

# Hubert L.Dreyfus, Stuart E. Dreyfus

From Socrates to Expert Systems:

The Limits and Dangers of Calculative Rationality

[http://socrates.berkeley.edu/~hdreyfus/html/paper\\_socrates.html](http://socrates.berkeley.edu/~hdreyfus/html/paper_socrates.html)

\*\*\*\*\*



## Hubert Dreyfus

Professor of Philosophy in the Graduate School

- [Home](#)
- [Courses](#)
- [Selected Papers](#)
- [Selected Books](#)
- [C.V.](#)
- [Dreydegger.org](#)
- [Phil. Faculty](#)
- [Dept. Philosophy](#)
- [UC Berkeley](#)

### From Socrates to Expert Systems: The Limits and Dangers of Calculative Rationality

Hubert L. Dreyfus  
Stuart E. Dreyfus

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. Now with the millennium only three years away, it is time for a retrospective evaluation of the attempt to program computers to be intelligent like HAL in the movie 2001.

Actual AI research began auspiciously around 1955 with Allen Newell and Herbert Simon's work at the RAND Corporation. Newell and Simon proved that computers could do more than calculate. They demonstrated that computers were physical symbol systems whose symbols 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 intelligence. Thus the information-processing model of the mind was born. But, looking back over these fifty years, it seems that theoretical AI with its promise of a robot like HAL appears to be a perfect example of what

Imre Lakatos has called a "degenerating research program".

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. Newell and Simon's early work on problem solving was, indeed, impressive, and by 1965 Artificial Intelligence had turned into a flourishing research program, thanks to a series of micro-world successes such as Terry Winograd's SHRDLU, a program that could respond to English-like commands by moving simulated, idealized blocks. The field had its own Ph.D. programs, professional societies and gurus. It looked like all one had to do was extend, combine, and render more realistic the micro-worlds and one would have genuine artificial intelligence. Marvin Minsky, head of the M.I.T. AI Laboratory, predicted in 1967 that "within a generation the problem of creating 'artificial intelligence' will be substantially solved."

Then, rather suddenly, the field ran into unexpected difficulties. The trouble started, as far as we can tell, around 1970 with the failure of attempts to program children's story understanding. The programs lacked the intuitive common sense of a four-year old. And no one knew what to do about it. [Insert birthday party, frames, and the commonsense knowledge problem.]

[An old philosophical dream was at the heart of the problem. AI is based on an idea which 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. Given the nature of computers, AI took up the search for such formal rules and representations. Common-sense-intuition had to be understood as some vast collection of rules and facts.]

It simply turned out to be much harder than one

expected to formulate, let alone formalize, the required theory of common sense. It was not, as Minsky had hoped, just a question of cataloguing a few hundred thousand facts. The common sense knowledge problem became the center of concern. Minsky's 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." [CYC. Can you whistle and eat at the same time?]

Given this impasse, it made sense a year later for researchers to return to microworlds - domains isolated from everyday common-sense intuition - and try to develop theories of at least such isolated domains. This is actually what happened - with the added realization that such isolated domains need not be games like chess nor micro-worlds like Winograd's blocks world, but could, instead, be skill domains like disease diagnosis or spectrograph analysis.

Thus, from the frustrating field of AI has recently emerged a new field called knowledge engineering, which by limiting its goals has applied AI research in ways that actually work in the real world. The result is the so called expert system, enthusiastically promoted in Edward Feigenbaum's book The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World. Feigenbaum spells out the goal:

The 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.

What the knowledge engineers claim to have discovered is that in areas which are cut off from everyday common sense and social intercourse, all a machine needs in order to behave like an expert is specialized knowledge of two types:

The facts of the domain - the widely shared knowledge ... that is written in textbooks and journals of the field {and] heuristic knowledge, which is the knowledge of good practice and good

judgment in a field.

Using both kinds of knowledge Feigenbaum developed a program called DENDRAL. It takes the data generated by a mass spectrograph and deduces from this data the molecular structure of the compound being analyzed. Another program, MYCIN, takes the results of blood tests such as the number of red cells, white cells, sugar in the blood, etc. and comes up with a diagnosis of which blood disease is responsible for this condition. It even gives an estimate of the reliability of its own diagnosis. In their narrow areas, such programs give impressive performances. They seem to confirm the observation of Leibniz, the grandfather of expert systems. He observed that:

[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.*

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..." 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 for deducing conclusions from a host of facts. So we would expect that if the rules which 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, inferential, and rooted in experiential knowledge. ... Experts build up a repertory of working rules of thumb, or "heuristics," that, combined with book knowledge, make them expert practitioners.

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 culture 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 an 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." 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 which everyone considers pious. Socrates gets annoyed and demands that Euthyphro, then, tell him his rules 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. They also could not articulate the principles underlying their expertise. Socrates therefore concluded that none of these experts knew anything and he didn't know anything either.

That might well have been the end of Western philosophy, but Plato admired Socrates and saw 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 said, learned the principles 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 acted. 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.

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 cannot say how we do it. So nothing has changed. Only now 2400 years later, thanks to Feigenbaum and his colleagues, we have a new name for what Socrates and Plato were doing: knowledge acquisition research.

But although philosophers and knowledge engineers have become convinced that expertise is based on applying sophisticated heuristics 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.

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," and Feigenbaum comments with Socratic annoyance: "At this point, knowledge threatens to become ten thousand special cases."

There are also other hints of trouble. Ever since the inception of Artificial Intelligence, researchers have been trying to produce artificial experts by programming the computer to follow the rules used by masters in various domains. Yet, although computers are faster and more accurate than people in applying rules, master-level performance has remained out of reach.

[Arthur Samuel's work is typical. In 1947, when electronic computers were just being developed, Samuel, then at IBM, decided to write a checker playing program. He elicited heuristic rules from checker masters and programmed a computer to follow these rules. The resulting checkers program is not only the first and one of the best expert systems

ever built; it is also a perfect example of the way fact turns into fiction in AI. Feigenbaum, for example, reports that "by 1961 [Samuel's program] played championship checkers, and it learned and improved with each game." In fact, Samuel said in an interview at Stanford University, where he is a retired professor, that the program did once defeat a state champion, but the champion "turned around and defeated the program in six mail games." According to Samuel, after 35 years of effort, "the program is quite capable of beating any amateur player and can give better players a good contest." It is clearly no champion. Samuel is still bringing in expert players for help but he "fears he may be reaching the point of diminishing returns." This does not lead him to question the view that the masters the program cannot beat are using heuristic rules; rather, like Plato and Feigenbaum, Samuel thinks that the experts are poor at recollecting their compiled heuristics. "The experts do not know enough about the mental processes involved in playing the game," he says.

The same story is repeated in every area of expertise, even in areas unlike checkers where expertise requires storing large numbers of facts, which should give an advantage to the computer.] In each area where there are experts with years of experience, the computer can do better than the beginner, and can even exhibit useful competence, but it cannot rival the very experts whose facts and supposed heuristics it is processing with incredible speed and unerring accuracy.

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 when he achieves expertise. We must be prepared to abandon the traditional view that run from Plato to Piaget and Chomsky that a beginner starts with specific cases and, as he 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. Since we are all experts in many areas, we have the necessary data.

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 learned new skills by instruction.

### Stage 1: Novice

Normally, the instruction process begins with the instructor decomposing the task environment into context-free features which 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.

For purposes of illustration, let us consider two variations: a bodily or motor skill and an intellectual skill. 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 now can refer to these new *situational aspects*, as well as to the objectively defined *nonsituational features* recognizable by the inexperienced novice.

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. He learns to observe the demeanor as well as position and velocity of pedestrians or other drivers. He can, for example, distinguish the behavior of a distracted or drunken driver from that of an impatient but alert one. Engine sounds and behavior styles cannot be adequately captured by words, so words cannot take the place of a few choice examples in learning such distinctions.

With experience, the chess beginner learns to recognize such situational aspects of positions as 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 that the learner is able to recognize 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. By restricting attention to only a few of the vast number of possibly relevant features and aspects, such a choice of a perspective makes decision making easier.

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

situation.

Coping thus becomes frightening rather than exhausting. Prior to this stage, if the learned rules didn't work out, the performer could rationalize that he hadn't been given adequate rules rather than feel remorse because of his mistake. Now, however, the learner feels responsible for disasters. Of course, often at this stage, 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, after taking into account speed, surface condition, criticality of time, etc., 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 being honked at 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. 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-following stance of the beginner. While it might seem that this involvement would interfere with detached rule-testing and so would inhibit further skill development, in fact just the opposite seems to be the case. As we shall soon see, if the detached rule-following stance of the novice and advanced beginner is replaced by involvement, one is set for further advancement, while resistance to the acceptance of risk and responsibility can lead to stagnation and ultimately to boredom and regression. [Patricia Benner's work.]

[For example, a competent driver no longer merely follows rules designed to enable him or her to operate a vehicle safely and courteously, but begins a trip by selecting a goal. If a driver wishes to get somewhere very quickly, for example, comfort and courtesy play a diminished role in the selection of maneuvers and slightly greater risks might be accepted. Driving in this manner, the driver feels pride if the trip is completed quickly and uneventfully, and remorse follows arrest or near collision. Should the trip involve, say, an incident in which the driver passes another car dangerously so that only quick action by the other driver prevents an accident, the competent driver can respond to this experience in one of two qualitatively different ways. One response would be for the driver to decide that one should never rush and so the driver would modify the rule used to decide to hurry. Or, perhaps, the frightened driver would modify the rule for conditions for safe passing so that the driver only passes under exceedingly safe circumstances. These would be the approaches of the driver doomed to timidity and fear, and, by our definition, to incompetence. Or, responding quite differently, the driver might accept the deeply felt consequences of the act but not detachedly ask himself what went wrong and, especially, why. If one does this, it is likely that one won't be quite so likely to hurry in the future or to pass in similar situations, but one has a much better chance of ultimately becoming, with enough frightening or, preferably, rewarding experiences, a relaxed and expert driver.]

#### Stage 4: Proficient

If events are experienced with involvement as the learner practices her skill, 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 principles, will thus 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 and intuitive behavior replaces reasoned responses.

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 easier and less stressful as the learner simply sees what needs to be achieved rather than deciding, by a calculative procedure, which of several possible alternatives should be selected. There is less doubt that what one is trying to accomplish is appropriate when the goal is simply obvious rather than the winner of a complex competition. In fact, at the moment of involved intuitive response, there can be no doubt, since doubt comes only with detached evaluation.

Remember that the involved, experienced performer sees goals and salient aspects, but not what to do to achieve these goals. This is inevitable since there are far fewer ways of seeing what is going on than there are ways of responding. Thus, the proficient performer, after seeing the goal and the important features of the situation, must still *decide* what to do. To decide, he falls back on detached rule-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 spends additional time *considering* the speed, angle of bank, and felt gravitational forces, in order to decide whether the car's speed is excessive.

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, for example, know that she should attack, but she must calculate how best to do so.

#### Stage 5: Expert

The *proficient performer*, immersed in the world of his skillful activity, *sees* what needs to be done, but must *decide* how to do it. The *expert* not only sees

what needs to be achieved; thanks to a vast repertoire of situational discriminations he sees 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 situations, all seen as similar by the proficient performer, those situations requiring one action from those demanding another. That is, with enough experience in a variety of situations, all seen from the same perspective but requiring different tactical decisions, the brain of the expert performer gradually decomposes this class of situations into subclasses, each of which shares the same action. This allows the immediate intuitive situational response that is characteristic of expertise.

The expert chess player, classed as an international master or 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 Stuart 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 can add numbers and simultaneously recognize and respond to 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 50,000 types of positions.

Driving probably involves the ability to discriminate a similar number of typical situations. The expert driver not only feels when slowing down on an off ramp is required; he simply performs the appropriate action. What must be done, simply is done.

We can see now that a beginner calculates 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 does not calculate. He does not solve problems. He does not even think. He just does what normally works and, of course, it normally works.

The description of skill acquisition I have presented enables us to understand why the knowledge engineers from Socrates, [to Samuel], 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 and Feigenbaum feared he might be doing -- discriminating thousands of special 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. 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 he has stored his experience of the actual outcomes of tens of thousands of situations.

[This in turn explains the common sense knowledge problem. The basis of common sense is our skill for coping with everyday materials. It is a knowing-how, not as knowing-that. (Example common sense physics. Children play with water for years building up the necessary thousands of typical cases.) This would explain why research in AI has been stalled and why we should expect the attempt to make intelligent computers by using rules and features to be abandoned by the end of this century.]

[In this idealized account of skillful expert coping it might seem that experts needn't 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 rules for choosing goals or 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 events 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 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 phenomena suggest that the expert uses intuition not calculation even in reflection.

If this were merely an academic discussion, we 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 sell out stock in expert systems companies. Indeed, it turns out that would have been a good idea, since they have almost all gone out of business. But we cannot be so casual. The Socratic picture of 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 danger that in the future skill and expertise will be lost through over reliance on calculative rationality. Today, as always, individual decision-makers understand 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 isolated elements 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 and principles 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 and 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 calculatively 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.]

## Hubert Dreyfus

<http://globetrotter.berkeley.edu/people5/Dreyfus/dreyfus-con0.html>

\*\*\*\*\*

## Conversations with History: Hubert Dreyfus

[Conversations with History](#): Institute of International Studies, UC Berkeley




See the  
[Conversation  
s with History  
Blog](#)

See a webcast  
of this  
interview:



[Dreyfus](#)



This interview is part of the Institute's "[Conversations with History](#)" series, and uses Internet technology to share with the public Berkeley's distinction as a global forum for ideas.

Welcome to a Conversation With History. I'm [Harry Kreisler](#) of the Institute of International Studies. Our guest today is Hubert Dreyfus, who is Professor of Philosophy at U.C. Berkeley. His many publications include *What Computers Still Can't Do* and *On the Internet*. He recently won the Barwise Prize of the American Philosophical Association, an award given for bringing philosophy and the study of computers under one roof, and he is a past president of the American Philosophical Association.

1. [Background](#) ... leaving Terre Haute ... physics ... discovering philosophy ... dyslexia ... graduate education ... trying to find the right question ... AI students at MIT
2. [Teaching Philosophy](#) ... the idea of a calling ... being called to teach ... younger brother ... encountering AI's challenge to philosophy ... analytics vs. phenomenologists
3. [The Question of the Body](#) ... Greek emphasis on the mind ... Pascal ... Kirkegaard ... Heidegger ... thinking about skills ... getting stuck in the rules
4. [More Thoughts on Teaching](#) ... the value of accidents ... learning while teaching ... risks and mistakes
5. [Merleau-Ponty](#) ... Continental roots in Heidegger ... coping ... getting an optimal grip ... roots in Aristotle ... going beneath Heidegger



6. [Artificial Intelligence](#) ... intercorporeality vs. cyberspace ... rules vs. unique response ... AI theory rooted in Cartesian view ... "physical symbol systems" ... the importance of the body ... AI inherits a lemon ... hopeless case
7. [The Disembodied Internet](#) ... Merleau-Ponty and intercorporeality ... Samuel Todes and body structure ... the telephone
8. [Human Interaction: The Interview](#) ... going with the flow ... skillful coping and mastery ... advice for students ... trusting luck



© Copyright 2005, Regents of the University of California

Site questions: Email [iis\\_webmgr](mailto:iis_webmgr@berkeley.edu) at [berkeley.edu](http://berkeley.edu).