Spinoza Lecture I

From Socrates to Artificial Intelligence: The Limits of Rule-based Rationality

Hubert L. Dreyfus

Introduction

In these Spinoza Lectures, I'd like to talk about skills – simple skills like driving, game playing, and generally getting around in the world. I'll argue that such skills are the basis of *all intelligibility*. Or, to put it more simply, knowing-how is the basis of all knowing-that. Skills are the basis of the propositional knowledge studied by philosophers since the beginning, and of that branch of psychology, now appropriately called Cognitive Science. Even current interest in whether perception is pre-propositional does not see, as Merleau-Ponty did, that even perception is a skill.

Skills have been ignored in philosophy from Plato to Nietzsche (with the exception of Aristotle) because philosophers assumed without argument, that skills were based on implicit theories. As Leibniz put it:

> [T]he most important observations and turns of skill in all sorts of trades and professions are as yet unwritten. This fact is proved by experience when, passing from theory to practice, we desire to accomplish something. Of course, we can also write up this practice, since it is at bottom *just another theory* more complex and particular.[1]

Heidegger was the first to make explicit that skills are not based on theories, and to place skillful coping at the center of his philosophy. He analyzed *techne* (everyday craftsman's expertise) in Division One of Being and Time, and *phronesis* (ethical expertise) in Division Two. He held that everyday skills contained our understanding of being and were so pervasive as to be invisible, or as he put it, the understanding of being in our everyday skills is the nearest to us and so the furthest away. He further held that the everyday understanding of being in our everyday coping practices is the proper object of phenomenology since phenomenology studies what does not show itself but can be made to show itself, and is the basis of that which does show itself.

But even Heidegger did not work out in detail the development and nature of know-how. So, in this lecture, I'll lay out the phenomenology of everyday skillful coping and suggest how, understanding coping requires us to reconsider the role of detached, rule-based reason in our everyday lives. Next time, I'll discuss practical mastery and how it allows us to understand practical wisdom and an even higher forms of skill.

For now, the best way to see what skills are and are not is to follow the rise and fall of so called expert systems. That requires a bit of background history.

It has been half a century since the computer burst upon the world along with promises that it would soon be programmed to be intelligent, and the related promise or threat that we would soon learn to understand ourselves as computers. In 1947 Alan Turing predicted that there would be intelligent computers by the end of the century. In l968 Marvin Minsky, an AI pioneer at MIT advising Stanley Kubrick who was then making 2001, said, "In 30 years [i.e. by l998] we should have machines whose intelligence is comparable to man's… By 2001 it will be very easy to make computers that appear to understand you and appear to converse with you." Now that 2001 has come and gone, with no sign of computer intelligence, it's time for a retrospective evaluation of the failure of attempts to program computers to be intelligent like HAL in Kubrick's film.

AI research began auspiciously half a century ago when Allen Newell and Herbert Simon showed that computers could do more than calculate. They demonstrated that computers were physical symbol systems whose zeros and ones could be made to stand for anything, including features of the real world, and whose programs could be used as rules for relating these features. In this way computers could be used to simulate certain important aspects of intelligence. Thus the information-processing model of the mind was born.

Then, rather suddenly, the field ran into unexpected difficulties. The trouble started around l970 with the failure of attempts to program children's story understanding. The best AI programs lacked the intuitive common sense of a four-year old.

An old philosophical assumption was at the heart of the problem. What is now called symbolic AI, or sometimes GOFAI [Good Old-fashion AI], is based on an idea that has been around in philosophy since Descartes, that all understanding consists in forming and using appropriate symbolic representations. For Descartes these were complex descriptions built up out of primitive ideas or elements. Kant added the important idea that all concepts were rules. Frege showed that rules could be formalized so that they could be manipulated without intuition or interpretation. Then, naturally, given the formal nature of operations in a computer, AI researchers took up the search for the formal rules and representations underlying intelligence. Symbolic AI thus, in effect, turned analytic philosophy into a research program.

But it turned out to be much harder than anyone expected to formulate, let alone formalize, a theory of common sense. It was not, as Minsky had hoped, just a question of cataloguing a few hundred thousand facts. Just common sense physics, for example, which tries to formalize our sense of how things stick, and pour, and roll, and break, etc. turned out to be too hard. Moreover, what is called the frame problem – how to limit computation only to those facts relevant to the current problem could not be solved in any general way. By the early 1970ies, the common sense knowledge problem became the center of concern. Minsky's optimistic mood changed completely in the course of fifteen years. In 1982 he told a reporter: "The AI problem is one of the hardest science has ever undertaken."[2]

Looking back over these fifty years, it seems that GOFAI, with its promise of a robot like HAL, is a perfect example of what Imre Lakatos has called a "degenerating research program".[3] A degenerating research program is one that starts out with a successful approach to a new domain, but which then runs into unexpected problems it cannot solve, and is finally abandoned by its practitioners. The failure of the research program based on symbolic representations is still reverberating through Cognitive Science. Jerry Fodor, one of the founders of computational psychology, wrote last year that "what our Cognitive Science has found out about the mind is mostly that we don't know how it works.[4]

Given this impasse in AI, it made sense for AI researchers to turn to microworlds - domains isolated from everyday common-sense intuition - and try to develop theories of the skills used in such isolated domains as disease diagnosis or spectrograph analysis.

Thus in the 1980ies, from the frustrating field of AI emerged a new field called Knowledge Engineering, which, by limiting its goals, attempted to use symbolic rules and representations to develop programs that actually worked in the real world. The result was the so-called expert system, enthusiastically promoted in Edward Feigenbaum's book <u>The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World</u>**.**[5] Feigenbaum spells out the goal: "Our machines will have reasoning power: they will automatically engineer vast amounts of knowledge to serve whatever purpose humans propose, from medical diagnosis to product design, from management decisions to education."[6]

What the knowledge engineers claimed to have discovered is that in areas which are cut off from everyday common sense and in which relevance is fixed beforehand, all a machine needs in order to behave like an expert is specialized knowledge of two types: "The facts of the domain that are written in the textbooks and journals of the field, [and] heuristic knowledge, which is the knowledge of good practice and good judgment in a field.” [7]

Using both kinds of knowledge, Feigenbaum developed a program called DENDRAL. It took the data generated by a mass spectrograph and deduced from this data the molecular structure of the compound being analyzed. Another program, MYCIN, took the results of blood tests such as the number of red cells, white cells, sugar in the blood, etc. and came up with a diagnosis of which blood disease is responsible for this condition. In their narrow areas, such programs gave impressive performances.

And, indeed, isn't the success of expert systems just what one would expect? If we agree with Feigenbaum that: "almost all the thinking that professionals do is done by reasoning... "[8] we can see that, once computers are used for reasoning and not just computation, they should be as good or better than we are at following rules and deducing conclusions from a host of facts. So we would expect that, if the rules that an expert has acquired from years of experience could be extracted and programmed, the resulting program would exhibit expertise. Again, Feigenbaum puts the point very clearly: “[T]he matters that set experts apart from beginners, are symbolic and inferential. ... Experts build up a repertory of working rules of thumb, or heuristics, that, combined with book knowledge, make them expert practitioners.”[9] So, since each expert already has a repertory of rules in his mind, all the expert system builder need do is get the rules out of the expert and program them into a computer.

This view is not new. In fact, it goes back to the beginning of Western philosophy when the first philosopher, Socrates, stalked around Athens looking for experts in order to draw out and test their rules. In one of his earliest dialogues, the Euthyphro, Plato tells us of such an encounter between Socrates and Euthyphro, a religious prophet and so a self-proclaimed expert on pious behavior. Socrates asks Euthyphro to tell him how to recognize piety: "I want to know what is characteristic of piety ... to use as a standard whereby to judge your actions and those of other men," he says. But, instead of revealing his piety-recognizing heuristic, Euthyphro does just what every expert does when cornered by Socrates. He gives him examples from his field of expertise, in this case mythical situations in the past in which men and gods have done things that everyone agrees are pious. Socrates gets annoyed and demands that Euthyphro tell him his rule for recognizing these cases as examples of piety, but although Euthyphro claims he knows how to tell pious acts from impious ones, he cannot state the rules which generate his judgments. Socrates ran into the same problem with craftsmen, poets and even statesmen.

None could articulate the rules underlying their expertise. Socrates therefore concluded that none of these experts knew anything and, since he could articulate any rules for his own practices, he concluded that he didn't know anything either.

That might well have been the end of Western philosophy, but Plato admired Socrates and understood his problem. So he developed an account of what caused the difficulty. Experts, at least in areas involving non-empirical knowledge such as morality and mathematics, had, in another life, Plato claimed, learned the rules involved, but they had forgotten them. The role of the philosopher was to help such moral and mathematical experts recollect the principles on which they were acting.

Knowledge engineers would now say that the rules experts - even experts in empirical domains - use have been put in a part of their mental computers where they work automatically. Feigenbaum says: “When we learned how to tie our shoes, we had to think very hard about the steps involved ... Now that we've tied many shoes over our lifetime, that knowledge is ‘compiled,’ to use the computing term for it; it no longer needs our conscious attention”.[10] On this Platonic view, the rules are there functioning in the expert's mind whether he is conscious of them or not. How else could one account for the fact that the expert can still perform the task? After all, we can still tie our shoes, even though we no longer can say how we do it. So nothing has changed. Only now, 2400 years after Socrates, thanks to Feigenbaum and his colleagues, we have a new name for what Socrates was doing. Instead of aiding recollection, we can now say he was doing knowledge acquisition research.[11]

But although philosophers and knowledge engineers have long been convinced that expertise is based on applying sophisticated heuristic rules to masses of facts, there are few available rules. As Feigenbaum explains: “[A]n expert's knowledge is often ill-specified or incomplete because the expert himself doesn't always know exactly what it is he knows about his domain”.[12] Indeed, when Feigenbaum suggests to an expert the rules the expert seems to be using, he gets a Euthyphro-like response. "That's true, but if you see enough patients/rocks/chip designs/instruments readings, you see that it isn't true after all,"[13] and Feigenbaum comments with Socratic annoyance: "At this point, knowledge threatens to become ten thousand special cases."[14]

And so, despite the promises and enthusiasm of hundreds of knowledge engineers, there are no systems programmed to follow the rules used by experts in various domains that behave like experts. Indeed, although computers are faster and more accurate than people in storing facts and applying rules, expert performance has remained out of reach. In each area where there are experts with years of experience, the rule-following computer can do better than the beginner, and can even exhibit useful competence, but it cannot rival the very experts whose facts and supposed heuristic rules it is processing with incredible speed and unerring accuracy.

The success of Deep Blue and Deep Junior may seem exceptions to this claim but they are not. Chess programs only reached Grand Master level when computers were so fast they could look at 200 million moves a second, and thereby look as many as fifteen moves ahead. Then they did not need to rely on rules obtained from masters who, it turns out, can consider only 200 to 300 moves.

In the face of this impasse, in spite of the authority and influence of Plato and 2400 years of philosophy, we must take a fresh look at what a skill is and what the expert acquires as he achieves expertise. We must be prepared to abandon the traditional view that runs from Plato to Descartes, to Leibniz to Kant, to Piaget to Chomsky to Feigenbaum, that a beginner starts with specific cases and, as he or she becomes more proficient, abstracts and interiorizes more and more sophisticated rules. It might turn out that skill acquisition moves in just the opposite direction: from abstract rules to particular cases.

Many of our skills are acquired at an early age by trial and error or by imitation, but to make the phenomenology of skillful behavior as clear as possible let's look at how, as adults we learn new skills by instruction.[15] I'll use examples from driving and chess – a bodily, practical skill and an intellectual skill -- to suggest the generality of the model. As Heidegger says, we should let the phenomenon show itself as it is in itself, so let's try.

Stage 1: Novice

Normally, the instruction process begins with the instructor decomposing the task environment into context-free *features* that the beginner can recognize without previous experience in the task domain. The beginner is then given rules for determining actions on the basis of these features, like a computer following a program.

The student automobile driver learns to recognize such interpretation-free features as speed (indicated by the speedometer) and is given rules such as shift to second when the speedometer needle points to ten miles an hour.

The novice chess player learns a numerical value for each type of piece regardless of its position, and the rule: "Always exchange if the total value of pieces captured exceeds the value of pieces lost." The player also learns to seek center control when no advantageous exchanges can be found, and is given a rule defining center squares and a rule for calculating extent of control. Most beginners are notoriously slow players, as they attempt to remember all these features and rules.

Stage 2: Advanced Beginner

As the novice gains experience actually coping with real situations, he begins to note, or an instructor points out, perspicuous examples of meaningful additional *aspects* of the situation. After seeing a sufficient number of examples, the student learns to recognize these new aspects. Instructional *maxims* can then refer to these new situational aspects, as well as to the objectively defined non-situational features recognizable by the inexperienced novice. Unlike a rule, a maxim requires that one already have some understanding of the domain to which the maxim applies.

The advanced beginner driver, using (situational) engine sounds as well as (non-situational) speed in his gear-shifting rules, learns the maxim: Shift up when the motor sounds like it is racing and down when it sounds like its straining. Engine sounds cannot be adequately captured by a list of features. In general, features cannot take the place of a few choice examples in learning such distinctions.

With experience, the chess beginner learns to recognize such situational aspects as an overextended position, a weakened king's side, or a strong pawn structure despite the lack of a precise and situation-free definition. The player can then follow maxims such as: attack a weakened king's side.

Stage 3: Competence

With more experience, the number of potentially relevant elements and procedures that the learner is able to recognize and follow becomes overwhelming. At this point, since a sense of what is important in any particular situation is missing, performance becomes nerve-wracking and exhausting, and the student may well wonder how anyone ever masters the skill.

To cope with this overload and to achieve competence, people learn through instruction or experience, to devise a plan or choose a perspective. The perspective then determines which elements of the situation should be treated as important and which ones can be ignored. Restricting attention to only a few of the vast number of possibly relevant features and aspects, makes decision making easier.

Naturally, to avoid mistakes, the competent performer seeks new rules and reasoning procedures to decide which plan or perspective to adopt. But such rules are not as easy to come by as are the rules and maxims in manuals and lectures. There are just too many situations differing from each other in too many subtle ways. More situations, in fact, than can be named or precisely defined, so no one can prepare for the learner a list of possible situations and what to do or look for in each. Competent performers, therefore, must decide for themselves in each situation what plan or perspective to adopt and when to adopt it without being sure that it will be appropriate.

Given this uncertainty, coping becomes frightening rather than exhausting. Prior to this stage, if the learned rules didn't work out, the performer, rather than feeling remorse for his mistakes, could rationalize that he hadn't been given adequate rules. But since at this stage, the result depends crucially on the perspective adopted by the learner, the learner feels responsible for how things turn out. Often the choice leads to confusion and failure, but sometimes, things work out well, and the competent performer experiences a kind of elation unknown to the beginner. Thus, learners find themselves on an emotional roller coaster.

A competent driver leaving the freeway on an off-ramp curve, learns to pay attention to the speed of the car, not whether to shift gears. After taking into account speed, surface condition, criticality of time, etc., he may decide he is going too fast. He then has to decide whether to let up on the accelerator, remove his foot altogether, or step on the brake and precisely when to do so. He is relieved if he gets through the curve without mishap and shaken if he begins to go into a skid.

The class A chess player, here classed as competent, may decide after studying a position that her opponent has weakened his king's defenses so that an attack against the king is a viable goal. If she chooses to attack, she can ignore features involving weaknesses in her own position created by the attack as well as the loss of pieces not essential to the attack. Pieces defending the enemy king become salient and their removal is all that matters. Since pieces not involved in the attack are being lost, the timing of the attack is critical. If she attacks too soon or too late her pieces will have been lost in vain and she will almost surely lose the game. Successful plans induce euphoria, while mistakes are felt in the pit of the stomach.

As the competent performer become more and more emotionally involved in his tasks, it becomes increasingly difficult to draw back and to adopt the detached rule and maxim following stance of the advanced beginner. Hilary Putnam, in his Spinoza Lecture, claimed that philosophy has, from Socrates on called for “reflective transcendence.” He holds that the moral of the Euthyphro is that, if one is to act wisely, one must “stand back” and be critical. But while it might seem to traditional philosophers like Putnam that involvement would interfere with detached reasoning and so lead to irrational decisions and inhibit further skill development, in fact just the opposite seems to be the case. Patricia Benner has studied nurses at each stage of skill acquisition. She finds that, unless the trainee stays emotionally involved and accepts the joy of a job well done, as well as the remorse of mistakes, he or she will not develop further, and will eventually burn out trying to keep track of all the features and aspects, rules and maxims that modern medicine requires.

In general, the attempt to stand back leads to stagnation and ultimately to boredom and regression.[16] Indeed, only if the detached rule-following stance of the novice and advanced beginner is replaced by involvement, is the learner capable of further skill advancement.

Stage 4: Proficient

If, however, the learner gives up the detached deliberations of the first three stages and experiences events with involvement, the resulting positive and negative experiences will strengthen successful responses and inhibit unsuccessful ones. The performer's theory of the skill, as represented by rules and features, will then gradually be replaced by situational discriminations accompanied by associated responses. Proficiency seems to develop if, and only if, experience is assimilated in this atheoretical way, that is, only if spontaneous behavior and emotional feedback replaces reasoned responses. Then, as the brain of the performer acquires the ability to discriminate among a variety of situations, each entered into with concern and involvement, appropriate plans spring to mind and certain aspects of the situation stand out as important without the learner standing back and choosing those plans or deciding to adopt that perspective. Action becomes less stressful as the learner simply sees what needs to be done rather than deciding, by a calculative procedure, which of several possible alternatives should be selected.

To understand this stage of skill acquisition, we must note that the involved, experienced performer sees goals and salient aspects, but not what to do to achieve these goals. This is inevitable because the proficient performer’s repertoire of situations he can discriminate is too crude. He has not yet had enough experiences to be able to pair each of the many possible responses he can make with the specific type of situation in which that response would be appropriate. Thus, the proficient performer, after *spontaneously* seeing the issue and the important aspects of the situation, must still *decide* what to do, and to decide he must fall back on detached rule-and-maxim following.

The proficient driver, approaching a curve on a rainy day, may feel in the seat of his pants that he is going dangerously fast. He must then decide whether to apply the brakes or merely to reduce pressure on the accelerator by some selected amount. Valuable time may be lost while he is working out a decision, but the proficient driver is certainly more likely to negotiate the curve safely than the competent driver who must first consider the speed, angle of bank, and felt gravitational forces, in order to decide whether he is going too fast.

The proficient chess player, who is classed a Master, can recognize almost immediately a large repertoire of types of positions. She then deliberates to determine which move will best achieve her goal. She may *sense*, for example, that she should attack, but she must *calculate* how best to do so.

Stage 5: Expert

The expert, however, not only sees what needs to be achieved; thanks to a vast repertoire of situational discriminations, he sees immediately how to achieve his goal. The ability to make more subtle and refined discriminations is what distinguishes the expert from the proficient performer. The expert has learned to distinguish among many types of situation, all seen from the same perspective but requiring different tactical decisions, those situations requiring one response from those demanding another. This allows the immediate intuitive situational response that is characteristic of expertise.

The chess Grandmaster, experiences a compelling sense of the issue and the best move. Excellent chess players can play at the rate of 5 to 10 seconds a move and even faster without any serious degradation in performance. At this speed they must depend almost entirely on intuition and hardly at all on analysis and comparison of alternatives.

A few years ago my brother and I performed an experiment in which an International Master, Julio Kaplan, was required to add numbers presented to him audibly at the rate of about one number per second as rapidly as he could, while at the same time playing five-second-a-move chess against a slightly weaker, but master level player. Even with his analytical mind completely occupied by adding numbers, Kaplan more than held his own against the Master in a series of games. Deprived of the time necessary to see problems or construct plans, Kaplan still produced fluid and coordinated play.

Kaplan's performance seems somewhat less amazing when one realizes that a chess position is as meaningful, interesting, and important to a professional chess player as a face in a receiving line is to a professional politician. Almost anyone could add numbers and simultaneously recognize and respond to familiar faces, even though each face will never exactly match the same face seen previously, and politicians can recognize thousands of faces, just as Julio Kaplan can recognize thousands of chess positions similar to ones previously encountered. The number of classes of discriminable situations, built up on the basis of experience, must be immense. It has been estimated that a master chess player can distinguish roughly 100,000 types of positions.

Driving probably involves the ability to discriminate a similar number of typical situations. The expert driver, not only feels in the seat of his pants when speed is the issue; he knows how to perform the appropriate action without calculating and comparing alternatives. On the off-ramp, his foot simply lifts off the accelerator or applies the appropriate pressure to the brake. What must be done, simply is done.

We can now see that a beginner reasons using rules and facts just like a heuristically programmed computer, but that with talent and a great deal of involved experience, the beginner develops into an expert who intuitively sees what to do without recourse to rules. The tradition has given an accurate description of the beginner and of the expert facing an unfamiliar situation, but normally an expert doesn't calculate. He doesn't solve problems. He doesn't think. He doesn't even need to be conscious of what he is doing. As basketball player, Larry Bird, says: "[A lot of the] things I do on the court are just reactions to situations …. A lot of time, I've passed the basketball and not realized I've passed it until a moment or so later."

The expert just does what normally works and, of course, it normally works. As Aristotle says, the expert "straightway" does "the appropriate thing, at the appropriate time, in the appropriate way."

Conclusion:

The description of skill acquisition I've presented enables us to understand why knowledge engineers from Socrates to Feigenbaum have had such trouble getting the expert to articulate the rules he is using. The expert is simply not following any rules! He is doing just what Socrates feared he might be doing -- discriminating thousands of typical cases.

This in turn explains why expert systems are never as good as experts. If one asks an expert for the rules he is using one will, in effect, force the expert to regress to the level of a beginner and state the rules he learned in school. Thus, instead of using rules he no longer remembers, as the knowledge engineers suppose, the expert is forced to remember rules he no longer uses.[17] If one programs these rules into a computer, one can use the speed and accuracy of the computer and its ability to store and access millions of facts to outdo a human beginner using the same rules. But such systems are at best competent. No amount of rules and facts can capture the knowledge an expert has when the neural network in his brain has been tuned by his experience of the actual outcomes of tens of thousands of situations.

We can also see why the common sense knowledge problem stumped AI researchers and why, as Fodor sees, the rule based models of how the mind and brain work cannot explain intelligence. The basis of common sense is our skill for coping with everyday materials and events. It is a knowing-how, not a knowing-that. For example, in acquiring a knowledge of common sense physics, children needs years of experience. It is striking how long they play with water and mud, presumably not in order to abstract rules for how objects and liquids behave, but to learn to discriminate the necessary thousands of typical cases. And, of course, children need to have even more experience to learn to cope skillfully with everyday situations and understand stories about them.

The primacy of skillful coping over theory thus explains why what I predicted in the sixties, thanks to the work of Heidegger and Merleau-Ponty, not what Minsky predicted, actually happened. By 2001, there were no intelligent robots like HAL. Rather, the computational model of mind in Cognitive Psychology was in serious trouble, and the AI and Expert Systems research programs, taken over from philosophers such as Leibniz, who thought that skills were obviously unconscious theories, had been almost totally abandoned.

The following material was not delivered in my lecture but can be included if you wish.

Given this idealized account of skillful expert coping, it might seem that experts don't need to think and are always right. Such, of course, is not the case. While most expert performance is ongoing and nonreflective, the best of experts, when time permits, think before they act. Normally, however, they don't think about their reasons for choosing possible actions, since, if they did, they, would regress to the competent level. Rather, they reflect upon the goal or perspective that seems evident to them and upon the action that seems appropriate to achieving that goal.

Let us call the kind of inferential reasoning exhibited by the novice, advanced beginner and competent performer as they apply and improve their theories and rules, "calculative rationality", and what experts exhibit when they have time, "deliberative rationality." Deliberative rationality is detached, reasoned observation of one's intuitive, practice-based behavior with an eye to challenging, and perhaps improving, intuition without replacing it by the purely theory-based action of the novice, advanced beginner or competent performer.

For example, sometimes, due to a sequence of events, one is led to see a situation from an inappropriate perspective. Seeing an event in one way rather than some other almost-as-reasonable way, can lead to seeing a subsequent event in a way quite different from how that event would have been interpreted had the second perspective been chosen. After several such choices one can have a totally different view of the situation than one would have had if, at the start, a different reasonable perspective had been chosen. Getting locked into a particular perspective when another one is equally reasonable or more reasonable is called "tunnel vision." An expert will try to protect against this by trying to see the situation in alternative ways, sometimes through reflection and sometimes by consulting others and trying to be sympathetic to their perhaps differing views. The important point for this lecture is that the expert uses intuition not calculation even in reflection.

If this were merely an academic discussion, I could conclude here, simply correcting the traditional account of expertise by replacing calculative with deliberative rationality; if it were merely a matter of business, we could all sell our stock in expert systems companies. Indeed, it turns out that would have been a good idea, since they have all failed. But we can’t be so casual. The Socratic picture of rule-based reason underlies a general movement towards calculative rationality in our culture, and that movement brings with it great dangers.

The increasingly bureaucratic nature of society is heightening the risk that, in the future, skill and expertise will be lost through over reliance on calculative rationality. Today, as always, individual decision-makers stay involved and respond to their situation intuitively as described in the highest levels of my skill acquisition model. But when more than one person is involved in a decision, the success of science and the availability of computers tend to favor the detached mode of problem description characteristic of calculative rationality. One wants a decision that affects the public to be explicit and logical so that rational discussion can be directed toward the relevance and validity of the rules and features used in the analysis. But, as we have seen, with experience comes a decreasing concern with accurate assessment of isolated elements. In evaluating elements, experts have no expertise.

For example, judges and ordinary citizens serving on our juries are beginning to distrust anything but "scientific" evidence. A ballistics expert who testified only that he had seen thousands of bullets and the gun barrels that had fired them, and that there was absolutely no doubt in his mind that the bullet in question had come from the gun offered in evidence, would be ridiculed by the opposing attorney and disregarded by the jury. Instead, the expert has to talk about the individual marks on the bullet and the gun and connect them by rules showing that only the gun in question could so mark the bullet. But in this he is no expert. If he is experienced in legal proceedings, he will know how to construct arguments that convince the jury, but he does not tell the court what he intuitively knows, for he will be evaluated by the jury on the basis of his "scientific" rationality, not in terms of his past record of good judgment. As a result some wise but honest experts are ignored, while non-expert authorities who are experienced at producing convincing legal testimony are much sought after. The same thing happens in psychiatric hearings, medical proceedings, and other situations where technical experts testify. Form becomes more important than content.

It is ironic that judges hearing a case will expect expert witnesses to rationalize their testimony, for, when rendering a decision involving conflicting conceptions of what is the central issue in a case and therefore what is the appropriate guiding precedent, judges will rarely if ever attempt to explain their choice of precedents. They presumably realize that they know more than they can explain and that ultimately unrationalized intuition must guide their decision-making, yet lawyers and juries seldom accord witnesses the same prerogative.

In each of these areas and many more, calculative rationality, which is sought for good reasons, means a loss of expertise. But in facing the complex issues before us we need all the wisdom we can find. Therefore, society must clearly distinguish its members who have intuitive expertise from those who have only calculative rationality. It must encourage its children to cultivate their intuitive capacities in order that they may achieve expertise, not encourage them to reason by rule and thereby become human logic machines. In general, to preserve expertise we must foster intuition all levels of decision making, otherwise wisdom will become an endangered species of knowledge.

---

[1] Leibniz, Selections, ed. Philip Wiener (New York: Scribner, 1951), p. 48. My italics.

[2] Gina Kolata, "How Can Computers Get Common Sense?", Science, Vol. 217, 24 Sept. 1982, p. 1237.

[3] Imre Lakatos, Philosophical Papers, ed. John Worrall, Cambridge University Press, 1978.

[4] Jerry Fodor, The mind doesn't work that way, (Cambridge, MA: MIT. Press, 2001) p. 100.

[5] Edward Feigenbaum and Pamela McCorduck, The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World, Addison-Wesley, l983.

[6] Ibid, p. 56.

[7] Ibid. pp. 76-77.

[8] Ibid. p. l8.

[9] Ibid., p. 64.

[10] Ibid., p. 55.

[11] Ibid., p. 79.

[12] Ibid., p. 85.

[13] Ibid., p. 82.

[14] Ibid.

[15] For a detailed treatment of the phenomenology of skill acquisition, see H. Dreyfus and S. Dreyfus, Mind Over Machine, Free Press, 1988.

[16] Patricia Benner has described this phenomenon in From Novice to Expert: Excellence and Power in Clinical Nursing Practice, Addison-Wesley, l984, p. 164. Furthermore, failure to take risks leads to rigidity rather than the flexibility we associate with expertise. When a risk-averse person makes an inappropriate decision and consequently finds himself in trouble, he tries to characterize his mistake by describing a certain class of dangerous situations and then makes a rule to avoid them in the future. To take an extreme example, if a driver, hastily pulling out of a parking space, is side-swiped by an oncoming car he mistakenly took to be approaching too slowly to be a danger, he may resolve to follow the rule, never pull out if there is a car approaching. Such a rigid response will make for safe driving in a certain class of cases, but it will block further skill refinement. In this case, it will prevent acquiring the skill of flexibly pulling out of parking places. In general, if one seeks to follow general rules one will not get beyond competence. Progress is only possible if, responding quite differently, the driver accepts the deeply felt consequences of his action without detachedly ask himself what went wrong and why. If he does this, he is less likely to pull out too quickly in the future, but he has a much better chance of ultimately becoming, with enough frightening or, preferably, rewarding experiences, a flexible, skilled driver.

One might object that this account has the role of emotions reversed; that the more the beginner is emotionally committed to learning the better, while an expert could be, and, indeed, often should be, coldly detached and rational in his practice. This is no doubt true, but the beginner's job is to follow the rules and gain experience, and it is merely a question of motivation whether he is involved or not. Furthermore, the novice is not emotionally involved in *choosing* an action, even if he is involved in its outcome. Only at the level of competence is there an emotional investment in the *choice of action.* Then emotional involvement seems to play an essential role in switching over from what one might roughly think of as a left-hemisphere analytic approach to a right-hemisphere holistic one. Of course, not just any emotional reaction such as enthusiasm, or fear of making a fool of oneself, or the exultation of victory, will do. What matters is taking responsibility for one's successful and unsuccessful choices, even brooding over them; not just feeling good or bad about winning or losing, but replaying one's performance in one's mind step by step or move by move. The point, however, is not to *analyze* one's mistakes and insights, but just to *let them sink in*. Experience shows that only then will one become an expert. After one becomes an expert one can rest on one's laurels and stop this kind of obsessing, but if one is to be the kind of expert that goes on learning, one has to go on dwelling emotionally on what critical choices one has made and how they affected the outcome.

[17] Or else the expert will try to resist, as Socrates found with Euthyphro, by offering examples of typical cases he remembers from his advanced beginner stage.