

Placement learning

What's so good about placement learning?

M Luntley

Expertise and judgment

Placement learning seems a good thing. It's where you pick up all the things they didn't and couldn't teach you in medical school. It's where you finally learn what you need to learn to acquire expertise in your chosen field. Expertise requires not only the knowledge found in medical textbooks and disseminated in lecture halls, it requires experiential knowledge. This is the knowledge that is distinctive of professional expertise. It is difficult to articulate and relies as much on intuition as following explicit rules. And that's why some of what you learn on placement is stuff they couldn't teach you in medical school, for it cannot be abstracted from the contingencies of practice.

The concept of experiential knowledge gives a satisfying explanation of why placement learning is a good thing. But the idea is problematic. There are conflicting intuitions about it. First, it is knowledge that is embedded in practice, it is contextualised to situation. That is why you need to go into placement to acquire it. Second, the knowledge learnt on placement transfers from case to case. The expert medic is not a one-hit wonder. Expertise enables quality performance time after time. Therefore, whatever you pick up on placement cannot be contextualised. If it was, it would not transfer from case to case. And if it did not transfer, how could you legitimise performance?¹ The first intuition supports the case for placement learning, the second takes it away. The puzzle is how to satisfy both intuitions at once. The solution is to make sense of the cognitive transfer from case to case.

What you learn from treating one patient makes your treatment of other patients better than it would otherwise have been. That's the cognitive transfer. But if that is right, what becomes of the concept of experiential knowledge, the idea that there is something contextualised about the wisdom acquired in situ on the ward and in the community? The literature on expertise does not help much. The idea of tacit knowledge,² of the trajectory from rule based novice to intuitive expert,³ the concept of experiential learning and the role of judgment in professional life have been much

described.⁵ But none of these descriptions really get to grips with the conceptual heart of our puzzle. Let's make the point more precise.

Expert systems model expertise in terms of explicit rules codified in decision trees. The idea that placement learning is a good thing because it is the site of experiential knowledge challenges this decision tree analysis of expertise. There is more to being an expert than internalising rules, otherwise nothing is left to the idea of autonomous professional judgment. You might as well replace the medic with a technician following the rules in the decision tree. On the other hand, how else can we make sense of the cognitive transfer in case by case learning other than in terms of rules that apply from one case to the other? If cognitive transfer is understood in terms of rules, then all the talk about experiential knowledge is simply hot air about knowledge that is difficult to formulate but which could in principle be made explicit. The platitudinous answer to my question is that placement learning allows you to learn how to put theory into practice. But if the decision tree analysis of expertise were right, there would be no need for this. You could simply enrol for the class that taught the decision tree rules. We need a better model of cognitive transfer.

Here is an alternative account of what goes on in case by case learning. The transfer is not mediated by rules, it is mediated by character. By "character" I mean the set of perceptual skills by which an individual couples with the environment. Having expertise is having a set of specialist discriminative skills that enable you to couple with features of the environment that novices miss. The key concept here is perceptual coupling with the environment, for it introduces a model of expertise which is not rule based.⁶⁻⁹ The idea is best illustrated with a simple example from robotics.

The traditional way of making a robot behave in a patterned way is to write an explicit rule for the behaviour in its instruction set. If you want a robot to turn on the central heating when the outside temperature drops to 8°C then

write that instruction into its programming, and provide it with a means for detecting the outside temperature. Actually, there is a much simpler way of doing this, you automate your central heating system with a thermostat. It's low grade artificial intelligence, but it illustrates a basic idea. A heating system with a thermostat behaves in a regular fashion. Whenever the temperature drops, the heating kicks in. It does this without having an explicit instruction for when to switch on for it is in dynamic coupling with the environment that links it directly with the relevant conditions. The behaviour of the system continually adapts to how things are outside. This is smart, but very low grade smart. An automated heating system does nothing more interesting than switch itself on and off and at the right temperature. It is not a general problem solving device. It solves just one problem, but very accurately and efficiently without the need for explicit rules. It turns out, however, that you can get robots to do quite a lot of things by providing them with choice environmental couplings rather than treating them as general problem solvers that require explicit rules for all the problems they treat.¹⁰ General problem solving devices need explicit rules in order to behave in a patterned way. Specialised devices do not need rules, they need appropriate forms of coupling that enable them to respond dynamically to specific environmental conditions. They are fitted for their environmental niche and operate efficiently in virtue of their coupling with that niche. The upshot is the same, behaviour gets to have the desired pattern to it.

So here's a thought: Are medical experts general problem solvers with an internalisation of all the rules required for handling whatever patients throw at them? Or, are they specialised devices equipped for responding to just the sorts of events for which they've been trained? Either way, they will behave in ways that are patterned.

The analogy with a thermostat is not very flattering. A more sophisticated version of the idea is that expertise consists in sets of perceptual skills for coupling with features of situations that the non-expert does not see. These perceptual skills have typical responses learnt with them. What matters in placement learning is not the assimilation of rules that have yet to be made explicit. What matters is that you hone and practise the perceptual skills by which you fine tune your responses to conditions your patients present. Much medical thinking is not rule driven, for it does not require general problem solving. It requires specific and focused problem solving. Expertise consists in a repertoire of perceptual couplings for specific problem solving, it is not an exercise in general problem solving.

There's an obvious objection to this: medical experts are general problem solvers—we all are. Fair point, but irrelevant. Of course medical experts are general problem solvers—they diagnose, do surgery, and many know how to turn on their central heating systems. The point that is worth examining is simply this: expertise in many professional fields is much more concerned with skills for specific problem solving than it is with general context-independent problem solving. It is the latter that can be articulated in general rules and analysed by the decision trees of expert systems. Task specific problem solving is more to do with having the perceptual skills to couple with the particularities that the environment throws at you and responding accordingly. It is about particular responses that bring stability by balancing conflicting goals—an area of our cognitive abilities that thus far has resisted analysis in the decision tree structures of expert systems.¹¹

Here is a second obvious objection. In thinking of expertise in terms of task specific skills, are we not endorsing a behaviourism that reduces the skills of the trained medic to a repertoire of stimulus-response pairings? Nothing could be further from the truth. Indeed, the model I am promoting offers real space for autonomous judgment and the capacity for innovation, rather than slavishly following the rule book.

The environments with which expert doctors engage are complex. They involve complex systems, such as people, whose gross behaviour cannot be predicted in a linear fashion from knowledge of how the elements are functioning. A patient may present with blood pressure with a diastolic >100. Their history may suggest this is a problem. There is, however, no drug you can prescribe that acts directly to reduce their blood pressure. There is a range of medications that will act on different elements of their physiology—adrenaline level, water retention, etc, that may have the desired effect on the gross behaviour. Expertise involves judging from among a repertoire of interventions on particular elements whose complex interactions with one another

produce the gross behaviour in which you are interested. Given the complexity of the system, there is and can be no general rule to tell you what to do to control their blood pressure. It requires judgment and judgment, like perceptual coupling, cannot be analysed in terms of rule following. The technician following the rule book has no scope for judgment, no capacity to respond to the particularities of the present case. The technician's response is prescribed by the rule. The expert with a repertoire of task specific skills has the capacity for judgment and novelty in handling the particularities of the individual patient. The expert does not turn to the rule book, they turn to the particularities of the case and make a judgment about what intervention to make next to try to bring the patient to some sort of stability. This is a complex, subtle and nuanced process which can look almost invisible from the perspective of the decision tree analysis.

The idea that perceptual knowledge is important in advanced expertise is familiar.¹² It figures in the highest levels in most taxonomies tracking the path from novice to expert.¹³ The appeal to perceptual knowledge on its own fails to answer my question, for perhaps all that perception supplies is a classification of a condition that has not yet been fully articulated but which, in principle, could be? If that were so, perception is simply a short cut to knowledge that fits a general rule. The model of perceptual knowledge I have sketched in terms of environmental coupling is quite different.

Perceptual coupling is a form of discrimination, of seeing things as similar. It is a case of spotting the patterns and responding in a patterned way. In doing this you classify and, having classified, you know what to do. But which comes first, the response or the classification? The model that treats perception as a short cut to rule based knowledge treats the classification as prior, the perceptual response is secondary. The model of perceptual coupling allows a different priority. Classifications, the patterns in our judgments and behaviour, emerge from our responses. The patterns are not all fixed in advance. They emerge from the dynamic interaction of the system's elements and our

interventions on them. And if that is anywhere near right, the rules set down in protocols for good practice cannot be prescriptive. They can, at best, be summaries of what has worked. What will work next will require expertise in coupling and judgment. And that's what you acquire in placement learning.

Postgrad Med J 2002;**78**:637–638

Author's affiliation

Department of Philosophy, University of Warwick, Coventry CV4 7AL, UK

Correspondence to: Professor Luntley;
michael.luntley@warwick.ac.uk

REFERENCES

- 1 **Crook J.** How do expert mental health nurses make on-the-spot clinical decisions? A review of the literature. *Journal of Psychiatric and Mental Health Nursing* 2001;**8**:1–5.
- 2 **Schon D.** *The reflective practitioner.* New York: Basic Books, 1983.
- 3 **Benner P.** *From novice to expert: excellence and power in clinical nursing practice.* London: Addison Wesley, 1984.
- 4 **Eraut M.** *Developing professional knowledge and competence.* London: Falmer Press, 1994.
- 5 **Hager P.** Know-how and workplace practical judgement. *Journal of Philosophy of Education* 2000;**34**:281–96.
- 6 **Luntley M.** *Performance, pay and professionals.* Impact pamphlet No 2. London: British Society for Philosophy of Education, 2000.
- 7 **Luntley M.** *Articulating practice.* *Proceedings of HAN annual conference 2000.* Milton Keynes: Open University, 2001.
- 8 **Luntley M.** Agency and our tacit sense of things. *Library of living philosophers—Marjorie Grene.* La Salle, IL: Open Court, 2002.
- 9 **Luntley M.** Knowing how to manage—expertise and embedded knowledge. *Reason in Practice* 2003;**3** (in press).
- 10 **Clark A.** *Being there.* Cambridge, MA: MIT Press, 1998 (gives numerous examples).
- 11 **Cowan R.** Expert systems: aspects of and limitations to the codifiability of knowledge. *Research Policy* 2001;**30**:1355–72.
- 12 **Elstein A, Schwartz A.** Clinical problem solving and diagnostic decision making: selective review of the cognitive literature. *BMJ* 2002;**324**:729–32
- 13 **Dreyfus H, Dreyfus S.** *Mind over machine: the power of human intuition and expertise in the era of the computer.* Oxford: Blackwell, 1986.