

Explorations in Experimental Epistemology

Robert W. Lawler, LawlerEmail@gmail.com
P. O. Box 726, Lake Geneva WI 53147 USA
for Oliver G. Selfridge, with gratitude and admiration

Abstract

This paper discusses the epistemology of learning in two different dimensions.

First, **Learning through Interaction** focuses on learning with concrete methods in a particular domain (learning strategies for playing Tic tic toe) with a machine learning study grounded in a human psychological case study (see "The Articulation of Complementary Roles," chapter 4 in Lawler, 1985).

The methods of learning in the computer modeling employ learning by example (Winston, 197n) and learning by debugging (Sussman, 197n). Codifying results and re-representing those results characterizes the learning (as a function of initial knowledge state, learning mechanisms invoked, and opponent actions) as a "network of genetic descent." This permits two novelties.

Although grounded in an individual's study, the modeling shifts focus to an epistemological question, "what kinds and paths of learning are possible." The outcomes are a new principle, "the learnability of a domain is the result of all the possible cases of concrete learning through particular experiences," and a new criterion, that the co-generability of related but variant knowledge forms is what makes learning possible in any particular domain.

Second, **Escaping from Particularity**, motivated by the inability of the machine learning model to encompass symmetry, confronts directly the issue of true novelty in learning in the "strategic aside," *Nil Ex Nihilo*. Explaining the interaction of different modalities of mind as resulting in an effective abstraction by redescription, the proposed 'Multi-Modal Mind' is advanced as a context in which Redescriptive Abstraction serves as a precursor of Piaget's Reflective Abstraction and as a viable candidate for the human mind's general developmental mechanism. This may suffice not only for Piaget's "spiral of learning" but more importantly, may help us understand "the helix of mind arising."

Keywords

learning; epistemology; abstraction; machine learning; strategy learning; case study;

Explorations in Experimental Epistemology

Robert W. Lawler

During six years at the MIT Artificial Intelligence Laboratory with Minsky and Papert, I collected detailed case study material about the learning of three of my children.[1] On returning from several exciting years at Mitterand's World Center for Computers and Human Development, with Papert and Negroponte, at Minsky's urging I joined Oliver G. Selfridge in a new AI Research Group at GTE's Fundamental Research Laboratory.[2] There, I had the opportunity, with the guidance of Selfridge and some help from Bud Frawley, to construct learning models reflecting ideas and observations from my case studies.

That modeling permitted solving some long-standing problems in cognition and epistemology. Questions answered positively in the interpretation and modeling of the LC2 complete Tic tac toe corpus [3] are:

- how one can apply “learning by example”[4] and “learning by debugging”[5] to what people do.
- why some kinds of knowledge can be learned from concrete experience while other kinds of knowledge cannot.
- how one can escape from the particularity of experiential knowledge to more abstract and general knowledge.

But first "why TicTacToe?" you might ask, and "What does 'complete' mean?" A grand theme of the Newell-Simon AI initiative had been the nature of expertise, as realized in Chess Mastery. 3T (as Selfridge referred to TicTacToe) was a game of competitive strategy within the reach of an unschooled 6 year old. The corpus was complete in preserving every game she played in the study from its first introduction through its end.

Learning through Interaction

If learning is an adaptive developmental mechanism, that adaptation must come from interaction with the everyday world. To follow the natural learning of common sense knowledge, a case-based approach is best to track unpredictable learning. Further, if learning is a process of changing one state of a cognitive system to another, based on interactions, then representations of that process in computing terms are most appropriate. But how can one use case material in modeling? As one uses boundary conditions to specify the particular form of a general solution to a differential equation. Such is the use made of the psychological studies here, as a foundation for the representations used in the models, and as justification for focusing on key issues: the centrality of egocentricity [6] in self-construction and the particularity of the naïve agent's knowledge.

The 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. Must we claim that learning happens in people the same way? No. Building such models is an exploration of the possible, according to a specification of what dimensions of consideration might be important. 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 can generate all possible interactions through which learning might occur, including some we first imagine are not important, we can explore alternate paths and the suite of relationships among elements of the ensemble.

Considering All the Possibilities

The model began as a verbally formulated analysis of one child's learning strategic play at tic tac toe (The Articulation of Complementary Roles, Chapter 4 in Lawler, 1985). It was continued in a constructive mode through developing a computer-embodied model, SLIM (Strategy Learner, Interactive Model. (See Lawler and Selfridge, 1986.) The latter was based on search through the space of possible interactions between one programmed agent (SLIM) having some of the characteristics of the psychological subject and a second, REO, a programmed Reasonably Expert Opponent. REO is "expert" only in the sense of being able to apply uniformly a set of cell preference rules for tactical play.[7]

Strategies for achieving specific forks are the knowledge structures of SLIM. Each has three parts: a Goal pattern, a plan of Actions, and a set of Constraints on those actions (each triple is thereby a GAC). I simulated operation of such structures in a program where SLIM plays tic tac toe against variations of REO. Applying these strategies leads to moves that often result in winning or losing; this in turn 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 the GACs are interrelated by the ways modifications can map from one to another.[8]

To evaluate specific learning mechanisms in particular cases, one must go beyond counting outcomes; 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.[9] 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. Suppose that SLIM begins with the objective of developing a fork represented by the Goal pattern {1 3 9} and will proceed with moves in the sequence plan [1 9 3] (see Figure 1). SLIM (A) moves first to cell 1. REO (1) prefers the center cell (5), and moves there. SLIM moves (B) in cell 9. The plan is followed until REO's second move (2) is to cell 3.[10]




Figure 1: cells, a plan, and one played game

SLIM's plan is blocked. The strategic goal {1 3 9} is given over -- but the game is not ended. SLIM, playing tactically with the same set of rules as REO, moves into cell 7, the only remaining corner cell. Unknowingly, SLIM has created a fork symmetrical to its fork-goal. SLIM can not recognize the fork. It has not the knowledge to do so. What happens ? REO blocks one of SLIM's two ways-to-win, choosing cell 4. SLIM, playing tactically, recognizes that it can win and moves into cell 8. This is the key juncture. SLIM recognizes "winning without expecting to do so" as a special circumstance. Even more, SLIM assumes that it has won through creating an unrecognized fork (otherwise REO would have blocked the win). SLIM takes the pattern of its first three moves as 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 for future play.[11]

The complete set of results involves consideration of all paths of 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 (shown in Figure 2). The tree with strategy three as top node may be taken as typical. Play in five specific games beginning with only GAC 3 known, generates the other five central GACs. 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. The specialness of the six central nodes is a consequence of co-generability. Some of those are directly generable, can generate each other (such as GACs 1 and 2) ; they are reciprocally generable (solid lines). Some lead to each other through intermediaries (GACs 1 & 3); they are cyclically generable (dashed).

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 forks are learned, and the learning mechanisms permitted in the simulations. These descent networks are summaries of results.

The experimental epistemology of SLIM begins with a focus on detail

- in the analysis of specific cases and
- in the analysis of the interaction of objects or agents with their context.

The basic principle applied is to try all cases and construct an interpretation of them. To predict the learning of a specific strategy by a human subject, one would need to know what strategies are already known, how the opponent's play would create opportunities for surprising wins for the subject, and what learning methods are in the subject's repertoire. Knowing these things in the machine case is what permits examination of the epistemological space of learnable strategies. In the analysis of SLIM, one begins with lists of games won without a plan. One then reformulates the relations between prototypes and generated plans into trees of fork plan transformations, which are the trees of descent. The learning algorithms are the functional mechanisms effecting the transformations. Aggregation 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 (the genetic descent network, figures 3).[12]




Figure 2: Plans Learnable from the Top Node Plan




Figure 3: derivation tress folded into a network

The learnability analysis of this paper introduces two novelties: a new principle and a new criterion. Start with an epistemological stance instead of a psychological one. Here one is not so much interested in what a particular child did as an individual. Yet the individual case provides boundary conditions for modeling with a question "if the details of at least one natural case have such characteristics, what kinds and paths of learning are possible?" The question is of general interest if one admits that particularity and egocentricity are common characteristics of novice thought.

SLIM started with the general principle that learning happens through interaction. The model is "mental;" it represents 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. The new principle is that the learnability of a domain is the result of all the possible cases of concrete learning through particular experiences. Co-generativity permits each central strategy (1- 6) to be learned no matter which is adopted as a prototype fork. This suggested a new criterion, that co-generability of related but variant knowledge forms is what makes learning possible in any particular domain.[13] Such contribute significantly to the learnability of a domain because they are mutually reinforcing. Peripheral strategies[14] are rarely learned because they can be learned in few ways. Thus, one can characterize the learnability of a domain as a function of particular interactions among agents based on the connectedness of possible paths of strategy learning. That is what the genetic descent network does. Furthermore, these methods and representations even make it possible to judge that knowledge of a given domain is more learnable than another.

Escaping from Particularity

Most people seem comfortable with the symmetries of 3T, which greatly simplify their analyses and strategies. But the GAC representation, with its cell specific definition of pattern and plan elements, is entirely different and has no way of engaging with common sense notions of symmetry at all.

If we ask where symmetry enters such highly particular descriptions, the answer MUST involve abstraction, but which form of those kinds possible? Abstraction by feature-based classification is the most commonly recognized form, but there are others. Piaget emphasizes a kind of abstraction, focusing more on what one does rather than on qualities one attributes to external things. This reflexive abstraction is a functional analysis of the genesis of some knowledge,[15] as presented elegantly in Bourbaki's description of the generality of axiomatic systems:

"A mathematician who tries to carry out a proof thinks of a well-defined mathematical object, which he is studying just at this moment. If he now believes that he has found a proof, he notices then, as he carefully examines all the sequences of inference, that only very few of the special properties in the object at issue have really played any significant role in the proof. It is consequently possible to carry out the same proof also for other objects possessing only those properties which had to be used. Here lies the simple idea of the axiomatic method: instead of explaining which objects should be examined, one has to specify only the properties of the objects which are to be used. These properties are placed as axioms at the start. It is no longer necessary to explain what the objects that should be studied really are...."

N. Bourbaki, in Fang, p. 69.

Robust data argue that well articulated, reflexive forms of thought are less accessible to children than adults. The possibility that mature, reflexive abstraction is unavailable to naive minds raises this theoretical question: what process of functional abstraction could precede such fully articulated reflexive abstraction; could such a precursor be the kernel from which such a mature form of functional abstraction may grow?

Nil Ex Nihilo: a strategic aside

Genesis once told us "God said 'Let there be light,' and there was light." Cosmologists now say there was no visible light for 300,000 years after the beginning of our universe. They explain the novelty of light's appearance this way: after a period of 'inflation,' the expanding low entropy plasma was so hot that all particles were unstable and there were constant conversions of energy into matter and back again; the density was so great that everything always collided with everything else. As temperature fell through those 300,000 years, protons began to hold captured electrons, matter and anti-matter cross-annihilated, and the quantity of particles dropped enough that photons could escape the collisions and gravity of the still expanding plasma. THEN there was light.[16]

Photons existed from the beginning, in the parts and interactions that made up the whole. Light was revealed by no longer being obscured.

About human learning, we know that it begins with coordination of sensory and motor nerve impulses. On a later, larger scale, the need to coordinate systems of sensory-motor interactions, e.g. "eye and hand" is clear. Why should not the interplay of body-system related, interior representation schemes be invoked to explain processes of thought in pre-linguistic and even in language capable minds? Such is the strategy behind the "Multi-Modal Mind." Nothing comes from nothing. The truly novel is manifest when released from what previously obscured it.

The Multi-modal Mind

Let us discriminate among the major components of the sensori-motor system and their cognitive descendents, even while assuming the preeminence of that system as the basis of mind. Imagine the entire sensori-motor system of the body as made up of a few large, related, but distinct sub-systems, each characterized by the special states and motions of the major body parts, thus:

Body Parts	S-M Subsystem	Major Operations
Trunk	Somatic	Being here; being touched
Legs	Locomotive	Moving from here to there
Head-eyes	Capital/visual	Looking at that there
Arms-hands	Manipulative	Touching; changing that
Ears	Aural	Hearing sounds; Language
Mouth	Oral	Making sounds; Language

We assume the representations of mind remain profoundly affected by the modality of interactions with experience through which it was developed. One implication is that the representations built through experience will involve different objects and relations, among themselves and with externals of the world, which will depend upon the particular mode of experience. Even if atomic units of description (e.g. condition action rules) are shared between modes, the entities which are the salient objects of concern and action are different, and they are in relation to each other only through learned correspondences. This general description of mind contrasts with the more uniformitarian visions which dominate psychology today. These major modal groupings of information structures are imagined to be populated with clusters of related cognitive structures, called "micro-views," with two distinct characters. Some are "task-based" and developed through prior experiences with the external world; others, with a primary character of controlling elements, develop from the relationships and interactions of these disparate, internal micro-views. The issue of cognitive development is cast into a framework of developing control structure within a system of originally competing micro-views.[17]

Redescriptive Abstraction

I propose that the multi-modal structure of the human mind permits development of a significant precursor to reflexive abstraction. The interaction of different modes of the mind in processes of explaining unanticipated outcomes of behavior can alter the operational interpretation and solution of a problem. Eventually, a change of balance can effectively substitute an alternative representation for the original; this could occur if the alternative representation is the more effective in formulating and coping with the encountered problem. In terms of the domain of our explorations and our representations, there is no escape from the particularity of the GAC representation unless some other description is engaged. A description of the same circumstance, rooted in a different mode of experience, would surely have both enough commonality and difference to provide an alternative, applicable description. I identify the GAC absolute grid as one capturing important characteristics of the coordinated visual-manipulative mode;[18] other descriptions based on the somatic or locomotive subsystems of mind could provide alternative descriptions which would by their very nature permit escape from the particularity of the former.

Why should explanation be involved? Peirce argues that "doubt is the motor of thought" and that mental activity ceases when no unanswered questions remain.[19] Circumstances requiring explanation typically involve surprises; the immediate implication is that the result was neither intuitively obvious nor were there adequate processes of inference available beforehand to predict the outcome (at least none such were invoked).

We propose that a different set of functional descriptions, in another modal system, can provide explanation for a set of structures controlling ongoing activity. The initial purpose served by

alternative representations is explanation. Symmetry, however, is a salient characteristic of body centered descriptions; this is the basis of their explanatory power when applied where other descriptions are inadequate. Going beyond explanation, when such an alternative description is applied to circumvent frustrations encountered in play, one will have the alternate structure applied with an emergent purpose. Through such events, the interaction of multiple representations permits a concrete form of abstraction to develop, one emergent from the application of alternative descriptions. Where does symmetry come from? The projection of body centered representations over the visual-manipulative grid.[20]

Emergent Abstraction

If alternative representations can serve as explanation for surprises developed through play, and if they can serve as a bridge to break away from the rigid formulation of the GAC representation, it is not impossible to believe they may begin to provide dynamic guidance as well -- exactly of the sort found useful by adults in their play. When this occurs, the alternative description, useful initially as an explanation for the more particular system of primary experiences, will become the dominant system for play. Then the symmetry implicit in the body-centric imagery will become a salient characteristic of the player's thinking about tic tac toe as the highly specific formulations of early experience recede into the background. Abstraction has taken place -- because the descriptions of the body mode are implicitly less absolute in respect of space than are those supposed to operate with the GAC representation. But the abstraction is not by features, nor is it by the articulate analysis of reflexive abstraction, as described by Bourbaki. This is an emergent abstraction via REDESCRIPTION, a new kind of functional abstraction. Redescriptive abstraction is a primary example of the coadaptive development of cognitive structures. As a kind of functional abstraction which does not yet require reflexive analysis of actions taken within the same mode of representation, but merely the interpretation of actions in one mode in terms of possible, familiar actions in another mode,[21] it needs bear less of an inferential burden than would the more analytic reflexive abstraction described by Bourbaki.

Redescriptive Abstraction and Analogy

One might say that emergent abstraction via redescription is "merely analogy". I propose an antithetical view: emergent abstraction explains why analogy is so natural and so important in human cognition. Redescriptive abstraction is a primary operation of the multi-modal mind; it is the way we must think to explain surprises to ourselves. We judge analogy and metaphor important because redescriptive abstraction is subsumed under those names.

Further, I speculate it is THE essential general developmental mechanism. This process can be the bootstrap for ego-centric cognitive development because accomplished without reference to moves or actions of the other agent of play.

References

Bartlett, Frederic C. *Remembering*. Cambridge University Press, 1932.

Bourbaki, N. Excerpt from *The Architecture of Mathematics*, quoted on p.69 in J. Fang (1970), *Towards a Modern Theory of Mathematics*. Hauppauge, N.Y.: Paideia series in modern mathematics, Vol. 1.

Fann, K. T. *Peirce's Theory of Abduction*. The Hague: Martinus Nijhoff.

- Lawler, R. *Computer Experience and Cognitive Development*. Chichester, England, and New York: Ellis Horwood, Ltd. and John Wiley Inc., 1985.
- Lawler, Robert W. & Oliver G. Selfridge. 1986 Learning Concrete Strategies through Interaction. The *Seventh Cognitive Science Conference*. U.C.Irvine.
- Lawler, R. "Coadaptation and the Development of Cognitive Structures," in *Advances in Artificial Intelligence-II*, Eds. DuBoulay, Hogg, and Steels, North Holland, Amsterdam, 1986. (Papers from ECAI, 1986.)
- Lawler, R. "On the Merits of the Particular Case," in Lawler and Carley, *Case Studies and Computing*. Norwood, N.J.; Ablex Publishing, 1993.
- Lawler, R "Consider the Particular Case" *Journal of Mathematical Behavior*, 1998.
- Minsky, M. A Framework for Representing Knowledge, in *The Psychology of Computer Vision*, (P. Winston, ed.), McGraw Hill, 198n.
- Minsky, M. *The Society of Mind*, Simon & Schuster, NY, 1986.
- Peirce. Charles, S. *The Fixation of Belief*. In Chance, Love, and Logic (M. R. Cohen, ed.). Harcourt, Brace, and Co. 1923.; also George Braziller, Inc., 1956.
- Piaget, J. *Biology and Knowledge*. Chicago: University of Chicago Press, 1971.
- Piaget, J. *The Language and Thought of the Child*. NY: New American Library.
- Sussman, G. *A Procedural Model of Skill Acquisition*, Elsevier Press, 198n.
- Weyl, H. *Symmetry*. Princeton; Princeton University Press, 1952.
- Winston, P. "Learning Structural Descriptions from Examples" in *The Psychology of Computer Vision*, (P. Winston, ed.), McGraw Hill, 198n.

Notes:

1. The materials of these case studies, published and unpublished alike, are now being assembled for public access at the web site www.NLCSA.org, as reported in the paper "With Heart Upon My Sleeve," at Constructionism 2010.
2. Minsky said Oliver had the quickest mind of anyone he had ever known, that he had a genius for undertaking deep studies with simple computational models, and that his interest in children's learning was as committed as my own. Minsky was right on all counts, and I had the deep honor to become Selfridge's colleague for the rest of his life.
3. During its construction, I referred to this corpus as "The Intimate Study." It served as the basis of two books (Lawler 1985 and Lawler et al., 1986, and other papers). Tic tac toe interpretation was completed as a post doc at MIT.
4. As in Patrick Winston's thesis "Learning Structural Descriptions from Examples."
5. As in Gerald Sussman's thesis "A Procedural Model of Skill Acquisition."
6. See Piaget, *The Language and Thought of the Child*.
7. REO's preferences are common in tactics: first, win if possible; block at need; finally choose a free cell: preferring the center cell first, then any corner, and finally a side cell.

8. Subject to rotational symmetry and prototypical strategy preferences of the psychological subject, the generation of games was exhaustive. Analysis focused on won-games in which SLIM moved first in cell 1. See figure 1.
9. 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.
10. 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.
11. 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 version of this text.
12. The complete Genetic Descent Network, with additional GACS, is shown in its fullness on page 19 of *The Merits of the Particular Case*, in Lawler and Carley, 1996. In general, the process followed with these data is similar to Weyl's use of reformulation in his general description of the development of theoretic knowledge in *Symmetry and Bourbaki's description of the genesis of axiomatic systems in The Architecture of Mathematics*.
13. This principle would support stable knowledge in minds with reconstructive memories, such as Bartlett (1932) suggests humans have.
14. Such are those defined by other GACs discussed by Lawler and Selfridge (1986) and in "On the Merit of the Particular Case," Chapter 1 of *Case Study and Computing*.
15. Piaget contrasts reflexive abstraction with classificatory or Aristotelian abstraction (p.320 in *Biology and Knowledge*), demeaning the latter by referring to it as "simple".
16. An elegant popular presentation by Stephen Weinberg of this speculative science is "The First 3 Minutes."
17. This view of mind is presented and applied in "Cognitive Organization", Chapter 5 of Lawler, 1985. A more extensive discussion of micro views appears in Chapter 7.
18. The GAC description is in terms of external things seen by a person referring to it. The absolute reference assigning numbers to specific cells preserves a top-down, left to right organization. Notice however, that even if one's internal representation were different -- based perhaps on a manipulative mode of thought and representation -- the essential points of following arguments remain sound.

19. Peirce's position (presented lucidly in "The Fixation of Belief" but ubiquitous in his writing) was the primary observation leading me to focus on this theme. He uses the term doubt because his discussion is cast in terms of belief; mine, cast in terms of goals, finds its equivalent expression as surprise. Doubts require evidence for elimination (but see Peirce on this); surprises require explanations. Surprise is accessible to mechanical minds as the divergence between expectation and outcome under a specific framework of interpretation.

20. See "Coadaptation and the Development of Cognitive Structures" in DuBoulay et al. *Advances in Artificial Intelligence* for a bit more detail of this argument.

21. The point here is that the process is more like Peirce's abduction than any inductive process of learning. See "Deduction, Induction, and Hypothesis" for Peirce's introduction to this distinction or K. T. Fann's "Peirce's Theory of Abduction" for an analysis of Peirce's developing ideas on abduction.