can be extended by specific changes without abstract generalization. The question is of general interest if one admits that particularity and egocentricity are common characteristics of novice thought.

In analyzing the possibilities for learning, the basic principle applied is to try all cases and construct an interpretation of them. There are many paths of possible learning, some central and some peripheral. In SLIM, the criterion of centrality is co-generativity. We first aggregate results of all possibilities in a fully explicit manner. SLIM highlights co-generativity, represented by those connections in the central group (strategies 1- 6) permitting each one to be learned no matter which is adopted as the prototype fork. Peripheral strategies[13] are rarely learned because they can be learned in few ways. The

[13] These are detailed in "On the Merits of the Particular Case," in Case Study and Computing.

aggregation of trees is systematic and constructive though not formal: one pulls together the empirical results of exhaustive exploration (trees of descent in Figure 2) into a new representation scheme (the genetic networks of descent in Figure 3).

Through the aggregation and reformulation of results, a new generality was achieved which suggested an idea, call it a "learnability principle," that co-generatability of related but variant knowledge forms is what makes learning possible in any domain. This principle would support knowledge stability in minds with reconstructive memories, such as Bartlett (1932) suggests humans have. This idea - - identifying the learnability of a domain with its connectedness - - is a direct consequence of the knowledge state transformations being the learning algorithms of the system. One may paraphrase the situation thus: "you'll learn where Rome is, if all roads lead there."

In tic-tac-toe, we have focused on a two dimensional puzzle, size 3x3. At first glance, Rubik's Cube is similar, as a 3x3x3 puzzle. But further dimensions are relevant. Color enters consideration as well as constraints implicit in the interior mechanism. Let's look at the situation generally before plunging into puzzle solution.

Constructed Personal Knowledge: Mastering a Solution to Rubik's Cube

Over 350 million Rubik's cubes have been sold worldwide. When a cube is "solved," each of the six faces has a single color. There is one such state. There are many more unsolved states.[14] There are different methods for solving this puzzle; typically one solution is provided as a guide when a puzzle is purchased, e.g. a "7 Step Solution Guide." Many people take up this puzzle to prove they are "smart." [15] Often these cubes remain disordered and are set aside, and the confident problem solvers are defeated. Rubic's Cube remains an opportunity for a personal learning experience and more, an example of activities involved in learning to solve, to understand, and to remember a puzzle solution that many find too difficult.

The Challenge Rubik's Cube Presents

Rubik's 3x3x3 cube first appears to be an assembly of 27 different little cubes (call them "cubies,") linked together internally. The mechanism takes up the central cubie and the six face-center cubies. These seven never move (except by cube rotation). Thus there are twenty moveable pieces, i.e. 12 edge cubies and 8 corner cubies.

What makes the puzzle engaging:
As purchased, the cube comes in its solved state. This proves it can be solved. The cube can be held in your hand; it is even sold with a solution booklet. Everything about it is familiar, except the mechanism inside. Clever people conclude, "I can do this."
What makes the puzzle hard:
Multiple modalities: recognizing visual patterns is essential to solving the cube; planning requires foreseeing how sequences of moves change multi-facial patterns.
Sameness of views: it is easy to get lost in regards to orientation.
Sameness of operations: it is easy to lose track of "where you are" in move sequences.
Interactions: changing the state of any face, edge, or cubie ALWAYS impacts others.

[14] See Wikipedia general article on Rubik's cube.
[15] See Wikipedia "Rubik's Cube in Popular Culture."