

LEARNING CONCRETE STRATEGIES THROUGH INTERACTION

R.W. Lawler and O.G. Selfridge
GTE Laboratories, Inc., Waltham, Ma. 02254

Abstract

We discuss learning and the adaptive generation of concrete strategies through interactive experience. The domain is the game Tictactoe. The knowledge structures embodying strategies we represent as having three parts: a Goal, a sequence of Actions, and a set of Constraints on those actions (GAC). We simulate such structures in a program that plays Tictactoe against different kinds of opponents. Applying these strategies leads to moves that often result in winning or losing; which in turn leads to the creation of new structures, by modifying the current GACs. These modifications are controlled by a small set of specific rules, so that the GACs are related by the ways modifications can map from one to another. Subject to certain limitations, we do a complete exploration of certain classes of strategy. This learnability analysis takes guidance from previous cognitive studies of a human subject by Lawler. The simulations were performed on a Symbolics 3600 in LISP.

This work avoids abstraction in order to explore learning based on the modification of fully explicit strategies learned through particular experiences. The results are a catalog of specific experiences through which learning occurs within this system and a description of networks of descent of concrete strategies from one another. We conclude with a description of experience motivated analogy, a mechanism proposed for the development of limited interrelations between structures based on particular experiences.

Introduction

The ultimate aim of this work is to model computationally a sequence of developmental steps leading from naive knowledge to domain specific mastery. We start with a model that embodies three well recognized characteristics of children's thinking and with a representation chosen taking guidance from the rich corpus of a detailed case study.¹ The three primary characteristics are egocentricity, concreteness, and interactivity. EGOCENTRICITY we take to mean an original, exclusive focus on one's-own goals and plans. By CONCRETENESS we refer to the importance of particular experience in determining the descriptions of what things are. We take seriously the concreteness of young children's thought, and we are willing to explore how very specific structures function and develop. The self-centeredness of children's play, even their speech (see Piaget, 1926), is characteristic of the naive learner. This leads us to use information structures that focus almost entirely on the intentions and actions known by the player. INTERACTIVITY is central in the twin senses of learning from interaction with others and from the interaction of disparate structures within the self. The development of structures from such a basis is eventually to proceed to mastery through recognizably significant processes such as the following:

- learning new concrete goals and plans as elementary variations of predecessors
- learning constraints upon those plans
- redescribing goals based on the application of plan constraints to goal specifications
- recognizing the applicability of general ideas (such as symmetry)
- internalizing the opponent's role (thus permitting vicarious play)
- developing mental play (thus permitting search for a victory through the tree of possible games)
- achieving a classification-based mastery of possible play in the domain.

We now explore learning-by-interaction of concrete strategies through analysis of computer-based simulations. The puzzle is how to understand adaptivity through interaction with some "other" thing that is not itself well-comprehended.

¹ Chapter 4 of Lawler, 1985, is a detailed analysis of the complete history of one child's play of Tictactoe throughout a two year period. References to data and analyses from that study will be to "Lawler's subject."

Although this exploration takes guidance from the study of a single child, it is NOT a model of a single child's development. The immediate aim of the work reported here is to explore the interrelatedness of the set of concrete strategies, such as can be generated with minimal transformations of a given strategy through purpose governed interactions with an uncomprehended opponent.² We believe that development can follow from reflection upon the disparity between anticipated outcomes of plans and actual outcomes of particular experiences. The central task of this model is to modify its repertoire of strategies when surprising outcomes occur in playing particular games. The model proceeds by the recognition of new specific goals and the construction of new plans for achieving those goals. This kind of learning embodies the reflexive construction of concrete strategies.

Central ideas for us are the disparateness of structures and their functional lability.³ We apply the assumption of disparateness in our simulations by creating a separate object for each strategy; each object proposes its next move regardless of the proposals of every other strategy. Our analysis combines (1) a suitable representation of knowledge structures; (2) decisions made to limit the interactions to some manageable number; (3) proposals of learning mechanisms operating on the knowledge structures; and (4) the analysis of sequences of simulations. Simulation lets us specify and activate the structures and developmental sequences we consider important, creating thereby a kind of experimental epistemology. This in turn permits us to connect analysis of the particular context of an incident of learning to specific changes in an organization of self-modifying information structures.

Representation of Knowledge

We represent knowledge structures as having the parts necessary for adaptive functioning. Learning what to do is essential: GOALS are explicitly represented. Knowing how to achieve a goal is essential; ACTION PLANS are explicitly represented. Knowing when a planned action will work and when it won't is essential; CONSTRAINTS limiting application of actions are represented explicitly. The structure composed of this triad, a GAC (Goal, Action, Constraints), is our representation of a strategy for achieving a fork in Tictactoe. Goals are considered as a three element set of the learner's marks which take part in a fork. This is the first element of a strategy. Plans of three step length, which add the order of achieving goal steps, are represented as lists. Constraints on plans are two element sublists, the first element being the step of the plan to which the constraint attaches and the second being the set of cell numbers of the opponent's moves which defeated the plan in a previous game. In our simulations, SLIM (our Strategy Learner, Interactive Model) plays against REO (a relatively expert opponent). REO can win, block, and apply various rules of cell choice – though ignorant of any strategies of the sort SLIM is learning. Within the execution of our simulation, the structure of GAC 1 below will lead to the three games shown depending on the opponent's moves (letters are for SLIM's moves, numbers for REO's):

	GOAL	ACTION	CONSTRAINT
GAC 1:	{1 3 9}	[1 9 3]	< [3 { 2 5 8 }]>
win by plan	plan defeat	constrained draw	cell numbers
$\begin{array}{ c c c } \hline A & 3 & C \\ \hline & 1 & D \\ \hline 2 & & B \\ \hline \end{array}$	$\begin{array}{ c c } \hline A & & C \\ \hline 2 & & 1 & 3 \\ \hline & & B \\ \hline \end{array}$	$\begin{array}{ c c c } \hline A & & C & 3 \\ \hline 4 & & 1 & E \\ \hline D & & 2 & B \\ \hline \end{array}$	$\begin{array}{ c c c } \hline 1 & & 2 & 3 \\ \hline 4 & & 5 & 6 \\ \hline 7 & & 8 & 9 \\ \hline \end{array}$

The player SLIM has no notion of symmetry (ongoing work explores how it may be learned). Our representations and learning mechanisms are committed to cell-specificity; they are also self-centered,

² This line of research arose from the attempt to develop the kinds of concrete learning models proposed in Selfridge and Selfridge '85 in a direction less dependent on the external setting of goals.

³ Satinoff, 1978, presents a physiological example of disparate systems for thermal regulation. Lawler 1985 assumes the disparateness of cognitive structure in interpreting learning in the human case. The functional lability of structures names the effectiveness of a given structure for something other than the purpose which shaped its development. Jacob 1981 argues the importance of functional lability in the evolution of species. The idea is not forward in this article, but is central to the explanation of learning through experience generally. See Lawler 1985, 1985b.

focussing on the learner's own plans and knowledge (as they must since, by principle, SLIM begins not knowing what the opponent will do; SLIM does not have the ability to model or predict an opponent's moves in any abstract way).⁴

Here are the fourteen corner-opening GACs from each of which SLIM began one of its play-sequences against REO (they are grouped in symmetrical quartets; because GACs five and six are doubly symmetrical, there are only two members in its group):

GAC 1 {139}[193]<>	GAC 5 {137}[137]<>
GAC 2 {179}[197]<>	GAC 6 {137}[173]<>
GAC 3 {139}[139]<>	
GAC 4 {179}[179]<>	
GAC 7 {135}[153]<>	GAC 11 {125}[152]<>
GAC 8 {157}[157]<>	GAC 12 {145}[154]<>
GAC 9 {135}[135]<>	GAC 13 {125}[125]<>
GAC 10 {157}[175]<>	GAC 14 {145}[145]<>

The central question is: when SLIM starts from one specific GAC, which others will it discover through experience – given a set of well defined assumptions about learning mechanisms and the opponent, and learning mechanisms. We will describe the central elements of the domain, the opponent, and learning mechanisms.

The Domain

There are at least four ways of mobilizing knowledge to play Tictactoe. Most often, mature people think of playing forward in the tree of possible moves. A second kind of mature play is categorical: one may group all possible pairs of openings and responses, reducing them by symmetry to twelve unique game openings; of these twelve, seven permit direct wins for the initial player and the other five permit setting traps for the unwary opponent. A third way is more primitive than the first two; we call it tactical: TACTICAL play is purely state dependent; from the tokens already played, choices are made to win when possible, to block at need, and to choose one cell from those available, one choice at a time and with no forward play.

The fourth way we call strategic: it depends upon having a multi-move action-plan for achieving a fork, which distinguishes it from tactical play. Initially it completely ignores the opponent's moves, which distinguishes it from look ahead. It involves NO systematic knowledge of the game at all, and does NOT make ANY explicit use of symmetry in play.

Lawler's subject clearly exhibited use of the third and fourth ways, that is tactical and strategic play. SLIM was so programmed, preferring strategic play. But when SLIM's strategies are frustrated, it plays with structured tactics (to be defined shortly). Winning with strategies depends upon SLIM's making a winning move tactically after forking the opponent. SLIM's opponent, REO, ALWAYS plays tactically. Both SLIM and REO can recognize and achieve a win. REO always blocks any threat by SLIM (or one of two when forked). SLIM, not noticing its opponent's moves, initially does not block while executing a strategy – until a defeat establishes its vulnerability, and analysis leads to identifying and constraining execution of the vulnerable plan step.

SLIM plays in two modes, executing strategies when possible and resorting to tactical play when necessary (all strategies frustrated) or appropriate (a fork has been achieved). When playing tactically, SLIM blocks as effectively as REO.⁵ When a single strategy is abandoned because of a threat, play reverts to the tactical mode – and thus blocking is effective.

⁴ The general commitment to egocentric knowledge representation has psychological justification in this specific case. Lawler's subject suffered the defeat above trying to achieve the victory of GAC 1 (the only strategy she knew), not attending to her opponent's move nor anticipating any threat to her intended fork.

⁵ SLIM does not block while pursuing a strategy. This seeming inconsistency reflects the behavior of Lawler's subject. The fragmentation of knowledge implied by such behavior is implemented in the model through non-communicating instances of Zetalisp flavors.

SLIM's opponent REO plays tactically in three variations. They can be defined in terms of application of a set of production rules:⁶

<own-winning-pattern>->	play winning-cell	rule	type of play
<opponent-win-pattern>>	play blocking-cell	—	—
<center-cell free>	-> play center-cell	1:	highly structured play only
<corner-cell free>	-> play corner-cell	2:	highly structured and structured play
<any-cell free>	-> play any-free-cell	3:	un-structured play; last choice for others.

When all the last three rules function, the opponent plays in a HIGHLY STRUCTURED fashion. This means that rule 1 inhibits the choice of any corner or side cell when cell 5 is free. When rule 1 is not functioning, the opponent plays in a STRUCTURED fashion. When both rules 1 and 2 are disabled, the opponent plays in an UNSTRUCTURED fashion – moving in any free cell. Since our simulations generate the complete spaces of potential experiences, disabling the limitations implied by one specific rule constraining its successors generates an increased number of games in the space.

Spaces of Potential Experience

The space of possible Tictactoe games is of magnitude $2 \times 9!$, or 725,760. We have severely limited the portion of that space to be explored in order to permit our tracing how the space of executable games changes with increasing strategic knowledge. Domain oriented limitations we have applied are focussing on corner opening play with SLIM moving first. These choices are based on the significant opening advantage of corner-oriented play and on the important role we assign to winning games as the primary cue for learning a new strategy. These constraints still leave on the order of $8!$ (40,320) games. We ignore drawn games and have minimized the generation of redundant games (alternative responses to a fork and so forth). We have also omitted the analysis of eight-move losses and nine-move wins as too complex for SLIM; we believe these circumstances are too complex for analysis by naive human players of Tictactoe also. These limitations of choice restrict the games playable, but they also permit us to explore how very simple strategies can generate others which will in turn expand the space of potential experiences.

Beyond these limitations of choice, we apply a more ESSENTIAL limitation. The most natural constraint on play is its specific purpose, in the twin senses of having winning as a general objective and intending to apply knowledge available as a means to that end. If one knows only a single, very specific strategy – such as “move first in this corner, then that other, and finally in the third one there” – and is determined to apply that knowledge, play is very much limited. At the beginning, our computer learner SLIM had such commitments, just like Lawler's subject. Our computer simulations are based on generating the space of KNOWLEDGE-LIMITED POTENTIAL EXPERIENCES. We want to know how different new strategies can be derived from the various outcomes of a single predecessor strategy's play. These descendents we call the FAN-OUT from a strategy. We want to know where you can go from where you start out. Will there be convergence to some single end or residual differences in final states?

We generate the entire space of potential experiences which a learner could have and then analyze all those games where learning occurs through activation of the allowed mechanisms. The learning mechanisms are limited to changing only one element of a structure at a time. Learning through losing games leads to the addition of a constraint element in the third slot of a GAC (if there isn't any GAC to which the constraint may be attached – a frequent situation – no such learning can occur). Learning through winning creates a new GAC, often two. Each has a goal, an action-plan, and an initially empty list of constraints. Whenever a seven-move win occurs, it is always possible to infer the pattern of a new goal.

Differences between methods of GAC creation depend upon how the new plan to achieve the goal is constructed. Because SLIM moves first in cell 1, any new strategy will have plans different in either step two or step three. The first mechanism, plan terminal modification (PTM), changes the final element of a known plan to match the actual win that occurred in a specific game. PTM generates one new GAC, the method is dependable and is always tried before the second, Goal-guided Center Deletion (GCD); GCD is more complex in inference and less certain in value.

⁶ Newell and Simon, 1972, pg. 62 shows such a system in detail as an example.

GCD changes the center element of a three-step plan: once a goal has been inferred, SLIM attempts to locate a known goal with at least two elements matching those of the new goal. If there is one, the associated plan is taken as the template for constructing a new plan. The element of the old plan not in the new goal is deleted from the original. The element of the new goal lacking in the original plan must then be inserted. But where? SLIM is not sophisticated enough to know or find out which is appropriate. Consequently, two new GACs are generated with a common goal but with plans differing in the order of the last two steps. The human counterpart would be recognizing a specific new goal but being uncertain as to the order of the last two moves.

Simulations

Strategic play proceeds until it achieves its expected victory – or is blocked or unexpectedly defeated. The typical loss by a strategy is through its circumvention (REO lays down three in a row while SLIM pursues its plan). The application of constraints to plans turns such six-move defeats into draws. The essential analysis is that loss after fork completion implies that the opponent posed a threat after two moves. Since SLIM pays no attention to what the opponent was doing, it has no sense of the order of moves in which the REO's victory was achieved. Therefore, the decision to apply a constraint to an ongoing game tests for the opponent's having acquired any two of the three moves in the triad by which REO lately defeated SLIM. If this be "look-ahead", it is not a general capability but a specific functional prohibition motivated by a prior defeat. A general capability would require further development.

An unexpected win is the cue to begin reflection. SLIM assumes that a victory was won not received as a gift; therefore the moves preceding the win contain a fork. (This works because REO always wins when possible and blocks at need.) Since the most recent move by which a win was completed could not have been part of the fork, the remaining moves are not merely IN the fork, they COMPRISE the fork pattern exactly. A valid forking goal can be found for any seven-move win against a relatively expert opponent.

Play Against a Highly-Structured Opponent

SLIM's processes of learning can best be shown by examples. Against a highly structured opponent (one who ALWAYS moves in the center when it is empty, who ALWAYS prefers a corner move to a side cell), the space of potential experiences is very small, and thus the learning that occurs is quite limited. We exhibit the action of the learning mechanism in particular games. Consider the case where SLIM knows only the strategy represented by the plan [193]. Because of a highly structured opponent's dominant commitment to an initial cell 5 move, only two games are possible.⁷ The first is success of SLIM's strategic plan in game [1597326]. The second is a victory achieved after the blockage of SLIM's plan [1593726]. Because SLIM will prefer a corner cell to a side cell, when REO blocks his third plan step by moving in cell 3, SLIM achieves the win of GAC 2's plan [197].


In reflexive analysis after play, SLIM constructs GAC 2. The determination of the goal is 179. The plan derives from this analysis: a win implies a fork. The fork is the three moves preceding the winning move. The order of moves must be similar to the frustrated well-known plan [193], because it was blocked only at move three by a token still available for examination. The move substituted by tactical play for the blocked move remains equally present. The conclusion is that [197] is a good candidate plan for attachment to the goal 179. The GAC is formed by this terminal modification of a previously known plan.

A case of a win for which PTM cannot construct a plan is game [1597326], played with a strategic objective of achieving plan [153]. SLIM abandoned its plan at move two because it was blocked. The form of the fork actually achieved was GAC 1's goal 139. Changing the final step of plan [153] can in no way lead to a plan [193]. The plan was constructed from [153] as a template (only GAC 7 was known at the time) by goal-guided center deletion. In fact, a second plan, [139], was also constructed.

Although one can always infer a new goal from a win, one cannot always construct a new plan for that goal, as with GAC 7's play against a highly structured opponent in game [1593748]. The effective forking pattern is GAC 2, but to create a plan with elements 9 and 7 from a template [153] is not

⁷ We will represent games in the text as number strings within brackets. Each digit represents the cell number of a move. The first move is that of SLIM, the second REO's, and so on alternately.

possible by mechanisms that can change only a single element. This is Plan Generation Failure (PGF). In this case of strategies descended from play in the space of GAC 7's potential experience against a highly structured opponent, we can see in miniature the major processes that act in the more extended simulations discussed below.



This tree of structures descended from one strategy is its FAN OUT. Our simulations have documented the fan out for the fourteen GACs specified earlier. Against a highly structured opponent, SLIM is able to learn very little, as shown by the first run summary:

GACs	1	2	3	4	5	6	7	8	9	10	11	12	13	14
games	2	2	1	1	1	1	6	6	2	2	6	6	1	1
wins	2	2	0	0	0	0	2	2	0	0	2	2	0	0
losses	0	0	1	1	1	1	0	0	0	0	0	0	0	0
draws	0	0	0	0	0	0	4	4	2	2	4	4	1	1
new GACs	1	1	0	0	0	0	2	2	0	0	0	0	0	0
constraints	0	0	1	1	1	1	0	0	0	0	0	0	0	0

Playing Against an Opponent of Structured Preferences

Table 1 summarizes the results. Each section shows play begun with a specific GAC as the learning root. For each of the 14 GACs is shown (1) its fan-out of learnable strategies, (2) the specific games through which learning occurred, and (3) the mechanism by which the new structure was generated. SLIM learned everything possible to learn in two passes. Further, the first six GACs all converged to an essentially uniform result wherein they were all known; each of these six GACs initially generated 10 games, while in the final common state 156 games were generated by them. These six thus form a central collection of strategies whose specialness needs be emphasized. Strategy three may be taken as typical of the others. Play in five specific games generates the other five central GACs. A sixth game installs a constraint upon the root GAC of the set. The remaining eight GACs fall into three categories when seen as bases for further learning: feeders (7,8); teasers (9,10,11,12); and mules (13,14). Feeders are strategies that themselves fail but generate games permitting construction of strategies in the central set; consequently, they give access to all other central strategies. When they are frustrated and abandoned, teasers lead to other winning games for which SLIM is unable to construct any plan. Mules are strategies which can win some games against the opponent of structured preferences but which cannot generate wins of any other forking patterns and are thus sterile as generators of new structures.

The specialness of the six central nodes is a consequence of their symmetry in respect of co-generability. Some generate others directly by PTM or GCD. Some generate others only with the intervention of an intermediary GAC (this is the reason that two passes are necessary to complete the central set). Some of those directly generable can generate each other; they are reciprocally generable. Those which are remotely generable nonetheless lead to each other through intermediaries; they are cyclically generable. For these six central strategies, the trees of structure descent fold together into a connected network of descent whose relations of co-generativity are shown in the diagram below; notice that those strategies cyclically generable form the interconnection by GCD between themselves and others reciprocally generable by either PTM or GCD:

Table 1
**STRUCTURE GENERATION
 OPPONENT WITH STRUCTURED PREFERENCES**

ROOT GAC	RUN ID	GENRT GAME	INFRD PLAN	REASON	GOAL/ PLAN or CONSTRNT	LEARNG MECH	GAC ALTERED	COMMENTS
1 [193]	1.1	[1974325]	1	root	{137} [137]	GCD	5	Guided center deletion creates two GACs. Plan tail modification creates one GAC.
	1.2	"	"	"	" [173]	GCD	6	
	1.3	[1593748]	"	"	{179} [197]	PTM	2	
	2.1	[1732956]	5	01 x L3	{139} [139]	PTM	3	
	2.2	[1374958]	6	01 x L3	{179} [179]	PTM	4	
2 [197]	1.1	[1974325]	2	root	{137} [137]	GCD	5	At the end of the runs six central plans have been learned. No plans have been constrained. at all.
	1.2	"	"	"	" [173]	GCD	6	
	1.3	[1795326]	"	"	{179} [197]	PTM	1	
	2.1	[1732596]	5	01 x L3	{139} [139]	PTM	3	
	2.2	[1374958]	6	01 x Le	{179} [179]	PTM	4	
3 [139]	1.1	[1932745]	3	root	{137} [137]	PTM	5	First constraint made. Six experiences make one constraint and five new GACS.
	1.2	[153298]	"	"	<[9{258}]>	PCA	3	
	1.3	[1395748]	"	"	{179} [197]	GCD	2	
	1.4	"	"	"	" [179]	GCD	4	
	2.1	[1974325]	4	01 x L3	{137} [173]	PTM	6	
2.2	[1795326]	2	01 x L3	{139} [193]	PTM	1		
4 [179]	1.1	[1974325]	4	root	{137} [173]	PTM	6	Root plan constrained. Six plans completed.
	1.2	[1795326]	"	"	{139} [193]	GCD	1	
	1.3	"	"	"	" [139]	GCD	3	
	1.4	[157496]	"	"	<[9{456}]>	PCA	4	
	2.1	[1932745]	3	01 x L3	{137} [137]	PTM	5	
2.2	[1593748]	1	02 x L3	{179} [197]	PTM	2		
5 [137]	1.1	[1732956]	5	root	{139}[139]	PTM	3	Six plans completed; root plan constrained.
	1.2	[153278]	"	"	<[7{258}]>	PCA	5	
	1.3	[1395748]	"	"	{179} [179]	GCD	4	
	1.4	"	"	"	" [197]	GCD	2	
	2.1	[1974325]	4	01 x L3	{137} [173]	PTM	6	
2.2	[1795326]	2	01 x L3	{139} [193]	PTM	1		
6 [173]	1.1	[1795326]	6	root	{139} [139]	GCD	3	Symmetry in relation to the the learnability of plans is discussed in the text and show in the descent network of Figure 1.
	1.2	"	"	"	" [193]	GCD	1	
	1.3	[157436]	"	"	<[3{456}]>	PCA	6	
	1.4	[1374958]	"	"	{179} [170]	PTM	4	
	2.1	[1932745]	3	01 x L3	{137} [137]	PTM	5	
2.2	[1593748]	1	02 x L3	{179} [197]	PTM	2		


Abbreviations used in the table:

GCD: Goal-guided Center Deletion; 01,02: Opponent moves 1 and 2; PTM: Plan Terminal Modification; L2,L3: Learner moves 2 and 3; PCA: Plan Constraint Application; 01xL2: Move 01 blocks L2; PGF: Plan Generation Failure.

Table 1
**STRUCTURE GENERATION:
 OPPONENT WITH STRUCTURED PREFERENCES**

ROOT GAC	RUN ID	GENRT GAME	INFRD PLAN	REASON	GOAL/ PLAN or CONSTRNT	LEARNG MECH	GAC ALTERED	COMMENTS
7 [153]	1.1	[195738]	7	root	<{3{789}}>	PCA	7	These mechanisms don't succeed when a goal is not near known games; first example: failure
	1.2	[1597326]	"	"	{139} [139]	GCD	3	
	1.3	"	"	"	" [193]	GCD	1	
	1.4	[1593748]	"	"	{179} []	PGF	.	
	2.1	[1974325]	1	01 x L2	[137]	GCD	5	
	2.2	"	"	"	[173]	GCD	6	
	2.3	[1593748]	1	02 x L3	{179} [197]	PTM	2	
	2.4	[1374958]	3	01 x L2	{179} [197]	GCD	2	
	2.5	"	"	"	" [179]	GCD	4	
	8	1.1	[195376]	8	root	<{7{369}}>	PCA	
[157]	1.2	[1597326]	"	"	{139} []	PGF	.	[157][193][157]
	1.3	[1593748]	"	"	{179} [179]	GCD	4	[157][193][157]
	1.4	"	"	"	" [197]	GCD	2	[157][193][157]
	2.1	[1974325]	4	01 x L3	{137} [173]	PTM	6	[179][173][193]
	2.2	[1932745]	2	01 x L2	{137} [173]	GCD	6	[197][173][157]
	2.3	"	"	"	" [137]	GCD	5	[197][137][157]
	2.4	[1795326]	2	01 x L3	{139} [193]	PTM	1	[197][193][197]
	2.5	[1732926]	4	01 x L3	{139} [193]	GCD	1	[179][193][179]
2.6	"	"	"	" [139]	GCD	3	[179][139][179]	
9 [135]	1.1	[1395748]	9	root	{179} []	PGF	.	No learning follows trying this plan
	1.2	[1374958]	"	"	{179} []	PGF	.	
10 [175]	1.1	[1795326]	10	root	{139} []	PGF	.	No learning follows trying this plan
	1.2	[1732956]	"	"	" []	PGF	.	
11 [152]	1.1	[195748]	11	root	<{2{789}}>	PCA	11	First example: two constraints on a GAC; they need NOT force different games of or new learning
	1.2	[195326]	"	"	<{2{369}}>	PCA	11	
	1.3	[1597348]	"	"	{139} []	PGF	.	
	1.4	[1593748]	"	"	{179} []	PGF	.	
	2.1	[1597326]	"	no other	{139} []	PGF	.	
	2.2	[1593748]	"	"	{179} []	PGF	.	
12 [154]	1.1	[195748]	12	root	<{4{789}}>	PCA	12	Symmetrical case to failures of GAC 11.
	1.2	[195346]	"	"	<{4{369}}>	PCA	"	
	1.3	[1597326]	"	"	{193} []	PGF	.	
	1.4	[1593748]	"	"	{179} []	PGF	.	
13	1.1	[192356]	13	root	<{5{369}}>	PCA	13	The only learning.
14	1.1	[194758]	14	root	<{5{789}}>	PCA	14	The only learning.

Figure 1
STRATEGY DESCENT NETWORK: STRUCTURED OPPONENT




solid arrows show descent by PTM; dashed arrows show descent by GCD
 With the two mechanisms assumed and the experience of appropriate specific games, the central strategies will all be learned once any one is known.

Against the opponent of structured choices, the feeder strategies exemplify both plan generation failure and redundant generation. Thus the goal 179 for which SLIM can construct no plan in pass one is regenerated in pass two. In the second pass, the expanded repertoire of known plans permits construction of GAC 2's plan. They also show the value of distinguishing between SLIM's intended plan when a game began, the pattern of the realized fork, and the plan drawn from its repertoire which serves as the template for construction of the new plan. There is no necessary identification of the first with the third. The plans of GACs 9 through 12 are too dissimilar from the actual forks they generate to permit construction of a new plan for the goals SLIM infers.

Figure 2 summarizes the results of comparable simulations of play against an unstructured opponent.

Figure 2
STRATEGY DESCENT NETWORK: UNSTRUCTURED OPPONENT



solid arrows show descent by PTM; dashed arrows show descent by GCD
 GACs 13 and 14 appear twice to simplify the network drawing

GACs 1 to 6 still remain central, but the flexibility of the unstructured opponent permits SLIM to learn additional strategies. GACs 13 and 14 remain non-learnable through experience with this opponent. The specific reason is that SLIM's tactical preferences remained structured in these simulations. Thus SLIM, never trying games whose second move is to a side cell, will never win accidentally with a fork containing such patterns; consequently, SLIM can never learn such strategies.


The form of these descent networks is related to symmetry relations among forking patterns, but the networks include more; they reflect also the order in which the fork is achieved, the play of the opponent, and the specific learning mechanisms permitted in the simulations. Whereas fork pattern symmetry can be used as an alternative mechanism for proposing new forks and games to play, these descent networks are summaries of results. The question yet to be answered is whether or not the relations represented in these networks can serve as the basis for the discovery that pattern oriented symmetry functions with power in this domain.

A general observation is that the generation of constraints seems less important than we expected. Against the opponent of structured choices, only the root GAC is constrained. In one sense, this is an artifact of programming, but it points to a significant further issue. Our simulations generate the trees of potential experiences by expanding a list of possible games around the generation of sets of plausible moves according to the appropriate strategic or tactical choices. The set of plausible moves is collected from separate instances of a class of strategy objects which do not communicate with each other. During the first pass, only the root GAC is accessible to be constrained. During the second pass, the alternative choices offered by the three or four functional strategies are so effective that forking the opponent closes off the possibly vulnerable proposals of its simultaneously active strategies.

Because of this collection process, the proposal of one strategy could at first be blocked by a constraint; yet the "forbidden" game could be generated anyway by an identical cell proposal being advanced by some other strategy for its own plan. This interaction has been suppressed in these simulations – but it does suggest a future direction of development for models of this class: the communication of constraints upon the occurrence of a "forbidden defeat". This can serve as the basis for experience motivated analogy, a new mechanism for the development of organization where none formerly existed.

Experience Motivated Analogy

Beginning with GAC 3 as the root of its strategies (playing against the structured opponent), the first simulation set of games permits generating GACs 2, 4, and 5. A constraint is also put on GAC3. Consider play in the game beginning [1532...] – while these four GACs are simultaneously active – at move 5 (SLIM's third move):



GAC 3, with its established constraint on a third move to cell 9 recognizes the vulnerability of its plan and quits; it marks cell 9 as a forbidden move. This proposal results in a game being played by the tactical routines, which block REO's threat. At this same pass, GAC 5, whose plan is compatible with that of GAC 3 through move 2 and is unconstrained, proposes a move to cell 7. This results in a different game being generated and lost when REO moves next in cell 8. GAC 4 at move 5 proposes a move to cell 7, the second step of its plan. The plan of GAC 4 may be characterized as having been interrupted in execution and as being diluted by an extraneous move. This game loses to REO's threat as well. Gac 2 proposes a move to cell 9, the second move of its plan, which is also interrupted and diluted. This is exactly the move that is constrained in GAC 3. The game is generated and lost.

When these losses occur, the process of blame assignment begins: who proposed these losing moves? Each GAC justifies its behavior. GAC 5 admits to the cell 7 proposal, with the additional information that its plan was uninterrupted and pure (not diluted); SLIM sends a message to GAC 5 that it should constrain itself on the cell 7 move by the threatening pattern 258. GAC 4 also proposed cell 7, but its

plan was interrupted and diluted. SLIM informs GAC 4 to constrain itself at its cell 7 move (the second plan step) with the threatening pattern 258. GAC 3 forbade the move to cell 9. GAC 2 proposed the move anyway. SLIM marks GAC 2 as an "idiot". (Etymologically, the term's Greek meaning refers to someone who only knows or cares about his own affairs.) SLIM informs GAC 2 that it should constrain itself by subordination to GAC 3, which forbid the move. The implication is that at its second move, GAC 2 should request constraints from GAC 3 and apply them in its own proposal process.

The significance of this procedure is in specifying a particular set of experiences which justify the judgment that GAC 2 is sufficiently similar to GAC 3 that it SHOULD use analogy to determine the vulnerability of its own plan. The first consequence is improved play. A second is that should GAC 3 constrain itself by some other threatening pattern – for example the perpendicular threat 456, GAC 2 would inherit that protection as well. This subordination creates an interrelationship among structures learned through experience without any specific control structure being imposed at the time of the generation of the structure.

Conclusions and Frontier Problems

Our analysis of a part of the space of strategies exemplifies the complexity that arises from experiential and interactive learning. The strategies and the learning mechanisms are simple, but the simulations show the surprising power that they can exhibit.

A modest extension of this model leads to goal redescription through application of plan-constraints to goal specifications, calling upon a Sussmanesque inversion of a failure-description to generate a new type of term (the hole or empty-cell) for use in the goal description. Networks of descent describe sets of richly interrelated structures created through processes of reflexive concrete construction. We are now exploring whether two other types of reflexive processes, acting upon these structures as objects, will be adequate to recognize the applicability of symmetry in this domain and ultimately to create a complete classification of games adequate to exhibit domain mastery. The first process classifies objects by specifiable features; the second classifies objects by analysis of their function and genesis. Piaget (1971) calls these processes Aristotelian classification and reflective abstraction.

A different frontier represents the effort to explore how to structure the model so that the components exhibit the sort of functional lability that characterizes human intelligence and seems essential to begin development of a model of the play of the opponent.⁸ Furthest from our current grasp is exploring the extent to which experience motivated analogy and other forms of developing control relationships permit the establishing through experience of coherent interaction among a set of initially disparate structures. The extension of such a model as this to other domains will be worthwhile attempting when it successfully traverses its path to mastery. It is our expectation that the concrete representations developed will tend to be domain limited but that the processes for the analysis of concrete strategies will be less bound to particulars.

Acknowledgements

Thanks to Bud Frawley for creating the control structure which made these simulations possible. Our colleagues Golden, Cooperman, Fawcett, and Braunegg have provided friendly criticism and help. Marie McPartland-Conn made it possible to have a readable copy of this text.

References

- Jacob, F., *The Possible and the Actual*. New York: Pantheon Books, 1982.
Lawler, R., *Computer Experience and Cognitive Development*. Chichester, England and New York: Ellis Horwood Ltd. and John Wiley, The Halstead Press, 1985.
Lawler, R. "The Internalization of External Processes," unpublished, 1985b.
Newell, A. and Simon, H. *Human Problem Solving* Englewood Cliffs, N.J.: Prentice Hall, 1972.
Piaget, J., *The Language and Thought of the Child*. New York: New American Library, 1974.
Piaget, J., *Biology and Knowledge* Chicago: University of Chicago Press, 1971.
Satinoff, E., "Neural Organization and Evolution of Thermal Regulation in Mammals," *Science* 201, pp. 15-22, 7 July 1978.
Selfridge, M. G. R. and Selfridge, O. G. *How Children Learn to Count: a computer model*. 1985.

⁸ A discussion of this work appears in *The Internalization of External Processes*, Lawler 1985b.