Chapter 22
Newly-Acquired Complex Performance Competence and Medical Errors

Alan Lesgold

Introduction

Increasingly, the acquisition of competence is defined using learning progressions [1]. For conceptual bodies of knowledge, such progressions are reasonably straightforward. They state stages of understanding and capability that people pass through on the path to expertise. Generally, once each stage is fully mastered, performance at that stage is relatively less demanding of cognitive processing resources, leaving some capacity free to notice cues for routine required actions. In mission critical areas, including many areas of medicine, competence includes not only knowing how to deal with situations but also being reliable, while exercising that expertise, in carrying out critical routines (e.g., hand washing) even when overloaded when complex problems that must be solved. This chapter considers the circumstances during the course of progressing to expertise under which there is a danger of routine but critical actions being omitted and then discusses possible ways to minimize the likelihood of critical omissions.

Newly-Acquired Complex Performance Competence and Medical Errors

In everyday life, we tend to assume that with practice people become more competent. This even is the case for certain critical kinds of competence, such as driving a car. We assume that it is the novice who will miss stop signs and not respond quickly enough to a potential accident situation. In this chapter, I raise the possibility that certain kinds of critical performances may vary non-monotonically

A. Lesgold, PhD
School of Education, University of Pittsburgh, Pittsburgh, PA 15260, USA
e-mail: al@pitt.edu

during the course of learning, sometimes becoming less reliable for a period after previously having been pretty well established. To establish this argument, it is necessary to accept that learning of complex performance capability proceeds in stages, i.e., that it involves learning progressions.

The study of science learning was the source for the notion of learning progressions [1, 2]. A learning progression is an account of the stages that a learner goes through in gaining expertise. Perhaps the best known learning progression is Piaget's stages of cognitive development, which specifies the stages a child goes through in becoming more able to gain understanding from new situations, progressively gaining the ability to observe, then to manipulate, then to plan abstractly a set of manipulations that might help in understanding a new set of situations. Many learning progressions, though, involve smaller and somewhat more concrete domains of competence, such as understanding electrical circuits or understanding how to diagnose cardiovascular disease or knowing how to evaluate and treat traumatic blows to the head.

Non-monotone Aspects of Competence Development

So, for example, specific areas of medical diagnosis and treatment knowledge may pass through several stages as that knowledge develops. Students learn enough anatomy and physiology to be able to understand how a disease develops and progresses, after which they learn how to reason through a specific case to diagnose that disease. The left side of Fig. 22.1 illustrates this progression. It also can happen that a student might learn a rule that is conceptually less completely grounded but still embodies the high probability that a particular cluster of symptoms indicates the likelihood of a particular disease. This is shown on the right side of Fig. 22.1.

Going even further, when all goes well, these two capabilities – to quickly recognize a disease from its symptoms and to diagnose it through reasoning about what could produce the presenting information – become coordinated, so that correct diagnoses come quickly to mind but also are reflected upon to be sure that they make sense in the case at hand. Note that these stages may be reached independently for different disease clusters and symptom clusters. One might, for example, become adept at dealing with one specialty like cardiology without becoming as well prepared in another like orthopedics. Indeed, the very presence of so many medical specialists is an indication that these stages are not stages of overall medical competence but rather for coherent subsets of medical practice.

Ordinarily, one would see progression from the novice stage through the intermediate levels to the expert stage as representing improvement in medical knowledge. Each stage, when fully attained, after all, means added diagnostic capability. Interestingly, though, sometimes short-term setbacks occur along the path to greater capability. In some of my own work on radiological expertise, this was the case [3, 4]. Indeed, I can recall situations over the course of a year or two where the same person diagnosed a particular X-ray image correctly early in residency and

Fig. 22.1 Sample learning progression
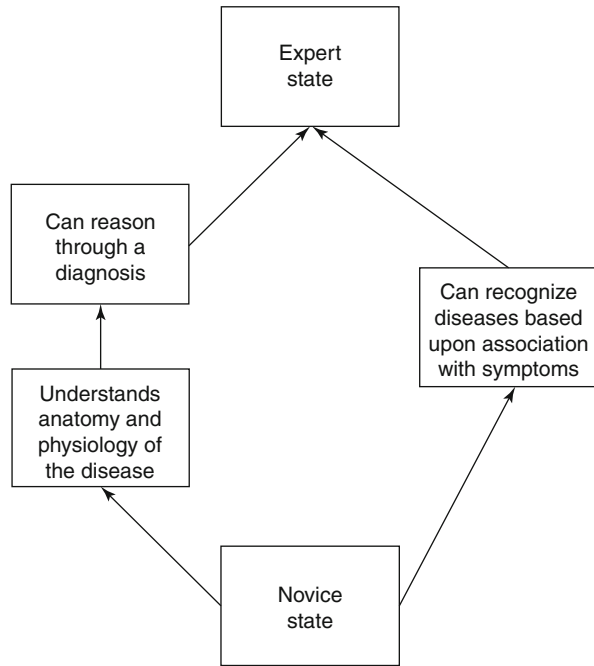
incorrectly after months of additional experience. While setbacks are temporary and overall competence keeps growing, such setbacks do occur. Perhaps the most well-known brief setback occurs as children learn tenses of verbs. It is not unusual for a child who has just learned that past tense verbs often end in –ed to revert from saying "went" to saying "goed". A range of developmental progressions are well documented to include brief setbacks [5]. These reversals of apparent competence are interesting because they may, as suggested below, be openings for increased medical error. I first discuss current views about non-monotone competence development and then consider its implications.

Three general explanations have been advanced for non-monotone developmental occurrences [6], and these also should be considered for non-monotone competence acquisition. First, and particularly relevant to medical learning, acquiring a more systematic approach to problem solving might lead to a small number of cases where the new approach fails while a more superficial approach might succeed. For example, I might correctly diagnose a particular instance of Disease X because the case matches the experience of my Aunt Maude. When I learn more about Disease X, I might learn that its standard symptoms more often mean that the patient has Disease Y but not yet know enough to recognize and understand why the symptoms of the patient like Aunt Maude indicate Disease X. That could make me incorrect in diagnosing a case like hers until I learn even more and become able to correctly differentiate the situations in which the less common situation arises.

A second possibility according to Siegler is that the newly learned capability may overload cognitive capacity until parts of it become automated. What used to

Table 22.1 Principles of knowledge and competence development

1. People are endowed with an innate information-processing system
2. People form higher units from lower units. In other words the learning system is hierarchical.
3. Higher units serve as components for still-higher units.
4. There is a bias to process using highest-formed units
5. If, for some reason, higher units are not available, lower-level units are utilized.

be shoot-from-the-hip recognition may suffer when deeper diagnostic capability has just been acquired, simply because the new inferential capability isn't automated enough to fit within limited human processing capacity. In such a situation, a resident who "follows his instincts" might be correct in a diagnosis while he might fail if put in a situation in which his diagnosis must be defended. This second possibility is the one to which I return below.

Siegler suggests a third possibility as well. This is that different aspects of competence may grow at different rates, allowing one aspect to overshadow another with the other being dominant later. In infant development, for example, leg length and leg muscle strength develop on slightly different tracks, so when the leg grows faster than its muscles, apparent balance capability may briefly be lost. In medical learning, knowledge of different mechanisms may similarly show uncoordinated development, leading to diagnoses favoring whichever area of practice has been dealt with most recently, especially for medical students and interns/residents on rotating assignments.

The prevailing research view [6] is that the underlying accumulation of knowledge is monotone, i.e., that further learning or development does not destroy knowledge, even if certain capabilities may temporarily decrease. Nonetheless, for the purpose of patient safety, understanding setbacks is important. Before getting to that, it is worth considering what the basic principles of the development of knowledge are in the first place, since we may be better able to anticipate how cognitive overload will express itself if we consider those principles.

While many different sets of learning principles can be found, when considering the long-term development of medical expertise, it is worth attending to principles originating in the developmental psychology world. For example, Table 22.1 lists a set of principles [7] we might consider (in quoting these principles, I have replaced the word "infant" with the word "people" because of the focus of this chapter). The first three principles explain a little of how the learning at different stages is combined to create higher-order processing units. Most notably for medicine, direct statistical association of symptoms to diagnoses and deep understanding of the mechanisms behind diseases get integrated into higher-order units that encompass both knowledge sources, enabling experts both to quickly recognize diseases and to check their recognition against a set of expectations generated from their deeper understanding [8–10].

The fourth and fifth principles in Table 22.1 are especially important in understanding performance failures. Once a resident has acquired a reasonable level of

Table 22.2 Stages in formation of medical expertise [9]

1. Development of elaborate declarative networks explaining the causes and consequences of disease in terms of general underlying pathophysiological processes
2. Encapsulation of these declarative networks into a limited number of diagnostic labels, syndromes or high-level, simplified causal models, explaining signs and symptoms
3. Transition into illness scripts through the acquisition of experience-based, contextual or enabling conditions knowledge
4. Storage of interpreted instances of these scripts as exemplars of the particular illness

ability to reason about the meaning of symptom clusters in a given situation, we can expect him to use that ability. Sometimes, though, when that ability is just developed, the cognitive load imposed by deeper reasoning can interfere with successful recognition-based performance. More broadly, it can interfere with a range of clinical behaviors that otherwise might occur close to automatically. It is this paradox, that deeper understanding of why one's recognition-based decisions are right can interfere with making and acting upon those recognitions until the deeper understanding is automated, that I suggest merits a bit more attention. After all, at least in hospital settings, much of medical care is delivered by new physicians who have just acquired much of the knowledge they use every day.

This is the fundamental point I wish to remind about in this chapter. It is common in discussing errors, both in aviation and in medicine, to cite "human error" as the cause, implying that an actor at the scene should have tried harder. As the effort to reduce error has matured in each area where human performance is critical, we have learned that some human error, while predictable, cannot be contained by just pushing people to work harder. Shooting soldiers on guard duty who fell asleep did not make camps more secure in the eighteenth and nineteenth centuries. Blaming pilots who, generally, died as a consequence of the error being identified, did not make flying safer. We cannot expect that simply manipulating incentives for the erring actor will make medical errors less likely. Rather, changes in training and especially engineering of patient care environments – the topic of some chapters in this volume – are essential.

In considering the status of medical expertise in hospital settings, it is useful to keep in mind how that expertise develops. Table 22.2 quotes four stages of developing expertise put forward by Schmidt and Rikers [9]. These stages pretty much ignore the acquisition of specific recognition for symptom clusters that occurs alongside knowledge-driven recognition of diseases, but they nicely unpack some of the ways in which knowledge-driven diagnostic skill develops and the later stages in which it is integrated with memory of specific cases that become exemplars.

For purposes of this discussion, what is important is that a lot of learning takes place after medical students and new physicians have acquired both substantial understanding of disease and its manifestations – and consequently after much of their time as hospital house staff. Moreover, each stretch of existing knowledge produces a period of increased cognitive activity, including more extensive inference from primary medical knowledge [11], which can exhaust the cognitive capabilities of the new physician.

Cognitive Overload and Medical Errors

Given the concepts sketched above, it may be worth exploring some of the implications of cognitive overload that occurs as medical knowledge is expanding. Perhaps the most important implication is that initial demonstration of mastery within a restricted situation may be an overestimate of the reliability of knowledge. Sometimes this is mundane. For example, one might observe, in a restricted set of situations, that a resident always washes his hands upon entering a patient room. Even so, we might expect that on occasion, when the resident is extremely overloaded mentally, he may forget to wash before touching the patient. Given that this can occur even when the disposition to do the right thing is present, extensive learning has occurred, and mastery has been demonstrated, it makes sense to provide an efficient and effective means of reminding the resident to wash.

Many such approaches to reminding have been tried. Some are likely to fail because the reminders themselves become so commonplace as to not intrude into consciousness when one is overloaded. Others are more effective because they intrude more into consciousness. For example, at least for compromised patients who require masks and gowns, the placement of a rack near the room door with all of the apparel that is needed likely also will prompt hand washing, simply because it intrudes so completely. Interestingly, this might be a situation where well-meant efforts to move the rack of gloves, masks, and gowns out of the way to facilitate movement of equipment and patients in and out of a room could decrease the effectiveness of the rack as an intrusive warning to engage in actions that, in easy cases, might be automatic and assumed.

There are, of course, other situations in which errors occur that are more complex. Here again, the first approach to consider is probably to assure that the patient environment intrusively reminds health care workers to do the right thing. Intrusion is critical if cognitive overload is the problem, since attentional field is decreased under conditions of overload, so routine warnings not only are habituated to but also lose effectiveness since they may not be noticed when the cognitive capabilities of a health care worker are overloaded. Another useful form of intrusion is paraprofessional help. A culture in which it is acceptable for a nurse to remind a doctor about basic practices will likely do better in maintaining those practices, since while the doctor may be concentrating on a hard diagnosis, the nurse may not simultaneously be as overloaded.

While incentives to reduce errors may not be effective when arranged for the key actor in a medical situation, since that actor already may be overloaded cognitively, such incentives may work when provided to other health care workers who may not be as overloaded and hence more able to remind the key actor to carry out a required action such as hand washing. More broadly, though, it should be noted that what incentives do, in essence, is elevate one action to be more likely than others. Under conditions of cognitive overload, simply making a required routine act take over consciousness and interrupt more complex thinking will assure that the routine act is carried out but interfere with the action that depends upon the cognition that
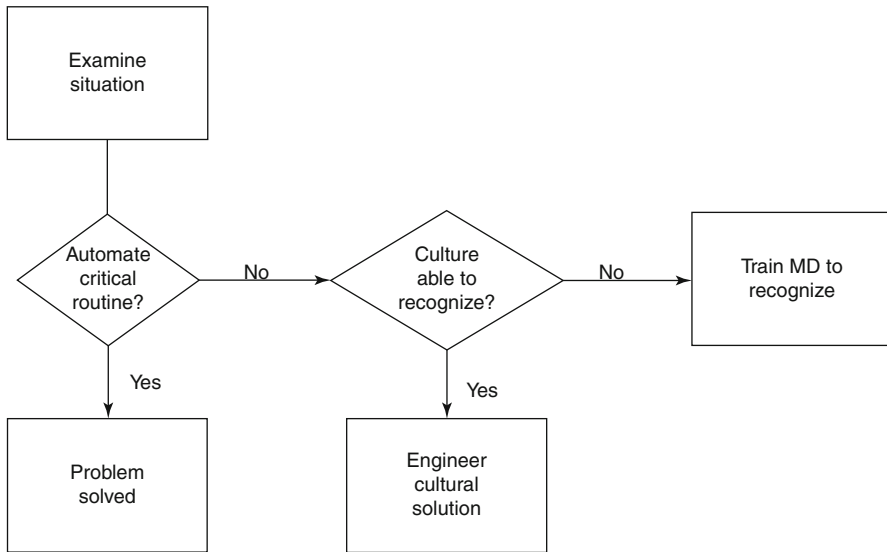
Fig. 22.2 Decision tree for preventing overload errors

produced the overload. This suggests that reminders alone may not always do the job. They will work best when they prompt a needed response, that response is highly automated (and thus not demanding of substantial cognitive resources), and the environment is engineered to best support an automated response.

When considering how best to assure that critical actions are performed, then, three basic possibilities should be considered, as suggested by Fig. 22.2. The best option, as just discussed, often is to automate the assurance of a critical function or make its achievement minimally demanding of cognitive resources. An example of this is the checklist. When a work protocol uses a checklist, there is a high level of certainty that each step in the checklist will be executed, at least when the step is understood by the work team. Moreover, a checklist serves as a temporary memory for work in progress, so the execution of a critical step will not erase the group's memory for steps that need to follow. For this reason, checklists are extremely useful. However, not all checklists are effective, and it is essential to design checklist and associated training well if they are to work [12]. For example, in order for a checklist to solve a problem, there has to be a trigger for its use. For example, the takeoff checklist used by pilots only works for takeoffs. If there is an emergent event in the air, it may or may not trigger a checklist type of protocol. Similarly, hospitals have code protocols (which really are somewhat more elaborated versions of checklists), which help assure that important actions are not overlooked in defined code situations. Such protocols work when there is a triggering event that causes them to be entered.

However, there are circumstances in which there is no triggering event, and there are also lapses in noticing the critical event and thereby triggering use of a protocol.

The simple hand washing case is an example of this. The cognitive overload that sometimes occurs in hospital settings tends to result in a person not noticing the appropriate trigger for hand washing. This can be overcome perhaps as suggested above, by making the cues for hand washing more intrusive. In essence, that is a variation of automating or engineering a solution to the problem. Sometimes, though, that is not completely possible. In that case, perhaps the next possibility to consider is enculturation of social processes that assure the triggering of appropriate routines.

Enculturation of Reliability

The aviation industry has gone through several generations of crew team training, and that training increasingly includes schemes of team protocols that assure that important routines are not ignored when complex situations occur [13]. Clearly, similar approaches are possible in medicine, and this book's chapters provide glimpses of this possibility.

One important element of such training is the development of shared understanding by the medical team of the effects of cognitive overload and the circumstances under which it is likely to arise. Another key element is likely to be identification of who in the team is most likely to not be overloaded and hence capable of assuring that critical routines are not missed. Finally, a team needs to practice engaging complex situations and moving into "crisis mode." In that crisis mode, part of the overt team activity is to split the needed tasks in ways that assure that critical tasks are not given to someone overloaded by even harder tasks, at least to the extent possible.

The aviation industry does this kind of training routinely, even though there are no fixed teams in aviation – individual pilots and cabin attendants are assigned individually, with only a modest force toward team continuity produced by union work rules and the accidents of particular personnel with seniority wanting to work together. Nonetheless, the team training works and, while the number of emergency cases has been extremely low both before and after training, such training seems to have reduced air disasters [14]. So, it seems worthwhile to consider medical team training across job levels (i.e., physicians, nurses, technicians) as an important part of an effort to remove medical errors.

It is important to realize, though, that situations of cognitive overload in the medical world are not restricted to emergency situations such as code calls. Aircraft crew deal with one flight at a time, and a flight lasts, on average, a couple hours. Medical practice is organized to involve almost continual parallel processing. Nurses have multiple patients to care for. While doctors in some specialty areas can make sequential rounds, seeing a patient at a time, even then it is not unusual for a doctor to see a patient, request data that takes time to collect, such as a lab test, and then be interrupted when the data become available. And, of course, there are practice areas, such as emergency department work, where parallel processing is pretty much the order of the day, at least at times. Even during office hours, it is not unusual for a physician to have overlap in the overall processing of patients, and the office team certainly has such overlap – a nurse or aide is positioning one patient to

be seen after the patient the physician is seeing at the moment while also handling the orders of the patient already seen and being on call to assist the physician with the patient currently being seen.

One approach, consistent with some of the work reported elsewhere in this volume, is to do the needed team training and related training of individual health care roles on the assumption that medical settings are continually in an overload mode. In the airplane cockpit, most of the time activity is being handled by the plane itself, with one pilot monitoring the situation to detect potential issues. In contrast, medical care systems, with their emphasis on staffing efficiency, generally have most or all staff assigned enough critical work to pretty much fill their cognitive capacity, at least much of the time. Therefore, it is less critical for such staff to learn when to go into team collaboration mode and more critical for teams to have permanent arrangements in place to assure that cognitive overload does not lead to omission of critical actions.

How this can happen will vary as a function of how urgent certain actions are. Let us first consider situations in which it is problematic if certain actions are omitted for extended periods but where moment-to-moment changes do not generally require action. For example, after I had my hip replaced, it was important for someone to check periodically, but not continually, for signs of blood clots. If such checks did not occur in the past hour, the risk would be extremely low, but if there was no check for a day or two, this could be more problematic. Presumably, this kind of problem of things not being noticed in a timeline of a few hours could be handled by a checklist. So, for example, an electronic patient record system might prompt the duty nurse every few hours to verify that a set of checks had been made – such prompting already occurs for various routine actions. While this is partly an engineering solution rather than a training solution, it also is common to train nursing staff as well as resident physicians to ask a set of check questions whenever visiting a given patient – or every few hours when visits are continual or frequent.

It would be consistent with the results reported or referenced in this volume to develop further research on the efficacy of monitoring checklists of this kind. While much of checklist work has been done to assure that relevant actions are taken at a point of treatment, extending the concept to checklists that can be applied to monitor whether the longer term course of patient care is free of critical omissions makes a lot of sense. Personal experience suggests that for "quality of life" issues like food service in hospitals, this is pretty routine, but perhaps it is less well developed or systematic for medical actions that should occur regularly but are not critical at any given moment. So, for example, surgical residents routinely check themselves for certain complications on daily visits, as part of implicit or explicit protocols, but they may be less likely to check for the pattern of overall health care team attention to certain issues over the past day.

The Role of Reminders

It also can safely be predicted that scheduled actions will occur reliably but that unscheduled actions and checks may have higher chance of omission. While some

of these unscheduled checks can be scheduled, by prompting a nurse to enter explicitly the results of a check just as there is prompting to deliver medication, there may be team activities that could supplement this, since there will never be a truly complete list of all the things that a nurse or physician should notice in a patient. In particular, team meetings might include discussion of what needs to be checked for in a particular class of patients, whether those checks are occurring reliably, and what evidence supports that belief. While much of this may be routine and part of training for different health care professionals, it also is likely that problems might emerge in such discussions and that discussion of those problems might lead to changes in systemic prompts for certain checks.

As noted above, there is a danger in automating reminders about routine checks. A check that tends to be made automatically requires minimal cognitive resources. While it is possible that without an automated reminder and under cognitive overload, the trigger that should prompt a check may fail to be attended, it also is possible, as discussed in this volume, that an extended array of automated reminders may produce cognitive overload and lack of attention to all of the warnings that might be posted. This volume reports work on dashboards for patient information display, and all of the problems associated with such dashboards also are present in any collection of automated warning systems meant to assure that routine checks are made and routine actions taken.

More broadly, the management of all of the routine as well as alarming data that is generated by or observable in a patient is itself a major source of cognitive overload. Systems that simply remind health care staff about checks needed or situations meriting a response are likely to contribute to cognitive overload and hence exacerbate the problem of omission of needed care activity. This volume includes discussion of efforts to improve dashboard displays so that the most relevant information is most salient and information is organized in manageable ways. While intelligent display management and prioritization of relevant information has great potential for improving the reliability and success of patient care, though, more may be required.

Specifically, some of the intelligence in data management will likely need to be provided by the health care worker, to supplement what can be done by machines. No matter how good automated prioritizing of warnings gets, it will not be perfect, and it will be critical to train personnel to find ways to work together to assure good outcomes.

Conclusion

To summarize, past research and the findings in this volume suggest that management of cognitive overload is a key requirement to assure that critical but routine actions are taken when needed. Engineering work environments and team work patterns is an important way to better prompt such actions and to assure that they are taken, when possible, by the health care worker whose overall activity will be least

affected by the added cognitive load. Training teams to distribute work to minimize overload and to react positively to reminders from colleagues likely will contribute to improved reliability of critical actions. In some cases, though, it also will be necessary to train health care teams to review their own performance of the routine and mundane but critical and to consider ways to improve it. The work reported in this volume makes considerable progress on the research needed to elaborate and confirm the efficacy of these key steps.

Discussion Questions

1. Managing cognitive load is an important consideration in efficient and effective management of critical care activities. What impact do you think the health care technology will have in managing cognitive load?
2. What new skills do you think will be needed for competent performance in complex domain, as heath care technology becomes a part of our everyday clinical practice?
3. What aspects of training in team collaboration will be most useful in assuring reliability of health care in situations where some team members will experience high cognitive load?

References

1. Smith CL, Wiser M, Anderson CW, Krajcik J. Implications of research on children's learning for standards and assessment: a proposed learning progression for matter and the atomic molecular theory. Meas Interdiscipl Res Perspect. 2006;4:1–98.
2. Stevens SY, Shin N, Krajcik JS. Towards a model for the development of an empirically tested learning progression. In: Learning progressions in science conference, Iowa City, 2009.
3. Lesgold A, Rubinson H, Feltovich P, Glaser R, Klopfer D, Wang Y. Expertise in a complex skill: diagnosing X-ray pictures. In: Chi M, Glaser R, Farr M, editors. The nature of expertise. Hillsdale: Erlbaum; 1988. p. 311–42.
4. Lesgold AM. Acquiring expertise. In: Anderson JR, Kosslyn SM, editors. Tutorials in learning and memory: essays in honor of Gordon Bower. San Francisco: Freeman; 1984. p. 31–60.
5. Cashon CH, Cohen LB. Beyond U-shaped development in infants' processing of faces: an information-processing account. J Cogn Dev. 2004;5:59–80.
6. Siegler RS. U-shaped interest in U-shaped development – and what it means. J Cogn Dev. 2004;5:1–10.
7. Cohen LB, Chaput HH, Cashon CH. A constructivist model of infant cognition. Cogn Dev. 2002;17:1323–43.
8. Boshuizen HP, Schmidt HG. On the role of biomedical knowledge in clinical reasoning by experts, intermediates and novices. Cogn Sci. 1992;16:153–84.
9. Schmidt HG, Rikers RM. How expertise develops in medicine: knowledge encapsulation and illness script formation. Med Educ. 2007;41:1133–9.
10. Patel VL, Glaser R, Arocha JF. Cognition and expertise: acquisition of medical competence. Clin Invest Med. 2000;23:256–60.

11. Patel VL, Groen GJ, Arocha JF. Medical expertise as a function of task difficulty. Mem Cogn. 1990;18:394–406.
12. Hales B, Terblanche M, Fowler R, Sibbald W. Development of medical checklists for improved quality of patient care. Int J Qual Health Care. 2008;20:22–30.
13. Helmreich RL, Merritt AC, Wilhelm JA. The evolution of crew resource management training in commercial aviation. Int J Aviat Psychol. 1999;9:19–32.
14. Flin R, O'Connor P, Mearns K. Crew resource management: improving team work in high reliability industries. Team Perform Manag. 2002;8(3/4):68–78.