

Organization Studies

<http://oss.sagepub.com>

Peripheral Vision: Expertise in Real World Contexts

Hubert L. Dreyfus and Stuart E. Dreyfus

Organization Studies 2005; 26; 779

DOI: 10.1177/0170840605053102

The online version of this article can be found at:
<http://oss.sagepub.com/cgi/content/abstract/26/5/779>

Published by:

 SAGE Publications

<http://www.sagepublications.com>

On behalf of:



[European Group for Organizational Studies](#)

Additional services and information for *Organization Studies* can be found at:

Email Alerts: <http://oss.sagepub.com/cgi/alerts>

Subscriptions: <http://oss.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>



Peripheral Vision

Expertise in Real World Contexts

Hubert L. Dreyfus and Stuart E. Dreyfus

Hubert L. Dreyfus
University of
California,
Berkeley, USA

Stuart E. Dreyfus
University of
California,
Berkeley, USA

Abstract

In this paper we describe a five-stage phenomenological model of skill acquisition, of which expertise is the highest stage. Contrary to the claims of knowledge engineers, we argue that expertise in general, and medical expertise in particular, cannot be captured in rule-based expert systems, since expertise is based on the making of immediate, unreflective situational responses; intuitive judgment is the hallmark of expertise. Deliberation is certainly used by experts, if time permits, but it is done for the purpose of improving intuition, not replacing it. The best way to avoid mistakes is to take responsibility for them when they occur, rather than try to prevent them by foolproof rules. In bureaucratic societies, however, there is the danger that expertise may be diminished through over-reliance on calculative rationality.

Keywords: skills, expertise, rules, knowledge engineering, rationality, intuition

I take it that you are concerned with how to preserve a high level of skilled performance, while minimizing the risk of serious accidents. As a philosopher, it seems to me that, before we can discuss this question productively, it is important to have in mind a model of what a skill is that is true to the phenomenon. To that end, I'll briefly lay out a five-stage model of skill acquisition. In the context of this meeting, focusing on accident prevention, I'll highlight three issues that stand out as especially important: (1) To what extent can expertise be captured in rule-based expert systems? (2) How should the learner and the expert (who is always a learner) relate to mistakes when they do occur so as to learn from them? (3) How are reflection and deliberation related to expert performance?

Background: History of Expert Systems

In the 1980s, after the failure of Artificial Intelligence (AI), there emerged a new field called knowledge engineering, which by limiting its goals sought to apply AI research in ways that actually worked in the real world. The result was 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:

Organization
Studies
26(5): 779–792
ISSN 0170–8406
Copyright © 2005
SAGE Publications
(London,
Thousand Oaks,
CA & New Delhi)

'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.' (Feigenbaum and McCorduck 1983: 56)

What the knowledge engineers claimed to have discovered is that in areas that 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.' (Feigenbaum and McCorduck 1983: 76–77)

An early success using this approach was called MYCIN. This expert system took the results of blood tests and came up with a diagnosis of which blood disease was responsible for this condition. It even gave an estimate of the reliability of its own diagnosis. 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' (Feigenbaum and McCorduck 1983: 18), 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 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, 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.' (Feigenbaum and McCorduck 1983: 64)

So, since each expert already has a repertory of rules in mind, all the expert system builder needs 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 that 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.

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.' (Feigenbaum and McCorduck 1983: 55)

On this view, the rules are there functioning in the expert's mind whether he or she 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, 2,400 years later, thanks to Feigenbaum and his colleagues, we have a new name for what Socrates and Plato were doing: 'knowledge acquisition research' (Feigenbaum and McCorduck 1983: 79).

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.' (Feigenbaum and McCorduck 1983: 85)

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' (Feigenbaum and McCorduck 1983: 82).

But ever since the inception of AI, 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.

The same story is repeated in every area of expertise. 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 2,400 years of philosophy, we must take a fresh look at what a skill is and what the expert acquires when he or she achieves expertise. We must be prepared to abandon the traditional view that runs from Plato to Piaget and Chomsky 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. Since we are all experts in many areas, we have the necessary data, so let's turn to our experience.

A Five-Stage Model of the Acquisition of Expertise

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.¹

For purposes of illustration, I'll consider two variations: a motor skill (driving), which presumably everyone here has acquired, and an intellectual skill (chess), which has been studied by learning theorists. I'll then generalize to teaching, which is my own skill. Once these cases have been laid out, I'll be in a position to risk some comments concerning diagnostic and surgical skills, looking forward to your input correcting an outsider's misconceptions. As I'll argue later in more detail, one only learns by making mistakes.

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 the desired skill. 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 domain-independent features as speed (indicated by the speedometer) and is given rules such as shift to second when the speedometer needle points to 10. 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 one for calculating extent of control.

In the classroom and lecture hall, the teacher supplies the facts and procedures that need to be learned in order for the student to begin to develop an understanding of some particular domain. The student learns to recognize the features and follow the procedures by drill and practice.

But merely following rules will produce poor performance in the real world. A car stalls if one shifts too soon on a hill or when the car is heavily loaded; a chess player who always exchanges to gain points is sure to be the

victim of a sacrifice by the opponent who gives up valuable pieces to gain a tactical advantage. The student needs not only the facts but also an understanding of the context in which that information makes sense.

Stage 2: Advanced Beginner

As the novice gains experience actually coping with real situations and begins to develop an understanding of the relevant context, he or she begins to note, or an instructor points out, perspicuous examples of meaningful additional aspects of the situation or domain. 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*, recognized on the basis of experience, as well as to the objectively defined non-situational *features* recognizable by the novice.

The advanced beginner driver uses (situational) engine sounds as well as (non-situational) speed in deciding when to shift. He learns the maxim: 'Shift up when the motor sounds like it's racing and down when it sounds like it's straining.' Engine sounds cannot be adequately captured by a list of features, so features cannot take the place of a few choice examples in learning the relevant distinctions.

With experience, the chess beginner learns to recognize overextended positions and how to avoid them. Similarly, she begins to recognize such situational aspects of positions as a weakened king's side or a strong pawn structure, despite the lack of precise and situation-free definitions. The player can then follow maxims such as: attack a weakened king's side. Unlike a rule, a maxim requires that one already have some understanding of the domain to which the maxim applies.²

At school, the instructor takes on the role of a coach who helps the student pick out and recognize the relevant aspects that organize and make sense of the material. Though aspects can be presented to passive students in front of their terminals, it is more efficient for the students to attempt to use the maxims that have been given them, while the instructor points out aspects of the current situation to the student as the student encounters them.

Still, at this stage, learning, whether it takes place at a distance or face to face, can be carried on in a detached, analytic frame of mind, as the student follows instructions and is given examples. But to progress further seems to require a special kind of involvement.

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-racking and exhausting, and the student might well wonder how anybody 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,

that then determines which elements of the situation or domain must be treated as important and which ones can be ignored. As students learn to restrict themselves to only a few of the vast number of possibly relevant features and aspects, understanding and decision-making become easier.

Naturally, to avoid mistakes, the competent performer seeks 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 given beginners in manuals and lectures. Indeed, in any skill domain the performer encounters a vast number of situations differing from each other in subtle ways. There are, in fact, more situations than can be named or precisely defined, so no one can prepare for the learner a list of types of possible situations and what to do or look for in each. Students, therefore, must decide for themselves in each situation what plan or perspective to adopt without being sure that it will turn out to be appropriate.

Given this uncertainty, coping becomes frightening rather than merely exhausting. Prior to this stage, if the rules don't work, the performer, rather than feeling remorse for his or her mistakes, can rationalize that he or she hadn't been given adequate rules. But, since at this stage, the result depends on the learner's choice of perspective, the learner feels responsible for his or her choice. Often, the choice leads to confusion and failure. But sometimes things work out well, and the competent student then experiences a kind of elation unknown to the beginner.

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., the driver 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 perform any of these actions. 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 ignores 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. 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 attacks induce euphoria, while mistakes are felt in the pit of the stomach.

If we were disembodied beings, pure minds free of our messy emotions, our responses to our successes and failures would lack this seriousness and excitement. Like a computer we would have goals and succeed or fail to achieve them but, as John Haugeland once said of chess machines that have been programmed to win, they are good at attaining their goal, but when it comes to winning, they don't give a damn. For embodied, emotional beings like us, however, success and failure do matter. So the learner is naturally frightened, elated, disappointed, or discouraged by the results of his or her

choice of perspective. And, as the competent student becomes more and more emotionally involved in the task, it becomes increasingly difficult to draw back and adopt the detached maxim-following stance of the advanced beginner.

But why let learning be infected with all that emotional stress? Haven't we in the West, since the Stoics, and especially since Descartes, learned to make progress by mastering our emotions and being as detached and objective as possible? Wouldn't rational motivations, objective detachment, honest evaluation, and hard work be the best way to acquire expertise?

While it might seem that involvement could only interfere with detached rule testing, and so would inevitably 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, resistance to involvement and risk leads to stagnation and ultimately to boredom and regression.³

Since students tend to imitate the teacher as model, teachers can play a crucial role in whether students will withdraw into being disembodied minds or become more and more emotionally involved in the learning situation. If the teacher is detached and computer-like, the students will be too. Conversely, if the teacher shows involvement in the way he or she pursues the truth, considers daring hypotheses and interpretations, is open to students' suggestions and objections, and is ready to be shown wrong, the students will be more likely to let their own successes and failures matter to them, and rerun the choices that lead to successful or unsuccessful outcomes.

In the classroom and lecture hall the stakes are less dramatic than the risk of a car accident while driving or of losing an important game of chess. Still, there is the possibility of taking the risk of proposing and defending an idea and finding out whether it fails or flies.

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 is consequently in trouble, he or she tries to characterize the mistake by describing a certain class of dangerous situations and making a rule to avoid them in the future. To take an extreme example, if our 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 to 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 asking 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 involvement 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 or her 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 or she is involved or not. Furthermore, the novice is not emotionally involved in *choosing* an action, even if he or she 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 who goes on learning, one has to go on dwelling emotionally on what critical choices one has made and how they affected the outcome.

Stage 4: Proficiency

Only if the detached, information-consuming stance of the novice, advanced beginner, and distance learner is replaced by involvement, is the student set for further advancement. Then, the resulting positive and negative emotional experiences will strengthen successful responses and inhibit unsuccessful ones, and the performer's theory of the skill, as represented by rules and principles, will gradually be replaced by situational discriminations, accompanied by associated responses. Proficiency seems to develop if, and only if, experience is assimilated in this embodied, atheoretical way. Only then do intuitive reactions replace reasoned responses.

As usual, this can be seen most clearly in cases of action. As the performer acquires the ability to discriminate among a variety of situations, each entered into with involvement, plans are evoked and certain aspects 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 done rather than using a calculative procedure to select one of several possible alternatives. When the goal is simply obvious, rather than the winner of a complex competition, there is less doubt as to whether what one is trying to accomplish is appropriate.

To understand this stage of skill acquisition we must 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 reacting. The proficient performer simply has not yet had enough experience with the outcomes of the wide variety of possible responses to each of the situations he or she can now discriminate, to react automatically. Thus, the proficient performer, after spontaneously seeing the point and the important aspects of the current situation, must still *decide* what to do. And to decide, he or she 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 by some specific amount on the accelerator. Valuable time may be lost while making 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 know, for example, that she should attack, but she must calculate how best to do so.

A student at this level sees the question that needs to be answered but has to figure out what the answer is.

Stage 5: Expertise

The *proficient performer*, immersed in the world of skillful activity, *sees* what needs to be done, but *decides* how to do it. The *expert* not only sees what needs to be achieved; thanks to a vast repertoire of situational discriminations, he or she also sees immediately how to achieve the goal. Thus, the ability to make more subtle and refined discriminations is what distinguishes the expert from the proficient performer. Among many situations, all seen as similar with respect to plan or perspective, the expert has learned to distinguish those situations requiring one reaction 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 gradually decomposes this class of situations into subclasses, each of which requires a specific response. 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. It has been estimated that an expert chess player can distinguish roughly 50,000 types of positions. For much expert performance, the number of classes of discriminable situations, built up on the basis of experience, must be comparatively large.

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 and applies the appropriate pressure to the brake. What must be done, simply is done. As Aristotle says, the expert 'straightway' does 'the appropriate thing, at the appropriate time, in the appropriate way'.

The student, who has mastered the material, immediately sees the solution to the current problem.

What is the role of the teacher at this stage? A student learns by small random variations on what he or she is doing, and then checking to see whether or not performance has improved. Of course, it would be better for learning if these small random variations were not random — if they were sensible deviations. If the learner watches someone good at doing something, that could limit the learner's random trials to the more promising ones. So observation and imitation of the activity of an expert can replace a random search for better ways to act. In general, this is the advantage of being an apprentice. Its importance is particularly clear in professional schools.

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, or solve problems, or even think. He or she just does what normally works and, of course, it normally works.

The description of skill acquisition I have presented enables us to understand why knowledge engineers from Socrates to Feigenbaum have had such trouble getting the expert to articulate the rules being used. The expert is simply not following any rules! He or she is doing just what Socrates and Feigenbaum feared — 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 or she is using, one will, in effect, force the expert to regress to the level of a beginner and state the rules learned in school. Thus, instead of using rules he or she no longer remembers, as the knowledge engineers suppose, the expert is forced to remember rules he or she 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 or she has stored experience of the actual outcomes of tens of thousands of situations.

Adapting this Skill Model to Medicine

This model poses a problem, however. How come many people grow up to be expert drivers, but only a relative few become accomplished chess players, violinists, basketball players, etc.? The answer seems to be that there are at least two kinds of skills: *crude skills*, like walking and driving, and *subtle skills* like music, sports and, what is more relevant here, surgery. That is, in

driving and walking, it is generally easy to do what is required since here is a larger margin for error, there is time to make corrections, and the results are not irreversible. However, in chess, music, and sports, a tiny difference in what one does can make a huge difference in the result, so being an expert requires learning to make subtle discriminations. Also, one has to act quickly and, usually, doesn't get a chance to correct one's mistakes on the fly as a driver does. In surgery, I imagine, subtle differences are as critical, or even more critical, than in chess, the timing is faster than in normal driving, and the chances to recoup are as limited, or even more limited, than in sports.

It is characteristic of crude skills like driving that one can perform expertly while thinking about something else, whereas subtle skills like music, sports, and surgery, require intense concentration. In neither type of skill is there any time for reflection and deliberation. Indeed, the less one is monitoring what one is doing — the more one is in the flow, or 'playing out of one's head' as athletes say, or 'in the groove', as musicians would have it — the better the performance.

The Role of Deliberation

In day-to-day hospital care, however, it seems to me that, although sometimes an unreflective immediate response is called for, there would often be time for reflection and deliberation. Here what is crucial is engaging in the appropriate kind of reflection and deliberation so as to preserve one's expertise.

While most expert performance is ongoing and unreflective, 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.

The increasingly bureaucratic nature of society is heightening the danger that in future skill and expertise will be lost through over-reliance on calculative rationality. Today, as always, individual decision-makers understand and respond to their situations 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. In medicine, I suspect that there are life-and-death decisions so tied up with questions of fairness and legal responsibility that one has to follow agreed-upon rules even if the decisions reached may not be as refined as they would be if one were allowed to trust one's intuitive expertise.

Conclusions

A detailed look at what a skill is and how one acquires it leads us to the following conclusions relevant to the concerns of this meeting.

1. Once we understand that a skill is never produced by interiorizing the rules that make up the theory of a domain, we see that expert systems should not be introduced in the hope of reducing human error while

- preserving expert performance. Not that expert systems might not be able to reduce human error, but the resulting rule-based performance will never be more than competent. In some areas competence may be all that is required, or it may turn out that, for bureaucratic reasons, rules must be followed. Still, even when a fail-safe set of rules is in place, it is important that it be overridable if it flies in the face of the intuitions of those with long experience in the domain.
2. The best way to avoid mistakes is to take responsibility for them when they occur, not to try to prevent them by foolproof rules. All cooperative skilled behavior, of which medicine would seem to be a prime example, runs the risk that those involved will try to avoid group disapproval by trying not to take responsibility for mistakes. Since learning depends on taking responsibility for the mistakes of the team and especially for one's role in whatever went wrong, it is important to establish an ethos of accepting and even dwelling upon one's responsibility for failures as well as one's successes.
 3. Reflection and deliberation should certainly be sought whenever there is time and serious consequences are involved. But recourse to detached calculative deliberation should be avoided since, when one becomes detached, expert intuitive judgment is lost.

Notes

This paper was jointly written by Hubert and Stuart Dreyfus. For stylistic reasons, however, we have chosen to present it from Hubert's point of view. A version of this paper was presented at a meeting on 'Advancing Medical Education' at the University of Chicago Pritzker School of Medicine, 25 October 2001.

- 1 For a detailed treatment of the phenomenology of skill acquisition, see Dreyfus and Dreyfus (1982).
- 2 See Michael Polanyi's *Personal Knowledge* (Polanyi 1958).
- 3 Patricia Benner has described this phenomenon in Benner (1984: 164).

References

Benner, P.
1984 *From novice to expert: Excellence and power in clinical nursing practice*. Reading, MA: Addison-Wesley.

Dreyfus, H., and S. Dreyfus
1982 *Mind over machine*. New York: Free Press.

Feigenbaum, E., and P. McCorduck
1983 *The Fifth Generation: Artificial Intelligence and Japan's computer challenge to the world*. Reading, MA: Addison-Wesley.

Polanyi, M.
1958 *Personal knowledge*. London: Routledge & Kegan Paul.

Hubert L. Dreyfus

Hubert L. Dreyfus is Professor of Philosophy at the University of California, Berkeley, and author of, among others, *Being-in-the World: A Commentary on Heidegger's Being and Time, Division I*; *What Computers Still Can't Do*; and *Mind Over Machine* (with Stuart Dreyfus).

Address: Department of Philosophy, 314 Moses Hall 2390, University of California, Berkeley, CA 94720–2390, USA.

E-mail: dreyfus@cogsci.berkeley.edu

**Stuart E.
Dreyfus**

Stuart E. Dreyfus is Professor Emeritus of Industrial Engineering & Operations Research at the University of California, Berkeley, and author of, among others, *Artificial Neural Networks: Making a Mind Versus Modeling the Brain* (with Hubert Dreyfus); and *Mind Over Machine* (with Hubert Dreyfus).
E-mail: dreyfus@ieor.berkeley.edu